Hyperparameter Optimization in Black-box Image Processing using Differentiable Proxies

ETHAN TSENG, FELIX YU, and YUTING YANG, Princeton University, United States
FAHIM MANNAN and KARL ST. ARNAUD, Algolux, Canada
DEREK NOWROUZEZAHRAI, McGill University, Canada
JEAN-FRANÇOIS LALONDE, Université Laval, Canada
FELIX HEIDE, Princeton University, United States and Algolux, Canada

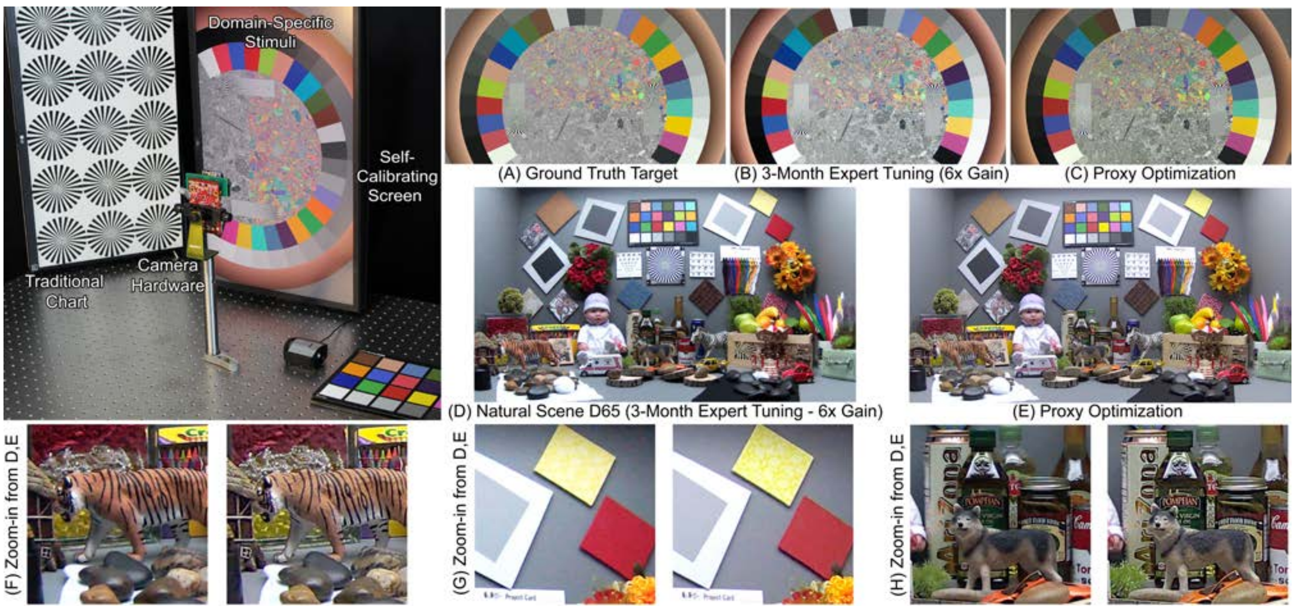

Fig. 1. We combine a radiometrically calibrated screen with a *hardware-in-the-loop* camera to generate stimulus/processed image pairs. We then train a *differentiable proxy model* on the displayed target (A), before using the model to optimize application-specific imaging tasks. Here, we optimize hyperparameters of a hardware ISP (ARM Mali C71) for perceptual accuracy. Our prototype improves image quality (C & E), compared to a three-month-long tuning by imaging experts (B & D). Our method is automatic and generates results in under an hour, enabling rapid prototyping for domain-specific imaging systems.

Nearly every commodity imaging system we directly interact with, or indirectly rely on, leverages power efficient, application-adjustable black-box hardware image signal processing (ISPs) units, running either in dedicated hardware blocks, or as proprietary software modules on programmable hardware. The configuration parameters of these black-box ISPs often have complex interactions with the output image, and must be adjusted prior to deployment according to application-specific quality and performance metrics. Today, this search is commonly performed *manually* by "golden eye" experts or algorithm developers leveraging domain expertise. We present a *fully automatic* system to optimize the parameters of black-box hardware and software image processing pipelines according to any arbitrary (i.e., application-specific) metric. We leverage a *differentiable* mapping between the configuration space and evaluation metrics, parameterized by a convolutional neural network that we train in an end-to-end fashion with imaging hardware in-the-loop. Unlike prior art, our *differentiable proxies* allow for high-dimension parameter search with stochastic first-order optimizers, without explicitly modeling any lower-level image processing transformations. As such, we can efficiently optimize black-box image processing pipelines for a variety of imaging applications, reducing application-specific configuration times from months to hours. Our optimization method is fully automatic, even with black-box hardware in the loop. We validate our method on experimental data for real-time display applications, object detection, and extreme low-light imaging. The proposed approach outperforms manual search qualitatively and quantitatively for all domain-specific applications tested. When applied to traditional denoisers, we demonstrate that—just by changing hyperparameters—traditional algorithms can outperform recent deep learning methods by a substantial margin on recent benchmarks.

CCS Concepts: • **Computing methodologies → Computational photography**.

Additional Key Words and Phrases: image processing

ACM Reference Format:
Ethan Tseng, Felix Yu, Yuting Yang, Fahim Mannan, Karl St. Arnaud, Derek Nowrouzezahrai, Jean-François Lalonde, and Felix Heide. 2019. Hyperparameter Optimization in Black-box Image Processing using Differentiable

0730-0301/2019/7-ART27 $15.00
https://doi.org/10.1145/3306346.3322996

Handheld Multi-Frame Super-Resolution

BARTLOMIEJ WRONSKI, IGNACIO GARCIA-DORADO, MANFRED ERNST, DAMIEN KELLY, MICHAEL KRAININ, CHIA-KAI LIANG, MARC LEVOY, and PEYMAN MILANFAR, Google Research

Fig. 1. We present a multi-frame super-resolution algorithm that supplants the need for demosaicing in a camera pipeline by merging a burst of raw images. We show a comparison to a method that merges frames containing the same-color channels together first, and is then followed by demosaicing (**top**). By contrast, our method (**bottom**) creates the full RGB directly from a burst of raw images. This burst was captured with a hand-held mobile phone and processed on device. Note in the third (red) inset that the demosaiced result exhibits aliasing (Moiré), while our result takes advantage of this aliasing, which changes on every frame in the burst, to produce a merged result in which the aliasing is gone but the cloth texture becomes visible.

Compared to DSLR cameras, smartphone cameras have smaller sensors, which limits their spatial resolution; smaller apertures, which limits their light gathering ability; and smaller pixels, which reduces their signal-to-noise ratio. The use of color filter arrays (CFAs) requires demosaicing, which further degrades resolution. In this paper, we supplant the use of traditional demosaicing in single-frame and burst photography pipelines with a multi-frame super-resolution algorithm that creates a complete RGB image directly from a burst of CFA raw images. We harness natural hand tremor, typical in handheld photography, to acquire a burst of raw frames with small offsets. These frames are then aligned and merged to form a single image with red, green, and blue values at every pixel site. This approach, which includes no explicit demosaicing step, serves to both increase image resolution and boost signal to noise ratio. Our algorithm is robust to challenging scene conditions: local motion, occlusion, or scene changes. It runs at 100 milliseconds per 12-megapixel RAW input burst frame on mass-produced mobile phones. Specifically, the algorithm is the basis of the *Super-Res Zoom* feature, as well as the default merge method in *Night Sight* mode (whether zooming or not) on Google's flagship phone.

Authors' address: Bartlomiej Wronski, bwronski@google.com; Ignacio Garcia-Dorado, ignaciod@google.com; Manfred Ernst, ernstm@google.com; Damien Kelly, damienkelly@google.com; Michael Krainin, mkrainin@google.com; Chia-Kai Liang, ckliang@google.com; Marc Levoy, levoy@google.com; Peyman Milanfar, milanfar@google.com Google Research, 1600 Amphitheatre Parkway, Mountain View, CA, 94043.

0730-0301/2019/7-ART28
https://doi.org/10.1145/3306346.3323024

CCS Concepts: • **Computing methodologies → Computational photography**; **Image processing**.

Additional Key Words and Phrases: computational photography, super-resolution, image processing, photography

ACM Reference Format:
Bartlomiej Wronski, Ignacio Garcia-Dorado, Manfred Ernst, Damien Kelly, Michael Krainin, Chia-Kai Liang, Marc Levoy, and Peyman Milanfar. 2019. Handheld Multi-Frame Super-Resolution. *ACM Trans. Graph.* 38, 4, Article 28 (July 2019), 18 pages. https://doi.org/10.1145/3306346.3323024

1 INTRODUCTION

Smartphone camera technology has advanced to the point that taking pictures with a smartphone has become the most popular form of photography [CIPA 2018; Flickr 2017]. Smartphone photography offers high portability and convenience, but many challenges still exist in the hardware and software design of a smartphone camera that must be overcome to enable it to compete with dedicated cameras.

Foremost among these challenges is limited spatial resolution. The resolution produced by digital image sensors is limited not only by the physical pixel count (e.g., 12-megapixel camera), but also by the presence of color filter arrays (CFA)[1] like the Bayer CFA [Bayer 1976]. Given that human vision is more sensitive to green, a quad of pixels in the sensor usually follows the Bayer pattern RGGB; i.e., 50% green, 25% red, and 25% blue. The final full-color image is generated from the spatially undersampled color channels through an interpolation process called demosaicing [Li et al. 2008].

[1] Also known as a color filter mosaic (CFM).

Local Light Field Fusion: Practical View Synthesis with Prescriptive Sampling Guidelines

BEN MILDENHALL*, University of California, Berkeley
PRATUL P. SRINIVASAN*, University of California, Berkeley
RODRIGO ORTIZ-CAYON, Fyusion Inc.
NIMA KHADEMI KALANTARI, Texas A&M University
RAVI RAMAMOORTHI, University of California, San Diego
REN NG, University of California, Berkeley
ABHISHEK KAR, Fyusion Inc.

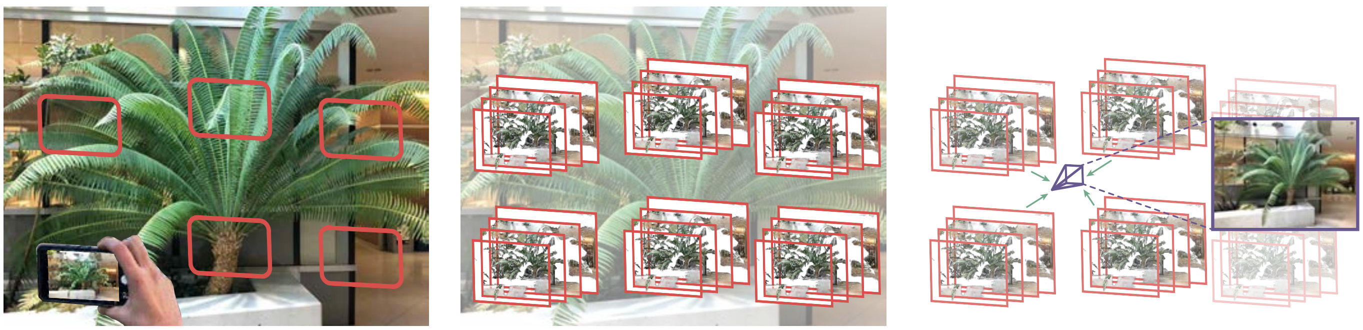

Fast and easy handheld capture with guideline: closest object moves at most D pixels between views

Promote sampled views to local light field via layered scene representation

Blend neighboring local light fields to render novel views

Fig. 1. We present a simple and reliable method for view synthesis from a set of input images captured by a handheld camera in an irregular grid pattern. We theoretically and empirically demonstrate that our method enjoys a prescriptive sampling rate that requires 4000× fewer input views than Nyquist for high-fidelity view synthesis of natural scenes. Specifically, we show that this rate can be interpreted as a requirement on the pixel-space disparity of the closest object to the camera between captured views (Section 3). After capture, we expand all sampled views into layered representations that can render high-quality local light fields. We then blend together renderings from adjacent local light fields to synthesize dense paths of new views (Section 4). Our rendering consists of simple and fast computations (homography warping and alpha compositing) that can generate new views in real-time.

We present a practical and robust deep learning solution for capturing and rendering novel views of complex real world scenes for virtual exploration. Previous approaches either require intractably dense view sampling or provide little to no guidance for how users should sample views of a scene to reliably render high-quality novel views. Instead, we propose an algorithm for view synthesis from an irregular grid of sampled views that first expands each sampled view into a local light field via a multiplane image (MPI) scene representation, then renders novel views by blending adjacent local light fields. We extend traditional plenoptic sampling theory to derive a bound that specifies precisely how densely users should sample views of a given scene when using our algorithm. In practice, we apply this bound to capture and render views of real world scenes that achieve the perceptual quality of Nyquist rate view sampling while using up to 4000× fewer views. We demonstrate our approach's practicality with an augmented reality smartphone app that guides users to capture input images of a scene and viewers that enable realtime virtual exploration on desktop and mobile platforms.

CCS Concepts: • **Computing methodologies → Image-based rendering**.

Additional Key Words and Phrases: view synthesis, plenoptic sampling, light fields, image-based rendering, deep learning

ACM Reference Format:
Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. 2019. Local Light Field Fusion: Practical View Synthesis with Prescriptive Sampling Guidelines. *ACM Trans. Graph.* 38, 4, Article 29 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3322980

*Denotes equal contribution

0730-0301/2019/7-ART29 $15.00
https://doi.org/10.1145/3306346.3322980

1 INTRODUCTION

The most compelling virtual experiences completely immerse the viewer in a scene, and a hallmark of such experiences is the ability to view the scene from a close interactive distance. This is currently possible with synthetically rendered scenes, but this level of intimacy has been very difficult to achieve for virtual experiences of real world scenes.

Synthetic Defocus and Look-Ahead Autofocus for Casual Videography

XUANER ZHANG, University of California, Berkeley
KEVIN MATZEN, Facebook Research
VIVIEN NGUYEN, University of California, Berkeley
DILLON YAO, University of California, Berkeley
YOU ZHANG, Chapman University, Independent Filmmaker
REN NG, University of California, Berkeley

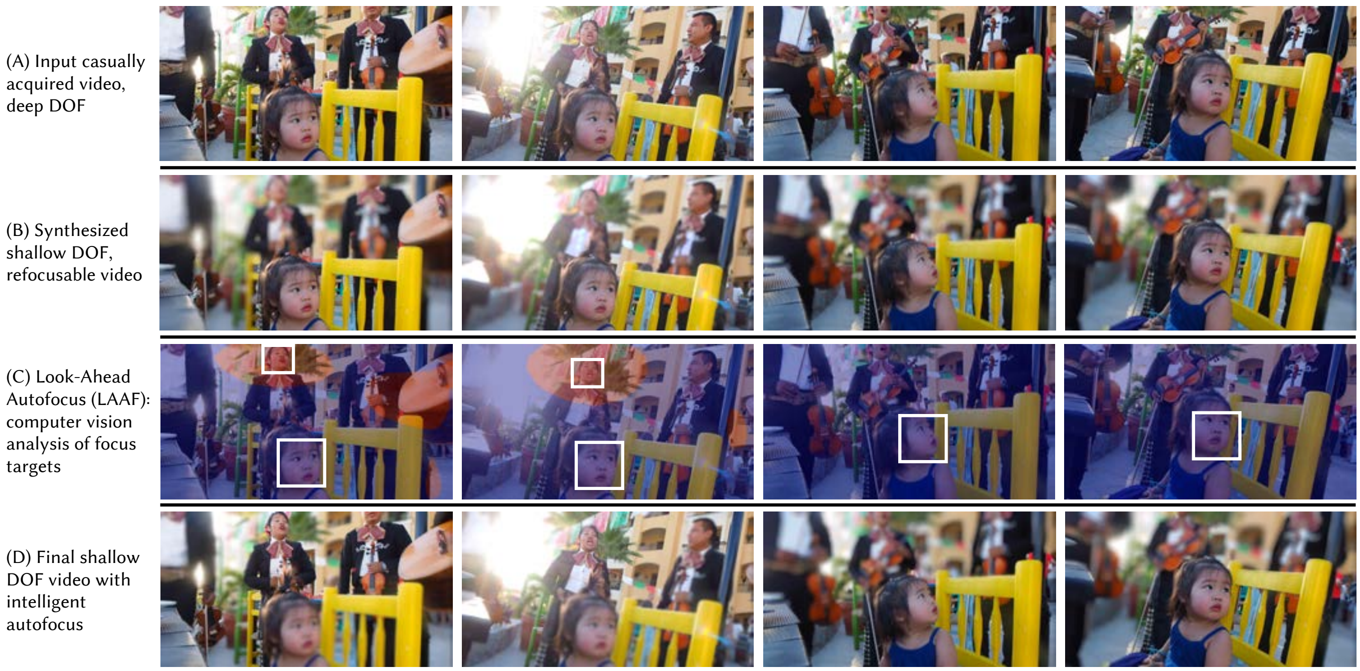

Fig. 1. We present a new approach to pursue cinema-like focus in casual videography, with shallow depth of field (DOF) and accurate focus that isolates the subject. We start with (A) a deep DOF video shot with a small lens aperture. We use a new combination of machine learning, physically-based rendering, and temporal filtering to synthesize (B) a shallow DOF, refocusable video. We also present a novel Look-Ahead Autofocus (LAAF) framework that uses computer vision to (C) analyze upcoming video frames for focus targets. Here, for example, we see face detection (white boxes) and localization of who is speaking/singing [Owens and Efros 2018] (heat map). The result is shallow DOF video (D), where LAAF tracks focus on the singer to start, and transitions focus to the child as the camera pans away from the musicians. The LAAF framework makes future-aware decisions to drive focus tracking and transitions at each frame. This presents a new framework to solve the fundamental realtime limitations of camera-based video autofocus systems.

In cinema, large camera lenses create beautiful shallow depth of field (DOF), but make focusing difficult and expensive. Accurate cinema focus usually relies on a script and a person to control focus in realtime. Casual videographers often crave cinematic focus, but fail to achieve it. We either sacrifice shallow DOF, as in smartphone videos; or we struggle to deliver accurate focus, as in videos from larger cameras. This paper is about a new approach in the pursuit of cinematic focus for casual videography. We present a system that synthetically renders refocusable video from a deep DOF video shot with a smartphone, and analyzes *future* video frames to deliver context-aware autofocus for the current frame. To create refocusable video, we extend recent machine learning methods designed for still photography, contributing a new dataset for machine training, a rendering model better suited to cinema focus, and a filtering solution for temporal coherence. To choose focus accurately for each frame, we demonstrate autofocus that looks at upcoming video frames and applies AI-assist modules such as motion, face, audio and saliency detection. We also show that autofocus benefits from machine learning and a large-scale video dataset with focus annotation, where we use our RVR-LAAF GUI to create this sizable dataset efficiently. We deliver, for example, a shallow DOF video where the autofocus transitions onto each person *before* she begins to speak. This is impossible for

0730-0301/2019/7-ART30 $15.00
https://doi.org/10.1145/3306346.3323015

Visual Smoothness of Polyhedral Surfaces

DAVIDE PELLIS, TU Wien
MARTIN KILIAN, TU Wien
FELIX DELLINGER, TU Wien
JOHANNES WALLNER, TU Graz
HELMUT POTTMANN, KAUST

Representing smooth geometric shapes by polyhedral meshes can be quite difficult in situations where the variation of edges and face normals is prominently visible. Especially problematic are saddle-shaped areas of the mesh, where typical vertices with six incident edges are ill suited to emulate the more symmetric smooth situation. The importance of a faithful discrete representation is apparent for certain special applications like freeform architecture, but is also relevant for simulation and geometric computing.

In this paper we discuss what exactly is meant by a good representation of saddle points, and how this requirement is stronger than a good approximation of a surface plus its normals. We characterize good saddles in terms of the normal pyramid in a vertex.

We show several ways to design meshes whose normals enjoy small variation (implying good saddle points). For this purpose we define a discrete energy of polyhedral surfaces, which is related to a certain total absolute curvature of smooth surfaces. We discuss the minimizers of both functionals and in particular show that the discrete energy is minimal not for triangle meshes, but for principal quad meshes. We demonstrate our procedures for optimization and interactive design by means of meshes intended for architectural design.

CCS Concepts: • **Computing methodologies → Shape modeling**; *Optimization algorithms.*

Additional Key Words and Phrases: Smoothness, polyhedral surfaces, variation of normals, architectural geometry, total absolute curvature

ACM Reference Format:
Davide Pellis, Martin Kilian, Felix Dellinger, Johannes Wallner, and Helmut Pottmann. 2019. Visual Smoothness of Polyhedral Surfaces. *ACM Trans. Graph.* 38, 4, Article 31 (July 2019), 11 pages. https://doi.org/10.1145/3306346.3322975

1 INTRODUCTION

The approximation of a smooth surface by a triangle mesh, or more generally, by a polyhedral mesh, is a basic task of geometry processing. It is a bit surprising that the following aspect has not received more attention in the past, namely the proper representation of saddle-shaped surfaces by meshes. Here 'proper' means that the shape of the immediate neighbourhood of a vertex in the mesh should resemble the shape of a small piece of smooth surface. This criterion can to some extent be expressed by the requirement that dihedral angles in the mesh are small.

Authors' addresses: Davide Pellis, TU Wien; Martin Kilian, TU Wien; Felix Dellinger, TU Wien; Johannes Wallner, TU Graz; Helmut Pottmann, KAUST.

0730-0301/2019/7-ART31 $15.00
https://doi.org/10.1145/3306346.3322975

Fig. 1. A reflective surface reveals the deficiencies in mesh fairness which we quantify by means of a discrete fairness energy based on edge lengths and dihedral angles. This energy its minimal for principal quad meshes (right). In this special case both meshes approximate a reference surface of small total variation of the normal vector field, so all visible deficiencies are not caused by the reference shape, but by the way this shape is meshed.

Meshes which are defective in this respect cause not only problems with the computation of discrete differential quantities, but their visual appearance can be very different from the underlying smooth reference surface. This is particularly true for meshes used for freeform architectural skins with reflective materials.

1.1 Overview and contributions

We start in § 2 with a discussion of the inevitable deficiencies of discrete surfaces which are meant to approximate smooth surfaces. We identify the shape of the normal pyramid as a key difference between the discrete case and the smooth case. We propose that a mesh faithfully models a smooth surface if, in negatively curved areas, it exhibits so-called *good saddles.* We go on to show that then normal pyramids are free of self-intersections. This is a criterion for visual smoothness of the mesh (in particular it implies small dihedral angles). We introduce a class of highly-constrained *principal-symmetric* meshes which in some sense optimally represent saddle-shaped surface geometry. The different mesh properties we study can be symbolically arranged in the following order,

$$\boxed{\text{principal-symmetric meshes}} \Longrightarrow \boxed{\text{meshes with good saddles}} \Longrightarrow \boxed{\text{intersection-free normal pyramids}}$$

with the more restrictive properties to the left.

To quantify the visual smoothness of a mesh, in § 3 we turn to a mesh energy measuring the total variation of the normal vectors. It essentially is the sum of edge lengths times dihedral angles. A small energy implies that normal pyramids cannot be too badly behaved.

Progressive Embedding

HANXIAO SHEN, New York University
ZHONGSHI JIANG, New York University
DENIS ZORIN, New York University
DANIELE PANOZZO, New York University

Tutte embedding is one of the most common building blocks in geometry processing algorithms due to its simplicity and provable guarantees. Although provably correct in infinite precision arithmetic, it fails in challenging cases when implemented using floating point arithmetic, largely due to the induced exponential area changes.

We propose *Progressive Embedding*, with similar theoretical guarantees to Tutte embedding, but more resilient to the rounding error of floating point arithmetic. Inspired by progressive meshes, we collapse edges on an invalid embedding to a valid, simplified mesh, then insert points back while maintaining validity. We demonstrate the robustness of our method by computing embeddings for a large collection of disk topology meshes.

By combining our robust embedding with a variant of the matchmaker algorithm, we propose a general algorithm for the problem of mapping multiply connected domains with arbitrary hard constraints to the plane, with applications in texture mapping and remeshing.

CCS Concepts: • **Computing methodologies → Shape modeling**.

Additional Key Words and Phrases: Embedding, Mesh Parametrization, Robust Geometry Processing, Locally Injective Maps

ACM Reference Format:
Hanxiao Shen, Zhongshi Jiang, Denis Zorin, and Daniele Panozzo. 2019. Progressive Embedding. *ACM Trans. Graph.* 38, 4, Article 32 (July 2019), 13 pages. https://doi.org/10.1145/3306346.3323012

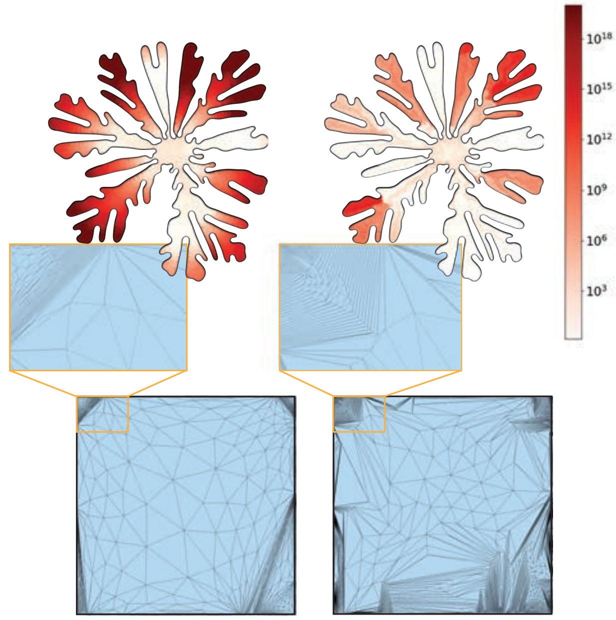

Fig. 1. The Tutte embedding of this Hele-Shaw polygon (left) contains 46 flipped triangles, due to numerical rounding errors. Our progressive embedding (right) produces a valid embedding, without any inverted element and with lower distortion. The colors represent the distortion of the triangles, measured using the symmetric Dirichlet energy.

1 INTRODUCTION

Piecewise linear surface-to-plane maps, or parametrizations, are ubiquitous in computer graphics, geometry processing, mechanical engineering, and scientific visualization. Depending on the applications, the maps are required to exhibit different properties, most commonly, low distortion, local injectivity, and global bijectivity.

The last two properties are challenging to guarantee for discrete maps. Most algorithms with guarantees use Tutte embedding as a component. Tutte embedding is a construction that is guaranteed to create bijective mappings under minimal assumptions, if both domains are simply connected and the target planar domain is convex. However, the guarantee only holds if the computation is performed in arbitrary precision rather than floating point arithmetic, as it is commonly done. Failure due to floating point approximation is not as uncommon as one would assume, as the algorithm is likely to create an extreme variation of scale and aspect ratios in complex mapping cases. To quantitatively evaluate this issue, we computed Tutte embeddings on 2718 models (all the genus 0 models from Thingi10k [Zhou and Jacobson 2016]) using double precision, and observed 80 failures. To the best of our knowledge, this problem has not been addressed before in the literature.

This rate of failure is problematic for batch processing large geometrical collections (for example for processing geometric deep learning datasets) or when the embedding has to be computed many times (for example in cross-parametrization [Kraevoy et al. 2003; Schreiner et al. 2004]). In these scenarios, a failure rate of 2.9% may not be tolerable, since it is not realistic to manually fix hundreds of problematic cases, and if failure happens on large meshes with millions of triangles it might not even be possible to fix them by hand.

A simple solution to this problem is the use of multi-precision (or rational) arithmetic [Granlund 2018]: if enough bits are used to represent the mantissa and exponent of the floating point representation, Tutte embedding will succeed, since the solution of a linear system can be computed exactly. However, the result in high precision is not directly usable by downstream applications, and requires to be rounded (or "snapped" [Halperin and Packer 2002]) to floating point coordinates. This is a surprisingly challenging problem for which, to the best of our knowledge, no solution applicable to our setting exists (Section 3.1).

Authors' addresses: Hanxiao Shen, New York University, hanxiao@cs.nyu.edu; Zhongshi Jiang, New York University, zhongshi@cims.nyu.edu; Denis Zorin, New York University, dzorin@cs.nyu.edu; Daniele Panozzo, panozzo@nyu.edu, New York University.

0730-0301/2019/7-ART32 $15.00
https://doi.org/10.1145/3306346.3323012

Atlas Refinement with Bounded Packing Efficiency

HAO-YU LIU, University of Science and Technology of China, China
XIAO-MING FU*, University of Science and Technology of China, China
CHUNYANG YE, University of Science and Technology of China, China
SHUANGMING CHAI, University of Science and Technology of China, China
LIGANG LIU, University of Science and Technology of China, China

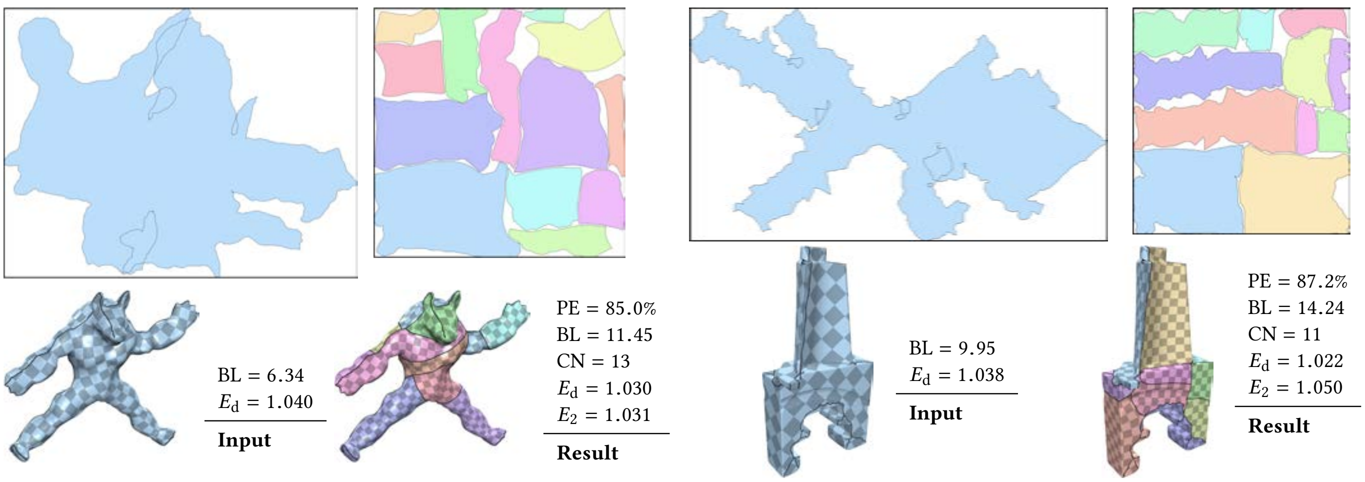

Fig. 1. Refining input parameterized charts to obtain packing efficiency that is greater than or equal to the given bounds. Compared to the inputs that have low packing efficiency and are not bijective, our method significantly improves the packing efficiency and ensures the bijection. We set the packing efficiency bounds as 80% for these two models. PE, BL, CN, E_d, and E_2 represent the packing efficiency, the boundary length, the number of charts, the symmetric Dirichlet distortion metric with respect to the input 3D surface, and the symmetric Dirichlet distortion metric with respect to the input atlas, respectively.

We present a novel algorithm to refine an input atlas with bounded packing efficiency. Central to this method is the use of the axis-aligned structure that converts the general polygon packing problem to a rectangle packing problem, which is easier to achieve high packing efficiency. Given a parameterized mesh with no flipped triangles, we propose a new angle-driven deformation strategy to transform it into a set of axis-aligned charts, which can be decomposed into rectangles by the motorcycle graph algorithm. Since motorcycle graphs are not unique, we select the one balancing the trade-off between the packing efficiency and chart boundary length, while maintaining bounded packing efficiency. The axis-aligned chart often contains greater distortion than the input, so we try to reduce the distortion while bounding the packing efficiency and retaining bijection. We demonstrate the efficacy of our method on a data set containing over five thousand complex models. For all models, our method is able to produce packed atlases with bounded packing efficiency; for example, when the packing efficiency bound is set to 80%, we elongate the boundary length by an average of 78.7% and increase the distortion by an average of 0.0533%. Compared to state-of-the-art methods, our method is much faster and achieves greater packing efficiency.

CCS Concepts: • **Computing methodologies → Shape modeling**.

Additional Key Words and Phrases: atlas refinement, bounded packing efficiency, axis-aligned chart

ACM Reference Format:
Hao-Yu Liu, Xiao-Ming Fu, Chunyang Ye, Shuangming Chai, and Ligang Liu. 2019. Atlas Refinement with Bounded Packing Efficiency. *ACM Trans. Graph.* 38, 4, Article 33 (July 2019), 13 pages. https://doi.org/10.1145/3306346.3323001

*The corresponding author

Authors' addresses: Hao-Yu Liu, University of Science and Technology of China, China, optexeon@mail.ustc.edu.cn; Xiao-Ming Fu, University of Science and Technology of China, China, fuxm@ustc.edu.cn; Chunyang Ye, University of Science and Technology of China, China, yechyang@mail.ustc.edu.cn; Shuangming Chai, University of Science and Technology of China, China, kfckfckf@mail.ustc.edu.cn; Ligang Liu, University of Science and Technology of China, China, lgliu@ustc.edu.cn.

0730-0301/2019/7-ART33 $15.00
https://doi.org/10.1145/3306346.3323001

1 INTRODUCTION

In computer graphics, atlases are constructed by packing 2D parameterized charts into rectangular texture image domains. They are commonly used to store surface signals, such as colors, normals, and textures. Due to the regularity of texture image grids, they are beneficial for some geometric processing tasks, such as signal smoothing and sharpening, texture stitching, geodesic distance computation, and line integral convolution [Prada et al. 2018].

A good atlas usually satisfies the following properties: (1) the parameterization of each chart is bijective and contains as little as possible isometric distortion; (2) the packed parameterized charts

Weaving Geodesic Foliations

JOSH VEKHTER, The University of Texas at Austin
JIACHENG ZHUO, The University of Texas at Austin
LUISA F GIL FANDINO, The University of Texas at Austin
QIXING HUANG, The University of Texas at Austin
ETIENNE VOUGA, The University of Texas at Austin

Fig. 1. We present new algorithms for computing geodesic foliations on discrete surfaces. A geodesic foliation on a six-fold branched cover of the Stanford bunny (*left*) describes the ribbon layout of a triaxial weave (*center*), which we fabricate out of birch veneer (*right*).

We study discrete geodesic foliations of surfaces—foliations whose leaves are all approximately geodesic curves—and develop several new variational algorithms for computing such foliations. Our key insight is a relaxation of vector field integrability in the discrete setting, which allows us to optimize for curl-free unit vector fields that remain well-defined near singularities and robustly recover a scalar function whose gradient is well aligned to these fields. We then connect the physics governing surfaces woven out of thin ribbons to the geometry of geodesic foliations, and present a design and fabrication pipeline for approximating surfaces of arbitrary geometry and topology by triaxially-woven structures, where the ribbon layout is determined by a geodesic foliation on a sixfold branched cover of the input surface. We validate the effectiveness of our pipeline on a variety of simulated and fabricated woven designs, including an example for readers to try at home.

CCS Concepts: • **Computing methodologies → Mesh geometry models**; Continuous simulation; • **Mathematics of computing →** *Geometric topology*.

Authors' addresses: Josh Vekhter, The University of Texas at Austin, 2317 Speedway, Austin, TX, 78712; Jiacheng Zhuo, The University of Texas at Austin; Luisa F Gil Fandino, The University of Texas at Austin; Qixing Huang, The University of Texas at Austin; Etienne Vouga, The University of Texas at Austin.

0730-0301/2019/7-ART34 $15.00
https://doi.org/10.1145/3306346.3323043

Additional Key Words and Phrases: Discrete Differential Geometry. Vector Field Design. Digital Fabrication. Texture Synthesis.

ACM Reference Format:
Josh Vekhter, Jiacheng Zhuo, Luisa F Gil Fandino, Qixing Huang, and Etienne Vouga. 2019. Weaving Geodesic Foliations. *ACM Trans. Graph.* 38, 4, Article 34 (July 2019), 22 pages. https://doi.org/10.1145/3306346.3323043

1 INTRODUCTION

Birds in the family *Ploceidae* have been weaving their nests since the Miocene; humans likewise have been using weaving to construct baskets, mats, and other two and three-dimensional aspects of the built environment for at least 10,000 years. More recently, composites woven out of non-traditional materials like carbon fiber and shape memory alloy have been gaining popularity in a number of domains, where improved material properties, as well as cost-effectiveness and the increasing capability of looms to weave in three dimensions, make it a particularly attractive technique for fabrication. Figure 2 shows some modern applications of weaving to architecture, art, chemical engineering, and medicine.

Away from singularities, fibers of the weave can be grouped into nearly-parallel families—two families at right angles in the case of simple weaves (plain, twill) and three interleaved at sixty degree angles for *triaxial* weaves. In their simplest form, all of these weave patterns are planar; curvature can be introduced into the woven structure by varying the spacing between consecutive members of the ribbon families, varying the angles at which two families cross, or inserting singularities or *dislocations* into the weave pattern. Such

Gaussian-Product Subdivision Surfaces

REINHOLD PREINER, Graz University of Technology, Austria
TAMY BOUBEKEUR, LTCI, Telecom ParisTech, Paris-Saclay University, France
MICHAEL WIMMER, TU Wien, Austria

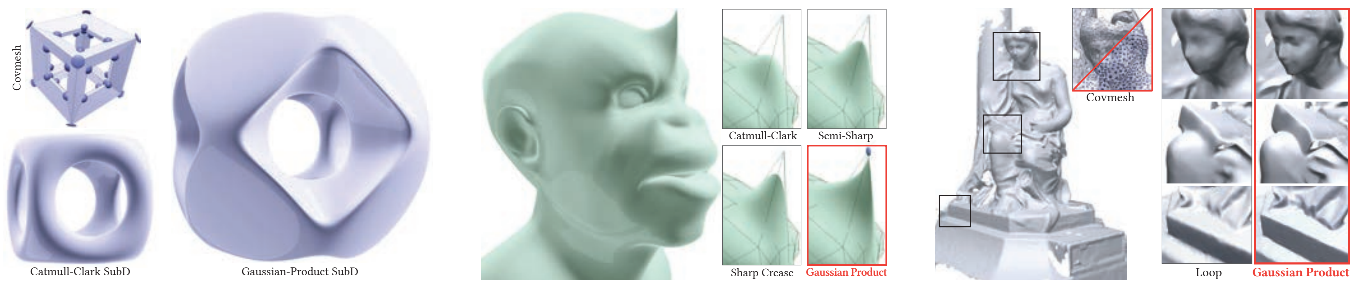

(a) Covariance Meshes (b) Surface Modeling (c) Reconstruction

Fig. 1. We introduce a new nonlinear subdivision surface model, which is based on a control-mesh representation encoding Gaussian covariances in its vertices (a). Our surface definition relies on a refinement using Gaussian products, providing enhanced shape control on the smooth limit surface. This significantly widens the space of possible shapes stemming from a given control mesh, enabling better modeling features, (semi-)sharpness and concavities without changing the base connectivity (b). In addition, our representation naturally integrates into surface-reconstruction pipelines, recovering high- and mid-frequency structures even from a low-resolution control mesh (c).

Probabilistic distribution models like Gaussian mixtures have shown great potential for improving both the quality and speed of several geometric operators. This is largely due to their ability to model large fuzzy data using only a reduced set of atomic distributions, allowing for large compression rates at minimal information loss. We introduce a new surface model that utilizes these qualities of Gaussian mixtures for the definition and control of a parametric smooth surface. Our approach is based on an enriched mesh data structure, which describes the probability distribution of spatial surface locations around each vertex via a Gaussian covariance matrix. By incorporating this additional covariance information, we show how to define a smooth surface via a nonlinear probabilistic subdivision operator based on products of Gaussians, which is able to capture rich details at fixed control mesh resolution. This entails new applications in surface reconstruction, modeling, and geometric compression.

CCS Concepts: • **Computing methodologies → Parametric curve and surface models**; *Mixture models*; *Mesh geometry models.*

Additional Key Words and Phrases: Gaussian mixtures, covariance mesh, subdivision surfaces, triangulation

ACM Reference Format:
Reinhold Preiner, Tamy Boubekeur, and Michael Wimmer. 2019. Gaussian-Product Subdivision Surfaces. *ACM Trans. Graph.* 38, 4, Article 35 (July 2019), 11 pages. https://doi.org/10.1145/3306346.3323026

Authors' addresses: Reinhold Preiner, r.preiner@cgv.tugraz.at, Graz University of Technology, CGV, Graz, Austria; Tamy Boubekeur, LTCI, Telecom ParisTech, Paris-Saclay University, IDS, Paris, France; Michael Wimmer, TU Wien, Vienna, Austria.

0730-0301/2019/7-ART35
https://doi.org/10.1145/3306346.3323026

1 INTRODUCTION

For the efficient processing of fuzzy geometric data like noisy point sets, probabilistic distribution models such as Gaussian mixtures have recently shown great potential for tasks like registration, filtering or resampling. This is largely due to their ability to model large fuzzy data using only a reduced set of atomic distributions, allowing for large compression rates at minimal information loss. Due to this compactness, it is desirable to be able to define a surface directly on such a sparse model, and avoid the need to expand to larger representations (e.g., meshes or point clouds) for further processing and rendering. In fact, such a sparse representation is also highly interesting for modeling applications, which aim at defining complex shapes using simple base representations. So far, there have been some attempts to define a probabilistic surface along the ridge contour of the probability density function (pdf) of a dense Gaussian mixture. However, this contour degenerates when the mixture is compressed to large anisotropic Gaussians, where discontinuities appear; also, the resulting surface definition is not amenable for further processing or modeling tasks.

In this paper, we introduce a new probabilistic surface representation that allows defining continuous, artifact-free surfaces even for a sparse set of Gaussians, while still closely resembling the ridge of their pdf. Our surface definition is based on a polygonization of the individual Gaussian components, resulting in a new, enriched mesh model that carries anisotropic covariance information at its vertices, called *covariance mesh*. The key idea to turn this representation into a continuous surface is to depart from the linear combination of individual Gaussians used in mixture models, and instead consider their *joint probability*, leading to a new interpolation method between Gaussians based on a product formulation. To apply this

Hierarchical Russian Roulette for Vertex Connections

YUSUKE TOKUYOSHI, SQUARE ENIX CO., LTD., Japan
TAKAHIRO HARADA, Advanced Micro Devices, Inc., USA

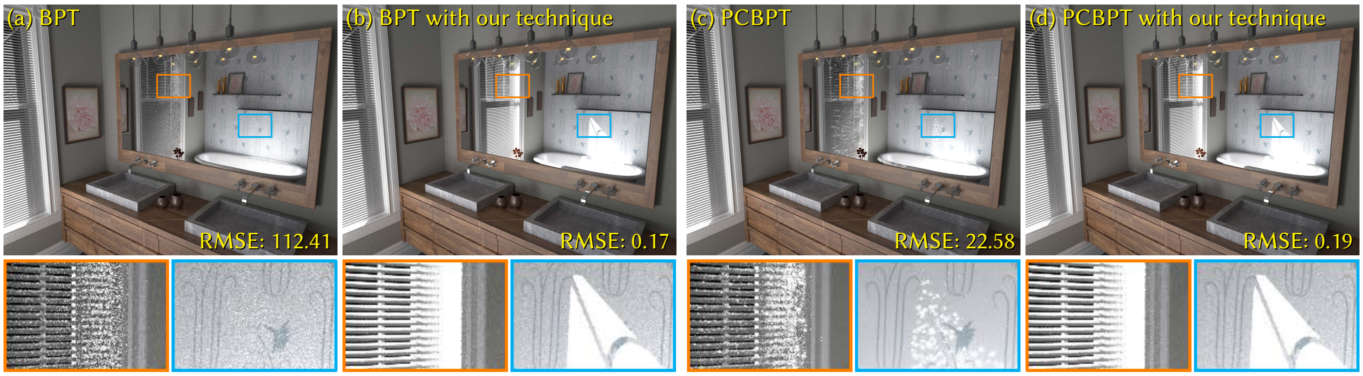

Fig. 1. Equal-time (15 min) comparison of rendering with and without our connection technique (1600×1200 screen resolution). This Bathroom scene is lit by caustics, and the BRDF of the mirror is extremely glossy (GGX roughness: 0.0001). BPT (a) and probabilistic connection BPT (PCBPT) (c) produce intensive variance for specular-diffuse-glossy and glossy-diffuse-glossy reflections shown in closeups. Our technique reduces this variance significantly (b, d).

While bidirectional path tracing is a well-established light transport algorithm, many samples are required to obtain high-quality results for specular-diffuse-glossy or glossy-diffuse-glossy reflections especially when they are highly glossy. To improve the efficiency for such light path configurations, we propose a *hierarchical Russian roulette* technique for vertex connections. Our technique accelerates a huge number of Russian roulette operations according to an approximate scattering lobe at an eye-subpath vertex for many cached light-subpath vertices. Our method dramatically reduces the number of random number generations needed for Russian roulette by introducing a hierarchical rejection algorithm which assigns random numbers in a top-down fashion. To efficiently reject light vertices in each hierarchy, we also introduce an efficient approximation of anisotropic scattering lobes used for the probability of Russian roulette. Our technique is easy to integrate into some existing bidirectional path tracing-based algorithms which cache light-subpath vertices (e.g., probabilistic connections, and vertex connection and merging). In addition, unlike existing many-light methods, our method does not restrict multiple importance sampling strategies thanks to the simplicity of Russian roulette. Although the proposed technique does not support perfectly specular surfaces, it significantly improves the efficiency for caustics reflected on extremely glossy surfaces in an unbiased fashion.

CCS Concepts: • **Computing methodologies → Ray tracing**.

Additional Key Words and Phrases: global illumination, bidirectional path tracing, vertex connection, Russian roulette, light culling

Authors' addresses: Yusuke Tokuyoshi, SQUARE ENIX CO., LTD., 6-27-30 Shinjuku, Shinjuku-ku, Tokyo, 160-8430, Japan, yusuke.tokuyoshi@gmail.com; Takahiro Harada, Advanced Micro Devices, Inc. USA, Takahiro.Harada@amd.com.

0730-0301/2019/7-ART36 $15.00
https://doi.org/10.1145/3306346.3323018

ACM Reference Format:
Yusuke Tokuyoshi and Takahiro Harada. 2019. Hierarchical Russian Roulette for Vertex Connections. *ACM Trans. Graph.* 38, 4, Article 36 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3323018

1 INTRODUCTION

Monte Carlo light transport simulation is widely used for photorealistic rendering nowadays, however developing a robust algorithm for various scenes is still a challenging problem. Bidirectional path tracing (BPT) [Lafortune and Willems 1993; Veach and Guibas 1994] is a well-established light transport algorithm which constructs various paths by connecting subpaths traced from a light source and eye (i.e., *light subpath* and *eye subpath*). By using multiple importance sampling (MIS) [Veach and Guibas 1995], the combination of such path sampling techniques reduces the estimation error significantly. However, BPT produces high variance for specular-diffuse-glossy (SDG) or glossy-diffuse-glossy (GDG) paths (Fig. 1a) because of a lack of sampling techniques suitable for these path configurations. Hence, a large number of samples are necessary to render high-quality images for highly glossy surfaces. Probabilistic connections [Popov et al. 2015] increase the sample count by reusing hundreds of light subpaths, however it is still insufficient and produces splotch-like artifacts due to correlated variance (Fig. 1c). This paper tackles the problem of reusing millions of light subpaths for BPT.

To improve the efficiency of connections in BPT, we propose an acceleration technique for *Russian roulette* operations [Arvo and Kirk 1990] according to an approximate scattering lobe at an eye vertex (i.e., vertex of an eye subpath). Russian roulette was used to reject unimportant connections to reduce the number of shadow rays [Veach 1998]. However, this stochastic rejection was still expensive when many light subpaths were cached and reused for every eye vertex to increase the number of path samples. This is because many light vertices (i.e., vertices of light subpaths) are candidates to

Optimal Multiple Importance Sampling

IVO KONDAPANENI*, Charles University, Prague
PETR VÉVODA*, Charles University, Prague and Render Legion, a. s.
PASCAL GRITTMANN, Saarland University
TOMÁŠ SKŘIVAN, IST Austria
PHILIPP SLUSALLEK, Saarland University and DFKI
JAROSLAV KŘIVÁNEK, Charles University, Prague and Render Legion, a. s.

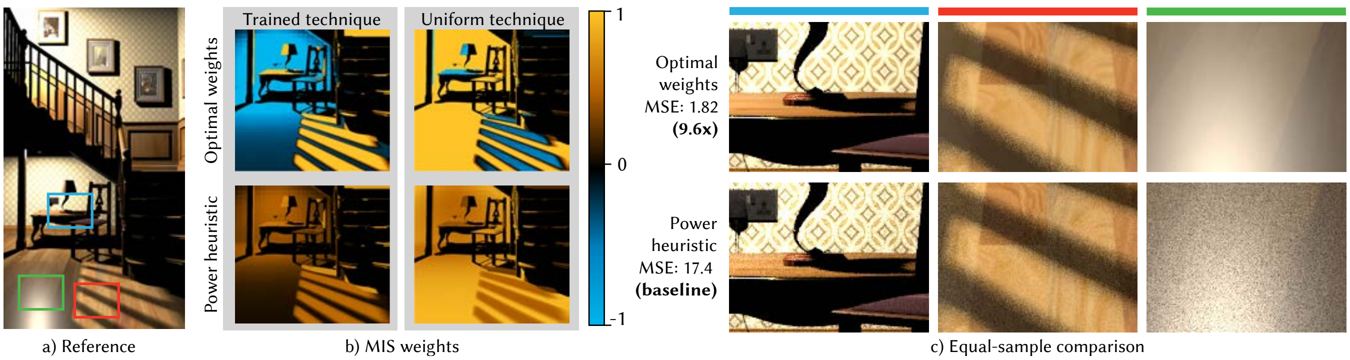

Fig. 1. Equal-sample comparison (20 per technique per pixel) of direct illumination estimated by an MIS combination of two light sampling techniques (*Trained* and *Uniform*, see Sec. 8.2 for details) with our optimal weights (top row) and the power heuristic (bottom row). The false-color images b) show per-pixel average MIS weight values as determined by the two weighting strategies. Unlike any of the existing MIS weighting heuristics, the optimal weights can have *negative* values, which provides additional opportunity for variance reduction, leading to an overall 9.6 times lower error per sample taken than the power heuristic in this scene.

Multiple Importance Sampling (MIS) is a key technique for achieving robustness of Monte Carlo estimators in computer graphics and other fields. We derive optimal weighting functions for MIS that provably minimize the variance of an MIS estimator, given a set of sampling techniques. We show that the resulting variance reduction over the balance heuristic can be higher than predicted by the variance bounds derived by Veach and Guibas, who assumed only non-negative weights in their proof. We theoretically analyze the variance of the optimal MIS weights and show the relation to the variance of the balance heuristic. Furthermore, we establish a connection between the new weighting functions and control variates as previously applied to mixture sampling. We apply the new optimal weights to integration problems in light transport and show that they allow for new design considerations when choosing the appropriate sampling techniques for a given integration problem.

Authors' addresses: Ivo Kondapaneni, Charles University, Prague; Petr Vévoda, Charles University, Prague, Render Legion, a. s. Prague; Pascal Grittmann, Saarland University, Saarbrücken; Tomáš Skřivan, IST Austria, Vienna; Philipp Slusallek, Saarland University, Saarbrücken, DFKI, Saarbrücken; Jaroslav Křivánek, Charles University, Prague, Render Legion, a. s. Prague.
*Ivo Kondapaneni and Petr Vévoda share the first authorship of this work.

0730-0301/2019/7-ART37 $15.00
https://doi.org/10.1145/3306346.3323009

CCS Concepts: • **Mathematics of computing → Probability and statistics**; • **Computing methodologies → Rendering**.

Additional Key Words and Phrases: Monte Carlo integration, Multiple Importance Sampling, combined estimators

ACM Reference Format:
Ivo Kondapaneni, Petr Vévoda, Pascal Grittmann, Tomáš Skřivan, Philipp Slusallek, and Jaroslav Křivánek. 2019. Optimal Multiple Importance Sampling. *ACM Trans. Graph.* 38, 4, Article 37 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3323009

1 INTRODUCTION

Monte Carlo (MC) integration is an essential tool in light transport simulation [Pharr et al. 2016; Veach 1997] and other fields of science and engineering [Kalos and Whitlock 2008]. An inherent problem of MC integration is its slow convergence, which is why numerous variance reduction schemes have been proposed, notably importance sampling. Its extension, known as *multiple importance sampling* (MIS) [Veach and Guibas 1995], is particularly versatile as it enables combining different sampling techniques in a robust way to form better MC estimates.

In the context of light transport simulation, MIS has served as a cornerstone for robust bidirectional path sampling [Georgiev et al. 2012a; Hachisuka et al. 2012; Křivánek et al. 2014; Popov et al. 2015; Veach and Guibas 1995], Markov chain Monte Carlo light transport [Gruson et al. 2016; Hachisuka et al. 2014; Šik et al. 2016], adaptive path sampling (path guiding) [Herholz et al. 2016; Müller et al. 2017; Vorba et al. 2014], or in isolated integration problems such as direct

Ellipsoidal Path Connections for Time-Gated Rendering

ADITHYA PEDIREDLA, Rice University, USA and Carnegie Mellon University, USA
ASHOK VEERARAGHAVAN, Rice University, USA
IOANNIS GKIOULEKAS, Carnegie Mellon University, USA

During the last decade, we have been witnessing the continued development of new time-of-flight imaging devices, and their increased use in numerous and varied applications. However, physics-based rendering techniques that can accurately simulate these devices are still lacking: while existing algorithms are adequate for certain tasks, such as simulating transient cameras, they are very inefficient for simulating time-gated cameras because of the large number of wasted path samples. We take steps towards addressing these deficiencies, by introducing a procedure for efficiently sampling paths with a predetermined length, and incorporating it within rendering frameworks tailored towards simulating time-gated imaging. We use our open-source implementation of the above to empirically demonstrate improved rendering performance in a variety of applications, including simulating proximity sensors, imaging through occlusions, depth-selective cameras, transient imaging in dynamic scenes, and non-line-of-sight imaging.

CCS Concepts: • **Computing methodologies → Ray tracing**.

Additional Key Words and Phrases: time-gated cameras, transient imaging.

ACM Reference Format:
Adithya Pediredla, Ashok Veeraraghavan, and Ioannis Gkioulekas. 2019. Ellipsoidal Path Connections for Time-Gated Rendering. *ACM Trans. Graph.* 38, 4, Article 38 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3323016

1 INTRODUCTION

Time-of-flight (ToF) imaging is revolutionizing a large array of applications including robotics, autonomous navigation, atmospheric sciences, medicine, and even entertainment and human-computer interaction. Unlike conventional intensity sensors, ToF sensors exploit the fact that the speed of light is finite, and record information about the time it takes for photons to travel from a light source to the sensor. This information can then be used to recover geometric and material properties of the scene the photon traveled through, including even parts of the scene that are not directly visible to the sensor because of opaque or translucent occluders.

Within computer graphics, imaging and rendering have consistently advanced in lockstep: As new imaging sensors are being developed, rendering algorithms that can accurately and efficiently simulate measurements of such sensors quickly follow. In turn, the availability of rendering algorithms greatly enhances the ability of the sensor designers to optimize their designs, resulting in novel sensors of superior performance. In particular, as state-of-the-art machine learning tools are increasingly used for both designing the sensors [Marco et al. 2017a] and processing their outputs [Gruber et al. 2019]. Efficient rendering tools are necessary to generate the large, diverse, and realistic datasets needed for training.

Existing physically-accurate rendering algorithms are primarily tailored to conventional, steady-state rendering. As a consequence, they can be suboptimal for simulating ToF sensors, with their performance varying considerably between different ToF rendering tasks. To characterize the performance characteristics of existing algorithms, we will broadly classify ToF rendering tasks into two categories. The first category includes tasks such as simulating *continuous-wave time-of-flight cameras*, which accumulate all photons with a weight that depends on their time of travel; as well as *transient cameras*, which aggregate all photons but separate them into a sequence of images, each recording contributions only from photons with a specific time of travel. Existing steady-state rendering algorithms remain efficient for simulation tasks in this category, as the majority of paths they generate will have a non-zero contribution, regardless of their pathlength. Tasks in this category have generally been the main focus of previous research on ToF rendering; for instance, Jarabo et al. [Jarabo et al. 2014] improve rendering performance by introducing path-sampling schemes and reconstruction techniques tailored to the transient imaging setting.

The second type of ToF rendering tasks involves simulating images that accumulate contributions only from a small subset of photons, whose time of travel is within some narrow interval. These *time-gated* rendering tasks arise in a large number of practical situations; examples include time-gated sensors used as proximity detectors, gated laser ranging cameras, as well as situations where transient imaging is performed in dynamic scenes such as outdoors environments or tissue with blood flow. Unfortunately, existing rendering algorithms cannot be used for efficient time-gated rendering: the vast majority of the paths generated by these algorithms end up being rejected, for having length outside the narrow range accumulated by the simulated sensor.

Given this challenge and the importance of the problem, we focus on developing efficient Monte Carlo algorithms for time-gated rendering tasks. At the core of the challenge is the fact that current physics-based rendering algorithms cannot sample paths that satisfy specific constraints on their length. We take first steps towards addressing this fundamental limitation, by making the following technical contributions: We develop a path-sampling technique that generates paths with a predetermined target length. This technique is a variant of standard bidirectional path tracing [Lafortune and Willems 1996; Veach and Guibas 1995a], which we term *bidirectional path tracing with ellipsoidal connections*. This allows us to simulate time-gated sensors without having to reject an excessive number of sample paths. We develop the mathematical machinery

Authors' addresses: Adithya Pediredla, Rice University, USA, adithya.k.pediredla@rice.edu, Carnegie Mellon University, USA, apedired@andrew.cmu.edu; Ashok Veeraraghavan, Rice University, USA, vashok@rice.edu; Ioannis Gkioulekas, Carnegie Mellon University, USA, igkioule@andrew.cmu.edu.

0730-0301/2019/7-ART38 $15.00
https://doi.org/10.1145/3306346.3323016

A Monte Carlo Framework for Rendering Speckle Statistics in Scattering Media

CHEN BAR, Department of Electrical Engineering, Technion, Israel
MARINA ALTERMAN, Department of Electrical Engineering, Technion, Israel
IOANNIS GKIOULEKAS, Robotics Institute, Carnegie Mellon University, USA
ANAT LEVIN, Department of Electrical Engineering, Technion, Israel

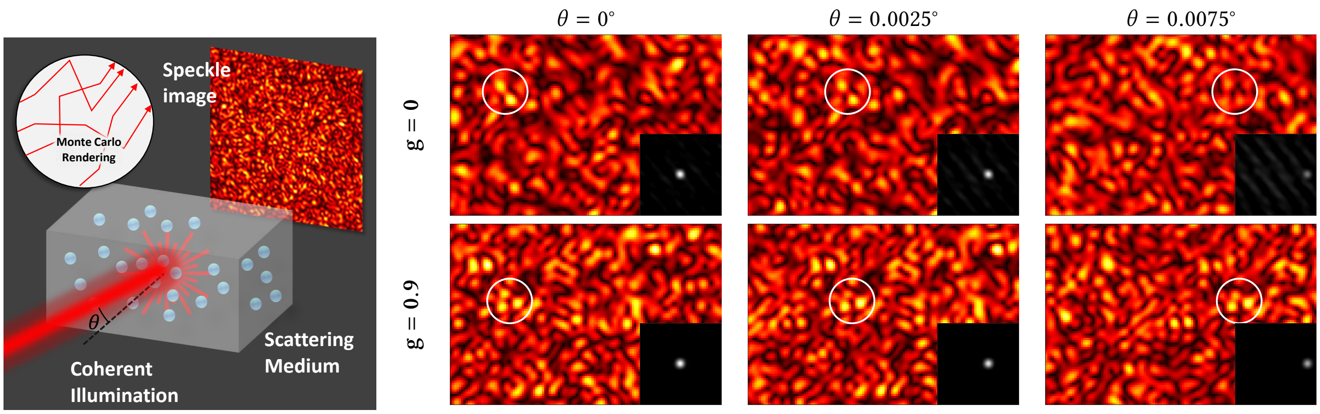

Fig. 1. **Simulation of memory effect in scattering.** Coherent images of translucent materials typically involve highly-fluctuating speckle structure. Despite their semi-random structure, speckles have strong statistical properties. For example, the memory effect property states that, as one tilts the illumination direction (setup at left), the resulting speckles shift. This property is at the core of multiple computational imaging applications. The memory effect is valid over a limited angular range that depends on material properties. Due to the absence of analytical formulas, it is generally necessary to measure this angular range for materials of interest empirically in the lab. We present a Monte Carlo rendering approach for simulating physically-accurate speckle images, as well as their statistics, as a function of material parameters. The figure shows speckle images rendered by our algorithm for a few illumination directions, as well as their auto-correlation (black insets), demonstrating the speckle shift property. As the angle difference increases, the correlation decays, and the decay rate is different for different material parameters—in this case, materials with Henyey-Greenstein (HG) phase functions of different parameters g. For the isotropic scattering case, $g = 0$, the pattern similarity is lost at the third column, whereas for the forward scattering case, $g = 0.9$, correlation is preserved. We verify the accuracy of our algorithm against an exact, yet computationally heavy, wave solver, as well as against analytical formulas derived under limiting assumptions.

We present a Monte Carlo rendering framework for the physically-accurate simulation of speckle patterns arising from volumetric scattering of coherent waves. These noise-like patterns are characterized by strong statistical properties, such as the so-called memory effect. These properties are at the core of imaging techniques for applications as diverse as tissue imaging, motion tracking, and non-line-of-sight imaging. Our rendering framework can replicate these properties computationally, in a way that is orders of magnitude more efficient than alternatives based on directly solving the wave equations. At the core of our framework is a path-space formulation for the covariance of speckle patterns arising from a scattering volume, which we derive from first principles. We use this formulation to develop two Monte Carlo rendering algorithms, for computing speckle covariance as well as directly speckle fields. While approaches based on wave equation solvers require knowing the microscopic position of wavelength-sized scatterers, our approach takes as input only bulk parameters describing the statistical distribution of these scatterers inside a volume. We validate the accuracy of our framework by comparing against speckle patterns simulated using wave equation solvers, use it to simulate memory effect observations that were previously only possible through lab measurements, and demonstrate its applicability for computational imaging tasks.

CCS Concepts: • **Computing methodologies → Ray tracing**.

Additional Key Words and Phrases: Scattering, Speckle statistics.

ACM Reference Format:
Chen Bar, Marina Alterman, Ioannis Gkioulekas, and Anat Levin. 2019. A Monte Carlo Framework for Rendering Speckle Statistics in Scattering Media. *ACM Trans. Graph.* 38, 4, Article 39 (July 2019), 22 pages. https://doi.org/10.1145/3306346.3322950

Authors' addresses: Chen Bar, Department of Electrical Engineering, Technion, Haifa, Israel; Marina Alterman, Department of Electrical Engineering, Technion, Haifa, Israel; Ioannis Gkioulekas, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA; Anat Levin, Department of Electrical Engineering, Technion, Haifa, Israel.

0730-0301/2019/7-ART39 $15.00
https://doi.org/10.1145/3306346.3322950

Rodent: Generating Renderers without Writing a Generator

ARSÈNE PÉRARD-GAYOT, Saarland University
RICHARD MEMBARTH, DFKI & Saarland University
ROLAND LEISSA, Saarland University
SEBASTIAN HACK, Saarland University
PHILIPP SLUSALLEK, DFKI & Saarland University

Monte-Carlo Renderers must generate many color samples to produce a noise-free image, and for each of those, they must evaluate complex mathematical models representing the appearance of the objects in the scene. These models are usually in the form of shaders: Small programs that are executed during rendering in order to compute a value for the current sample.

Renderers often compile and optimize shaders just before rendering, taking advantage of the knowledge of the scene. In principle, the entire renderer could benefit from a-priori code generation. For instance, scheduling can take advantage of the knowledge of the scene in order to maximize hardware usage. However, writing such a configurable renderer eventually means writing a compiler that translates a scene description into machine code.

In this paper, we present a framework that allows generating entire renderers for CPUs and GPUs without having to write a dedicated compiler: First, we provide a rendering library in a functional/imperative language that elegantly abstracts the individual rendering concepts using higher-order functions. Second, we use *partial evaluation* to combine and specialize the individual components of a renderer according to a particular scene.

Our results show that the renderers we generate outperform equivalent high-performance implementations written with state-of-the-art ray tracing libraries on the CPU and GPU.

CCS Concepts: • **Computing methodologies → Rendering**; • **Software and its engineering → Compilers**.

Additional Key Words and Phrases: Rendering, Ray-tracing, Generator, Partial Evaluation

ACM Reference Format:
Arsène Pérard-Gayot, Richard Membarth, Roland Leißa, Sebastian Hack, and Philipp Slusallek. 2019. Rodent: Generating Renderers without Writing a Generator. *ACM Trans. Graph.* 38, 4, Article 40 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3322955

Authors' addresses: Arsène Pérard-Gayot, Saarland University, Saarland Informatics Campus E1 1, Saarbrücken, Germany, 66123, perard@cg.uni-saarland.de; Richard Membarth, DFKI & Saarland University, Saarland Informatics Campus D3 2, Saarbrücken, Germany, 66123, membarth@dfki.de; Roland Leißa, Saarland University, Saarland Informatics Campus E1 3, Saarbrücken, Germany, 66123, leissa@cs.uni-saarland.de; Sebastian Hack, Saarland University, Saarland Informatics Campus E1 3, Saarbrücken, Germany, 66123, hack@cs.uni-saarland.de; Philipp Slusallek, DFKI & Saarland University, Saarland Informatics Campus D3 2, Saarbrücken, Germany, 66123, slusallek@dfki.de.

0730-0301/2019/7-ART40 $15.00
https://doi.org/10.1145/3306346.3322955

1 INTRODUCTION

Monte-Carlo Renderers are naturally complex and configurable pieces of software: They typically incorporate powerful renderer-specific languages to describe 3D scenes. Most of those languages allow the designer to define small programs that define the appearance of an object, called *shaders* [Pixar 1988; Sony Pictures Imageworks 2017]. Before rendering, the renderer *compiles* these shader programs into either machine code or an internal representation. During shader compilation, most renderers perform standard code optimizations, such as constant folding or dead code elimination. At this stage, some renderers also apply *specialization* [Guenter et al. 1995; McCool et al. 2002; Sons et al. 2014]: Since the scene is known, some computations in the shaders might have become superfluous and can thus be eliminated. For instance, the renderer can avoid fetching and interpolating vertex attributes if those attributes are not present or unused, or remove redundant texture lookups.

Even when shaders are specialized, concrete implementations of rendering systems suffer from the same problem as many other high-performance codes: In terms of software engineering, all components (e.g. shading system, integrator, primitive intersection and acceleration structure traversal) are ideally decoupled from each other and implemented in a high-level, abstract way that focuses on the algorithmic properties to make code reusable and easier to understand. However, to achieve this decoupling and abstraction one typically has to pay a performance price. Essentially, every code abstraction technique that the programmer uses needs a compiler optimization to remove it. Because compilers usually do not succinctly remove all of these abstractions, programmers sacrifice genericity for performance. Code is specialized manually and generic code is tainted by target-dependent artifacts: Developers use SIMD intrinsics to vectorize their algorithms, or write CUDA code that can only run on a GPU.

A solution to this problem is to optimize across rendering library layers. If the scene description is known at every layer during compilation, constants can be folded, branches pruned, and the renderer becomes more efficient, because it is specialized for that scene. Of course, not all scene parameters will enable performance critical optimizations, and some parameters may be unknown at compile-time, like the camera position. In practice, knowing the *type* of each object is enough. However, to generate a specialized renderer from a scene description would normally require to write a dedicated, *domain-specific compiler* that can not only specialize shaders, but also the other components. From the scene description and rendering parameters, this compiler would generate machine code that corresponds to a renderer in that specific configuration.

Interactive Hand Pose Estimation using a Stretch-Sensing Soft Glove

OLIVER GLAUSER and SHIHAO WU, ETH Zurich, Switzerland
DANIELE PANOZZO, New York University, USA
OTMAR HILLIGES and OLGA SORKINE-HORNUNG, ETH Zurich, Switzerland

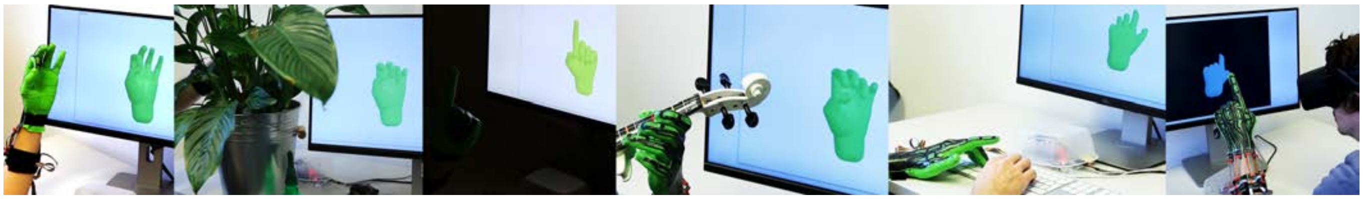

Fig. 1. Our stretch-sensing soft glove captures hand poses in real time and with high accuracy. It functions in diverse and challenging settings, like heavily occluded environments or changing light conditions, and lends itself to various applications. All images shown here are frames from recorded live sessions.

We propose a stretch-sensing soft glove to interactively capture hand poses with high accuracy and without requiring an external optical setup. We demonstrate how our device can be fabricated and calibrated at low cost, using simple tools available in most fabrication labs. To reconstruct the pose from the capacitive sensors embedded in the glove, we propose a deep network architecture that exploits the spatial layout of the sensor itself. The network is trained only once, using an inexpensive off-the-shelf hand pose reconstruction system to gather the training data. The per-user calibration is then performed on-the-fly using only the glove. The glove's capabilities are demonstrated in a series of ablative experiments, exploring different models and calibration methods. Comparing against commercial data gloves, we achieve a 35% improvement in reconstruction accuracy.

CCS Concepts: • **Human-centered computing → Interaction devices**; • **Computing methodologies → Motion capture**; *Machine learning*; Computer graphics.

Additional Key Words and Phrases: hand tracking, data glove, sensor array, stretch-sensing

ACM Reference Format:
Oliver Glauser, Shihao Wu, Daniele Panozzo, Otmar Hilliges, and Olga Sorkine-Hornung. 2019. Interactive Hand Pose Estimation using a Stretch-Sensing Soft Glove. *ACM Trans. Graph.* 38, 4, Article 41 (July 2019), 15 pages. https://doi.org/10.1145/3306346.3322957

1 INTRODUCTION

Hands are our primary means to manipulate physical objects and communicate with each other. Many applications such as gaming, robotics, biomechanical analysis, rehabilitation and emerging human-computer interaction paradigms such as augmented and virtual reality (AR/VR) critically depend on accurate means to recover the full hand pose even under dexterous articulation. These challenging applications require that a hand tracking solution fulfills the following requirements: 1) it must be *real-time*, 2) it should work in a variety of environments and settings, and 3) it should be minimally invasive in terms of user instrumentation.

In many applications, hand pose is recovered via commercial motion capture systems (MoCap) such as Vicon [2019], but these require expensive infrastructure and markers placed on the user. Marker-less approaches to the task of hand pose estimation include multiple cameras [Ballan et al. 2012; Tompson et al. 2014; Oikonomidis et al. 2011b], or more recently, a single depth camera [Oberweger and Lepetit 2017; Oberweger et al. 2015; Tang et al. 2014; Wan et al. 2016] or even monocular camera [Spurr et al. 2018; Iqbal et al. 2018; Cai et al. 2018; Mueller et al. 2018; Zimmermann and Brox 2017]. Despite this significant progress, vision-based methods require externally mounted cameras with the whole hand visible in the image. This limitation presents a practical barrier for many applications, in particular those where heavy occlusions can be expected, such as while interacting with an object, wearing gloves or other items of clothing or while working in cluttered environments. Thus camera-based techniques are limited to applications with a controlled environment and impose physical constraints on immersive user experiences.

Mounting sensors directly onto the user's hand removes the need for direct line-of-sight and can improve robustness and reliability. Not surprisingly, a variety of glove-like devices have been proposed in research (e.g. [Chossat et al. 2015]) and are available commercially (e.g., [Cyb 2019; Man 2019]). Such approaches typically leverage inertial measurement units (IMUs), bend sensors, strain sensors or combinations thereof to capture local bone transformations. While potentially accurate, placing a sufficient amount of sensing elements on a glove in order to capture all the degrees-of-freedom (DoFs) of the hand is challenging due to space constraints. Hence, most existing solutions use fewer sensors than there are DoFs in the human hand. This inherently restricts the reconstruction fidelity.

Authors' addresses: Oliver Glauser; Shihao Wu, ETH Zurich, Switzerland, {oliver.glauser, shihao.wu}@inf.ethz.ch; Daniele Panozzo, New York University, USA, panozzo@nyu.edu; Otmar Hilliges; Olga Sorkine-Hornung, ETH Zurich, Switzerland, {otmar.hilliges, olga.sorkine}@inf.ethz.ch.

0730-0301/2019/7-ART41
https://doi.org/10.1145/3306346.3322957

Learning to Fly: Computational Controller Design for Hybrid UAVs with Reinforcement Learning

JIE XU, Massachusetts Institute of Technology
TAO DU, Massachusetts Institute of Technology
MICHAEL FOSHEY, Massachusetts Institute of Technology
BEICHEN LI, Massachusetts Institute of Technology
BO ZHU, Dartmouth College
ADRIANA SCHULZ, University of Washington
WOJCIECH MATUSIK, Massachusetts Institute of Technology

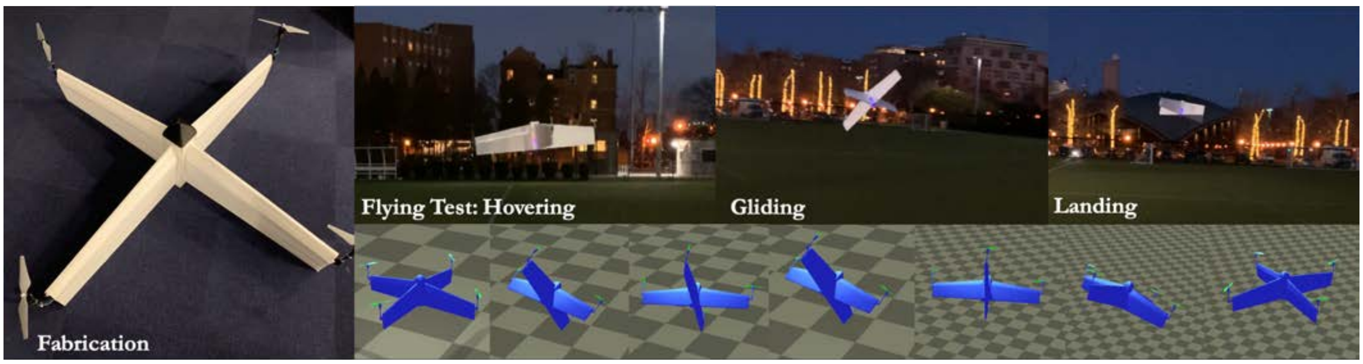

Fig. 1. We propose a computational approach to automatically design a mode-free, model-agnostic neural network controller for customized hybrid UAVs (left). Our novel reinforcement learning framework enables the training of a controller in a stochastically enhanced physics simulator (right bottom) and the crossing of the gap between simulation and reality (right top).

Hybrid unmanned aerial vehicles (UAV) combine advantages of multicopters and fixed-wing planes: vertical take-off, landing, and low energy use. However, hybrid UAVs are rarely used because controller design is challenging due to its complex, mixed dynamics. In this paper, we propose a method to automate this design process by training a mode-free, model-agnostic neural network controller for hybrid UAVs. We present a neural network controller design with a novel error convolution input trained by reinforcement learning. Our controller exhibits two key features: First, it does not distinguish among flying modes, and the same controller structure can be used for copters with various dynamics. Second, our controller works for real models without any additional parameter tuning process, closing the gap between virtual simulation and real fabrication. We demonstrate the efficacy of the proposed controller both in simulation and in our custom-built hybrid UAVs (Figure 1, 8). The experiments show that the controller is robust to exploit the complex dynamics when both rotors and wings are active in flight tests.

CCS Concepts: • **Computing methodologies → Control methods**; *Reinforcement learning*; *Physical simulation.*

Additional Key Words and Phrases: hybrid UAVs, neural network controllers

ACM Reference Format:
Jie Xu, Tao Du, Michael Foshey, Beichen Li, Bo Zhu, Adriana Schulz, and Wojciech Matusik. 2019. Learning to Fly: Computational Controller Design for Hybrid UAVs with Reinforcement Learning. *ACM Trans. Graph.* 38, 4, Article 42 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3322940

Authors' addresses: Jie Xu, Massachusetts Institute of Technology, jiex@csail.mit.edu; Tao Du, Massachusetts Institute of Technology; Michael Foshey, Massachusetts Institute of Technology; Beichen Li, Massachusetts Institute of Technology; Bo Zhu, Dartmouth College; Adriana Schulz, University of Washington; Wojciech Matusik, Massachusetts Institute of Technology.

0730-0301/2019/7-ART42 $15.00
https://doi.org/10.1145/3306346.3322940

1 INTODUCTION

Multicopters are becoming increasingly popular due to their flight flexibility and stable hovering capability. Fixed-wing airplanes, on the other hand, are more energy efficient during level flight, making them better vehicles for long distance flights. In order to leverage the advantages of both designs, new hybrid UAVs which equip multicopters with a pair of fixed wings have piqued the interest of the aircraft community [Bapst et al. 2015; Hugh Stone et al. 2008; Oosedo et al. 2013; Ritz and D'Andrea 2017]. These vehicles can quickly switch between hover and flight modes, allowing for highly stable, energy efficient flights. There are many ways to combine

Designing Chain Reaction Contraptions from Causal Graphs

ROBIN ROUSSEL, University College London
MARIE-PAULE CANI, LIX, École Polytechnique
JEAN-CLAUDE LÉON, Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LJK
NILOY J. MITRA, University College London and Adobe Research

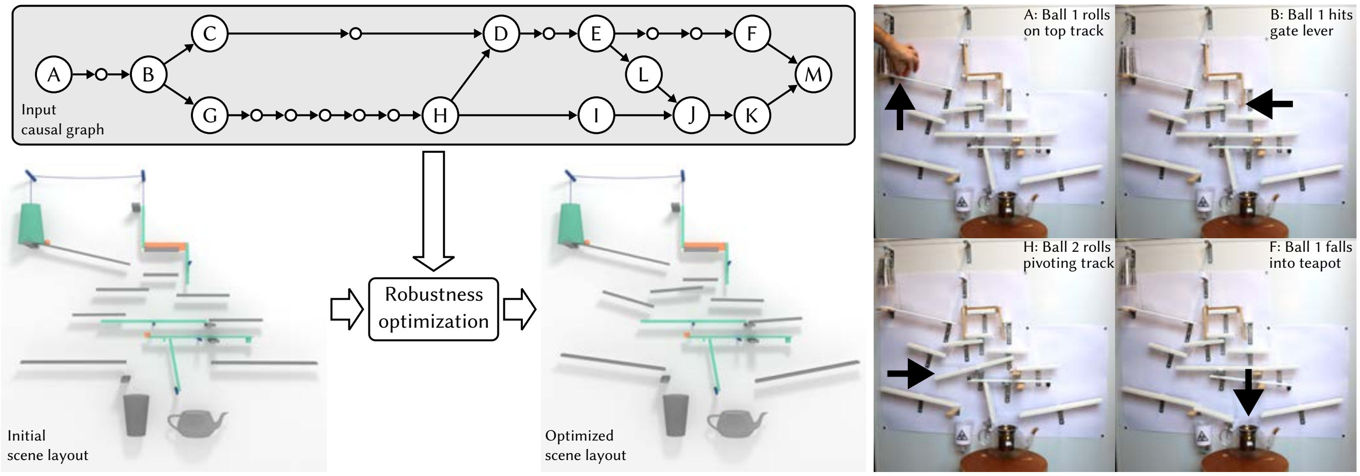

Fig. 1. Our system takes as input an initial scene layout associated with a causal graph of expected events. It then combines simulation, search and learning to build a success probability measure with respect to layout perturbations, and optimizes the layout for robustness against manual placement errors during assembly. The optimized layout is then exported as a guide sheet and used to successfully assemble complex chain reactions in the physical world.

Chain reaction contraptions, commonly referred to as Rube Goldberg machines, achieve simple tasks in an intentionally complex fashion via a cascading sequence of events. They are fun, engaging and satisfying to watch. Physically realizing them, however, involves hours or even days of manual trial-and-error effort. The main difficulties lie in predicting failure factors over long chains of events and robustly enforcing an expected causality between parallel chains, especially under perturbations of the layout. We present a computational framework to help design the layout of such contraptions by optimizing their robustness to possible assembly errors. Inspired by the active learning paradigm in machine learning, we propose a generic sampling-based method to progressively approximate the *success probability distribution* of a given scenario over the design space of possible scene layouts. The success or failure of any given simulation is determined from a user-specified causal graph enforcing a time ordering between expected events. Our method scales to complex causal graphs and high dimensional design spaces by dividing the graph and scene into simpler sub-scenarios. The aggregated success probability distribution is subsequently used to optimize the entire layout. We demonstrate the use of our framework through a range of real world examples of increasing complexity, and report significant improvements over alternative approaches. Code and fabrication diagrams are available on the project page.

CCS Concepts: • **Computing methodologies → Shape analysis**; *Model verification and validation*; Physical simulation.

Additional Key Words and Phrases: computational design, robust design, causal graphs, chain reactions, success probability distribution

ACM Reference Format:
Robin Roussel, Marie-Paule Cani, Jean-Claude Léon, and Niloy J. Mitra. 2019. Designing Chain Reaction Contraptions from Causal Graphs. *ACM Trans. Graph.* 38, 4, Article 43 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3322977

Authors' addresses: Robin Roussel, University College London, robin.roussel.15@ucl.ac.uk; Marie-Paule Cani, LIX, École Polytechnique, marie-paule.cani@polytechnique.edu; Jean-Claude Léon, Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LJK, jean-claude.leon@ense3.grenoble-inp.fr; Niloy J. Mitra, University College London and Adobe Research, n.mitra@ucl.ac.uk.

0730-0301/2019/7-ART43 $15.00
https://doi.org/10.1145/3306346.3322977

1 INTRODUCTION

Chain reaction contraptions, also known as Rube Goldberg machines, achieve simple functions from intentionally complex sequences of events (see Figure 2). These machines sit at the intersection of entertainment, art and engineering; as such, they are featured in TV commercials [Honda 2003], exhibited as art pieces [Fischli and Weiss 1987], and used for educational purposes in classrooms and science fairs [Kim and Park 2012]. A particularly compelling aspect of these setups is the careful management of risk: a chain of events

A null-scattering path integral formulation of light transport

BAILEY MILLER*, Dartmouth College, USA
ILIYAN GEORGIEV*, Autodesk, United Kingdom
WOJCIECH JAROSZ, Dartmouth College, USA

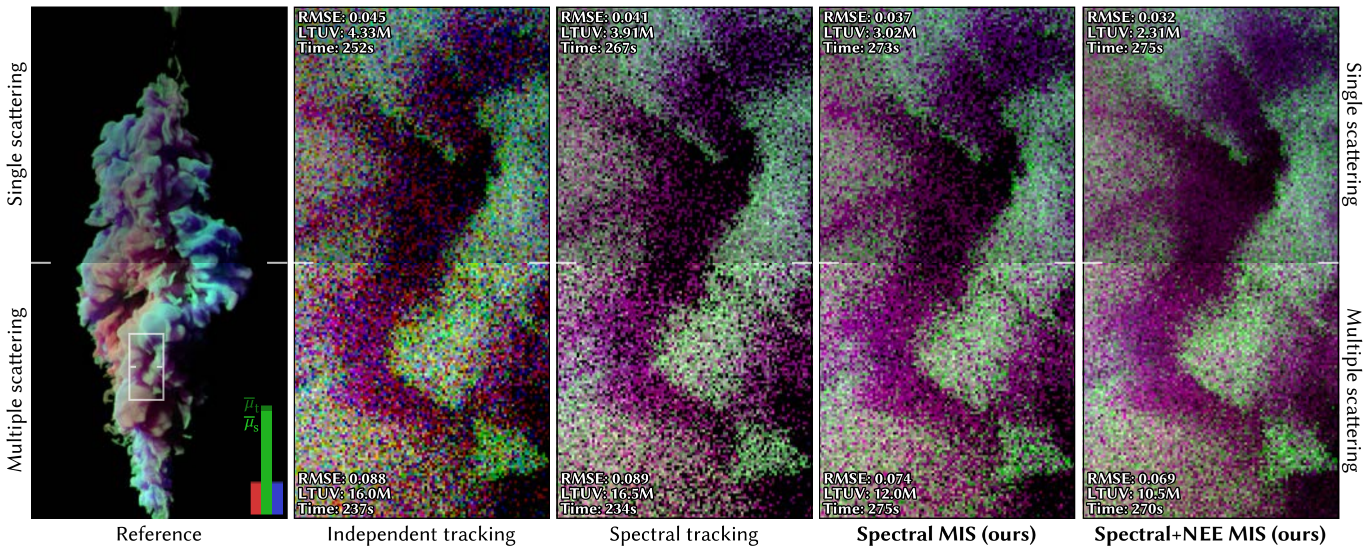

Fig. 1. Unbiased rendering of spectrally and spatially varying participating media could previously be accomplished using delta tracking separately for each color component (leftmost inset), but this leads to strong color noise. Spectral tracking [Kutz et al. 2017] can reduce this noise by rendering all color components together (middle left inset), but at the cost of sampling distances based on the densest color component of the medium. Our theoretical framework allows us to leverage different sampling techniques across color components (middle right), or exploit next-event estimation (NEE) (far right), and combine these into a more robust, lower-variance estimator via multiple importance sampling (MIS). See Tables 1 and 2 for descriptions of the methods and the medium.

Unbiased rendering of general, heterogeneous participating media currently requires using null-collision approaches for estimating transmittance and generating free-flight distances. A long-standing limitation of these approaches, however, is that the corresponding path pdfs cannot be computed due to the black-box nature of the null-collision rejection sampling process. These techniques therefore cannot be combined with other sampling techniques via multiple importance sampling (MIS), which significantly limits their robustness and generality. Recently, Galtier et al. [2013] showed how to derive these algorithms directly from the radiative transfer equation (RTE). We build off this generalized RTE to derive a *path integral* formulation of null scattering, which reveals the sampling pdfs and allows us to devise new, express existing, and combine complementary unbiased techniques via MIS. We demonstrate the practicality of our theory by combining, for the first time, several path sampling techniques in spatially and spectrally varying media, generalizing and outperforming the prior state of the art.

CCS Concepts: • **Computing methodologies → Ray tracing**.

Additional Key Words and Phrases: global illumination, light transport, participating media, null scattering, Monte Carlo integration

ACM Reference Format:
Bailey Miller, Iliyan Georgiev, and Wojciech Jarosz. 2019. A null-scattering path integral formulation of light transport. *ACM Trans. Graph.* 38, 4, Article 44 (July 2019), 13 pages. https://doi.org/10.1145/3306346.3323025

*Authors with equal contribution.

Authors' addresses: Bailey Miller, bailey.m.miller@dartmouth.edu, Dartmouth College, 9 Maynard Street, Hanover, NH, 03755, USA; Iliyan Georgiev, iliyan.georgiev@autodesk.com, Autodesk, 17 Broadwick Street, London, W1F 0DE, United Kingdom; Wojciech Jarosz, wojciech.k.jarosz@dartmouth.edu, Dartmouth College, 9 Maynard Street, Hanover, NH, 03755, USA.

0730-0301/2019/7-ART44 $15.00
https://doi.org/10.1145/3306346.3323025

1 INTRODUCTION

The world around us is filled with participating media which volumetrically attenuates and scatters light as it travels from light sources to our eyes. While important in many fields, simulating this transport efficiently and accurately is unfortunately a notoriously difficult problem, since it requires solving not only the rendering equation [Kajiya 1986; Immel et al. 1986], but also its volumetric generalization, the radiative transfer equation [Chandrasekhar 1960].

Monte Carlo path sampling methods such as (bidirectional) path tracing [Kajiya 1986; Lafortune and Willems 1993; Veach and Guibas 1995] and its volumetric variants [Lafortune and Willems 1996; Georgiev et al. 2013] have been investigated in academia for decades

Fractional Gaussian Fields for Modeling and Rendering of Spatially-Correlated Media

JIE GUO*, State Key Lab for Novel Software Technology, Nanjing University
YANJUN CHEN†, State Key Lab for Novel Software Technology, Nanjing University
BINGYANG HU†, State Key Lab for Novel Software Technology, Nanjing University
LING-QI YAN, University of California, Santa Barbara
YANWEN GUO*, State Key Lab for Novel Software Technology, Nanjing University
YUNTAO LIU, State Key Lab for Novel Software Technology, Nanjing University

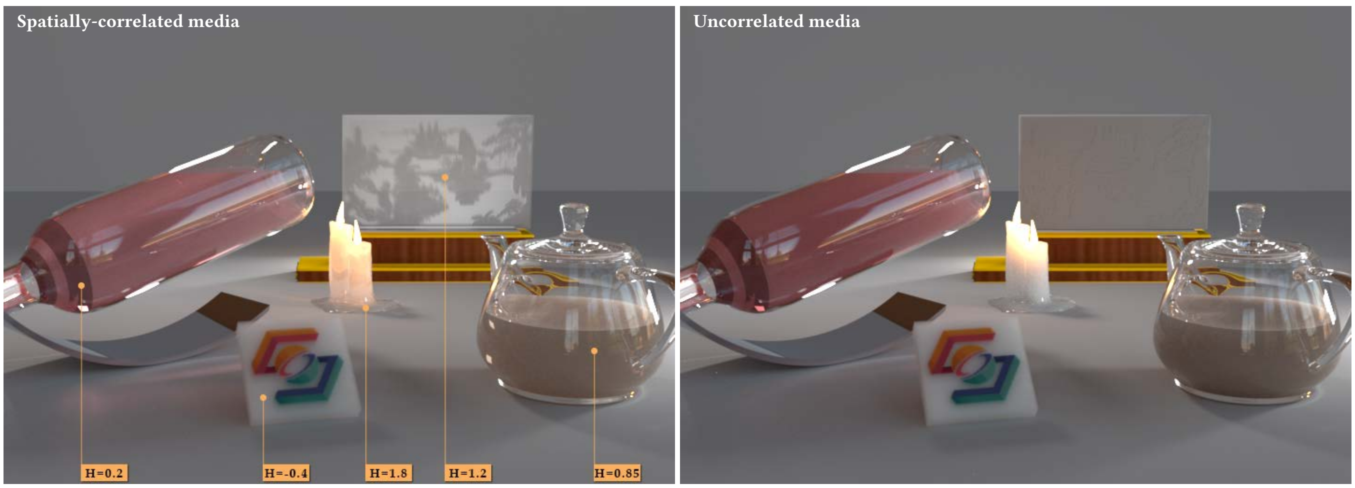

Fig. 1. We render a complex scene containing several spatially-correlated media, demonstrating that our method is able to reproduce a wide range of appearances stemming from short-range to long-range correlations and support macroscopic heterogeneity (left). A reference generated by the classical transport theory is provided for a comparison (right).

Transmission of radiation through spatially-correlated media has demonstrated deviations from the classical exponential law of the corresponding uncorrelated media. In this paper, we propose a general, physically-based method for modeling such correlated media with non-exponential decay of transmittance. We describe spatial correlations by introducing the Fractional Gaussian Field (FGF), a powerful mathematical tool that has proven useful in many areas but remains under-explored in graphics. With the FGF, we study the effects of correlations in a unified manner, by modeling both high-frequency, noise-like fluctuations and k-th order fractional Brownian motion (fBm) with a stochastic continuity property. As a result, we are able to reproduce a wide variety of appearances stemming from different types of spatial correlations. Compared to previous work, our method is the first that addresses both short-range and long-range correlations using physically-based fluctuation models. We show that our method can simulate different extents of randomness in spatially-correlated media, resulting in a smooth transition in a range of appearances from exponential falloff to complete transparency. We further demonstrate how our method can be integrated into an energy-conserving RTE framework with a well-designed importance sampling scheme and validate its ability compared to the classical transport theory and previous work.

*Corresponding authors
†Both authors contributed equally to the paper

Authors' addresses: Jie Guo, State Key Lab for Novel Software Technology, Nanjing University, guojie@nju.edu.cn; Yanjun Chen, State Key Lab for Novel Software Technology, Nanjing University, cujooyer@gmail.com; Bingyang Hu, State Key Lab for Novel Software Technology, Nanjing University, fhymyang@gmail.com; Ling-Qi Yan, University of California, Santa Barbara, lingqi@cs.ucsb.edu; Yanwen Guo, State Key Lab for Novel Software Technology, Nanjing University, ywguo@nju.edu.cn; Yuntao Liu, State Key Lab for Novel Software Technology, Nanjing University, windspectator@gmail.com.

0730-0301/2019/7-ART45 $15.00
https://doi.org/10.1145/3306346.3323031

CCS Concepts: • **Computing methodologies → Reflectance modeling**.

Additional Key Words and Phrases: Non-exponential transmittance, Volume rendering, Correlation, Random field, Importance sampling

ACM Reference Format:
Jie Guo, Yanjun Chen, Bingyang Hu, Ling-Qi Yan, Yanwen Guo, and Yuntao Liu. 2019. Fractional Gaussian Fields for Modeling and Rendering of Spatially-Correlated Media. *ACM Trans. Graph.* 38, 4, Article 45 (July 2019), 13 pages. https://doi.org/10.1145/3306346.3323031

Photon surfaces for robust, unbiased volumetric density estimation

XI DENG*, SHAOJIE JIAO*, BENEDIKT BITTERLI, and WOJCIECH JAROSZ, Dartmouth College, USA

Fig. 1. We compare the equal-time variance of different estimators in a scene containing participating media. We show the full light transport in the scene (left), single scattering (middle/right, top half) and multiple scattering volumetric transport (middle/right, bottom half). Our estimators (middle) provide significant variance reduction compared to prior density estimators (right) at equal render time.

We generalize photon planes to *photon surfaces*: a new family of unbiased volumetric density estimators which we combine using multiple importance sampling. To derive our new estimators, we start with the extended path integral which duplicates the vertex at the end of the camera and photon subpaths and couples them using a blurring kernel. To make our formulation unbiased, however, we use a delta kernel to couple these two end points. Unfortunately, sampling the resulting singular integral using Monte Carlo is impossible since the probability of generating a contributing light path by independently sampling the two subpaths is zero. Our key insight is that we can eliminate the delta kernel and make Monte Carlo estimation practical by integrating any three dimensions analytically, and integrating only the remaining dimensions using Monte Carlo. We demonstrate the practicality of this approach by instantiating a collection of estimators which analytically integrate the distance along the camera ray and two arbitrary sampling dimensions along the photon subpath (e.g., distance, direction, surface area). This generalizes photon planes to curved "photon surfaces", including new "photon cone", "photon cylinder", "photon sphere", and multiple new "photon plane" estimators. These estimators allow us to handle light paths not supported by photon planes, including single scattering, and surface-to-media transport. More importantly, since our estimators have complementary strengths due to analytically integrating different dimensions of the path integral, we can combine them using multiple importance sampling. This combination mitigates singularities present in individual estimators, substantially reducing variance while remaining fully unbiased. We demonstrate our improved estimators on a number of scenes containing homogeneous media with highly anisotropic phase functions, accelerating both multiple scattering and single scattering compared to prior techniques.

CCS Concepts: • **Computing methodologies → Ray tracing**.

*Joint first authors.

Authors' address: Xi Deng, xi.deng.gr@dartmouth.edu; Shaojie Jiao, shaojie.jiao.gr@dartmouth.edu; Benedikt Bitterli, benedikt.bitterli.gr@dartmouth.edu; Wojciech Jarosz, wojciech.k.jarosz@dartmouth.edu, Dartmouth College, 9 Maynard Street, Hanover, NH, 03755, USA.

0730-0301/2019/7-ART46 $15.00
https://doi.org/10.1145/3306346.3323041

Additional Key Words and Phrases: global illumination, light transport, participating media, photon density estimation

ACM Reference Format:
Xi Deng, Shaojie Jiao, Benedikt Bitterli, and Wojciech Jarosz. 2019. Photon surfaces for robust, unbiased volumetric density estimation. *ACM Trans. Graph.* 38, 4, Article 46 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3323041

1 INTRODUCTION

Accurate and efficient simulation of light transport in participating media is a challenging yet important problem in fields ranging from the movie industry, medical imaging, video games and even nuclear reactor design. Solving the governing rendering equation [Immel et al. 1986; Kajiya 1986] and equation of radiative transfer [Chandrasekhar 1960] has been the subject of substantial research over the past decades [Cerezo et al. 2005; Novák et al. 2018a,b], resulting in a wealth of rendering algorithms that tackle this problem.

Unbiased rendering methods such as path tracing [Kajiya 1986] are among the oldest forms of Monte Carlo (MC) light transport, and they remain popular [Pharr et al. 2016; Christensen and Jarosz 2016; Fascione et al. 2017] due to their simplicity and ability to render images where the only error is noise. Naive path tracing converges poorly in difficult lighting scenarios, so subsequent variants such as bidirectional path tracing (BPT) [Veach and Guibas 1994, 1997; Lafortune and Willems 1993] augment it with a diverse set of complementary path sampling strategies which can be combined using multiple importance sampling (MIS) [Veach and Guibas 1995].

Photon density estimation using points [Jensen 1996; Jensen and Christensen 1998] or beams [Jarosz et al. 2008, 2011a] gains efficiency partly from the path reuse permitted by the density estimation framework. This efficiency, however, comes at the cost of bias, which manifests as blurring in the image. In the same spirit as BPT, recent work has sought to combine some of these estimators with unbiased methods [Georgiev et al. 2012; Hachisuka et al. 2012] and each other [Křivánek et al. 2014] using MIS to leverage their complementary strengths. Unfortunately, this is made difficult by the fact that photon mapping and unbiased path sampling approaches operate in path spaces of different dimension, and the resulting combination remains biased. Most recently, Bitterli and Jarosz [2017]

Creating Impactful Characters: Correcting Human Impact Accelerations using High Rate IMUs in Dynamic Activities

CALVIN KUO, University of British Columbia, Canada
ZIHENG LIANG, University of British Columbia, Canada
YE FAN, University of British Columbia, Canada and Vital Mechanics Research, Canada
JEAN-SÉBASTIEN BLOUIN, University of British Columbia, Canada
DINESH K. PAI, University of British Columbia, Canada and Vital Mechanics Research, Canada

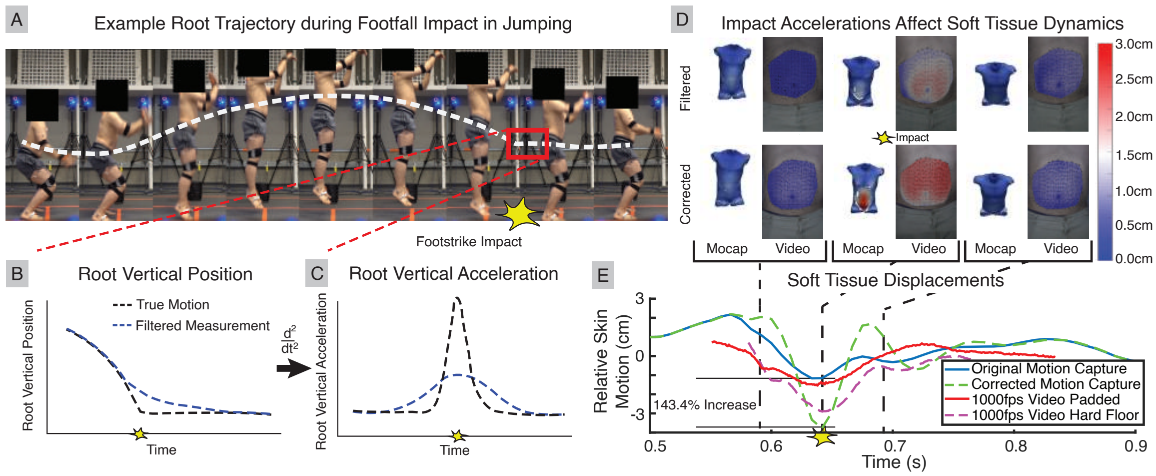

Fig. 1. **Correcting impact accelerations for impactful characters:** Motion capture of dynamic activities such as (A) jumping estimate generalized coordinates such as the (B) pelvis vertical position well, but can filter (C) accelerations at the impact. (D) Correcting the impact acceleration magnitude with a sparse set of higher rate IMUs can produce more dynamic soft tissue animations, as soft tissue motion is excited by body accelerations. (E) Corrections to jumping impacts produced a 143.4% increase in soft tissue motion at the stomach for a generic male with an increase in impact acceleration of 73.6%. Similar increases in soft tissue motion from filtered motion capture to corrected accelerations is observed using high speed video of a subject jumping on a soft pad to represent the effect of low-pass filters in motion capture, and on a hard surface to represent corrected ground-truth accelerations.

Human motion capture using video-based or sensor-based methods gives animators the capability to directly translate complex human motions to create lifelike character animations. Advances in motion capture algorithms have improved their accuracy for estimating human generalized motion coordinates (joint angles and body positions). However, the traditional motion capture pipeline is not well suited to measure short duration, high acceleration impacts, such as running and jumping footstrikes. While high acceleration impacts have minimal influence on generalized coordinates, they play a big role in exciting soft tissue dynamics.

Here we present a method for correcting motion capture trajectories using a sparse set of inertial measurement units (IMUs) collecting at high sampling rates to produce more accurate impact accelerations without sacrificing accuracy of the generalized coordinates representing gross motions. We demonstrate the efficacy of our method by correcting human motion captured experimentally using commercial motion capture systems with high rate IMUs sampling at 400Hz during basketball jump shots and running. With our method, we automatically corrected 185 jumping impacts and 1266 running impacts from 5 subjects. Post correction, we found an average increase of 84.6% and 91.1% in pelvis vertical acceleration and ankle dorsiflexion velocity respectively for basketball jump shots, and an average increase of 110% and 237% in pelvis vertical acceleration and ankle plantarflexion velocity respectively for running. In both activities, pelvis vertical position and ankle angle had small corrections on average below 2.0cm and 0.20rad respectively. Finally, when driving a human rig with soft tissue dynamics using corrected motions, we found a 143.4% and 11.2% increase in soft tissue oscillation amplitudes in basketball jump shots and running respectively. Our methodology can be generalized to correct impact accelerations for other body segments, and provide new tools to create realistic soft tissue animations during dynamic activities for more lifelike characters and better motion reconstruction for biomechanical analyses.

Authors' addresses: Calvin Kuo, University of British Columbia, Vancouver, BC, Canada, calvin.kuo@ubc.ca; Ziheng Liang, University of British Columbia, Vancouver, BC, Canada, zhliang@cs.ubc.ca; Ye Fan, University of British Columbia, Vancouver, BC, Canada , Vital Mechanics Research, Vancouver, BC, Canada, yefan@cs.ubc.ca; Jean-Sébastien Blouin, University of British Columbia, Vancouver, BC, Canada, jean-sebastien.blouin@ubc.ca@ubc.ca; Dinesh K. Pai, University of British Columbia, Vancouver, BC, Canada , Vital Mechanics Research, Vancouver, Canada, pai@cs.ubc.ca.

0730-0301/2019/7-ART47 $15.00
https://doi.org/10.1145/3306346.3322978

InteractionFusion: Real-time Reconstruction of Hand Poses and Deformable Objects in Hand-object Interactions

HAO ZHANG, BNRist and School of Software, Tsinghua University, China
ZI-HAO BO, BNRist and School of Software, Tsinghua University, China
JUN-HAI YONG, BNRist and School of Software, Tsinghua University, China
FENG XU, BNRist and School of Software, Tsinghua University, China

Hand-object interaction is challenging to reconstruct but important for many applications like HCI, robotics and so on. Previous works focus on either the hand or the object while we jointly track the hand poses, fuse the 3D object model and reconstruct its rigid and nonrigid motions, and perform all these tasks in real time. To achieve this, we first use a DNN to segment the hand and object in the two input depth streams and predict the current hand pose based on the previous poses by a pre-trained LSTM network. With this information, a unified optimization framework is proposed to jointly track the hand poses and object motions. The optimization integrates the segmented depth maps, the predicted motion, a spatial-temporal varying rigidity regularizer and a real-time contact constraint. A nonrigid fusion technique is further involved to reconstruct the object model. Experiments demonstrate that our method can solve the ambiguity caused by heavy occlusions between hand and object, and generate accurate results for various objects and interacting motions.

CCS Concepts: • **Computing methodologies** → **Shape modeling**.

Additional Key Words and Phrases: hand tracking, hand-object interaction, non-rigid motion, model reconstruction

ACM Reference Format:
Hao Zhang, Zi-Hao Bo, Jun-Hai Yong, and Feng Xu. 2019. InteractionFusion: Real-time Reconstruction of Hand Poses and Deformable Objects in Hand-object Interactions. *ACM Trans. Graph.* 38, 4, Article 48 (July 2019), 11 pages. https://doi.org/10.1145/3306346.3322998

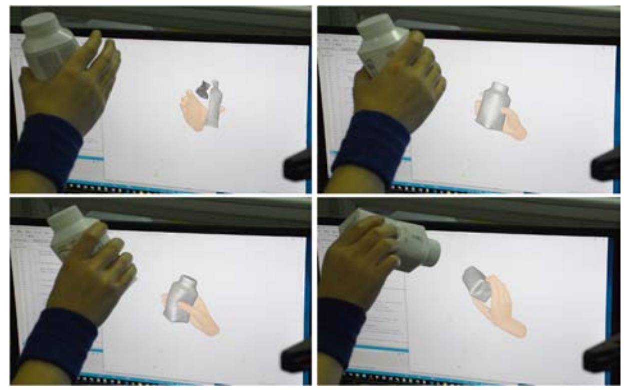

Fig. 1. Real-time reconstruction of hand-object interactions. The top two figures indicate the fusion of the object geometry from incomplete to complete. The bottom two figures show the nonrigid deformation and rigid motion of the object caused by interactions. More results are shown in the result section and the accompanying video.

1 INTRODUCTION

Reconstructing 3D motions of human hands is an important topic in computer vision and graphics due to its numerous applications in Human-Computer Interaction (HCI), robotics, rehabilitation, behavior analysis, virtual and augmented reality, and so on. There are many works focusing on the tracking of isolated hands, which is useful for a few applications like gesture recognition and control. However, human hands are majorly used for interacting with the environment or manipulating objects. Comparing with isolated hands, the reconstruction of interacting motions is more eagerly required.

Authors' addresses: Hao Zhang, BNRist and School of Software, Tsinghua University, China, zhanghao16@mails.tsinghua.edu.cn; Zi-Hao Bo, BNRist and School of Software, Tsinghua University, China; Jun-Hai Yong, BNRist and School of Software, Tsinghua University, China; Feng Xu, BNRist and School of Software, Tsinghua University, 30 Shuangqing Rd, Haidian Qu, Beijing Shi, China, feng-xu@tsinghua.edu.cn.

0730-0301/2019/7-ART48 $15.00
https://doi.org/10.1145/3306346.3322998

In literature, besides pure hand tracking [Han et al. 2018; Taylor et al. 2016; Tkach et al. 2017], there are many works focusing on hand-object interactions. Some works assume known object shape and majorly focus on the interacting motions of hands [Ballan et al. 2012; Hamer et al. 2009; Kyriazis and Argyros 2013, 2014; Oikonomidis et al. 2011; Sridhar et al. 2016]. These techniques need to pre-model the objects and do not allow shape changes of objects. Recently, some works can handle objects with articulated motions [Tzionas et al. 2016] or even nonrigid motions [Petit et al. 2018; Tsoli and Argyros 2018]. To solve the object motions, these techniques require offline processing and initial templates of the objects, which limits their applications in reality. Techniques for in hand reconstruction aim to reconstruct the shape of in hand objects and thus do not require templates, but they only handle static shapes, not motions of objects. In general, there is no existing work that simultaneously reconstructs the 3D motion of human hands and the shapes of the manipulated objects as well as their motions (rigid or nonrigid), not to mention achieving all of them in real time.

There are some key challenges for the aforementioned full reconstruction of hand-object interactions. Hand tracking itself is difficult due to the ambiguity caused by complex motions, lack of geometry/texture features and self-occlusion. When hand interacts with objects, the ambiguity further increases due to the much heavier occlusions. The same situation happens to the objects, which may always be occluded by fingers and the palm. Furthermore, for objects, we need to not only solve their rigid and nonrigid motions,

Real-time Pose and Shape Reconstruction of Two Interacting Hands With a Single Depth Camera

FRANZISKA MUELLER, MPI Informatics, Saarland Informatics Campus
MICAH DAVIS, Universidad Rey Juan Carlos
FLORIAN BERNARD and OLEKSANDR SOTNYCHENKO, MPI Informatics, Saarland Informatics Campus
MICKEAL VERSCHOOR, MIGUEL A. OTADUY, and DAN CASAS, Universidad Rey Juan Carlos
CHRISTIAN THEOBALT, MPI Informatics, Saarland Informatics Campus

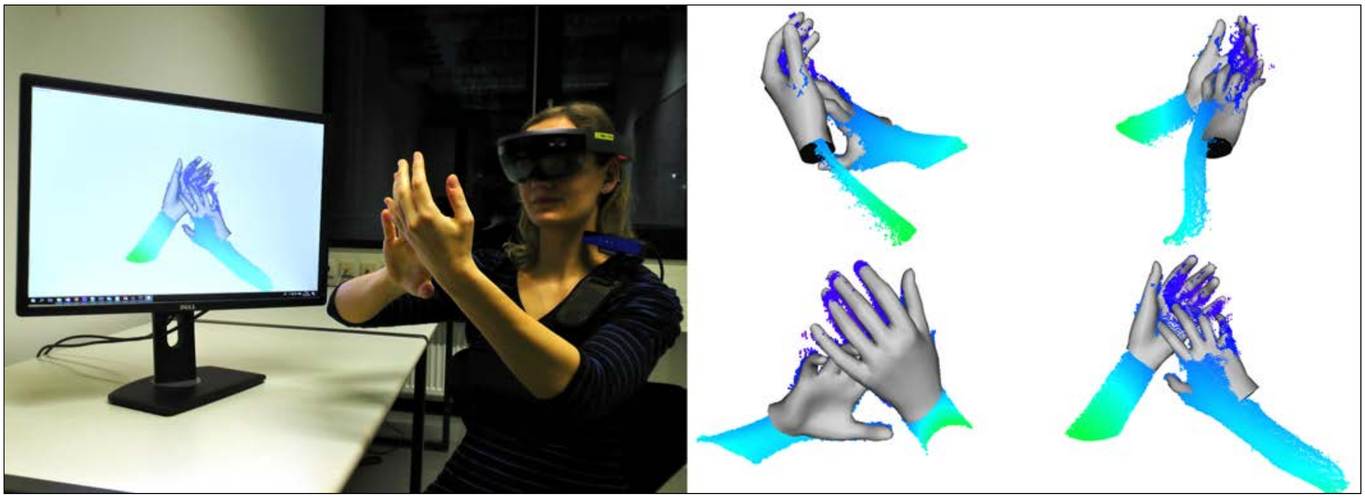

Fig. 1. We present a method that estimates the pose and shape of two interacting hands in real time from a single depth camera. On the left we show an AR setup with a shoulder-mounted depth camera. On the right we show the depth data and the estimated 3D hand pose and shape from four different views.

We present a novel method for real-time pose and shape reconstruction of two strongly interacting hands. Our approach is the first two-hand tracking solution that combines an extensive list of favorable properties, namely it is marker-less, uses a single consumer-level depth camera, runs in real time, handles inter- and intra-hand collisions, and automatically adjusts to the user's hand shape. In order to achieve this, we embed a recent parametric hand pose and shape model and a dense correspondence predictor based on a deep neural network into a suitable energy minimization framework. For training the correspondence prediction network, we synthesize a two-hand dataset based on physical simulations that includes both hand pose and shape annotations while at the same time avoiding inter-hand penetrations. To achieve real-time rates, we phrase the model fitting in terms of a nonlinear least-squares problem so that the energy can be optimized based on a highly efficient GPU-based Gauss-Newton optimizer. We show state-of-the-art results in scenes that exceed the complexity level demonstrated by previous work, including tight two-hand grasps, significant inter-hand occlusions, and gesture interaction.[1]

CCS Concepts: • **Computing methodologies → Tracking**; *Computer vision*; Neural networks.

Additional Key Words and Phrases: hand tracking, hand pose estimation, two hands, depth camera, computer vision

ACM Reference Format:
Franziska Mueller, Micah Davis, Florian Bernard, Oleksandr Sotnychenko, Mickeal Verschoor, Miguel A. Otaduy, Dan Casas, and Christian Theobalt. 2019. Real-time Pose and Shape Reconstruction of Two Interacting Hands With a Single Depth Camera. *ACM Trans. Graph.* 38, 4, Article 49 (July 2019), 13 pages. https://doi.org/10.1145/3306346.3322958

Authors' addresses: Franziska Mueller, MPI Informatics, Saarland Informatics Campus, frmueller@mpi-inf.mpg.de; Micah Davis, Universidad Rey Juan Carlos, davis.micah@urjc.es; Florian Bernard, fbernard@mpi-inf.mpg.de; Oleksandr Sotnychenko, MPI Informatics, Saarland Informatics Campus, osotnych@mpi-inf.mpg.de; Mickeal Verschoor, mickeal.verschoor@urjc.es; Miguel A. Otaduy, miguel.otaduy@urjc.es; Dan Casas, Universidad Rey Juan Carlos, dan.casas@urjc.es; Christian Theobalt, MPI Informatics, Saarland Informatics Campus, theobalt@mpi-inf.mpg.de.

0730-0301/2019/7-ART49 $15.00
https://doi.org/10.1145/3306346.3322958

1 INTRODUCTION

The marker-less estimation of hand poses is a challenging problem that has received a lot of attention in the vision and graphics communities. The relevance of the problem is owed to the fact that hand pose recognition plays an important role in many application

[1]project website: https://handtracker.mpi-inf.mpg.de/projects/TwoHands/

Accurate Markerless Jaw Tracking for Facial Performance Capture

GASPARD ZOSS, DisneyResearch|Studios, ETH Zurich
THABO BEELER, DisneyResearch|Studios
MARKUS GROSS, DisneyResearch|Studios, ETH Zurich
DEREK BRADLEY, DisneyResearch|Studios

Fig. 1. We present a method to accurately track the jaw during facial performance capture, without the need for attaching markers or tracking teeth.

We present the first method to accurately track the invisible jaw based solely on the visible skin surface, without the need for any markers or augmentation of the actor. As such, the method can readily be integrated with off-the-shelf facial performance capture systems. The core idea is to learn a non-linear mapping from the skin deformation to the underlying jaw motion on a dataset where ground-truth jaw poses have been acquired, and then to retarget the mapping to new subjects. Solving for the jaw pose plays a central role in visual effects pipelines, since accurate jaw motion is required when retargeting to fantasy characters and for physical simulation. Currently, this task is performed mostly manually to achieve the desired level of accuracy, and the presented method has the potential to fully automate this labour intense and error prone process.

Authors' addresses: Gaspard Zoss, DisneyResearch|Studios, ETH Zurich, gaspard.zoss@disneyresearch.com; Thabo Beeler, DisneyResearch|Studios, thabo.beeler@disneyresearch.com; Markus Gross, DisneyResearch|Studios, ETH Zurich, gross@disneyresearch.com; Derek Bradley, DisneyResearch|Studios, derek.bradley@disneyresearch.com.

0730-0301/2019/7-ART50 $15.00
https://doi.org/10.1145/3306346.3323044

CCS Concepts: • **Computing methodologies** → *Motion capture*; *Motion processing*.

Additional Key Words and Phrases: Jaw Tracking, Markerless, Data Driven Animation, Facial Performance Capture, Acquisition

ACM Reference Format:
Gaspard Zoss, Thabo Beeler, Markus Gross, and Derek Bradley. 2019. Accurate Markerless Jaw Tracking for Facial Performance Capture. *ACM Trans. Graph.* 38, 4, Article 50 (July 2019), 8 pages. https://doi.org/10.1145/3306346.3323044

1 INTRODUCTION

Generating realistic facial animation has always been a central ingredient in the creation of digital characters for computer games, visual effects for film and other virtual experiences. A very important component of the face is the underlying jaw, as the jaw's motion is often used to control the deformation of the face surface using methods like skinning. For this reason, most facial animation rigs contain an explicit jaw rig. While oftentimes the jaw is rigged with a few simple rotation and transformation controls, we have also seen recent advances in the design and control mechanisms of jaw rigs built from real-world capture data [Yang et al. 2018; Zoss et al. 2018].

Parametrization Quantization with Free Boundaries for Trimmed Quad Meshing

MAX LYON, RWTH Aachen University
MARCEL CAMPEN, Osnabrück University
DAVID BOMMES, University of Bern
LEIF KOBBELT, RWTH Aachen University

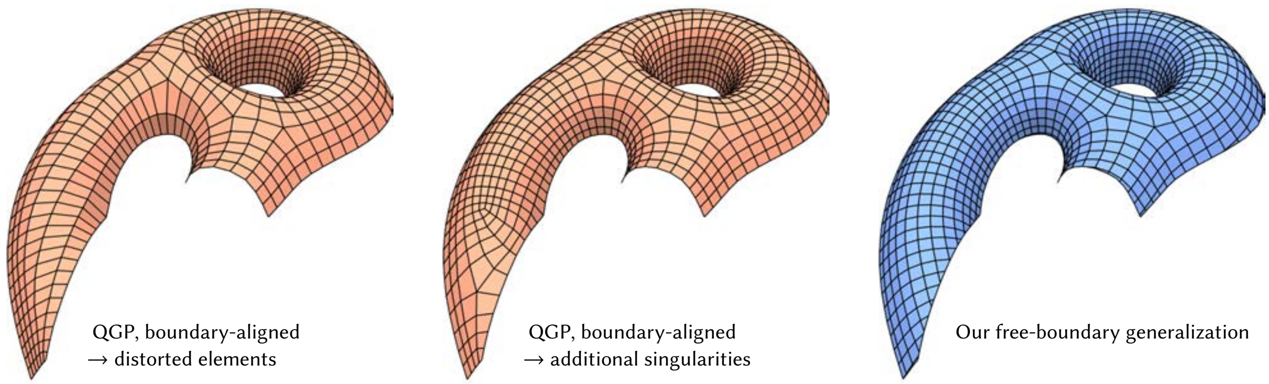

Fig. 1. When generating quad meshes for given surfaces, alignment of the mesh to the surface boundary may or may not be relevant, depending on the use case. Enforcing boundary alignment when this is not necessary, needlessly leads to lower mesh quality, e.g. distorted elements (left) or additional irregular vertices (center). *Trimmed quad meshing* with non-aligned boundaries (right), enabled by our free-boundary generalization of integer grid map quantization (QGP, [Campen et al. 2015]), avoids these issues and yields meshes of higher quality in such cases.

The generation of quad meshes based on surface parametrization techniques has proven to be a versatile approach. These techniques quantize an initial seamless parametrization so as to obtain an integer grid map implying a pure quad mesh. State-of-the-art methods following this approach have to assume that the surface to be meshed either has no boundary, or has a boundary which the resulting mesh is supposed to be aligned to. In a variety of applications this is not desirable and non-boundary-aligned meshes or grid-parametrizations are preferred. We thus present a technique to robustly generate integer grid maps which are either boundary-aligned, non-boundary-aligned, or partially boundary-aligned, just as required by different applications. We thereby generalize previous work to this broader setting. This enables the reliable generation of trimmed quad meshes with partial elements along the boundary, preferable in various scenarios, from tiled texturing over design and modeling to fabrication and architecture, due to fewer constraints and hence higher overall mesh quality and other benefits in terms of aesthetics and flexibility.

CCS Concepts: • **Computing methodologies → Computer graphics**; **Mesh models**; **Mesh geometry models**; *Shape modeling*.

Authors' addresses: Max Lyon, Visual Computing Institute, RWTH Aachen University; Marcel Campen, Osnabrück University; David Bommes, Computer Science, University of Bern; Leif Kobbelt, Visual Computing Institute, RWTH Aachen University.

0730-0301/2019/7-ART51 $15.00
https://doi.org/10.1145/3306346.3323019

Additional Key Words and Phrases: motorcycle graph, seamless parametrization, seamless texturing, trimmed NURBS

ACM Reference Format:
Max Lyon, Marcel Campen, David Bommes, and Leif Kobbelt. 2019. Parametrization Quantization with Free Boundaries for Trimmed Quad Meshing. *ACM Trans. Graph.* 38, 4, Article 51 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3323019

1 INTRODUCTION

Integer grid maps have been introduced as a versatile tool for the generation of high quality quad meshes based on surface parametrization [Bommes et al. 2013a, 2009; Kälberer et al. 2007; Tong et al. 2006]. Focus has often been on surfaces without boundary, or surfaces with boundary where quad edges coincide with the boundary everywhere (*boundary-aligned quad meshes*).

For a variety of applications, e.g. in simulation, texturing, structural and architectural design (cf. Section 2), mesh or grid map alignment to the surface boundary is neither necessary nor beneficial – rather, it brings in needless distortion: the mesh could be of higher quality (in terms of structural regularity, element shape, element sizing, feature or curvature alignment) if it was not forced to align with the boundary.

Unfortunately, straightforward application of the state-of-the-art parametrization quantization algorithm for the robust generation of integer grid maps [Campen et al. 2015] in a scenario *without* or with only *partial*, selective boundary alignment leads to a number of critical issues:

TriWild: Robust Triangulation with Curve Constraints

YIXIN HU, New York University, USA
TESEO SCHNEIDER, New York University, USA
XIFENG GAO, Florida State University, USA
QINGNAN ZHOU, Adobe Research, USA
ALEC JACOBSON, University of Toronto, Canada
DENIS ZORIN, New York University, USA
DANIELE PANOZZO, New York University, USA

We propose a robust 2D meshing algorithm, *TriWild*, to generate curved triangles reproducing smooth feature curves, leading to coarse meshes designed to match the simulation requirements necessary by applications and avoiding the geometrical errors introduced by linear meshes. The robustness and effectiveness of our technique are demonstrated by batch processing an SVG collection of 20k images, and by comparing our results against state of the art linear and curvilinear meshing algorithms. We demonstrate for our algorithm the practical utility of computing diffusion curves, fluid simulations, elastic deformations, and shape inflation on complex 2D geometries.

CCS Concepts: • **Mathematics of computing → Mesh generation**.

Additional Key Words and Phrases: Mesh Generation, Curved Triangulation, Robust Geometry Processing

ACM Reference Format:
Yixin Hu, Teseo Schneider, Xifeng Gao, Qingnan Zhou, Alec Jacobson, Denis Zorin, and Daniele Panozzo. 2019. TriWild: Robust Triangulation with Curve Constraints. *ACM Trans. Graph.* 38, 4, Article 52 (July 2019), 15 pages. https://doi.org/10.1145/3306346.3323011

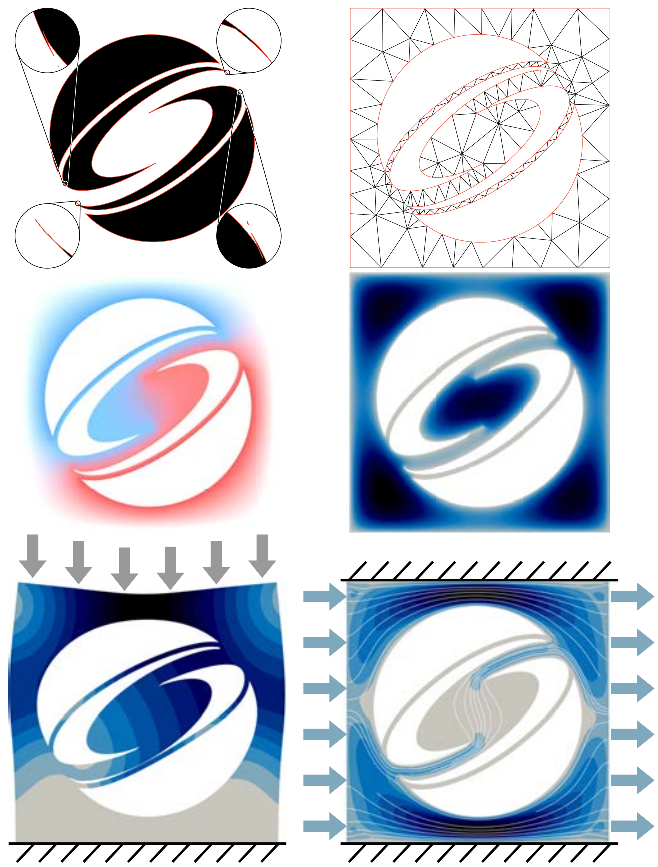

Fig. 1. The official ACM SIGGRAPH logo (www.siggraph.org/about/logos) is converted into a curved triangle mesh. We use the mesh to compute diffusion curves (Laplace), inflate surface (bilaplace), deform elastic bodies (Neo-Hooke), and simulate fluid flow (Stokes). Note that the imperfections in the input (shown in the closeups) are automatically healed by our method.

1 INTRODUCTION

Triangle meshing is at the core of a large fraction of two-dimensional computer graphics and computer aided engineering applications, most commonly, used to solve PDEs or optimization problems on 2D domains, in the context of physical simulation, geometric modeling, animation and nonphotorealistic rendering. Major efforts have been invested in robustly generating meshes with linear edges with good geometric quality. However, the restriction to linear meshes makes precise reproduction of simple curved shapes, such as a Bézier curve, impossible independently of the resolution used, resulting in artifacts and/or excessive refinement in applications ranging from physical simulation to nonphotorealistic rendering. Curved meshes, i.e. meshes with curved edges, are an effective solution to this problem: the idea is to use curved triangles instead of linear ones, providing significantly superior geometric approximation of a shape using a mesh of a particular size. In most cases, the lower triangle count leads to an overall more efficient simulation for a given desired accuracy [Braess 2007; Ciarlet and Raviart 1972; Scott 1973, 1975]. A simple 2D example is shown in Figure 2, which has a geometric error of 2% of the overall area when 236 linear triangles are used, and the error can be reduced to numerical zero with the same number of curved triangles with a cubic Lagrangian geometric map (Figure 2). While the use of curved meshes is well established in the FEM literature (with a few applications in graphics [Boyé et al. 2012; Mezger et al. 2009]), the automatic generation of these meshes is rarely considered, and the few existing methods we tested have a high failure rate on real-world examples (Section 5).

Authors' addresses: Yixin Hu, New York University, USA, yixin.hu@nyu.edu; Teseo Schneider, New York University, USA, teseo.schneider@nyu.edu; Xifeng Gao, Florida State University, USA, gao@cs.fsu.edu; Qingnan Zhou, Adobe Research, USA, qzhou@adobe.com; Alec Jacobson, University of Toronto, Canada, jacobson@cs.toronto.edu; Denis Zorin, New York University, USA, dzorin@cs.nyu.edu; Daniele Panozzo, New York University, USA, panozzo@nyu.edu.

0730-0301/2019/7-ART52 $15.00
https://doi.org/10.1145/3306346.3323011

Finding Hexahedrizations for Small Quadrangulations of the Sphere

KILIAN VERHETSEL, Université catholique de Louvain, Belgique
JEANNE PELLERIN, Université catholique de Louvain, Belgique
JEAN-FRANÇOIS REMACLE, Université catholique de Louvain, Belgique

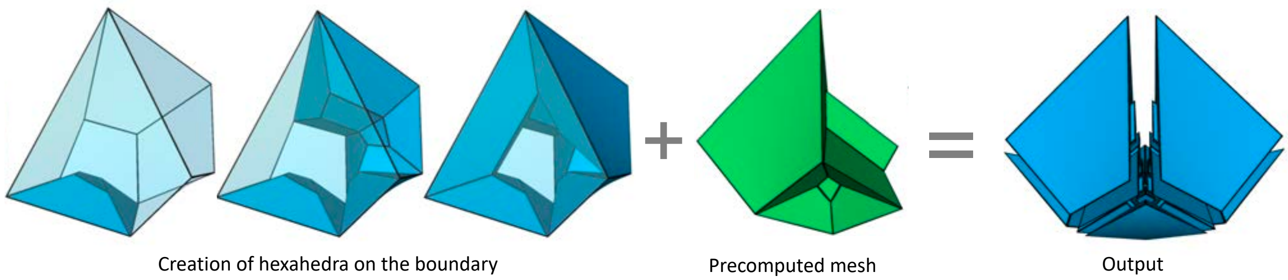

Fig. 1. Given a quadrangulation of the topological sphere, our algorithm creates hexahedra on the boundary until the unmeshed cavity matches the boundary of a pre-computed hex mesh that is merged to obtain the final combinatorial hexahedral mesh.

This paper tackles the challenging problem of constrained hexahedral meshing. An algorithm is introduced to build combinatorial hexahedral meshes whose boundary facets exactly match a given quadrangulation of the topological sphere. This algorithm is the first practical solution to the problem. It is able to compute small hexahedral meshes of quadrangulations for which the previously known best solutions could only be built by hand or contained thousands of hexahedra. These challenging quadrangulations include the boundaries of transition templates that are critical for the success of general hexahedral meshing algorithms.

The algorithm proposed in this paper is dedicated to building combinatorial hexahedral meshes of small quadrangulations and ignores the geometrical problem. The key idea of the method is to exploit the equivalence between quad flips in the boundary and the insertion of hexahedra glued to this boundary. The tree of all sequences of flipping operations is explored, searching for a path that transforms the input quadrangulation Q into a new quadrangulation for which a hexahedral mesh is known. When a small hexahedral mesh exists, a sequence transforming Q into the boundary of a cube is found; otherwise, a set of pre-computed hexahedral meshes is used.

A novel approach to deal with the large number of problem symmetries is proposed. Combined with an efficient backtracking search, it allows small shellable hexahedral meshes to be found for all even quadrangulations with up to 20 quadrangles. All 54, 943 such quadrangulations were meshed using no more than 72 hexahedra. This algorithm is also used to find a construction to fill arbitrary domains, thereby proving that any ball-shaped domain bounded by n quadrangles can be meshed with no more than 78 n hexahedra. This very significantly lowers the previous upper bound of 5396 n.

CCS Concepts: • **Computing methodologies** → **Mesh geometry models**; • **Mathematics of computing** → *Permutations and combinations*; *Combinatorial optimization.*

Additional Key Words and Phrases: hex-meshing, shelling, symmetry

ACM Reference Format:
Kilian Verhetsel, Jeanne Pellerin, and Jean-François Remacle. 2019. Finding Hexahedrizations for Small Quadrangulations of the Sphere. *ACM Trans. Graph.* 38, 4, Article 53 (July 2019), 13 pages. https://doi.org/10.1145/3306346.3323017

Authors' addresses: Kilian Verhetsel, Université catholique de Louvain, Avenue Georges Lemaître 4-6, Louvain-la-Neuve, 1348, Belgique, kilian.verhetsel@uclouvain.be; Jeanne Pellerin, Université catholique de Louvain, Avenue Georges Lemaître 4-6, Louvain-la-Neuve, 1348, Belgique, jeanne.pellerin@total.com; Jean-François Remacle, Université catholique de Louvain, Avenue Georges Lemaître 4-6, Louvain-la-Neuve, 1348, Belgique, jean-francois.remacle@uclouvain.be.

0730-0301/2019/7-ART53 $15.00
https://doi.org/10.1145/3306346.3323017

1 INTRODUCTION

Volumetric mesh generation is a required step for engineering analysis. Robust algorithms are able to automatically produce a tetrahedral mesh constrained to have a given triangulation as its boundary, e.g. [Si 2015]. However, subdivisions into hexahedra (cube-like cells) are often preferred over tetrahedrizations for their good numerical properties such as a better convergence with fewer elements [Shepherd and Johnson 2008] and faster assembly times [Remacle et al. 2016]. Yet, the hexahedral meshing problem, and more particularly the boundary constrained variant, remains open to this date.

Finding solutions to the boundary constrained hex-meshing problem is crucial for hex-meshing algorithms that use a few simple templates to reduce the complexity of the general meshing problem to a small set of inputs (*e.g.* [Mitchell 1999; Yamakawa and Shimada 2002]). More importantly, hex-dominant mesh generation techniques usually leave small cavities unmeshed [Yamakawa and Shimada 2003] and filling them is one of the missing pieces to the more general problem of all-hex mesh generation.

This paper introduces an algorithm that solves the combinatorial constrained hex-meshing problem for small quadrangulation of

Harmonic Triangulations

MARC ALEXA, TU Berlin, Germany

We introduce the notion of harmonic triangulations: a harmonic triangulation simultaneously minimizes the Dirichlet energy of all piecewise linear functions. By a famous result of Rippa, Delaunay triangulations are the harmonic triangulations of planar point sets. We prove by explicit counterexample that in 3D a harmonic triangulation does not exist in general. However, we show that bistellar flips are harmonic: if they decrease Dirichlet energy for one set of function values, they do so for all. This observation gives rise to the notion of locally harmonic triangulations. We demonstrate that locally harmonic triangulations can be efficiently computed, and efficiently reduce sliver tetrahedra. The notion of harmonic triangulation also gives rise to a scalar measure of the quality of a triangulation, which can be used to prioritize flips and optimize the position of vertices. Tetrahedral meshes generated by optimizing this function generally show better quality than Delaunay-based optimization techniques.

CCS Concepts: • **Mathematics of computing** → **Mesh generation**; Discrete optimization; • **Computing methodologies** → *Mesh geometry models*; • **Theory of computation** → Computational geometry.

Additional Key Words and Phrases: Dirichlet energy, simplicial meshes, tetrahedral mesh optimization

ACM Reference Format:
Marc Alexa. 2019. Harmonic Triangulations. *ACM Trans. Graph.* 38, 4, Article 54 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3322986

1 INTRODUCTION

The Delaunay triangulation of a point set is a fundamental tool in geometry processing and mesh generation. The reason for this is that it can be computed efficiently in any dimension and the geometric properties of the elements are well defined.

In the plane, Delaunay triangulations satisfy several optimality properties, for example, the minimal interior angle is maximized. A remarkable property is the connection to Dirichlet energy: for the given points in the plane attach a function value to each point. A triangulation of the points gives rise to a piecewise linear function (we present a representation of this function in Section 3). We may ask: which triangulation minimizes the Dirichlet energy of the PL function? The surprising result by Rippa [1990] is that the Delaunay triangulation minimizes Dirichlet energy *independent of the attached function values*.

This property, just as many other optimality properties of planar Delaunay triangulations, fails to generalize to three or more dimensions [Musin 1997]. The starting point of this work was the question: 'Are the good geometric properties of Delaunay triangulations really due to the classic definitions (see Section 2), or could small Dirichlet energy be the reason?' Our central idea is, consequently, to characterize and define harmonic triangulations as minimizers of Dirichlet energy – see Definition 3.

Our main results are that (1) minimization of Dirichlet energy independent of function values induces a partial order on the triangulations, which is given by the partial order of the positive semi-definite quadratic forms of the Dirichlet energy; and (2) bistellar flips are consistent with this order, in the sense that the effect of a flip on the Dirichlet energy is independent of the function values. This means that flipping can be used to generate local minima of Dirichlet energy in the flip-graph of triangulations. As minimizers of Dirichlet energy we suggest to call such triangulations *harmonic*. The details and additional geometric insights are discussed in Section 4

An important (and to our knowledge open) question is if a globally harmonic triangulation exists. In two dimensions this is the case, since the Delaunay triangulations minimizes Dirichlet energy independent of function values. In Section 6.1 we give an explicit counterexample for three dimensions. This means flipping towards locally harmonic triangulations is the best we can do, as the different locally harmonic triangulations are usually not comparable, meaning the comparison of their Dirichlet energies depends on the choice of function values.

There are many functions mapping from the quadratic form of Dirichlet energy to a scalar that are consistent with the partial order. Such functions can be used to optimize not just the combinatorics but also the vertex positions – similar in spirit to optimal Delaunay triangulations [Chen and Xu 2004]. We derive the gradients of a natural functions (see Section 5) and use it to optimize vertex positions in a descent scheme (Section 7.3).

Based on our observations we develop basic algorithms (see Section 7) for the generation of locally harmonic triangulations by flipping and potentially moving the vertices. These algorithms are easy to implement given frameworks for robust geometric computations and data structures that support bistellar flips on simplicial complexes. In a series of experiments we demonstrate the favorable properties of these algorithms: harmonic flipping to a locally harmonic triangulation is orders of magnitude faster than sliver exudation [Cheng et al. 2000]. The geometric properties of locally harmonic triangulations are similar or better than techniques based on Delaunay triangulations.

We conclude that harmonic triangulations are a useful new tool for triangulating and optimizing point sets in three dimensions. There are multiple applications for this tool and important follow-up explorations, some of which are discussed in Section 9.

2 BACKGROUND AND RELATED WORK

Given a point set $\{\mathbf{x}_i\}$ in $\mathbb{R}^d$, the Delaunay triangulation can be characterized in several seemingly different but intimately related ways. All of the following characterizations are independent of

Author's address: Marc Alexa, TU Berlin, Sekretariat MAR 6-6, Marchstr. 23, 10587 Berlin, Germany, marc.alexa@tu-berlin.de.

0730-0301/2019/7-ART54 $15.00
https://doi.org/10.1145/3306346.3322986

Navigating Intrinsic Triangulations

NICHOLAS SHARP, Carnegie Mellon University
YOUSUF SOLIMAN, Caltech
KEENAN CRANE, Carnegie Mellon University

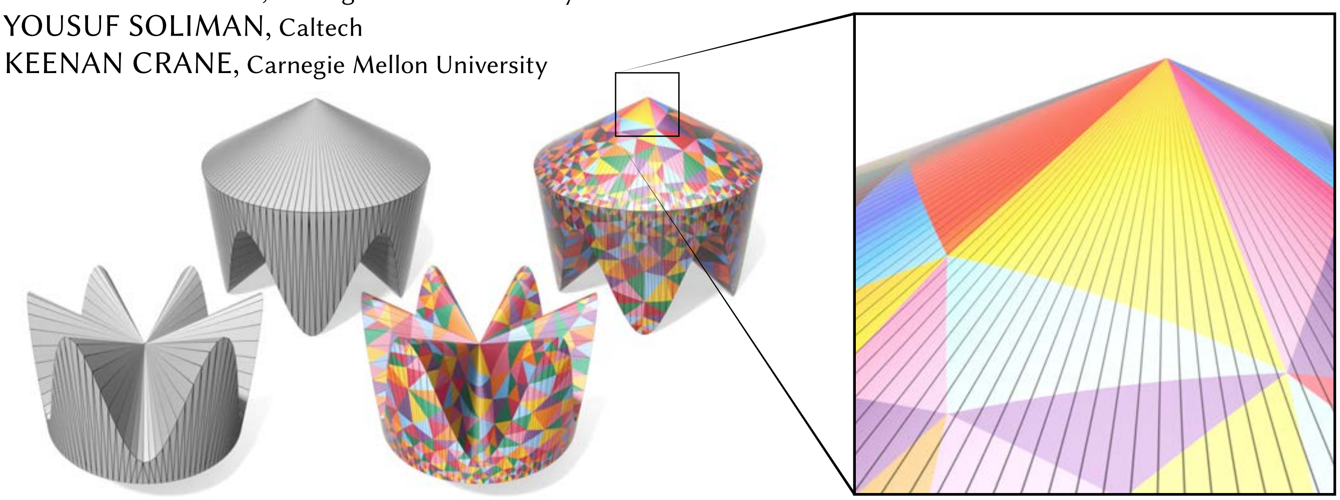

Fig. 1. Our data structure makes it possible to treat a crude input mesh *(left)* as a high-quality *intrinsic triangulation (right)* while exactly preserving the original geometry. Existing algorithms can be run directly on the new triangulation as though it is an ordinary triangle mesh. Here, a mesh with tiny input angles becomes a geometrically identical Delaunay triangulation with angles no smaller than 30°—a feat impossible for traditional, extrinsic remeshing.

We present a data structure that makes it easy to run a large class of algorithms from computational geometry and scientific computing on extremely poor-quality surface meshes. Rather than changing the geometry, as in traditional remeshing, we consider *intrinsic triangulations* which connect vertices by straight paths along the exact geometry of the input mesh. Our key insight is that such a triangulation can be encoded implicitly by storing the direction and distance to neighboring vertices. The resulting *signpost data structure* then allows geometric and topological queries to be made on-demand by tracing paths across the surface. Existing algorithms can be easily translated into the intrinsic setting, since this data structure supports the same basic operations as an ordinary triangle mesh (vertex insertions, edge splits, *etc.*). The output of intrinsic algorithms can then be stored on an ordinary mesh for subsequent use; unlike previous data structures, we use a constant amount of memory and do not need to explicitly construct an *overlay mesh* unless it is specifically requested. Working in the intrinsic setting incurs little computational overhead, yet we can run algorithms on extremely degenerate inputs, including all manifold meshes from the *Thingi10k* data set. To evaluate our data structure we implement several fundamental geometric algorithms including intrinsic versions of Delaunay refinement and optimal Delaunay triangulation, approximation of Steiner trees, adaptive mesh refinement for PDEs, and computation of Poisson equations, geodesic distance, and flip-free tangent vector fields.

Authors' addresses: Nicholas Sharp, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA, 15213; Yousuf Soliman, Caltech; Keenan Crane, Carnegie Mellon University.

0730-0301/2019/7-ART55 $15.00
https://doi.org/10.1145/3306346.3322979

CCS Concepts: • **Mathematics of computing → Mesh generation**.

Additional Key Words and Phrases: remeshing, discrete differential geometry

ACM Reference Format:
Nicholas Sharp, Yousuf Soliman, and Keenan Crane. 2019. Navigating Intrinsic Triangulations. *ACM Trans. Graph.* 38, 4, Article 55 (July 2019), 16 pages. https://doi.org/10.1145/3306346.3322979

1 INTRODUCTION

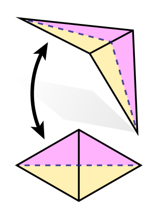

The geometry of a polyhedron has little to do with the way it is triangulated. For instance, flipping a diagonal of a triangulated cube does not change its shape; in general, any two neighboring faces of a triangulation can be laid out flat and connected along the opposite diagonal (see inset). Although the new edge looks bent when drawn on the surface, each triangle is still described by three ordinary edge lengths. Such *intrinsic triangulations* effectively provide "scaffolding" on top of a fixed geometric space: no information about shape is lost by changing the way vertices are connected. However, the choice of triangulation can have significant impact on the behavior of algorithms.

Intrinsic triangulations of geometric spaces have a long history in mathematics, but have seen limited use in practical algorithms: existing data structures support only simple *edge flips*, precluding their use for general geometry processing. Yet a full-blown intrinsic data structure is quite powerful, since it decouples the triangulation used to describe the domain from the one used to implement algorithms on that domain. Hence, rather than trying to make algorithms more robust one at a time, we can immediately run a large class of *existing* algorithms on low-quality inputs, with little to no modification.

Beyond Trilinear Interpolation: Higher Quality for Free

BALÁZS CSÉBFALVI, Budapest University of Technology and Economics, Hungary

In volume-rendering applications, it is a de facto standard to reconstruct the underlying continuous function by using trilinear interpolation, and to estimate the gradients for the shading computations by calculating central differences on the fly. In a GPU implementation, this requires seven trilinear texture samples: one for the function reconstruction, and six for the gradient estimation. In this paper, for the first time, we show that the six additional samples can be used not just for gradient estimation, but for significantly improving the quality of the function reconstruction as well. As the additional arithmetic operations can be performed in the shadow of the texture fetches, we can achieve this quality improvement for free without reducing the rendering performance at all. Therefore, our method can completely replace the standard trilinear interpolation in the practice of GPU-accelerated volume rendering.

CCS Concepts: • **Computing methodologies → Volumetric models**; **Image processing**; **Texturing**.

Additional Key Words and Phrases: GPU-accelerated volume rendering, Catmull-Rom spline interpolation.

ACM Reference Format:
Balázs Csébfalvi. 2019. Beyond Trilinear Interpolation: Higher Quality for Free. *ACM Trans. Graph.* 38, 4, Article 56 (July 2019), 8 pages. https://doi.org/10.1145/3306346.3323032

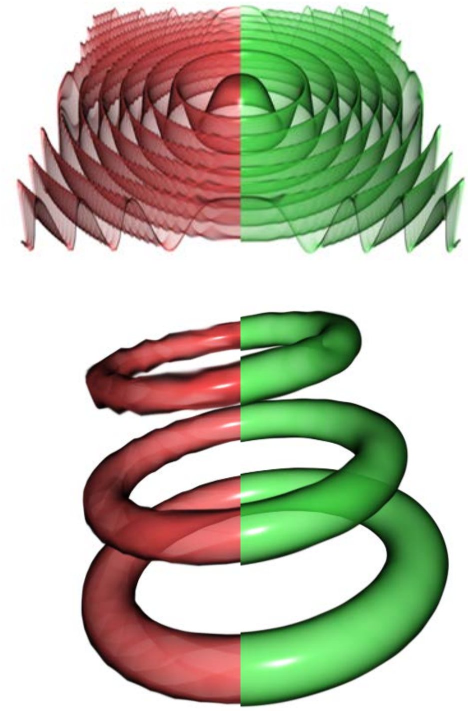

Fig. 1. Semitransparent rendering of implicit surfaces using trilinear interpolation (red) and our method (green) for reconstructing the underlying trivariate function from its discrete sampled representation. Our method significantly reduces the artifacts introduced by the trilinear interpolation but still guarantees the same rendering efficiency. The gradients are estimated from six trilinear samples in both cases, but our method reuses these samples for improving the quality of the function reconstruction.

1 INTRODUCTION

Volumetric data sets are usually obtained by sampling a trivariate function that represents a continuous 3D phenomenon. In order to visualize the original continuous phenomenon, the underlying function needs to be reconstructed between the sample positions. This requires a convolution of the discrete samples by an interpolation filter if the sample positions are assumed to be defined by the grid points of a regular sampling grid. In the practice of volume rendering, trilinear interpolation is one of the most popular resampling techniques, as it represents a widely accepted trade-off between image quality and rendering speed. This is especially true for GPU-accelerated volume rendering [Cabral et al. 1994; Engel et al. 2006; Krüger and Westermann 2003; Westermann and Ertl 1998], where a hardwired implementation of trilinear texture fetching is available. Although higher-order filters, such as the tricubic B-spline [Marschner and Lobb 1994; Mitchell and Netravali 1988] or the tricubic Catmull-Rom spline [Catmull and Rom 1974; Keys 1981], guarantee higher image quality [Marschner and Lobb 1994], they are significantly slower to evaluate even if their state-of-the-art GPU implementations [Csébfalvi 2018; Ruijters et al. 2008; Sigg and Hadwiger 2005] are used. In this paper, we propose a resampling technique that is as efficient as a trilinear interpolation combined with on-the-fly central differencing, but still provides much higher reconstruction quality (see Figure 1). Therefore, it can potentially become a new standard tool for volume resampling.

Author's address: Balázs Csébfalvi, Budapest University of Technology and Economics, Department of Control Engineering and Information Technology, Magyar tudósok körútja 2, Budapest, H-1117, Hungary, cseb@iit.bme.hu.

0730-0301/2019/7-ART56 $15.00
https://doi.org/10.1145/3306346.3323032

2 RELATED WORK

Theoretically, higher-order reconstruction filtering has been thoroughly studied in the literature taking both sampling-theoretical [Li et al. 2004; Marschner and Lobb 1994; Mitchell and Netravali 1988; Theußl et al. 2000] and approximation-theoretical aspects [Blu et al. 1999; Blu and Unser 1999a,b; Condat et al. 2005; Csébfalvi 2008; Möller et al. 1997, 1998] into account. However, only few papers propose fast GPU implementations for higher-order filters. In 2005, Sigg and Hadwiger published a pioneering work [Sigg and Hadwiger 2005] on an efficient evaluation of tricubic B-spline filtering on the

Procedural Phasor Noise

THIBAULT TRICARD* and SEMYON EFREMOV*, Université de Lorraine, CNRS, Inria, LORIA
CÉDRIC ZANNI, Université de Lorraine, CNRS, Inria, LORIA
FABRICE NEYRET, CNRS, Université Grenoble Alpes, Grenoble INP, LJK, INRIA
JONÀS MARTÍNEZ, Université de Lorraine, CNRS, Inria, LORIA
SYLVAIN LEFEBVRE, Université de Lorraine, CNRS, Inria, LORIA

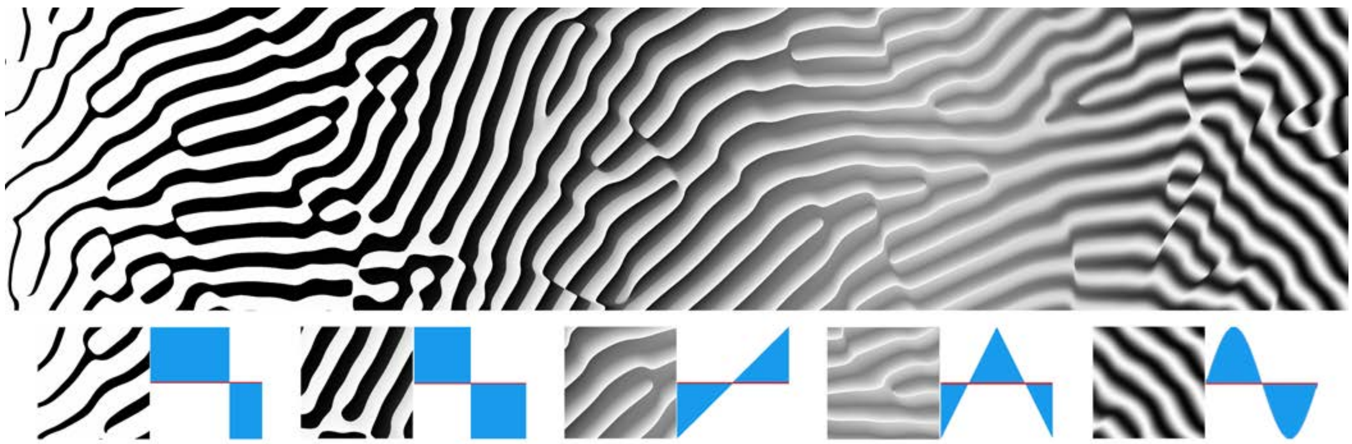

Fig. 1. High-contrast patterns produced by our approach. Note how the profile of the oscillations smoothly transition from a rectangular wave (20% black), to a square wave, to a triangular profile and finally a sine wave. At the same time, the orientation of the waves changes from left to right. The field visualized here is purely procedural. It is obtained by feeding our phasor noise into periodic profile functions (shown in blue), that are interpolated from left to right.

Procedural pattern synthesis is a fundamental tool of Computer Graphics, ubiquitous in games and special effects. By calling a single procedure in every pixel – or voxel – large quantities of details are generated at low cost, enhancing textures, producing complex structures within and along surfaces. Such procedures are typically implemented as pixel shaders.

We propose a novel procedural pattern synthesis technique that exhibits desirable properties for modeling highly contrasted patterns, that are especially well suited to produce surface and microstructure details. In particular, our synthesizer affords for a precise control over the profile, orientation and distribution of the produced stochastic patterns, while allowing to grade all these parameters spatially.

Our technique defines a stochastic smooth phase field – *a phasor noise* – that is then fed into a periodic function (e.g. a sine wave), producing an oscillating field with prescribed main frequencies and preserved contrast oscillations. In addition, the profile of each oscillation is directly controllable (e.g. sine wave, sawtooth, rectangular or any 1D profile). Our technique builds upon a reformulation of Gabor noise in terms of a *phasor field* that affords for a clear separation between local intensity and phase.

Applications range from texturing to modeling surface displacements, as well as multi-material microstructures in the context of additive manufacturing.

CCS Concepts: • **Computing methodologies → Texturing**.

Additional Key Words and Phrases: procedural, textures, noise, Gabor, pattern, texture synthesis

ACM Reference Format:
Thibault Tricard, Semyon Efremov, Cédric Zanni, Fabrice Neyret, Jonàs Martínez, and Sylvain Lefebvre. 2019. Procedural Phasor Noise. *ACM Trans. Graph.* 38, 4, Article 57 (July 2019), 13 pages. https://doi.org/10.1145/3306346.3322990

*Joint first authors

Authors' addresses: Thibault Tricard; Semyon Efremov, Université de Lorraine, CNRS, Inria, LORIA, Nancy, France, F-54000; Cédric Zanni, Université de Lorraine, CNRS, Inria, LORIA, Nancy, France, F-54000; Fabrice Neyret, CNRS, Université Grenoble Alpes, Grenoble INP, LJK, INRIA; Jonàs Martínez, Université de Lorraine, CNRS, Inria, LORIA, Nancy, France, F-54000; Sylvain Lefebvre, Université de Lorraine, CNRS, Inria, LORIA, Nancy, France, F-54000.

0730-0301/2019/7-ART57 $15.00
https://doi.org/10.1145/3306346.3322990

1 INTRODUCTION

The Computer Graphics community is in a perpetual quest to extend the range of patterns that can be generated at low memory and computational costs, from a single procedure. Most techniques rely on *procedural noises* [Lagae et al. 2010], which generate scalar fields with prescribed frequency content. These base noises are then combined through various functions to produce interesting patterns [Ebert et al. 2003].

A key limitation of standard procedural noises is the lack of direct control over the local characteristics of the patterns, such as contrast,

TileGAN: Synthesis of Large-Scale Non-Homogeneous Textures

ANNA FRÜHSTÜCK, KAUST
IBRAHEEM ALHASHIM, KAUST
PETER WONKA, KAUST

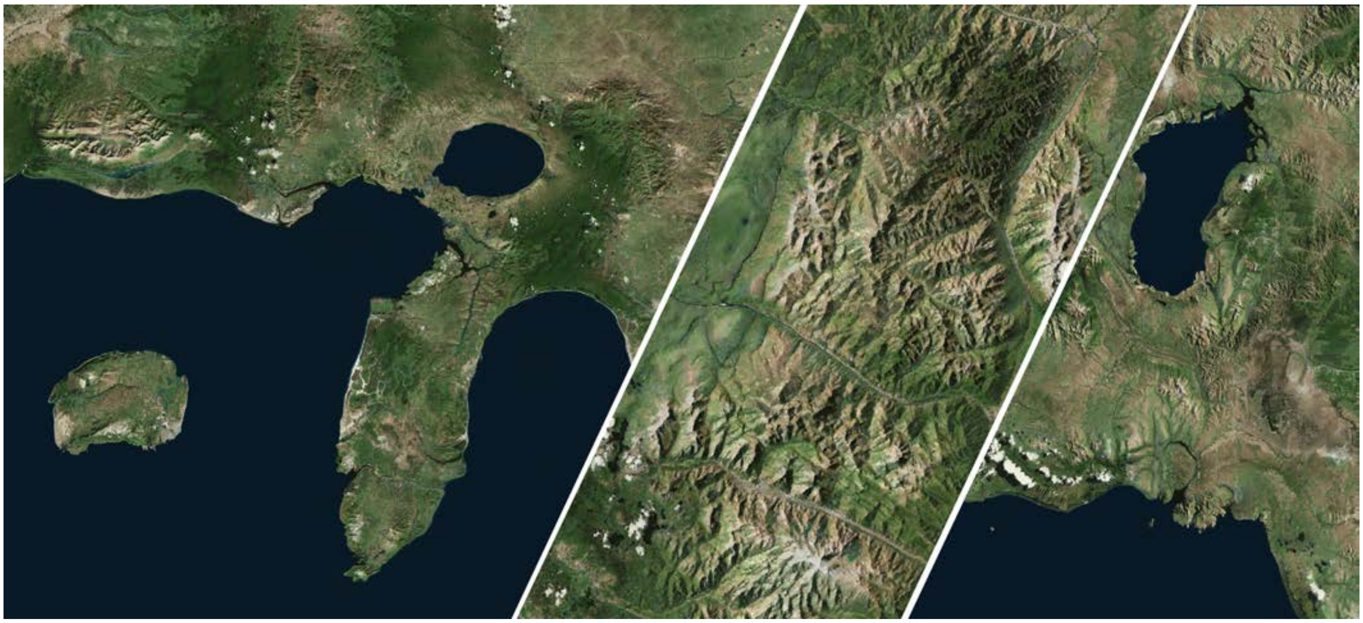

Fig. 1. TileGAN can synthesize large-scale textures with rich details. We show aerial images at different levels of detail generated using our framework, which allows for interactive texture editing. Our results contain a broad diversity of features at multiple scales and can be several hundreds of megapixels in size.

We tackle the problem of texture synthesis in the setting where many input images are given and a large-scale output is required. We build on recent generative adversarial networks and propose two extensions in this paper. First, we propose an algorithm to combine outputs of GANs trained on a smaller resolution to produce a large-scale plausible texture map with virtually no boundary artifacts. Second, we propose a user interface to enable artistic control. Our quantitative and qualitative results showcase the generation of synthesized high-resolution maps consisting of up to hundreds of megapixels as a case in point.

CCS Concepts: • **Computing methodologies → Computer graphics**; **Texturing**.

Additional Key Words and Phrases: Texture Synthesis, Image Generation, Deep Learning, Generative Adversarial Networks

Authors' addresses: Anna Frühstück, KAUST, anna.fruehstueck@kaust.edu.sa; Ibraheem Alhashim, KAUST, ibraheem.alhashim@kaust.edu.sa; Peter Wonka , KAUST, pwonka@gmail.com — Bldg 1, Al Khawarizmi, 4700 KAUST, Thuwal 23955-6900, Kingdom of Saudi Arabia.

0730-0301/2019/7-ART58
https://doi.org/10.1145/3306346.3322993

ACM Reference Format:
Anna Frühstück, Ibraheem Alhashim, and Peter Wonka. 2019. TileGAN: Synthesis of Large-Scale Non-Homogeneous Textures. *ACM Trans. Graph.* 38, 4, Article 58 (July 2019), 11 pages. https://doi.org/10.1145/3306346.3322993

1 INTRODUCTION

Example-based texture synthesis is the task of generating textures that look similar to a given input example. The visual features of the input texture should be faithfully reproduced while maintaining both small-scale as well as global characteristics of the exemplar.

In this paper, we are interested in synthesizing large-scale textures that consist of multiple megapixels (see Fig. 1). The first challenge in large-scale texture synthesis is to process a large amount of input data. This is crucial because without a considerable amount of reference data, any generated output will not have a lot of variability and lack features at multiple scales. Such a synthesized output could be large-scale, but will be very homogeneous and boring or repetitive. Recent work in parametric texture synthesis using generative adversarial networks (GANs) seems ideally suited to tackle this challenge and we build on a recent GAN architecture that can generate high-quality results when trained on natural textures [Karras et al. 2018a]. The second challenge in large-scale texture synthesis is how to generate large-scale output data. This is the core topic of this

Semantic Photo Manipulation with a Generative Image Prior

DAVID BAU, MIT CSAIL and MIT-IBM Watson AI Lab
HENDRIK STROBELT, IBM Research and MIT-IBM Watson AI Lab
WILLIAM PEEBLES, MIT CSAIL
JONAS WULFF, MIT CSAIL
BOLEI ZHOU, The Chinese University of Hong Kong
JUN-YAN ZHU, MIT CSAIL
ANTONIO TORRALBA, MIT CSAIL and MIT-IBM Watson AI Lab

Fig. 1. Our proposed method enables several new interactive photo manipulations in which a user edits a photo with high-level concepts rather than pixel colors. Our deep generative model can synthesize new content that follows both the user's intention and the natural image statistics. *T*op: Given simple user strokes, our method can automatically synthesize and manipulate different objects while adjusting the surrounding context to match. *B*ottom: Our users can edit the visual appearance of objects directly, such as changing the appearance of rooftops or trees. Photos from the LSUN dataset [Yu et al. 2015].

Despite the recent success of GANs in synthesizing images conditioned on inputs such as a user sketch, text, or semantic labels, manipulating the high-level attributes of an existing natural photograph with GANs is challenging for two reasons. First, it is hard for GANs to precisely reproduce an input image. Second, after manipulation, the newly synthesized pixels often do not fit the original image. In this paper, we address these issues by adapting the image prior learned by GANs to image statistics of an individual image. Our method can accurately reconstruct the input image and synthesize new content, consistent with the appearance of the input image. We demonstrate our interactive system on several semantic image editing tasks, including synthesizing new objects consistent with background, removing unwanted objects, and changing the appearance of an object. Quantitative and qualitative comparisons against several existing methods demonstrate the effectiveness of our method.

CCS Concepts: • **Computing methodologies → Image representations**; **Neural networks**; **Image manipulation**.

Additional Key Words and Phrases: image editing, generative adversarial networks, deep learning, vision for graphics

ACM Reference Format:
David Bau, Hendrik Strobelt, William Peebles, Jonas Wulff, Bolei Zhou, Jun-Yan Zhu, and Antonio Torralba. 2019. Semantic Photo Manipulation with a Generative Image Prior. *ACM Trans. Graph.* 38, 4, Article 59 (July 2019), 11 pages. https://doi.org/10.1145/3306346.3323023

1 INTRODUCTION

The whirlwind of progress in deep learning has produced a steady stream of promising generative models [Goodfellow et al. 2014; Karras et al. 2018] that render natural scenes increasingly indistinguishable from reality and provide an intuitive way to generate realistic imagery given high-level user inputs [Bau et al. 2019; Wang et al. 2018].

Authors' addresses: David Bau, MIT CSAIL and MIT-IBM Watson AI Lab, davidbau@csail.mit.edu; Hendrik Strobelt, IBM Research and MIT-IBM Watson AI Lab, hendrik.strobelt@ibm.com; William Peebles, MIT CSAIL, wisp@csail.mit.edu; Jonas Wulff, MIT CSAIL, jwulff@csail.mit.edu; Bolei Zhou, The Chinese University of Hong Kong, bzhou@ie.cuhk.edu.hk; Jun-Yan Zhu, MIT CSAIL, junyanz@csail.mit.edu; Antonio Torralba, MIT CSAIL and MIT-IBM Watson AI Lab, torralba@csail.mit.edu.

0730-0301/2019/7-ART59
https://doi.org/10.1145/3306346.3323023

The Face of Art: Landmark Detection and Geometric Style in Portraits

JORDAN YANIV, Tel Aviv University & The Interdisciplinary Center, Herzliya
YAEL NEWMAN, Tel Aviv University & The Interdisciplinary Center, Herzliya
ARIEL SHAMIR, The Interdisciplinary Center, Herzliya

Fig. 1. Top: landmark detection results on artistic portraits with different styles allows to define the geometric style of an artist. Bottom: results of style transfer of portraits using various artists' geometric style including Modigliani, Picasso, Keane, Leger and Foujita.
From left to right: Portrait of Bindo Altoviti, 1515 by Raphael courtesy WikiArt [Public Domain] via (http://bit.ly/2HzoPyz), Gypsy Woman with a Baby, 1919 by Amedeo Modigliani courtesy WikiArt [Public Domain] via (http://bit.ly/2EbAWkn), Two Women with Rice Cakes and Swords, 1844-1845 by Utagawa Kunisada courtesy Van Gogh Museum [Public Domain] via (http://bit.ly/2JUuFh1), Little Girl with Doll, 1918 by Tsuguharu Foujita courtesy WikiArt [Public Domain US] via (http://bit.ly/2Q82xbf), Portrait of the Composer Anton von Webern, 1914 by Oskar Kokoschka courtesy WikiArt [Public Domain US] via (http://bit.ly/30AuxsS), Woman with Peanuts, 1962 ©Estate of Roy Lichtenstein courtesy Image-Duplicator [Fair Use] via (http://bit.ly/2HA3DIF). Natural face images from [Minear and Park 2004], used with permission.

Facial Landmark detection in natural images is a very active research domain. Impressive progress has been made in recent years, with the rise of neural-network based methods and large-scale datasets. However, it is still a challenging and largely unexplored problem in the artistic portraits domain. Compared to natural face images, artistic portraits are much more diverse. They contain a much wider style variation in both geometry and texture and are more complex to analyze. Moreover, datasets that are necessary to train neural networks are unavailable.

We propose a method for artistic augmentation of natural face images that enables training deep neural networks for landmark detection in artistic portraits. We utilize conventional facial landmarks datasets, and transform their content from natural images into "artistic face" images. In addition, we use a feature-based landmark correction step, to reduce the dependency between the different facial features, which is necessary due to position and shape variations of facial landmarks in artworks. To evaluate our landmark detection framework, we created an "Artistic-Faces" dataset, containing 160 artworks of various art genres, artists and styles, with a large variation in both geometry and texture. Using our method, we can detect facial features in artistic portraits and analyze their geometric style. This allows the definition of signatures for artistic styles of artworks and artists, that encode both the geometry and the texture style. It also allows us to present a geometric-aware style transfer method for portraits.

Authors' addresses: Jordan Yaniv, Tel Aviv University & The Interdisciplinary Center, Herzliya, jordanya@mail.tau.ac.il; Yael Newman, Tel Aviv University & The Interdisciplinary Center, Herzliya, yaelnewman@mail.tau.ac.il; Ariel Shamir, The Interdisciplinary Center, Herzliya, arik@idc.ac.il.

0730-0301/2019/7-ART60 $15.00
https://doi.org/10.1145/3306346.3322984

CCS Concepts: • **Computing methodologies** → *Image processing*; Image representations; *Non-photorealistic rendering*; *Neural networks.*

Additional Key Words and Phrases: facial landmark detection , neural networks, artistic image augmentation, geometry aware style transfer

ACM Reference Format:
Jordan Yaniv, Yael Newman, and Ariel Shamir. 2019. The Face of Art: Landmark Detection and Geometric Style in Portraits. *ACM Trans. Graph.* 38, 4, Article 60 (July 2019), 15 pages. https://doi.org/10.1145/3306346.3322984

1 INTRODUCTION

Portraiture has been an important part of art going back as far as 5000 years ago to ancient Egypt. Before the invention of photography, a painted, sculpted, or drawn portrait was the only way

Distortion-Free Wide-Angle Portraits on Camera Phones

YICHANG SHIH, WEI-SHENG LAI, and CHIA-KAI LIANG, Google

(a) A wide-angle photo with distortions on subjects' faces. (b) Distortion-free photo by our method.

Fig. 1. (a) A group selfie taken by a wide-angle 97° field-of-view phone camera. The perspective projection renders unnatural look to faces on the periphery: they are stretched, twisted, and squished. (b) Our algorithm restores all the distorted face shapes and keeps the background unaffected.

Photographers take wide-angle shots to enjoy expanding views, group portraits that never miss anyone, or composite subjects with spectacular scenery background. In spite of the rapid proliferation of wide-angle cameras on mobile phones, a wider field-of-view (FOV) introduces a stronger perspective distortion. Most notably, faces are stretched, squished, and skewed, to look vastly different from real-life. Correcting such distortions requires professional editing skills, as trivial manipulations can introduce other kinds of distortions. This paper introduces a new algorithm to undistort faces without affecting other parts of the photo. Given a portrait as an input, we formulate an optimization problem to create a content-aware warping mesh which *locally* adapts to the stereographic projection on facial regions, and seamlessly evolves to the perspective projection over the background. Our new energy function performs effectively and reliably for a large group of subjects in the photo. The proposed algorithm is fully automatic and operates at an interactive rate on the mobile platform. We demonstrate promising results on a wide range of FOVs from 70° to 120°.

CCS Concepts: • **Computing methodologies** → **Computational photography**; **Image processing**.

Additional Key Words and Phrases: Content-Aware Warping, Computational Photography, Perspective Correction.

ACM Reference Format:
YiChang Shih, Wei-Sheng Lai, and Chia-Kai Liang. 2019. Distortion-Free Wide-Angle Portraits on Camera Phones. *ACM Trans. Graph.* 38, 4, Article 61 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3322948

Authors' address: YiChang Shih, yichang@google.com; Wei-Sheng Lai, wlai24@ucmerced.edu; Chia-Kai Liang, ckliang@google.com, 1600 Amphitheatre Parkway, Mountain View, CA, Google.

0730-0301/2019/7-ART61
https://doi.org/10.1145/3306346.3322948

1 INTRODUCTION

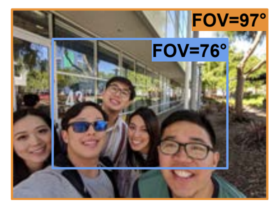

Empowered with extra peripheral vision to "see", a wide-angle lens is incredibly capable of capturing commemorative moments filled with more people and landscapes, and has been widely used for wedding, sport, landscape, and street photography. For mobile phones, as illustrated in the thumbnail at right, zooming out from normal to wide FOV enables group selfies with friends and families using a handheld camera. Adopting wide-angle cameras has been a recent trend among premium phones. For example, LG G6 has a 100° FOV front camera and a 125° FOV rear camera.

Unfortunately, a wide-angle lens distorts faces when projecting the surrounding world onto a flat image. It leads to unnatural, wider, asymmetric, and unpleasant faces as shown in the inset above and Fig. 1a, and gives misleading impressions of the subjects. In the modern mobile era, people take and share portrait shots, selfies, and group selfies by phones all the time [Izadinia et al. 2015]. However, in our study, all wide-angle camera phones suffer from either perspective distortion, or fish-eye like artifacts that bend straight edges on buildings, facades, interiors, and window frames.

We present an automatic algorithm to reverse perspective distortion on portraitures, so that everyone in the photo looks natural and real. Given an input image, we compute the subject mask to assign per-vertex weights on a coarse mesh over the input image. Then, the core of our approach formulates energy terms that encourage facial vertices to locally emulate the stereographic projection, a conformal mapping between a sphere and a plane, for distortion restoration. Our output combines both the stereographic and perspective projections on a single image. The energy function encourages smooth transitions between the two conflicting projections at face boundary. Different from existing works in generic

Wallpaper Pattern Alignment along Garment Seams

KATJA WOLFF and OLGA SORKINE-HORNUNG, ETH Zurich, Switzerland

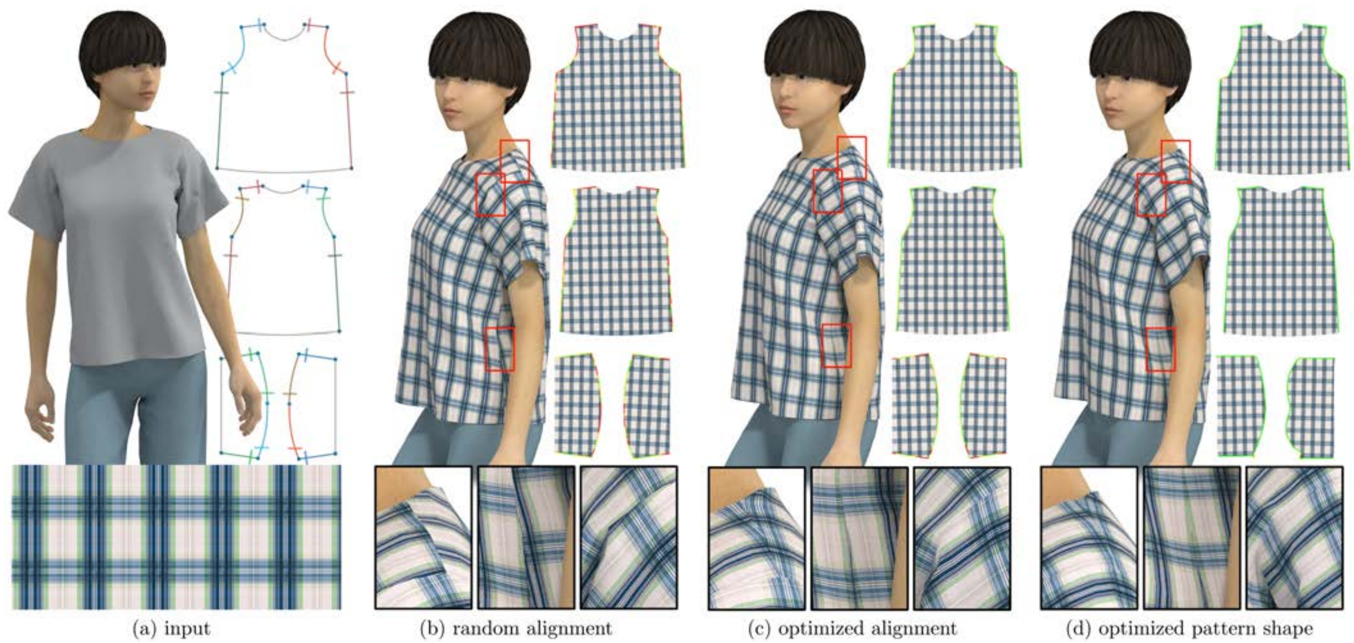

Fig. 1. Our method supports the production of high quality textured garment designs by optimizing the continuity and symmetry properties of the texture across seams. Our input is a sewing pattern and the desired fabric texture, or a 2D wallpaper pattern (a). A random placement of the sewing pattern on the fabric is likely to result in obvious misalignment of the texture on the seams, which is visually displeasing (b). Our method computes an optimal placement, such that misalignments of the repetitive wallpaper pattern are minimized (c). For most garment shapes, it is impossible to achieve a perfectly seamless appearance without changing the shape of the sewing pattern itself. Therefore we additionally provide the option to slightly alter the sewing pattern shape to obtain a much better texture fit along the seams (d). The improved visual appearance is noticeable in the 3D simulation of the garment, while the shape remains close to the original design. See also Fig. 3. The mismatch at each seam edge is visualized on a color scale from green (0 cm) to red (2 cm).

Despite recent developments towards on-demand, individualized garment design and fabrication, the majority of processes in the fashion industry are still inefficient and heavily dependent on manual work. A significant amount of recent research in this area has been focused on supporting designers to digitally create sewing patterns and shapes, but there is little work on textured fabrics. Aligning textile patterns like stripes or plaid along garment seams requires an experienced tailor and is thus reserved only for expensive, high-end garments. We present an interactive algorithm for automatically aligning repetitive textile patterns along seams for a given garment, allowing a user to make design choices at each step of our pipeline. Our approach is based on the 17 wallpaper groups and the symmetries they exhibit. We exploit these symmetries to optimize the alignment of the sewing pattern with the textured fabric for each of its pieces, determining where to cut the fabric. We optionally alter the sewing pattern slightly for a perfect fit along seams, without visibly changing the 3D shape of the garment. The pieces can then be cut automatically by a CNC or laser cutter. Our approach fits within the pipeline of digital garment design, eliminating the difficult, manual step of aligning and cutting the garment pieces by hand.

CCS Concepts: • **Computing methodologies → Computer graphics**; **Shape modeling**; **Mesh geometry models**; *Texturing*.

Additional Key Words and Phrases: computational fabrication, garment modeling, wallpaper patterns, N-RoSy fields, non-convex optimization

ACM Reference Format:
Katja Wolff and Olga Sorkine-Hornung. 2019. Wallpaper Pattern Alignment along Garment Seams. *ACM Trans. Graph.* 38, 4, Article 62 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3322991

Authors' address: Katja Wolff, katja.wolff@inf.ethz.ch; Olga Sorkine-Hornung, sorkine@inf.ethz.ch, ETH Zurich, Switzerland.

0730-0301/2019/7-ART62
https://doi.org/10.1145/3306346.3322991

1 INTRODUCTION

Growing interest in fabrication processes in the computer science and engineering community, coupled with the dropping costs of machinery and material, have stimulated a lot of research on computational fabrication. The possibility to quickly produce real world-objects from digital models has transformed many sectors like 3D printing, architecture, and various industrial processes. Though

Visual Knitting Machine Programming

VIDYA NARAYANAN*, Carnegie Mellon University
KUI WU*, University of Utah
CEM YUKSEL, University of Utah
JAMES MCCANN, Carnegie Mellon University

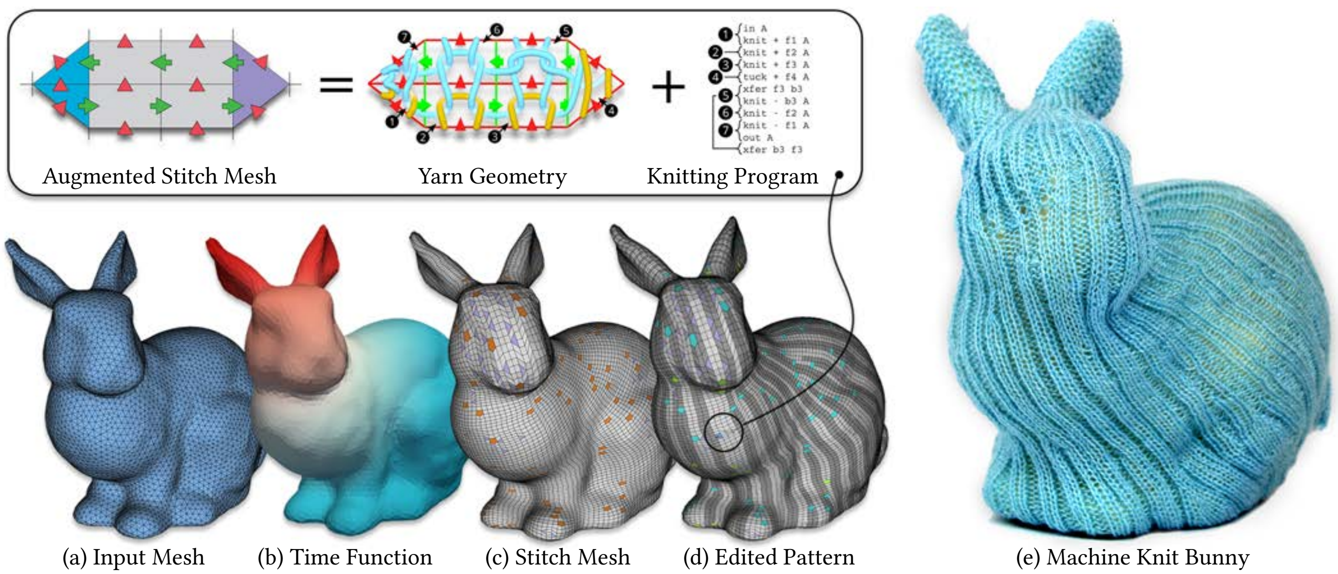

Fig. 1. ***Stages of our visual knit programming system:*** *(a) Our system begins with an input mesh; (b) generates a knitting time function; (c) remeshes the surface to create an augmented stitch mesh; (d) allows the user to interactively edit and add patterns, textures, and colorwork; and (e) generates instructions for fabrication on an industrial knitting machine. At the core of our interface is the augmented stitch mesh, which associates yarn geometry, dependency information, and a knitting program with each face.*

Industrial knitting machines are commonly used to manufacture complicated shapes from yarns; however, designing patterns for these machines requires extensive training. We present the first general visual programming interface for creating 3D objects with complex surface finishes on industrial knitting machines. At the core of our interface is a new, augmented, version of the stitch mesh data structure. The augmented stitch mesh stores low-level knitting operations per-face and encodes the dependencies between faces using directed edge labels. Our system can generate knittable augmented stitch meshes from 3D models, allows users to edit these meshes in a way that preserves their knittability, and can schedule the execution order and location of each face for production on a knitting machine. Our system is general, in that its knittability-preserving editing operations are sufficient to transform between any two machine-knittable stitch patterns with the same orientation on the same surface. We demonstrate the power and flexibility of our pipeline by using it to create and knit objects featuring a wide range of patterns and textures, including intarsia and Fair Isle colorwork; knit and purl textures; cable patterns; and laces.

CCS Concepts: • **Computing methodologies → Mesh geometry models**; • **Applied computing** → *Computer-aided manufacturing*.

Additional Key Words and Phrases: automatic knitting, fabrication, stitch meshes

ACM Reference Format:
Vidya Narayanan, Kui Wu, Cem Yuksel, and James McCann. 2019. Visual Knitting Machine Programming. *ACM Trans. Graph.* 38, 4, Article 63 (July 2019), 13 pages. https://doi.org/10.1145/3306346.3322995

*Co-first authors; equal contribution.

Authors' addresses: Vidya Narayanan, Carnegie Mellon University; Kui Wu, University of Utah; Cem Yuksel, University of Utah; James McCann, Carnegie Mellon University.

0730-0301/2019/7-ART63 $15.00
https://doi.org/10.1145/3306346.3322995

1 INTRODUCTION

Computer-controlled knitting machines are powerful tools for computer-aided fabrication, and are widely used in the garment and accessory industries. When properly programmed, they can turn yarns into soft 3D surfaces in a wide range of shapes, textures, and colors. Knitting machines create these objects by using a small vocabulary of operations which manipulate loops on their *needle beds*, two long rows of loop storage locations. Once programmed,

Computational Peeling Art Design

HAO LIU[*], University of Science and Technology of China, China
XIAO-TENG ZHANG[*], University of Science and Technology of China, China
XIAO-MING FU[†], University of Science and Technology of China, China
ZHI-CHAO DONG, University of Science and Technology of China, China
LIGANG LIU[†], University of Science and Technology of China, China

(a) A parrot standing on a branch **(b)** Lizard **(c)** Fern

Fig. 1. Peeling artworks generated by our design system. Crafting along the drawn curves on the citruses (top left), the citruses are unfolded into a parrot standing on a branch (a), a lizard (b), or a fern (c).

Some artists peel citrus fruits into a variety of elegant 2D shapes, depicting animals, plants, and cartoons. It is a creative art form, called *Citrus Peeling Art*. This art form follows the conservation principle, i.e., each shape must be created using one entire peel. Central to this art is finding optimal cut lines so that the citruses can be cut and unfolded into the desired shapes. However, it is extremely difficult for users to imagine and generate cuts for their desired shapes. To this end, we present a computational method for citrus peeling art designs. Our key insight is that instead of solving the difficult cut generation problem, we map a designed input shape onto a citrus in an attempt to cover the entire citrus and use the mapped boundary to generate the cut paths. Sometimes, a mapped shape is unable to completely cover a citrus. Consequently, we have developed five customized ways of interaction that are used to rectify the input shape so that it is suitable for citrus peeling art. The mapping process and user interactions are iteratively conducted to satisfy a user's design intentions. A large number of experiments, including a formative user study, demonstrate the capability and practicability of our method for peeling art design and construction.

CCS Concepts: • **Computing methodologies → Shape modeling**.

Additional Key Words and Phrases: peeling art, mesh cutting, deformation, parameterizations, mapping

ACM Reference Format:
Hao Liu, Xiao-Teng Zhang, Xiao-Ming Fu, Zhi-Chao Dong, and Ligang Liu. 2019. Computational Peeling Art Design. *ACM Trans. Graph.* 38, 4, Article 64 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3323000

[*]Joint first authors
[†]The corresponding authors

Authors' addresses: Hao Liu, University of Science and Technology of China, China, pllh@mail.ustc.edu.cn; Xiao-Teng Zhang, University of Science and Technology of China, China, zxt2015@mail.ustc.edu.cn; Xiao-Ming Fu, University of Science and Technology of China, China, fuxm@ustc.edu.cn; Zhi-Chao Dong, University of Science and Technology of China, China, dzc206@mail.ustc.edu.cn; Ligang Liu, University of Science and Technology of China, China, lgliu@ustc.edu.cn.

0730-0301/2019/7-ART64 $15.00
https://doi.org/10.1145/3306346.3323000

1 INTRODUCTION

Computational art assists users to easily visualize, design, and create generative art in algorithmic and programmatic ways [Bickel et al. 2018; Wood et al. 2016]. *Citrus Peeling Art*, created by Yoshihiro Okada [Okada 2010], allows users to peel a citrus fruit along guidance lines and unfold the whole citrus into desired 2D shapes, as shown in Fig. 2. [1] This fascinating art form offers an opportunity for enjoyment and creativity and has been widely used in children's education.

However, only a limited number of pieces have been created so far [Okada 2010]. To create new samples of this art form, a designer needs a fertile imagination and a large number of trial-and-error experiments, which is non-trivial and time consuming. Therefore, it is highly challenging for general users to peel citruses into customized shapes.

[1]More examples can be found at https://youtube.com/user/yosigirito/videos.

Neural Volumes: Learning Dynamic Renderable Volumes from Images

STEPHEN LOMBARDI, Facebook Reality Labs
TOMAS SIMON, Facebook Reality Labs
JASON SARAGIH, Facebook Reality Labs
GABRIEL SCHWARTZ, Facebook Reality Labs
ANDREAS LEHRMANN, Facebook Reality Labs
YASER SHEIKH, Facebook Reality Labs

Fig. 1. Renderings of objects captured and modeled by our system. The input to our method consists of synchronized and calibrated multi-view video. We build a dynamic, volumetric representation of the scene by training an encoder-decoder network end-to-end using a differentiable ray marching algorithm.

Modeling and rendering of dynamic scenes is challenging, as natural scenes often contain complex phenomena such as thin structures, evolving topology, translucency, scattering, occlusion, and biological motion. Mesh-based reconstruction and tracking often fail in these cases, and other approaches (e.g., light field video) typically rely on constrained viewing conditions, which limit interactivity. We circumvent these difficulties by presenting a learning-based approach to representing dynamic objects inspired by the integral projection model used in tomographic imaging. The approach is supervised directly from 2D images in a multi-view capture setting and does not require explicit reconstruction or tracking of the object. Our method has two primary components: an encoder-decoder network that transforms input images into a 3D volume representation, and a differentiable ray-marching operation that enables end-to-end training. By virtue of its 3D representation, our construction extrapolates better to novel viewpoints compared to screen-space rendering techniques. The encoder-decoder architecture learns a latent representation of a dynamic scene that enables us to produce novel content sequences not seen during training. To overcome memory limitations of voxel-based representations, we learn a dynamic irregular grid structure implemented with a warp field during ray-marching. This structure greatly improves the apparent resolution and reduces grid-like artifacts and jagged motion. Finally, we demonstrate how to incorporate surface-based representations into our volumetric-learning framework for applications where the highest resolution is required, using facial performance capture as a case in point.

CCS Concepts: • **Computing methodologies → Neural networks**; **Rendering**; **Volumetric models**.

Additional Key Words and Phrases: Volumetric Rendering, Volume Warping, Ray Potentials, Differentiable Ray Marching

ACM Reference Format:
Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. 2019. Neural Volumes: Learning Dynamic Renderable Volumes from Images. *ACM Trans. Graph.* 38, 4, Article 65 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3323020

1 INTRODUCTION

Polygon meshes are an extremely popular representation for 3D geometry in photo-realistic scenes. Mesh-based representations efficiently model solid surfaces and can be paired with sophisticated reflectance functions to generate compelling renderings of natural scenes. In addition, there has been significant progress recently in optimization techniques to support real-time ray-tracing, allowing for interactivity and immersion in demanding applications such as Virtual Reality (VR). However, little of the interactive photo-real content available today is data-driven because many real-world phenomena are challenging to reconstruct and track with high fidelity. State-of-the-art motion capture systems struggle to handle complex occlusions (e.g., running hands through one's hair), to account for reflectance variability (e.g., specularities in the sheen of a moving object), or to track topological evolution in dynamic participating

Authors' addresses: Stephen Lombardi, Facebook Reality Labs, stephen.lombardi@fb.com; Tomas Simon, Facebook Reality Labs, tomas.simon@fb.com; Jason Saragih, Facebook Reality Labs, jason.saragih@fb.com; Gabriel Schwartz, Facebook Reality Labs, gabe.schwartz@oculus.com; Andreas Lehrmann, Facebook Reality Labs, asml@fb.com; Yaser Sheikh, Facebook Reality Labs, yasers@fb.com.

0730-0301/2019/7-ART65
https://doi.org/10.1145/3306346.3323020

Deferred Neural Rendering: Image Synthesis using Neural Textures

JUSTUS THIES, Technical University of Munich
MICHAEL ZOLLHÖFER, Stanford University
MATTHIAS NIESSNER, Technical University of Munich

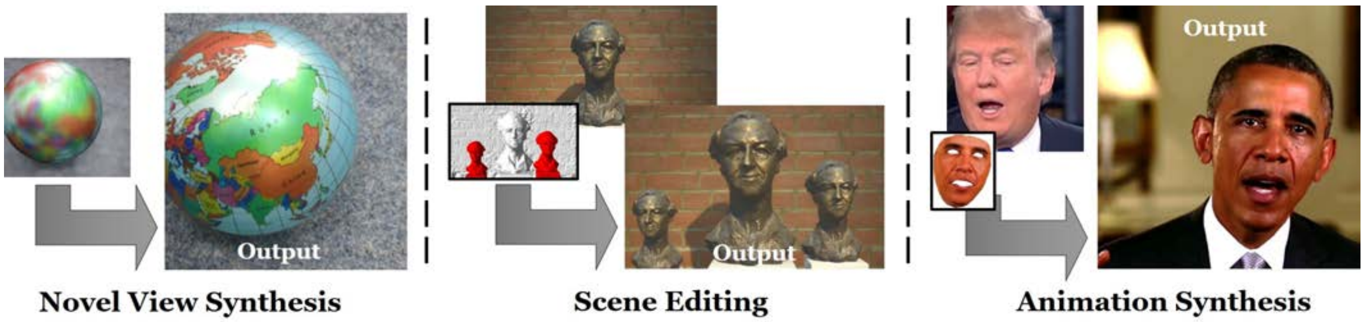

Fig. 1. We present an image synthesis approach that learns object-specific neural textures which can be interpreted by a neural renderer. Our approach can be trained end-to-end with real data, allowing us to re-synthesize novel views of static objects, edit scenes, as well as re-render dynamic animated surfaces [0].

The modern computer graphics pipeline can synthesize images at remarkable visual quality; however, it requires well-defined, high-quality 3D content as input. In this work, we explore the use of imperfect 3D content, for instance, obtained from photo-metric reconstructions with noisy and incomplete surface geometry, while still aiming to produce photo-realistic (re-)renderings. To address this challenging problem, we introduce *Deferred Neural Rendering*, a new paradigm for image synthesis that combines the traditional graphics pipeline with learnable components. Specifically, we propose *Neural Textures*, which are learned feature maps that are trained as part of the scene capture process. Similar to traditional textures, neural textures are stored as maps on top of 3D mesh proxies; however, the high-dimensional feature maps contain significantly more information, which can be interpreted by our new deferred neural rendering pipeline. Both neural textures and deferred neural renderer are trained end-to-end, enabling us to synthesize photo-realistic images even when the original 3D content was imperfect. In contrast to traditional, black-box 2D generative neural networks, our 3D representation gives us explicit control over the generated output, and allows for a wide range of application domains. For instance, we can synthesize temporally-consistent video re-renderings of recorded 3D scenes as our representation is inherently embedded in 3D space. This way, neural textures can be utilized to coherently re-render or manipulate existing video content in both static and dynamic environments at real-time rates. We show the effectiveness of our approach in several experiments on novel view synthesis, scene editing, and facial reenactment, and compare to state-of-the-art approaches that leverage the standard graphics pipeline as well as conventional generative neural networks.

Authors' addresses: Justus Thies, Technical University of Munich, justus.thies@tum.de; Michael Zollhöfer, Stanford University, zollhoefer@cs.stanford.edu; Matthias Nießner, Technical University of Munich, niessner@tum.de.

0730-0301/2019/7-ART66 $15.00
https://doi.org/10.1145/3306346.3323035

CCS Concepts: • **Computing methodologies → Computer vision**; **Computer graphics**.

Additional Key Words and Phrases: neural rendering, neural texture, novel view synthesis, facial reenactment

ACM Reference Format:
Justus Thies, Michael Zollhöfer, and Matthias Nießner. 2019. Deferred Neural Rendering: Image Synthesis using Neural Textures. *ACM Trans. Graph.* 38, 4, Article 66 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3323035

1 INTRODUCTION

The current computer graphics pipeline has evolved during the last decades, and is now able to achieve remarkable rendering results. From a fixed function pipeline around the rasterization unit, the graphics pipeline has turned into a programmable rendering pipeline. Based on this pipeline the Siggraph community has established rendering techniques that now achieve nearly photo-realistic imagery. While the visuals are stunning, a major drawback of these classical approaches is the need of well-defined input data, including a precise definition of the surface geometry, the underlying material properties, and the scene illumination. In movie or video game productions, this underlying 3D content is manually-created by skilled artists in countless working hours. An alternative is to obtain 3D content from real-world scenes by using 3D reconstruction techniques. However, given the inherent limitations of state-of-the-art 3D reconstruction approaches, such as noisy, oversmoothed geometry or occlusions, the obtained 3D content is imperfect. From this captured content, it is nearly impossible to re-synthesize photo-realistic images with the existing computer graphics pipeline and rendering techniques.

In this work, we assume that captured 3D content will always suffer from reconstruction artifacts in one way or another. Rather than aiming to fix the artifacts in the 3D content, we propose to

[0] Third Party Material: Trump, MSNBC (https://youtu.be/Tsh_V3U7EfU) [Fair Use]
Obama, The Obama White House (https://youtu.be/d-VaUaTF3_k) [Public Domain]

VR Facial Animation via Multiview Image Translation

SHIH-EN WEI, JASON SARAGIH, TOMAS SIMON, ADAM W. HARLEY*, STEPHEN LOMBARDI, MICHAL PERDOCH, ALEXANDER HYPES, DAWEI WANG, HERNAN BADINO, and YASER SHEIKH
Facebook Reality Labs

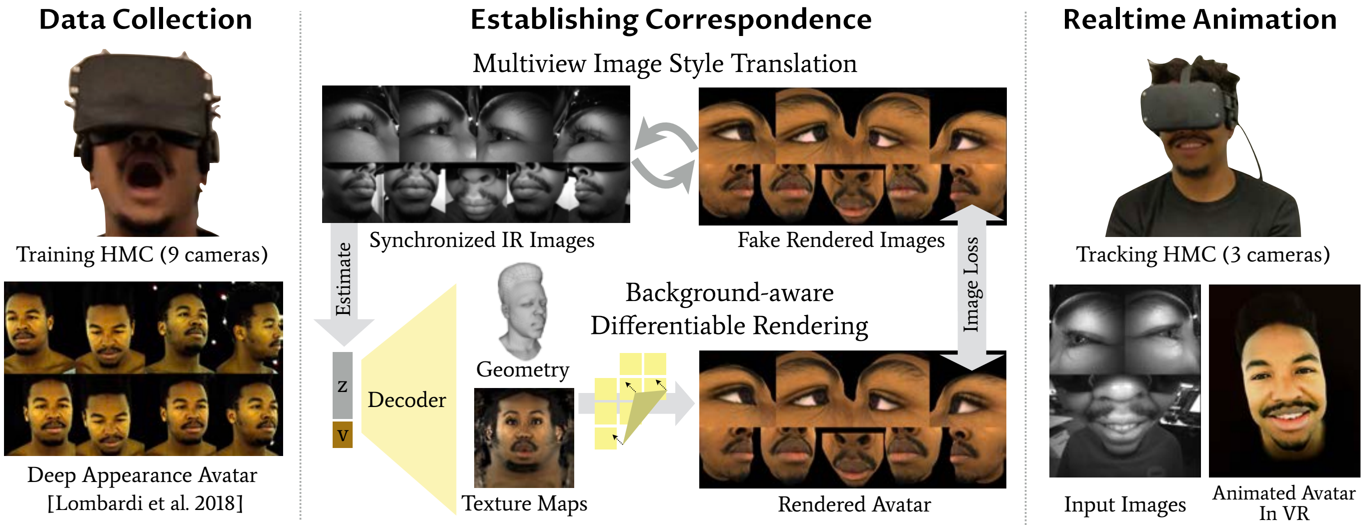

Fig. 1. We present a **VR realtime facial animation** system with headset mounted cameras (HMC) which augment a standard head mounted display (HMD). Our method establishes precise correspondence between 9 camera images from a *training* HMC and the parameters of a photorealistic avatar. Finally, we use a common subset of 3 cameras on a *tracking* HMC to animate the avatar in realtime.

A key promise of Virtual Reality (VR) is the possibility of remote social interaction that is more immersive than any prior telecommunication media. However, existing social VR experiences are mediated by inauthentic digital representations of the user (i.e., stylized avatars). These stylized representations have limited the adoption of social VR applications in precisely those cases where immersion is most necessary (e.g., professional interactions and intimate conversations). In this work, we present a bidirectional system that can animate avatar heads of both users' full likeness using consumer-friendly headset mounted cameras (HMC). There are two main challenges in doing this: unaccommodating camera views and the image-to-avatar domain gap. We address both challenges by leveraging constraints imposed by multiview geometry to establish precise image-to-avatar correspondence, which are then used to learn an end-to-end model for real-time tracking. We present designs for a *training* HMC, aimed at data-collection and model building, and a *tracking* HMC for use during interactions in VR. Correspondence between the avatar and the HMC-acquired images are automatically found through self-supervised multiview image translation, which does not require manual annotation or one-to-one correspondence between domains. We evaluate the system on a variety of users and demonstrate significant improvements over prior work.

CCS Concepts: • **Human-centered computing** → **Virtual reality**; • **Computing methodologies** → **Computer vision**; **Unsupervised learning**; **Animation**.

Additional Key Words and Phrases: Face Tracking, Unsupervised Image Style Transfer, Differentiable Rendering

ACM Reference Format:
Shih-En Wei, Jason Saragih, Tomas Simon, Adam W. Harley, Stephen Lombardi, Michal Perdoch, Alexander Hypes, Dawei Wang, Hernan Badino, and Yaser Sheikh. 2019. VR Facial Animation via Multiview Image Translation. *ACM Trans. Graph.* 38, 4, Article 67 (July 2019), 16 pages. https://doi.org/10.1145/3306346.3323030

1 INTRODUCTION

Virtual Reality (VR) has seen increased ubiquity in recent years. This has opened up the possibility for remote collaboration and interaction that is more engaging and immersive than achievable through other media. Concurrently, there has been great progress in generating accurate digital doubles and avatars. Driven by the gaming and movie industries, a number of compelling demonstrations of state of the art systems have recently attracted interest in the community [Epic Games 2017; Hellblade 2018; Magic Leap 2018; Seymour et al. 2017; Unreal Engine 4 2018]. These systems show

*Currently at Carnegie Mellon University, work done while at Facebook Reality Labs.

Authors' address: Shih-En Wei, shih-en.wei@fb.com; Jason Saragih, jason.saragih@fb.com; Tomas Simon, tomas.simon@fb.com; Adam W. Harley, aharley@cmu.edu; Stephen Lombardi, stephen.lombardi@fb.com; Michal Perdoch, michal.perdoch@oculus.com; Alexander Hypes, alexander.hypes@oculus.com; Dawei Wang, dawei.wang@oculus.com; Hernan Badino, hernan.badino@oculus.com; Yaser Sheikh, yasers@fb.com, Facebook Reality Labs, Pittsburgh, PA.

0730-0301/2019/7-ART67
https://doi.org/10.1145/3306346.3323030

Text-based Editing of Talking-head Video

OHAD FRIED, Stanford University
AYUSH TEWARI, Max Planck Institute for Informatics
MICHAEL ZOLLHÖFER, Stanford University
ADAM FINKELSTEIN, Princeton University
ELI SHECHTMAN, Adobe
DAN B GOLDMAN
KYLE GENOVA, Princeton University
ZEYU JIN, Adobe
CHRISTIAN THEOBALT, Max Planck Institute for Informatics
MANEESH AGRAWALA, Stanford University

Fig. 1. We propose a novel text-based editing approach for talking-head video. Given an edited transcript, our approach produces a realistic output video in which the dialogue of the speaker has been modified and the resulting video maintains a seamless audio-visual flow (i.e. no jump cuts).

Editing talking-head video to change the speech content or to remove filler words is challenging. We propose a novel method to edit talking-head video based on its transcript to produce a realistic output video in which the dialogue of the speaker has been modified, while maintaining a seamless audio-visual flow (i.e. no jump cuts). Our method automatically annotates an input talking-head video with phonemes, visemes, 3D face pose and geometry, reflectance, expression and scene illumination per frame. To edit a video, the user has to only edit the transcript, and an optimization strategy then chooses segments of the input corpus as base material. The annotated parameters corresponding to the selected segments are seamlessly stitched together and used to produce an intermediate video representation in which the lower half of the face is rendered with a parametric face model. Finally, a recurrent video generation network transforms this representation to a photorealistic video that matches the edited transcript. We demonstrate a large variety of edits, such as the addition, removal, and alteration of words, as well as convincing language translation and full sentence synthesis.

CCS Concepts: • **Information systems** → *Video search*; *Speech / audio search*; • **Computing methodologies** → *Computational photography*; *Reconstruction*; *Motion processing*; *Graphics systems and interfaces.*

Additional Key Words and Phrases: Text-based video editing, talking heads, visemes, dubbing, face tracking, face parameterization, neural rendering.

ACM Reference Format:
Ohad Fried, Ayush Tewari, Michael Zollhöfer, Adam Finkelstein, Eli Shechtman, Dan B Goldman, Kyle Genova, Zeyu Jin, Christian Theobalt, and Maneesh Agrawala. 2019. Text-based Editing of Talking-head Video. *ACM Trans. Graph.* 38, 4, Article 68 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3323028

Authors' addresses: Ohad Fried, Stanford University; Ayush Tewari, Max Planck Institute for Informatics; Michael Zollhöfer, Stanford University; Adam Finkelstein, Princeton University; Eli Shechtman, Adobe; Dan B Goldman; Kyle Genova, Princeton University; Zeyu Jin, Adobe; Christian Theobalt, Max Planck Institute for Informatics; Maneesh Agrawala, Stanford University.

0730-0301/2019/7-ART68 $15.00
https://doi.org/10.1145/3306346.3323028

1 INTRODUCTION

Talking-head video – framed to focus on the face and upper body of a speaker – is ubiquitous in movies, TV shows, commercials, YouTube video logs, and online lectures. Editing such pre-recorded video is challenging, but can be needed to emphasize particular content, remove filler words, correct mistakes, or more generally match the editor's intent. Using current video editing tools, like Adobe Premiere, skilled editors typically scrub through raw video footage to find relevant segments and assemble them into the desired story. They must carefully consider where to place cuts so as to minimize disruptions to the overall audio-visual flow.

Anisotropic Elasticity for Inversion-Safety and Element Rehabilitation

THEODORE KIM, Pixar Animation Studios
FERNANDO DE GOES, Pixar Animation Studios
HAYLEY IBEN, Pixar Animation Studios

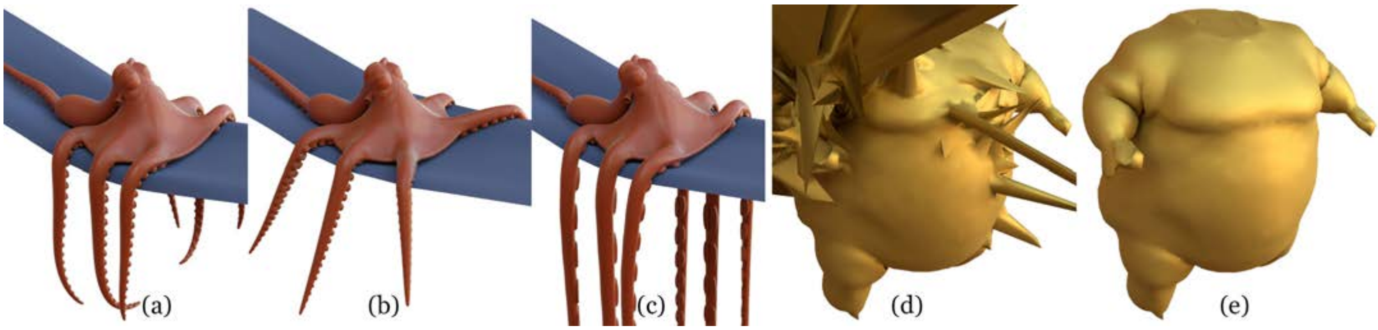

Fig. 1. Our analysis of anisotropic hyperelastic energies yields a novel, robust, and inversion-safe anisotropic energy. Our energy allow fibers along the tentacles (a) to be robustly stiffened by 100× (b) and softened by 100× (c). Our analysis also produces an anisotropic rehabilitation approach to divergent, badly-conditioned simulations (d) that, when applied, allows them to converge (e).

We present an analysis of anisotropic hyperelasticity, specifically transverse isotropy, that obtains closed-form expressions for the eigendecompositions of many common energies. We then use these to build fast and concise Newton implementations. We leverage our analysis in two separate applications. First, we show that existing anisotropic energies are not inversion-safe, and contain spurious stable states under large deformation. We then propose a new anisotropic strain invariant that enables the formulation of a novel, robust, and inversion-safe energy. The new energy fits completely within our analysis, so closed-form expressions are obtained for its eigensystem as well. Secondly, we use our analysis to *rehabilitate* badly-conditioned finite elements. Using this method, we can robustly simulate large deformations even when a mesh contains degenerate, zero-volume elements. We accomplish this by swapping the badly-behaved isotropic direction with a well-behaved anisotropic term. We validate our approach on a variety of examples.

CCS Concepts: • **Computing methodologies → Physical simulation**.

Additional Key Words and Phrases: Anisotropy, element degeneracy, elasticity, physics-based simulation

ACM Reference Format:
Theodore Kim, Fernando De Goes, and Hayley Iben. 2019. Anisotropic Elasticity for Inversion-Safety and Element Rehabilitation. *ACM Trans. Graph.* 38, 4, Article 69 (July 2019), 15 pages. https://doi.org/10.1145/3306346.3323014

1 INTRODUCTION

Anisotropic energies are indispensable when simulating realistic phenomena such as muscles [Lee et al. 2018], plants [Wang et al. 2017], and cloth [Baraff and Witkin 1998], which exhibit directional effects that cannot be captured with isotropic energies alone. Isotropic energies have been extensively analyzed in computer graphics, with efficient methods proposed for both Newton-based [Smith et al. 2018; Stomakhin et al. 2012; Teran et al. 2005] and gradient-based [Bouaziz et al. 2014; Wang and Yang 2016] solvers. Analyses of anisotropic energies have received relatively less attention, with the important exceptions of linear [Cai 2016; Li and Barbič 2015] and spline-based [Xu et al. 2015] orthotropic materials. This knowledge gap is especially visible when incorporating anisotropic energies into Newton-type solvers, because only approximate or brute-force methods are available to project the Hessian back to semi-positive-definiteness.

We address this problem by analyzing the specific anisotropy case of transverse isotropy, i.e. a material that has been strengthened or weakened along one axis. Our analysis shows that many common anisotropic energies from graphics and biomechanics have closed-form eigendecompositions. In fact, we find that there is a family of anisotropic energies that all share the exact same eigenvectors, and the only variation appears in the eigenvalues. These results then enable the construction of fast and concise Newton implementations. We apply this analysis to two applications.

First, we observe that existing transversely isotropic energies are not inversion-safe [Irving et al. 2004], and contain spurious stable states that are especially prevalent under large deformation and high anisotropic stiffness. Under inversion, the anisotropic forces can overwhelm the isotropic forces and drive the simulation towards non-physical states. In order to introduce inversion-safety into these energies, we propose a new anisotropic strain invariant. By carefully avoiding a variety of singularities, we show that a fast, simple, robust, and inversion-safe anisotropic energy can be formulated. The energy is quadratic, so it can be used to introduce both anisotropic stiffening and *softening* to existing isotropic models (Fig. 1(a)-(c)). The energy

Authors' addresses: Theodore Kim, Pixar Animation Studios; Fernando De Goes, Pixar Animation Studios; Hayley Iben, Pixar Animation Studios.

0730-0301/2019/7-ART69
https://doi.org/10.1145/3306346.3323014

Decomposed Optimization Time Integrator for Large-Step Elastodynamics

MINCHEN LI, University of Pennsylvania & Adobe Research
MING GAO, University of Pennsylvania
TIMOTHY LANGLOIS, Adobe Research
CHENFANFU JIANG, University of Pennsylvania
DANNY M. KAUFMAN, Adobe Research

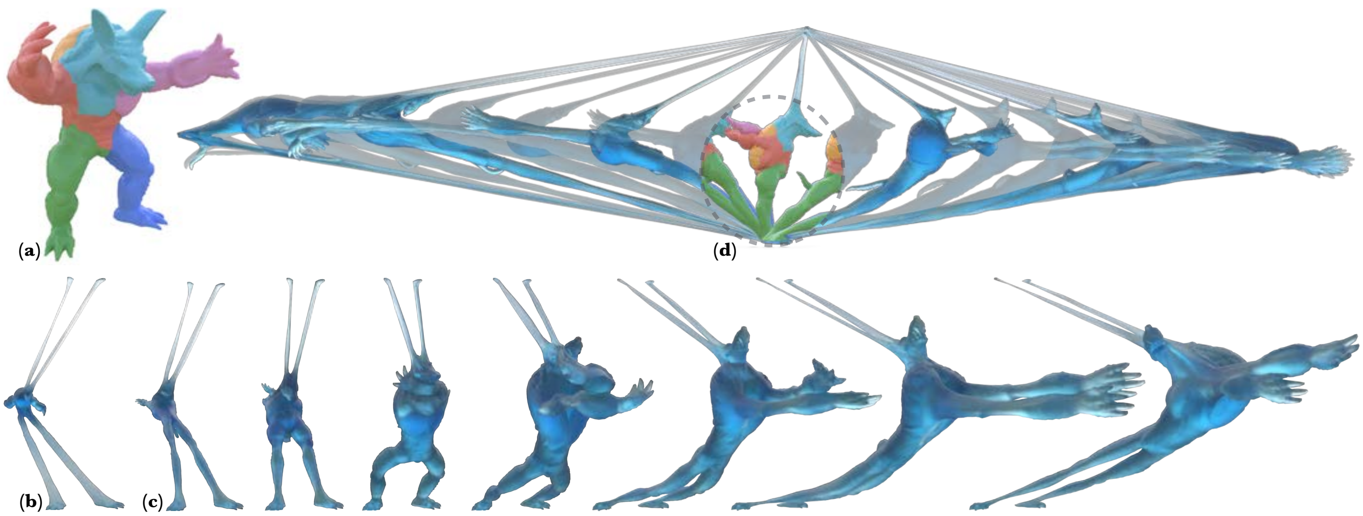

Fig. 1. ***Severe deformation dynamics simulated with large steps.*** **(a)** The Decomposed Optimization Time Integrator (DOT) decomposes spatial domains to generate high-quality simulations of nonlinear materials undergoing large-deformation dynamics. In **(b)** we apply DOT to rapidly stretch and pull an Armadillo backwards. We render in **(c)** a few frames of the resulting slingshot motion right after release. DOT efficiently solves time steps while achieving user-specified accuracies — even when stepping at frame-rate size steps; here at 25 ms. In **(d)** we emphasize the large steps taken by rendering all DOT-simulated time steps from the Armadillo's high-speed trajectory for the first few moments after release.

Simulation methods are rapidly advancing the accuracy, consistency and controllability of elastodynamic modeling and animation. Critical to these advances, we require efficient time step solvers that reliably solve all implicit time integration problems for elastica. While available time step solvers succeed admirably in some regimes, they become impractically slow, inaccurate, unstable, or even divergent in others — as we show here. Towards addressing these needs we present the Decomposed Optimization Time Integrator (DOT), a new domain-decomposed optimization method for solving the per time step, nonlinear problems of implicit numerical time integration. DOT is especially suitable for large time step simulations of deformable bodies with nonlinear materials and high-speed dynamics. It is efficient, automated, and robust at large, fixed-size time steps, thus ensuring stable, continued progress of high-quality simulation output. Across a broad range of extreme and mild deformation dynamics, using frame-rate size time steps with widely varying object shapes and mesh resolutions, we show that DOT always converges to user-set tolerances, generally well-exceeding and always close to the best wall-clock times across all previous nonlinear time step solvers, irrespective of the deformation applied.

CCS Concepts: • **Computing methodologies → Physical simulation**.

Additional Key Words and Phrases: Computational Optimization, Domain Decomposition

ACM Reference Format:
Minchen Li, Ming Gao, Timothy Langlois, Chenfanfu Jiang, and Danny M. Kaufman. 2019. Decomposed Optimization Time Integrator for Large-Step Elastodynamics. *ACM Trans. Graph.* 38, 4, Article 70 (July 2019), 10 pages. https://doi.org/10.1145/3306346.3322951

Authors' addresses: Minchen Li, University of Pennsylvania & Adobe Research, minchernl@gmail.com; Ming Gao, University of Pennsylvania, ming.gao07@gmail.com; Timothy Langlois, Adobe Research, tlangloi@adobe.com; Chenfanfu Jiang, University of Pennsylvania, cffjiang@seas.upenn.edu; Danny M. Kaufman, Adobe Research, kaufman@adobe.com.

0730-0301/2019/7-ART70 $15.00
https://doi.org/10.1145/3306346.3322951

Affine Interpolation in a Lie Group Framework

SUMUKH BANSAL, Dhirubhai Ambani Institute of Information and Communication Technology, India
ADITYA TATU, Dhirubhai Ambani Institute of Information and Communication Technology, India

Affine transformations are of vital importance in many tasks pertaining to motion design and animation. Interpolation of affine transformations is non-trivial. Typically, the given affine transformation is decomposed into simpler components which are easier to interpolate. This may lead to unintuitive results, while in some cases, such solutions may not work. In this work, we propose an interpolation framework which is based on a Lie group representation of the affine transformation. The Lie group representation decomposes the given transformation into simpler and meaningful components, on which computational tools like the exponential and logarithm maps are available in closed form. Interpolation exists for all affine transformations while preserving a few characteristics of the original transformation. A detailed analysis and several experiments of the proposed framework are included.

CCS Concepts: • **Computing methodologies** → *Animation.*

Additional Key Words and Phrases: Affine transformation interpolation, Lie groups, Lie Bodies.

ACM Reference Format:
Sumukh Bansal and Aditya Tatu. 2019. Affine Interpolation in a Lie Group Framework. *ACM Trans. Graph.* 38, 4, Article 71 (July 2019), 16 pages. https://doi.org/10.1145/3306346.3322997

1 INTRODUCTION

3D affine transformations play a pivotal role in many applications pertaining to computer vision, computer graphics and geometry processing. The need to interpolate affine transformations arises ubiquitously in applications involving animation design [Shoemake 1992; Whited et al. 2010], inverse kinematics [Der et al. 2006; Sumner et al. 2005], motion estimation and averaging [Zefran and Kumar 1998], image morphing, and robotics [Selig 2010]. It is well known that the set of matrices $\mathbb{T}$ representing affine transformations forms a Lie group[Ochiai and Anjyo 2013], interpolation on which is non-trivial. Moreover, as far as possible, the interpolated transformation should preserve the properties of the original affine transformation, such as orthogonality and area preservation.

Typical approaches decompose affine transformations into rotational and shear/scale components, after which each component is handled separately. Steady Affine Motion (SAM), a scheme proposed in [Rossignac and Vinacua 2011] for affine interpolation is found to maintain several desired properties. But in cases exhibiting large shear and rotation, the SAM interpolation does not exist.

The proposed approach decomposes any given 3D invertible, orientation-preserving affine transformation[1] into a series of intuitive transformations (each coming from a Lie group) needed to deform a tetrahedron into a fixed canonical tetrahedron. This gives a Lie group representation of the given affine transformation. Thus, an orientation-preserving 3D affine transformation can be represented as a mapping between two specific tetrahedrons, and conversely, a mapping between two oriented tetrahedrons with given correspondence as a unique 3D orientation-preserving affine transformation. This is a generalization of the Lie Bodies representation of 3D triangular meshes introduced by Freifeld *et al.* [Freifeld and Black 2012]. The Lie Bodies approach represents each triangle of a mesh via a specific set of transformations needed to deform a corresponding triangle in a given template mesh to the triangle under consideration. The approach proposed in this paper represents any orientation-preserving 3D affine transformation as a decomposition into three components: a 3D rigid transformation, uniform scaling, and a specific 3D-shear, refer to Figure 1. The interpolation of the affine transformation is obtained using interpolation of the three components. The advantage of this decomposition is that closed-form solutions are known for interpolations of the three components. Moreover, several properties of the original affine transformation are preserved by the proposed scheme.

To summarize, the contributions of the paper are:

(1) The proposed approach interpolates any orientation-preserving 3D affine transformation, unlike the state-of-the-art approach in [Rossignac and Vinacua 2011].
(2) We provide a detailed analysis of the proposed interpolation scheme and show that it has several desirable properties like: (a) preserves isometry, (b) preserves volume, and (c) yields a monotonic change in the volume.
(3) The proposed interpolation scheme can also interpolate two tetrahedrons (and as a special case, triangles) related by any orientation-preserving 3D affine transformation. The interpolation is unique given a correspondence and a vertex ordering. Variation with respect to vertex ordering is analyzed in detail in the paper.

In the next section, a review of the related work is provided. In Section 3, we provide details of the proposed representation of an affine transformation/tetrahedron, along with a discussion on its existence and uniqueness. Section 4 describes the proposed interpolation algorithm and properties of the given affine transformation that are preserved by the interpolations obtained. Two important invariance properties related to our approach are analyzed in Section 5, followed by experiments and results in Section 6. We conclude the paper in Section 7.

[1] An affine transformation $\mathbb{T} : \mathbb{R}^n \to \mathbb{R}^n$ is *orientation-preserving* if $det(\mathbb{T}) > 0$.

Authors' addresses: Sumukh Bansal, Dhirubhai Ambani Institute of Information and Communication Technology, Gandhinagar, Gujrat, India, sumukhbansal@gmail.com; Aditya Tatu, Dhirubhai Ambani Institute of Information and Communication Technology, Gandhinagar, India, aditya_tatu@daiict.ac.in.

0730-0301/2019/7-ART71 $15.00
https://doi.org/10.1145/3306346.3322997

Synthesis of Biologically Realistic Human Motion Using Joint Torque Actuation

YIFENG JIANG, Georgia Institute of Technology, USA
TOM VAN WOUWE, KU Leuven, Belgium
FRIEDL DE GROOTE, KU Leuven, Belgium
C. KAREN LIU, Georgia Institute of Technology, USA

Using joint actuators to drive the skeletal movements is a common practice in character animation, but the resultant torque patterns are often unnatural or infeasible for real humans to achieve. On the other hand, physiologically-based models explicitly simulate muscles and tendons and thus produce more human-like movements and torque patterns. This paper introduces a technique to transform an optimal control problem formulated in the muscle-actuation space to an equivalent problem in the joint-actuation space, such that the solutions to both problems have the same optimal value. By solving the equivalent problem in the joint-actuation space, we can generate human-like motions comparable to those generated by musculotendon models, while retaining the benefit of simple modeling and fast computation offered by joint-actuation models. Our method transforms constant bounds on muscle activations to nonlinear, state-dependent torque limits in the joint-actuation space. In addition, the metabolic energy function on muscle activations is transformed to a nonlinear function of joint torques, joint configuration and joint velocity. Our technique can also benefit policy optimization using deep reinforcement learning approach, by providing a more anatomically realistic action space for the agent to explore during the learning process. We take the advantage of the physiologically-based simulator, OpenSim, to provide training data for learning the torque limits and the metabolic energy function. Once trained, the same torque limits and the energy function can be applied to drastically different motor tasks formulated as either trajectory optimization or policy learning.

CCS Concepts: • **Computing methodologies** → **Animation**; **Physical simulation**; *Supervised learning*; *Reinforcement learning*.

Additional Key Words and Phrases: character animation, trajectory optimization, biomechanics, musculotendon modeling, muscle redundancy problem.

ACM Reference Format:
Yifeng Jiang, Tom Van Wouwe, Friedl De Groote, and C. Karen Liu. 2019. Synthesis of Biologically Realistic Human Motion Using Joint Torque Actuation. *ACM Trans. Graph.* 38, 4, Article 72 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3322966

Fig. 1. Top: A swing motion solved by trajectory optimization in the muscle-actuation space. Bottom: Our proposed method can solve the same task in the joint-actuation space, yielding similar motion but costing fewer iterations and less computation time.

1 INTRODUCTION

Realistic movement of virtual humans plays an integral role in bringing fictional figures to life in films and games, enhancing immersive experiences in VR, and recently, teaching robots how to physically interact with humans [Clegg et al. 2018]. While such applications in robotics and machine learning have created new research avenues for the field of character animation, they also introduce new problems that challenge the existing techniques. Most noticeably, the virtual humans interacting with robots must exhibit not only human-like movements, but also valid joint torques consistent with their physiological capability.

Virtual human is often modeled as articulated rigid bodies with actuated joints that directly and independently generate torques to drive the kinematic movement. While *joint-actuation* simplifies the modeling, simulation, and control of virtual humans, it often produces torque patterns that are unnatural or infeasible for real humans to achieve. Consequently, additional kinematic constraints or motion data are often needed to improve the naturalness of the kinematic trajectories. Alternatively, musculotendon models explicitly simulate the dynamics of muscles and tendons to drive the skeletal system. As such, models based on *muscle-actuation* are able to impose physiologically realistic constraints and energetic cost on the resultant torque trajectories, leading to more human-like

Authors' addresses: Yifeng Jiang, Georgia Institute of Technology, USA, yjiang340@gatech.edu; Tom Van Wouwe, KU Leuven, Belgium, tom.vanwouwe@kuleuven.be; Friedl De Groote, KU Leuven, Belgium, friedl.degroote@kuleuven.be; C. Karen Liu, Georgia Institute of Technology, USA, karenliu@cc.gatech.edu.

0730-0301/2019/7-ART72 $15.00
https://doi.org/10.1145/3306346.3322966

Scalable Muscle-Actuated Human Simulation and Control

SEUNGHWAN LEE, Seoul National University
MOONSEOK PARK, Seoul National University Bundang Hospital
KYOUNGMIN LEE, Seoul National University Bundang Hospital
JEHEE LEE, Seoul National University

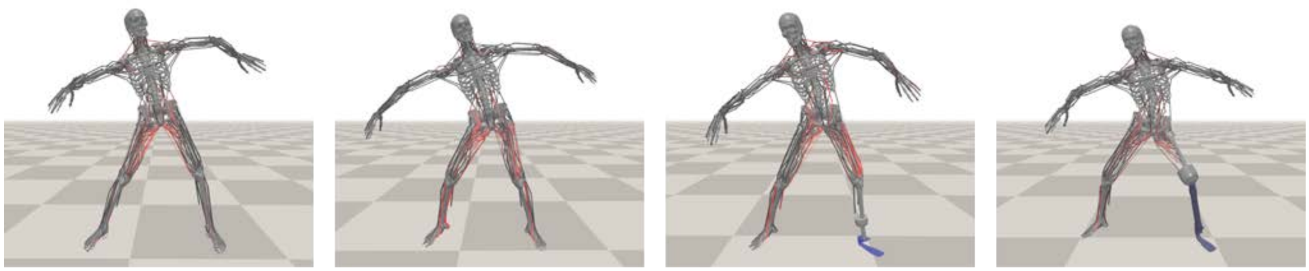

Fig. 1. Physics-based simulation and control of dynamic motor skills actuated by 284 to 346 musculotendon units. (Left to right) Musculoskeletal model with multi-segment feet, two-toe feet, and prosthetic legs.

Many anatomical factors, such as bone geometry and muscle condition, interact to affect human movements. This work aims to build a comprehensive musculoskeletal model and its control system that reproduces realistic human movements driven by muscle contraction dynamics. The variations in the anatomic model generate a spectrum of human movements ranging from typical to highly stylistic movements. To do so, we discuss scalable and reliable simulation of anatomical features, robust control of under-actuated dynamical systems based on deep reinforcement learning, and modeling of pose-dependent joint limits. The key technical contribution is a scalable, two-level imitation learning algorithm that can deal with a comprehensive full-body musculoskeletal model with 346 muscles. We demonstrate the predictive simulation of dynamic motor skills under anatomical conditions including bone deformity, muscle weakness, contracture, and the use of a prosthesis. We also simulate various pathological gaits and predictively visualize how orthopedic surgeries improve post-operative gaits.

CCS Concepts: • **Computing methodologies → Physical simulation**; **Motion processing**.

Additional Key Words and Phrases: Anatomical Human Modeling, Musculoskeletal Modeling, Deep Reinforcement Learning, Joint Range of Motion Modeling, Locomotion Control, Gait Analysis

Authors' addresses: Seunghwan Lee, Department of Computer Science and Engineering, Seoul National University, Seoul, lsw9021@mrl.snu.ac.kr; Moonseok Park, Department of Orthopaedic Surgery, Seoul National University Bundang Hospital, Seoul, pmsmed@gmail.com; Kyoungmin Lee, Department of Orthopaedic Surgery, Seoul National University Bundang Hospital, Seoul, oasis100@snubh.org; Jehee Lee, Department of Computer Science and Engineering, Seoul National University, Seoul, jehee@mrl.snu.ac.kr.

0730-0301/2019/8-ART73
https://doi.org/

ACM Reference Format:
Seunghwan Lee, Moonseok Park, Kyoungmin Lee, and Jehee Lee. 2019. Scalable Muscle-Actuated Human Simulation and Control. *ACM Trans. Graph.* 38, 4, Article 73 (August 2019), 13 pages.

1 INTRODUCTION

Human motion is affected by many anatomical factors such as the geometry of bones, muscle conditions, fatigue, habits, and even emotion. Small changes in anatomical conditions often alter the overall motion and result in distinctive movement patterns of each individual. The musculoskeleton of a human body is a highly-complex dynamic system. The human body has over 600 muscles and the half of them participate in joint movements. Muscle contraction and relaxation are a dynamic process of activating and deactivating tension-generating sites within muscle fibers. The brain sends excitation signal through the nervous system to activate and deactivate individual muscles and thus coordinates full-body movements.

This work aims to build a comprehensive musculoskeletal model and its control system that reproduces realistic human movements driven by muscle contraction dynamics. The variations in the model generate a wide spectrum of human movements ranging from normal (or typical) movements to highly stylistic variants, even to pathologic ones as well. The key technical challenges include accurate and comprehensive musculoskeletal modeling, the scalable and reliable simulation of anatomical features, and the robust control of the under-actuated dynamical system. Our model includes most of the skeletal muscles that serve for moving major joints. Our simulation system reliably deals with muscle contraction dynamics and joint range of motion (ROM) induced by background elasticity of muscles. We also present a new control algorithm based on Deep Reinforcement Learning (DRL).

Recently, DRL has shown its potentials for the control of physically-simulated articulated figures. The control policy represented by

Physics-based Full-body Soccer Motion Control for Dribbling and Shooting

SEOKPYO HONG, Visual Media Lab, KAIST
DASEONG HAN, Handong Global University
KYUNGMIN CHO, Visual Media Lab, KAIST
JOSEPH S. SHIN (formerly SUNG YONG SHIN), Handong Global University and KAIST
JUNYONG NOH*, Visual Media Lab, KAIST

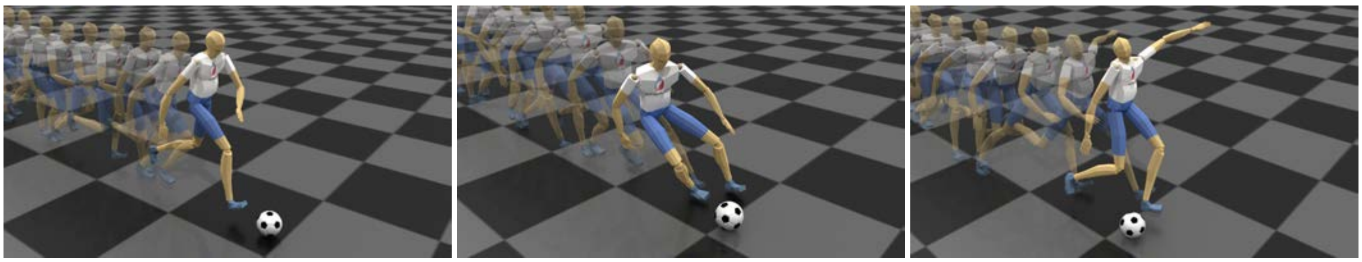

(a) Dribbling forward (b) Dribbling to the side (c) Shooting

Fig. 1. Soccer motions

Playing with a soccer ball is not easy even for a real human because of dynamic foot contacts with the moving ball while chasing and controlling it. The problem of online full-body soccer motion synthesis is challenging and has not been fully solved yet. In this paper, we present a novel motion control system that produces physically-correct full-body soccer motions: dribbling forward, dribbling to the side, and shooting, in response to an online user motion prescription specified by a motion type, a running speed, and a turning angle. This system performs two tightly-coupled tasks: data-driven motion prediction and physics-based motion synthesis. Given example motion data, the former synthesizes a reference motion in accordance with an online user input and further refines the motion to make the character kick the ball at a right time and place. Provided with the reference motion, the latter then adopts a Model Predictive Control (MPC) framework to generate a physically-correct soccer motion, by solving an optimal control problem that is formulated based on dynamics for a full-body character and the moving ball together with their interactions. Our demonstration shows the effectiveness of the proposed system that synthesizes convincing full-body soccer motions in various scenarios such as adjusting the desired running speed of the character, changing the velocity and the mass of the ball, and maintaining balance against external forces.

CCS Concepts: • **Computing methodologies → Animation**; **Physical simulation**.

Additional Key Words and Phrases: 3D character animation, physics-based simulation, motion control

ACM Reference Format:
Seokpyo Hong, Daseong Han, Kyungmin Cho, Joseph S. Shin (formerly Sung Yong Shin), and Junyong Noh. 2019. Physics-based Full-body Soccer Motion Control for Dribbling and Shooting. *ACM Trans. Graph.* 38, 4, Article 74 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3322963

1 INTRODUCTION

Playing with a soccer ball is not easy even for a real human. To dribble a soccer ball, for example, the player runs after the ball, following a time-varying trajectory while occasionally touching it with a foot in the desired direction and speed, which involves highly complex, subtle dynamics for full-body human motions and their interactions with the ball. Soccer motions have been studied extensively in robotics and character animation [Barrett et al. 2010; Choi et al. 2015, 2016; Hester et al. 2010; Jain and Liu 2009; Leottau et al. 2014; Sung et al. 2013; Yi et al. 2013]. However, fundamental issues still remain unsolved in the synthesis of even basic soccer motions such as dribbling and shooting: how to adjust the footsteps of the character to make it kick the ball with a foot at a right time and place, how to achieve physics-based online motion control over the character that interacts with a ball, and how to make the character balance against unexpected external forces while controlling the ball. These issues have not been settled yet although there are partial solutions for simple cases: kicking a ball in a stationary stance [Ha

*Corresponding author

Authors' addresses: Seokpyo Hong, Visual Media Lab, KAIST, sphorpin89@gmail.com; Daseong Han, Handong Global University, dshan@handong.edu; Kyungmin Cho, Visual Media Lab, KAIST, cluemaker@kaist.ac.kr; Joseph S. Shin (formerly Sung Yong Shin), Handong Global University and KAIST, josephsyshin@gmail.com; Junyong Noh, Visual Media Lab, KAIST, junyongnoh@kaist.ac.kr.

0730-0301/2019/7-ART74 $15.00
https://doi.org/10.1145/3306346.3322963

Learning Character-Agnostic Motion for Motion Retargeting in 2D

KFIR ABERMAN, Tel-Aviv University, AICFVE Beijing Film Academy
RUNDI WU, Peking University
DANI LISCHINSKI, Shandong University, Hebrew University of Jerusalem
BAOQUAN CHEN*, Peking University
DANIEL COHEN-OR, Tel-Aviv University

Analyzing human motion is a challenging task with a wide variety of applications in computer vision and in graphics. One such application, of particular importance in computer animation, is the retargeting of motion from one performer to another. While humans move in three dimensions, the vast majority of human motions are captured using video, requiring 2D-to-3D pose and camera recovery, before existing retargeting approaches may be applied. In this paper, we present a new method for retargeting video-captured motion between different human performers, without the need to explicitly reconstruct 3D poses and/or camera parameters.

In order to achieve our goal, we learn to extract, directly from a video, a high-level latent motion representation, which is invariant to the skeleton geometry and the camera view. Our key idea is to train a deep neural network to decompose temporal sequences of 2D poses into three components: motion, skeleton, and camera view-angle. Having extracted such a representation, we are able to re-combine motion with novel skeletons and camera views, and decode a retargeted temporal sequence, which we compare to a ground truth from a synthetic dataset.

We demonstrate that our framework can be used to robustly extract human motion from videos, bypassing 3D reconstruction, and outperforming existing retargeting methods, when applied to videos in-the-wild. It also enables additional applications, such as performance cloning, video-driven cartoons, and motion retrieval.

CCS Concepts: • **Computing methodologies** → **Motion processing**; **Neural networks**.

Additional Key Words and Phrases: Motion retargeting, autoencoder, motion analysis

ACM Reference Format:
Kfir Aberman, Rundi Wu, Dani Lischinski, Baoquan Chen, and Daniel Cohen-Or. 2019. Learning Character-Agnostic Motion for Motion Retargeting in 2D. *ACM Trans. Graph.* 38, 4, Article 75 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3322999

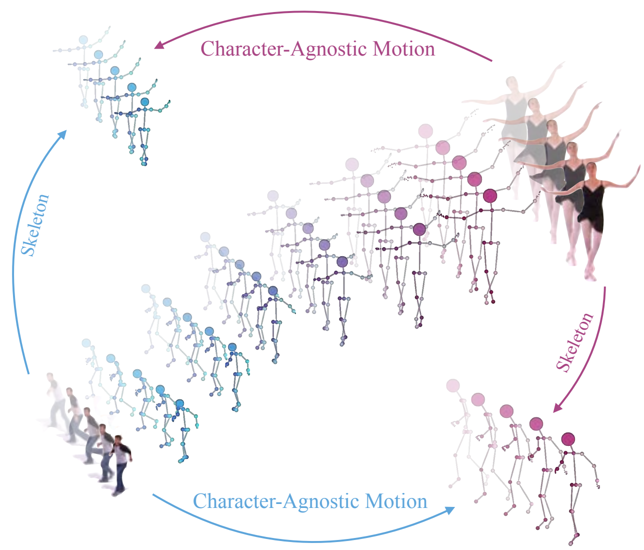

Fig. 1. Given two videos of different performers, our approach enables to extract character-agnostic motion from each video, and transfer it to a new skeleton and view angle (top-left and bottom-right), directly in 2D. In addition, separate latent representations for motion, skeleton, and view-angle are extracted, enabling control and interpolation of these parameters.

1 INTRODUCTION

Understanding and synthesizing human motion has been a central research topic in computer animation. Motion is inherently a 4D entity, commonly represented using a low-level encoding: as a temporal sequence of poses, specified as a set of joint positions and/or angles. Such a representation strongly depends on the skeleton and its geometric properties, such as the lengths of the limbs and their proportions. Thus, the same motion performed by two individuals with different skeletons might have significantly different representations. One might even argue that character-agnostic motion is a slippery and elusive notion, which is not completely well-defined.

In this work, we address the challenging problem of retargeting the video-captured motion of one human performer to another. In a nutshell, our approach is to extract an abstract, character- and camera-agnostic, latent representation of human motion directly from ordinary video. The extracted motion may then be applied to other, possibly very different, skeletons, and/or shown from new viewpoints.

The challenges that we face are twofold: First, the abstract motion representation that we seek is new and unknown, and thus we do not have the benefit of supervision. Second, working on video introduces an additional obstacle, as the joint trajectories are observed in 2D,

*Corresponding author

Authors' addresses: Kfir Aberman, Tel-Aviv University, AICFVE Beijing Film Academy; Rundi Wu, Peking University; Dani Lischinski, Shandong University, Hebrew University of Jerusalem; Baoquan Chen, Peking University; Daniel Cohen-Or, Tel-Aviv University.

0730-0301/2019/7-ART75 $15.00
https://doi.org/10.1145/3306346.3322999

Deep View Synthesis from Sparse Photometric Images

ZEXIANG XU, University of California, San Diego
SAI BI, University of California, San Diego
KALYAN SUNKAVALLI, Adobe Research
SUNIL HADAP*, Lab126, Amazon
HAO SU, University of California, San Diego
RAVI RAMAMOORTHI, University of California, San Diego

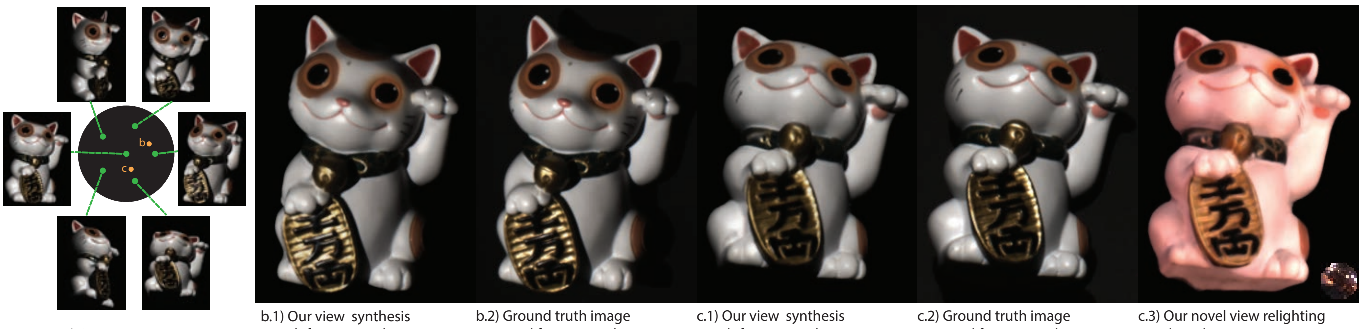

Fig. 1. We present a method to synthesize scene appearance from a novel view by interpolating only six wide-baseline images (a). We do this by using a structured setup to capture photometric images under directional lighting and interpolating them using a novel deep neural network. Our method can reproduce complex appearance effects like specularities, shadows, and occlusions (b.1,c.1) resulting in images that are close to ground truth captured images (b.2,c.2). These results can be combined with image-based relighting methods to visualize the scene under novel view and lighting (c.3).

The goal of light transport acquisition is to take images from a sparse set of lighting and viewing directions, and combine them to enable arbitrary relighting with changing view. While relighting from sparse images has received significant attention, there has been relatively less progress on view synthesis from a sparse set of "photometric" images—images captured under controlled conditions, lit by a single directional source; we use a spherical gantry to position the camera on a sphere surrounding the object. In this paper, we synthesize novel viewpoints across a wide range of viewing directions (covering a 60° cone) from a sparse set of just six viewing directions. While our approach relates to previous view synthesis and image-based rendering techniques, those methods are usually restricted to much smaller baselines, and are captured under environment illumination. At our baselines, input images have few correspondences and large occlusions; however we benefit from structured photometric images. Our method is based on a deep convolutional network trained to directly synthesize new views from the six input views. This network combines 3D convolutions on a plane sweep volume with a novel per-view per-depth plane attention map prediction network to effectively aggregate multi-view appearance. We train our network with a large-scale synthetic dataset of 1000 scenes with complex geometry and material properties. In practice, it is able to synthesize novel viewpoints for captured real data and reproduces complex appearance effects like occlusions, view-dependent specularities and hard shadows. Moreover, the method can also be combined with previous relighting techniques to enable changing both lighting and view, and applied to computer vision problems like multiview stereo from sparse image sets.

CCS Concepts: • **Computing methodologies → Image-based rendering**.

Additional Key Words and Phrases: appearance acquisition, novel view synthesis

ACM Reference Format:
Zexiang Xu, Sai Bi, Kalyan Sunkavalli, Sunil Hadap, Hao Su, and Ravi Ramamoorthi. 2019. Deep View Synthesis from Sparse Photometric Images. *ACM Trans. Graph.* 38, 4, Article 76 (July 2019), 13 pages. https://doi.org/10.1145/3306346.3323007

*This work was done prior to joining Amazon.

Authors' addresses: Zexiang Xu, University of California, San Diego, zexiangxu@cs.ucsd.edu; Sai Bi, University of California, San Diego, bisai@cs.ucsd.edu; Kalyan Sunkavalli, Adobe Research, sunkaval@adobe.com; Sunil Hadap, Lab126, Amazon, sunilhadap@acm.org; Hao Su, University of California, San Diego, haosu@eng.ucsd.edu; Ravi Ramamoorthi, University of California, San Diego, ravir@cs.ucsd.edu.

0730-0301/2019/7-ART76 $15.00
https://doi.org/10.1145/3306346.3323007

1 INTRODUCTION

A central problem in computer graphics and vision is to acquire images of a scene and reproduce its appearance under arbitrary lighting and viewpoint. This has traditionally been accomplished by densely sampling the scene's "reflectance field" [Debevec et al. 2000] and interpolating these images using a combination of image-based rendering and relighting methods. Recent work has demonstrated image-based relighting from sparse "photometric" images captured

Deep Reflectance Fields

High-Quality Facial Reflectance Field Inference from Color Gradient Illumination

ABHIMITRA MEKA, MPI Informatics, Saarland Informatics Campus, Google
CHRISTIAN HÄNE and ROHIT PANDEY, Google
MICHAEL ZOLLHÖFER, Stanford University
SEAN FANELLO, GRAHAM FYFFE, ADARSH KOWDLE, XUEMING YU, JAY BUSCH, JASON DOURGARIAN, PETER DENNY, SOFIEN BOUAZIZ, PETER LINCOLN, MATT WHALEN, GEOFF HARVEY, JONATHAN TAYLOR, SHAHRAM IZADI, ANDREA TAGLIASACCHI, and PAUL DEBEVEC, Google
CHRISTIAN THEOBALT, MPI Informatics, Saarland Informatics Campus
JULIEN VALENTIN and CHRISTOPH RHEMANN, Google

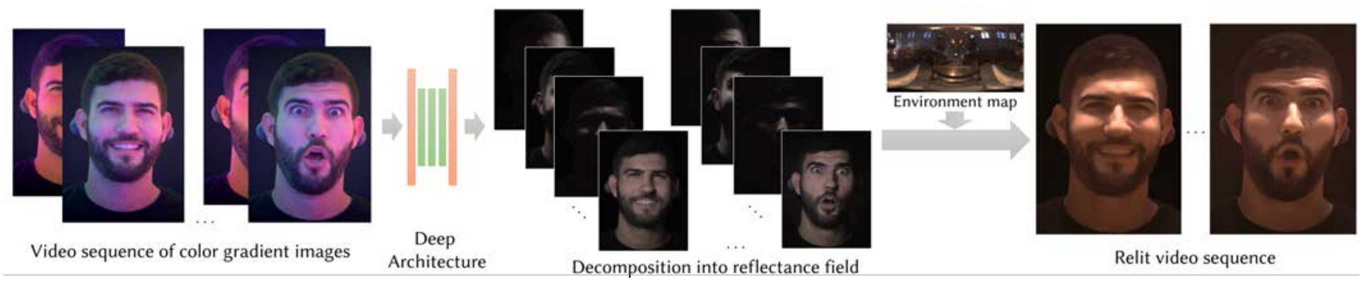

Fig. 1. Deep Reflectance Fields – given only *two* observations (color gradient images) of an actor, our method is able to relight the subject under any lighting condition. Our approach generalizes to unseen subjects, viewpoints, illumination conditions and can handle dynamic performances.

We present a novel technique to relight images of human faces by learning a model of facial reflectance from a database of 4D reflectance field data of several subjects in a variety of expressions and viewpoints. Using our learned model, a face can be relit in arbitrary illumination environments using only two original images recorded under spherical color gradient illumination. The output of our deep network indicates that the color gradient images contain the information needed to estimate the full 4D reflectance field, including specular reflections and high frequency details. While capturing spherical color gradient illumination still requires a special lighting setup, reduction to just two illumination conditions allows the technique to be applied to dynamic facial performance capture. We show side-by-side comparisons which demonstrate that the proposed system outperforms the state-of-the-art techniques in both realism and speed.

Authors' addresses: Abhimitra Meka, MPI Informatics, Saarland Informatics Campus, Google, ameka@mpi-inf.mpg.de; Christian Häne; Rohit Pandey, Google; Michael Zollhöfer, Stanford University; Sean Fanello; Graham Fyffe; Adarsh Kowdle; Xueming Yu; Jay Busch; Jason Dourgarian; Peter Denny; Sofien Bouaziz; Peter Lincoln; Matt Whalen; Geoff Harvey; Jonathan Taylor; Shahram Izadi; Andrea Tagliasacchi; Paul Debevec, Google; Christian Theobalt, MPI Informatics, Saarland Informatics Campus; Julien Valentin; Christoph Rhemann, crhemann@google.com, Google.

0730-0301/2019/7-ART77 $15.00
https://doi.org/10.1145/3306346.3323027

CCS Concepts: • **Computing methodologies → Computer vision**; **Machine learning**; *Rendering*.

Additional Key Words and Phrases: Reflectance estimation, machine learning

ACM Reference Format:
Abhimitra Meka, Christian Häne, Rohit Pandey, Michael Zollhöfer, Sean Fanello, Graham Fyffe, Adarsh Kowdle, Xueming Yu, Jay Busch, Jason Dourgarian, Peter Denny, Sofien Bouaziz, Peter Lincoln, Matt Whalen, Geoff Harvey, Jonathan Taylor, Shahram Izadi, Andrea Tagliasacchi, Paul Debevec, Christian Theobalt, Julien Valentin, and Christoph Rhemann. 2019. Deep Reflectance Fields: High-Quality Facial Reflectance Field Inference from Color Gradient Illumination. *ACM Trans. Graph.* 38, 4, Article 77 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3323027

1 INTRODUCTION

Modifying the lighting in a facial portrait image is a much sought after capability that would benefit many visual effects including portrait photography, and virtual or augmented reality applications. This relighting is particularly challenging since the facial appearance is the result of a complex interaction of light with the many materials that make up the skin, eyes, hair, teeth, and clothing, each of which have complex geometry and varying amounts of specular reflection and subsurface scattering. Further, ignoring or approximating these properties is especially perilous as humans are highly capable of detecting the subtle cues of realism in facial renderings. While today's computer graphics techniques can produce photo-realistic digital human models which can be rendered in any lighting and from any viewpoint, creating such models is still extremely laborious

Multi-view Relighting using a Geometry-Aware Network

JULIEN PHILIP, Université Côte d'Azur and Inria
MICHAËL GHARBI, Adobe
TINGHUI ZHOU, UC Berkeley
ALEXEI A. EFROS, UC Berkeley
GEORGE DRETTAKIS, Université Côte d'Azur and Inria

(a) our algorithm can relight a single-illumination drone video dynamically to synthesize a "time-lapse" effect

(b) single-view input

(c) three relit outputs: here we built the *proxy* geometry using internet photos of the same location

Fig. 1. Two applications of our multi-view relighting system. (a) We show five different frames from a drone video (copyright Namyeska youtu.be/JHeDP7_YBos used with permission) relit with a "time-lapse" effect of a rotating sun (see supplemental for the full video). A user can also relight a photograph of a known landmark (b) to different target lighting conditions (c). For this, we applied our algorithm to a collection of 50 internet images of the same location .

We propose the first learning-based algorithm that can relight images in a plausible and controllable manner given multiple views of an outdoor scene. In particular, we introduce a *geometry-aware* neural network that utilizes multiple geometry cues (normal maps, specular direction, etc.) and source and target shadow masks computed from a noisy *proxy geometry* obtained by multi-view stereo. Our model is a three-stage pipeline: two sub-networks refine the source and target shadow masks, and a third performs the final relighting. Furthermore, we introduce a novel representation for the shadow masks, which we call *RGB shadow images*. They reproject the colors from all views into the shadowed pixels and enable our network to cope with inacuraccies in the proxy and the non-locality of the shadow casting interactions. Acquiring large-scale multi-view relighting datasets for real scenes is challenging, so we train our network on photorealistic synthetic data. At train time, we also compute a noisy stereo-based geometric proxy, this time from the synthetic renderings. This allows us to bridge the gap between the real and synthetic domains. Our model generalizes well to real scenes. It can alter the illumination of drone footage, image-based renderings, textured mesh reconstructions, and even internet photo collections.

Authors' addresses: Julien Philip, Université Côte d'Azur and Inria, julien.philip@inria.fr; Michaël Gharbi, Adobe, mgharbi@adobe.com; Tinghui Zhou, UC Berkeley, tinghuiz@eecs.berkeley.edu; Alexei A. Efros, UC Berkeley, efros@eecs.berkeley.edu; George Drettakis, Université Côte d'Azur and Inria, George.Drettakis@inria.fr.

0730-0301/2019/7-ART78 $15.00
https://doi.org/10.1145/3306346.3323013

CCS Concepts: • **Computing methodologies → Image manipulation**.

Additional Key Words and Phrases: Image relighting, Multi-view, Deep Learning

ACM Reference Format:
Julien Philip, Michaël Gharbi, Tinghui Zhou, Alexei A. Efros, and George Drettakis. 2019. Multi-view Relighting using a Geometry-Aware Network. *ACM Trans. Graph.* 38, 4, Article 78 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3323013

1 INTRODUCTION

Changing the illumination of an outdoor image is a notoriously difficult problem that requires the lighting to be modified consistently across the image, and shadows to be removed and resynthesized for the new sun position [Duchêne et al. 2015; Tchou et al. 2004; Yu et al. 1999]. Cast shadows are particularly challenging because an occluder can be arbitrarily far from the point it shadows, or even out of view.

The basic premise of our approach is to use multi-view information and approximate 3D geometry to reason about non-local lighting interactions and guide the relighting task. We introduce the first learning-based algorithm that can relight multi-view datasets of outdoor scenes (Fig. 1), which have become a commodity thanks to smartphone cameras, large-scale internet photo collections and drone cinematography. Our model uses a neural network designed to exploit geometric cues. It includes a careful treatment of cast shadows and is trained solely on realistic synthetic renderings.

Single Image Portrait Relighting

TIANCHENG SUN, University of California, San Diego
JONATHAN T. BARRON and YUN-TA TSAI, Google Research
ZEXIANG XU, University of California, San Diego
XUEMING YU, GRAHAM FYFFE, CHRISTOPH RHEMANN, JAY BUSCH, and PAUL DEBEVEC, Google
RAVI RAMAMOORTHI, University of California, San Diego

(a) Input image and estimated lighting (b) Rendered images from our method under three novel illuminations

Fig. 1. Given only a single input image taken with a standard cellphone camera of a portrait (a), our model is able to quickly (160 ms.) generate new images of our human subject as though they are illuminated under new, previously-unseen lighting environments (b).

Lighting plays a central role in conveying the essence and depth of the subject in a portrait photograph. Professional photographers will carefully control the lighting in their studio to manipulate the appearance of their subject, while consumer photographers are usually constrained to the illumination of their environment. Though prior works have explored techniques for relighting an image, their utility is usually limited due to requirements of specialized hardware, multiple images of the subject under controlled or known illuminations, or accurate models of geometry and reflectance. To this end, we present a system for *portrait relighting*: a neural network that takes as input a *single* RGB image of a portrait taken with a standard cellphone camera in an unconstrained environment, and from that image produces a relit image of that subject as though it were illuminated according to any provided environment map. Our method is trained on a small database of 18 individuals captured under different directional light sources in a controlled light stage setup consisting of a densely sampled sphere of lights. Our proposed technique produces quantitatively superior results on our dataset's validation set compared to prior works, and produces convincing qualitative relighting results on a dataset of hundreds of real-world cellphone portraits. Because our technique can produce a 640×640 image in only 160 milliseconds, it may enable interactive user-facing photographic applications in the future.

CCS Concepts: • **Computing methodologies → Image-based rendering**; **Computational photography**; **Neural networks**.

Additional Key Words and Phrases: Portrait relighting, Image-based relighting, Light estimation.

ACM Reference Format:
Tiancheng Sun, Jonathan T. Barron, Yun-Ta Tsai, Zexiang Xu, Xueming Yu, Graham Fyffe, Christoph Rhemann, Jay Busch, Paul Debevec, and Ravi Ramamoorthi. 2019. Single Image Portrait Relighting. *ACM Trans. Graph.* 38, 4, Article 79 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3323008

1 INTRODUCTION

The rise of mobile computing has led to tremendous growth in the popularity of consumer digital photography, and one of the most popular and ubiquitous kinds of photos taken is the portrait: an image of a human subject's face or upper body. Portrait photography follows in the tradition of portrait painting, where since the renaissance artists have recognized how lighting can capture the depth and essence of the subject on a 2D canvas [Schütze 2015]. These ideas largely influenced professional portrait lighting techniques [Schriever 1909]. In the 18th and 19th centuries, portraits were often taken by professional photographers, who carefully considered and controlled the lighting of the scene in addition to the pose and appearance of their subject. In the modern age of "selfies" and candid photography, it is difficult or impossible for consumer photographers to control the lighting of their subject — one would

Authors' addresses: Tiancheng Sun, tis037@cs.ucsd.edu, University of California, San Diego; Jonathan T. Barron, barron@google.com; Yun-Ta Tsai, yuntatsai@google.com, Google Research; Zexiang Xu, zexiangxu@cs.ucsd.edu, University of California, San Diego; Xueming Yu, xuemingyu@google.com; Graham Fyffe, fyffe@google.com; Christoph Rhemann, crhemann@google.com; Jay Busch, jbusch@google.com; Paul Debevec, debevec@google.com, Google; Ravi Ramamoorthi, ravir@cs.ucsd.edu, University of California, San Diego.

0730-0301/2019/7-ART79 $15.00
https://doi.org/10.1145/3306346.3323008

Surface2Volume: Surface Segmentation Conforming Assemblable Volumetric Partition

CHRYSTIANO ARAÚJO*, University of British Columbia
DANIELA CABIDDU*, CNR IMATI
MARCO ATTENE, CNR IMATI
MARCO LIVESU, CNR IMATI
NICHOLAS VINING, University of British Columbia
ALLA SHEFFER, University of British Columbia

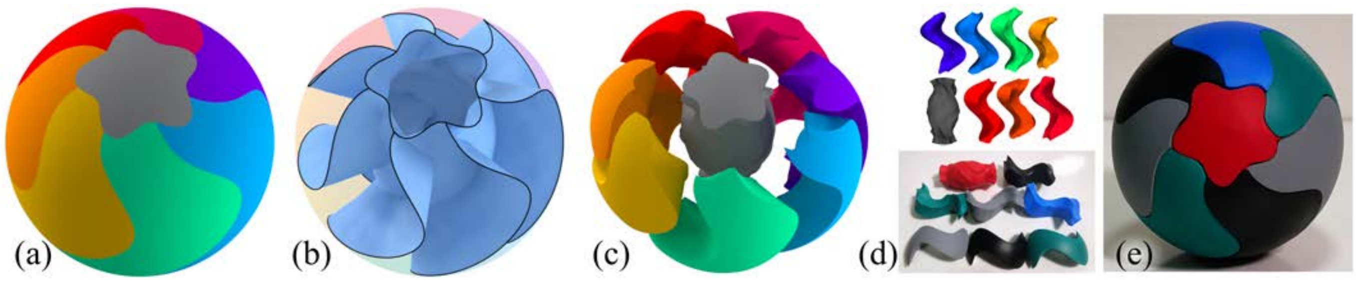

Fig. 1. Surface2Volume (left to right): (a) Input multi-color surface; (b) inner interfaces generated by Surface2Volume; (c) assemblable parts conforming to the surface-segmentation; (d) virtual (top) and fabricated (bottom) single-color parts; and (e) assembled target object.

Users frequently seek to fabricate objects whose outer surfaces consist of regions with different surface attributes, such as color or material. Manufacturing such objects in a single piece is often challenging or even impossible. The alternative is to partition them into single-attribute volumetric parts that can be fabricated separately and then assembled to form the target object. Facilitating this approach requires partitioning the input model into parts that *conform* to the surface segmentation and that can be moved apart with no collisions. We propose *Surface2Volume*, a partition algorithm capable of producing such *assemblable* parts, each of which is affiliated with a single attribute, the outer surface of whose assembly conforms to the input surface geometry and segmentation. In computing the partition we strictly enforce conformity with surface segmentation and assemblability, and optimize for ease of fabrication by minimizing part count, promoting part simplicity, and simplifying assembly sequencing. We note that computing the desired partition requires solving for three types of variables: per-part assembly trajectories, partition topology, i.e. the connectivity of the interface surfaces separating the different parts, and the geometry, or location, of these interfaces. We efficiently produce the desired partitions by addressing one type of variables at a time: first computing the assembly trajectories, then determining interface topology, and finally computing interface locations that allow parts assemblability. We algorithmically identify inputs that necessitate sequential assembly, and partition these inputs gradually by computing and disassembling a subset of assemblable parts at a time. We demonstrate our method's robustness and versatility by employing it to partition a range of models with complex surface segmentations into assemblable parts. We further validate our framework via output fabrication and comparisons to alternative partition techniques.

CCS Concepts: • **Computing methodologies → Mesh geometry models**; *Volumetric models.*

Additional Key Words and Phrases: shape decomposition, fabrication, assemblability

ACM Reference Format:
Chrystiano Araújo, Daniela Cabiddu, Marco Attene, Marco Livesu, Nicholas Vining, and Alla Sheffer. 2019. Surface2Volume: Surface Segmentation Conforming Assemblable Volumetric Partition. *ACM Trans. Graph.* 38, 4, Article 80 (July 2019), 16 pages. https://doi.org/10.1145/3306346.3323004

*Joint first authors

Authors' addresses: Chrystiano Araújo, University of British Columbia; Daniela Cabiddu, CNR IMATI; Marco Attene, CNR IMATI; Marco Livesu, CNR IMATI; Nicholas Vining, University of British Columbia; Alla Sheffer, University of British Columbia.

0730-0301/2019/7-ART80 $15.00
https://doi.org/10.1145/3306346.3323004

1 INTRODUCTION

Digital fabrication algorithms are successfully used to create real-life replicas of virtual models with uniform color and material. However, users often wish to create objects with non-uniform visible surface attributes, such as shapes whose outer surface consists of regions with different color, opacity, or texture (Figures 1, 2). Manufacturing such objects as a single solid necessitates the use of multi-attribute, or multi-material, fabrication methodologies, or after-the-fact surface painting. These approaches exhibit numerous limitations (Section 2). An appealing alternative is to decompose

CurviSlicer: Slightly curved slicing for 3-axis printers

JIMMY ETIENNE, Université de Lorraine, CNRS, Inria, LORIA
NICOLAS RAY, Université de Lorraine, CNRS, Inria, LORIA
DANIELE PANOZZO, New York University
SAMUEL HORNUS, Université de Lorraine, CNRS, Inria, LORIA
CHARLIE C. L. WANG, The Chinese University of Hong Kong
JONÀS MARTÍNEZ, Université de Lorraine, CNRS, Inria, LORIA
SARA MCMAINS, University of California, Berkeley
MARC ALEXA, Technische Universität Berlin
BRIAN WYVILL, University of Victoria
SYLVAIN LEFEBVRE, Université de Lorraine, CNRS, Inria, LORIA

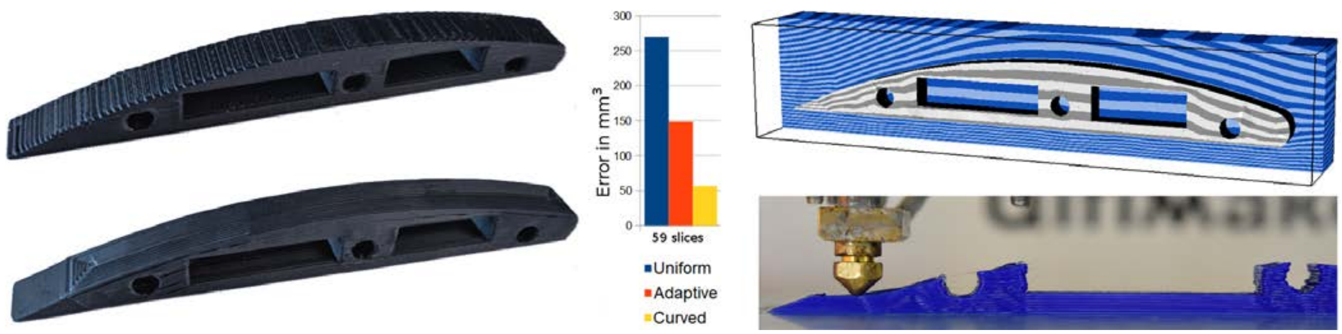

Fig. 1. Our technique curves deposition paths to improve parts printed with fused filament fabrication. Compared to state of the art adaptive slicing (top left) which is limited to planar layers, our print (bottom left) has a smooth surface finish while using the same number of layers (40). The reproduction accuracy is improved overall (middle graph), with a total volume error of 57mm^3 compared to the 149mm^3 of adaptive slicing. Our approach computes a continuous deformation of space (top right) under fabrication constraints (thicknesses, slope). The produced toolpaths are guaranteed to print without collisions on standard 3-axis 3D printers, here an Ultimaker2 (bottom right).

Most additive manufacturing processes fabricate objects by stacking planar layers of solidified material. As a result, produced parts exhibit a so-called staircase effect, which results from sampling slanted surfaces with parallel planes. Using thinner slices reduces this effect, but it always remains visible where layers *almost* align with the input surfaces.

In this research we exploit the ability of some additive manufacturing processes to deposit material slightly out of plane to dramatically reduce these artifacts. We focus in particular on the widespread Fused Filament Fabrication (FFF) technology, since most printers in this category can deposit along slightly curved paths, under deposition slope and thickness constraints.

Our algorithm curves the layers, making them either follow the natural slope of the input surface or on the contrary, make them intersect the surfaces at a steeper angle thereby improving the sampling quality. Rather than directly computing curved layers, our algorithm optimizes for a deformation of the model which is then sliced with a standard planar approach. We demonstrate that this approach enables us to encode all fabrication constraints, including the guarantee of generating collision-free toolpaths, in a convex optimization that can be solved using a QP solver.

We produce a variety of models and compare print quality between curved deposition and planar slicing.

CCS Concepts: • **Computing methodologies → Shape modeling**.

Additional Key Words and Phrases: curved slicing, additive manufacturing

ACM Reference Format:
Jimmy Etienne, Nicolas Ray, Daniele Panozzo, Samuel Hornus, Charlie C. L. Wang, Jonàs Martínez, Sara McMains, Marc Alexa, Brian Wyvill, and Sylvain Lefebvre. 2019. CurviSlicer: Slightly curved slicing for 3-axis printers. *ACM Trans. Graph.* 38, 4, Article 81 (July 2019), 11 pages. https://doi.org/10.1145/3306346.3323022

Authors' addresses: Jimmy Etienne, Université de Lorraine, CNRS, Inria, LORIA; Nicolas Ray, Université de Lorraine, CNRS, Inria, LORIA; Daniele Panozzo, New York University; Samuel Hornus, Université de Lorraine, CNRS, Inria, LORIA; Charlie C. L. Wang, The Chinese University of Hong Kong; Jonàs Martínez, Université de Lorraine, CNRS, Inria, LORIA; Sara McMains, University of California, Berkeley; Marc Alexa, Technische Universität Berlin; Brian Wyvill, University of Victoria; Sylvain Lefebvre, Université de Lorraine, CNRS, Inria, LORIA.

0730-0301/2019/7-ART81 $15.00
https://doi.org/10.1145/3306346.3323022

Star-Shaped Metrics for Mechanical Metamaterial Design

JONÀS MARTÍNEZ, Université de Lorraine, CNRS, Inria, LORIA
MÉLINA SKOURAS, Université Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK
CHRISTIAN SCHUMACHER, ETH Zürich
SAMUEL HORNUS, Université de Lorraine, CNRS, Inria, LORIA
SYLVAIN LEFEBVRE, Université de Lorraine, CNRS, Inria, LORIA
BERNHARD THOMASZEWSKI, Université de Montréal

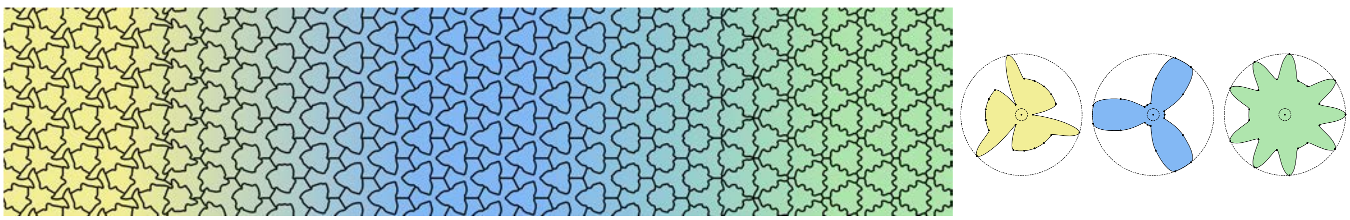

(a) (b)

Fig. 1. Our method generates a smoothly-graded pattern (a) when interpolating between three star-shaped distance functions (b) on a regular honeycomb lattice. Each distance function is compactly parameterized with polar coordinates, allowing for simple interpolation in metric space as indicated by color-coding.

We present a method for designing mechanical metamaterials based on the novel concept of Voronoi diagrams induced by star-shaped metrics. As one of its central advantages, our approach supports interpolation between arbitrary metrics. This capability opens up a rich space of structures with interesting aesthetics and a wide range of mechanical properties, including isotropic, tetragonal, orthotropic, as well as smoothly graded materials. We evaluate our method by creating large sets of example structures, provided as accompanying material. We validate the mechanical properties predicted by simulation through tensile tests on a set of physical prototypes.

CCS Concepts: • **Computing methodologies → Shape modeling**.

Additional Key Words and Phrases: mechanical metamaterial, voronoi diagram, star-shaped metrics

ACM Reference Format:
Jonàs Martínez, Mélina Skouras, Christian Schumacher, Samuel Hornus, Sylvain Lefebvre, and Bernhard Thomaszewski. 2019. Star-Shaped Metrics for Mechanical Metamaterial Design. *ACM Trans. Graph.* 38, 4, Article 82 (July 2019), 13 pages. https://doi.org/10.1145/3306346.3322989

Authors' addresses: Jonàs Martínez, Université de Lorraine, CNRS, Inria, LORIA, F-54000, Nancy, France; Mélina Skouras, Université Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000, Grenoble, France ; Christian Schumacher, ETH Zürich; Samuel Hornus, Université de Lorraine, CNRS, Inria, LORIA, F-54000, Nancy, France; Sylvain Lefebvre, Université de Lorraine, CNRS, Inria, LORIA, F-54000, Nancy, France; Bernhard Thomaszewski, Université de Montréal.

0730-0301/2019/7-ART82 $15.00
https://doi.org/10.1145/3306346.3322989

1 INTRODUCTION

Digital manufacturing technologies such as 3D printing and laser cutting enable us to fabricate designs with great geometric detail. One particular way of exploiting this capability is to create patterned sheet materials whose geometric structures can be tailored to control their macro-mechanical behavior.

A typical approach to model and analyze structured sheet materials is centered around the concept of a representative element—a tile—which is repeated, transformed and laid out such as to generate a regular spatial tiling. Changing the shape of the representative tile allows for controlling macro-mechanical properties such as isotropy or negative Poisson's ratios. Generalizing this material design principle from a single representative tile to *families* of tiles that can be combined in a spatially-varying manner opens the door to structures with progressively-graded material properties.

There are several challenges that arise when defining such families of tiles. First, the space of tilings that can be generated from the tile family should span a wide range of macro-mechanical properties. At the same time, for applications such as fashion design or architecture, the shape of the tiles and the resulting patterns should also be aesthetically pleasing. Second, the tiles must satisfy a number of geometric constraints to ensure that they are tileable (boundary compatibility) and fabricable (minimum feature thickness). Lastly, the space of tiles should support progressive, robust gradation between different material properties.

In this work, we propose a new approach for generating tilings that simultaneously satisfy the above desiderata. Rather than producing a specific discrete set of tiles, our method defines a continuous space of tile geometries. Each point in this space corresponds to a tile shape that can cover the plane without gaps or overlaps. Furthermore, tile geometries can be robustly interpolated along any

X-Shells: A new class of deployable beam structures

JULIAN PANETTA, EPFL
MINA KONAKOVIĆ-LUKOVIĆ, EPFL
FLORIN ISVORANU, EPFL
ETIENNE BOULEAU, Ingeni SA
MARK PAULY, EPFL

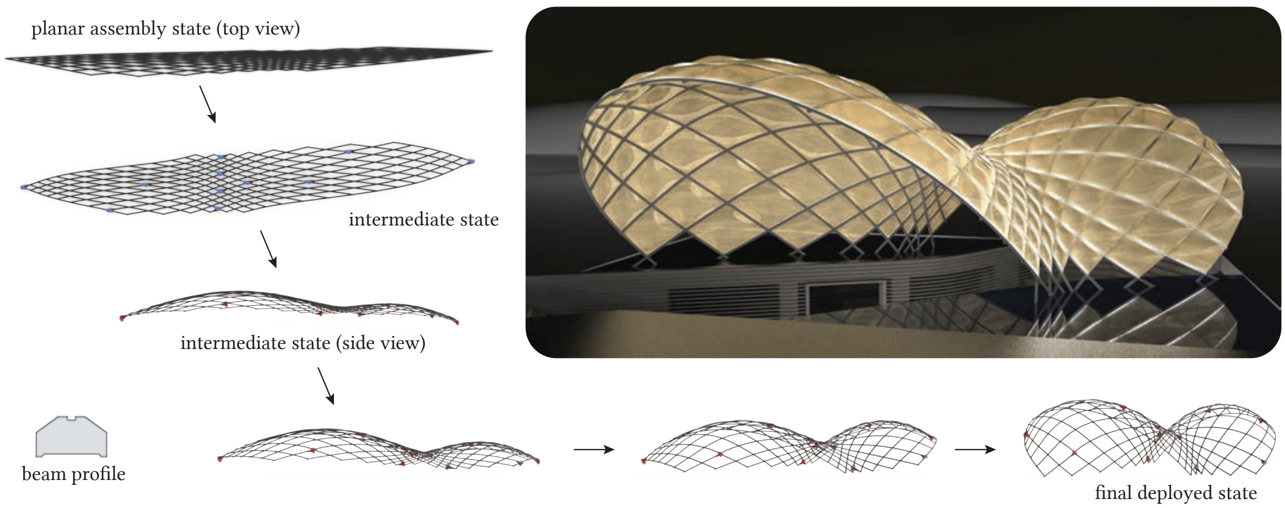

Fig. 1. An *X-shell* is a deformable mechanism that can be assembled from linear beam elements in a flat configuration and deployed to a desired 3D target form. Our algorithm computes the layout and parameters of the flexible beam network as well as a sparse pattern of actuation forces required to deploy the structure. The sequence illustrates the deployment process, where torque actuators at the joints are stylized in red and blue. The top right image shows a design study for a potential architectural application with additional cladding.

We present *X-shells*, a new class of deployable structures formed by an ensemble of elastically deforming beams coupled through rotational joints. An X-shell can be assembled conveniently in a flat configuration from standard elastic beam elements and then deployed through force actuation into the desired 3D target state. During deployment, the coupling imposed by the joints will force the beams to twist and buckle out of plane to maintain a state of static equilibrium. This complex interaction of discrete joints and continuously deforming beams allows interesting 3D forms to emerge. Simulating X-shells is challenging, however, due to unstable equilibria at the onset of beam buckling. We propose an optimization-based simulation framework building on a discrete rod model that robustly handles such difficult scenarios by analyzing and appropriately modifying the elastic energy Hessian. This real-time simulation method forms the basis of a computational design tool for X-shells that enables interactive design space exploration by varying and optimizing design parameters to achieve a specific design intent. We jointly optimize the assembly state and the deployed configuration to ensure the geometric and structural integrity of the deployable X-shell. Once a design is finalized, we also optimize for a sparse distribution of actuation forces to efficiently deploy it from its flat assembly state to its 3D target state. We demonstrate the effectiveness of our design approach with a number of design studies that highlight the richness of the X-shell design space, enabling new forms not possible with existing approaches. We validate our computational model with several physical prototypes that show excellent agreement with the optimized digital models.

CCS Concepts: • **Computing methodologies → Shape Modeling**; **Simulation**.

Additional Key Words and Phrases: grid shells, deployable structures, physics-based simulation, numerical optimization, computational design

ACM Reference Format:
Julian Panetta, Mina Konaković-Luković, Florin Isvoranu, Etienne Bouleau, and Mark Pauly. 2019. X-Shells: A new class of deployable beam structures. *ACM Trans. Graph.* 38, 4, Article 83 (July 2019), 15 pages. https://doi.org/10.1145/3306346.3323040

Authors' addresses: J. Panetta, EPFL, julian.panetta@epfl.ch; M. Konaković- Luković, EPFL, mina.konakovic@epfl.ch; F. Isvoranu, EPFL, florin.isvoranu@epfl. ch; E. Bouleau, Ingeni SA, etienne.bouleau@ingeni.ch; M. Pauly, EPFL, mark. pauly@epfl.ch.

0730-0301/2019/7-ART83 $15.00
https://doi.org/10.1145/3306346.3323040

Multi-Robot Collaborative Dense Scene Reconstruction

SIYAN DONG, Shandong University and AICFVE Beijing Film Academy
KAI XU*, National University of Defense Technology and AICFVE Beijing Film Academy
QIANG ZHOU, Shandong University and AICFVE Beijing Film Academy
ANDREA TAGLIASACCHI, Google Inc., University of Waterloo, and University of Victoria
SHIQING XIN, Shandong University
MATTHIAS NIESSNER, Technical University of Munich
BAOQUAN CHEN*, Peking University

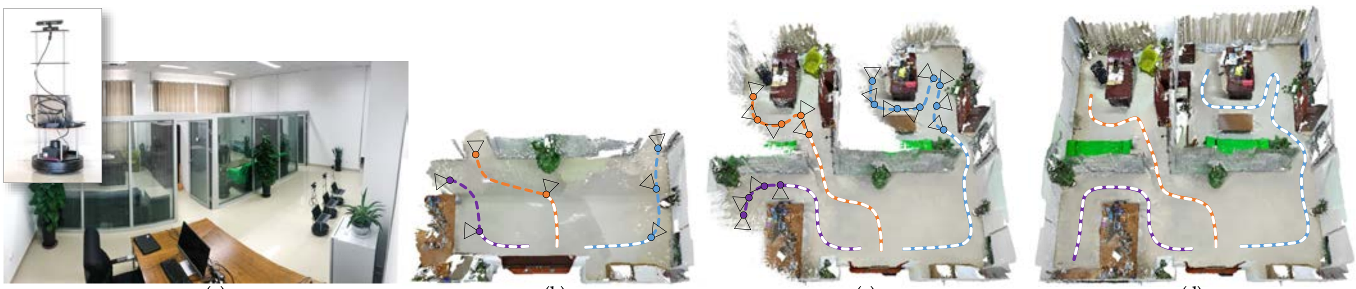

Fig. 1. We introduce an algorithm for multi-robot collaborative dense reconstruction of unknown indoor scenes (a). Given the partially scanned scene, we extract a set of task view points based on the uncertainty in the current reconstruction and assign them to the robots (b-c). The task assignment is formulated as an Optimal Mass Transport (OMT). For each robot, a smooth movement path is planned based on the tasks assigned to it. Our algorithm enables the robots to efficiently coordinate with each other, evenly distribute their scanning effort, and efficiently achieve a full coverage and high-quality reconstruction (d).

We present an autonomous scanning approach which allows multiple robots to perform collaborative scanning for dense 3D reconstruction of unknown indoor scenes. Our method plans scanning paths for several robots, allowing them to efficiently coordinate with each other such that the collective scanning coverage and reconstruction quality is maximized while the overall scanning effort is minimized. To this end, we define the problem as a dynamic task assignment and introduce a novel formulation based on Optimal Mass Transport (OMT). Given the currently scanned scene, a set of task views are extracted to cover scene regions which are either unknown or uncertain. These task views are assigned to the robots based on the OMT optimization. We then compute for each robot a smooth path over its assigned tasks by solving an approximate traveling salesman problem. In order to showcase our algorithm, we implement a multi-robot auto-scanning system. Since our method is computationally efficient, we can easily run it in real time on commodity hardware, and combine it with online RGB-D reconstruction approaches. In our results, we show several real-world examples of large indoor environments; in addition, we build a benchmark with a series of carefully designed metrics for quantitatively evaluating multi-robot autoscanning. Overall, we are able to demonstrate high-quality scanning results with respect to reconstruction quality and scanning efficiency, which significantly outperforms existing multi-robot exploration systems.

*Corresponding authors: Kai Xu (kevin.kai.xu@gmail.com) and Baoquan Chen (baoquan.chen@gmail.com).

Authors' addresses: Siyan Dong, Shandong University, AICFVE Beijing Film Academy; Kai Xu, National University of Defense Technology, AICFVE Beijing Film Academy; Qiang Zhou, Shandong University, AICFVE Beijing Film Academy; Andrea Tagliasacchi, Google Inc. University of Waterloo, University of Victoria; Shiqing Xin, Shandong University; Matthias Nießner, Technical University of Munich; Baoquan Chen, Peking University.

0730-0301/2019/7-ART84 $15.00
https://doi.org/10.1145/3306346.3322942

CCS Concepts: • **Computing methodologies** → *Shape analysis.*

Additional Key Words and Phrases: Autonomous scene reconstruction, multi-robot, collaborative reconstruction, Optimal Mass Transport

ACM Reference Format:
Siyan Dong, Kai Xu, Qiang Zhou, Andrea Tagliasacchi, Shiqing Xin, Matthias Nießner, and Baoquan Chen. 2019. Multi-Robot Collaborative Dense Scene Reconstruction. *ACM Trans. Graph.* 38, 4, Article 84 (July 2019), 16 pages. https://doi.org/10.1145/3306346.3322942

1 INTRODUCTION

Reconstructing and mapping indoor environments is critical to a large variety of applications, ranging from 3D content creation for augmented and virtual reality to localization for domestic robot navigation. On the hardware side, we have witnessed the emergence and proliferation of commodity range sensors (e.g., Microsoft Kinect, Intel RealSense, etc.) that capture depth data in real-time. On the software side, researchers have made incredible progress developing online RGB-D reconstruction methods [Izadi et al. 2011; Newcombe et al. 2011] that are able to reconstruct large environments [Chen et al. 2013; Nießner et al. 2013] along with robust camera tracking [Dai et al. 2017; Whelan et al. 2015; Zhang et al. 2014].

A Symmetric Objective Function for ICP

SZYMON RUSINKIEWICZ, Princeton University, USA

The Iterative Closest Point (ICP) algorithm, commonly used for alignment of 3D models, has previously been defined using either a point-to-point or point-to-plane objective. Alternatively, researchers have proposed computationally-expensive methods that directly minimize the distance function between surfaces. We introduce a new symmetrized objective function that achieves the simplicity and computational efficiency of point-to-plane optimization, while yielding improved convergence speed and a wider convergence basin. In addition, we present a linearization of the objective that is exact in the case of exact correspondences. We experimentally demonstrate the improved speed and convergence basin of the symmetric objective, on both smooth models and challenging cases involving noise and partial overlap.

CCS Concepts: • **Computing methodologies → Shape analysis**.

Additional Key Words and Phrases: 3D Registration, 3D Alignment, Iterative Closest Point

ACM Reference Format:
Szymon Rusinkiewicz. 2019. A Symmetric Objective Function for ICP. *ACM Trans. Graph.* 38, 4, Article 85 (July 2019), 7 pages. https://doi.org/10.1145/3306346.3323037

1 INTRODUCTION

Registration of 3D shapes is a key step in both 3D model creation (from scanners or computer vision systems) and shape analysis. For rigid-body alignment based purely on geometry (as opposed to RGB-D), the most common methods are based on variants of the Iterative Closest Point (ICP) algorithm [Besl and McKay 1992]. In this method, points are repeatedly selected from one model, their nearest points on the other model (given the current best-estimate rigid-body alignment) are selected as correspondences, and an incremental transformation is found that minimizes distances between point pairs. The algorithm eventually converges to a local minimum of surface-to-surface distance.

Because ICP-like algorithms can be made efficient and reliable, they have become widely adopted. As a result, researchers have focused on both addressing the shortcomings of ICP and extending it to new settings such as color-based registration and non-rigid alignment. One particular class of improvements has focused on the loss function that is optimized to obtain an incremental transformation. For example, as compared to the original work of Besl and McKay, which minimized *point-to-point* distance, the method of Chen and Medioni [1992] minimized the distance between a point on one mesh and a plane containing the matching point and perpendicular to its normal. This *point-to-plane* objective generally results in faster convergence to the correct alignment and greater ultimate accuracy, though it does not necessarily increase the basin of convergence. Work by Fitzgibbon [2001], Mitra et al. [2004], and Pottmann et al. [2006] showed that both point-to-point and point-to-plane minimization may be thought of as approximations to minimizing the squared Euclidean distance function of the surface, and they presented algorithms that achieved greater convergence speed and stability, albeit at the cost of greater computational complexity and/or auxiliary data structures.

This paper proposes a symmetrized version of the point-to-plane objective for use in ICP, incorporating two key ideas. First, the plane in which the error is minimized is based on the surface normals of *both* points in the corresponding pair. Second, the optimization is performed in a "stationary" coordinate system, while both meshes are moved in opposite directions. These changes require a relatively small modification to the optimization problem being performed, and almost no increased computation per iteration, but result in improved convergence of ICP.

The reason for this improvement is that the symmetric objective is minimized whenever the pair of points lies on a second-order (constant-curvature) patch of surface, rather than being minimized only if the points are on a plane. Thus, we gain some of the same benefits as second-order distance function minimization methods, but without explicit computation of second-order surface properties, or the need for volumetric data structures to store an approximation to the squared Euclidean distance function.

In addition to the primary contribution of the new symmetric objective, we also introduce an alternative approach to linearization of rotations that allows us to reduce the optimization to a linear least-squares problem, while still solving for the exact transformation when correspondences are exact. We conduct experiments that demonstrate both greater per-iteration error reduction and an increase in the convergence basin for our proposed method.

2 RELATED WORK

Since the original ICP algorithms by Besl and McKay [1992] and Chen and Medioni [1992], there have been significant efforts to improve convergence and stability. For a comprehensive overview of many variants, see the surveys by Rusinkiewicz and Levoy [2001], Díez et al. [2015], and Pomerleau et al. [2015]. Much of this work focuses on finding better correspondences (e.g., by matching local surface properties or descriptors), performing outlier-tolerant optimization, or generalizing to non-rigid deformation. Here we focus specifically on methods that modify the objective function and/or the strategy for minimizing it.

Segal et al. [2009] generalize ICP to associate a probabilistic model (in practice, a covariance matrix) with each point. This al-

Author's address: Szymon Rusinkiewicz, Department of Computer Science, Princeton University, Princeton, NJ, USA.

0730-0301/2019/7-ART85 $15.00
https://doi.org/10.1145/3306346.3323037

Warp-and-Project Tomography for Rapidly Deforming Objects

GUANGMING ZANG, RAMZI IDOUGHI, RAN TAO, GILLES LUBINEAU, PETER WONKA, and WOLFGANG HEIDRICH, King Abdullah University of Science And Technology, Saudi Arabia

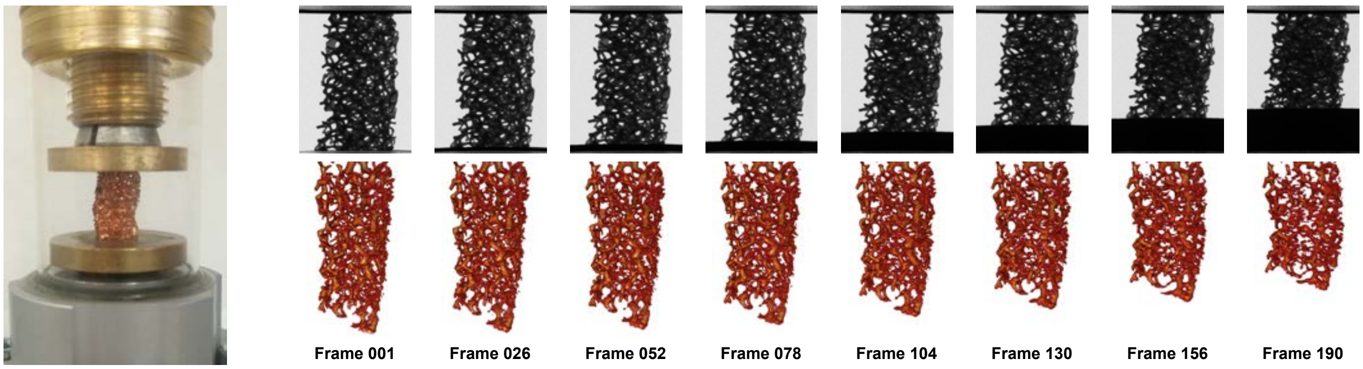

Fig. 1. We introduce a CT reconstruction method for objects that undergo rapid deformation during the scan. Shown here is a copper foam crumpling under a compressive force during the scan. The whole complex animation is reconstructed using only 192 projection images that all correspond to different deformation states of the foam.

Computed tomography has emerged as the method of choice for scanning complex shapes as well as interior structures of stationary objects. Recent progress has also allowed the use of CT for analyzing deforming objects and dynamic phenomena, although the deformations have been constrained to be either slow or periodic motions.

In this work we improve the tomographic reconstruction of time-varying geometries undergoing faster, non-periodic deformations. Our method uses a warp-and-project approach that allows us to introduce an essentially continuous time axis where consistency of the reconstructed shape with the projection images is enforced for the specific time and deformation state at which the image was captured. The method uses an efficient, time-adaptive solver that yields both the moving geometry as well as the deformation field.

We validate our method with extensive experiments using both synthetic and real data from a range of different application scenarios.

CCS Concepts: • **Computing methodologies** → *Computer graphics*; *3D imaging*; *Motion capture*.

Additional Key Words and Phrases: X-ray computed tomography, 4D reconstruction, deformation capture

ACM Reference Format:
Guangming Zang, Ramzi Idoughi, Ran Tao, Gilles Lubineau, Peter Wonka, and, Wolfgang Heidrich. 2019. Warp-and-Project Tomography for Rapidly Deforming Objects. *ACM Trans. Graph.* 38, 4, Article 86 (July 2019), 13 pages. https://doi.org/10.1145/3306346.3322965

Authors' address: Guangming Zang, guangming.zang@kaust.edu.sa; Ramzi Idoughi, ramzi.idoughi@kaust.edu.sa; Ran Tao, ran.tao@kaust.edu.sa; Gilles Lubineau, gilles.lubineau@kaust.edu.sa; Peter Wonka, peter.wonka@kaust.edu.sa;
Wolfgang Heidrich, wolfgang.heidrich@kaust.edu.sa, King Abdullah University of Science And Technology, Thuwal, 23955-6900, Saudi Arabia.

0730-0301/2019/7-ART86
https://doi.org/10.1145/3306346.3322965

1 INTRODUCTION

X-ray computed tomography (CT) is widely used in computer graphics and computer vision and even more frequently in medicine, biology and material science as a non-destructive imaging technique, able to reveal inner structures of the studied object. Until recently, X-ray CT was only used to scan static objects from different viewing angles. The need for a dynamic tomography reconstruction arises in applications where the scanned object undergoes deformation, or if the target of study is the motion itself. Even in the static case, tomography is often an ill-posed problem that requires hundreds of projections to reconstruct high-quality volumes. When the scanned object undergoes deformation, the number of projections for each deformation state is often insufficient for reconstructing each state with a traditional reconstruction methods. This makes the dynamic tomography reconstruction a highly challenging task.

Recently, Zang et al. [2018b] proposed a non-parametric Space-Time tomographic method (ST-tomography) to scan and analyze deforming objects and dynamic phenomena. While this method resulted in marked improvement of the state of the art, it does suffer from several shortcomings that we address in this work: First, ST-tomography was conceived for the case of relatively slow and smooth motion fields, where the deformation is negligible for short sequences of $\approx 10-60$ successive frames. Second, the method relies on an explicit tradeoff between spatial and temporal reconstruction quality. Finally, the temporal sampling is uniform, resulting in wasted computational effort for slow moving periods, as well as poor reconstruction quality for fast moving periods.

In this work, we propose a new warp-and-project approach for dynamic tomographic reconstruction. This new method, inspired by ST-tomography, relaxes the assumption of slow deformations in order to reconstruct objects with larger motion even between successive projections, where the ST-Tomography fails. Our proposed

Symmetric Moving Frames

ETIENNE CORMAN, Carnegie Mellon University / University of Toronto
KEENAN CRANE, Carnegie Mellon University

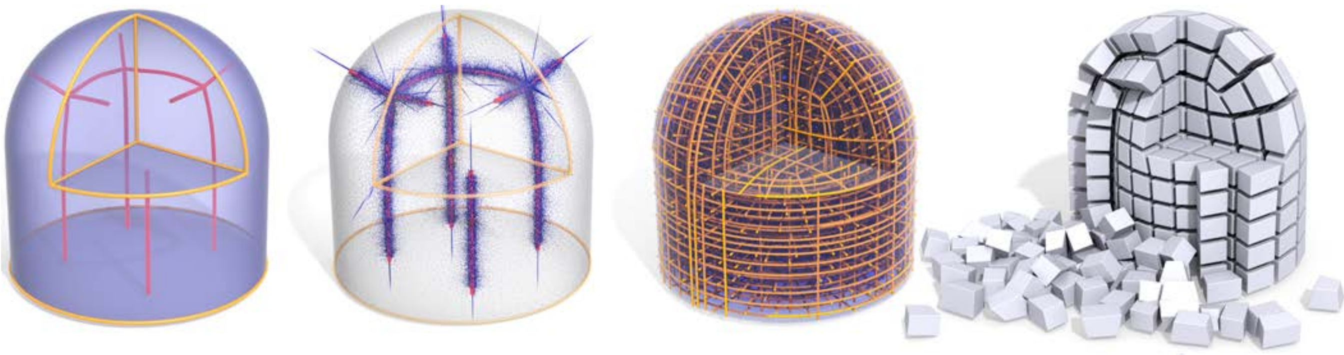

Fig. 1. Given a collection of singular and feature curves on a volumetric domain *(far left)*, we compute the smoothest rotational derivative that winds around these curves *(center left)*, and describes a symmetric 3D cross field *(center right)* which can be directly used for hexahedral meshing *(far right)*.

A basic challenge in field-guided hexahedral meshing is to find a spatially-varying frame that is adapted to the domain geometry and is continuous up to symmetries of the cube. We introduce a fundamentally new representation of such *3D cross fields* based on Cartan's method of moving frames. Our key observation is that cross fields and ordinary frame fields are locally characterized by identical conditions on their *Darboux derivative*. Hence, by using derivatives as the principal representation (and only later recovering the field itself), one avoids the need to explicitly account for symmetry during optimization. At the discrete level, derivatives are encoded by skew-symmetric matrices associated with the edges of a tetrahedral mesh; these matrices encode arbitrarily large rotations along each edge, and can robustly capture singular behavior even on coarse meshes. We apply this representation to compute 3D cross fields that are as smooth as possible everywhere but on a prescribed network of singular curves—since these fields are adapted to curve tangents, they can be directly used as input for field-guided mesh generation algorithms. Optimization amounts to an easy nonlinear least squares problem that behaves like a convex program in the sense that it always appears to produce the same result, independent of initialization. We study the numerical behavior of this procedure, and perform some preliminary experiments with mesh generation.

CCS Concepts: • **Mathematics of computing → Mesh generation**.

Additional Key Words and Phrases: cross fields, discrete differential geometry

ACM Reference Format:
Etienne Corman and Keenan Crane. 2019. Symmetric Moving Frames. *ACM Trans. Graph.* 38, 4, Article 87 (July 2019), 16 pages. https://doi.org/10.1145/3306346.3323029

Authors' addresses: Etienne Corman, Carnegie Mellon University / University of Toronto; Keenan Crane, Carnegie Mellon University, Pittsburgh, PA, 15213.

0730-0301/2019/7-ART87 $15.00
https://doi.org/10.1145/3306346.3323029

0 INTRODUCTION

A *hexahedral mesh* decomposes a solid region of three-dimensional space into six-sided cells; such meshes play an important role in numerical algorithms across geometry processing and scientific computing. An attractive approach to mesh generation is to first construct a *guidance field* oriented along features of interest, then extract a mesh aligned with this field. However, there are major open questions about how to even *represent* such fields in a way that is compatible with the demands of hexahedral meshing—the most elementary of which is how to identify frames that differ by rotational symmetries of the cube. These so-called *3D cross fields* allow one to encode networks of singular features (Fig. 1, far left), which are critical to achieving good element quality.

In differential geometry, *Cartan's method of moving frames* provides a rich theory for spatially-varying coordinate frames, but to date has not been used for hexahedral meshing—perhaps because, classically, it does not consider fields with local rotational symmetry (like cross fields). In this paper we show how the theory of moving frames can be naturally applied in the symmetric case, and how to incorporate constraints needed for hexahedral meshing, namely, adaptation to a network of *singular curves* which correspond to mesh edges of irregular degree. Specifically, we consider the following problem: given a domain and a valid singularity network, find the smoothest 3D cross field compatible with this network. Here, a *valid* network means one that is compatible with the global topology of some hexahedral mesh, as recently studied by Liu et al. [2018].

Computationally, our method amounts to solving an augmented version of *Cartan's second structure equation*

$$\boxed{d\omega - \omega \wedge \omega = 0.}$$

Much as the curl-free condition $\nabla \times X = 0$ characterizes vector fields X that can be locally expressed as the gradient of a scalar potential, the structure equation characterizes differential 1-forms ω which are the *Darboux derivative* of some spatially-varying frame field.

Optimal Transport-Based Polar Interpolation of Directional Fields

JUSTIN SOLOMON, Massachusetts Institute of Technology*
AMIR VAXMAN, Utrecht University*

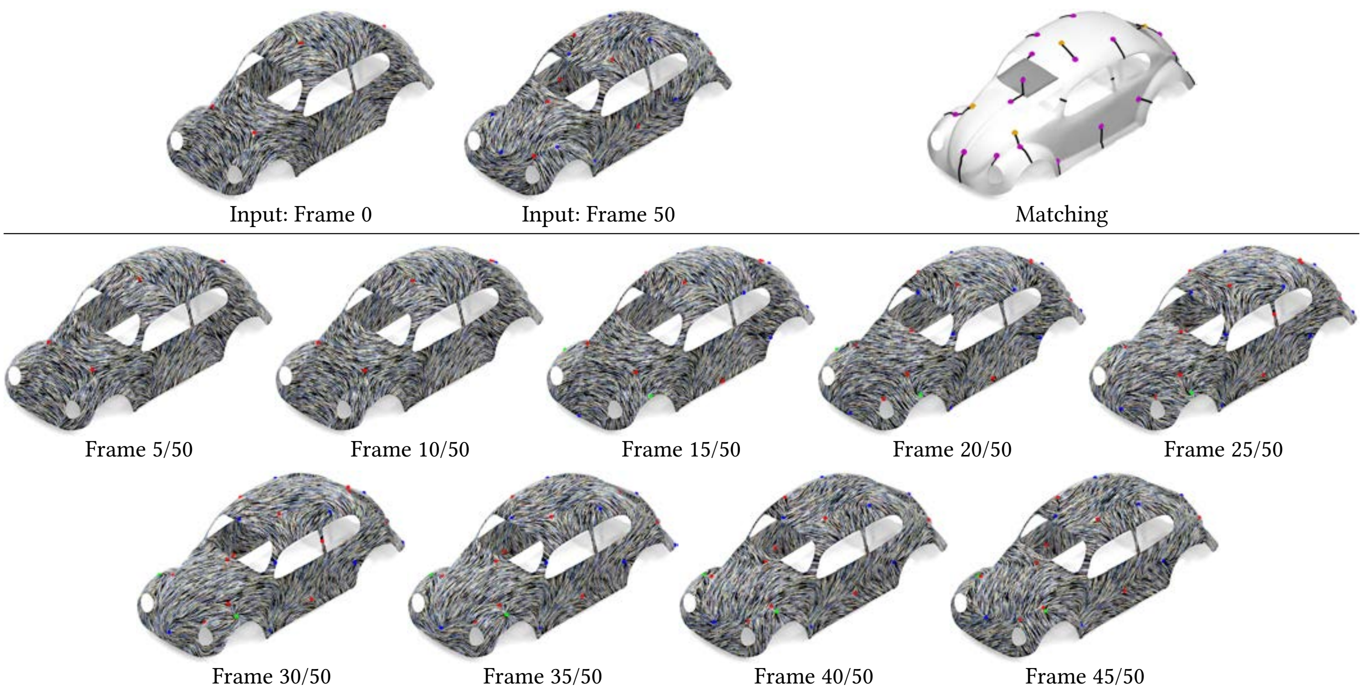

Fig. 1. Our algorithm uses optimal transport to match the singular points (upper right) of two input line fields (upper left) and uses this matching to generate a time-varying interpolant. Our technique interpolates the fields smoothly in time by sliding singularities along the domain rather than having them appear and disappear, merging them with boundary curves as needed. Singularities are color-coded according to index (green for $+2/N$, blue for $+1/N$, red for $-1/N$); purple and orange singularity colors in the matching image distinguish the two input frames.

We propose an algorithm that interpolates between vector and frame fields on triangulated surfaces, designed to complement field design methods in geometry processing and simulation. Our algorithm is based on a *polar* construction, leveraging a conservation law from the Hopf-Poincaré theorem to match singular points using ideas from optimal transport; the remaining detail of the field is interpolated using straightforward machinery. Our model is designed with topology in mind, sliding singular points along the surface rather than having them appear and disappear, and it caters to all surface topologies, including boundary and generator loops.

* Joint first authors.
Authors' addresses: Justin Solomon, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology*, jsolomon@mit.edu; Amir Vaxman, Department of Information and Computing Sciences, Utrecht University*, A.Vaxman@uu.nl.

0730-0301/2019/7-ART88 $15.00
https://doi.org/10.1145/3306346.3323005

CCS Concepts: • **Computing methodologies → Shape analysis**; • **Mathematics of computing →** *Interpolation.*

Additional Key Words and Phrases: Optimal transport, frame fields

ACM Reference Format:
Justin Solomon and Amir Vaxman. 2019. Optimal Transport-Based Polar Interpolation of Directional Fields. *ACM Trans. Graph.* 38, 4, Article 88 (July 2019), 13 pages. https://doi.org/10.1145/3306346.3323005

1 INTRODUCTION

Vector and frame fields are ubiquitous in geometry processing and physically-based animation. Used to guide physical motion, texture synthesis, meshing, and more exotic tasks like path planning, algorithms that design and process fields find application in most geometry-oriented stages of the graphics pipeline.

Field processing algorithms typically fall into two major categories. First, *field design* algorithms in geometry processing place fields on a domain given a sparse set of user constraints, such as placement of singular points or directionality. Second, *simulation* algorithms use vector fields to capture quantities like momentum and

SPOT: Sliced Partial Optimal Transport

NICOLAS BONNEEL, Univ. Lyon, CNRS
DAVID COEURJOLLY, Univ. Lyon, CNRS

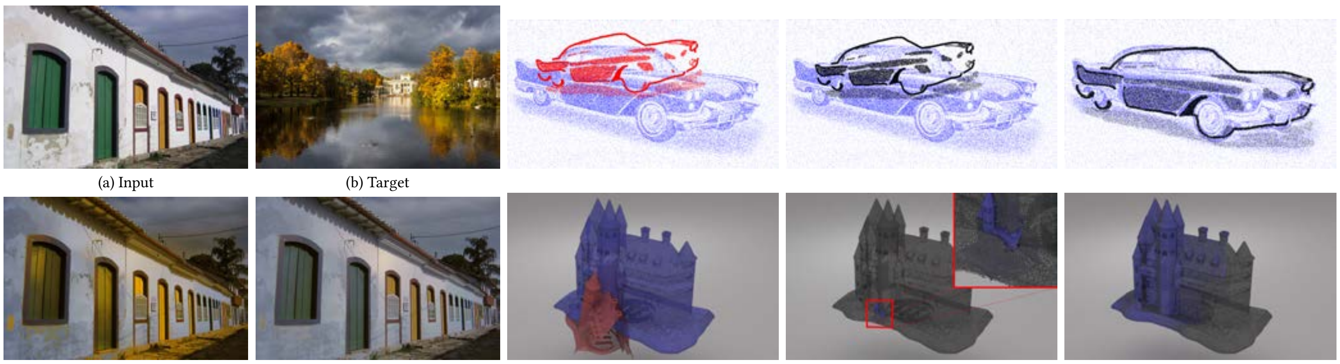

(a) Input (b) Target (c) Full histogram matching (d) Partial histogram matching (e) Point clouds (f) Euclidean ICP (g) Our method

Fig. 1. Our technique allows to partially match point sets within an optimal transport framework. **Left.** In color matching applications such as that of PitiÃľ et al. [Pitié et al. 2005], matching all pixels of an input image (a) seen as 3D points in an RGB space to all pixels of a target image (b) can lead to erroneous transfers (c) due to mismatched content (here, trees of the target are not present in the input, and distort colors in the output). Instead, we match the input image (a) to a subset of an enlarged target image (b), thus effectively preventing spurious colors (d). **Right.** Given a source 2-d or 3-d point cloud (in red) and a target one (in blue) of different sizes (e), we match them with a similarity transform. While classical iterative closest point (ICP) fails (f), our approach, called *fast iterative sliced transport*, achieves robust registrations (g). *Facade by Phil Whitehouse [CC BY 2.0] via Flickr (http://flic.kr/p/48kgPR) and palace by Neil Williamson [CC BY-SA 2.0] (http://flic.kr/p/NJ6Vxq).*

Optimal transport research has surged in the last decade with wide applications in computer graphics. In most cases, however, it has focused on the special case of the so-called "balanced" optimal transport problem, that is, the problem of optimally matching positive measures of equal total mass. While this approach is suitable for handling probability distributions as their total mass is always equal to one, it precludes other applications manipulating disparate measures. Our paper proposes a fast approach to the optimal transport of constant distributions supported on point sets of different cardinality via one-dimensional slices. This leads to one-dimensional partial assignment problems akin to alignment problems encountered in genomics or text comparison. Contrary to one-dimensional balanced optimal transport that leads to a trivial linear-time algorithm, such partial optimal transport, even in 1-d, has not seen any closed-form solution nor very efficient algorithms to date. We provide the first efficient 1-d partial optimal transport solver. Along with a quasilinear time problem decomposition algorithm, it solves 1-d assignment problems consisting of up to millions of Dirac distributions within fractions of a second in parallel. We handle higher dimensional problems via a slicing approach, and further extend the popular iterative closest point algorithm using optimal transport – an algorithm we call *Fast Iterative Sliced Transport*. We illustrate our method on computer graphics applications such a color transfer and point cloud registration.

Authors' addresses: Nicolas Bonneel, Univ. Lyon, CNRS, nicolas.bonneel@liris.cnrs.fr; David Coeurjolly, Univ. Lyon, CNRS, david.coeurjolly@liris.cnrs.fr.

0730-0301/2019/7-ART89 $15.00
https://doi.org/10.1145/3306346.3323021

CCS Concepts: • **Computing methodologies → Image manipulation**; **Point-based models**.

Additional Key Words and Phrases: Optimal transport, sequence alignment

ACM Reference Format:
Nicolas Bonneel and David Coeurjolly. 2019. SPOT: Sliced Partial Optimal Transport. *ACM Trans. Graph.* 38, 4, Article 89 (July 2019), 13 pages. https://doi.org/10.1145/3306346.3323021

1 INTRODUCTION

Optimal transport is a popular mathematical framework for manipulating positive measures, and in particular, in most cases studied so far, *probability measures.* It has become widespread in computer graphics [Bonneel et al. 2011; Nader and Guennebaud 2018; Solomon et al. 2015] and machine learning [Arjovsky et al. 2017; Deshpande et al. 2018; Kolouri et al. 2018] for its ability to compare histograms or to produce compelling interpolations between probability distributions. Within this framework, transporting a histogram f towards another g is often seen as moving a pile of sand shaped by the graph of f towards a hole shaped by g at minimal cost. Stopping the motion of the sand at some intermediate time defines a measure in-between f and g, a concept called *displacement interpolation*, or, when dealing with more than two input measures, *Wasserstein barycenter*. This computationally difficult problem has received recent attention, with fast solutions for various specific cases when the size of the hole matches that of the pile of sand, the so-called

MeshCNN: A Network with an Edge

RANA HANOCKA, Tel Aviv University
AMIR HERTZ, Tel Aviv University
NOA FISH, Tel Aviv University
RAJA GIRYES, Tel Aviv University
SHACHAR FLEISHMAN, Amazon
DANIEL COHEN-OR, Tel Aviv University

Polygonal meshes provide an efficient representation for 3D shapes. They explicitly capture both shape surface and topology, and leverage non-uniformity to represent large flat regions as well as sharp, intricate features. This non-uniformity and irregularity, however, inhibits mesh analysis efforts using neural networks that combine convolution and pooling operations. In this paper, we utilize the unique properties of the mesh for a direct analysis of 3D shapes using *MeshCNN*, a convolutional neural network designed specifically for triangular meshes. Analogous to classic CNNs, MeshCNN combines specialized convolution and pooling layers that operate on the mesh edges, by leveraging their intrinsic geodesic connections. Convolutions are applied on edges and the four edges of their incident triangles, and pooling is applied via an edge collapse operation that retains surface topology, thereby, generating new mesh connectivity for the subsequent convolutions. MeshCNN learns which edges to collapse, thus forming a task-driven process where the network exposes and expands the important features while discarding the redundant ones. We demonstrate the effectiveness of MeshCNN on various learning tasks applied to 3D meshes.

CCS Concepts: • **Computing methodologies → Neural networks**; **Shape analysis**.

Additional Key Words and Phrases: Geometric Deep Learning, Shape Analysis, Convolutional Neural Network, Shape Segmentation

ACM Reference Format:
Rana Hanocka, Amir Hertz, Noa Fish, Raja Giryes, Shachar Fleishman, and Daniel Cohen-Or. 2019. MeshCNN: A Network with an Edge. *ACM Trans. Graph.* 38, 4, Article 90 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3322959

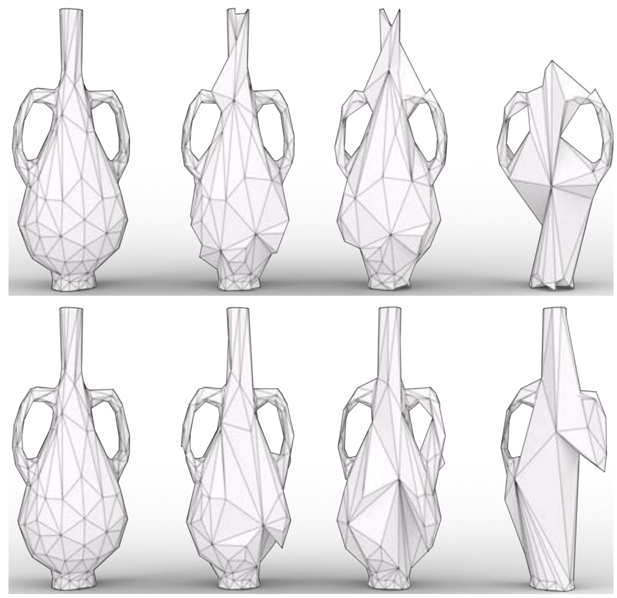

Fig. 1. Mesh pooling operates on irregular structures and adapts spatially to the task. Unlike geometric simplification (removes edges with a minimal geometric distortion), mesh pooling delegates which edges to collapse to the network. Top row: MeshCNN trained to classify whether a vase has a *handle*, bottom row: trained on whether there is a *neck* (top-piece).

1 INTRODUCTION

Three dimensional shapes are prevalent in the field of computer graphics, but also a major commodity in related fields such as computer vision and computational geometry. Shapes around us, and in particular those describing natural entities, are commonly composed of continuous surfaces. For computational reasons, and to facilitate data processing, various discrete approximations for 3D shapes have been suggested and utilized to represent shapes in an array of applications. A favorite of many, the polygonal mesh representation, or mesh, for short, approximates surfaces via a set of 2D polygons in 3D space [Botsch et al. 2010]. The mesh provides an efficient, non-uniform representation of the shape. On the one hand, only a small number of polygons are required to capture large, simple, surfaces. On the other hand, representation flexibility supports a higher resolution where needed, allowing a faithful reconstruction, or portrayal, of salient shape features that are often geometrically intricate. Another advantageous characteristic of the mesh is the inherent ability to encode connectivity information. This forms a comprehensive representation of the underlying surface.

These advantages are apparent in comparison to another popular option: the point cloud representation. Despite its simplicity and direct relation to common data acquisition techniques (scanning), the point cloud representation falls short when a higher quality and preservation of sharp shape features are required.

Authors' addresses: Rana Hanocka, Tel Aviv University; Amir Hertz, Tel Aviv University; Noa Fish, Tel Aviv University; Raja Giryes, Tel Aviv University; Shachar Fleishman, Amazon; Daniel Cohen-Or, Tel Aviv University.

0730-0301/2019/7-ART90 $15.00
https://doi.org/10.1145/3306346.3322959

SAGNet: Structure-aware Generative Network for 3D-Shape Modeling

ZHIJIE WU, Shenzhen University
XIANG WANG, Shenzhen University
DI LIN, Shenzhen University
DANI LISCHINSKI, The Hebrew University of Jerusalem
DANIEL COHEN-OR, Shenzhen University
HUI HUANG*, Shenzhen University

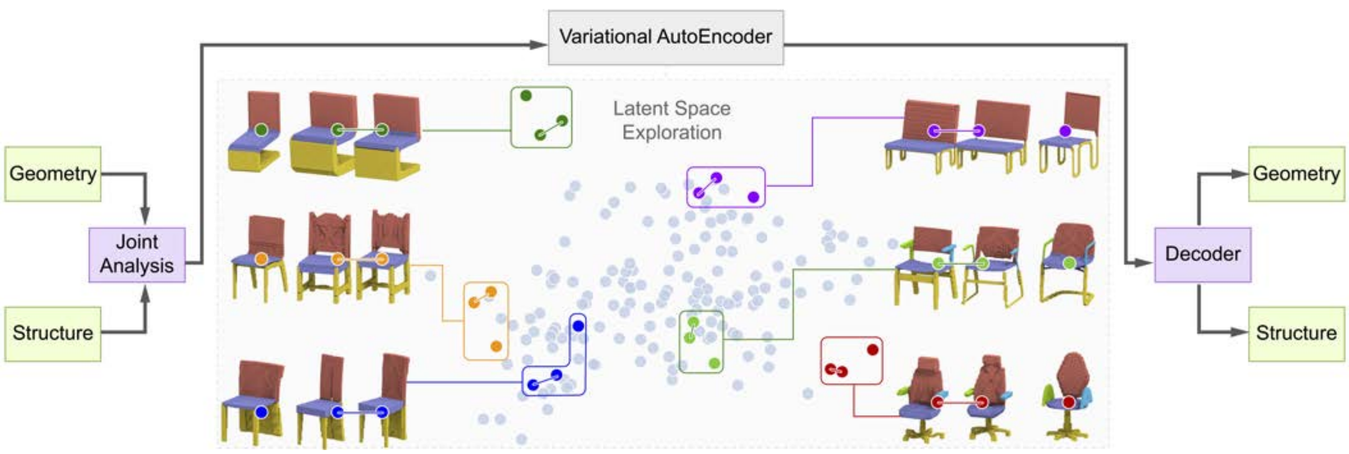

Fig. 1. Our generative model jointly analyzes the structure and geometry of shapes, encoding them into a single latent code. The highlighted triplets above demonstrate that, in this joint latent space, pairs of nearby points represent models that are close to each other in both geometry and structure, while stepping away from the pair introduces differences in either geometry, or structure, or both. Note that all the shapes shown here are voxel-based, and the bounding boxes of their parts are hidden on purpose for a clearer visualization.

We present SAGNet, a structure-aware generative model for 3D shapes. Given a set of segmented objects of a certain class, the geometry of their parts and the pairwise relationships between them (the structure) are jointly learned and embedded in a latent space by an autoencoder. The encoder intertwines the geometry and structure features into a single latent code, while the decoder disentangles the features and reconstructs the geometry and structure of the 3D model. Our autoencoder consists of two branches, one for the structure and one for the geometry. The key idea is that during the analysis, the two branches exchange information between them, thereby learning the dependencies between structure and geometry and encoding two augmented features, which are then fused into a single latent code. This explicit intertwining of information enables separately controlling the geometry and the structure of the generated models. We evaluate the performance of our method and conduct an ablation study. We explicitly show that encoding of shapes accounts for both similarities in structure and geometry. A variety of quality results generated by SAGNet are presented.

CCS Concepts: • **Computing methodologies** → **Computer graphics**; **Shape modeling**; **Shape analysis**.

Additional Key Words and Phrases: geometric modeling, shape analysis, data-driven synthesis, generative network, variational autoencoder

ACM Reference Format:
Zhijie Wu, Xiang Wang, Di Lin, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. 2019. SAGNet: Structure-aware Generative Network for 3D-Shape Modeling. *ACM Trans. Graph.* 38, 4, Article 91 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3322956

1 INTRODUCTION

Modeling of 3D shapes is a central problem in computer graphics. In recent years more attention has been given to structure-aware modeling techniques, where the relations among parts are carefully considered. Analyzing the structure provides a high-level understanding of the shape, and it goes beyond the low-level analysis of the local geometry [Mitra et al. 2014]. The structure of a shape can be inferred from a single instance, but analyzing a family of shapes that share some similar structural characteristics can yield a much more powerful representation [Fish et al. 2014]. However, such an

*Corresponding author: Hui Huang (hhzhiyan@gmail.com)

Authors' addresses: Zhijie Wu, Shenzhen University; Xiang Wang, Shenzhen University; Di Lin, Shenzhen University; Dani Lischinski, The Hebrew University of Jerusalem; Daniel Cohen-Or, Shenzhen University; Hui Huang, College of Computer Science & Software Engineering, Shenzhen University.

0730-0301/2019/7-ART91 $15.00
https://doi.org/10.1145/3306346.3322956

iMapper: Interaction-guided Scene Mapping from Monocular Videos

ARON MONSZPART, University College London
PAUL GUERRERO, University College London
DUYGU CEYLAN, Adobe Research
ERSIN YUMER, Uber ATG
NILOY J. MITRA, University College London and Adobe Research

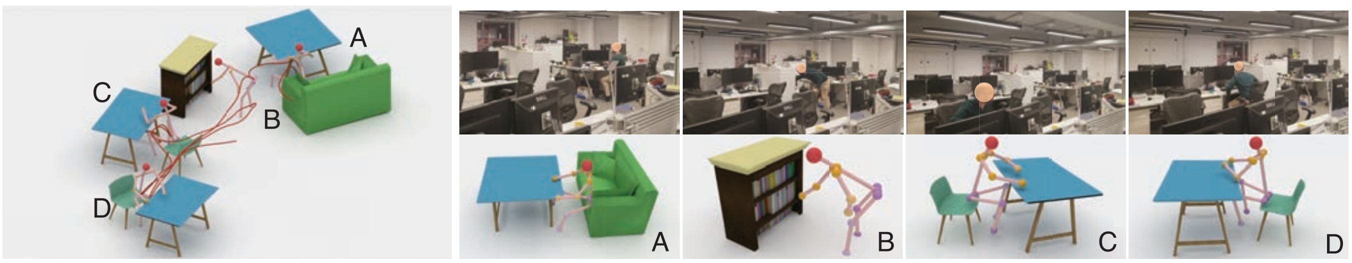

Fig. 1. **Interaction-guided scene mapping.** We present iMapper that discovers potential human-object *interactions* in an input monocular video and utilizes them to infer the object layout of the recorded scene containing medium to heavy occlusion. We show the final generated 3D scene as well as recovered interactions (Scene13). Note that although *only* the 2D human joint detection (left) is available to our algorithm, here we additionally show reference video frames (corresponding to A, B, C) to help judge the original scene layout. Please refer to supplementary video.

Next generation smart and augmented reality systems demand a computational understanding of monocular footage that captures humans in physical spaces to reveal plausible object arrangements and human-object interactions. Despite recent advances, both in scene layout and human motion analysis, the above setting remains challenging to analyze due to regular occlusions that occur between objects and human motions. We observe that the *interaction* between object arrangements and human actions is often strongly correlated, and hence can be used to help recover from these occlusions. We present iMapper, a data-driven method to identify such human-object interactions and utilize them to infer layouts of occluded objects. Starting from a monocular video with detected 2D human joint positions that are potentially noisy and occluded, we first introduce the notion of *interaction-saliency* as space-time snapshots where informative human-object interactions happen. Then, we propose a global optimization to retrieve and fit interactions from a database to the detected salient interactions in order to best explain the input video. We extensively evaluate the approach, both quantitatively against manually annotated ground truth and through a user study, and demonstrate that iMapper produces plausible scene layouts for scenes with medium to heavy occlusion. Code and data are available on the project page.

Authors' addresses: Aron Monszpart, Department of Computer Science, University College London, aron.monszpart.12@ucl.ac.uk; Paul Guerrero, Department of Computer Science, University College London, paul.guerrero@ucl.ac.uk; Duygu Ceylan, Adobe Research, ceylan@adobe.com; Ersin Yumer, Uber ATG, meyumer@gmail.com; Niloy J. Mitra, Department of Computer Science, University College London and Adobe Research, n.mitra@ucl.ac.uk.

0730-0301/2019/7-ART92 $15.00
https://doi.org/10.1145/3306346.3322961

CCS Concepts: • **Computing methodologies → Scene understanding**; **Shape inference**; **Shape analysis**; *Motion capture*; *Perception.*

Additional Key Words and Phrases: shape analysis, interaction, scene layout, monocular video, occlusion

ACM Reference Format:
Aron Monszpart, Paul Guerrero, Duygu Ceylan, Ersin Yumer, and Niloy J. Mitra. 2019. iMapper: Interaction-guided Scene Mapping from Monocular Videos. *ACM Trans. Graph.* 38, 4, Article 92 (July 2019), 15 pages. https://doi.org/10.1145/3306346.3322961

1 INTRODUCTION

Computational understanding of monocular videos that capture human-object interactions in physical spaces is critical for many emerging fields such as virtual and augmented reality, smart home systems, assisted living, and robotics. Such applications require access to object arrangements embedded in the physical spaces along with the common human-object interactions performed in such spaces. For example, our future personal robot assistants should know our working habits along with the supporting object arrangements in our living rooms or workspaces. Hence, a joint understanding of scenes and human actions from the input feed is necessary.

While both scene understanding and human performance analysis are highly popular research areas, traditionally, researchers have tackled them as two separate problems. On the one hand, *scene estimation methods* such as Kinect Fusion [Newcombe et al. 2011] and Bundle Fusion [Dai et al. 2017b] can produce high-quality static indoor reconstructions, while the likes of DynamicFusion [Newcombe et al. 2015] can capture non-rigidly deforming scenes by fusing depth information across space and time. These methods, however, require the sensor to be manually moved to see around occlusions making

On the Accurate Large-scale Simulation of Ferrofluids

LIBO HUANG, TORSTEN HÄDRICH, DOMINIK L. MICHELS, KAUST

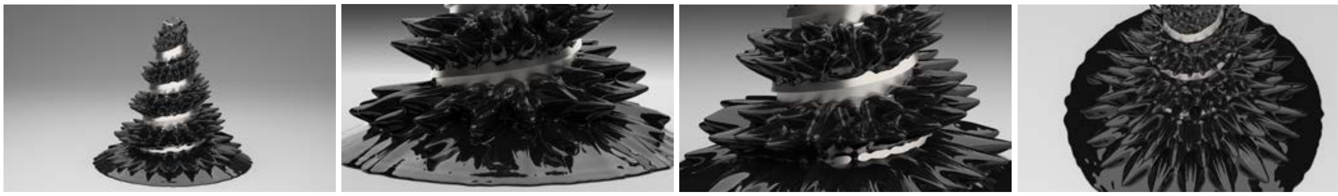

Fig. 1. Our approach can accurately reproduce the observation that real ferrofluid is literally climbing up a steel helix placed above a strong electromagnet. This figure shows the final results of our simulation of this scenario rendered from different viewpoints.

We present an approach to the accurate and efficient large-scale simulation of the complex dynamics of ferrofluids based on physical principles. Ferrofluids are liquids containing magnetic particles that react to an external magnetic field without solidifying. In this contribution, we employ smooth magnets to simulate ferrofluids in contrast to previous methods based on the finite element method or point magnets. We solve the magnetization using the analytical solution of the smooth magnets' field, and derive the bounded magnetic force formulas addressing particle penetration. We integrate the magnetic field and force evaluations into the fast multipole method allowing for efficient large-scale simulations of ferrofluids. The presented simulations are well reproducible since our approach can be easily incorporated into a framework implementing a Fast Multipole Method and a Smoothed Particle Hydrodynamics fluid solver with surface tension. We provide a detailed analysis of our approach and validate our results against real wet lab experiments. This work can potentially open the door for a deeper understanding of ferrofluids and for the identification of new areas of applications of these materials.

CCS Concepts: • **Computing methodologies → Physical simulation**.

Additional Key Words and Phrases: Computational Electromagnetics, Fast Multipole Method (FMM), Ferrofluids, Fluid Mechanics, Large-scale Simulations, Magnetic Fluids, Maxwell's equations, Natural Phenomena, Navier-Stokes Equations, Numerical Simulations, Smoothed Particle Hydrodynamics (SPH), Validation Experiments.

ACM Reference Format:
Libo Huang, Torsten Hädrich, Dominik L. Michels. 2019. On the Accurate Large-scale Simulation of Ferrofluids. *ACM Trans. Graph.* 38, 4, Article 93 (July 2019), 15 pages. https://doi.org/10.1145/3306346.3322973

1 INTRODUCTION

Ferrofluids were invented in the early 1960s by NASA scientists to pump fuel into spacecrafts in low gravity environments without mechanical actions. Today, these fluids are mostly known to the public for its peculiar behavior when interacting with magnetic fields. Because of the interesting geometric structure, they are widely used in artwork (as, *e.g.,* in the SIGGRAPH Art Gallery [Kodama 2008]), science exhibitions, and in several advertisements. Due to the complexity of the movements of their spikes, it becomes impossible for an artist to create a lively animation by intuition. A fast and accurate ferrofluid simulator can not only be used for animation, but also to explore, and perhaps better understand, the interplay between the magnetic field, the geometry of the external object, and the resulting surface shape/topology of the ferrofluid. However, realistic simulations of the ferrofluid dynamics are not yet available. In this contribution, for the first time, we achieved the level of realism to bring its fascination from reality to the virtual world. With such a simulation tool at hand, we can further design the magnetic field and external objects in the virtual world to create the desired shapes in reality.

Next to their artistic applications, ferrofluids are employed in different scenarios described in the literature [Nochetto et al. 2016a; Raj et al. 1995] ranging from applications in acoustics, instrumentation, lubrication, vacuum technology, vibration damping [Miwa et al. 2003], to radar absorbing materials [Vinoy and Jha 1996]. Several applications are also touching the field of visual computing, for example, ferrofluids are used for magnetic resonance imaging contrast enhancement and the construction of adaptive deformable mirrors [Brousseau et al. 2007]. Additional applications can be found in the field of micro- and nanoelectronics [Hartshorne et al. 2004; Zahn 2001].

The interesting dynamics of ferrofluids are caused by the interplay of an external magnetic field and the surface tension, forming a pattern of characteristic spikes. Figure 2 illustrates the influence of the magnetic field strength and surface tension on the shape of the spikes: a strong surface tension force smooths the contour, while a strong field strength increases the height of the spikes. These effects are further coupled with fluid motion finally leading to a highly complex dynamical behavior. Given this degree of complexity, the large-scale numerical simulation of ferrofluids can potentially open the door for a deeper understanding of ferrofluids and for the identification of new areas of applications of these materials.

In this contribution, we aim for the large-scale simulation of ferrofluid dynamics. The simulation of magnetism-based effects is

Author's address: Libo Huang, Torsten Hädrich, Dominik L. Michels, KAUST, Visual Computing Center, Thuwal 23955, KSA.

0730-0301/2019/7-ART93
https://doi.org/10.1145/3306346.3322973

An Adaptive Variational Finite Difference Framework for Efficient Symmetric Octree Viscosity

RYAN GOLDADE and YIPENG WANG, University of Waterloo, Canada
MRIDUL AANJANEYA, Rutgers University, USA
CHRISTOPHER BATTY, University of Waterloo, Canada

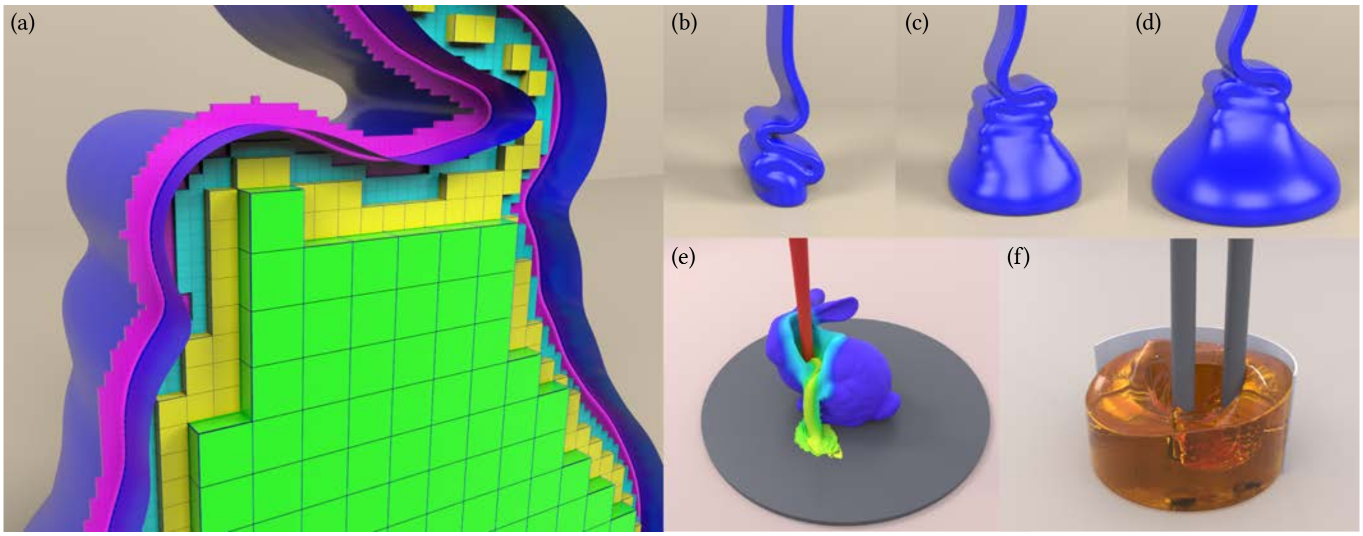

Fig. 1. Our adaptive viscosity discretization is constructed on a graded octree structure (a), and achieves speed-up factors for the linear solve ranging from 3.8 to 9.4 compared to the regular grid approach. Our method supports rotational effects observed with a buckling sheet of viscous liquid (b-d), spatially varying viscosity coefficients (e), and kinematic objects (f).

While pressure forces are often the bottleneck in (near-)inviscid fluid simulations, viscosity can impose orders of magnitude greater computational costs at lower Reynolds numbers. We propose an implicit octree finite difference discretization that significantly accelerates the solution of the free surface viscosity equations using adaptive staggered grids, while supporting viscous buckling and rotation effects, variable viscosity, and interaction with scripted moving solids. In experimental comparisons against regular grids, our method reduced the number of active velocity degrees of freedom by as much as a factor of 7.7 and reduced linear system solution times by factors between 3.8 and 9.4. We achieve this by developing a novel adaptive variational finite difference methodology for octrees and applying it to the optimization form of the viscosity problem. This yields a linear system that is symmetric positive definite by construction, unlike naive finite difference/volume methods, and much sparser than a hypothetical finite element alternative. Grid refinement studies show spatial convergence at first order in L_∞ and second order in L_1, while the significantly smaller size of the octree linear systems allows for the solution of viscous forces at higher effective resolutions than with regular grids. We demonstrate the practical benefits of our adaptive scheme by replacing the regular grid viscosity step of a commercial liquid simulator (Houdini) to yield large speed-ups, and by incorporating it into an existing inviscid octree simulator to add support for viscous flows. Animations of viscous liquids pouring, bending, stirring, buckling, and melting illustrate that our octree method offers significant computational gains and excellent visual consistency with its regular grid counterpart.

CCS Concepts: • **Computing methodologies → Physical simulation**.

Additional Key Words and Phrases: viscosity, liquid, octree, variational, symmetry

ACM Reference Format:
Ryan Goldade, Yipeng Wang, Mridul Aanjaneya, and Christopher Batty. 2019. An Adaptive Variational Finite Difference Framework for Efficient Symmetric Octree Viscosity. *ACM Trans. Graph.* 38, 4, Article 94 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3322939

Authors' addresses: Ryan Goldade, rgoldade@uwaterloo.ca; Yipeng Wang, University of Waterloo, 200 University Ave W., Waterloo, ON, N2L 3G1, Canada, yipeng.wang@uwaterloo.ca; Mridul Aanjaneya, Rutgers University, 110 Frelinghuysen Road, Piscataway, NJ, 08854, USA, mridul.aanjaneya@rutgers.edu; Christopher Batty, University of Waterloo, 200 University Ave W., Waterloo, ON, N2L 3G1, Canada, christopher.batty@uwaterloo.ca.

0730-0301/2019/7-ART94 $15.00
https://doi.org/10.1145/3306346.3322939

1 INTRODUCTION

Pressure and viscosity are the two fundamental internal forces that govern the motion of all Newtonian fluids; a stable and efficient

Mixing Sauces: A Viscosity Blending Model for Shear Thinning Fluids

KENTARO NAGASAWA*, The University of Tokyo
TAKAYUKI SUZUKI, The University of Tokyo
RYOHEI SETO, Kyoto University
MASATO OKADA, The University of Tokyo
YONGHAO YUE*, The University of Tokyo, Aoyama Gakuin University

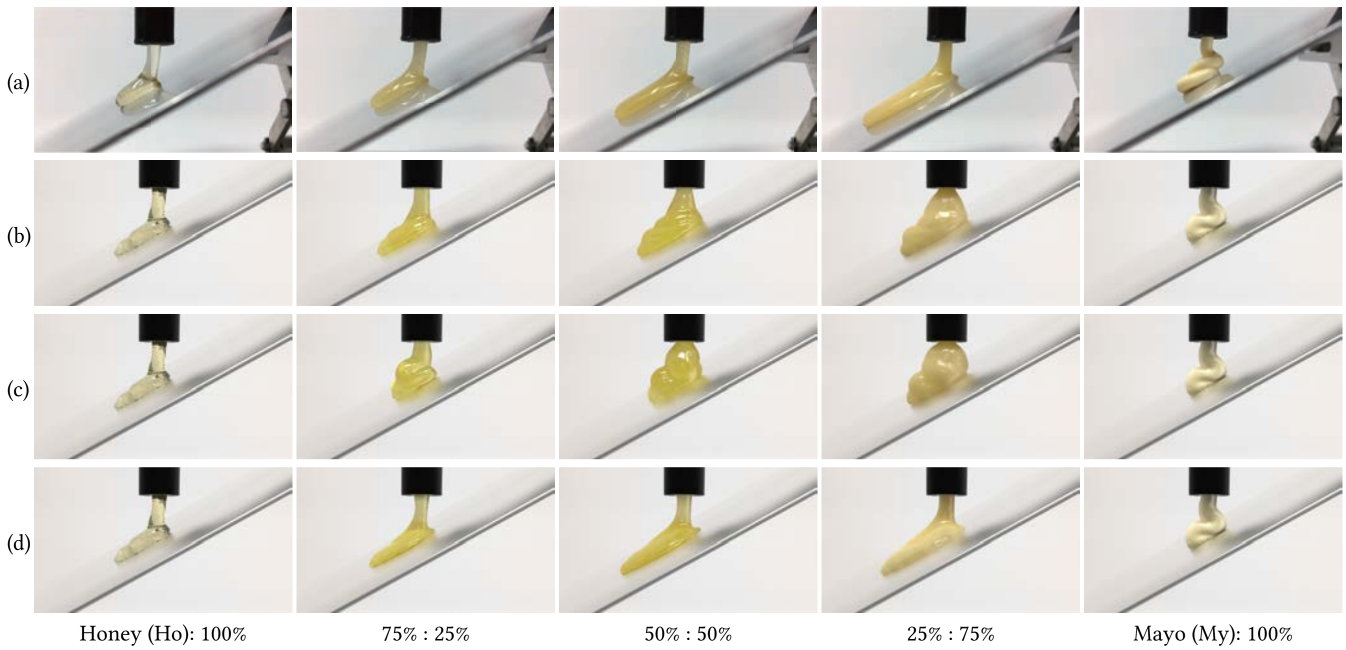

Fig. 1. Blending honey and mayo. (a) A captured footage of the mixtures with various mixing ratios. The material appears 'stagnant' when it is a pure mayo or a pure honey, but their mixtures flow much more smoothly. (b) With an MPM simulation using the linearly blended material parameters, the flow all looks stagnant irrespective of the mixing ratio. (c) The flow still looks stagnant when we initialize the material points with either the property of pure mayo or pure honey, with their volume fraction set according to the mixing ratio. (d) Our nonlinear blending model generates the mixed material property only using the material properties of the pure constituents and their mixing ratios, with which we can reproduce the smoothly flowing behaviors of the blended states.

The materials around us usually exist as *mixtures* of constituents, each constituent with possibly a different *elasto-viscoplastic* property. How can we describe the material property of such a mixture is the core question of this paper. We propose a nonlinear blending model that can capture intriguing flowing behaviors that can differ from that of the individual constituents (Fig. 1). We used a laboratory device, *rheometer*, to measure the flowing properties of various fluid-like foods, and found that an *elastic Herschel-Bulkley model* has nice agreements with the measured data even for the mixtures of these foods. We then constructed a blending model such that it qualitatively agrees with the measurements and is closed in the parameter space of the elastic Herschel-Bulkley model. We provide validations through comparisons between the measured and estimated properties using our model, and comparisons between simulated examples and captured footages. We show the utility of our model for producing interesting behaviors of various mixtures.

CCS Concepts: • **Computing methodologies → Physical simulation**;

Additional Key Words and Phrases: Physical simulation, viscoplasticity, nonlinear blending, material point method, shear thinning fluids

ACM Reference Format:
Kentaro Nagasawa, Takayuki Suzuki, Ryohei Seto, Masato Okada, and Yonghao Yue. 2019. Mixing Sauces: A Viscosity Blending Model for Shear Thinning Fluids. *ACM Trans. Graph.* 38, 4, Article 95 (July 2019), 17 pages. https://doi.org/10.1145/3306346.3322947

*Authors contributed equally.

Authors' addresses: K. Nagasawa, T. Suzuki, M. Okada, Kiban 701, Kashiwa-no-ha 5-1-5, Kashiwa-shi, Chiba, 277-8561, Japan; R. Seto, A4-113, Katsura-Campus, Sakyo-ku, Kyoto University, Kyoto, 615-8510, Japan; Y. Yue, O-525, Fuchinobe 5-10-1, Chuo-ku, Sagamihara, Kanagawa, 252-5258, Japan.

0730-0301/2019/7-ART95 $15.00
https://doi.org/10.1145/3306346.3322947

SurfaceBrush: From Virtual Reality Drawings to Manifold Surfaces

ENRIQUE ROSALES, University of British Columbia, Canada and Universidad Panamericana, México
JAFET RODRIGUEZ, Universidad Panamericana, México
ALLA SHEFFER, University of British Columbia, Canada

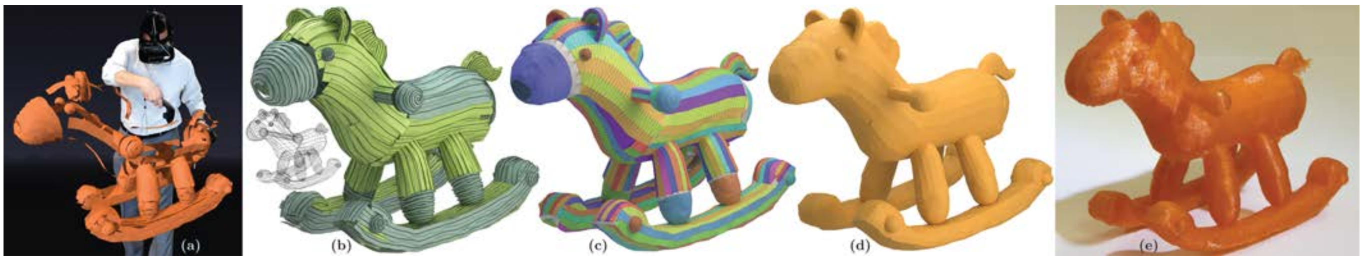

Fig. 1. (a) Drawing 3D strokes using a VR brush. (b) Completed 3D brush-stroke drawing with central stroke polylines drawn in black and ribbon color reflecting normal orientation (green for front, turquoise for back, inset shows poylines alone). (c) triangle mesh strips connecting adjacent stroke polylines (multicolor), and gray triangle strips that complete the surface connecting differently directed stroke groups. (d) Final output. (e) Fabricated model. Input drawing: © Jafet Rodriguez.

Popular Virtual Reality (VR) tools allow users to draw varying-width, ribbon-like 3D brush strokes by moving a hand-held controller in 3D space. Artists frequently use dense collections of such strokes to draw virtual 3D shapes. We propose *SurfaceBrush*, a surfacing method that converts such VR drawings into user-intended manifold free-form 3D surfaces, providing a novel approach for modeling 3D shapes. The inputs to our method consist of dense collections of artist-drawn stroke ribbons described by the positions and normals of their central polylines, and ribbon widths. These inputs are highly distinct from those handled by existing surfacing frameworks and exhibit different sparsity and error patterns, necessitating a novel surfacing approach. We surface the input stroke drawings by identifying and leveraging local coherence between nearby artist strokes. In particular, we observe that strokes intended to be adjacent on the artist imagined surface often have similar tangent directions along their respective polylines. We leverage this local stroke direction consistency by casting the computation of the user-intended manifold surface as a constrained matching problem on stroke polyline vertices and edges. We first detect and smoothly connect adjacent similarly-directed sequences of stroke edges producing one or more manifold partial surfaces. We then complete the surfacing process by identifying and connecting adjacent similarly directed edges along the borders of these partial surfaces. We confirm the usability of the SurfaceBrush interface and the validity of our drawing analysis via an observational study. We validate our stroke surfacing algorithm by demonstrating an array of manifold surfaces computed by our framework starting from a range of inputs of varying complexity, and by comparing our outputs to reconstructions computed using alternative means.

CCS Concepts: • **Computing methodologies → Mesh geometry models**; *Virtual reality*.

Additional Key Words and Phrases: Virtual Reality, 3D drawing, surface modeling, surface reconstruction

ACM Reference Format:
Enrique Rosales, Jafet Rodriguez, and Alla Sheffer. 2019. SurfaceBrush: From Virtual Reality Drawings to Manifold Surfaces. *ACM Trans. Graph.* 38, 4, Article 96 (July 2019), 15 pages. https://doi.org/10.1145/3306346.3322970

1 INTRODUCTION

Humans frequently communicate 3D shapes via 2D sketches or drawings, inspiring the development of modeling interfaces that employ such drawings as inputs [Olsen et al. 2009]. Virtual Reality (VR) systems support real-time capture and visualization of human 3D gestures enabling users to draw surfaces directly in 3D space (Figure 1a). Using such drawings as input for 3D modeling can sidestep the main algorithmic challenge of 2D sketch-based modeling methods – the need to derive 3D information from a 2D input. Effectively leveraging the opportunity provided by VR interfaces requires modeling frameworks capable of processing the types of 3D drawings users can comfortably provide using these tools. Our observations show that artists using the VR medium frequently depict complex free-form 3D geometries using collections of dense, ruled surface *brush strokes* traced in 3D space (Figure 1b) [Sketchfab 2018]. Our *SurfaceBrush* framework algorithmically converts VR brush stroke drawings into manifold surface meshes describing the user-intended geometry (Figure 1d), enabling downstream applications such as 3D printing (Figure 1e).

Users of VR systems, such as TiltBrush [2018] or Drawing on Air [Keefe et al. 2007], trace strokes using a handheld controller. These systems then automatically translate controller locations into polyline stroke vertex positions and controller orientations into stroke normals. They subsequently render the captured input as

Authors' addresses: Enrique Rosales, University of British Columbia, Canada, albertr@cs.ubc.ca, Universidad Panamericana, Facultad de Ingeniería, Zapopan, Jalisco, 45010, México; Jafet Rodriguez, Universidad Panamericana, Facultad de Ingeniería, Zapopan, Jalisco, 45010, México, arodrig@up.edu.mx; Alla Sheffer, University of British Columbia, Canada, sheffa@cs.ubc.ca.

0730-0301/2019/7-ART96 $15.00
https://doi.org/10.1145/3306346.3322970

Perceptual Rasterization for Head-mounted Display Image Synthesis

SEBASTIAN FRISTON, University College London
TOBIAS RITSCHEL, University College London
ANTHONY STEED, University College London

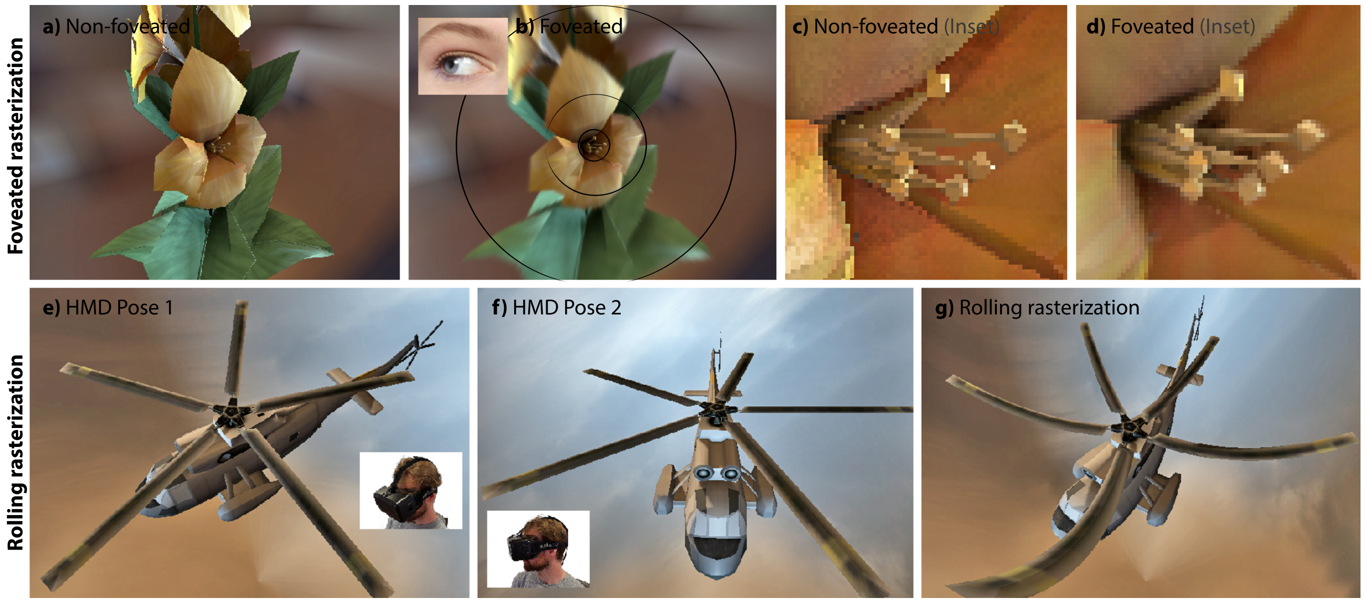

Fig. 1. Perceptual rasterization is a generalization of classic rasterization to the requirements of HMDs such as foveation *(top row)* and rolling image formation *(bottom row)*. On an HMD, most pixels appear in the periphery **(a)**. We rasterize images with continuously-varying pixel density **(b)**. A zoom of the the foveated area shows how a common same-shading-effort image has aliasing **(c)**, while our result benefits from higher pixel density, resulting in super-sampling **(d)**. In common rasterization, each pixel on the display is effectively sampled at the same simulation time ($t = 0$ for the first frame **(e)** and $t = 1$ for the next frame **(f)**). When displayed on a "rolling" HMD display, where pixels are illuminated at different points in time, latency is introduced: the rightmost pixel is outdated by ca. 16 ms. Our rolling rasterization **(g)** allows spatially-varying time: starting at $t = 0$ on the left of the image and increasing to 1 on the right.

We suggest a rasterization pipeline tailored towards the needs of HMDs, where latency and field-of-view requirements pose new challenges beyond those of traditional desktop displays. Instead of image warping for low latency, or using multiple passes for foveation, we show how both can be produced directly in a single perceptual rasterization pass. We do this with per-fragment ray-casting. This is enabled by derivations of tight space-time-fovea pixel bounds, introducing just enough flexibility for the requisite geometric tests, but retaining most of the simplicity and efficiency of the traditional rasterizaton pipeline. To produce foveated images, we rasterize to an image with spatially varying pixel density. To compensate for latency, we extend the image formation model to directly produce "rolling" images where the time at each pixel depends on its display location. Our approach overcomes limitations of warping with respect to disocclusions, object motion and view-dependent shading, as well as geometric aliasing artifacts in other foveated rendering techniques. A set of perceptual user studies demonstrates the efficacy of our approach.

CCS Concepts: • **Computing methodologies** → **Rasterization**; *Perception*; *Virtual reality*.

Additional Key Words and Phrases: foveated rendering, latency, perceptually-based, virtual reality

ACM Reference Format:
Sebastian Friston, Tobias Ritschel, and Anthony Steed. 2019. Perceptual Rasterization for Head-mounted Display Image Synthesis. *ACM Trans. Graph.* 38, 4, Article 97 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3323033

Authors' addresses: Sebastian Friston, University College London, sebastian.friston@ucl.ac.uk; Tobias Ritschel, University College London, t.ritschel@ucl.ac.uk; Anthony Steed, University College London, a.steed@ucl.ac.uk.

0730-0301/2019/7-ART97 $15.00
https://doi.org/10.1145/3306346.3323033

1 INTRODUCTION

The use cases of HMDs have requirements beyond those of typical desktop display-based systems. Completely subsuming the user's vision, the HMD and system driving it must maintain low and predictable latency to facilitate a sense of agency and avoid serious negative consequences such as breaks-in-presence [Slater 2002], simulator sickness [Buker et al. 2012], and reduced performance [Ellis

Luminance-Contrast-Aware Foveated Rendering

OKAN TARHAN TURSUN, Max Planck Institute for Informatics
ELENA ARABADZHIYSKA-KOLEVA, Saarland University, MMCI and Max Planck Institute for Informatics
MAREK WERNIKOWSKI, West Pomeranian University of Technology
RADOSŁAW MANTIUK, West Pomeranian University of Technology
HANS-PETER SEIDEL, Max Planck Institute for Informatics
KAROL MYSZKOWSKI, Max Planck Institute for Informatics
PIOTR DIDYK, Università della Svizzera Italiana

Fig. 1. Current foveated rendering techniques (left) use a fixed quality decay for peripheral vision. While this can be a conservative solution, it does not provide a full computational benefit. Our technique (right) performs content-adaptive foveation and relaxes the quality requirements for content for which the sensitivity of the human visual system at large eccentricities degrades faster. Image by pixel2013 / Pixabay.

Current rendering techniques struggle to fulfill quality and power efficiency requirements imposed by new display devices such as virtual reality headsets. A promising solution to overcome these problems is foveated rendering, which exploits gaze information to reduce rendering quality for the peripheral vision where the requirements of the human visual system are significantly lower. Most of the current solutions model the sensitivity as a function of eccentricity, neglecting the fact that it also is strongly influenced by the displayed content. In this work, we propose a new luminance-contrast-aware foveated rendering technique which demonstrates that the computational savings of foveated rendering can be significantly improved if local luminance contrast of the image is analyzed. To this end, we first study the resolution requirements at different eccentricities as a function of luminance patterns. We later use this information to derive a low-cost predictor of the foveated rendering parameters. Its main feature is the ability to predict the parameters using only a low-resolution version of the current frame, even though the prediction holds for high-resolution rendering. This property is essential for the estimation of required quality before the full-resolution image is rendered. We demonstrate that our predictor can efficiently drive the foveated rendering technique and analyze its benefits in a series of user experiments.

CCS Concepts: • **Computing methodologies → Perception**; **Rendering**; **Image manipulation**;

Additional Key Words and Phrases: foveated rendering, perception

ACM Reference Format:
Okan Tarhan Tursun, Elena Arabadzhiyska-Koleva, Marek Wernikowski, Radosław Mantiuk, Hans-Peter Seidel, Karol Myszkowski, and Piotr Didyk. 2019. Luminance-Contrast-Aware Foveated Rendering. *ACM Trans. Graph.* 38, 4, Article 98 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3322985

Authors' addresses: Okan Tarhan Tursun, Max Planck Institute for Informatics, okan.tursun@mpi-inf.mpg.de; Elena Arabadzhiyska-Koleva, Saarland University, MMCI, Max Planck Institute for Informatics, earabadz@mpi-inf.mpg.de; Marek Wernikowski, West Pomeranian University of Technology, mwernikowski@wi.zut.edu.pl; Radosław Mantiuk, West Pomeranian University of Technology, rmantiuk@wi.zut.edu.pl; Hans-Peter Seidel, Max Planck Institute for Informatics, hpseidel@mpi-inf.mpg.de; Karol Myszkowski, Max Planck Institute for Informatics, karol@mpi-inf.mpg.de; Piotr Didyk, Università della Svizzera Italiana, piotr.didyk@usi.ch.

0730-0301/2019/7-ART98 $15.00
https://doi.org/10.1145/3306346.3322985

1 INTRODUCTION

New display designs, such as virtual and augmented reality glasses, may revolutionize the way we interact with the virtual and the real worlds. While virtual reality (VR) enables us to experience new, unknown environments which we would not be able to explore otherwise, augmented reality (AR) allows us to enrich reality with digital information. For these technologies to succeed, the visual quality delivered by the new devices has to first meet the capabilities and the requirements of the human visual system (HVS). In par-

Foveated AR: Dynamically-Foveated Augmented Reality Display

JONGHYUN KIM, YOUNGMO JEONG*, MICHAEL STENGEL, KAAN AKŞIT, RACHEL ALBERT, BEN BOUDAOUD, TREY GREER, JOOHWAN KIM, WARD LOPES, ZANDER MAJERCIK, PETER SHIRLEY, JOSEF SPJUT, MORGAN MCGUIRE, and DAVID LUEBKE, NVIDIA, United States

Fig. 1. Display results from our Foveated AR prototype. By tracking the user's gaze direction (red cross), the system dynamically provides high-resolution inset images to the foveal region and low-resolution large-FOV images to the periphery. The system supports accommodation cues; the magenta and blue zoom-in panels show optical defocus of real objects together with foveated display of correctly defocus-blurred synthetic objects. Red dashed discs highlight the foveal vs peripheral display regions. A monocular wearable prototype (functional but manually actuated) illustrates the compact optical path.

We present a near-eye augmented reality display with resolution and focal depth dynamically driven by gaze tracking. The display combines a traveling microdisplay relayed off a concave half-mirror magnifier for the high-resolution foveal region, with a wide field-of-view peripheral display using a projector-based Maxwellian-view display whose nodal point is translated to follow the viewer's pupil during eye movements using a traveling holographic optical element. The same optics relay an image of the eye to an infrared camera used for gaze tracking, which in turn drives the foveal display location and peripheral nodal point. Our display supports accommodation cues by varying the focal depth of the microdisplay in the foveal region, and by rendering simulated defocus on the "always in focus" scanning laser projector used for peripheral display. The resulting family of displays significantly improves on the field-of-view, resolution, and form-factor tradeoff present in previous augmented reality designs. We show prototypes supporting 30, 40 and 60 cpd foveal resolution at a net 85° × 78° field of view per eye.

CCS Concepts: • **Hardware → Communication hardware, interfaces and storage**; **Displays and imagers**.

*Also with Seoul National University.

Authors' address: Jonghyun Kim, jonghyunk@nvidia.com; Youngmo Jeong, youngmo.snu@gmail.com; Michael Stengel, mstengel@nvidia.com; Kaan Akşit, kaksit@nvidia.com; Rachel Albert, ralbert@nvidia.com; Ben Boudaoud, bboudaoud@nvidia.com; Trey Greer, tgreer@nvidia.com; Joohwan Kim, sckim@nvidia.com; Ward Lopes, wlopes@nvidia.com; Zander Majercik, amajercik@nvidia.com; Peter Shirley, pshirley@nvidia.com; Josef Spjut, jspjut@nvidia.com; Morgan McGuire, mcguire@nvidia.com; David Luebke, dluebke@nvidia.com, NVIDIA, 2788 San Tomas Expressway, Santa Clara, CA, United States, 95051.

0730-0301/2019/7-ART99 $15.00
https://doi.org/10.1145/3306346.3322987

Additional Key Words and Phrases: foveated,varifocal,augmented,display

ACM Reference Format:
Jonghyun Kim, Youngmo Jeong, Michael Stengel, Kaan Akşit, Rachel Albert, Ben Boudaoud, Trey Greer, Joohwan Kim, Ward Lopes, Zander Majercik, Peter Shirley, Josef Spjut, Morgan McGuire, and David Luebke. 2019. Foveated AR: Dynamically-Foveated Augmented Reality Display. *ACM Trans. Graph.* 38, 4, Article 99 (July 2019), 15 pages. https://doi.org/10.1145/3306346.3322987

1 INTRODUCTION

Augmented reality (AR) aims to present virtual objects at real-world positions and orientations, without compromising the viewer's natural vision. These objects may be rendered photorealistically, or more practically, like the emissive ghosts depicted by science fiction "holographic" projections. Also called mixed reality, AR allows richer and more natural interaction than displays such as Google Glass, which simply superimpose 2D content in a "heads-up display" style akin to aviation and automotive windshield displays, or the video-passthrough smartphone AR applications.

The enabling technology for widespread consumer AR is a high-fidelity head-mounted display (HMD), sometimes termed a near-eye display (NED), that is comfortable both in visual properties and form factor. Such AR displays offer a tantalizing replacement for smartphone and computer screens, but in practice all designs involve tradeoffs between field of view (FOV), resolution, eye box (the "sweet spot" within which a viewer's pupil will perceive a correct image), correct focus cues, and form factor. The greatest challenge in HMD design is not in optimizing any individual metric, but instead *simultaneously* providing a wide FOV, variable focus, high resolution, a wide eye box, and a slim form factor.

Meeting this challenge requires significant new advances in the underlying technology. Inherent physical constraints preclude evolving traditional, larger displays to hit AR targets. The requirement that an AR display be see-through constrains the form factor and

Vidgets: Modular Mechanical Widgets for Mobile Devices

CHANG XIAO, Columbia University
KARL BAYER, Snap Inc.
CHANGXI ZHENG, Columbia University
SHREE K. NAYAR, Snap Inc.

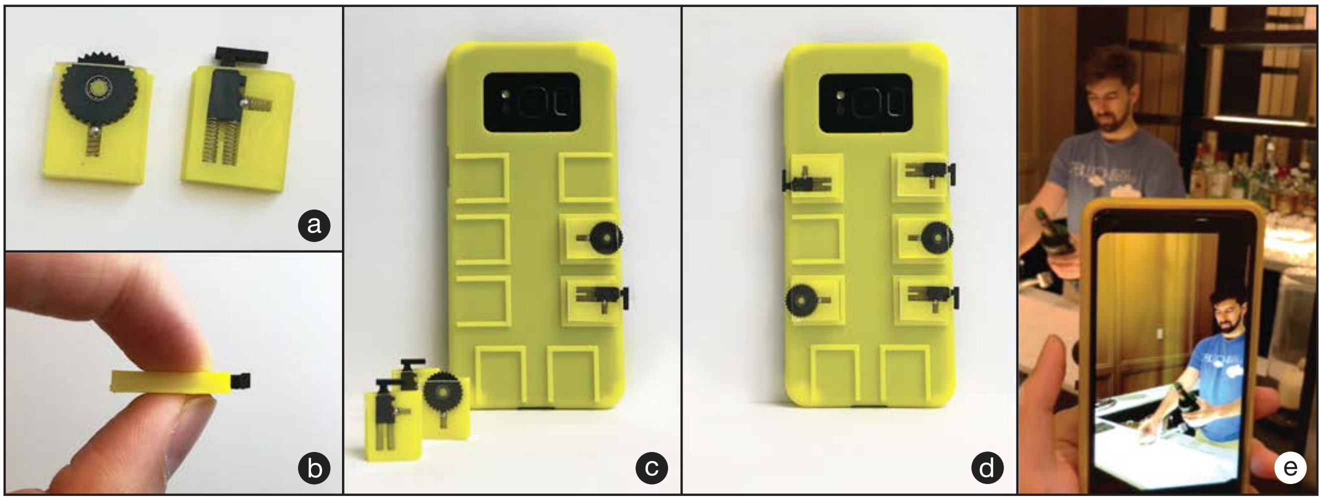

Fig. 1. **Vidgets widgets.** We design a family of physical widgets that are modular, power efficient, (a) compact, and (b) thin. (c) These widgets can be attached to a smartphone's protective case in a modular fashion. (d) The widgets can be used to construct a wide range of user interfaces and adapt to individual's preferences. (e) As just one example, here we use a Vidgets knob to zoom a camera's field of view and a Vidgets button to trigger the capture. In this way, the user can easily take a photo with a desired focus using a single hand, without needing another hand to pinch on the screen. We demonstrate several more applications in the paper.

We present *Vidgets*, a family of mechanical widgets, specifically push buttons and rotary knobs that augment mobile devices with tangible user interfaces. When these widgets are attached to a mobile device and a user interacts with them, the widgets' nonlinear mechanical response shifts the device slightly and quickly, and this subtle motion can be detected by the accelerometer commonly equipped on mobile devices. We propose a physics-based model to understand the nonlinear mechanical response of widgets. This understanding enables us to design tactile force profiles of these widgets so that the resulting accelerometer signals become easy to recognize. We then develop a lightweight signal processing algorithm that analyzes the accelerometer signals and recognizes how the user interacts with the widgets in real time. Vidgets widgets are low-cost, compact, reconfigurable, and power efficient. They can form a diverse set of physical interfaces that enrich users' interactions with mobile devices in various practical scenarios. We demonstrate their use in three applications: photo capture with single-handed zoom, control of mobile games, and making a playable mobile music instrument.

CCS Concepts: • **Human-centered computing** → *Haptic devices*.

Additional Key Words and Phrases: mechanical widget, accelerometer

ACM Reference Format:
Chang Xiao, Karl Bayer, Changxi Zheng, and Shree K. Nayar. 2019. Vidgets: Modular Mechanical Widgets for Mobile Devices. *ACM Trans. Graph.* 38, 4, Article 100 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3322943

Authors' addresses: Chang Xiao, Columbia University, chang@cs.columbia.edu; Karl Bayer, Snap Inc. kbayer@snap.com; Changxi Zheng, Columbia University, cxz@cs.columbia.edu; Shree K. Nayar, Snap Inc. snayar@snap.com.

0730-0301/2019/7-ART100 $15.00
https://doi.org/10.1145/3306346.3322943

1 INTRODUCTION

The quest to have the maximum possible screen size and thinnest possible bezel is dominating today's smartphone market. Both Apple and Samsung removed the Home button and HTC has removed all physical buttons from their latest smartphones. Hardware widgets such as buttons on mobile devices appear to be going extinct.

Yet, there are still many aspects of physical widgets that their digital counterparts have not improved on, or even equalled: Physical widgets give us a reassuring feeling of control, a touching sense of the phone's orientation, and are provably more efficient for applications such as gaming [Chu and Wong 2011; Lee and Oulasvirta 2016; Zaman et al. 2010]. Physical widgets also often complement software widgets, in cases where it is just inconvenient for the user to touch the screen—for example, when the user's hands are wet or

Tangent-Space Optimization for Interactive Animation Control

LOÏC CICCONE, ETH Zürich
CENGIZ ÖZTIRELI, DisneyResearch|Studios
ROBERT W. SUMNER, DisneyResearch|Studios and ETH Zürich

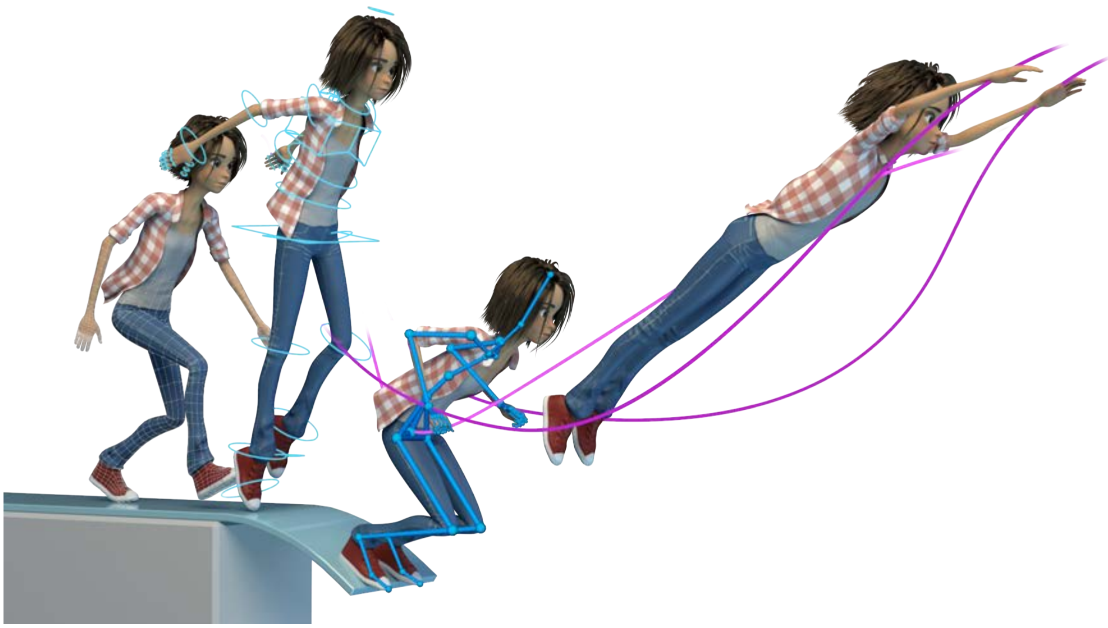

Fig. 1. Left: In traditional animation the character's deformations are driven by rig controls (light blue), which provide fine control but require a granular interaction. Middle: Our system provides an armature (dark blue) that coordinately drives the rig elements for a more natural and flexible character manipulation. Right: We introduce a curve representation (purple) for easily controlling the interpolation between key poses without adding any keyframe.

Character animation tools are based on a keyframing metaphor where artists pose characters at selected keyframes and the software automatically interpolates the frames inbetween. Although the quality of the interpolation is critical for achieving a fluid and engaging animation, the tools available to adjust the result of the automatic inbetweening are rudimentary and typically require manual editing of spline parameters. As a result, artists spend a tremendous amount of time posing and setting more keyframes. In this pose-centric workflow, animators use combinations of forward and inverse kinematics. While forward kinematics leads to intuitive interpolations, it does not naturally support positional constraints such as fixed contact points. Inverse kinematics can be used to fix certain points in space at keyframes, but can lead to inferior interpolations, is slow to compute, and does not allow for positional contraints at non-keyframe frames. In this paper, we address these problems by formulating the control of interpolations with positional constraints over time as a space-time optimization problem in the tangent space of the animation curves driving the controls. Our method has the key properties that it (1) allows the manipulation of positions and orientations over time, extending inverse kinematics, (2) does not add new keyframes that might conflict with an artist's preferred keyframe style, and (3) works in the space of artist editable animation curves and hence integrates seamlessly with current pipelines. We demonstrate the utility of the technique in practice via various examples and use cases.

CCS Concepts: • **Computing methodologies** → **Animation**; *Graphics systems and interfaces.*

Additional Key Words and Phrases: Interpolation, Inverse kinematics

ACM Reference Format:
Loïc Ciccone, Cengiz Öztireli, and Robert W. Sumner. 2019. Tangent-Space Optimization for Interactive Animation Control. *ACM Trans. Graph.* 38, 4, Article 101 (July 2019), 10 pages. https://doi.org/10.1145/3306346.3322938

Authors' addresses: Loïc Ciccone, ETH Zürich; Cengiz Öztireli, DisneyResearch|Studios; Robert W. Sumner, DisneyResearch|Studios, ETH Zürich.

0730-0301/2019/7-ART101 $15.00
https://doi.org/10.1145/3306346.3322938

1 INTRODUCTION

Character animation software provides the tools, algorithms, and interfaces that artists use to breath life into animated characters. Contemporary software uses a keyframing metaphor inspired by classic hand-drawn animation: Artists pose characters at selected

Vibration-Minimizing Motion Retargeting for Robotic Characters

SHAYAN HOSHYARI, Disney Research and University of British Columbia
HONGYI XU, Disney Research
ESPEN KNOOP, Disney Research
STELIAN COROS, ETH Zurich
MORITZ BÄCHER, Disney Research

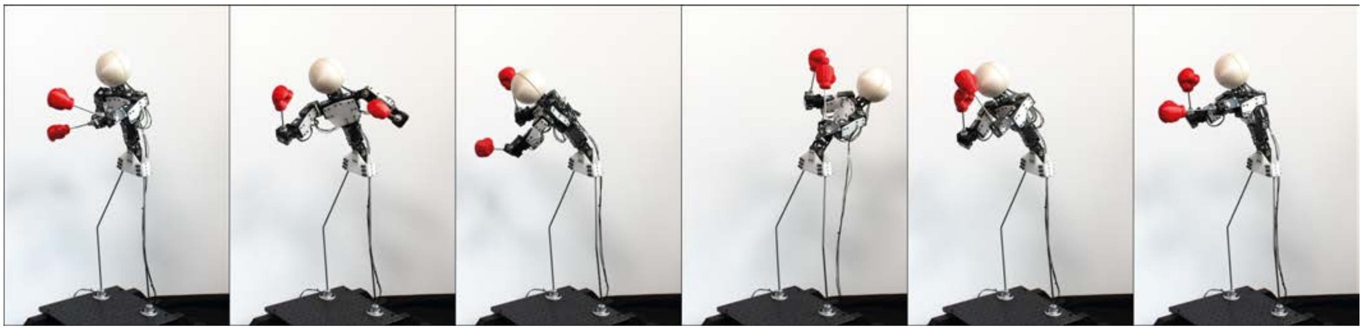

Fig. 1. We present a method for retargeting fast and dynamic animations onto physical robot characters, where the motor trajectories are optimized in order to suppress unwanted structural vibrations and match the artistic intent as closely as possible. We demonstrate our approach on a range of examples, including, as seen here, a boxing sequence with fast punches, blocks, and dodges.

Creating animations for robotic characters is very challenging due to the constraints imposed by their physical nature. In particular, the combination of fast motions and unavoidable structural deformations leads to mechanical oscillations that negatively affect their performances. Our goal is to automatically transfer motions created using traditional animation software to robotic characters while avoiding such artifacts. To this end, we develop an optimization-based, dynamics-aware motion retargeting system that adjusts an input motion such that visually salient low-frequency, large amplitude vibrations are suppressed. The technical core of our animation system consists of a differentiable dynamics simulator that provides constraint-based two-way coupling between rigid and flexible components. We demonstrate the efficacy of our method through experiments performed on a total of five robotic characters including a child-sized animatronic figure that features highly dynamic drumming and boxing motions.

CCS Concepts: • **Computing methodologies** → *Physical simulation.*

Additional Key Words and Phrases: animation retargeting, robotic characters, dynamics, model reduction, vibration minimization, adjoint method

Authors' addresses: Shayan Hoshyari, Disney Research and University of British Columbia; Hongyi Xu, Disney Research; Espen Knoop, Disney Research; Stelian Coros, ETH Zurich; Moritz Bächer, Disney Research.

0730-0301/2019/7-ART102 $15.00
https://doi.org/10.1145/3306346.3323034

ACM Reference Format:
Shayan Hoshyari, Hongyi Xu, Espen Knoop, Stelian Coros, and Moritz Bächer. 2019. Vibration-Minimizing Motion Retargeting for Robotic Characters. *ACM Trans. Graph.* 38, 4, Article 102 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3323034

1 INTRODUCTION

Ever since Leonardo da Vinci's times, children and adults alike have been fascinated by mechanical systems that are designed to generate natural movements. Over the centuries, da Vinci's first automatons—the Mechanical Lion and Knight—have evolved into lifelike animatronic figures that are routinely deployed in museums, theme parks and movie studios across the world. And today, thanks to the advent of affordable, easy-to-use digital fabrication technologies and electromechanical components, the process of creating compelling robotic characters is easily accessible to anyone.

Keyframing techniques are commonly used to breathe life into animatronic characters. While these techniques are conceptually identical to those employed in *Computer Animation*, creating vibrant motions for *real-world* characters introduces unique challenges. These challenges stem from the physical characteristics of an animatronic figure's design. It is easy, for example, to design virtual characters that have as many degrees of freedom as needed. The design of their robotic counterparts, on the other hand, must balance motion versatility against the constraints imposed by the size, weight and placement of its mechanical components. Furthermore, the motions of real-world characters are strictly bound to the laws of physics. While the idealized limbs of a virtual character are perfectly rigid, for example, structural deformations are unavoidable in physical systems. Unfortunately, the combination of dynamic

PuppetMaster: Robotic Animation of Marionettes

SIMON ZIMMERMANN, ETH Zurich
ROI PORANNE, ETH Zurich and University of Haifa
JAMES M. BERN, ETH Zurich
STELIAN COROS, ETH Zurich

We present a computational framework for robotic animation of real-world string puppets. Also known as marionettes, these articulated figures are typically brought to life by human puppeteers. The puppeteer manipulates rigid handles that are attached to the puppet from above via strings. The motions of the marionette are therefore governed largely by gravity, the pull forces exerted by the strings, and the internal forces arising from mechanical articulation constraints. This seemingly simple setup conceals a very challenging and nuanced control problem, as marionettes are, in fact, complex coupled pendulum systems. Despite this, in the hands of a master puppeteer, marionette animation can be nothing short of mesmerizing. Our goal is to enable autonomous robots to animate marionettes with a level of skill that approaches that of human puppeteers. To this end, we devise a predictive control model that accounts for the dynamics of the marionette and kinematics of the robot puppeteer. The input to our system consists of a string puppet design and a target motion, and our trajectory planning algorithm computes robot control actions that lead to the marionette moving as desired. We validate our methodology through a series of experiments conducted on an array of marionette designs and target motions. These experiments are performed both in simulation and using a physical robot, the human-sized, dual arm ABB YuMi® IRB 14000.

CCS Concepts: • **Theory of computation → Nonconvex optimization**; • **Computing methodologies → Physical simulation**; • **Computer systems organization → Robotic control**;

Additional Key Words and Phrases: Computer graphics, robotics, puppeteering, sensitivity analysis

ACM Reference Format:
Simon Zimmermann, Roi Poranne, James M. Bern, and Stelian Coros. 2019. PuppetMaster: Robotic Animation of Marionettes. *ACM Trans. Graph.* 38, 4, Article 103 (July 2019), 11 pages. https://doi.org/10.1145/3306346.3323003

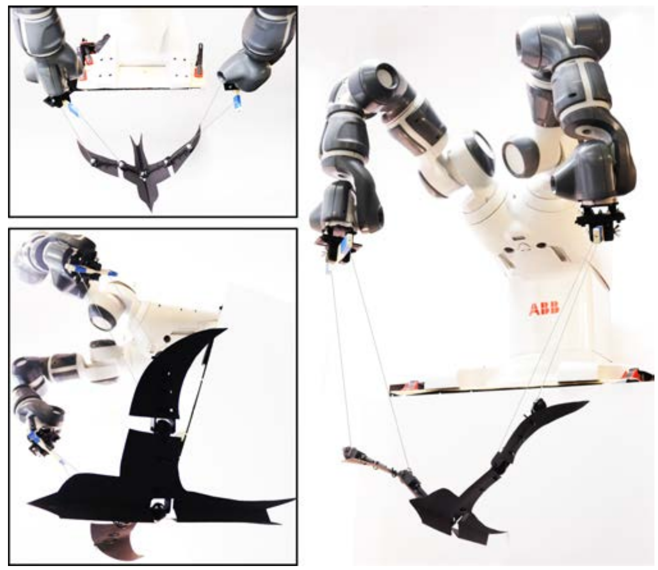

Fig. 1. A robot (ABB YuMi® IRB 14000) controlling a bird-shaped, string-driven marionette. This type of puppet is notoriously difficult to animate. Our control framework generates optimal motion trajectories for the robot puppeteer such that the marionette performs a user-specified motion.

1 INTRODUCTION

Marionettes are articulated, string-actuated puppets that have provided a medium for animation in performance arts since Ancient Greece. In the hands of a skilled puppeteer, marionettes produce motions that are incredibly expressive, fluid and compelling. However, their deceptively natural movements conceal the fact that marionettes are very challenging to control. Marionettes are under-actuated, high-dimensional, highly non-linear coupled pendulum systems. They are driven by gravity, the tension forces generated by a small number of cables, and the internal forces arising from mechanical articulation constraints. As such, the map between the actions of a puppeteer and the motions performed by the marionette is notoriously unintuitive, and mastering this unique art form takes unfaltering dedication and a great deal of practice.

With the long term goal of endowing robots with human-level dexterity when it comes to manipulating complex physical systems, we present *PuppetMaster* – a physics-based motion planning framework for robotic animation of marionettes. As illustrated in Fig. 2, the input to our computational puppeteering system consists of: 1) a kinematic description of the robot puppeteer (i.e. the hierarchical arrangement of actuators and rigid links); 2) the design of the marionette, which includes the articulated puppet itself, the control handles that the robot will be manipulating, as well as the strings that attach the puppet to the handles; and 3) a target motion that the marionette should aim to reproduce. Our core technical contribution is a novel trajectory optimization method built upon a physics simulator using an implicit time-stepping scheme. To compute derivatives of the system's forward dynamics, we show how to apply sensitivity analysis techniques to *entire motion trajectories*, as well as how to effectively exploit the specific sparsity structure imposed by the time domain. Our mathematical model forms the

Authors' addresses: Simon Zimmermann, ETH Zurich; Roi Poranne, ETH Zurich and University of Haifa; James M. Bern, ETH Zurich; Stelian Coros, ETH Zurich.

0730-0301/2019/7-ART103 $15.00
https://doi.org/10.1145/3306346.3323003

RedMax: Efficient & Flexible Approach for Articulated Dynamics

YING WANG, Texas A&M University
NICHOLAS J. WEIDNER, Texas A&M University
MARGARET A. BAXTER, Texas A&M University
YURA HWANG, Texas A&M University
DANNY M. KAUFMAN, Adobe Research
SHINJIRO SUEDA, Texas A&M University

(a) (b) (c) (d)

Fig. 1. (a) Near linear time evaluation for a cable-stayed Bridge. (b) Near linear time evaluation for a deployable Umbrella. (c) Starfish, showing fully two-way coupled integration between articulated and deformable bodies. (d) Klann walker, with rapid evaluation of internal friction within joints.

It is well known that the dynamics of articulated rigid bodies can be solved in $O(n)$ time using a recursive method, where n is the number of joints. However, when elasticity is added between the bodies (*e.g.,* damped springs), with linearly implicit integration, the stiffness matrix in the equations of motion breaks the tree topology of the system, making the recursive $O(n)$ method inapplicable. In such cases, the only alternative has been to form and solve the system matrix, which takes $O(n^3)$ time. We propose a new approach that is capable of solving the linearly implicit equations of motion in near linear time. Our method, which we call RedMax, is built using a combined reduced/maximal coordinate formulation. This hybrid model enables direct flexibility to apply arbitrary combinations of constraints and contact modeling in both reduced and maximal coordinates, as well as mixtures of implicit and explicit forces in either coordinate representation. We highlight RedMax's flexibility with seamless integration of deformable objects with two-way coupling, at a standard additional cost. We further highlight its flexibility by constructing an efficient internal (joint) and external (environment) frictional contact solver that can leverage bilateral joint constraints for rapid evaluation of frictional articulated dynamics.

CCS Concepts: • **Computing methodologies** → **Physical simulation**.

Additional Key Words and Phrases: Physical simulation, Rigid body dynamics, Constraints, Contact, Friction

Authors' addresses: Ying Wang, Texas A&M University; Nicholas J. Weidner, Texas A&M University; Margaret A. Baxter, Texas A&M University; Yura Hwang, Texas A&M University; Danny M. Kaufman, Adobe Research; Shinjiro Sueda, Texas A&M University.

0730-0301/2019/7-ART104 $15.00
https://doi.org/10.1145/3306346.3322952

ACM Reference Format:
Ying Wang, Nicholas J. Weidner, Margaret A. Baxter, Yura Hwang, Danny M. Kaufman, and Shinjiro Sueda. 2019. RedMax: Efficient & Flexible Approach for Articulated Dynamics. *ACM Trans. Graph.* 38, 4, Article 104 (July 2019), 10 pages. https://doi.org/10.1145/3306346.3322952

1 INTRODUCTION

Articulated rigid body dynamics has many applications in various disciplines, including biomechanics, robotics, aerospace, and computer graphics. It has been extensively studied starting in the 1960s (*e.g.,* [Roberson 1966]), but it was not until the 1980s that an $O(n)$ algorithm, where n is the number of joints or bodies, became widely known [Featherstone 1983]. This algorithm and its variants are based on a recursive formulation, where various quantities are computed recursively based on the tree structure of the mechanism. Around the same time, an alternative $O(n^3)$ method based on matrix factorization was also developed [Walker and Orin 1982], which, according to De Jalon and Bayo [2012], can outperform the $O(n)$ recursive method when n is small (< 10). Although some important mechanisms, such as serial manipulators, have only a few joints, for many applications in computer graphics, n can be quite large—even a single hand has $n \geq 15$. Therefore, there are still many cases where $O(n)$ methods are still preferred over $O(n^3)$ methods.

The story changes when implicit elasticity is added between *arbitrary* bodies rather than only between immediate neighbors. Examples of such scenarios include: simply attaching damped springs between pairs of bodies; architectural design with cables [Whiting et al. 2012; Deuss et al. 2014]; musculoskeletal simulations with line-based forces [Delp et al. 2007; Wang et al. 2012]; and deployable folding mechanisms [Demaine and O'Rourke 2008; Zhou et al. 2014]. Using the linearly implicit integrator commonly used in graphics [Baraff and Witkin 1998], the $O(n)$ recursive method no longer

Spectral Coarsening of Geometric Operators

HSUEH-TI DEREK LIU, University of Toronto, Canada
ALEC JACOBSON, University of Toronto, Canada
MAKS OVSJANIKOV, École Polytechnique, France

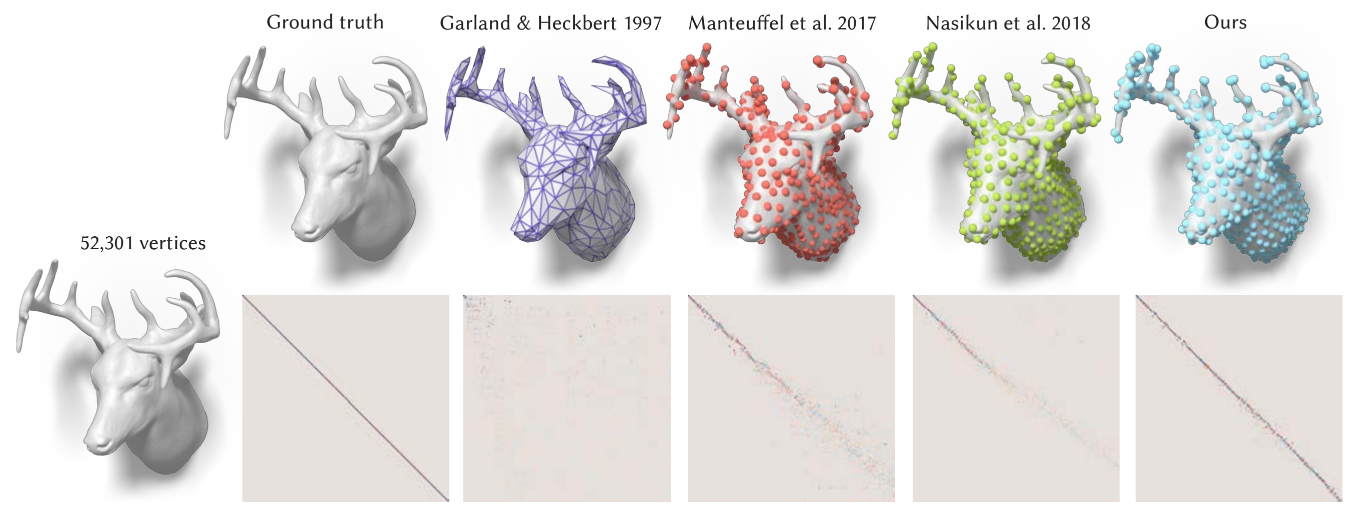

Fig. 1. There are many ways to coarsen a 52,301×52,301 sparse anisotropic Laplace matrix down to a sparse 500×500 matrix: simplify the mesh [Garland and Heckbert 1997] and rediscretize; apply algebraic multigrid coarsening [Manteuffel et al. 2017]; or approximate using radial-basis functions [Nasikun et al. 2018]. We introduce a way to measure how well the coarse operator maintains the original operator's eigenvectors (bottom row). The visualization shows deviation from a diagonal matrix indicating poor eigenvector preservation. In response, we introduce an optimization to coarsen geometric operators while preserving eigenvectors and maintaining sparsity and positive semi-definiteness.

We introduce a novel approach to measure the behavior of a geometric operator before and after coarsening. By comparing eigenvectors of the input operator and its coarsened counterpart, we can quantitatively and visually analyze how well the spectral properties of the operator are maintained. Using this measure, we show that standard mesh simplification and algebraic coarsening techniques fail to maintain spectral properties. In response, we introduce a novel approach for *spectral coarsening*. We show that it is possible to significantly reduce the sampling density of an operator derived from a 3D shape without affecting the low-frequency eigenvectors. By marrying techniques developed within the algebraic multigrid and the functional maps literatures, we successfully coarsen a variety of isotropic and anisotropic operators while maintaining sparsity and positive semi-definiteness. We demonstrate the utility of this approach for applications including operator-sensitive sampling, shape matching, and graph pooling for convolutional neural networks.

Authors' addresses: Hsueh-Ti Derek Liu, University of Toronto, 40 St. George Street, Toronto, ON, M5S 2E4, Canada; Alec Jacobson, University of Toronto, 40 St. George Street, Toronto, ON, M5S 2E4, Canada; Maks Ovsjanikov, École Polytechnique, 1 Rue Honoré d'Estienne d'Orves, Palaiseau, 91120, France.

0730-0301/2019/7-ART105 $15.00
https://doi.org/10.1145/3306346.3322953

CCS Concepts: • **Computing methodologies** → **Shape analysis**; • **Mathematics of computing** → *Computations on matrices.*

Additional Key Words and Phrases: geometry processing, numerical coarsening, spectral geometry

ACM Reference Format:
Hsueh-Ti Derek Liu, Alec Jacobson, and Maks Ovsjanikov. 2019. Spectral Coarsening of Geometric Operators. *ACM Trans. Graph.* 38, 4, Article 105 (July 2019), 13 pages. https://doi.org/10.1145/3306346.3322953

1 INTRODUCTION

Geometry processing relies heavily on building matrices to represent linear operators defined on geometric domains. While typically sparse, these matrices are often too large to work with efficiently when defined over high resolution representations. A common solution is to simplify or coarsen the domain. However, matrices built from coarse representations often do not behave the same way as their fine counterparts leading to inaccurate results and artifacts when resolution is restored. Quantifying and categorizing *how* this behavior is different is not straightforward and most often coarsening is achieved through operator-oblivious remeshing. The common appearance-based or geometric metrics employed by remeshers, such as the classical quadratic error metric [Garland and Heckbert 1997] can have very little correlation to maintaining operator behavior.

Tensor Maps for Synchronizing Heterogeneous Shape Collections

QIXING HUANG, The University of Texas at Austin, USA
ZHENXIAO LIANG, The University of Texas at Austin, USA
HAOYUN WANG, Tsinghua University, USA
SIMIAO ZUO, The University of Texas at Austin, USA
CHANDRAJIT BAJAJ, The University of Texas at Austin, USA

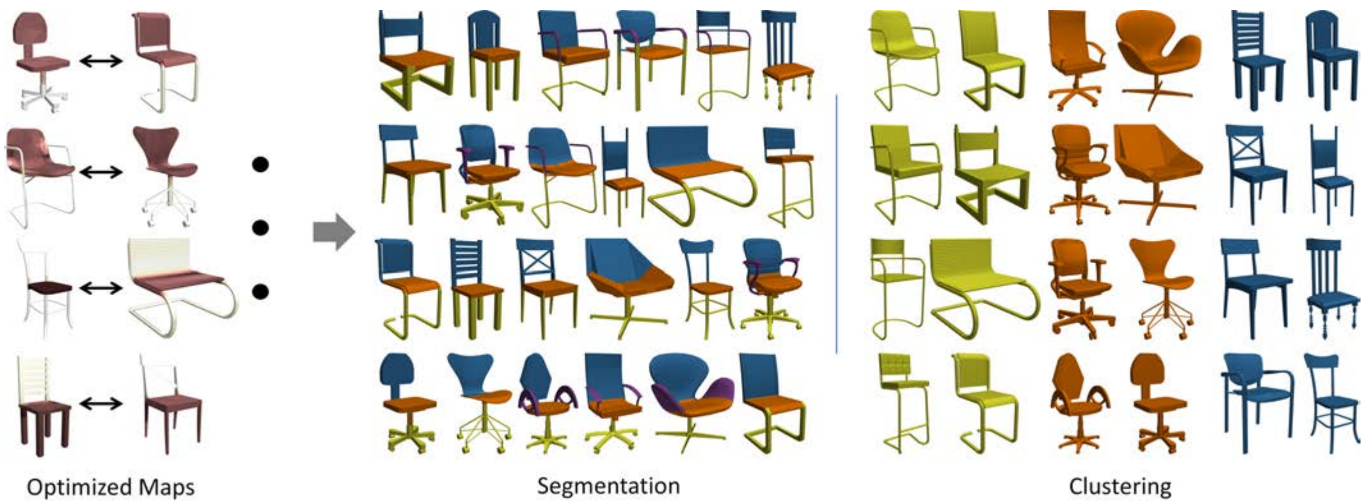

Fig. 1. A prototypical example of a heterogeneous shape model collection with highly consistent correspondence maps generated from our tensor approach and optimized across the entire collection. These maps simultaneously encodes structural and functional similarities and variances. The corresponding regions of each object in the shape collections is shown with matching color. The induced shape maps additionally enable applications in shape segmentation and sub-region weighted shape co-clustering.

Establishing high-quality correspondence maps between geometric shapes has been shown to be the fundamental problem in managing geometric shape collections. Prior work has focused on computing efficient maps between pairs of shapes, and has shown a quantifiable benefit of joint map synchronization, where a collection of shapes are used to improve (denoise) the pairwise maps for consistency and correctness. However, these existing map synchronization techniques place very strong assumptions on the input shapes collection such as all the input shapes fall into the same category and/or the majority of the input pairwise maps are correct. In this paper, we present a multiple map synchronization approach that takes a heterogeneous shape collection as input and simultaneously outputs consistent dense pairwise shape maps. We achieve our goal by using a novel tensor-based representation for map synchronization, which is efficient and robust than all prior matrix-based representations. We demonstrate the usefulness of this approach across a wide range of geometric shape datasets and the applications in shape clustering and shape co-segmentation.

CCS Concepts: • **Computer systems organization** → **Embedded systems**; *Redundancy*; Robotics; • **Networks** → Network reliability;

Additional Key Words and Phrases: machine learning, numerical analysis, optimization, object scanning/acquisition, shape analysis and shape matching and retrieval

ACM Reference Format:
Qixing Huang, Zhenxiao Liang, Haoyun Wang, Simiao Zuo, and Chandrajit Bajaj. 2019. Tensor Maps for Synchronizing Heterogeneous Shape Collections. *ACM Trans. Graph.* 38, 4, Article 106 (July 2019), 18 pages. https://doi.org/10.1145/3306346.3322944

Authors' addresses: Qixing Huang, The University of Texas at Austin, 2317 Speedway, Austin, TX, 78712, USA; Zhenxiao Liang, The University of Texas at Austin, 2317 Speedway, Austin, TX, 78712, USA; Haoyun Wang, Tsinghua University, 2317 Speedway, Austin, TX, 78712, USA; Simiao Zuo, The University of Texas at Austin, 2317 Speedway, Austin, TX, 78712, USA; Chandrajit Bajaj, The University of Texas at Austin, 2317 Speedway, Austin, TX, 78712, USA.

0730-0301/2019/7-ART106 $15.00
https://doi.org/10.1145/3306346.3322944

1 INTRODUCTION

Digital shape collections are a rich resource of information for diverse data driven applications. Developing effective tools to analyze and organize them is a central research problem in geometry processing and machine learning. Prior related papers have focused on computing shape correspondence maps across all pairs of shapes in

Stylizing Video by Example

ONDŘEJ JAMRIŠKA, Czech Technical University in Prague, Faculty of Electrical Engineering
ŠÁRKA SOCHOROVÁ, Czech Technical University in Prague, Faculty of Electrical Engineering
ONDŘEJ TEXLER, Czech Technical University in Prague, Faculty of Electrical Engineering
MICHAL LUKÁČ, Adobe Research
JAKUB FIŠER, Adobe Research
JINGWAN LU, Adobe Research
ELI SHECHTMAN, Adobe Research
DANIEL SÝKORA, Czech Technical University in Prague, Faculty of Electrical Engineering

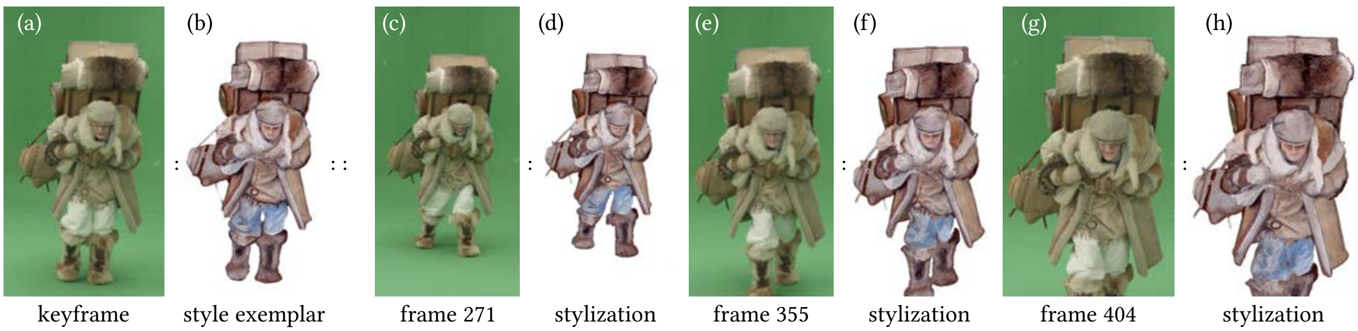

Fig. 1. An example of a stylized sequence produced by our approach. One frame from the sequence is selected as a keyframe (a) and a corresponding style exemplar is painted using watercolor (b). Then, for the rest of the sequence (c, e, g) our technique produces stylized output (d, f, h) which preserves the artistic attributes of the specified style exemplar, reflects structural changes in the target video, and maintains temporal coherence. Video frames (a, c, e, g) courtesy of © MAUR film, style exemplar (b) courtesy of © Pavla Sýkorová, used with permission.

We introduce a new example-based approach to video stylization, with a focus on preserving the visual quality of the style, user controllability and applicability to arbitrary video. Our method gets as input one or more keyframes that the artist chooses to stylize with standard painting tools. It then automatically propagates the stylization to the rest of the sequence. To facilitate this while preserving visual quality, we developed a new type of guidance for state-of-art patch-based synthesis, that can be applied to any type of video content and does not require any additional information besides the video itself and a user-specified mask of the region to be stylized. We further show a temporal blending approach for interpolating style between keyframes that preserves texture coherence, contrast and high frequency details. We evaluate our method on various scenes from real production setting and provide a thorough comparison with prior art.

CCS Concepts: • **Computing methodologies → Motion processing**; **Image processing**.

Additional Key Words and Phrases: style transfer

ACM Reference Format:
Ondřej Jamriška, Šárka Sochorová, Ondřej Texler, Michal Lukáč, Jakub Fišer, Jingwan Lu, Eli Shechtman, and Daniel Sýkora. 2019. Stylizing Video by Example. *ACM Trans. Graph.* 38, 4, Article 107 (July 2019), 11 pages. https://doi.org/10.1145/3306346.3323006

Authors' addresses: Ondřej Jamriška, Czech Technical University in Prague, Faculty of Electrical Engineering, jamriond@fel.cvut.cz; Šárka Sochorová, Czech Technical University in Prague, Faculty of Electrical Engineering, sochosar@fel.cvut.cz; Ondřej Texler, Czech Technical University in Prague, Faculty of Electrical Engineering, texleond@fel.cvut.cz; Michal Lukáč, Adobe Research, lukac@adobe.com; Jakub Fišer, Adobe Research, fiser@adobe.com; Jingwan Lu, Adobe Research, jlu@adobe.com; Eli Shechtman, Adobe Research, elishe@adobe.com; Daniel Sýkora, Czech Technical University in Prague, Faculty of Electrical Engineering, sykorad@fel.cvut.cz.

0730-0301/2019/7-ART107 $15.00
https://doi.org/10.1145/3306346.3323006

1 INTRODUCTION

In the past decades, advances in computer graphics led to a revolution in the art of animation, giving birth to an entirely new branch of animation which is three-dimensional, and includes photorealistic lighting effects and physically accurate simulation. Together with lighting, material, and performance capture, the production pipelines of animated video now resemble live-action production more closely than traditional animation. An unfortunate side effect of this is that, due to production and technical considerations, there is a "style gap" between traditional and 3D animation, where the latter has its own distinct look, and it has so far been impossible to convincingly reproduce the look of the former using the aforementioned production pipelines. Currently, there are no automated methods that could use live-action performance capture to produce

Interactive and Automatic Navigation for 360° Video Playback

KYOUNGKOOK KANG, DGIST
SUNGHYUN CHO, DGIST

(a) 360° video (b) Normal Field-of-View (NFoV) video converted from (a) (c) Camera path for converting (a) into (b)

Fig. 1. Interactive and automatic navigation of 360° video*. Our system computes an optimal camera path that shows salient areas in a 360° video (a), and plays a NFoV video based on the path in an online manner. Users can interactively change the viewing direction while watching a video, and the system instantly updates the camera path according to user interaction. The y-axis in (c) is the horizontal angle in a 360° video. The solid blue line in (c) illustrates the camera path computed by our system, and the dotted line indicates the moment when the viewing direction is changed by user interaction. The solid blue line on the right of the dotted line is the updated path after user interaction. The original video is from the Pano2vid dataset [Su et al. 2016], which is originally from Youtube (fydc16RTfCo ©my360planet - Johannes Löffelmann). In the rest of the paper, we simply identify the name of the dataset, and the Youtube ID for each video.

A common way to view a 360° video on a 2D display is to crop and render a part of the video as a normal field-of-view (NFoV) video. While users can enjoy natural-looking NFoV videos using this approach, they need to constantly make manual adjustment of the viewing direction not to miss interesting events in the video. In this paper, we propose an interactive and automatic navigation system for comfortable 360° video playback. Our system finds a virtual camera path that shows the most salient areas through the video, generates a NFoV video based on the path, and plays it in an online manner. A user can interactively change the viewing direction while watching a video, and the system instantly updates the path reflecting the intention of the user. To enable online processing, we design our system consisting of an offline pre-processing step, and an online 360° video navigation step. The pre-processing step computes optical flow and saliency scores for an input video. Based on these, the online video navigation step computes an optimal camera path reflecting user interaction, and plays a NFoV video in an online manner. For improved user experience, we also introduce optical flow-based camera path planning, saliency-aware path update, and adaptive control of the temporal window size. Our experimental results including user studies show that our system provides more pleasant experience of watching 360° videos than existing approaches.

CCS Concepts: • **Computing methodologies** → **Computational photography**.

Additional Key Words and Phrases: 360° video, spherical panorama, user interaction, 360° video navigation

Authors' addresses: Kyoungkook Kang, DGIST, kkang@dgist.ac.kr; Sunghyun Cho, DGIST, scho@dgist.ac.kr.

0730-0301/2019/7-ART108 $15.00
https://doi.org/10.1145/3306346.3323046

ACM Reference Format:
Kyoungkook Kang and Sunghyun Cho. 2019. Interactive and Automatic Navigation for 360° Video Playback. *ACM Trans. Graph.* 38, 4, Article 108 (July 2019), 11 pages. https://doi.org/10.1145/3306346.3323046

1 INTRODUCTION

360° videos that record all directions at once are recently gaining popularity and getting widely available with the recent advent of virtual reality applications. Facebook and YouTube now support 360° videos so that users can easily watch 360° videos on mobile devices and computers. 360° cameras such as Samsung Gear 360, LG 360, and GoPro Fusion 360 have been rapidly emerging, enabling even casual users to easily create 360° videos. Moreover, commercially available virtual reality headsets such as Oculus Rift, Samsung Gear VR, HTC Vive, and PlayStation VR are getting popular too, accelerating production of new 360° video contents.

Because of the nature of 360° videos that record all directions at once, the most comfortable way for a viewer to watch such a video would be to wear a head-mounted display (HMD), and turn his/her head and body around to watch different directions. However, HMDs are not always available, and it is cumbersome to wear a HMD to watch a video every time. A more common scenario is to view a 360° video on a 2D display such as a computer screen or a smartphone.

One way to view a 360° video on a 2D display is to project the entire video using equirectangular projection so that a viewer can watch all directions at once. However, equirectangular projection introduces severe geometrical distortions (Fig. 1(a)). A more popular way widely used in many applications such as YouTube is to crop a part of a video, and render it as a normal field-of-view (NFoV) video. While watching a video, a viewer can manually change the viewing direction, e.g., by dragging a mouse or by orienting his/her phone. Unfortunately, it is uncomfortable either because a viewer needs to constantly adjust the viewing direction as the positions of

Computational Design of Fabric Formwork

XIAOTING ZHANG, Boston University, USA
GUOXIN FANG, Delft University of Technology, The Netherlands
MELINA SKOURAS, Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, France
GWENDA GIESELER, Delft University of Technology, The Netherlands
CHARLIE C.L. WANG, The Chinese University of Hong Kong, China
EMILY WHITING, Boston University, USA

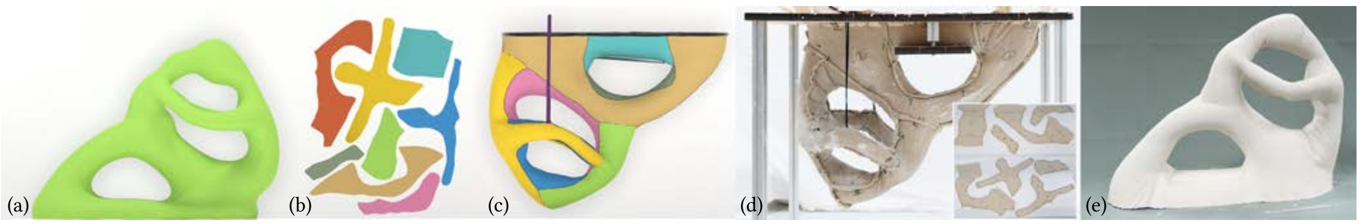

Fig. 1. A fertility model designed and fabricated using our computational approach. For a target 3D model (a), our system can automatically compute a set of flat panels (b) that can be sewn together to serve as fabric containers to form a target shape by pressure of liquid plaster poured in – see (c) for the simulation under force equilibrium of membrane tension, liquid pressure and external supports. The generated flat panels are used to conduct the physical fabrication of fabric formwork (d). After drying and unwrapping the fabric container, a sculpture with the designed target shape has been fabricated (e).

We present an inverse design tool for fabric formwork – a process where flat panels are sewn together to form a fabric container for casting a plaster sculpture. Compared to 3D printing techniques, the benefit of fabric formwork is its properties of low-cost and easy transport. The process of fabric formwork is akin to molding and casting but having a soft boundary. Deformation of the fabric container is governed by force equilibrium between the pressure forces from liquid fill and tension in the stretched fabric. The final result of fabrication depends on the shapes of the flat panels, the fabrication orientation and the placement of external supports. Our computational framework generates optimized flat panels and fabrication orientation with reference to a target shape, and determines effective locations for external supports. We demonstrate the function of this design tool on a variety of models with different shapes and topology. Physical fabrication is also demonstrated to validate our approach.

CCS Concepts: • **Computing methodologies → Shape modeling**; *Modeling and simulation*; • **Applied computing → Computer-aided design**.

Additional Key Words and Phrases: computational design, fabric formwork, shape optimization, fabrication, casting

Authors' addresses: Xiaoting Zhang; Emily Whiting, Department of Computer Science, Boston University, USA; Guoxin Fang; Gwenda Gieseler, Department of Design Engineering, Delft University of Technology, the Netherlands; Melina Skouras, Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, France; Charlie C.L. Wang, Department of Mechanical and Automation Engineering, The Chinese University of Hong Kong, China.

0730-0301/2019/7-ART109 $15.00
https://doi.org/10.1145/3306346.3322988

ACM Reference Format:
Xiaoting Zhang, Guoxin Fang, Melina Skouras, Gwenda Gieseler, Charlie C.L. Wang, and Emily Whiting. 2019. Computational Design of Fabric Formwork. *ACM Trans. Graph.* 38, 4, Article 109 (July 2019), 13 pages. https://doi.org/10.1145/3306346.3322988

1 INTRODUCTION

Despite rapid advances in 3D printing, fabricating large, durable and high quality objects is still impractical for all but highly trained experts. Although 3D printing has been demonstrated experimentally to produce house-sized structures with concrete [Khoshnevis 2004] and standard thermoplastics [Bogue 2013], these fabrication methods require enormous time, financial investment, and non-standard equipment such as oversized gantries. Traditional casting processes with rigid molds can capture detailed geometry, but present restrictive constraints to allow mold removal, and molds are similarly expensive to create for large models [Groover 2011].

The goal of this paper is to lay the algorithmic foundation for large scale fabrication with conventional, low cost materials. We propose a fundamentally different approach to large scale construction through computational design of *fabric formwork*. The fabric formwork process is akin to molding and casting. Using the flexibility of a textile membrane as a container (typically nylon, polyester, or polypropylene), fabric formwork provides a method to form natural tension geometries [West 2016]. Containers are fabric panels attached at seams by sewing. Wet concrete is poured into the fabric container, the fluid pressure and membrane boundary then work together to produce the structural shape after solidification. The deformation of every point on the panel is governed by force equilibrium: the pressure force from the liquid concrete balanced by tension in the stretched textile. The benefit is a low-cost and easily transportable fabrication process. The fabric panels can be created

Volume-Aware Design of Composite Molds

THOMAS ALDERIGHI, Università di Pisa and ISTI-CNR, Italy
LUIGI MALOMO, ISTI-CNR, Italy
DANIELA GIORGI, ISTI-CNR, Italy
BERND BICKEL, IST, Austria
PAOLO CIGNONI, ISTI-CNR, Italy
NICO PIETRONI, University of Technology Sidney, Australia and ISTI-CNR, Italy

(a) (b) (c)

Fig. 1. Our technique for casting extremely complex shapes up to (a) objects made up of multiple entangled pieces. (b) We design a cut layout in the mold volume, which is defined by a parting surface membrane (blue) and a set of additional membranes to enforce moldability (red). (c) We assemble a composite mold made up of a hard plastic shell and a soft silicone part, and then fabricate the given object using liquid casting techniques.

We propose a novel technique for the automatic design of molds to cast highly complex shapes. The technique generates composite, two-piece molds. Each mold piece is made up of a hard plastic shell and a flexible silicone part. Thanks to the thin, soft, and smartly shaped silicone part, which is kept in place by a hard plastic shell, we can cast objects of unprecedented complexity. An innovative algorithm based on a volumetric analysis defines the layout of the internal cuts in the silicone mold part. Our approach can robustly handle thin protruding features and intertwined topologies that have caused previous methods to fail. We compare our results with state of the art techniques, and we demonstrate the casting of shapes with extremely complex geometry.

CCS Concepts: • **Computing methodologies** → *Shape modeling*.

Additional Key Words and Phrases: fabrication, mold design, casting

ACM Reference Format:
Thomas Alderighi, Luigi Malomo, Daniela Giorgi, Bernd Bickel, Paolo Cignoni, and Nico Pietroni. 2019. Volume-Aware Design of Composite Molds. *ACM Trans. Graph.* 38, 4, Article 110 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3322981

Authors' addresses: Thomas Alderighi, Università di Pisa , ISTI-CNR, Italy; Luigi Malomo, ISTI-CNR, Italy; Daniela Giorgi, ISTI-CNR, Italy; Bernd Bickel, IST, Austria; Paolo Cignoni, ISTI-CNR, Italy; Nico Pietroni, University of Technology Sidney, Australia , ISTI-CNR, Italy.

0730-0301/2019/7-ART110 $15.00
https://doi.org/10.1145/3306346.3322981

1 INTRODUCTION

Casting is a well-established manufacturing technique with many appealing properties: a wide range of materials, such as concrete, plaster, plastic resins, or edibles (e.g., chocolate) can be used, and re-usable molds reduce production costs, while allowing for the efficient reproduction of multiple copies. However, the geometric complexity of shapes that are fabricable by casting is much more limited in comparison to 3D printing, and a significant amount of effort has to be invested to design moldable shapes and the mold itself.

Recently, a new generation of computational design algorithms have been introduced to support the creation of re-usable molds [Alderighi et al. 2018; Herholz et al. 2015; Malomo et al. 2016; Nakashima et al. 2018]. These approaches opened up new avenues for reproducing shapes, by simplifying the production process and making the design workflow accessible to non-experts. This has been an essential step towards democratizing molding; however, reproducing complex free-form geometry by casting remains challenging and is constrained by the available computational design tools.

In this paper, we propose a novel method for the automatic design of composite molds for objects of complex geometry and topology (Figure 1.a). Our molds are made up of an elastic silicone part surrounded by a hard plastic shell (Figure 1.c). In art reproduction and

Geometry-Aware Scattering Compensation for 3D Printing

DENIS SUMIN*, Max Planck Institute for Informatics, Germany
TOBIAS RITTIG*, Charles University, Czech Republic
VAHID BABAEI, Max Planck Institute for Informatics, Germany
THOMAS NINDEL, Charles University, Czech Republic
ALEXANDER WILKIE, Charles University, Czech Republic
PIOTR DIDYK, Università della Svizzera italiana, Switzerland
BERND BICKEL, IST Austria, Austria
JAROSLAV KŘIVÁNEK, Charles University, Czech Republic
KAROL MYSZKOWSKI, Max Planck Institute for Informatics, Germany
TIM WEYRICH, University College London, United Kingdom

Fig. 1. A demonstration of 3D printouts obtained with our method. The Earth diameter is 5 cm.

Commercially available full-color 3D printing allows for detailed control of material deposition in a volume, but an exact reproduction of a target surface appearance is hampered by the strong subsurface scattering that causes nontrivial volumetric cross-talk at the print surface. Previous work showed how an iterative optimization scheme based on accumulating absorptive materials at the surface can be used to find a volumetric distribution of print materials that closely approximates a given target appearance.

In this work, we first revisit the assumption that pushing the absorptive materials to the surface results in minimal volumetric cross-talk. We design a full-fledged optimization on a small domain for this task and confirm this previously reported heuristic. Then, we extend the above approach that is critically limited to color reproduction on planar surfaces, to arbitrary 3D shapes. Our method enables high-fidelity color texture reproduction on 3D prints by effectively compensating for internal light scattering within arbitrarily shaped objects. In addition, we propose a content-aware gamut mapping that significantly improves color reproduction for the pathological case of thin geometric features. Using a wide range of sample objects with complex textures and geometries, we demonstrate color reproduction whose fidelity is superior to state-of-the-art drivers for color 3D printers.

CCS Concepts: • **Computing methodologies → Reflectance modeling**; **Volumetric models**; • **Applied computing → Computer-aided manufacturing**.

Additional Key Words and Phrases: computational fabrication, appearance reproduction, appearance enhancement, sub-surface light transport, volumetric optimization, gradient rendering

ACM Reference Format:
Denis Sumin, Tobias Rittig, Vahid Babaei, Thomas Nindel, Alexander Wilkie, Piotr Didyk, Bernd Bickel, Jaroslav Křivánek, Karol Myszkowski, and Tim Weyrich. 2019. Geometry-Aware Scattering Compensation for 3D Printing. *ACM Trans. Graph.* 38, 4, Article 111 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3322992

*Denis Sumin and Tobias Rittig share the first authorship of this work.

Authors' addresses: Denis Sumin, Max Planck Institute for Informatics, Saarbrücken, Germany, denis.sumin@mpi-inf.mpg.de; Tobias Rittig, Charles University, Prague, Czech Republic, tobias@cgg.mff.cuni.cz; Vahid Babaei, Max Planck Institute for Informatics, Saarbrücken, Germany; Thomas Nindel, Charles University, Prague, Czech Republic; Alexander Wilkie, Charles University, Prague, Czech Republic; Piotr Didyk, Università della Svizzera italiana, Lugano, Switzerland; Bernd Bickel, IST Austria, Maria Gugging, Austria; Jaroslav Křivánek, Charles University, Prague, Czech Republic; Karol Myszkowski, Max Planck Institute for Informatics, Saarbrücken, Germany; Tim Weyrich, University College London, London, United Kingdom.

0730-0301/2019/7-ART111 $15.00
https://doi.org/10.1145/3306346.3322992

LayerCode: Optical Barcodes for 3D Printed Shapes

HENRIQUE TELES MAIA, Columbia University
DINGZEYU LI, Adobe Research
YUAN YANG, Columbia University
CHANGXI ZHENG, Columbia University

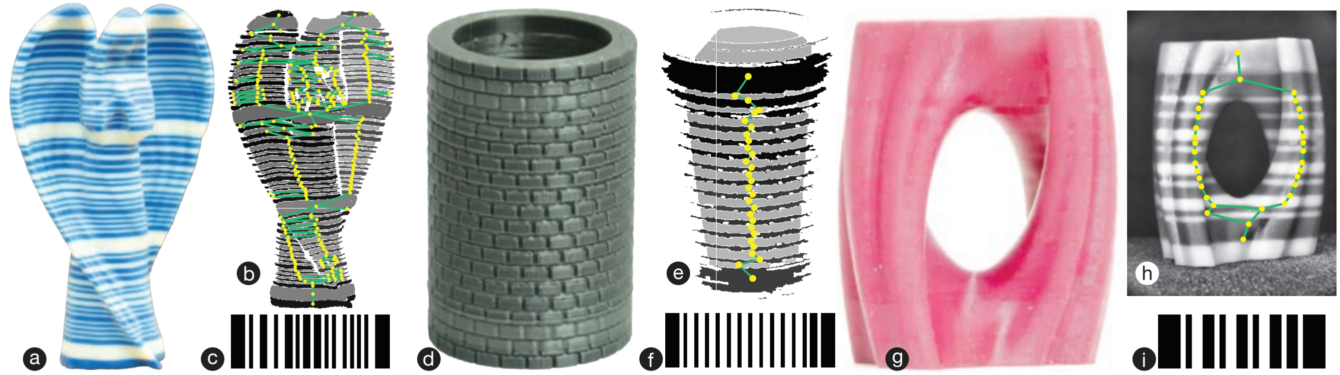

Fig. 1. **LayerCode tags** are deployed in 3D printed objects through two-color printing (a), variable layer heights (d), and near-infrared steganography (g). In the first case (a), the LayerCode tag is visible; in the second (d), the tag is less visible; and in the third (g) it is completely invisible, but still machine-readable. Just like reading a barcode, we capture an image of each object, and our decoding algorithm processes the image to create a decoding graph (b, e, h), from which a linear barcode is recovered (c, f, i). In this case, the corresponding LayerCode bit string reveals a 24-bit code repeated 3 times in (a), a 24-bit code repeated once in (d), and a 12-bit code repeated once in (g).

With the advance of personal and customized fabrication techniques, the capability to embed information in physical objects becomes evermore crucial. We present *LayerCode*, a tagging scheme that embeds a carefully designed barcode pattern in 3D printed objects as a deliberate byproduct of the 3D printing process. The LayerCode concept is inspired by the structural resemblance between the parallel black and white bars of the standard barcode and the universal layer-by-layer approach of 3D printing. We introduce an encoding algorithm that enables the 3D printing layers to carry information without altering the object geometry. We also introduce a decoding algorithm that reads the LayerCode tag of a physical object by just taking a photo. The physical deployment of LayerCode tags is realized on various types of 3D printers, including Fused Deposition Modeling printers as well as Stereolithography based printers. Each offers its own advantages and tradeoffs. We show that LayerCode tags can work on complex, nontrivial shapes, on which all previous tagging mechanisms may fail. To evaluate LayerCode thoroughly, we further stress test it with a large dataset of complex shapes using virtual rendering. Among 4,835 tested shapes, we successfully encode and decode on more than 99% of the shapes.

Authors' addresses: Henrique Teles Maia, Columbia University, henrique@cs.columbia.edu; Dingzeyu Li, Adobe Research, dinli@adobe.com; Yuan Yang, Columbia University, yy2664@columbia.edu; Changxi Zheng, Columbia University, cxz@cs.columbia.edu.

0730-0301/2019/7-ART1 $15.00
https://doi.org/10.1145/3306346.3322960

CCS Concepts: • **Hardware → Emerging interfaces**; • **Mathematics of computing → Graph algorithms**;

Additional Key Words and Phrases: 3D printing, information embedding, fabrication, physical hyperlinks

ACM Reference Format:
Henrique Teles Maia, Dingzeyu Li, Yuan Yang, and Changxi Zheng. 2019. LayerCode: Optical Barcodes for 3D Printed Shapes. *ACM Trans. Graph.* 38, 4, Article 1 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3322960

1 INTRODUCTION

Invented 45 years ago, the optical barcode has become an indispensable minutiae in today's digital era. The design is simple, e.g. black and white bars printed on a flat surface, but its use is ubiquitous. From package delivery and airplane boarding to inventory management and patient identification, the barcode serves as a link that bridges physical artifacts to modern digital systems.

In this work, we rethink barcodes in the context of additive manufacturing, popularly known as 3D printing. 3D printing offers a quick way of making customized, complex shaped objects. Unlike a mass-produced product which by design has a reserved flat surface region to host barcodes, 3D printed shapes are often complex and curved: thin features, slender threads, and holes are not uncommon. As a result, traditional barcodes cannot be placed on such objects.

Recent years have seen a few approaches proposed toward embedding optical tags in 3D printed objects, on the surface [Kikuchi et al. 2018], beneath the surface [Li et al. 2017] and inside the objects [Willis and Wilson 2013]. However, these approaches either require specialized (and expensive) hardware to read the tags or

Direct Delta Mush Skinning and Variants

BINH HUY LE, SEED - Electronic Arts
JP LEWIS, Google AI

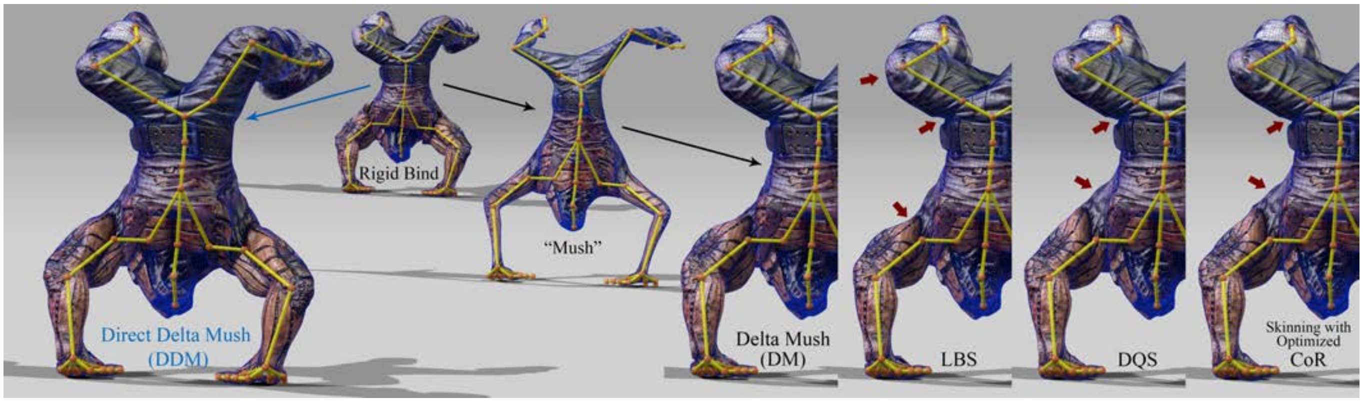

Fig. 1. The skinned model (left) is produced directly from the “unrigged” rigid bind model using our Direct Delta Mush algorithm. DDM can produce equivalent results to the Delta Mush algorithm but uses a direct local computation rather than the iterated global “mush” runtime smoothing of DM. The DM and DDM algorithms both provide greatly simplified authoring. They do not have the bulge and cleft artifacts common to other methods, which are prominent in the under-arm and hip regions (respectively) in this example (red arrows). DDM offers further advantages over DM, as described in the paper.

A significant fraction of the world’s population have experienced virtual characters through games and movies, and the possibility of online VR social experiences may greatly extend this audience. At present, the skin deformation for interactive and real-time characters is typically computed using geometric skinning methods. These methods are efficient and simple to implement, but obtaining quality results requires considerable manual “rigging" effort involving trial-and-error weight painting, the addition of virtual helper bones, etc. The recently introduced Delta Mush algorithm largely solves this rig authoring problem, but its iterative computational approach has prevented direct adoption in real-time engines.

This paper introduces Direct Delta Mush, a new algorithm that simultaneously improves on the efficiency and control of Delta Mush while generalizing previous algorithms. Specifically, we derive a direct rather than iterative algorithm that has the same ballpark computational form as some previous geometric weight blending algorithms. Straightforward variants of the algorithm are then proposed to further optimize computational and storage cost with insignificant quality losses. These variants are equivalent to special cases of several previous skinning algorithms.

Our algorithm simultaneously satisfies the goals of reasonable efficiency, quality, and ease of authoring. Further, its explicit decomposition of rotational and translational effects allows independent control over bending versus twisting deformation, as well as a skin sliding effect.

Authors’ addresses: Binh Huy Le, SEED - Electronic Arts, Redwood City, CA, bbinh85@gmail.com; JP Lewis, Google AI, San Francisco, CA, noisebrain@gmail.com.

0730-0301/2019/7-ART113 $15.00
https://doi.org/10.1145/3306346.3322982

CCS Concepts: • **Computing methodologies → Animation**.

Additional Key Words and Phrases: skinning, skeletal animation, delta mush, real time, deformation, character animation

ACM Reference Format:
Binh Huy Le and JP Lewis. 2019. Direct Delta Mush Skinning and Variants. *ACM Trans. Graph.* 38, 4, Article 113 (July 2019), 13 pages. https://doi.org/10.1145/3306346.3322982

1 INTRODUCTION

Typically characters are the main focus of any movie or game. Major characters are often humans or animals, and thus are articulated models with rigid bones underlying deformable flesh and skin. Other objects in the scene such as trees can deform and may also be represented with a similar underlying approach. A key focus in all these cases is getting the deformation right.

A character deformation method suitable for games and interactive applications such as animation should have the following characteristics: (1) speed, (2) quality, (3) simplicity of setup and authoring. Existing approaches to character deformation can be very broadly classified into geometric skinning and simulation approaches. Simulation approaches produce the highest quality but may be less suitable in terms of criteria (1) and (3). Regarding speed, simulation effects are not justified when nearly the same effect can be produced with a cheaper method. It should be remembered that character deformation is just one of many things that must be computed within the frame interval at typical frame rates of 24fps (movie animation), 60fps (games) or 120fps (VR). Other tasks include various rendering steps, gameplay AI, collision detection, other types of physics, etc. Simulation approaches are also not ideal in terms of simplicity. The rig may require constructing additional components

NeuroSkinning: Automatic Skin Binding for Production Characters with Deep Graph Networks

LIJUAN LIU, NetEase Fuxi AI Lab
YOUYI ZHENG, State Key Lab of CAD&CG, Zhejiang University
DI TANG, NetEase Fuxi AI Lab
YI YUAN, NetEase Fuxi AI Lab
CHANGJIE FAN, NetEase Fuxi AI Lab
KUN ZHOU, State Key Lab of CAD&CG, Zhejiang University

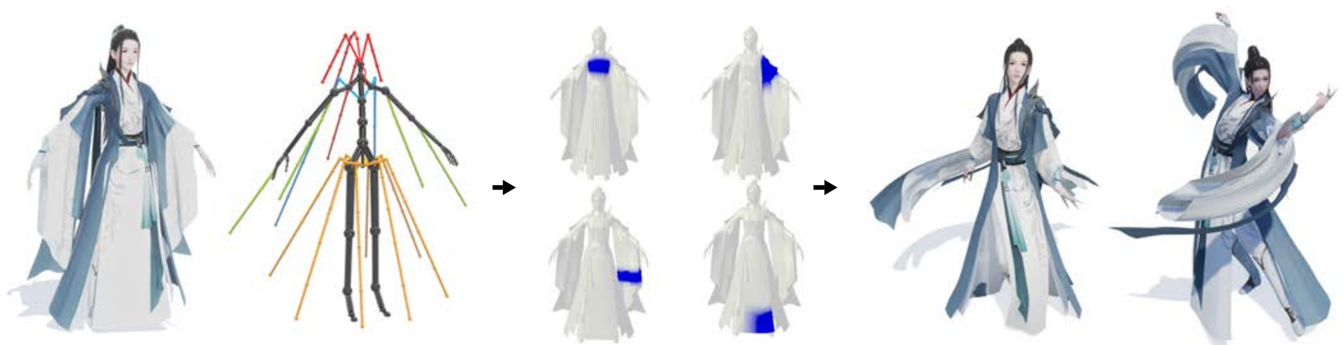

Fig. 1. Given a character mesh and its associated skeleton hierarchy in rest pose (left), our method automatically predicts skin weights (middle), which can produce high-quality deformations comparable to those generated by artist-painted weights (right).

We present a deep-learning-based method to automatically compute skin weights for skeleton-based deformation of production characters. Given a character mesh and its associated skeleton hierarchy in rest pose, our method constructs a graph for the mesh, each node of which encodes the mesh-skeleton attributes of a vertex. An end-to-end deep graph convolution network is then introduced to learn the mesh-skeleton binding patterns from a set of character models with skin weights painted by artists. The network can be used to predict the skin weight map for a new character model, which describes how the skeleton hierarchy influences the mesh vertices during deformation. Our method is designed to work for non-manifold meshes with multiple disjoint or intersected components, which are common in game production and require complex skeleton hierarchies for animation control. We tested our method on the datasets of two commercial games. Experiments show that the predicted skin weight maps can be readily applied to characters in the production pipeline to generate high-quality deformations.

*Corresponding authors: Youyi Zheng (youyizheng@zju.edu.cn) and Kun Zhou (kunzhou@acm.org).
Authors' addresses: Lijuan Liu, NetEase Fuxi AI Lab; Youyi Zheng, State Key Lab of CAD&CG, Zhejiang University, Hangzhou, 310058; Di Tang, NetEase Fuxi AI Lab; Yi Yuan, NetEase Fuxi AI Lab; Changjie Fan, NetEase Fuxi AI Lab; Kun Zhou, State Key Lab of CAD&CG, Zhejiang University, Hangzhou, 310058.

0730-0301/2019/7-ART114 $15.00
https://doi.org/10.1145/3306346.3322969

CCS Concepts: • **Computing methodologies → Neural networks**; **Animation**.

Additional Key Words and Phrases: Animation, Skinning, Deep Learning

ACM Reference Format:
Lijuan Liu, Youyi Zheng, Di Tang, Yi Yuan, Changjie Fan, and Kun Zhou. 2019. NeuroSkinning: Automatic Skin Binding for Production Characters with Deep Graph Networks. *ACM Trans. Graph.* 38, 4, Article 114 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3322969

1 INTRODUCTION

Skinning techniques, such as Linear Blend Skinning (LBS) [Kavan and Žára 2005] and Dual Quaternion Skinning (DQS) [Kavan et al. 2007], are widely used in game production and supported by standard modeling software. In a typical skinning pipeline, an artist first creates a mesh model and specifies a skeleton hierarchy for the model. Skin weights are then painted onto the mesh to indicate how the mesh vertices deform with the skeleton. For complex character models in production, it is a time-consuming process to paint skin weights to produce satisfactory deformations even for professionals. For example, it took more than eight hours for an artist to paint the skin weights for the character shown in Fig. 1.

A few techniques have been proposed to automatically compute skinning weights using heat diffusion [Bang et al. 2015; Baran and Popović 2007; Wareham and Lasenby 2008], elastic deformer [Kavan and Sorkine 2012], or geodesic voxel binding [Dionne and de Lasa 2013]. They either work only for manifold meshes or assume that a skeleton bone has high influence weights on mesh vertices near to it.

Hand Modeling and Simulation Using Stabilized Magnetic Resonance Imaging

BOHAN WANG, University of Southern California, USA
GEORGE MATCUK, University of Southern California, USA
JERNEJ BARBIČ, University of Southern California, USA

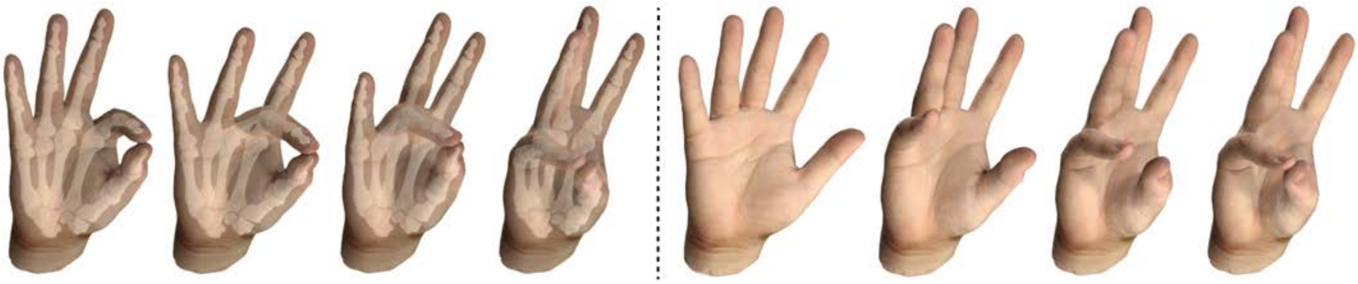

Fig. 1. **Opposition of the thumb:** Our method produces an accurate data-driven skeleton mesh kinematic "rig". Here, we used our rig to reliably reproduce the well-known opposition of the thumb to all other 4 fingers. The solid soft tissue was computed using a FEM simulation attached to the articulated bone meshes. The fingers do not merely touch but also orient to be co-planar at the contact location, subject to biomechanical limits. The right-most four images show representative FEM frames of the opposition between the thumb and the pinky finger.

We demonstrate how to acquire complete human hand bone anatomy (meshes) in multiple poses using magnetic resonance imaging (MRI). Such acquisition was previously difficult because MRI scans must be long for high-precision results (over 10 minutes) and because humans cannot hold the hand perfectly still in non-trivial and badly supported poses. We invent a manufacturing process whereby we use lifecasting materials commonly employed in film special effects industry to generate hand molds, personalized to the subject, and to each pose. These molds are both ergonomic and encasing, and they stabilize the hand during scanning. We also demonstrate how to efficiently segment the MRI scans into individual bone meshes in all poses, and how to correspond each bone's mesh to same mesh connectivity across all poses. Next, we interpolate and extrapolate the MRI-acquired bone meshes to the entire range of motion of the hand, producing an accurate data-driven animation-ready rig for bone meshes. We also demonstrate how to acquire not just bone geometry (using MRI) in each pose, but also a matching highly accurate surface geometry (using optical scanners) in each pose, modeling skin pores and wrinkles. We also give a soft tissue Finite Element Method simulation "rig", consisting of novel tet meshing for stability at the joints, spatially varying geometric and material detail, and quality constraints to the acquired skeleton kinematic rig. Given an animation sequence of hand joint angles, our FEM soft tissue rig produces quality hand surface shapes in arbitrary poses in the hand range of motion. Our results qualitatively reproduce important features seen in the photographs of the subject's hand, such as similar overall organic shape and fold formation.

CCS Concepts: • **Computing methodologies → Physical simulation**.

Additional Key Words and Phrases: hand, MRI, animation, simulation, modeling, anatomy, deformable objects, FEM, optical scanning

ACM Reference Format:
Bohan Wang, George Matcuk, and Jernej Barbič. 2019. Hand Modeling and Simulation Using Stabilized Magnetic Resonance Imaging. *ACM Trans. Graph.* 38, 4, Article 115 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3322983

Authors' addresses: Bohan Wang, University of Southern California, Los Angeles, CA, USA, bohanwan@usc.edu; George Matcuk, University of Southern California, Los Angeles, CA, USA, matcuk@usc.edu; Jernej Barbič, University of Southern California, Los Angeles, CA, USA, jnb@usc.edu.

0730-0301/2019/7-ART115 $15.00
https://doi.org/10.1145/3306346.3322983

1 INTRODUCTION

Hands and their modeling and animation are of paramount importance in many applications. In computer games and film, clothing may occlude the body of the characters, but the hands are often exposed and important for the story. In engineering, hand anatomical models can be used to design tools and equipment. In computer vision, anatomically based modeling can improve the tracking of hands, because it provides better statistical priors on hand shapes. In healthcare, accurate hand shapes and motion can be used to design better finger and partial hand prosthetics, and better tools for surgeons. As is well-known and argued by medical authorities in the field [Kapandji 2009], hand's dexterity stems from a thumb opposing four fingers (Figure 1), making the hand perfectly suited for precise grasping, lifting, climbing and daily manipulations of objects. To improve realism, virtual hands should be modeled similarly to biological hands, and this requires building precise anatomical and kinematic models of real human hands.

Unfortunately, complex hand biomechanics today is not modeled, measured or resolved in any quality way. While there are several computational models of the hand kinematics, few attempted to

Wave-Based Non-Line-of-Sight Imaging using Fast f–k Migration

DAVID B. LINDELL, Stanford University
GORDON WETZSTEIN, Stanford University
MATTHEW O'TOOLE, Carnegie Mellon University

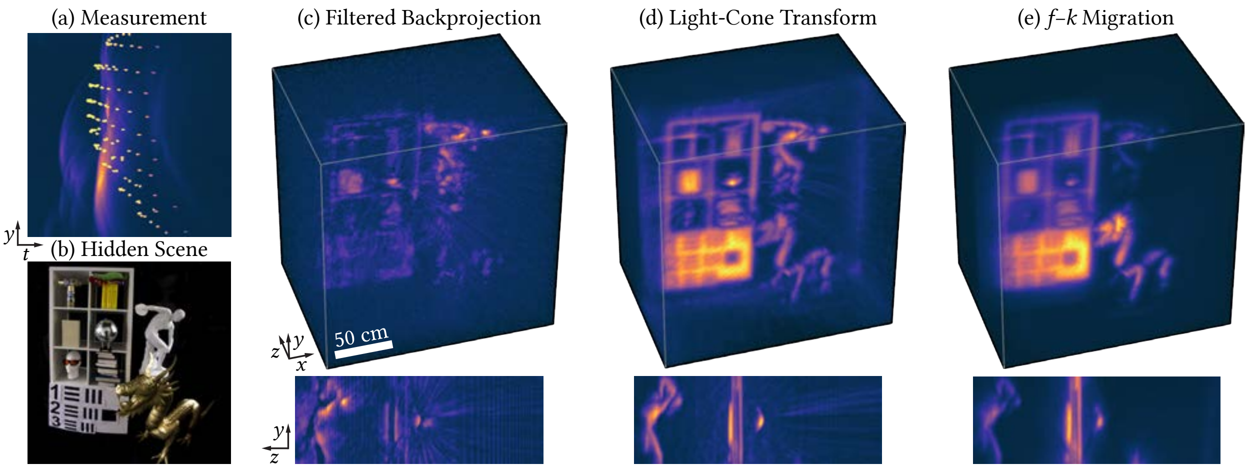

Fig. 1. Non-line-of-sight (NLOS) reconstructions of a hidden, room-sized scene. **(a-b)** One approach to NLOS imaging is to capture time-resolved measurements sampled across a visible surface, and reconstruct the 3D shape and reflectance of the hidden scene. A disco ball produces the bright dots seen in the measurements of indirect light transport (a), and other diffuse and glossy objects produce the streaks. **(c)** Among the various methods for reconstructing shape and reflectance from these measurements, filtered backprojection (FBP) is conceptually one of the simpler methods; it involves a delay-and-sum (backprojection) operation of the time-resolved measurements, followed by a heuristic high-pass filter on the result. **(d)** The light-cone transform (LCT) is a fast reconstruction algorithm that produces more accurate reconstructions in less time, but requires making restrictive assumptions on light transport (*e.g.*, assumes the scene only contains diffuse objects). **(e)** In this paper, we introduce f–k migration, an algorithm that is both fast and versatile, for NLOS imaging. The wave-based nature of this inverse method is unique in being robust to objects with diverse and complex reflectance properties, such as the glossy dragon, the diffuse statue, and the reflective disco ball shown in this scene. All volumes are rendered as maximum intensity projections.

Imaging objects outside a camera's direct line of sight has important applications in robotic vision, remote sensing, and many other domains. Time-of-flight-based non-line-of-sight (NLOS) imaging systems have recently demonstrated impressive results, but several challenges remain. Image formation and inversion models have been slow or limited by the types of hidden surfaces that can be imaged. Moreover, non-planar sampling surfaces and non-confocal scanning methods have not been supported by efficient NLOS algorithms. With this work, we introduce a wave-based image formation model for the problem of NLOS imaging. Inspired by inverse methods used in seismology, we adapt a frequency-domain method, f–k migration, for solving the inverse NLOS problem. Unlike existing NLOS algorithms, f–k migration is both fast and memory efficient, it is robust to specular and other complex reflectance properties, and we show how it can be used with non-confocally scanned measurements as well as for non-planar sampling surfaces. f–k migration is more robust to measurement noise than alternative methods, generally produces better quality reconstructions, and is easy to implement. We experimentally validate our algorithms with a new NLOS imaging system that records room-sized scenes outdoors under indirect sunlight, and scans persons wearing retroreflective clothing at interactive rates.

CCS Concepts: • **Computing methodologies** → *3D imaging*; *Computational photography*.

Additional Key Words and Phrases: computational photography, time-of-flight imaging, non-line-of-sight imaging

ACM Reference Format:
David B. Lindell, Gordon Wetzstein, and Matthew O'Toole. 2019. Wave-Based Non-Line-of-Sight Imaging using Fast f–k Migration. *ACM Trans. Graph.* 38, 4, Article 116 (July 2019), 13 pages. https://doi.org/10.1145/3306346.3322937

Authors' addresses: David B. Lindell, lindell@stanford.edu; Gordon Wetzstein, gordon.wetzstein@stanford.edu, Department of Electrical Engineering, Stanford University, Stanford CA, 94305, USA; Matthew O'Toole, mpotoole@cmu.edu, Robotics Institute & Computer Science Department, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, USA.

0730-0301/2019/7-ART116 $15.00
https://doi.org/10.1145/3306346.3322937

1 INTRODUCTION

Conventional 3D imaging systems based on the time-of-flight principle measure the time it takes a light pulse to travel along a direct path from a source, to a visible object, and back to a sensor. Non-line-of-sight (NLOS) imaging, on the other hand, uses multi-bounce light

Compact Snapshot Hyperspectral Imaging with Diffracted Rotation

DANIEL S. JEON, SEUNG-HWAN BAEK, and SHINYOUNG YI, KAIST
QIANG FU, XIONG DUN, and WOLFGANG HEIDRICH, KAUST
MIN H. KIM, KAIST

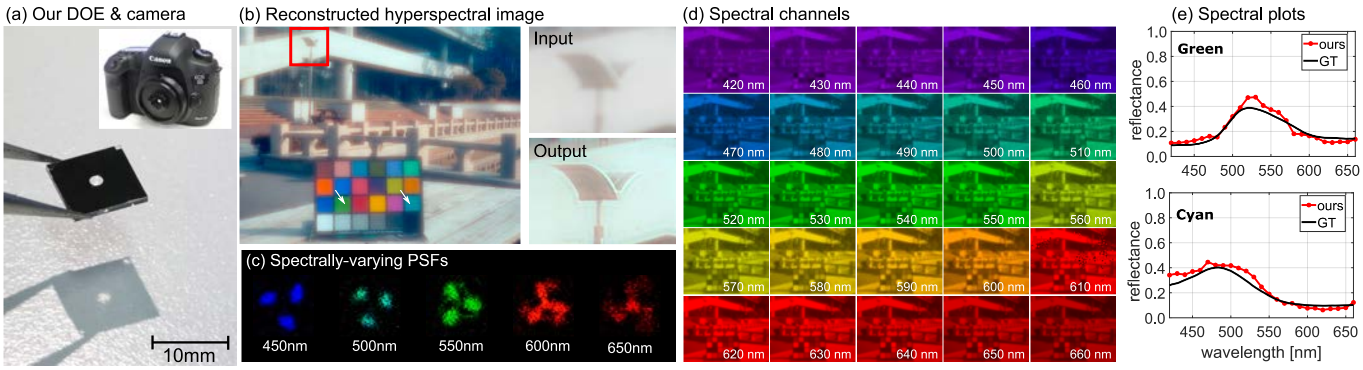

Fig. 1. We propose a compact, diffraction-based snapshot hyperspectral imaging method with a novel diffractive optical element attached to a conventional, bare image sensor. Our method replaces the common optical elements in hyperspectral imaging (prism, coded mask, relay and imaging lenses) with a single optical element. Our single DOE-based camera is coupled with a data-driven spectral reconstruction method that can restore faithful spectral information from spectrally-varying point spread functions. (a) Our fabricated DOE (inset) and a DSLR camera, installed with the DOE for spectral imaging. (b) Reconstructed hyperspectral image from real input. (c) Spectrally-varying PSFs measured per wavelength. (d) Corresponding captured spectral channels. (e) Spectral plots of two patches from the captured ColorChecker, compared to the ground truth.

Traditional snapshot hyperspectral imaging systems include various optical elements: a dispersive optical element (prism), a coded aperture, several relay lenses, and an imaging lens, resulting in an impractically large form factor. We seek an alternative, minimal form factor of snapshot spectral imaging based on recent advances in diffractive optical technology. We thereupon present a compact, diffraction-based snapshot hyperspectral imaging method, using only a novel diffractive optical element (DOE) in front of a conventional, bare image sensor. Our diffractive imaging method replaces the common optical elements in hyperspectral imaging with a single optical element. To this end, we tackle two main challenges: First, the traditional diffractive lenses are not suitable for color imaging under incoherent illumination due to severe chromatic aberration because the size of the point spread function (PSF) changes depending on the wavelength. By leveraging this wavelength-dependent property alternatively for hyperspectral imaging, we introduce a novel DOE design that generates an anisotropic shape of the spectrally-varying PSF. The PSF size remains virtually unchanged, but instead the PSF shape rotates as the wavelength of light changes. Second, since there is no dispersive element and no coded aperture mask, the ill-posedness of spectral reconstruction increases significantly. Thus, we propose an end-to-end network solution based on the unrolled architecture of an optimization procedure with a spatial-spectral prior, specifically designed for deconvolution-based spectral reconstruction. Finally, we demonstrate hyperspectral imaging with a fabricated DOE attached to a conventional DSLR sensor. Results show that our method compares well with other state-of-the-art hyperspectral imaging methods in terms of spectral accuracy and spatial resolution, while our compact, diffraction-based spectral imaging method uses only a single optical element on a bare image sensor.

CCS Concepts: • **Computing methodologies → Hyperspectral imaging**.

Additional Key Words and Phrases: hyperspectral imaging, diffraction

ACM Reference Format:
Daniel S. Jeon, Seung-Hwan Baek, Shinyoung Yi, Qiang Fu, Xiong Dun, Wolfgang Heidrich, and Min H. Kim. 2019. Compact Snapshot Hyperspectral Imaging with Diffracted Rotation. *ACM Trans. Graph.* 38, 4, Article 117 (July 2019), 13 pages. https://doi.org/10.1145/3306346.3322946

1 INTRODUCTION

Hyperspectral imaging has been utilized in various sensing applications, such as biomedical inspection, material classification, material appearance acquisition, digital heritage preservation, forensic science, etc. [Kim and Rushmeier 2011; Kim et al. 2012b, 2014; Nam and Kim 2014]. Based on *geometrical* optics, various hyperspectral imaging systems have been developed for snapshot imaging of dynamic objects and include various optical elements: a dispersive optical element (prism or diffraction grating), a coded aperture mask, several relay lenses, and an objective imaging lens. The dimensions of a typical compressive hyperspectral imager are larger than those of a conventional camera; for instance, its length is greater than a meter [Kim et al. 2012a; Lee and Kim 2014; Lin et al. 2014]. Actual imaging applications are limited to laboratory environments.

Authors' addresses: Daniel S. Jeon; Seung-Hwan Baek; Shinyoung Yi, KAIST, School of Computing, Daejeon, South Korea, 34141; Qiang Fu; Xiong Dun; Wolfgang Heidrich, KAUST, Visual Computing Center, Thuwal, 23955-6900; Min H. Kim, KAIST, School of Computing, Daejeon, South Korea, 34141, corresponding_author:minhkim@kaist.ac.kr.

0730-0301/2019/7-ART117 $15.00
https://doi.org/10.1145/3306346.3322946

Silly Rubber: An Implicit Material Point Method for Simulating Non-equilibrated Viscoelastic and Elastoplastic Solids

YU FANG, University of Pennsylvania
MINCHEN LI, University of Pennsylvania
MING GAO, University of Pennsylvania
CHENFANFU JIANG, University of Pennsylvania

Fig. 1. Our method easily captures the compliant and slowly-recovering behavior of a viscoelastic solid silicone rubber (left), as well as controllable flowing and merging behavior of creamy dense fluid foam (right) in a unified framework.

Simulating viscoelastic polymers and polymeric fluids requires a robust and accurate capture of elasticity and viscosity. The computation is known to become very challenging under large deformations and high viscosity. Drawing inspirations from return mapping based elastoplasticity treatment for granular materials, we present a finite strain integration scheme for general viscoelastic solids under arbitrarily large deformation and non-equilibrated flow. Our scheme is based on a predictor-corrector exponential mapping scheme on the principal strains from the deformation gradient, which closely resembles the conventional treatment for elastoplasticity and allows straightforward implementation into any existing constitutive models. We develop a new Material Point Method that is fully implicit on both elasticity and inelasticity using augmented Lagrangian optimization with various preconditioning strategies for highly efficient time integration. Our method not only handles viscoelasticity but also supports existing elastoplastic models including Drucker-Prager and von-Mises in a unified manner. We demonstrate the efficacy of our framework on various examples showing intricate and characteristic inelastic dynamics with competitive performance.

CCS Concepts: • **Computing methodologies → Physical simulation**.

Additional Key Words and Phrases: Material Point Method (MPM), augmented Lagrangian, viscoelasticity, elastoplasticity

Authors' addresses: Yu Fang, University of Pennsylvania; Minchen Li, University of Pennsylvania; Ming Gao, University of Pennsylvania; Chenfanfu Jiang, University of Pennsylvania.

0730-0301/2019/7-ART118 $15.00
https://doi.org/10.1145/3306346.3322968

ACM Reference Format:
Yu Fang, Minchen Li, Ming Gao, and Chenfanfu Jiang. 2019. Silly Rubber: An Implicit Material Point Method for Simulating Non-equilibrated Viscoelastic and Elastoplastic Solids. *ACM Trans. Graph.* 38, 4, Article 118 (July 2019), 13 pages. https://doi.org/10.1145/3306346.3322968

1 INTRODUCTION

Many real-world materials and natural phenomena involve inelastic deformation. Common examples include metal bending under pressure, sand flowing down a slope, and whipped cream forming a captivating shape on a cake. In computer graphics, since the pioneering work of Terzopoulos et al. [1988; 1987], the study of viscoelasticity, plasticity, and fracture has attracted increasing attention from numerous researchers, not only due to the ubiquity of these behaviors but also due to the complex and intriguing dynamics they exhibit.

Recently, the Material Point Method (MPM) [Sulsky et al. 1995] has gained popularity for inelastic materials due to its flexible geometric representation using meshless particles. Since inelastic objects tend to undergo extreme deformation with permanent shape change from their rest configurations, representing them with meshes as in the Finite Element Method (FEM) imposes the additional difficulties of mesh distortion and tangling. MPM, on the other hand, uses particles to track the history of all strain and stress states and relies on a background grid to accurately evaluate derivatives in force computations. The superior treatment of extreme deformations and topology changes has enabled MPM to be used for simulating many inelastic phenomena in computer graphics, showcased by the successful simulation of materials such as snow [Gaume et al. 2018; Stomakhin et al. 2013], sand [Daviet and Bertails-Descoubes 2016; Klár et al. 2016; Yue et al. 2018], sponge and foam [Ram et al. 2015], and non-Newtonian fluids [Yue et al. 2015].

CD-MPM: Continuum Damage Material Point Methods for Dynamic Fracture Animation

JOSHUAH WOLPER, University of Pennsylvania
YU FANG, University of Pennsylvania
MINCHEN LI, University of Pennsylvania
JIECONG LU, University of Pennsylvania
MING GAO, University of Pennsylvania
CHENFANFU JIANG, University of Pennsylvania

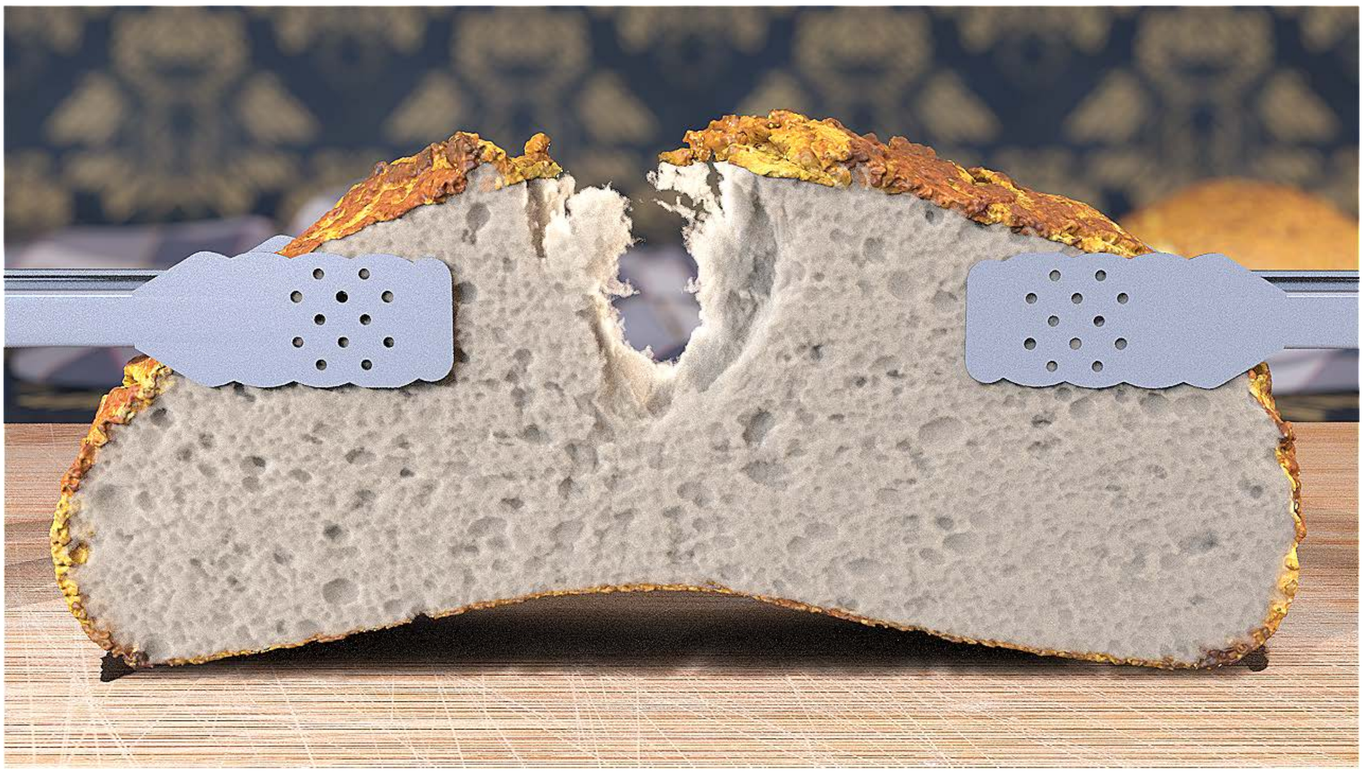

Fig. 1. **Breaking Bread.** We tear a slice of bread with more than 11 million PFF-MPM particles, revealing intricate fracture patterns and natural dynamics.

We present two new approaches for animating dynamic fracture involving large elastoplastic deformation. In contrast to traditional mesh-based techniques, where sharp discontinuity is introduced to split the continuum at crack surfaces, our methods are based on Continuum Damage Mechanics (CDM) with a variational energy-based formulation for crack evolution. Our first approach formulates the resulting dynamic material damage evolution with a Ginzburg-Landau type phase-field equation and discretizes it with the Material Point Method (MPM), resulting in a coupled momentum/damage solver rooted in phase field fracture: PFF-MPM. Although our PFF-MPM approach achieves convincing fracture with or without plasticity, we also introduce a return mapping algorithm that can be analytically solved for a wide range of general non-associated plasticity models, achieving more than two times speedup over traditional iterative approaches. To demonstrate the efficacy of the algorithm, we also develop a Non-Associated Cam-Clay (NACC) plasticity model with a novel fracture-friendly hardening scheme. Our NACC plasticity paired with traditional MPM composes a second approach to dynamic fracture, as it produces a breadth of organic, brittle

Authors' addresses: Joshuah Wolper, University of Pennsylvania; Yu Fang, University of Pennsylvania; Minchen Li, University of Pennsylvania; Jiecong Lu, University of Pennsylvania; Ming Gao, University of Pennsylvania; Chenfanfu Jiang, University of Pennsylvania.

0730-0301/2019/7-ART119 $15.00
https://doi.org/10.1145/3306346.3322949

Implicit Untangling: A Robust Solution for Modeling Layered Clothing

THOMAS BUFFET, Inria Rhônes-Alpes, France and LIX, Ecole Polytechnique, CNRS, IP Paris, France
DAMIEN ROHMER, LIX, Ecole Polytechnique, CNRS, IP Paris, France
LOÏC BARTHE, University of Toulouse - IRIT - CNRS, France
LAURENCE BOISSIEUX, Inria Rhônes-Alpes, France
MARIE-PAULE CANI, LIX, Ecole Polytechnique, CNRS, IP Paris, France

Fig. 1. Our method is able, from a set of garments possibly exhibiting deep interpenetration, to compute an untangled state, ie. a guaranteed intersection-free configuration, even while considering the thicknesses of cloth layers. Animations can then be launched.

We propose a robust method for untangling an arbitrary number of cloth layers, possibly exhibiting deep interpenetrations, to a collision-free state, ready for animation. Our method relies on an intermediate, implicit representation to solve the problem: the user selects a few garments stored in a library together with their implicit approximations, and places them over a mannequin while specifying the desired order between layers. The intersecting implicit surfaces are then combined using a new family of N-ary composition operators, specially designed for untangling layers. Garment meshes are finally projected to the deformed implicit surfaces in linear time, while best preserving triangles and avoiding loss of details.

Each of the untangling operators computes the target surface for a given garment in a single step, while accounting for the order between cloth layers and their individual thicknesses. As a group, they guarantee an intersection-free output configuration. Moreover, a weight can be associated with each layer to tune their relative influence during untangling, such as leather being less deformed than cloth. Results for each layer then reflect the combined effect of the other layers, enabling us to output a plausible configuration in contact regions. As our results show, our method can be used to generate plausible, new static shapes of garments when underwear has been added, as well as collision-free configurations enabling a user to safely launch animations of arbitrarily complex layered clothing.

CCS Concepts: • **Computing methodologies → Collision Detection**; **Volumetric Models**.

Additional Key Words and Phrases: Implicit surfaces, Cloth untangling, Collision processing, Cloth animation

ACM Reference Format:
Thomas Buffet, Damien Rohmer, Loïc Barthe, Laurence Boissieux, and Marie-Paule Cani. 2019. Implicit Untangling: A Robust Solution for Modeling Layered Clothing . *ACM Trans. Graph.* 38, 4, Article 120 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3323010

Authors' addresses: Thomas Buffet, Inria Rhônes-Alpes , Montbonnot-Saint-Martin, France, 38330, 655 Avenue de l'Europe , LIX, Ecole Polytechnique, CNRS, IP Paris, France, 91120, Palaiseau, 1 rue Honoré d'Estienne d'Orves Alan Turing building; Damien Rohmer, LIX, Ecole Polytechnique, CNRS, IP Paris, France, 91120, Palaiseau, 1 rue Honoré d'Estienne d'Orves Alan Turing building; Loïc Barthe, University of Toulouse - IRIT - CNRS, France, Toulouse; Laurence Boissieux, Inria Rhônes-Alpes , Montbonnot-Saint-Martin, France, 38330, 655 Avenue de l'Europe; Marie-Paule Cani, LIX, Ecole Polytechnique, CNRS, IP Paris, France, 91120, Palaiseau, 1 rue Honoré d'Estienne d'Orves Alan Turing building.

0730-0301/2019/7-ART120 $15.00
https://doi.org/10.1145/3306346.3323010

1 INTRODUCTION

The efficient animation of dressed virtual characters is an essential step in 3D production pipelines. Tremendous efforts have been dedicated to achieve high quality dynamics for individual pieces of clothing, based on physically-based simulation models and efficient collision detection and response. Extending these animations to multiple layers of garments is however difficult, starting with the

Learning to Optimize Halide with Tree Search and Random Programs

ANDREW ADAMS, Facebook AI Research
KARIMA MA, UC Berkeley
LUKE ANDERSON, MIT CSAIL
RIYADH BAGHDADI, MIT CSAIL
TZU-MAO LI, MIT CSAIL
MICHAËL GHARBI, Adobe
BENOIT STEINER, Facebook AI Research
STEVEN JOHNSON, Google
KAYVON FATAHALIAN, Stanford University
FRÉDO DURAND, MIT CSAIL
JONATHAN RAGAN-KELLEY, UC Berkeley

We present a new algorithm to automatically schedule Halide programs for high-performance image processing and deep learning. We significantly improve upon the performance of previous methods, which considered a limited subset of schedules. We define a parameterization of possible schedules much larger than prior methods and use a variant of beam search to search over it. The search optimizes runtime predicted by a cost model based on a combination of new derived features and machine learning. We train the cost model by generating and featurizing hundreds of thousands of random programs and schedules. We show that this approach operates effectively with or without autotuning. It produces schedules which are on average almost twice as fast as the existing Halide autoscheduler without autotuning, or more than twice as fast with, and is the first automatic scheduling algorithm to significantly outperform human experts on average.

CCS Concepts: • **Computing methodologies** → **Image processing**; • **Software and its engineering** → **Domain specific languages**.

Additional Key Words and Phrases: optimizing compilers, Halide

ACM Reference Format:
Andrew Adams, Karima Ma, Luke Anderson, Riyadh Baghdadi, Tzu-Mao Li, Michaël Gharbi, Benoit Steiner, Steven Johnson, Kayvon Fatahalian, Frédo Durand, and Jonathan Ragan-Kelley. 2019. Learning to Optimize Halide with Tree Search and Random Programs. *ACM Trans. Graph.* 38, 4, Article 121 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3322967

Authors' addresses: Andrew Adams, Facebook AI Research, andrew.b.adams@gmail.com; Karima Ma, UC Berkeley, karima_ma@berkeley.edu; Luke Anderson, MIT CSAIL, lukea@mit.edu; Riyadh Baghdadi, MIT CSAIL, baghdadi@mit.edu; Tzu-Mao Li, MIT CSAIL, tzumao@mit.edu; Michaël Gharbi, Adobe, mgharbi@adobe.com; Benoit Steiner, Facebook AI Research, benoitsteiner@fb.com; Steven Johnson, Google, srj@google.com; Kayvon Fatahalian, Stanford University, kayvonf@cs.stanford.edu; Frédo Durand, MIT CSAIL, fredo@csail.mit.edu; Jonathan Ragan-Kelley, UC Berkeley, jrk@berkeley.edu.

0730-0301/2019/7-ART121 $15.00
https://doi.org/10.1145/3306346.3322967

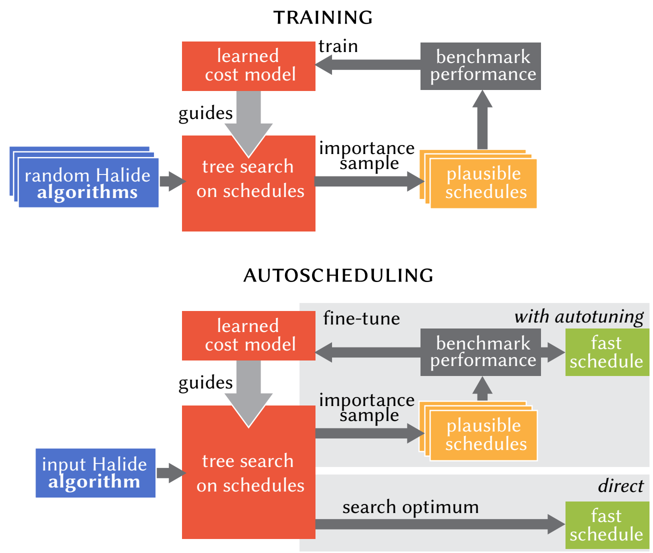

Fig. 1. We generate schedules for Halide programs using tree search over the space of schedules (Sec. 3) guided by a learned cost model and optional autotuning (Sec. 4). The cost model is trained by benchmarking thousands of randomly-generated Halide programs and schedules (Sec. 5). The resulting code significantly outperforms prior work and human experts (Sec. 6).

1 INTRODUCTION

Image processing and deep learning are pervasive. They are computationally intense, and implementations often have to be highly optimized by experts, at great cost, to be usable in practice. The Halide programming language has proven to be a powerful tool for this task because it separates the *algorithm* — what you want to compute — from the *schedule* — how you want to compute it, including choices about memory locality, redundant computation, and parallelism [Ragan-Kelley et al. 2012, 2013]. While Halide makes it easy to try different schedules, writing schedules that achieve high performance is hard: it requires expertise in hardware architecture and optimization, and even then, the space of possible schedules is enormous and their performance can be difficult to predict.

Automating the synthesis of high-performance schedules is sorely needed, but prior methods are limited in multiple ways [Mullapudi et al. 2016, 2015; Sioutas et al. 2018]. First, by design they only consider a small subset of all possible schedules. They generally work by modestly generalizing specific schedule templates or idioms, which are difficult to compose or extend to capture the countless other potentially fruitful choices for each application and target architecture. Second, they explore their key choices using specialized search procedures, which are tightly coupled to the family of schedules they consider. Third, they navigate this space using hand-designed cost models which struggle to accurately predict performance on real machines. These cost models can be tuned to guide the search to good performance on a few specific applications at a time, but they

KleinPAT: Optimal Mode Conflation For Time-Domain Precomputation Of Acoustic Transfer

JUI-HSIEN WANG, Stanford University
DOUG L. JAMES, Stanford University

We propose a new modal sound synthesis method that rapidly estimates all acoustic transfer fields of a linear modal vibration model, and greatly reduces preprocessing costs. Instead of performing a separate frequency-domain Helmholtz radiation analysis for each mode, our method partitions vibration modes into *chords* using optimal *mode conflation*, then performs a single time-domain wave simulation for each chord. We then perform *transfer deconflation* on each chord's time-domain radiation field using a specialized QR solver, and thereby extract the frequency-domain transfer functions of each mode. The precomputed transfer functions are represented for fast far-field evaluation, e.g., using multipole expansions. In this paper, we propose to use a single scalar-valued Far-field Acoustic Transfer (FFAT) cube map. We describe a GPU-accelerated vector wavesolver that achieves high-throughput acoustic transfer computation at accuracy sufficient for sound synthesis. Our implementation, KleinPAT, can achieve hundred- to thousand-fold speedups compared to existing Helmholtz-based transfer solvers, thereby enabling large-scale generation of modal sound models for audio-visual applications.

CCS Concepts: • **Computing methodologies** → *Modeling and simulation*; *Physical simulation*.

Additional Key Words and Phrases: Computer animation, sound synthesis, GPU, modal models.

ACM Reference Format:
Jui-Hsien Wang and Doug L. James. 2019. KleinPAT: Optimal Mode Conflation For Time-Domain Precomputation Of Acoustic Transfer. *ACM Trans. Graph.* 38, 4, Article 122 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3322976

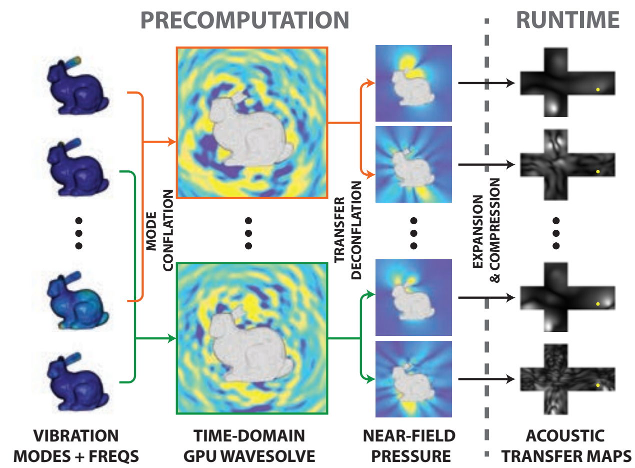

Fig. 1. **KleinPAT Overview:** Our method can evaluate acoustic transfer for 292 modes of this plastic bunny using only 6 time-domain wave simulations constructed by optimally conflating modes into 6 chords. The resulting sounds fields are deconflated to estimate the 292 transfer fields, then approximated with far-field acoustic transfer (FFAT) cube maps for real-time evaluation. This precomputed acoustic transfer (PAT) preprocess is over 2000x faster than traditional BEM-based approaches for the bunny.

PROLOGUE

Fritz Heinrich Klein created the first known all-interval twelve-tone row, the *Mutterakkord* (Mother chord), for his chamber-orchestra composition Die Maschine in 1921. Mother chord (see inset figure) is arranged so that it contains one instance of each interval within the octave [Wikipedia 2019]. Our KleinPAT algorithm optimally arranges different modal *tones* of a vibrating 3D object into *chords*, which are then *played* together by a time-domain vector wavesolver in order to efficiently estimate all acoustic transfer fields. All together now...

Authors' addresses: Jui-Hsien Wang, Stanford University, jui-hsien.wang@stanford.edu; Doug L. James, Stanford University.

0730-0301/2019/7-ART122 $15.00
https://doi.org/10.1145/3306346.3322976

1 INTRODUCTION

Physics-based modal sound synthesis is an effective and popular method for synthesizing plausible rigid-body contact sounds in computer animation and interactive virtual environments. Modal sound models use linear vibration modes to effectively characterize the surface vibrations of struck objects. The sound amplitude of each mode is described by its *acoustic transfer function*, whose spatial structure accounts for perceptually important wave effects like diffraction and radiation efficiency.

Unfortunately obtaining the acoustic transfer functions is a computational drag, since it involves solving the frequency-domain wave equation (Helmholtz equation) once for each vibration mode—there can be hundreds of audible modes even for simple objects. Precomputed Acoustic Transfer (PAT) methods were introduced to accelerate the otherwise costly *runtime* evaluation of transfer by precomputing each mode's exterior radiation field, and representing it in a form optimized for real-time evaluation. However the precomputation of acoustic transfer is still a computational bottleneck, even when using modern Helmholtz BEM solvers; hours to days of transfer precomputation can be required on large computing clusters, for objects of even small to moderate size, with higher frequency modes being more expensive. Currently, the high cost of

Physically-based Statistical Simulation of Rain Sound

SHIGUANG LIU, Tianjin University, China
HAONAN CHENG, Tianjin University, China
YIYING TONG, Michigan State University, USA

A typical rainfall scenario contains tens of thousands of dynamic sound sources. A characteristic of the large-scale scene is the strong randomness in raindrop distribution, which makes it notoriously expensive to synthesize such sounds with purely physical methods. Moreover, the raindrops hitting different surfaces (liquid or various solids) can emit distinct sounds, for which prior methods with unified impact sound models are ill-suited.

In this paper, we present a physically-based statistical simulation method to synthesize realistic rain sound, which respects surface materials. We first model the raindrop sound with two mechanisms, namely the initial impact and the subsequent pulsation of entrained bubbles. Then we generate material sound textures (MSTs) based on a specially designed signal decomposition and reconstruction model. This allows us to distinguish liquid surface with bubble sound and different solid surfaces with MSTs. Furthermore, we build a basic rain sound (BR-sound) bank with the proposed raindrop sound clustering method based on a statistical model, and design a sound source activator for simulating spatial propagation in an efficient manner. This novel method drastically decreases the computational cost while producing convincing sound results. Various experiments demonstrate the effectiveness of our sound simulation model.

CCS Concepts: • **Computing methodologies** → *Computer graphics*; • **Applied computing** → *Sound and music computing*.

Additional Key Words and Phrases: sound synthesis, basic rain sound, physical and statistical, rain animation, procedural audio

ACM Reference Format:
Shiguang Liu, Haonan Cheng, and Yiying Tong. 2019. Physically-based Statistical Simulation of Rain Sound. *ACM Trans. Graph.* 38, 4, Article 123 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3323045

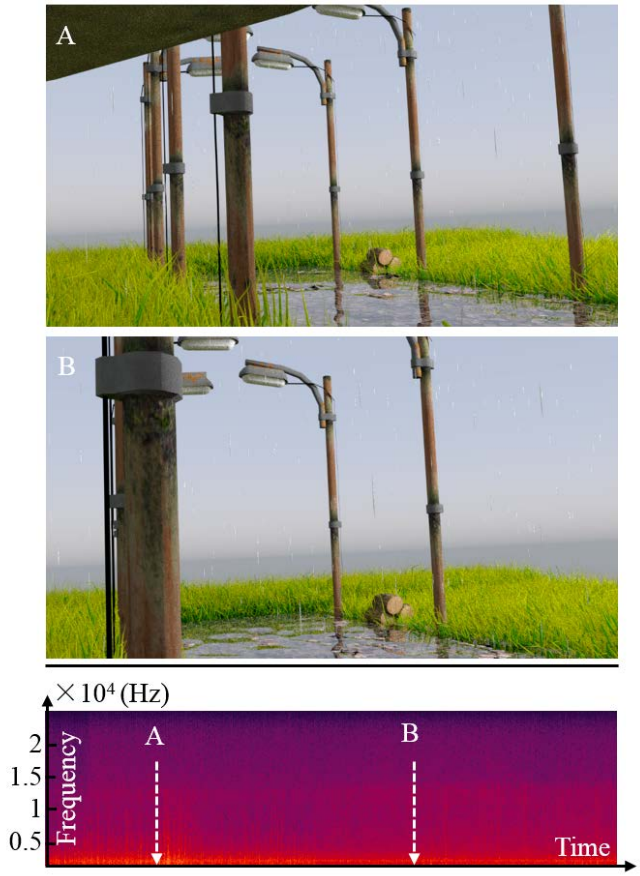

Fig. 1. **Audiovisual simulation of walking along a country road in the rain.** Frames A and B illustrate two positions of the listener. From A to B, as the listener gets closer to the street lights, the metal sound effects become more and more apparent. Our method is able to capture this dynamic change of rain sound as shown in the spectrogram on the bottom row.

1 INTRODUCTION

Rain sound, as a common environmental sound, is widely used in virtual scenes such as animated movies and computer games. It can greatly enhance one's sense of immersion in virtual scenes. There are two obvious characteristics of a rainfall scenario: large scale with a large number of raindrops, and the heterogeneity of interaction surfaces in the scene, which make synthesis of accurate rain sound rather computationally demanding. In this paper, we propose a method that can efficiently resolve these challenges and synthesize realistic surface material-aware rain sound. We show an example produced by our sound synthesis pipeline in Figure 1. (See more in § 6.)

Recent advances in the sound simulation have enabled the synthesis of rain sound and these methods can be classified into two categories, non-physically-based methods and physically-based methods. The non-physically based methods usually employ signal processing to adjust the recorded sound for authenticity. However, for the methods in this category, as discussed in [Verron and Drettakis 2012; Zita 2003], there are two major limitations: *i*) the indirect correspondence with the parameters of the physical animation model, and *ii*) the need for artificial adjustment. Both limitations can be overcome by physically-based methods with direct matching of

Authors' addresses: Shiguang Liu, Tianjin University, China, lsg@tju.edu.cn; Haonan Cheng, Tianjin University, China, tjuchn@tju.edu.cn; Yiying Tong, Michigan State University, USA, ytong@msu.edu.

0730-0301/2019/7-ART123 $15.00
https://doi.org/10.1145/3306346.3323045

Variational Implicit Point Set Surfaces

ZHIYANG HUANG, Washington University in St. Louis, USA
NATHAN CARR, Adobe Systems, USA
TAO JU, Washington University in St. Louis, USA

We propose a new method for reconstructing an implicit surface from an un-oriented point set. While existing methods often involve non-trivial heuristics and require additional constraints, such as normals or labelled points, we introduce a direct definition of the function from the points as the solution to a constrained quadratic optimization problem. The definition has a number of appealing features: it uses a single parameter (parameter-free for exact interpolation), applies to any dimensions, commutes with similarity transformations, and can be easily implemented without discretizing the space. More importantly, the use of a global smoothness energy allows our definition to be much more resilient to sampling imperfections than existing methods, making it particularly suited for sparse and non-uniform inputs.

CCS Concepts: • **Computing methodologies → Point-based models**; *Volumetric models.*

Additional Key Words and Phrases: surface reconstruction, implicit surfaces, point clouds, radial basis functions

ACM Reference Format:
Zhiyang Huang, Nathan Carr, and Tao Ju. 2019. Variational Implicit Point Set Surfaces. *ACM Trans. Graph.* 38, 4, Article 124 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3322994

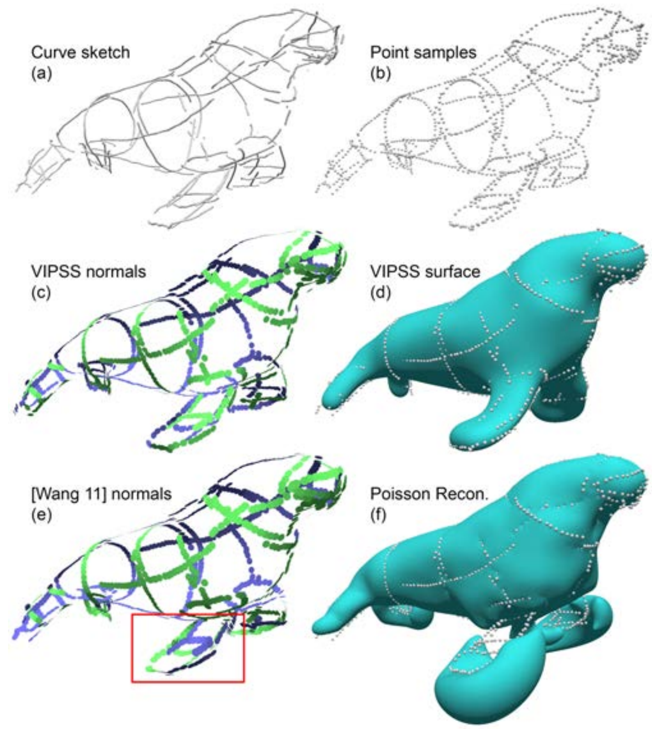

Fig. 1. Given sparse, non-uniform, noisy and un-oriented points (b) sampled from a set of unstructured 3D curves (a), our variational definition (VIPSS with $\lambda = 0.003$) simultaneously produces oriented normals (c) and a smooth approximating surface (d). The input is challenging for state-of-the-art normal estimation methods such as [Wang et al. 2011], which fails around sparsely sampled thin features (the flippers) (e). Incorrect normals lead to poor reconstructions using existing implicit methods such as Screened Poisson [Kazhdan and Hoppe 2013] (f, fitting weight $\alpha = 0.5$).

1 INTRODUCTION

Constructing a curve or surface that interpolates or approximates a given set of 2D or 3D points is one of the fundamental problems in geometric modeling. A common representation of the reconstructed output is the zero-level set of some smooth implicit function. This representation naturally ensures a smooth and closed manifold. In addition, an implicit function enables a range of applications such as boolean operations and collision detection.

While extensively studied [Berger et al. 2017], implicit modeling from points remains a difficult problem. A fundamental challenge is that the constant zero function, although meaningless for reconstruction, perfectly meets our goal: the function is smooth and its zero-level set (which includes the entire space) interpolates any input points. A common approach to avoid this trivial solution is to introduce additional constraints, such as normals at the input points or additional spatial locations with inside/outside labels. If such constraints are not available as part of the input or provided by the user, they will need to be inferred from the input data *prior to* reconstruction.

This two-stage paradigm - constraint estimation followed by reconstruction - has a number of drawbacks. The use of multiple disparate methods, each carrying its own set of parameters, makes parameter-tuning a challenging task and complicates the analysis (e.g., how the output changes with the input). More importantly, methods for constraint estimation are completely unaware of the quality of the reconstructed surface. Lacking any better guidance, current estimation methods (e.g., for normals) rely on local point neighborhoods, which are often unreliable when the points are sparse or non-uniformly distributed. Errors in constraint estimation, in turn, lead to poorly reconstructed surfaces (e.g., Figure 1 (e,f)).

We propose a direct definition of an implicit function from an un-oriented point set. Unlike the above-mentioned two-stage paradigm, our definition integrates constraint estimation and surface

Authors' addresses: Zhiyang Huang, Washington University in St. Louis, St. Louis, MO, USA; Nathan Carr, Adobe Systems, San Jose, CA, USA; Tao Ju, Washington University in St. Louis, St. Louis, MO, USA.

0730-0301/2019/7-ART124 $15.00
https://doi.org/10.1145/3306346.3322994

Sample-based Monte Carlo Denoising using a Kernel-Splatting Network

MICHAËL GHARBI, Adobe and MIT CSAIL
TZU-MAO LI, MIT CSAIL
MIIKA AITTALA, MIT CSAIL
JAAKKO LEHTINEN, Aalto University and NVIDIA
FRÉDO DURAND, MIT CSAIL

Fig. 1. State-of-the-art pixel-based Monte Carlo denoising algorithms (right) struggle with very noisy inputs rendered with a low sample count (left). Our method (middle) works with the *samples* directly, it uses a *splatting* approach, and is trained using deep learning. This makes it possible to appropriately handle various components of the illumination (indirect lighting, specular reflection, motion blur, depth of field, etc) more effectively.

Denoising has proven to be useful to efficiently generate high-quality Monte Carlo renderings. Traditional pixel-based denoisers exploit summary statistics of a pixel's sample distributions, which discards much of the samples' information and limits their denoising power. On the other hand, sample-based techniques tend to be slow and have difficulties handling general transport scenarios. We present the first convolutional network that can learn to denoise Monte Carlo renderings directly from the samples. Learning the mapping between samples and images creates new challenges for the network architecture design: the order of the samples is arbitrary, and they should be treated in a permutation invariant manner. To address these challenges, we develop a novel kernel-predicting architecture that *splats* individual samples onto nearby pixels. Splatting is a natural solution to situations such as motion blur, depth-of-field and many light transport paths, where it is easier to predict which pixels a sample contributes to, rather than a *gather* approach that needs to figure out, for each pixel, which samples (or nearby pixels) are relevant. Compared to previous state-of-the-art methods, ours is robust to the severe noise of low-sample count images (e.g. 8 samples per pixel) and yields higher-quality results both visually and numerically. Our approach retains the generality and efficiency of pixel-space methods while enjoying the expressiveness and accuracy of the more complex sample-based approaches.

Authors' addresses: Michaël Gharbi, Adobe and MIT CSAIL, mgharbi@adobe.com; Tzu-Mao Li, MIT CSAIL, tzumao@mit.edu; Miika Aittala, MIT CSAIL, miika@mit.edu; Jaakko Lehtinen, Aalto University and NVIDIA, jaakko.lehtinen@aalto.fi; Frédo Durand, MIT CSAIL, fredo@mit.edu.

0730-0301/2019/7-ART125 $15.00
https://doi.org/10.1145/3306346.3322954

CCS Concepts: • **Computing methodologies → Neural networks**; **Ray tracing**; *Image processing*.

Additional Key Words and Phrases: Monte Carlo denoising, deep learning, data-driven methods, convolutional neural networks

ACM Reference Format:
Michaël Gharbi, Tzu-Mao Li, Miika Aittala, Jaakko Lehtinen, and Frédo Durand. 2019. Sample-based Monte Carlo Denoising using a Kernel-Splatting Network. *ACM Trans. Graph.* 38, 4, Article 125 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3322954

1 INTRODUCTION

Because Monte Carlo methods rely on stochastic point samples of an intricate integrand, they often suffer from noise. This motivates Monte Carlo denoising techniques, which broadly fall into two categories. *Sample-based* techniques keep track of the individual samples, while *pixel-based* methods work directly on the rendered image. Most methods operate in pixel space (e.g. [Bako et al. 2017; Bitterli et al. 2016]). In addition to the noisy input radiance, they usually exploit first and second order statistics of auxiliary buffers (like depth, normal, albedo, etc) [McCool 1999]. We argue that, in many lighting configurations, simple per-pixel aggregates can under-represent the complexity of the local light transport phenomena, in particular because the distribution is often multimodal.

We present a new Monte Carlo denoising technique that leverages the power of deep learning in the following key manners compared to previous denoising methods:

- Rather than work on pixel-based representations, our input is the raw set of Monte Carlo samples, which we argue allows our method to appropriately handle information of different nature. In particular, depth of field and motion blur generate

Deep Convolutional Reconstruction For Gradient-Domain Rendering

MARKUS KETTUNEN, Aalto University
ERIK HÄRKÖNEN, Aalto University
JAAKKO LEHTINEN, Aalto University and Nvidia

Fig. 1. Comparison of the primal-domain denoisers NFOR [Bitterli et al. 2016] and KPCN [Bako et al. 2017] to our gradient-domain reconstruction NGPT from very noisy equal-time inputs (8 samples for ours and 20 for others). Generally outperforming the comparison methods, our results show that gradient sampling is useful also in the context of non-linear neural image reconstruction, often resolving e.g. shadows better than techniques that do not make use of gradients.

It has been shown that rendering in the gradient domain, i.e., estimating finite difference gradients of image intensity using correlated samples, and combining them with direct estimates of pixel intensities by solving a screened Poisson problem, often offers fundamental benefits over merely sampling pixel intensities. The reasons can be traced to the frequency content of the light transport integrand and its interplay with the gradient operator. However, while they often yield state of the art performance among algorithms that are based on Monte Carlo sampling alone, gradient-domain rendering algorithms have, until now, not generally been competitive with techniques that combine Monte Carlo sampling with post-hoc noise removal using sophisticated non-linear filtering.

Drawing on the power of modern convolutional neural networks, we propose a novel reconstruction method for gradient-domain rendering. Our technique replaces the screened Poisson solver of previous gradient-domain techniques with a novel dense variant of the U-Net autoencoder, additionally taking auxiliary feature buffers as inputs. We optimize our network to minimize a perceptual image distance metric calibrated to the human visual system. Our results significantly improve the quality obtained from gradient-domain path tracing, allowing it to overtake state-of-the-art comparison techniques that denoise traditional Monte Carlo samplings. In particular, we observe that the correlated gradient samples — that offer information about the smoothness of the integrand unavailable in standard Monte Carlo sampling — notably improve image quality compared to an equally powerful neural model that does not make use of gradient samples.

Authors' addresses: Markus Kettunen, Aalto University, markus.kettunen@aalto.fi; Erik Härkönen, Aalto University, erik.harkonen@aalto.fi; Jaakko Lehtinen, Aalto University and Nvidia, jaakko.lehtinen@aalto.fi.

0730-0301/2019/7-ART126 $15.00
https://doi.org/10.1145/3306346.3323038

CCS Concepts: • **Computing methodologies → Neural networks**; **Ray tracing**.

Additional Key Words and Phrases: gradient-domain rendering, gradient-domain reconstruction, screened poisson, ray tracing

ACM Reference Format:
Markus Kettunen, Erik Härkönen, and Jaakko Lehtinen. 2019. Deep Convolutional Reconstruction For Gradient-Domain Rendering. *ACM Trans. Graph.* 38, 4, Article 126 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3323038

1 INTRODUCTION

Realistic image synthesis seeks to produce realistic virtual photographs by computationally solving the Rendering Equation [Kajiya 1986], often by randomly sampling paths that carry light from the light sources to the sensor. Rendering with too few samples leaves the image with visually distracting noise. Unsurprisingly, practical applications constantly struggle with striking a balance between the complexity of content (slower, more noise) and available computational resources.

Since many Monte Carlo samples are required for a high quality image, this leaves four main approaches for making rendering faster: (1) making samples faster to evaluate (e.g. GPU rendering, ray tracing hardware, optimized low-level algorithms), (2) sharing contributions between nearby paths (e.g. photon mapping), (3) being clever in choosing the light paths to sample (e.g. Bidirectional Path Tracing, adaptive importance samplers), and, finally, (4) denoising or reconstruction, attempting to produce a better picture out of the samples by relying on various smoothness assumptions or analytic models of the transport phenomena being modeled.

Despite a long history, and continuous research progress in all of these areas, significant problems remain. Only naturally, the quality obtained by "more pure" techniques that rely on few assumptions or heuristics tends to lag behind those that assume more. For instance,

A Learned Shape-Adaptive Subsurface Scattering Model

DELIO VICINI, Ecole Polytechnique Fédérale de Lausanne (EPFL)
VLADLEN KOLTUN, Intel Labs
WENZEL JAKOB, Ecole Polytechnique Fédérale de Lausanne (EPFL)

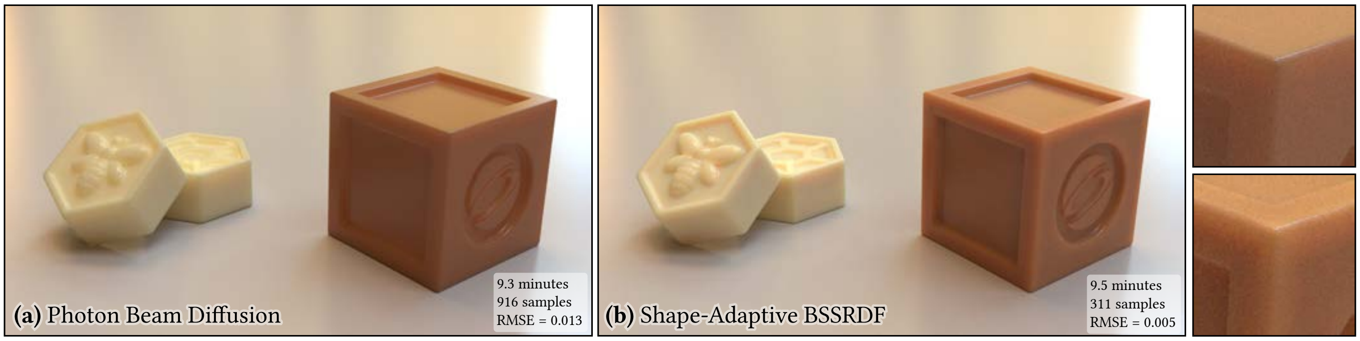

Fig. 1. Rendering of translucent soap blocks with significant anisotropy ($g = 0.9$). **(a)** Dipole diffusion-based models such as Photon Beam Diffusion [Habel et al. 2013] yield overall flat appearance due to their internal assumption of isotropy and plane-parallel light transport. **(b)** Our learned subsurface scattering model adapts to both geometry and anisotropy, producing more realistic appearance and lower error compared to a path-traced reference. Accounting for the geometry leads to increased scattering around the silhouette and overall higher contrast in regions with geometric detail. Our method has constant sampling weights, hence renderings converge with fewer samples. Please see the supplemental material for an interactive version of this figure including error maps.

Subsurface scattering, in which light refracts into a translucent material to interact with its interior, is the dominant mode of light transport in many types of organic materials. Accounting for this phenomenon is thus crucial for visual realism, but explicit simulation of the complex internal scattering process is often too costly. BSSRDF models based on analytic transport solutions are significantly more efficient but impose severe assumptions that are almost always violated, e.g. planar geometry, isotropy, low absorption, and spatio-directional separability. The resulting discrepancies between model and usage lead to objectionable errors in renderings, particularly near geometric features that violate planarity.

This article introduces a new shape-adaptive BSSRDF model that retains the efficiency of prior analytic methods while greatly improving overall accuracy. Our approach is based on a conditional variational autoencoder, which learns to sample from a reference distribution produced by a brute-force volumetric path tracer. In contrast to the path tracer, our autoencoder directly samples outgoing locations on the object surface, bypassing a potentially lengthy internal scattering process.

The distribution is conditional on both material properties and a set of features characterizing geometric variation in a neighborhood of the incident location. We use a low-order polynomial to model the local geometry as an implicitly defined surface, capturing curvature, thickness, corners, as well as cylindrical and toroidal regions. We present several examples of objects with challenging medium parameters and complex geometry and compare to ground truth simulations and prior work.

Authors' addresses: Delio Vicini, Ecole Polytechnique Fédérale de Lausanne (EPFL), delio.vicini@epfl.ch; Vladlen Koltun, Intel Labs, vladlen.koltun@intel.com; Wenzel Jakob, Ecole Polytechnique Fédérale de Lausanne (EPFL), wenzel.jakob@epfl.ch.

0730-0301/2019/7-ART127 $15.00
https://doi.org/10.1145/3306346.3322974

CCS Concepts: • **Computing methodologies → Reflectance modeling**.

ACM Reference Format:
Delio Vicini, Vladlen Koltun, and Wenzel Jakob. 2019. A Learned Shape-Adaptive Subsurface Scattering Model. *ACM Trans. Graph.* 38, 4, Article 127 (July 2019), 15 pages. https://doi.org/10.1145/3306346.3322974

1 INTRODUCTION

A variety of materials ranging from organic substances like skin or fruits to liquids and translucent stone exhibit significant subsurface light transport, where light that enters the surface at one point exits some distance away. This process is generally modeled using the radiative transfer equation (RTE) and involves long sequences of scattering interactions within a medium that fills the object's interior. Although Monte Carlo methods can be used to solve the resulting integration problem, this tends to be impracticably expensive for many real-world materials due to their high albedo, anisotropy, and density.

For instance, Narasimhan et al. [2006] measured the scattering parameters of a large number of materials and report that milk has an albedo of $\alpha \approx 0.99959$, meaning that light will scatter an expected number of 2439 times before being absorbed. Furthermore, scattering is strongly forward-peaked ($g \approx 0.9$), and light will thus tend to penetrate deeply into the object.

An alternative approach for rendering subsurface scattering without the need for costly internal scattering simulations involves the notion of a *bidirectional scattering-surface reflectance distribution function* (BSSRDF). BSSRDF models directly encode surface-to-surface transport and are typically based on analytic solutions of simplified light transport problems, such as solutions of the

Efficient and Conservative Fluids Using Bidirectional Mapping

ZIYIN QU*, AICFVE and University of Pennsylvania
XINXIN ZHANG*, AICFVE
MING GAO, University of Pennsylvania
CHENFANFU JIANG, University of Pennsylvania
BAOQUAN CHEN, Peking University

In this paper, we introduce BiMocq^2, an unconditionally stable, pure Eulerian-based advection scheme to efficiently preserve the advection accuracy of all physical quantities for long-term fluid simulations. Our approach is built upon the method of characteristic mapping (MCM). Instead of the costly evaluation of the temporal characteristic integral, we evolve the mapping function itself by solving an advection equation for the mappings. Dual mesh characteristics (DMC) method is adopted to more accurately update the mapping. Furthermore, to avoid visual artifacts like instant blur and temporal inconsistency introduced by re-initialization, we introduce multi-level mapping and back and forth error compensation. We conduct comprehensive 2D and 3D benchmark experiments to compare against alternative advection schemes. In particular, for the vortical flow and level set experiments, our method outperforms almost all state-of-art hybrid schemes, including FLIP, PolyPic and Particle-Level-Set, at the cost of only two Semi-Lagrangian advections. Additionally, our method does not rely on the particle-grid transfer operations, leading to a highly parallelizable pipeline. As a result, more than 45× performance acceleration can be achieved via even a straightforward porting of the code from CPU to GPU.

CCS Concepts: • **Computing methodologies → Continuous models**; **Physical simulation**.

Additional Key Words and Phrases: Fluid Simulation, Conservative Advection, MCM,

ACM Reference Format:
Ziyin Qu, Xinxin Zhang, Ming Gao, Chenfanfu Jiang, and Baoquan Chen. 2019. Efficient and Conservative Fluids Using Bidirectional Mapping. *ACM Trans. Graph.* 38, 4, Article 128 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3322945

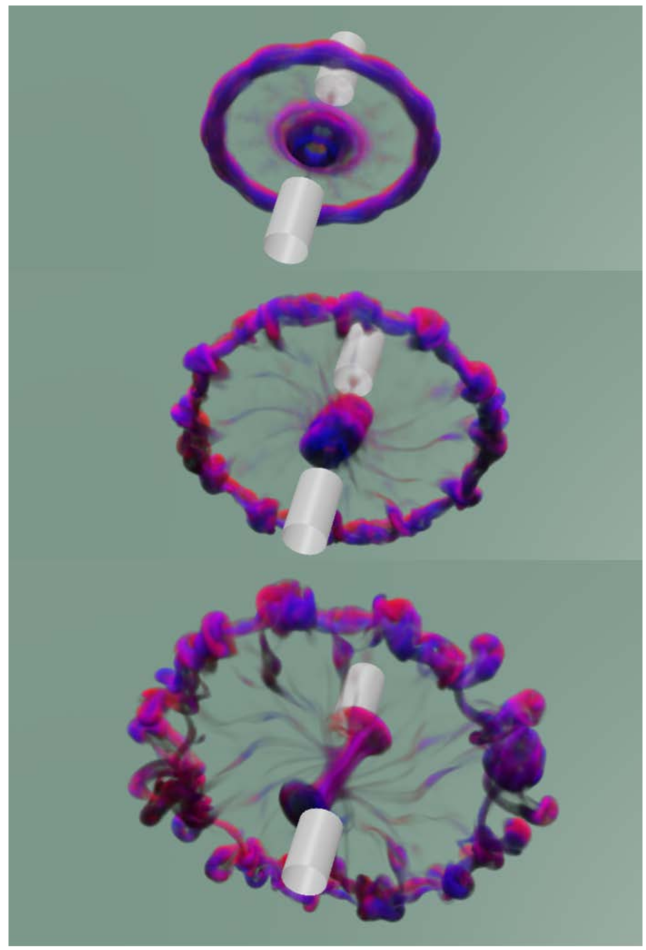

Fig. 1. Frames of vortex ring colliding at Re=2000 simulated using BiMocq^2, from top to bottom: frame 80, frame 140, and frame 280. The proposed method reproduces the whole process where colliding vortex rings stretch and form the out-shooting jets from the side of the rings.

1 INTRODUCTION

Eulerian based fluid simulations have achieved great success in reproducing a wide range of phenomena in computer graphics, such as liquids [Aanjaneya et al. 2017; Enright et al. 2002b; Foster and Fedkiw 2001], smoke and fire [Fedkiw et al. 2001; Nguyen et al. 2002; Rasmussen et al. 2003; Setaluri et al. 2014]. The framework to solve the Navier-Stokes equation in a time splitting manner, especially the critical insight to treat the nonlinear advection term using Semi-Lagrangian scheme [Stam 1999], results in its most significant advantage of stability, hence ease of control, along with its biggest criticism: the numerical diffusion. For more details about the background, we refer to the book by Bridson [2008].

*Z. Qu, X. Zhang are joint first authors.

Authors' addresses: Ziyin Qu, AICFVE and University of Pennsylvania, ziyinq@seas.upenn.edu; Xinxin Zhang, AICFVE, zhangshinshin@gmail.com; Ming Gao, University of Pennsylvania, minggao@seas.upenn.edu; Chenfanfu Jiang, University of Pennsylvania, cffjiang@seas.upenn.edu; Baoquan Chen, Peking University, baoquan@pku.edu.cn.

0730-0301/2019/7-ART128 $15.00
https://doi.org/10.1145/3306346.3322945

On Bubble Rings and Ink Chandeliers

MARCEL PADILLA, TU Berlin
ALBERT CHERN, TU Berlin
FELIX KNÖPPEL, TU Berlin
ULRICH PINKALL, TU Berlin
PETER SCHRÖDER, Caltech

We introduce variable thickness, viscous vortex filaments. These can model such varied phenomena as underwater bubble rings or the intricate "chandeliers" formed by ink dropping into fluid. Treating the evolution of such filaments as an instance of Newtonian dynamics on a Riemannian configuration manifold we are able to extend classical work in the dynamics of vortex filaments through inclusion of viscous drag forces. The latter must be accounted for in low Reynolds number flows where they lead to significant variations in filament thickness and form an essential part of the observed dynamics. We develop and document both the underlying theory and associated practical numerical algorithms.

CCS Concepts: • **Mathematics of computing → Partial differential equations**; *Differential calculus*; *Algebraic topology*; • **Computing methodologies → Mesh models**; *Physical simulation*; • **Applied computing → Physics.**

Additional Key Words and Phrases: Differential geometry, physical modeling, geodesics, fluid simulation, vorticity methods, vortex filaments.

ACM Reference Format:
Marcel Padilla, Albert Chern, Felix Knöppel, Ulrich Pinkall, and Peter Schröder. 2019. On Bubble Rings and Ink Chandeliers. *ACM Trans. Graph.* 38, 4, Article 129 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3322962

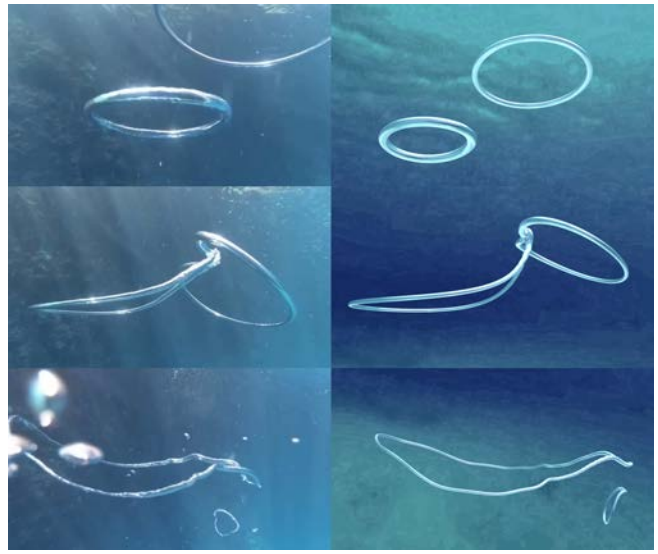

Fig. 1. Evolution of two underwater bubble rings connecting, showing frames from a video (left) and our simulation (right) with approximate ring size 2 m and circulation 2.5 $m^2 s^{-1}$. See also the video at time 1:12.

1 INTRODUCTION

The study of vortex dynamics is by now a classical subject in fluid dynamics [Saffman 1992] with roots all the way back to Cauchy [1815] (see [Frisch and Villone 2014] for an historical account). One of the central elements in this study is the vortex filament, first articulated by Helmholtz [1858] (see [Meleshko et al. 2012] for an historical account). A vortex filament is a closed curve (or beginning and ending on the boundary of the domain) carrying concentrated vorticity, generating a velocity field according to the Biot–Savart law [1820]. Vortex filaments are easily accessible through experiment [Rogers 1858], serve as useful abstractions in the analysis of flows [Shariff and Leonard 1992], and can serve as a basis for Lagrangian flow simulations in the high Reynolds number regime in computational fluid dynamics and computer graphics [Chorin 1990, 1993; Angelidis and Neyret 2005; Bernard 2006, 2009; Weißmann and Pinkall 2010; Liao et al. 2018] (to name but a few).

Comparatively less is known about the evolution of vortex filaments in low to moderate Reynolds number flows, where they can model beautiful phenomena as diverse as bubble rings [Turner 1957; Pedley 1968] (Figs. 1, 3) and the chandeliers formed by ink dropping into clear fluid [Tomlinson 1864; Thomson and Newall 1886] (Figs. 2, 4). In these settings viscous drag and buoyancy affect thickness variations along the filament. These have significant impact on the dynamics of the filament as a whole and must be modeled explicitly.

To capture these dynamics we develop a model for variable thickness vortex filaments subject to viscous drag and buoyancy. The dynamics without drag and buoyancy coincide with classic descriptions [Moore and Saffman 1972] according to the Biot–Savart [1820] and Bernoulli [1738] laws. To account for the effect of dissipative drag we first recast the classic evolution equations as geodesic equations on a suitably constructed configuration manifold [Kaluza 1921; Klein 1926]. Dissipative drag can then be included in a standard manner (as can other external forces such as buoyancy). This results in three coupled equations describing the deformation of the center curve of the filament (Eq. (18)), changes in the thickness along the filament (Eq. (15)), and changes in the circulation, *i.e.*, filament strength (Eq. (12b)), as functions of time. Using appropriate numerical methods, which we detail, we demonstrate the fidelity of our model through simulations and comparisons with experiments.

0730-0301/2019/7-ART129 $15.00
https://doi.org/10.1145/3306346.3322962

Fundamental Solutions for Water Wave Animation

CAMILLE SCHRECK, IST Austria
CHRISTIAN HAFNER, IST Austria
CHRIS WOJTAN, IST Austria

This paper investigates the use of fundamental solutions for animating detailed linear water surface waves. We first propose an analytical solution for efficiently animating circular ripples in closed form. We then show how to adapt the method of fundamental solutions (MFS) to create ambient waves interacting with complex obstacles. Subsequently, we present a novel wavelet-based discretization which outperforms the state of the art MFS approach for simulating time-varying water surface waves with moving obstacles. Our results feature high-resolution spatial details, interactions with complex boundaries, and large open ocean domains. Our method compares favorably with previous work as well as known analytical solutions. We also present comparisons between our method and real world examples.

CCS Concepts: • **Computing methodologies → Physical simulation**.

Additional Key Words and Phrases: water waves, dispersion relation, Helmholtz equation, wave equation, fundamental solution, equivalent sources, interactive fluid simulation

ACM Reference Format:
Camille Schreck, Christian Hafner, and Chris Wojtan. 2019. Fundamental Solutions for Water Wave Animation. *ACM Trans. Graph.* 38, 4, Article 130 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3323002

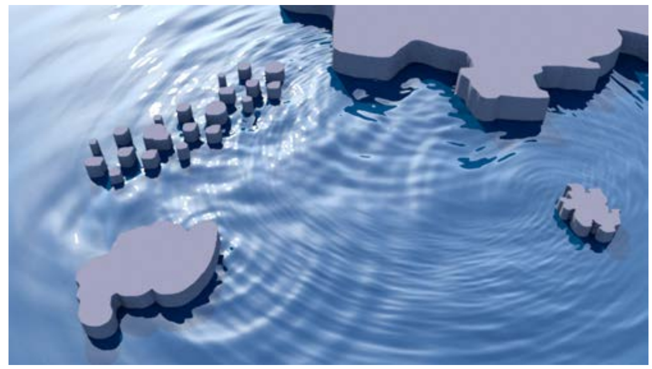

Fig. 1. We introduce the method of fundamental solutions to handle interactions between surface water waves with a large range of frequencies and complex detailed boundaries in an expansive open domain.

1 INTRODUCTION

This paper presents new strategies for efficiently animating water surface waves, especially for simulation in large, open domains with abundant high-frequency visual details and detailed boundaries. Such scenarios are particularly challenging for the state of the art in computer animation, which either rely on finite-difference approximations [Tessendorf 2004a], require a grid or mesh for the entire domain [Jeschke et al. 2018; Jeschke and Wojtan 2015], or require Lagrangian wave samples to fill the domain wherever details are present [Jeschke and Wojtan 2017; Yuksel et al. 2007]. Many of these approaches become prohibitively expensive for open (practically infinite) domains with extremely high-frequency capillary ripples. Our method aims to simulate detailed water ripples down to the level of individual raindrops and pebbles on a shoreline in an infinite, open ocean.

This work achieves detailed water animations by deriving fundamental solutions to the linear equations for water surface waves. Previous works used fundamental solutions (also referred to as "Green's functions") to compute efficient animations of elastic deformations [James and Pai 1999] and fracture [Hahn and Wojtan 2015, 2016; Zhu et al. 2015], sound propagation [James et al. 2006; Mehra et al. 2013], fluid vortices [Brochu et al. 2012; Cottet et al. 2000; Pfaff et al. 2012; Weißmann and Pinkall 2010], soap films [Da et al. 2015], three-dimensional liquids [Da et al. 2016; Keeler and Bridson 2014], keyframe animation [Barbič et al. 2012], and physically-based sculpting tools [De Goes and James 2017, 2018], to name a few. These fundamental solutions often enable unconditionally stable and extremely efficient analytical solutions for point sources, and linear combinations of them will produce mesh-free or boundary-only solutions (convenient for extremely large domains, and when high spatial frequencies make meshing difficult). However, to the best of our knowledge, our work is the first to investigate the use of fundamental solutions for animating water surface waves in Computer Graphics.

This work offers the following contributions to the state of the art in water animation:

- We introduce fundamental solutions for the animation of water surface waves in CG. Water waves have a frequency-dependent wave speed, so they present complications which are not solvable with previous techniques designed for linear elasticity and sound propagation. We introduce novel methods for handling this dispersive behavior.
- We apply analytical approximations to obtain a closed-form ripple function, which is useful for animating physically realistic rain drops in real-time.
- We introduce the use of *the method of fundamental solutions* for animating water waves, and we propose an interactive animation technique based on it. The resulting animation is not tied to a grid resolution, so we can animate infinite

Authors' addresses: Camille Schreck, IST Austria; Christian Hafner, IST Austria; Chris Wojtan, IST Austria.

0730-0301/2019/7-ART130 $15.00
https://doi.org/10.1145/3306346.3323002

Synthetic Silviculture: Multi-scale Modeling of Plant Ecosystems

MIŁOSZ MAKOWSKI, Adam Mickiewicz University
TORSTEN HÄDRICH, JAN SCHEFFCZYK, and DOMINIK L. MICHELS, KAUST
SÖREN PIRK, Google Brain
WOJTEK PAŁUBICKI, Adam Mickiewicz University

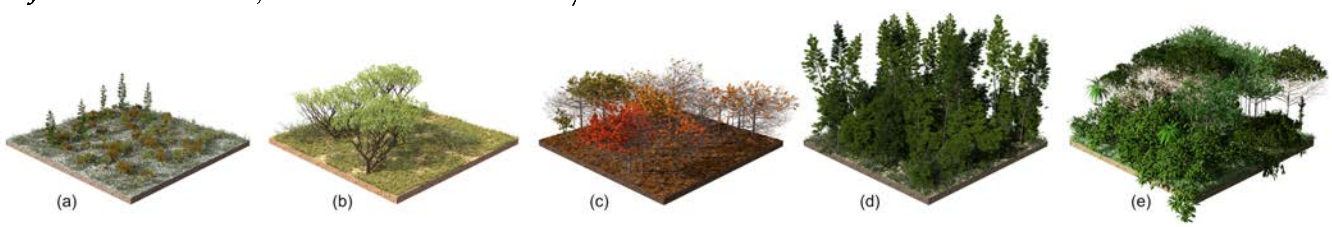

Fig. 1. Ecosystems of different biome types generated with our framework: a tundra with cold-adapted trees (a), a savanna with grass and acacia trees (b), a deciduous forest composed of maple trees (c), a boreal forest with tall pine trees (d), and a rain forest scene with a large variety of species (e). Our framework exploits inter- and intra-plant self-similarities to model plants and thereby allows us to interactively generate complex ecosystems.

Due to the enormous amount of detail and the interplay of various biological phenomena, modeling realistic ecosystems of trees and other plants is a challenging and open problem. Previous research on modeling plant ecologies has focused on representations to handle this complexity, mostly through geometric simplifications, such as points or billboards. In this paper we describe a multi-scale method to design large-scale ecosystems with individual plants that are realistically modeled and faithfully capture biological features, such as growth, plant interactions, different types of tropism, and the competition for resources. Our approach is based on leveraging inter- and intra-plant self-similarities for efficiently modeling plant geometry. We focus on the interactive design of plant ecosystems of up to 500K plants, while adhering to biological priors known in forestry and botany research. The introduced parameter space supports modeling properties of nine distinct plant ecologies while each plant is represented as a 3D surface mesh. The capabilities of our framework are illustrated through numerous models of forests, individual plants, and validations.

CCS Concepts: • **Computing methodologies → Interactive simulation**.

Additional Key Words and Phrases: Botanical Tree Models, Ecosystem Design, Natural Phenomena, Interactive Modeling, Plant Ecosystems, Self-Similarity, Self-Organization, Multi-Scale, Visual Models of Trees

ACM Reference Format:
Miłosz Makowski, Torsten Hädrich, Jan Scheffczyk, Dominik L. Michels, Sören Pirk, and Wojtek Pałubicki. 2019. Synthetic Silviculture: Multi-scale Modeling of Plant Ecosystems. *ACM Trans. Graph.* 38, 4, Article 131 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3323039

Authors' addresses: Miłosz Makowski, Adam Mickiewicz University, Umultowska 87, 61-614 Poznań, Poland; Torsten Hädrich; Jan Scheffczyk; Dominik L. Michels, KAUST, Visual Computing Center, Thuwal 23955, KSA; Sören Pirk, Google Brain, 1600 Amphitheatre Parkway, Mountain View, CA, 94043; Wojtek Pałubicki, Adam Mickiewicz University, Umultowska 87, 61-614 Poznań, Poland.

0730-0301/2019/7-ART131 $15.00
https://doi.org/10.1145/3306346.3323039

1 INTRODUCTION

Modeling trees and plants with all their nuances in structure and appearance is a thoroughly studied, and yet ongoing topic in visual computing. At the scale of plant ecosystems, this task is made even more challenging not only due to the enormous complexity of geometric detail, but also because of the intricacies of biological phenomena that affect how plants grow and interact. While existing approaches often model plant ecosystems by simplifying the involved geometry into more organized and hierarchical representations, such as voxels or layers, many applications benefit from detailed models that retain a plants' structure, faithful to all its defining features. This ranges from forests as content in 3D simulators and games, over urban and environmental planning applications, to research on ecosystems in forestry, and more recently the training of autonomous agents in virtual environments [Müller et al. 2018]. Detailed models of plant ecosystems, that even enable interactions with individual plants and their branches and leaves, play a key role in many situations. However, the computational costs for modeling these ecosystems is often beyond the capabilities of commodity rendering hardware. Furthermore, modeling the branching structures resulting from complex plant interactions as part of the growth process – inevitable for realistic ecosystems – requires in-depth knowledge of the biological processes, which is an intractable requirement for most content creators.

Traditionally, plant ecosystems are simulated by jointly generating plausible distributions of plant species and modeling their geometry [Deussen et al. 2002, 1998; Lane and Prusinkiewicz 2002]. Several approaches exists to capture the various levels of abstraction, such as volumetric textures [Bruneton and Neyret 2012], voxels [Jaeger and Teng 2003], or branch templates [Livny et al. 2011]. Choosing the appropriate level of detail scheme is critical for modeling plant ecosystems, and a few approaches have been proposed to enable simplifications, while also adhering to plant structure [Gumbau et al. 2011; Neubert et al. 2011]. Only more recently methods focus on realistic geometric representations for trees with an emphasize on individual parts [Xie et al. 2016; Zhang et al. 2017], and

PlanIT: Planning and Instantiating Indoor Scenes with Relation Graph and Spatial Prior Networks

KAI WANG, Brown University
YU-AN LIN, Brown University
BEN WEISSMANN, Brown University
MANOLIS SAVVA, Simon Fraser University
ANGEL X. CHANG, Simon Fraser University
DANIEL RITCHIE, Brown University

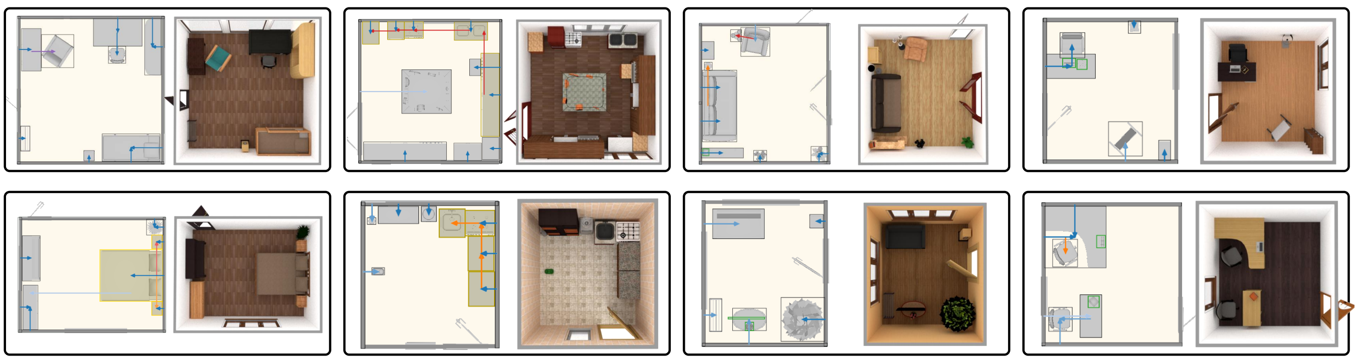

Fig. 1. We present PlanIT: a scene synthesis framework that unifies high-level *planning* of scene structure using a generative model over relation graphs with lower-level spatial *instantiation* using neural-guided search with spatial neural network modules. We first generate a relation graph with objects at the nodes and spatial or semantic relations at the edges (left images). Then, given the graph structure we select and place objects to instantiate a concrete 3D scene.

We present a new framework for interior scene synthesis that combines a high-level relation graph representation with spatial prior neural networks. We observe that prior work on scene synthesis is divided into two camps: object-oriented approaches (which reason about the set of objects in a scene and their configurations) and space-oriented approaches (which reason about what objects occupy what regions of space). Our insight is that the object-oriented paradigm excels at high-level *planning* of how a room should be laid out, while the space-oriented paradigm performs well at *instantiating* a layout by placing objects in precise spatial configurations. With this in mind, we present PlanIT, a layout-generation framework that divides the problem into two distinct *planning* and *instantiation* phases. PlanIT represents the "plan" for a scene via a relation graph, encoding objects as nodes and spatial/semantic relationships between objects as edges. In the planning phase, it uses a deep graph convolutional generative model to synthesize relation graphs. In the instantiation phase, it uses image-based convolutional network modules to guide a search procedure that places objects into the scene in a manner consistent with the graph. By decomposing the problem in this way, PlanIT generates scenes of comparable quality to those generated by prior approaches (as judged by both people and learned classifiers), while also providing the modeling flexibility of the intermediate relationship graph representation. These graphs allow the system to support applications such as scene synthesis from a partial graph provided by a user.

CCS Concepts: • **Computing methodologies** → **Computer graphics**; **Neural networks**; **Probabilistic reasoning**;

Additional Key Words and Phrases: indoor scene synthesis, object layout, neural networks, convolutional networks, deep learning, relationship graphs, graph generation

ACM Reference Format:
Kai Wang, Yu-an Lin, Ben Weissmann, Manolis Savva, Angel X. Chang, and Daniel Ritchie. 2019. PlanIT: Planning and Instantiating Indoor Scenes with Relation Graph and Spatial Prior Networks. *ACM Trans. Graph.* 38, 4, Article 132 (July 2019), 15 pages. https://doi.org/10.1145/3306346.3322941

Authors' addresses: Kai Wang, Brown University; Yu-an Lin, Brown University; Ben Weissmann, Brown University; Manolis Savva, Simon Fraser University; Angel X. Chang, Simon Fraser University; Daniel Ritchie, Brown University.

0730-0301/2019/7-ART132 $15.00
https://doi.org/10.1145/3306346.3322941

1 INTRODUCTION

People spend a large percentage of their lives indoors—in bedrooms, living rooms, kitchens, etc. As computer graphics reproduces the real world in increasing fidelity, the demand for virtual versions of such spaces also grows. Virtual and augmented reality experiences often take place in such environments. Online virtual interior design tools are available to help people redesign their own spaces [Planner5d 2017; RoomSketcher 2017]. Some furniture design companies now primarily advertise their products by rendering virtual scenes, as it is faster, cheaper, and more flexible to do so than to stage real-world

Content-aware Generative Modeling of Graphic Design Layouts

XINRU ZHENG[*], City University of Hong Kong
XIAOTIAN QIAO[*], City University of Hong Kong
YING CAO, City University of Hong Kong
RYNSON W.H. LAU, City University of Hong Kong

(a) Inputs (b) Content-aware layout generation (c) Adding user constraint

Fig. 1. Our probabilistic generative model can support content-aware layout generation. Given input images, design category, and keywords that summarize the text contents (which are either automatically extracted from the input text or directly provided by the user), as shown in (a), our method automatically generates multiple layouts that fit the visual and textual contents (b). The user may optionally express his/her design intent by roughly sketching some elements on a page, e.g., adding two image elements by sketching two green regions I_1 and I_2 and a headline element by sketching a red region H, as shown in the small diagram in (c). Our method will then generate a layout that matches the user's intent, i.e., two images and a header placed at the specified locations, as shown in the large diagram in (c). The input images are from Pexels.

Layout is fundamental to graphic designs. For visual attractiveness and efficient communication of messages and ideas, graphic design layouts often have great variation, driven by the contents to be presented. In this paper, we study the problem of content-aware graphic design layout generation. We propose a deep generative model for graphic design layouts that is able to synthesize layout designs based on the visual and textual semantics of user inputs. Unlike previous approaches that are oblivious to the input contents and rely on heuristic criteria, our model captures the effect of visual and textual contents on layouts, and implicitly learns complex layout structure variations from data without the use of any heuristic rules. To train our model, we build a large-scale magazine layout dataset with fine-grained layout annotations and keyword labeling. Experimental results show that our model can synthesize high-quality layouts based on the visual semantics of input images and keyword-based summary of input text. We also demonstrate that our model internally learns powerful features that capture the subtle interaction between contents and layouts, which are useful for layout-aware design retrieval.

CCS Concepts: • **Computing methodologies** → **Computer graphics**; *Machine learning approaches*; *Neural networks.*

Additional Key Words and Phrases: Graphic design, layout, content-aware, deep generative networks

ACM Reference Format:
Xinru Zheng, Xiaotian Qiao, Ying Cao, and Rynson W.H. Lau. 2019. Content-aware Generative Modeling of Graphic Design Layouts. *ACM Trans. Graph.* 38, 4, Article 133 (July 2019), 15 pages. https://doi.org/10.1145/3306346.3322971

1 INTRODUCTION

Layout is at the core of graphic designs, including magazines, posters, comics and webpages. A high-quality layout can benefit information presentation, guide reader attention and enhance visual attractiveness [Stribley 2015; Ying 2014]. Graphic design layout problems are receiving a growing interest in the graphics community in recent years. Some prior works try to model graphic design layouts for layout generation guided by style, perception and aesthetics [Cao et al. 2012, 2014; O'Donovan et al. 2014; Pang et al. 2016].

In graphic designs, layouts are especially created to frame contents (e.g., images and text) in order to present messages and ideas quickly and clearly. Therefore, rich layout variations in graphic designs are largely driven by the visual and textual contents to be presented [Prust 2010]. In other words, generating an effective graphic design layout requires understanding the visual content

[*]Xinru Zheng and Xiaotian Qiao are joint first authors.
Ying Cao is the corresponding author. This work was led by Rynson Lau.

Authors' addresses: Xinru Zheng, Department of Computer Science, City University of Hong Kong, xrzheng22@gmail.com; Xiaotian Qiao, Department of Computer Science, City University of Hong Kong, xt.qiao@my.cityu.edu.hk; Ying Cao, Department of Computer Science, City University of Hong Kong, caoying59@gmail.com; Rynson W.H. Lau, Department of Computer Science, City University of Hong Kong, rynson.lau@cityu.edu.hk.

0730-0301/2019/7-ART133 $15.00
https://doi.org/10.1145/3306346.3322971

Deep Inverse Rendering for High-resolution SVBRDF Estimation from an Arbitrary Number of Images

DUAN GAO, Tsinghua University and Microsoft Research Asia
XIAO LI, University of Science and Technology of China and Microsoft Research Asia
YUE DONG, Microsoft Research Asia
PIETER PEERS, College of William & Mary
KUN XU, Tsinghua University
XIN TONG, Microsoft Research Asia

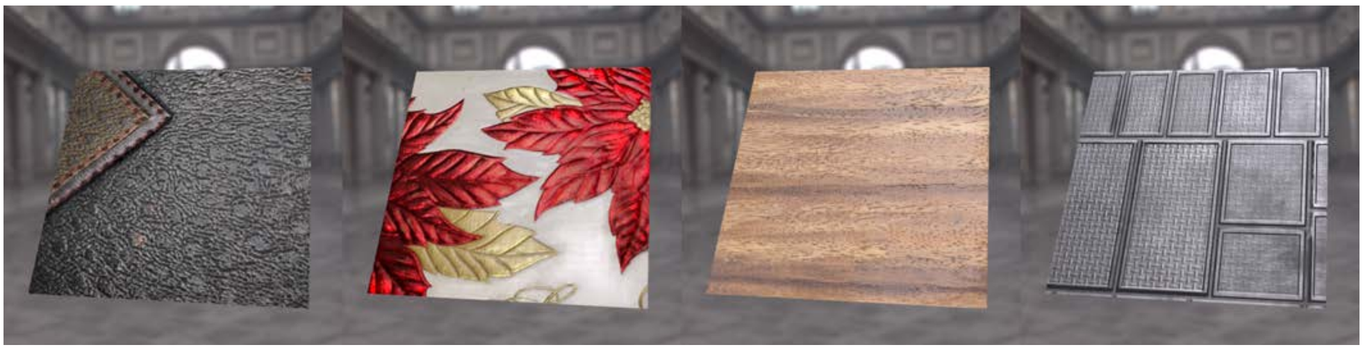

Fig. 1. Visualizations under natural lighting of four captured $1k$ resolution SVBRDFs estimated using our deep inverse rendering framework. The leather material (left) is reconstructed from just 2 input photographs captured with a mobile phone camera and flash, while the other materials are recovered from 20 input photographs.

In this paper we present a unified deep inverse rendering framework for estimating the spatially-varying appearance properties of a planar exemplar from an arbitrary number of input photographs, ranging from just a single photograph to many photographs. The precision of the estimated appearance scales from plausible when the input photographs fails to capture all the reflectance information, to accurate for large input sets. A key distinguishing feature of our framework is that it directly optimizes for the appearance parameters in a latent embedded space of spatially-varying appearance, such that no handcrafted heuristics are needed to regularize the optimization. This latent embedding is learned through a fully convolutional auto-encoder that has been designed to regularize the optimization. Our framework not only supports an arbitrary number of input photographs, but also at high resolution. We demonstrate and evaluate our deep inverse rendering solution on a wide variety of publicly available datasets.

CCS Concepts: • **Computing methodologies → Reflectance modeling**; *Image processing*.

Additional Key Words and Phrases: Material Capture, SVBRDF, Deep Learning, Auto-encoder.

ACM Reference Format:
Duan Gao, Xiao Li, Yue Dong, Pieter Peers, Kun Xu, and Xin Tong. 2019. Deep Inverse Rendering for High-resolution SVBRDF Estimation from an Arbitrary Number of Images. *ACM Trans. Graph.* 38, 4, Article 134 (July 2019), 15 pages. https://doi.org/10.1145/3306346.3323042

1 INTRODUCTION

Estimating the surface reflectance properties of a spatially-varying material is a challenging problem. Methods based on inverse rendering (e.g., [Dong et al. 2014; Hui et al. 2017]) can obtain accurate estimates for a sufficiently large number of input photographs. However, if the number of photographs is too low, such inverse rendering methods fail to produce plausible results. Recently, a number of techniques have been presented that, leveraging recent advances in deep learning, focus on achieving plausible results from just a single image [Deschaintre et al. 2018; Li et al. 2017, 2018a,b; Ye et al. 2018]. However, these methods fail to reproduce reflectance features that are ambiguous and/or not visible in the single input photograph. For example, specular features that are not excited by the incident lighting in the input photograph can only be inserted based on learned heuristics. Adding one or more photographs that provide

The first two authors contributed equally to this paper. Part of this work was performed while Pieter Peers visited Microsoft Research Asia.
Authors' addresses: D. Gao and K. Xu, Department of Computer Science and Technology, Tsinghua University, Beijing, China;, emails:gao-d17@mails.tsinghua.edu.cn, xukun@tsinghua.edu.cn; X. Li, University of Science and Technology of China, Hefei, China; email:pableetoli@gmail.com; Y. Dong and X. Tong, Microsoft Research Asia, Beijing, China; emails:yuedong, xtong@microsoft.com; P. Peers, Computer Science Department, College of William & Mary, Williamsburg, VA, 23187, USA; email:ppeers@siggraph.org.

0730-0301/2019/7-ART134 $15.00
https://doi.org/10.1145/3306346.3323042

A Similarity Measure for Material Appearance

MANUEL LAGUNAS, Universidad de Zaragoza, I3A, Spain
SANDRA MALPICA, Universidad de Zaragoza, I3A, Spain
ANA SERRANO, Universidad de Zaragoza, I3A, Spain
ELENA GARCES, Universidad Rey Juan Carlos, Spain
DIEGO GUTIERREZ, Universidad de Zaragoza, I3A, Spain
BELEN MASIA, Universidad de Zaragoza, I3A, Spain

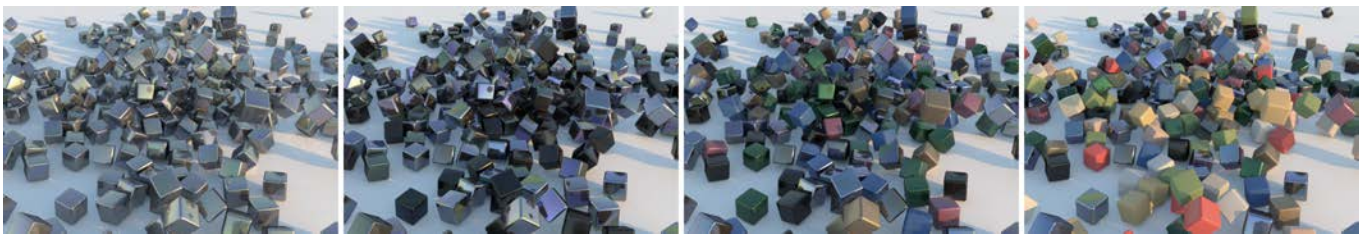

Fig. 1. The cubes in the leftmost image have all been rendered with the same aluminium material. Our similarity measure for material appearance can be used to automatically generate alternative depictions of the same scene, where the similarity of the materials varies in a controlled manner. The next three images show results with materials randomly chosen by progressively extending the search distance from the original aluminium, from similar in appearance to farther away materials within the same dataset.

We present a model to measure the similarity in appearance between different materials, which correlates with human similarity judgments. We first create a database of 9,000 rendered images depicting objects with varying materials, shape and illumination. We then gather data on perceived similarity from crowdsourced experiments; our analysis of over 114,840 answers suggests that indeed a shared perception of appearance similarity exists. We feed this data to a deep learning architecture with a novel loss function, which learns a feature space for materials that correlates with such perceived appearance similarity. Our evaluation shows that our model outperforms existing metrics. Last, we demonstrate several applications enabled by our metric, including appearance-based search for material suggestions, database visualization, clustering and summarization, and gamut mapping.

CCS Concepts: • **Computing methodologies** → *Appearance and texture representations*; *Perception*; • **Computer systems organization** → *Neural networks.*

Additional Key Words and Phrases: Material appearance, neural networks, physically based material perception

Authors' addresses: Manuel Lagunas, Universidad de Zaragoza, I3A, Spain, Zaragoza, mlagunas@unizar.es; Sandra Malpica, Universidad de Zaragoza, I3A, Spain, Zaragoza, smalpica@unizar.es; Ana Serrano, Universidad de Zaragoza, I3A, Spain, Zaragoza, anase@unizar.es; Elena Garces, Universidad Rey Juan Carlos, Spain, Madrid, elena.garces@urjc.es; Diego Gutierrez, Universidad de Zaragoza, I3A, Spain, Zaragoza, diegog@unizar.es; Belen Masia, Universidad de Zaragoza, I3A, Spain, Zaragoza, bmasia@unizar.es.

0730-0301/2019/7-ART135 $15.00
https://doi.org/10.1145/3306346.3323036

ACM Reference Format:
Manuel Lagunas, Sandra Malpica, Ana Serrano, Elena Garces, Diego Gutierrez, and Belen Masia. 2019. A Similarity Measure for Material Appearance. *ACM Trans. Graph.* 38, 4, Article 135 (July 2019), 12 pages. https://doi.org/10.1145/3306346.3323036

1 INTRODUCTION

Humans are able to recognize materials, compare their appearance, or even infer many of their key properties effortlessly, just by briefly looking at them. Many works propose classification techniques, although it seems clear that labels do not suffice to capture the richness of our subjective experience with real-world materials [Fleming 2017]. Unfortunately, the underlying perceptual process of material recognition is complex, involving many distinct variables; such process is not yet completely understood [Anderson 2011; Fleming 2014; Maloney and Brainard 2010].

Given the large number of parameters involved in our perception of materials, many works have focused on individual attributes (such as the perception of gloss [Pellacini et al. 2000; Wills et al. 2009], or translucency [Gkioulekas et al. 2015]), while others have focused on particular applications like material synthesis [Zsolnai-Fehér et al. 2018], editing [Serrano et al. 2016], or filtering [Jarabo et al. 2014]. However, the fundamentally difficult problem of establishing a *similarity measure for material appearance* remains an open problem. Material appearance can be defined as "the visual impression we have of a material" [Dorsey et al. 2010]; as such, it depends not only on the BRDF of the material, but also on external factors like lighting or geometry, as well as human judgement [Adelson 2001; Fleming 2014]. This is different from the common notion of image similarity (devoted to finding detectable differences between images, e.g., [Wang et al. 2004]), or from similarity in BRDF space (which has

Using Moments to Represent Bounded Signals for Spectral Rendering

CHRISTOPH PETERS, Karlsruhe Institute of Technology, Germany
SEBASTIAN MERZBACH, University of Bonn, Germany
JOHANNES HANIKA, Karlsruhe Institute of Technology and Weta Digital, Germany
CARSTEN DACHSBACHER, Karlsruhe Institute of Technology, Germany

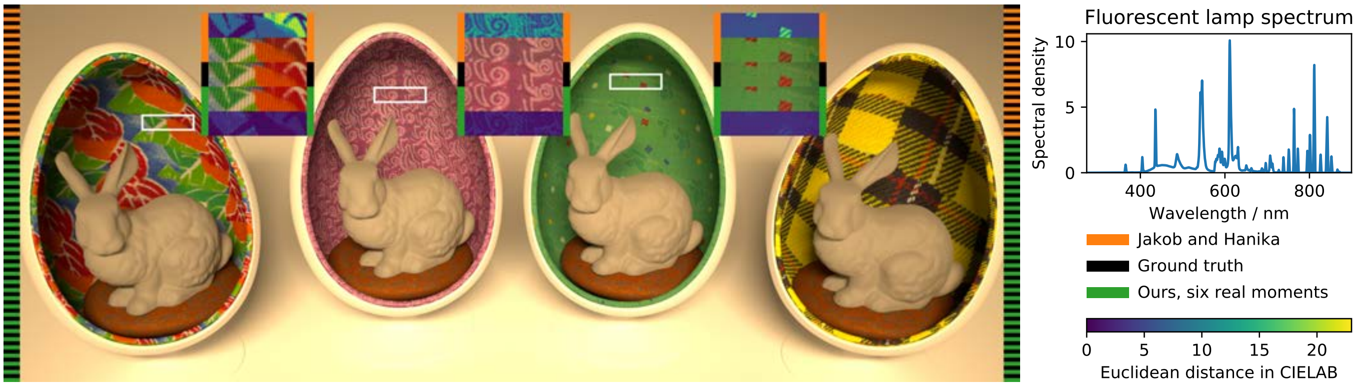

Fig. 1. We store six moments per texel in 8 bytes to represent the spectral albedo textures in this scene. Our novel reconstruction method (green stripes) uses this data to approximate the actual reflectance spectrum. Thus, color reproduction is accurate, even with challenging illuminants such as the fluorescent lamp used here. The interleaving in the upper part, the magnified insets and the insets with distances in CIELAB show how strongly state of the art tristimulus techniques [Jakob and Hanika, 2019] (orange stripes) deviate from ground truth (black stripes).

We present a compact and efficient representation of spectra for accurate rendering using more than three dimensions. While tristimulus color spaces are sufficient for color display, a spectral renderer has to simulate light transport per wavelength. Consequently, emission spectra and surface albedos need to be known at each wavelength. It is practical to store dense samples for emission spectra but for albedo textures, the memory requirements of this approach are unreasonable. Prior works that approximate dense spectra from tristimulus data introduce strong errors under illuminants with sharp peaks and in indirect illumination. We represent spectra by an arbitrary number of Fourier coefficients. However, we do not use a common truncated Fourier series because its ringing could lead to albedos below zero or above one. Instead, we present a novel approach for reconstruction of bounded densities based on the theory of moments. The core of our technique is our bounded maximum entropy spectral estimate. It uses an efficient closed form to compute a smooth signal between zero and one that matches the given Fourier coefficients exactly. Still, a ground truth that localizes all of its mass around a few wavelengths can be reconstructed adequately. Therefore, our representation covers the full gamut of valid reflectances. The resulting textures are compact because each coefficient can be stored in 10 bits. For compatibility with existing tristimulus assets, we implement a mapping from tristimulus color spaces to three Fourier coefficients. Using three coefficients, our technique gives state of the art results without some of the drawbacks of related work. With four to eight coefficients, our representation is superior to all existing representations. Our focus is on offline rendering but we also demonstrate that the technique is fast enough for real-time rendering.

CCS Concepts: • **Computing methodologies → Reflectance modeling**.

Additional Key Words and Phrases: spectral rendering, reflectance spectra, emission spectra, bounded signals, bounded maximum entropy spectral estimate, bounded MESE, Fourier coefficients, trigonometric moments

ACM Reference Format:
Christoph Peters, Sebastian Merzbach, Johannes Hanika, and Carsten Dachsbacher. 2019. Using Moments to Represent Bounded Signals for Spectral Rendering. *ACM Trans. Graph.* 38, 4, Article 136 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3322964

Authors' addresses: Christoph Peters, christoph.peters@kit.edu, Karlsruhe Institute of Technology, Am Fasanengarten 5, 76131, Karlsruhe, Germany; Sebastian Merzbach, merzbach@cs.uni-bonn.de, University of Bonn, Endenicher Allee 19a, 53115, Bonn, Germany; Johannes Hanika, hanika@kit.edu, Karlsruhe Institute of Technology, and Weta Digital, Karlsruhe, Germany; Carsten Dachsbacher, dachsbacher@kit.edu, Karlsruhe Institute of Technology, Karlsruhe, Germany.

0730-0301/2019/7-ART136 $15.00
https://doi.org/10.1145/3306346.3322964

1 INTRODUCTION

Color representation is one of the most fundamental problems in computer graphics. The ubiquitous tristimulus color spaces are adequate for color display, yet insufficient for accurate rendering. A physically based renderer should simulate light transport at each wavelength and thus dense spectra for light emission and surface reflectances need to be known. The accuracy of techniques that infer dense spectra from tristimulus data is limited as the problem is severely underconstrained [Jakob and Hanika, 2019, Meng et al., 2015, Otsu et al., 2018, Smits, 1999]. Especially for scenes with spiky

Accurate Appearance Preserving Prefiltering for Rendering Displacement-Mapped Surfaces

LIFAN WU, University of California, San Diego
SHUANG ZHAO, University of California, Irvine
LING-QI YAN, University of California, Santa Barbara
RAVI RAMAMOORTHI, University of California, San Diego

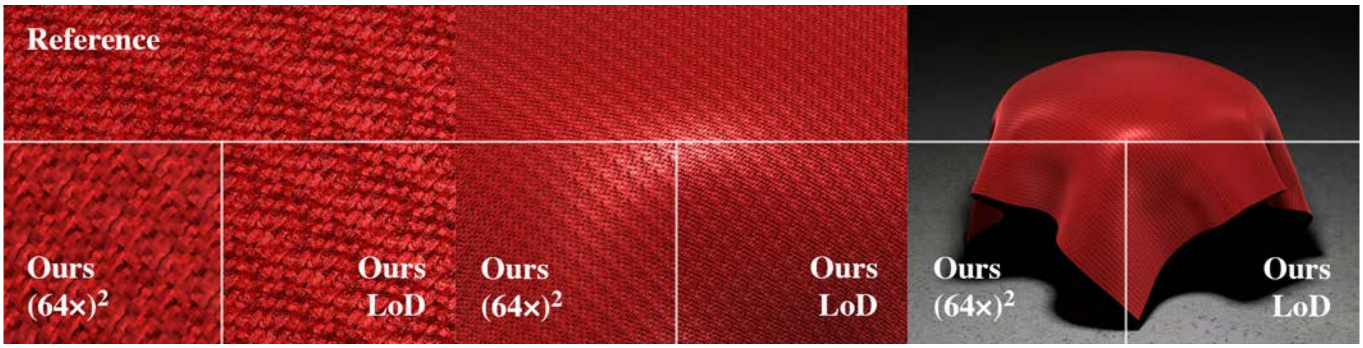

Fig. 1. We present a new approach to prefilter high-resolution displacement maps while preserving the input appearance. High-resolution displacement maps can produce rich geometric details (top-left) but they are difficult to prefilter. Our prefiltered model handles the change of shadowing, masking and interreflections caused by downsampling the displacement map. At a single scale, although the detailed micro-structures are different (see ours $(64\times)^2$ at left), our prefiltered model preserves the original appearance accurately when we view the object from a distance (see the right image). We can also combine our models at multiple downsampling scales to form a mipmap, enabling accurate and anti-aliased LoD rendering.

Prefiltering the reflectance of a displacement-mapped surface while preserving its overall appearance is challenging, as smoothing a displacement map causes complex changes of illumination effects such as shadowing-masking and interreflection. In this paper, we introduce a new method that prefilters displacement maps and BRDFs jointly and constructs SVBRDFs at reduced resolutions. These SVBRDFs preserve the appearance of the input models by capturing both shadowing-masking and interreflection effects. To express our appearance-preserving SVBRDFs efficiently, we leverage a new representation that involves spatially varying NDFs and a novel scaling function that accurately captures micro-scale changes of shadowing, masking, and interreflection effects. Further, we show that the 6D scaling function can be factorized into a 2D function of surface location and a 4D function of direction. By exploiting the smoothness of these functions, we develop a simple and efficient factorization method that does not require computing the full scaling function. The resulting functions can be represented at low resolutions (e.g., 4^2 for the spatial function and 15^4 for the angular function), leading to minimal additional storage. Our method generalizes well to different types of geometries beyond Gaussian surfaces. Models prefiltered using our approach at different scales can be combined to form mipmaps, allowing accurate and anti-aliased level-of-detail (LoD) rendering.

CCS Concepts: • **Computing methodologies → Rendering**.

Additional Key Words and Phrases: multi-resolution, level of detail, prefiltering, global illumination

ACM Reference Format:
Lifan Wu, Shuang Zhao, Ling-Qi Yan, and Ravi Ramamoorthi. 2019. Accurate Appearance Preserving Prefiltering for Rendering Displacement-Mapped Surfaces. *ACM Trans. Graph.* 38, 4, Article 137 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3322936

1 INTRODUCTION

High-resolution displacement maps are commonly used to describe detailed micro-geometries that can produce richly diverse appearances. Compared to normal mapping, displacement mapping is more physically consistent and can offer more realistic self-shadowing and silhouettes. However, such realism comes at the cost of difficult prefiltering: smoothing a displacement map usually weakens its intrinsic shadowing and results in brightened overall appearance. Therefore, rendering a high-resolution displacement map without introducing severe aliasing generally requires significant supersampling, which is computationally expensive.

Previous displacement mapping techniques such as LEAN [Olano and Baker 2010] and LEADR [Dupuy et al. 2013] can produce anti-aliased renderings of rough surfaces. However, they assume the normals of the input surfaces to have Beckmann distributions, which is usually violated in practice and fundamentally limits the accuracy

Authors' addresses: Lifan Wu, University of California, San Diego, liw086@eng.ucsd.edu; Shuang Zhao, University of California, Irvine, shz@ics.uci.edu; Ling-Qi Yan, University of California, Santa Barbara, lingqi@cs.ucsb.edu; Ravi Ramamoorthi, University of California, San Diego, ravir@cs.ucsd.edu.

0730-0301/2019/7-ART137 $15.00
https://doi.org/10.1145/3306346.3322936