Selective Padding for Polycube-based Hexahedral Meshing

Authors: Gianmarco Cherchi, Pierre Alliez, Riccardo Scateni, Max Lyon, David Bommes

First published: 03 January 2019

DOI: 10.1111/cgf.13593

URL: https://diglib.eg.org/handle/10.1111/cgf13593
https://doi.org/10.1111/cgf.13593

Abstract: Hexahedral meshes generated from polycube mapping often exhibit a low number of singularities but also poor-quality elements located near the surface. It is thus necessary to improve the overall mesh quality, in terms of the minimum scaled Jacobian (MSJ) or average SJ (ASJ). Improving the quality may be obtained via global padding (or pillowing), which pushes the singularities inside by adding an extra layer of hexahedra on the entire domain boundary. Such a global padding operation suffers from a large increase of complexity, with unnecessary hexahedra added. In addition, the quality of elements near the boundary may decrease. We propose a novel optimization method which inserts sheets of hexahedra so as to perform selective padding, where it is most needed for improving the mesh quality. A sheet can pad part of the domain boundary, traverse the domain and form singularities. Our global formulation, based on solving a binary problem, enables us to control the balance between quality improvement, increase of complexity and number of singularities. We show in a series of experiments that our approach increases the MSJ value and preserves (or even improves) the ASJ, while adding fewer hexahedra than global padding.

Filtered Quadrics for High Speed Geometry Smoothing and Clustering

Authors: Hélène Legrand, Jean-Marc Thiery, Tamy Boubekeur

First published: 12 February 2019

DOI: 10.1111/cgf.13597

URL: https://diglib.eg.org/handle/10.1111/cgf13597
https://doi.org/10.1111/cgf.13597

Abstract: Modern 3D capture pipelines produce dense surface meshes at high speed, which challenge geometric operators to process such massive data on-the-fly. In particular, aiming at instantaneous feature-preserving smoothing and clustering disqualifies global variational optimizers and one usually relies on high-performance parallel kernels based on simple measures performed on the positions and normal vectors associated with the surface vertices. Although these operators are effective on small supports, they fail at properly capturing larger scale surface structures. To cope with this problem, we propose to enrich the surface representation with filtered quadrics, a compact and discriminating range space to guide processing. Compared to normal-based approaches, this additional vertex attribute significantly improves feature preservation for fast bilateral filtering and mode-seeking clustering, while exhibiting a linear memory cost in the number of vertices and retaining the simplicity of convolutional filters. In particular, the overall performance of our approach stems from its natural compatibility with modern fine-grained parallel computing architectures such as graphics processor units (GPU). As a result, filtered quadrics offer a superior ability to handle a broad spectrum of frequencies and preserve large salient structures, delivering meshes on-the-fly for interactive and streaming applications, as well as quickly processing large data collections, instrumental in learning-based geometry analysis.

Stylized Image Triangulation

Authors: Kai Lawonn, Tobias Günther

First published: 18 July 2018

DOI: 10.1111/cgf.13526

URL: https://diglib.eg.org/handle/10.1111/cgf13526
https://doi.org/10.1111/cgf.13526

Abstract: The art of representing images with triangles is known as image triangulation, which purposefully uses abstraction and simplification to guide the viewer's attention. The manual creation of image triangulations is tedious and thus several tools have been developed in the past that assist in the placement of vertices by means of image feature detection and subsequent Delaunay triangulation. In this paper, we formulate the image triangulation process as an optimization problem. We provide an interactive system that optimizes the vertex locations of an image triangulation to reduce the root mean squared approximation error. Along the way, the triangulation is incrementally refined by splitting triangles until certain refinement criteria are met. Thereby, the calculation of the energy gradients is expensive and thus we propose an efficient rasterization-based GPU implementation. To ensure that artists have control over details, the system offers a number of direct and indirect editing tools that split, collapse and re-triangulate selected parts of the image. For final display, we provide a set of rendering styles, including constant colours, linear gradients, tonal art maps and textures. Finally, we demonstrate temporal coherence for animations and compare our method with existing image triangulation tools.

Learning A Stroke-Based Representation for Fonts

Authors: Elena Balashova, Amit H. Bermano, Vladimir G. Kim, Stephen DiVerdi, Aaron Hertzmann, Thomas Funkhouser

First published: 05 October 2018

DOI: 10.1111/cgf.13540

URL: https://diglib.eg.org/handle/10.1111/cgf13540
https://doi.org/10.1111/cgf.13540

Abstract: Designing fonts and typefaces is a difficult process for both beginner and expert typographers. Existing workflows require the designer to create every glyph, while adhering to many loosely defined design suggestions to achieve an aesthetically appealing and coherent character set. This process can be significantly simplified by exploiting the similar structure character glyphs present across different fonts and the shared stylistic elements within the same font. To capture these correlations, we propose learning a stroke-based font representation from a collection of existing typefaces. To enable this, we develop a stroke-based geometric model for glyphs, a fitting procedure to reparametrize arbitrary fonts to our representation. We demonstrate the effectiveness of our model through a manifold learning technique that estimates a low-dimensional font space. Our representation captures a wide range of everyday fonts with topological variations and naturally handles discrete and continuous variations, such as presence and absence of stylistic elements as well as slants and weights. We show that our learned representation can be used for iteratively improving fit quality, as well as exploratory style applications such as completing a font from a subset of observed glyphs, interpolating or adding and removing stylistic elements in existing fonts.

Efficient Computation of Smoothed Exponential Maps

Authors: Philipp Herholz, Marc Alexa

First published: 14 March 2019

DOI: 10.1111/cgf.13607

URL: https://doi.org/10.1111/cgf.13607

Abstract: Many applications in geometry processing require the computation of local parameterizations on a surface mesh at interactive rates. A popular approach is to compute local exponential maps, i.e. parameterizations that preserve distance and angle to the origin of the map. We extend the computation of geodesic distance by heat diffusion to also determine angular information for the geodesic curves. This approach has two important benefits compared to fast approximate as well as exact forward tracing of the distance function: First, it allows generating smoother maps, avoiding discontinuities. Second, exploiting the factorization of the global Laplace-Beltrami operator of the mesh and using recent localized solution techniques, the computation is more efficient even compared to fast approximate solutions based on Dijkstra's algorithm.

MegaViews: Scalable Many-View Rendering with Concurrent Scene-View Hierarchy Traversal

Authors: Timothy R. Kol, Pablo Bauszat, Sungkil Lee, Elmar Eisemann

First published: 12 July 2018

DOI: 10.1111/cgf.13527

URL: https://diglib.eg.org/handle/10.1111/cgf13527
https://doi.org/10.1111/cgf.13527

Abstract: We present a scalable solution to render complex scenes from a large amount of viewpoints. While previous approaches rely either on a scene or a view hierarchy to process multiple elements together, we make full use of both, enabling sublinear performance in terms of views and scene complexity. By concurrently traversing the hierarchies, we efficiently find shared information among views to amortize rendering costs. One example application is many-light global illumination. Our solution accelerates shadow map generation for virtual point lights, whose number can now be raised to over a million while maintaining interactive rates.

Real-time Human Shadow Removal in a Front Projection System

Authors: Jaedong Kim, Hyunggoog Seo, Seunghoon Cha, Junyong Noh

First published: 01 October 2018

DOI: 10.1111/cgf.13541

URL: https://diglib.eg.org/handle/10.1111/cgf13541
https://doi.org/10.1111/cgf.13541

Abstract: When a person is located between a display and an operating projector, a shadow is cast on the display. The shadow on the display may eliminate important visual information and therefore adversely affect the viewing experiences. There have been various attempts to remove the human shadow cast on

a projection display by using multiple projectors. While previous approaches successfully removed the shadow region when a person moderately moves around or stands stationary in front of the display, there is still an afterimage effect due to the lack of consideration of the limb motion of the person. We propose a new real-time approach to removing the shadow cast by a person who dynamically interacts with the display, making limb motions in a front projection system. The proposed method utilizes a human skeleton obtained from a depth camera to track the posture of the person which changes over time. A model that consists of spheres and conical frustums is constructed based on the skeleton information in order to represent volumetric information of the person being tracked. Our method precisely estimates the shadow region by projecting the volumetric model onto the display. In addition, employment of intensity masks that are built based on a distance field helps suppress the afterimage of the shadow that appears when the person moves abruptly. It also helps blend the projected overlapping images from different projectors and show one smoothly combined display. The experiment results verify that our approach removes the shadow of a person effectively in a front projection environment and is fast enough to achieve real-time performance.

Gradient-Guided Local Disparity Editing

Authors: Leonardo Scandolo, Pablo Bauszat, and Elmar Eisemann

First published: 06 September 2018

DOI: 10.1111/cgf.13537

URL: https://diglib.eg.org/handle/10.1111/cgf13537
https://doi.org/10.1111/cgf.13537

Abstract: Stereoscopic 3D technology gives visual content creators a new dimension of design when creating images and movies. While useful for conveying emotion, laying emphasis on certain parts of the scene, or guiding the viewer's attention, editing stereo content is a challenging task. Not respecting comfort zones or adding incorrect depth cues, for example depth inversion, leads to a poor viewing experience. In this paper, we present a solution for editing stereoscopic content that allows an artist to impose disparity constraints and removes resulting depth conflicts using an optimization scheme. Using our approach, an artist only needs to focus on important high-level indications that are automatically made consistent with the entire scene while avoiding contradictory depth cues and respecting viewer comfort.

Urban Walkability Design using Virtual Population Simulation

Authors: Caluadewa Dharshaka Tharindu Mathew, Paulo R. Knob, Soraia Raupp Musse, Daniel G. Aliaga

First published: 08 October 2018

DOI: 10.1111/cgf.13585

URL: https://diglib.eg.org/handle/10.1111/cgf13585
https://doi.org/10.1111/cgf.13585

Abstract: We present a system to generate a procedural environment that produces a desired crowd behaviour. Instead of altering the behavioural parameters of the crowd itself, we automatically alter the environment to yield such desired crowd behaviour. This novel inverse approach is useful both to crowd simulation in virtual environments and to urban crowd planning applications. Our approach tightly integrates and extends a space discretization crowd simulator with inverse procedural modelling. We extend crowd simulation by goal exploration (i.e. agents are initially unaware of the goal locations), variable-appealing sign usage and

several acceleration schemes. We use Markov chain Monte Carlo to quickly explore the solution space and yield interactive design. We have applied our method to a variety of virtual and real-world locations, yielding one order of magnitude faster crowd simulation performance over related methods and several fold improvement of crowd indicators.

Data-Driven Crowd Motion Control with Multitouch Gestures

Authors: Yijun Shen, Joseph Henry,He Wang,Edmond S. L. Ho, Taku Komura, Hubert P. H. Shum

First published: 26 March 2018

DOI: 10.1111/cgf.13333

URL: https://diglib.eg.org/handle/10.1111/cgf13333
https://doi.org/10.1111/cgf.13333

Abstract: Controlling a crowd using multi-touch devices appeals to the computer games and animation industries, as such devices provide a high-dimensional control signal that can effectively define the crowd formation and movement. However, existing works relying on pre-defined control schemes require the users to learn a scheme that may not be intuitive. We propose a data-driven gesture-based crowd control system, in which the control scheme is learned from example gestures provided by different users. In particular, we build a database with pairwise samples of gestures and crowd motions. To effectively generalize the gesture style of different users, such as the use of different numbers of fingers, we propose a set of gesture features for representing a set of hand gesture trajectories. Similarly, to represent crowd motion trajectories of different numbers of characters over time, we propose a set of crowd motion features that are extracted from a Gaussian mixture model. Given a run-time gesture, our system extracts the K nearest gestures from the database and interpolates the corresponding crowd motions in order to generate the run-time control. Our system is accurate and efficient, making it suitable for real-time applications such as real-time strategy games and interactive animation controls.

Denoising Deep Monte Carlo Renderings

Authors: Delio Vicini, David Adler, Jan Novák, Fabrice Rousselle, Brent Burley

First published: 27 August 2018

DOI: 10.1111/cgf.13533

URL: https://diglib.eg.org/handle/10.1111/cgf13533
https://doi.org/10.1111/cgf.13533

Abstract: We present a novel algorithm to denoise deep Monte Carlo renderings, in which pixels contain multiple colour values, each for a different range of depths. Deep images are a more expressive representation of the scene than conventional flat images. However, since each depth bin receives only a fraction of the flat pixel's samples, denoising the bins is harder due to the less accurate mean and variance estimates. Furthermore, deep images lack a regular structure in depth-the number of depth bins and their depth ranges vary across pixels. This prevents a straightforward application of patch-based distance metrics frequently used to improve the robustness of existing denoising filters. We address these constraints by combining a flat image-space non-local means filter operating on pixel colours with a deep cross-bilateral filter operating on auxiliary features (albedo, normal, etc.). Our approach significantly reduces noise in deep images while preserving their structure. To our best knowledge, our algorithm is the first to enable

efficient deep-compositing workflows with denoised Monte Carlo renderings. We demonstrate the performance of our filter on a range of scenes highlighting the challenges and advantages of denoising deep images.

Increasing of Spatial Resolutions of BTF Measurement with Scheimpflug Imaging

Authors: Vlastimil Havran, Jan Hošek, Šarka Němcová, Jiří Čáp

First published: 04 January 2019

DOI: 10.1111/cgf.13594

URL: https://diglib.eg.org/handle/10.1111/cgf13594
https://doi.org/10.1111/cgf.13594

Abstract: We present an improved way of acquiring spatially varying surface reflectance represented by a bidirectional texture function (BTF). Planar BTF samples are measured as images at several inclination angles which puts constraints on the minimum depth of field of cameras used in the measurement instrument. For standard perspective imaging, we show that the size of a sample measured and the achievable spatial resolution are strongly interdependent and limited by diffraction at the lens' aperture. We provide a formula for this relationship. We overcome the issue of the required depth of field by using Scheimpflug imaging further enhanced by an anamorphic attachment. The proposed optics doubles the spatial resolution of images compared to standard perspective imaging optics. We built an instrument prototype with the proposed optics that is portable and allows for measurement on site. We show rendered images using the visual appearance measured by the instrument prototype.

Controllable Image-Based Transfer of Flow Phenomena

Authors: Carles Bosch, Gustavo Patow

First published: 26 July 2018

DOI: 10.1111/cgf.13530

URL: https://diglib.eg.org/handle/10.1111/cgf13530
https://doi.org/10.1111/cgf.13530

Abstract: Modelling flow phenomena and their related weathering effects is often cumbersome due their dependence on the environment, materials and geometric properties of objects in the scene. Example-based modelling provides many advantages for reproducing real textures, but little effort has been devoted to reproducing and transferring complex phenomena. In order to produce realistic flow effects, it is possible to take advantage of the widespread availability of flow images on the Internet, which can be used to gather key information about the flow. In this paper, we present a technique that allows the transfer of flow phenomena between photographs, adapting the flow to the target image and giving the user flexibility and control through specifically tailored parameters. This is done through two types of control curves: a fitted theoretical curve to control the mass of deposited material, and an extended colour map for properly adapting to the target appearance. In addition, our method filters and warps the input flow in order to account for the geometric details of the target surface. This leads to a fast and intuitive approach to easily transfer phenomena between images, providing a set of simple and intuitive parameters to control the process.