Computer Graphics Forum Technical Papers in EG2018

Easy Generation of Facial Animation Using Motion Graphs

J. Serra[1,2], O. Cetinaslan[1], S. Ravikumar[2], V. Orvalho[1] and D. Cosker[2]
1. Faculdade de Ciências da Universidade do Porto, Portugal
2. University of Bath, UK

First published: 13 June 2017
DOI: 10.1111/cgf.13218

Abstract: Facial animation is a time-consuming and cumbersome task that requires years of experience and/or a complex and expensive set-up. This becomes an issue, especially when animating the multitude of secondary characters required, e.g. in films or video-games. We address this problem with a novel technique that relies on motion graphs to represent a landmarked database. Separate graphs are created for different facial regions, allowing a reduced memory footprint compared to the original data. The common poses are identified using a Euclidean-based similarity metric and merged into the same node. This process traditionally requires a manually chosen threshold, however, we simplify it by optimizing for the desired graph compression. Motion synthesis occurs by traversing the graph using Dijkstra's algorithm, and coherent noise is introduced by swapping some path nodes with their neighbours. Expression labels, extracted from the database, provide the control mechanism for animation. We present a way of creating facial animation with reduced input that automatically controls timing and pose detail. Our technique easily fits within video-game and crowd animation contexts, allowing the characters to be more expressive with less effort. Furthermore, it provides a starting point for content creators aiming to bring more life into their characters.

Enhanced Visualization of Detected 3D Geometric Differences

Gianpaolo Palma[1], Manuele Sabbadin[1,2], Massimiliano Corsini[1] and Paolo Cignoni[1]
1. Visual Computing Lab, ISTI, CNR, Pisa, Italy
2. University of Pisa, Italy

First published: 28 June 2017
DOI: 10.1111/cgf.13239

Abstract: The wide availability of 3D acquisition devices makes viable their use for shape monitoring. The current techniques for the analysis of time-varying data can efficiently detect actual significant geometric changes and rule out differences due to irrelevant variations (such as sampling, lighting and coverage). On the other hand, the effective visualization of such detected changes can be challenging when we want to show at the same time the original appearance of the 3D model. In this paper, we propose a dynamic technique for the effective visualization of detected differences between two 3D scenes. The presented approach, while retaining the original appearance, allows the user to switch between the two models in a way that enhances the geometric differences that have been detected as significant. Additionally, the same technique is able to visually hides the other negligible, yet visible, variations. The main idea is to use two distinct screen space time-based interpolation functions for the significant 3D differences and for the small variations to hide. We have validated the proposed approach in a user study on a different class of datasets, proving the objective and subjective effectiveness of the method.

Super-Resolution of Point Set Surfaces Using Local Similarities

Azzouz Hamdi-Cherif[1,2], Julie Digne[1,3] and Raphaëlle Chaine[1,3]
1. LIRIS, CNRS, France
2. INSA de Lyon, France
3. Université Lyon, France

First published: 29 June 2017
DOI: 10.1111/cgf.13216

Abstract: Three-dimensional scanners provide a virtual representation of object surfaces at some given precision that depends on many factors such as the object material, the quality of the laser ray or the resolution of the camera. This precision may even vary over the surface, depending, for example, on the distance to the scanner which results in uneven and unstructured point sets, with an uncertainty on the coordinates. To enhance the quality of the scanner output, one usually resorts to local surface interpolation between measured points. However, object surfaces often exhibit interesting statistical features such as repetitive geometric textures. Building on this property, we propose a new approach for surface super-resolution that detects repetitive patterns or self-similarities and exploits them to improve the scan resolution by aggregating scattered measures. In contrast with other surface super-resolution methods, our algorithm has two important advantages. First, when handling multiple scans, it does not rely on surface registration. Second, it is able to produce super-resolution from even a single scan. These features are made possible by a new local shape description able to capture differential properties of order above 2. By comparing those descriptors, similarities are detected and used to generate a high-resolution surface. Our results show a clear resolution gain over state-of-the-art interpolation methods.

On the Stability of Functional Maps and Shape Difference Operators

R. Huang[1], F. Chazal[1] and M. Ovsjanikov[2]
1. DataShape team, INRIA Saclay, France
2. LIX, Ecole Polytechnique, France

First published: 4 July 2017
DOI: 10.1111/cgf.13238

Abstract: In this paper, we provide stability guarantees for two frameworks that are based on the notion of functional maps—the framework of shape difference operators and the one of analyzing and visualizing the deformations between shapes. We consider two types of perturbations in our analysis: one is on the input shapes and the other is on the change in scale. In theory, we formulate and justify the robustness that has been observed in practical implementations of those frameworks. Inspired by our theoretical results, we propose a pipeline for constructing shape difference operators on point clouds and show numerically that the results are robust and informative. In particular, we show that both the shape difference operators and the derived areas of highest distortion are stable with respect to changes in shape representation and change of scale. Remarkably, this is in contrast with the well-known instability of the eigenfunctions of the Laplace–Beltrami operator computed on point clouds compared to those obtained on triangle meshes.

Olfaction and Selective Rendering

Carlo Harvey[1], Thomas Bashford-Rogers[2], Kurt Debattista[3], Efstratios Doukakis[3] and Alan Chalmers[3]
1. Birmingham City University, UK
2. UWE Bristol, UK

3. University of Warwick, UK

First published: 14 September 2017
DOI: 10.1111/cgf.13295

Abstract: Accurate simulation of all the senses in virtual environments is a computationally expensive task. Visual saliency models have been used to improve computational performance for rendered content, but this is insufficient for multi-modal environments. This paper considers cross-modal perception and, in particular, if and how olfaction affects visual attention. Two experiments are presented in this paper. Firstly, eye tracking is gathered from a number of participants to gain an impression about where and how they view virtual objects when smell is introduced compared to an odourless condition. Based on the results of this experiment a new type of saliency map in a selective-rendering pipeline is presented. A second experiment validates this approach, and demonstrates that participants rank images as better quality, when compared to a reference, for the same rendering budget.

Frame Rate vs Resolution: A Subjective Evaluation of Spatiotemporal Perceived Quality Under Varying Computational Budgets

K. Debattista[1], K. Bugeja[2], S. Spina[2], T. Bashford-Rogers[3] and V. Hulusic[4]
1. University of Warwick, UK
2. University of Malta, Malta
3. University of the West of England, UK
4. Université Paris-Saclay, France

First published: 28 September 2017
DOI: 10.1111/cgf.13302

Abstract: Maximizing performance for rendered content requires making compromises on quality parameters depending on the computational resources available . Yet, it is currently unclear which parameters best maximize perceived quality. This work investigates perceived quality across computational budgets for the primary spatiotemporal parameters of resolution and frame rate. Three experiments are conducted. Experiment 1 (n = 26) shows that participants prefer fixed frame rates of 60 frames per second (fps) at lower resolutions over 30 fps at higher resolutions. Experiment 2 (n = 24) explores the relationship further with more budgets and quality settings and again finds 60 fps is generally preferred even when more resources are available. Experiment 3 (n = 25) permits the use of adaptive frame rates, and analyses the resource allocation across seven budgets. Results show that while participants allocate more resources to frame rate at lower budgets the situation reverses once higher budgets are available and a frame rate of around 40 fps is achieved. In the overall, the results demonstrate a complex relationship between frame rate and resolution's effects on perceived quality. This relationship can be harnessed, via the results and models presented, to obtain more cost-effective virtual experiences.

Stereo-Consistent Contours in Object Space

Dennis R. Bukenberger, Katharina Schwarz and Hendrik P. A. Lensch
Eberhard Karls University, Tübingen, Germany

First published: 22 November 2017
DOI: 10.1111/cgf.13291

Abstract: Notebook scribbles, art or technical illustrations—line drawings are a simplistic method to visually communicate information. Automated line drawings often originate from virtual 3D models, but one cannot trivially experience their three-dimensionality. This paper introduces a novel concept to produce stereo-consistent line drawings of virtual 3D objects. Some contour lines do not only depend on an object's geometry, but also on the position of the observer. To accomplish consistency between multiple view positions, our approach exploits geometrical characteristics of 3D surfaces in object space. Established techniques for stereo-consistent line drawings operate on rendered pixel images. In contrast, our pipeline operates in object space using vector geometry, which yields many advantages: The position of the final viewpoint(s) is flexible within a certain window even after the contour generation, e.g., a stereoscopic image pair is only one possible application. Such windows can be concatenated to simulate contours observed from an arbitrary camera path. Various types of popular contour generators can be handled equivalently, occlusions are natively supported and stylization based on geometry characteristics is also easily possible.

Localized Manifold Harmonics for Spectral Shape Analysis

S. Melzi[1], E. Rodolà[2,3], U. Castellani[1] and M. M. Bronstein[2,4,5,6]

1. University of Verona, Italy
2. USI Lugano, Switzerland
3. Sapienza University of Rome, Italy
4. Tel Aviv University, Israel
5. Intel Perceptual Computing, Israel
6. IAS, TU Munich, Germany

First published: 27 November 2017
DOI: 10.1111/cgf.13309

Abstract: The use of Laplacian eigenfunctions is ubiquitous in a wide range of computer graphics and geometry processing applications. In particular, Laplacian eigenbases allow generalizing the classical Fourier analysis to manifolds. A key drawback of such bases is their inherently global nature, as the Laplacian eigenfunctions carry geometric and topological structure of the entire manifold. In this paper, we introduce a new framework for local spectral shape analysis. We show how to efficiently construct localized orthogonal bases by solving an optimization problem that in turn can be posed as the eigendecomposition of a new operator obtained by a modification of the standard Laplacian. We study the theoretical and computational aspects of the proposed framework and showcase our new construction on the classical problems of shape approximation and correspondence. We obtain significant improvement compared to classical Laplacian eigenbases as well as other alternatives for constructing localized bases.

Laplace–Beltrami Operator on Point Clouds Based on Anisotropic Voronoi Diagram

Hongxing Qin[1,2], Yi Chen[1,2], Yunhai Wang[3], Xiaoyang Hong[1,2], Kangkang Yin[4] and Hui Huang[5]

1. CKLCI, Chongqing University of Posts & Telecommunications, China
2. CCST, Chongqing University of Posts & Telecommunications, China
3. Shandong University, China
4. Simon Fraser University, Canada
5. Shenzhen University, China

First published: 13 December 2017
DOI: 10.1111/cgf.13315

Abstract: The symmetrizable and converged Laplace–Beltrami operator (Δ_M) is an indispensable tool for spectral geometrical analysis of point clouds. The Δ_M, introduced by Liu et al. [LPG12] is guaranteed to be symmetrizable, but its convergence degrades when it is applied to models with sharp features. In this paper, we propose a novel Δ_M, which is not only symmetrizable but also can handle the point-sampled surface containing significant sharp features. By constructing the anisotropic Voronoi diagram in the local tangential space, the Δ_M can be well constructed for any given point. To compute the area of anisotropic Voronoi cell, we introduce an efficient approximation by projecting the cell to the local tangent plane and have proved its convergence. We present numerical experiments that clearly demonstrate the robustness and efficiency of the proposed Δ_M for point clouds that may contain noise, outliers, and non-uniformities in thickness and spacing. Moreover, we can show that its spectrum is more accurate than the ones from existing Δ_M for scan points or surfaces with sharp features.

An Implicit SPH Formulation for Incompressible Linearly Elastic Solids
Andreas Peer, Christoph Gissler, Stefan Band and Matthias Teschner
University of Freiburg, Germany

First published: 14 December 2017
DOI: 10.1111/cgf.13317

Abstract: We propose a novel smoothed particle hydrodynamics (SPH) formulation for deformable solids. Key aspects of our method are implicit elastic forces and an adapted SPH formulation for the deformation gradient that—in contrast to previous work—allows a rotation extraction directly from the SPH deformation gradient. The proposed implicit concept is entirely based on linear formulations. As a linear strain tensor is used, a rotation-aware computation of the deformation gradient is required. In contrast to existing work, the respective rotation estimation is entirely realized within the SPH concept using a novel formulation with incorporated kernel gradient correction for first-order consistency. The proposed implicit formulation and the adapted rotation estimation allow for significantly larger time steps and higher stiffness compared to explicit forms. Performance gain factors of up to one hundred are presented. Incompressibility of deformable solids is accounted for with an ISPH pressure solver. This further allows for a pressure-based boundary handling and a unified processing of deformables interacting with SPH fluids and rigids. Self-collisions are implicitly handled by the pressure solver.

Field-Aligned Isotropic Surface Remeshing
Xingyi Du[1], Xiaohan Liu[2,3], Dong-Ming Yan[2,3], Caigui Jiang[4], Juntao Ye[2,3], and Hui Zhang[1]
1. Software School, Tsinghua University, China
2. NLPR, Institute of Automation, Chinese Academy of Sciences, China
3. University of Chinese Academy of Sciences, China
4. Max Planck Institute for Informatics, Germany

First published: 28 February 2018
DOI: 10.1111/cgf.13329

Abstract: We present a novel isotropic surface remeshing algorithm that automatically aligns the mesh edges with an underlying directional field. The alignment is achieved by minimizing an energy function that combines both centroidal Voronoi tessellation (CVT) and the penalty enforced by a six-way rotational symmetry field. The CVT term ensures uniform distribution of the vertices and high remeshing quality, and the field constraint enforces the directional alignment of the edges.

Experimental results show that the proposed approach has the advantages of isotropic and field-aligned remeshing. Our algorithm is superior to the representative state-of-the-art approaches in various aspects.

Direct Position-Based Solver for Stiff Rods

Crispin Deul[1], Tassilo Kugelstadt2, Marcel Weiler[1] and Jan Bender[2]
1. Graduate School CE, TU Darmstadt, Germany
2. RWTH Aachen University, Germany

First published: 6 March 2018
DOI: 10.1111/cgf.13326

Abstract: In this paper, we present a novel direct solver for the efficient simulation of stiff, inextensible elastic rods within the Position-Based Dynamics (PBD) framework. It is based on the XPBD algorithm, which extends PBD to simulate elastic objects with physically meaningful material parameters. XPBD approximates an implicit Euler integration and solves the system of non-linear equations using a non-linear Gauss-Seidel solver. However, this solver requires many iterations to converge for complex models and if convergence is not reached, the material becomes too soft. In contrast we use Newton iterations in combination with our direct solver to solve the non-linear equations which significantly improves convergence by solving all constraints of an acyclic structure (tree), simultaneously. Our solver only requires a few Newton iterations to achieve high stiffness and inextensibility. We model inextensible rods and trees using rigid segments connected by constraints. Bending and twisting constraints are derived from the well-established Cosserat model. The high performance of our solver is demonstrated in highly realistic simulations of rods consisting of multiple ten-thousand segments. In summary, our method allows the efficient simulation of stiff rods in the Position-Based Dynamics framework with a speedup of two orders of magnitude compared to the original XPBD approach.