Computer Graphics Forum Technical Papers in SGP 2019

Functional Maps Representation On Product Manifolds

Emanuele Rodolà, Zorah Lähner, Alex M. Bronstein, Michael M. Bronstein, Justin Solomon

Computer Graphics Forum, Volume 38, Issue 1, Pages 678-689, February 2019
First published: 21 January 2019
DOI: 10.1111/cgf.13598

Abstract: We consider the tasks of representing, analysing and manipulating maps between shapes. We model maps as densities over the product manifold of the input shapes; these densities can be treated as scalar functions and therefore are manipulable using the language of signal processing on manifolds. Being a manifold itself, the product space endows the set of maps with a geometry of its own, which we exploit to define map operations in the spectral domain; we also derive relationships with other existing representations (soft maps and functional maps). To apply these ideas in practice, we discretize product manifolds and their Laplace–Beltrami operators, and we introduce localized spectral analysis of the product manifold as a novel tool for map processing. Our framework applies to maps defined between and across 2D and 3D shapes without requiring special adjustment, and it can be implemented efficiently with simple operations on sparse matrices.

Solid Geometry Processing on Deconstructed Domains

Silvia Sellán, Herng Yi Cheng, Yuming Ma, Mitchell Dembowski, Alec Jacobson

Computer Graphics Forum, Volume 38, Issue 1, Pages 564-579, February 2019
First published: 26 December 2018
DOI: 10.1111/cgf.13592

Abstract. Many tasks in geometry processing are modelled as variational problems solved numerically using the finite element method. For solid shapes, this requires a volumetric discretization, such as a boundary conforming tetrahedral mesh. Unfortunately, tetrahedral meshing remains an open challenge and existing methods either struggle to conform to complex boundary surfaces or require manual intervention to prevent failure. Rather than create a single volumetric mesh for the entire shape, we advocate for solid geometry processing on deconstructed domains, where a large and complex shape is composed of overlapping solid subdomains. As each smaller and simpler part is now easier to tetrahedralize, the question becomes how to account for overlaps during problem modelling and how to couple solutions on each subdomain together algebraically. We explore how and why previous coupling methods fail, and propose a method that couples solid domains only along their boundary surfaces. We demonstrate the superiority of this method through empirical convergence tests and qualitative applications to solid geometry processing on a variety of popular second-order and fourth-order partial differential equations.

LSMAT Least Squares Medial Axis Transform

Daniel Rebain, Baptiste Angles, Julien Valentin, Nicholas Vining, Jiju Peethambaran, Shahram Izadi, Andrea Tagliasacchi

Computer Graphics Forum
First published: 30 January 2019
DOI: 10.1111/cgf.13599

Abstract. The medial axis transform has applications in numerous fields including visualization, computer graphics, and computer vision. Unfortunately, traditional medial axis transformations are usually brittle in the presence of outliers, perturbations and/or noise along the boundary of objects. To overcome this limitation, we introduce a new formulation of the medial axis transform which is naturally robust in the presence of these artefacts. Unlike previous work which has approached the medial axis from a computational geometry angle, we consider it from a numerical optimization perspective. In this work, we follow the definition of the medial axis transform as 'the set of maximally inscribed spheres'. We show how this definition can be formulated as a least squares relaxation where the transform is obtained by minimizing a continuous optimization problem. The proposed approach is inherently parallelizable by performing independent optimization of each sphere using Gauss–Newton, and its least-squares form allows it to be significantly more robust compared to traditional computational geometry approaches. Extensive experiments on 2D and 3D objects demonstrate that our method provides superior results to the state of the art on both synthetic and real-data.