Computer-Aided Design 102 (2018) 61–71

Contents lists available at ScienceDirect

Computer-Aided Design

journal homepage: www.elsevier.com/locate/cad

Average and variance of a quasi-parallel family of surfaces

Mukul Sati *, Jarek Rossignac

School of Interactive Computing, Georgia Tech, USA

ARTICLE INFO

Keywords:
Surface averaging
Projection iteration
Visualization
Solid tolerancing

ABSTRACT

We provide theoretical foundations and practical computational tools for the statistical analysis of the local disparity between a family of situated surfaces. We do not mean statistics on discrete measures, such as pairwise Hausdorff distance, of these surfaces, but instead local, shape-variability statistics for all points on these surfaces in a manner that generalizes the mean and variance of numbers. Given a family $\mathbb{F}$ of n input surfaces B_i, we wish to compute a surface, B, that is the average of the surfaces in $\mathbb{F}$ and to associate, with each point p of B, a variance value, $v(p)$, which is the average of its squared distances to the input surfaces and hence measures the local disparity between the surfaces of the family. We choose B_k as any one of the input surfaces in $\mathbb{F}$. We define B as union of all points p, each resulting from 'snapping' a different point p_k of B_k. Snapping followed by closest projections may be used to establish pointwise correspondence between all input surfaces. When this correspondence defines a homeomorphism between B and each B_i, and hence between each pair of input surfaces, we say that $\mathbb{F}$ is a quasi-parallel family of surfaces and that B is their average. In such valid configurations, B exhibits properties that one would expect from an average of surfaces and the resulting variance field over B may be used for analysis, optimization, and visualization. A sufficient condition for this validity is to require that each pair of surfaces in $\mathbb{F}$ be projection-homeomorphic.

In practice, we only snap the vertices of a triangulation T_k that approximates B_k. The snap produces a triangulation T that approximates B. We obtain a triangulation T_i of an input surface B_i by projecting the vertices of T onto it. We propose a practical, although partial validity test that compares T to each T_i.

1. Introduction

In algorithms that operate on large data sets, a computed average may serve as an abstraction (summary of a data set) to simplify visualization and to accelerate queries. We focus on data sets that comprise CAD models of a family of solids. Our goal is to provide a useful definition of the solid that is the average of the input solids. Furthermore, we wish to annotate its boundary by a field that indicates local variability between the boundaries of the input solids.

We expect that the proposed definition of the average of solids may play a critical role in the analysis of the yield of manufacturing processes, especially for CNC machined parts.

The executions of a manufacturing process plan produce parts that are very similar, although not identical. Traditionally, part inspection and manufacturing process control have been carried out by comparing a selection of measured "as-manufactured" solid models to a nominal "as-designed" solid model in the context of associated tolerances.

However, owing to the significant benefits brought about, it is an increasing trend in manufacturing to instead use an explicit "digital twin" [1] representation of the entire manufacturing pipeline.

We propose a definition and a practical computation (1) of the solid average of the "as-manufactured" models and (2) of a local variability field over its surface. We further suggest that the average might be used as a more accurate statistical representation of the output of the manufacturing pipeline, because it is independent of the nominal model.

Our proposed approach will make it possible to assess separately (1) the variability (uncertainty, random errors) of a manufacturing process and (2) its compliance with a toleranced "as-designed" model (systematic error). Furthermore, the emergence of open architectures (such as [2]) and of sensor-equipped CNC machines make it conceivable to update the average "as-manufactured" shape in near-realtime. Doing so, might, for example, help to improve the manufacturing process by identifying regions where the feed-rate of the cutting tool must be reduced (Fig. 1).

To enable this vision, we first need a definition and effective algorithms for the computation of the average of a family of shapes. In this paper, we propose a simple construction called 'Snap' that

* Corresponding author.
E-mail address: mukul@gatech.edu (M. Sati).

https://doi.org/10.1016/j.cad.2018.04.014

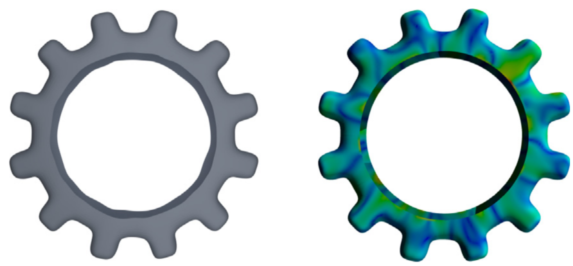

Fig. 1. One of the "as-manufactured" input solids (left) and the average of the family (right), color coded with the variance field. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

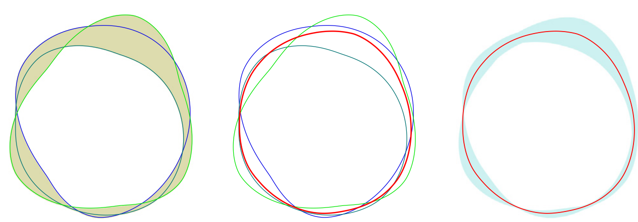

Fig. 2. A 2D illustration: A family of three input curves (proxies for surfaces) and the shaded gap between them (left). Their red average (center). Te variability shown by mapping standard deviation to variable thickness along the average (right). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

operates on a family $\mathbb{F}$ of n surfaces, $\{B_i\}$, each being the smooth, manifold boundary of a connected solid. Snap(B_k) considers one such surface, B_k as the 'seed' surface (e.g., in Fig. 3, $k = 3$). It maps each point, p_k, of B_k to a point $p = \text{Snap}(p_k)$.

As a mapping applied to all points of B_k, the Snap construction takes B_k to the point set B, which we call the 'snap set'. For 'quasi-parallel' configurations of inputs (for example, in which each pair of surfaces is projection-homeomorphic), B is a smooth, manifold surface. We call it the 'average surface', because it possesses properties that one would expect of the average of the input family. For example, B contains all points that lie in each B_i. B is manifold and parallel to the B_i surfaces in portions where the B_i's are parallel to each other (such as where they are each the offset by a constant distance of one of them). In these regions, each point p of B is the centroid of its closest projections $\{p_i\}$ onto the surfaces $\{B_i\}$. Our construction is also similarity-invariant and symmetric (independent of the order of the input surfaces). When $\mathbb{F}$ has only two surfaces, our definition of B matches the definition of the medial axis of the gap (Fig. 2) between them.

When does the Snap construction work? Sufficient and practically computable conditions for projection-homeomorphism between two surfaces are proposed in [3] and are based on comparing their Hausdorff distance to their minimum feature sizes. Pairwise projection-homeomorphism is sufficient, but not necessary for the existence of a valid average for the family and for the construction, by the algorithm proposed here, of a triangle mesh that approximates it. Thus, we use the term 'quasi-parallel family of surfaces' to describe the more general set of 'valid configurations' for which our algorithm produces a valid average, regardless of whether the family is pairwise projection-homeomorphic. However, we do not mean to imply that a valid average cannot exist for configurations for which it is not produced by our algorithm. This disparity is linked with the sampling nature of the proposed implementation and with the practical simplifications used for assessing validity from an approximating triangulation.

When the input surface family is quasi-parallel, B is a variable distance offset of B_k, and it establishes a homeomorphism between each pair of surfaces of the input family. It can be represented implicitly, storing a displacement field indicating the signed normal offset distance d by which point p_k of B_k must be displaced along its normal n_k, so as to reach the corresponding point p of B. Thus, the average surface is a valuable tool for signal and geometry processing tasks on $\mathbb{F}$, such as mapping or comparing surface attributes (e.g., textures).

Note that we do not assume any given correspondence between the situated input surfaces and that, in fact, the proposed solution may be used to establish such correspondences.

When computing the central tendency of numbers, one may choose to use the mean or the median. Similarly, we propose two options for the Snap construction — vAS and zAS. Their results are usually similar. They are identical when $\mathbb{F}$ contains only two surfaces.

> **vAS** produces a surface that is a subset of the 2-dimensional valley of the scalar function that, at each point, sums the squared distances to the B_i's.
> **zAS** produces a surface that is a subset of the zero-set of the scalar function that, at each point, sums the signed distances to the B_i's.

Snap works as follows. It initializes p to p_k. Then, it adjusts p by a short series of 'Moves' until p converges to a stable position, which, by definition, lies on Snap set B. Fig. 3 is a schematic overview of Snap, in a 2D setting. We analyze Snap theoretically, and provide experimental evidence that Snap converges in about 3 or 4 Moves to the Snap set.

In practice, we only snap a chosen subset of points p_k of B_k (see Fig. 3, where $k = 3$). When these are the vertices of a given triangulation T_k^* of B_k, snapping them produces a triangulation T of B.

Then, for each input surface B_i, we compute the closest projections of the vertices of T onto B_i to obtain a triangulation T_i of B_i. Hence, for each input surface B_i (including the surface B_k, whose sample triangulation T_k^* we use for seeding Snap), we create an approximating triangulation T_i. All these T_i have the same number of vertices as T, and, also, implicitly inherit the connectivity of T. T itself, has the same number of vertices as T_k^*, and inherits its connectivity.

We propose and analyze two variants of Snap, which differ in their symmetry and practical usefulness. Our first variant is symmetric in its constructions, but, T_k may be different from T_k^*. In the second variant, $T_k = T_k^*$.

We use these $n + 1$ triangulations (T and all T_i) as input for an a posteriori validity test, in which we check the compatibility between sets of the corresponding triangles.

We do not discuss verifying whether each triangulation T_i approximates the corresponding input surface B_i with sufficient geometric and topological accuracy, because (1) the suitable solution depends on how B_i is represented and on what closeness measure and validity definition is used, (2) this is an important problem that extends far beyond our scope, and (3) it is covered in various prior publications, for example those focused on surface simplification [4–6], on resampling and remeshing [7–9], or on reconstruction [10] and those focused on measuring the Hausdorff distance between two surfaces [11,12]. We advocate testing compatibility between T and each T_i and between each T_i/B_i pair by testing projection-homeomorphisms [3].

For studying statistical variability, at a point p of the computed average surface, B, one may compute the squared distances from p to each input shape, B_i, and use the average of these values as a measure of the local variance, $v(p)$, of the family $\mathbb{F}$. In the context of manufacturing, for inspection and online process improvement, the local variance facilitates novel visualizations to help engineers