ARTICLE IN PRESS

Computer-Aided Design ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Computer-Aided Design

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

Parameterizing and extending trimmed regions for tensor-product surface fitting☆

Márton Vaitkus *, Tamás Várady

Budapest University of Technology and Economics, Hungary

ARTICLE INFO

Keywords:
Surface fitting
Trimmed surfaces
Constrained parameterization
Labeling
Mesh extensions

ABSTRACT

Our goal is to approximate the data points of an irregular, trimmed triangular mesh by standard, tensor-product surfaces, such as NURBS. This is a difficult and ambiguous task driven by several parameters, including tolerances, knot vectors, smoothing terms, etc. One of the most crucial issues is the parameterization of the data points that will have a strong influence on the quality of the surface to be fitted.

We propose new techniques that attempt to establish some correspondence between the given trimmed region and the unknown four-sided surface. In the first phase a four-sided virtual guiding frame is created in 3D using labeled boundary segments. Then a 2D parameterization is computed that minimizes distortion while satisfying constraints implicitly defined by the labels and the frame. In the second phase new data points are inserted to fill in the domain rectangle, and an inverse mapping to 3D is performed. This leads to a four-sided mesh that extends the original trimmed region in a smooth and natural manner, and is well-suited for fitting an untrimmed tensor-product surface.

We will discuss several examples to illustrate the benefits of our techniques using a "black-box" surface fitter. Benefits include even curvature distribution, natural surface extensions in the vicinity of the trimming boundaries, avoiding wiggles for the full surface and a tight bounding box.

1. Introduction

The approximation of data points or triangle meshes by tensor-product parametric surfaces is an important task in geometric modeling and reverse engineering. While there are several applications where data points are arranged into a regular four-sided grid, our current interest is to approximate *trimmed regions* of irregular triangle meshes that are bounded by a multi-sided loop of boundary segments. While a regular structure implicitly determines four boundaries and implies a parameterization for the data points, in the trimmed case we have no explicit information about the orientation of the four boundaries to be created, and the surface geometry beyond the trimmed region. There are infinitely many tensor-product surfaces that approximate a set of given data points — we aim to produce surfaces that are well-suited for CAGD applications.

Fitting parametric surfaces is a thoroughly investigated, complex problem that depends on various parameters, including tolerances, knot vectors, smoothing terms and many others. In order to formulate surface fitting as a linear least-squares problem, appropriate (u, v) parameter values must be assigned to the data points.

It is impossible to pick a single best surface for a given trimmed region, and it is hard to prioritize the somewhat subjective, contradictory and application-dependent requirements. Nevertheless, there appears to be a general consensus among researchers (see e.g. [1–3]) and practitioners (see e.g. [4–7]) on the most important goals: (i) tolerance-driven fitting, (ii) relatively low number of control points, (iii) natural surface extensions in the vicinity of the boundary loops, (iv) controlled curvature distribution and no wiggling beyond the trimmed region, (v) moderate growth of the surface area and the bounding box with respect to the trimmed region and (vi) good alignment of the control grid with internal geometric features. Requirements (iii) and (iv) have particular significance when fitted surfaces are used for building and possibly redesigning B-rep models; in these cases intersection curves and fillet boundaries are often placed beyond the original trimming loops. (iv) and (v) help to prevent producing false intersection curves, and tighter boxes improve computational efficiency. A natural layout of the control net by (vi) permits using fewer control points for a prescribed tolerance, and also facilitates editing in the target CAD system. Several examples will demonstrate the importance of these requirements.

In this paper we are not going to deal with various surface fitting approaches, rather propose two useful "preprocessing" techniques, that strongly enhance the quality of the surface to be fitted.

The first is about setting the *orientation* of the surface. A 2D parameterization of the data points is computed, based on *labels*

☆ This special issue was edited by Scott Canann, Steven J. Owen & Hang Si.

* Corresponding author.
E-mail address: vaitkus@iit.bme.hu (M. Vaitkus).

https://doi.org/10.1016/j.cad.2017.11.008
0010-4485/© 2017 Elsevier Ltd. All rights reserved.

assigned to some of the boundary segments, and a corresponding *virtual guiding frame*. Accordingly the control points will be "well-aligned", and editing and optimizing the surface in the forthcoming phases can be accomplished in a natural manner. At the same time, weakly defined surface areas (see below) and the overall dimension of the surface remain relatively small. Once the preferred orientation is specified, a 2D parameterization follows, that respects these constraints and minimizes geometric distortion.

The second technique is about surface *extension*. After adding *supplementary data points* to fill the domain rectangle, an extended triangulation is created, and the data points are mapped back to optimized positions in 3D. In this way a four-sided mesh is obtained, that smoothly extrapolates the original trimmed region and is ready to be approximated by a tensor-product surface. The extension of the trimmed regions is critical for the stability of surface fitting. Without extension, the support of certain basis functions might contain only very few data points or none at all. Consequently the location of the corresponding *weak control points* will be defined in a very indeterminate manner, and the solution of the related badly conditioned linear least-squares system might yield a surface with unpredictable oscillations beyond the trimmed regions.

Although the above two techniques may be used independently, best results are achieved when these are applied in a successive manner. As an example, take the six-sided region in Fig. 1. Compare the surfaces with wiggling control points using an "ad-hoc" parameterization (Fig. 1a) and a simple ARAP-based parameterization (Fig. 1b) with the high quality surface produced by our proposed techniques (Fig. 1c). We will return to this example in Section 6.

Our approach has been found particularly useful in two application areas:

(i) In reverse engineering (RE), a CAD model is to be produced from a measured point cloud [8]. A crucial step of the process is to segment the triangular mesh, representing the object, into disjoint regions, that will later correspond to the individual faces of the final CAD model, see Fig. 2a. Each region will be approximated by a dedicated surface type, such as simple, regular surfaces (planes, quadrics, etc.), procedural surfaces (sweeps, lofts, fillets, etc.) and truly free-form shapes. In the last case the regions are likely to be fitted by tensor-product B-splines, as this surface representation is supported by most commercial CAD systems and data exchange standards such as IGES and STEP.

(ii) There are several tasks in surface design and related applications where non-standard representations are superior to tensor-product surfaces, nevertheless, the created geometry needs to be converted into NURBS format for downstream CAD/CAM applications. Typical examples include transfinite and control-point based multi-sided patches (e.g. [9,10], implicit surfaces (e.g. [11] or cell-based representations (e.g. [12]). In these cases, the regions we fit have several boundaries with an irregular topology without possessing an obvious tensor-product structure. This also motivated our research to find suitable trimmed surface representations with well-oriented boundaries and natural extensions. An example of a curve network based model is shown in Fig. 2b.

We note that there exists an alternative approach where trimmed regions are represented by smoothly connected tensor-product splines. To find optimal quadrilateral layouts, that match the trimming boundaries and ensure high-degree smoothness at extraordinary points [13] is a hard problem. The trimmed surface approach avoids these difficulties.

Our paper is structured as follows. After a brief overview on previous research in Section 2, we discuss the concept of labeling to orient the surface to be fitted and the workflow of our algorithm (Section 3). We describe the consecutive phases of constrained parameterization in Section 4, then using this intermediate data structure we discuss how to extend trimmed meshes in the 3D space (Section 5). We demonstrate the efficiency of our method by means of several comparative examples (Section 6). Suggestions for future work conclude the paper.

2. Previous work

The literature on tensor-product fitting algorithms is vast, we refer to the works [2,14] and the references therein. Below we will focus only on particular areas that have been found strongly related to our current project, including initial parameterization for surface fitting, mesh parameterization with constraints, handling weak control points and extrapolating meshes.

Initial data parameterization for tensor-product fitting. Initial parameterization methods usually fall into two major categories: projection-based and flattening-based methods.

In the former case data points are projected onto a *carrier surface* having some intrinsic parameterization. The carrier surface can be a plane, a cylinder, a Bézier or a Coons patch, depending on the complexity of the geometry; however the majority of published methods [15–17] tend to assume that all four boundary curves of the untrimmed patch are available. Pottmann and Leopoldseder [18] and Liu et al. [19] adapted the method of active contours/snakes by formulating the fitting problem as the direct minimization of a (squared) distance function between the spline and the point cloud, for which they describe an efficient approximation. While these methods apparently avoid the need for data parameterization, the problem of finding a good initial surface still remains.

In contrast, flattening-based methods map a triangle mesh formed by the data points to the (u, v) plane by minimizing some measure of geometric distortion. Mesh parameterization methods in our context has been described by Weiss et al. [2], but this work has preceded the development of modern free-boundary flattening methods, and orientation constraints were not addressed — see our discussion in the next section. More recently Lai et al. [20] and Wang and Zheng [21] parameterized triangle meshes in a "feature-sensitive" way to facilitate surface fitting, but the trimmed case is not touched upon.

We also mention a third class of methods that compute a parameterization by optimizing the geometry of its isolines [22–24]. The problem of point-cloud parameterization is also studied in the context of non-linear dimensionality reduction or manifold learning — see e.g. [25] for an application of manifold learning to trimmed fitting.

Constrained mesh parameterization. Mesh parameterization has an extensive literature — we refer to the survey [26]. Early works inspired by Tutte's method [27] flatten a simply connected mesh inside a fixed convex boundary, while more modern "free-boundary" methods minimize some geometric distortion measure instead. A large number of works aim to minimize angle distortion — such *conformal maps* are simple to compute and are widely used for graphics and remeshing, but prescribing constraints on them remains difficult, see [28] for a survey of the state-of-the-art. A distortion metric that is arguably more natural for CAD surfaces is *As-Rigid-As-Possible (ARAP)* [29], which measures the isometry of the mapping, and can be optimized efficiently in an alternating manner [30]. While conformal maps effectively treat surfaces as soap films, ARAP corresponds to an elastic deformation [31].

Methods for handling mesh parameterization constraints required by certain applications have also been researched. For texture mapping, positional constraints have to be satisfied, while ensuring low geometric distortion, see [26], Ch. 6.2. In parameterization-based quad meshing or layouting the isolines are aligned with principal curvatures and sharp features [13,32].

Please cite this article in press as: Vaitkus M., Várady T. Parameterizing and extending trimmed regions for tensor-product surface fitting. Computer-Aided Design (2017), https://doi.org/10.1016/j.cad.2017.11.008.