Technical University of Denmark

DTU

Cuttable Ruled Surface Strips for Milling

Steenstrup, Kasper Hornbak; Nørbjerg, Toke Bjerge; Søndergaard, Asbjørn ; Bærentzen, Jakob Andreas; Gravesen, Jens

DTU Library
Technical Information Center of Denmark

Cuttable Ruled Surface Strips for Milling

Kasper H. Steenstrup, Toke B. Nørbjerg, Asbjørn Søndergaard, Andreas Bærentzen, and Jens Gravesen

K. H. Steenstrup, T. B. Nørbjerg, A. Bærentzen, J. Gravesen
Technical University of Denmark, Denmark

khor@dtu.dk ✉
tono@dtu.dk
janba@dtu.dk
jgra@dtu.dk

A. Søndergaard
Odico Formwork Robotics, Denmark

asbjorn.sondergaard@aarch.dk

Abstract

This paper proposes a novel pre-processing method for industrial robotic CNC-milling. The method targets a hybrid machining process, in which the main bulk of material is removed through robotic hot or abrasive wire cutting, after which regular CNC-machining is employed for removal of the remaining material volume. Hereby, the roughing process is significantly sped up, reducing overall machining time. We compare our method to the convex hull and remove between 5% and 75% more material; on most models we obtain a 50% improvement. Our method ensures that no overcutting happens and that the result is cuttable by wire cutting.

1. Introduction

Recent years have seen a dramatic increase in the exploration of industrial robots for the purpose of architectural production (Kohler et al. 2014). While predominantly still a topic of research, some of these developments have recently matured into commercialisation, targeting the deployment of industrial robots for large-scale production (Søndergaard 2014). Within subtractive processes, the Denmark start-up Odico has been successfully bringing robotic hot-wire cutting to market.

CNC milling is a well-established process in industrial production of, particularly, foam-casting moulds, but also digitally produced stonework and bespoke timber manufacturing. While the process enables a very high degree of surface control and design freedom, it is also inherently limited by vast machining times for larger-scale applications that require the removal of large quantities of material, such as the machining of foam for ship hulls, wind turbine blades, or architectural structures. This adversity becomes significantly amplified when applied to hard materials, such as CNC milling of stone (Steinhagen et al. 2016). Wire cutting on the other hand, enables a dramatic reduction in production times, as volumetric artefacts can be produced in one or few swipes (McGee et al. 2013). However, here the precondition is that production geometries be described via ruled surfaces, which thus constrains the design freedom for the benefit of production efficiency. The wire-cutting methodology and its implications are extendable to abrasive wire sawing of, for instance, stoneworks, exemplified at the works of the Sagrada Familia cathedral (Burry 2016) as well as robotic abrasive wire sawing, as explored by Feringa and Søndergaard (2015).

The development of robotic hot-blade cutting (Søndergaard et al. 2016) provides a cost-effective and time-efficient manufacturing process for general curved foam geometries. However, this process is also constrained by the detail level achievable, and is inadequate for small surface details, while well suited for large-scale variations often deployed at industrial and architectural scale. In addition, so far, blade cutting is not applicable to non-foam materials.

The three processes milling, wire cutting, and blade cutting (see Figure 1) can be viewed as complementary, each covering a particular spectrum within subtractive machining. As such, an extension of the processes is to consider new ways for hybridisation.

One such possibility is the combination of milling and wire cutting, in which the latter is applied for removal of the initial volume which would otherwise be machined in CNC roughing processes. While CNC roughing is generally fast compared to CNC finishing, when assuming high target surface smoothness, the roughing process can represent substantial machining times when applied to voluminous subtractions. Additionally, certain architectural applications may enable the omission of surface finishing machining in favour of leaving roughed

Figure 1. Different cutting techniques: CNC-milling, hot wire cutting, and hot-blade cutting.

Figure 2. Creating an artificial concrete landscape in the urban harbour front of Copenhagen.

Figure 3. To the left a ruled surface defined by two curves. To the right a piecewise ruled surface defined by several curves.

surfaces for practical or aesthetic purposes. Two projects exemplifying this within the BladeRunner production portfolio, can be found in Feringa (2014). There, 212 m³ of expanded polystyrene foam were milled to achieve a three-dimensional guideline shape to be coated in-situ with 70–100 mm of polished spent, giving the final shape, see Figure 2. Only roughing processes were applied as for ensuring enhanced binding between the foam core and concrete shell, roughing representing approximately 92 direct machining hours.

Given the amount of machining hours spent on roughing, a hybrid approach would, for this case, have caused a reduction in processing time of between 69–72 %. In light of this finding, work was initiated to find a rationalisation algorithm that would cover any arbitrary three-dimensional shape with a set of non-convex ruled surfaces, such as to allow for a maximum of initial volume to be removed through wire cutting, while within the same robot cell shifting subsequently to aCNC tooling setup.

For this, we propose a method that combines fast wire cutting and precise CNC milling removing as much material as possible using a wire before the precise shape is milled. As we allow multiple cuts, the wire-cut surface is a piecewise ruled surface, and it can be considered as an approximation or *rationalisation* of the required surface. We can formulate the problem as follows: Given a surface, rationalise it with a piecewise ruled surface such that the rationalisation never intersects the original surface (no overcutting), and such that it can be manufactured by wire cutting. The latter implies that not only the rulings, but also the extension of the rulings never intersect the surface.

Usually a mould is composed by several blocks, and we do not consider the whole surface, but only a segment contained in a single block. As the final shape is milled, we do not need to consider any continuity conditions between the piecewise ruled rationalisations of the different segments.

A ruled surface is given by moving a line segment through space while it changes length and orientation. As a line segment is determined by its endpoints, the two curves described by the end points determine the surface uniquely, see Figure 3. A particular class of piecewise ruled surfaces is obtained by letting a polygon move through space while it changes shape, see Figure 3.

If the polygon at all times is on the outside of the original surface and furthermore is planar and convex, then we are certain that the extensions of the rulings never intersect the original surface, see Figure 4.

Piecewise ruled surfaces are well known in architecture: Flöry et al. (2012) describes a method to rationalise free form architecture, focussing on the smoothnessbetween rulings, and Flöry and Pottmann (2010) find areas where a good rationalisation can be done. Both papers use the asymptotic directions as guides for the rulings. In Wang and Elber (2014) large GPU-powered dynamic programming is used to minimise the distance between the original surface and the rationalisation.

Figure 4. Planar intersection of the surface and the piecewise ruled rationalisation. If the rulings turn less than 180° then the convexity of the rulings guarantees that the extended rulings never intersect the surface.

The paper by Elber and Fish (1997) constructs a piecewise ruled approximation of a free-form surface using Bézier surfaces and a subdivision scheme to get the approximation within tolerated error, but global accessibility is not guaranteed. In Elber (1995), a free-form surface is approximated by piecewise developable surfaces, by using a simple developable primitive and a subdivision technique.

Milling with a cylindrical tool produces piecewise ruled surfaces, so they have also been studied in this context. To improve the tool path Chiou (2004) shows that the error in the rough milling can be lowered by separating the ruled surface into multiple strips. The paper Chu and Chen (2006) constructs a piecewise ruled/developable rationalisation where a subdivision scheme is used if a tolerance is not met. Tool interference is taken care of, but only to the extent of a fixed axis flank milling tool. In Cao and Dong (2015) the one-sided Hausdorff distance is used to minimise the overcutting.

The paper by Julius et al. (2005) uses an iterative algorithm to automatically obtainrationalisation consisting of developable patches. In Jiang et al. (2014) user input is used to create a Lobel mesh which has the utility to create developable patches.

Our method distinguishes itself by accepting a general free-form surface as input and guaranteeing the wire does not cut into the model (overcutting), and that the rationalisation is cuttable by a wire.

2. Method

Given a surface $\mathbf{f}$, we want to minimise the distance between it and a piecewise ruled spline surface $\mathbf{s}$ of the form:

(1)

$$\mathbf{s}(u,v) = \sum_{i=1}^{k} \sum_{j=1}^{h} \beta_i^p(u) \, \beta_j^1(v) \, \mathbf{c}_{i,j} \,,$$

where $\mathbf{c} = \{\mathbf{c}_{i,j}\}$ is the set of control points and β_i^p is a B-spline of degree p. Observe that $\mathbf{s}$ is a piecewise ruled surface since the basis function β_i^1 has degree 1. We now discretise the piecewise ruled surface s by choosing a uniform grid $(\mathbf{u}_i, \mathbf{v}_j)$, $i = 1, \ldots, N$, $j = 1, \ldots, M$, in the parameter plane and we discretise the original surface $\mathbf{f}$ by sampling points $\mathbf{f}_{i,j}$, $i = 1, \ldots, N$, $j = 1, \ldots, M$, on the surface. How the sampling is done is explained in Section 2.3 below. We furthermore make sure that the v-knots are among the parameter values v_j, i.e. we have indices $1 = j_1 < j_2 < \cdots < j_h = M$ such that the knot vector in the v direction is $v_{j1}, \ldots, v_{jh}$. We measure the distance between $\mathbf{f}$ and $\mathbf{s}$ by the discrete square distance

(2)

$$\sum_{i=1}^{N} \sum_{j=1}^{M} \|\mathbf{f}_{i,j} - \mathbf{s}(u_i, v_j)\|^2,$$

2.1 Constraints

We need several constraints in the optimisation, which we now describe one by one.

One-sided approximation
To avoid overcutting, the rulings should all be on the outside of the model. So if $\mathbf{N}_{i,j}$ is the *outward* normal of $\mathbf{f}$ at the point $\mathbf{f}_{i,j}$ then we require that

(3)

$$\left(\mathbf{s}(u_i, v_i) - \mathbf{f}_{i,j}\right) \cdot \mathbf{N}_{i,j} \geq \varepsilon_1 \,, \quad \text{for all } i, j.$$

If $\varepsilon_1 = 0$, overcutting is only prevented at the sample points $\mathbf{f}_{i,j}$, but with $\varepsilon_1 > 0$ and sufficiently dense sampling overcutting is prevented. Alternatively, if $\mathbf{f}$ is a spline surface the local (or strong) convex hull property (Piegl and Tiller 2012, P3.22, P4.25) can be used to guarantee that no overcutting happens.

Planarity and convexity of rulings

The piecewise ruled surface $\mathbf{s}$ given by (1) can be considered as swept by a moving polygon, and we require that the polygon $\mathbf{s}(u_i, v_{j1}), \ldots, \mathbf{s}(u_i, v_{jh})$ is planar for all i. We now let $\mathbf{r}_{i,\ell} = \mathbf{s}(u_i, v_{j\ell+1}) - \mathbf{s}(u_i, v_{j\ell})$, i.e., it is one of the rulings. The difference in the other direction is denoted $\mathbf{w}_{i,\ell} = \mathbf{s}(u_{i+1}, v_{j\ell}) - \mathbf{s}(u_i, v_{j\ell})$. The cross product $\mathbf{n}_{i,\ell} = \mathbf{r}_{i,\ell} \times \mathbf{r}_{i,\ell+1}$ is a normal to the plane spanned by $\mathbf{r}_{i,\ell}$ and $\mathbf{r}_{i,\ell+1}$, see Figure 5. If all the normals $\mathbf{n}_{i,\ell}$, $\ell = 1, \ldots, h-1$ are parallel then the polygon is planar, and if they all point in the same direction then the polygon is convex or concave. We can formulate this condition as

(4)
$$\mathbf{n}_{i,\ell_1} \cdot \mathbf{n}_{i,\ell_2} = \left\| \mathbf{n}_{i,\ell_1} \right\| \left\| \mathbf{n}_{i,\ell_2} \right\|, \quad \ell_1, \ell_2 = 1, \ldots, h-1.$$

or to simplify it a bit

(5)
$$\mathbf{n}_{i,1} \cdot \mathbf{n}_{i,\ell} = \left\| \mathbf{n}_{i,1} \right\| \left\| \mathbf{n}_{i,\ell} \right\|, \quad \ell = 2, \ldots, h-1.$$

To rule out the possibility of a concave polygon, we require that the normal $\mathbf{n}_{i,\ell}$ points in roughly the same direction as $\mathbf{w}_{i,\ell}$. This can be formulated as

(6)
$$\mathbf{w}_{i,\ell} \cdot \mathbf{n}_{\ell,1} \geq \frac{1}{2} \left\| \mathbf{w}_{i,\ell} \right\| \left\| \mathbf{n}_{i,\ell} \right\|. \quad \ell = 1, \ldots, h-1.$$

or as (5) secures that $\mathbf{n}_{i,\ell}$ points in the same direction as $\mathbf{n}_{i,1}$ we can simplify it to

(7)
$$\mathbf{w}_{i,1} \cdot \mathbf{n}_{i,1} \geq \frac{1}{2} \left\| \mathbf{w}_{i,1} \right\| \left\| \mathbf{n}_{i,1} \right\|.$$

Figure 5. The discretised piecewise ruled surface from Figure 3. The points $\mathbf{s}(u_i, v_{j1}), \ldots, \mathbf{s}(u_i, v_{jh})$ form an instance of the moving polygon. The vector $\mathbf{r}_{i,\ell}$ is a leg in the polygon, i.e. a ruling. The vectors $\mathbf{w}_{i,1}, \ldots, \mathbf{w}_{i,h}$ goes from one polygon of rulings to the next.

Boundary

Ultimately the rationalised surface $\mathbf{s}$ is supposed to be cut from a block, which we assume has the form of a box with axis parallel sides, given by $a_1 \leq x \leq a_2$, $b_1 \leq y \leq b_2$, and $d_1 \leq z \leq d_2$. So no part of the boundary is allowed to be strictly inside the block. We furthermore assume that we will be cutting roughly in the x direction. So we require that

(8) (9)

$$c^x_{1,j} \leq a_1, \qquad c^x_{k,j} \geq a_2, \qquad j = 1, \ldots, h,$$
$$c^y_{i,1} \leq b_1, \qquad c^y_{i,h} \geq b_2, \qquad i = 1, \ldots, k.$$

where the superscript denotes the different components of the control points. In our implementation we start and end with the polygon on the block boundary, so in (8) the inequalities are replaced with equalities.

Limit the directions of the rulings

The rulings are not allowed to turn more than 180°, and this can be secured if the y-coordinate is a strictly increasing function. This is the case if it holds for the control polygon, and we formulate this as

(10)

$$c^y_{i+1,j} - c^y_{i,j} \geq \varepsilon_2,$$

where ε_2 is some small positive number. Strictly speaking we only need the difference to be non-negative, but using an $\varepsilon_2 > 0$ also prevents any ruling from collapsing into a single point.

2.2 The Optimisation Problem

We are now able to formulate the optimisation problem

(11)

$$\underset{\mathbf{c}}{\text{minimize}} \sum_{i=1}^{N} \sum_{j=1}^{M} \|\mathbf{f}_{i,j} - \mathbf{s}(u_i, v_j)\|^2,$$

such that

$$c_{1,j}^x = a_1, \quad c_{k,j}^x = a_2, \quad j = 1, \ldots, h,$$

$$c_{i,1}^y \leq b_1, \quad c_{i,h}^y \geq b_2, \quad i = 1, \ldots, k,$$

$$(\mathbf{s}(u_i, v_j) - \mathbf{f}_{i,j}) \cdot \mathbf{N}_{i,j} \geq \varepsilon_1, \quad i = 1, \ldots, k, \quad j = 1, \ldots, h,$$

$$\mathbf{n}_{i,1} \cdot \mathbf{n}_{i,\ell} = \|\mathbf{n}_{i,1}\| \|\mathbf{n}_{i,\ell}\|, \quad \ell = 2, \ldots, h-1,$$

$$\mathbf{w}_{i,1} \cdot \mathbf{n}_{i,1} \geq \frac{1}{2} \|\mathbf{w}_{i,1}\| \|\mathbf{n}_{i,1}\|,$$

$$c_{i+1,j}^y - c_{i,j}^y \geq \varepsilon_2,$$

where

(12) (13) (14)

$$\mathbf{n}_{i,\ell} = \mathbf{r}_{i,\ell} \times \mathbf{r}_{i,\ell+1}, \qquad \ell = 1, \ldots, h-2,$$

$$\mathbf{r}_{i,\ell} = \mathbf{s}(u_i, v_{j_{\ell+1}}) - \mathbf{s}(u_i, v_{j_\ell}), \qquad \ell = 1, \ldots, h-1,$$

$$\mathbf{w}_{i,\ell} = \mathbf{s}(u_{i+1}, v_{j_\ell}) - \mathbf{s}(u_i, v_{j_\ell}), \qquad i = 1, \ldots, k-1.$$

We solve this optimisation problem using the interior point method as in Wächter and Biegler (2005).

Figure 6. Illustration of the 3 steps that initialise the model into 900 points. Firstly: 10 of the 30 planes are shown, Secondly: 12 of the 30 intersection curves. Finally: 300 of the 900 discretisation points are shown.

Model 1 Model 2 Model 3

Model 4 Model 5

Figure 7. Five models rationalised by piece-wise ruled surfaces. Model one has been cut four times in one directions and four times in roughly the orthogonal direction. The four other models have been cut four times in one direction.

2.3 Initialisation

All that is left is to explain how we choose the sampling points $\mathbf{f}_{i,j}$ and initialise the optimisation.

First the coordinate system is chosen such that outward normal of the surface is roughly in the z-direction, i.e., upward. Then a cutting direction is chosen, and we create N parallel planes orthogonal to the cutting direction and uniformly spaced. For each plane the intersection curve with the original surface $\mathbf{f}$ is found. Finally, each intersection curve is discretised into M points. This is illustrated in Figure 6, where N and M both are 30.

3. Results

We have run our algorithm on the five models shown in Figure 7. Model 1 was cut in two different directions, four times in each directions, while Models 2–5 was cut four times in one direction. For comparison we have also calculated the convex hull, and the bounding box.

We have only considered cutting directions parallel to the sides of the blocks, but if we consider Model 2, diagonal cuts would be favourable. The depression in Model 4 poses a problem for wire cutting: No matter what direction a line at the depression has, it will cut into one of the 'mountains' at the edge. On the other hand, we see that even though we sweep the surface using a polygon with four legs the optimisation has put two of the legs on the same line. So we have in effect a polygon with only three legs and consequently only need three cuts to produce the rationalisation.

If we imagine the surface sitting inside the block and remove the outer part we are left with a solid object, see Figure 8. We have calculated the volume of that object in each case. The results are summarised in Table 1. We normalise all volumes with respect to the volume of the bounding box, i.e. we use Vol / Vol(BB). We show the volume of the model, but in the three other cases we show the volume that needs to be milled away, i.e., (Vol − Vol(Model))/ Vol(BB).

So for Model 1 we can see that the volume of the model is 49 % of the bounding box volume, that our method has left 4 % of the bounding box volume for milling, that the convex hull leaves 15 % of the bounding box volume for milling, and that the bounding box (doing nothing) leaves 51 % of the volume for milling. We also show how much more volume our method removes compared to the convex hull and the bounding box, respectively. That is we show (Vol − Vol(our))/ (Vol − Vol(Model)). For Model 1 this is 76 % and 93 %, respectively.

Figure 8. Model 1 with volume shown, the volume is created by intersection the bounding box with the surface.

No.	Model	Ruled	Convex Hull		Bounding Box	
1	0.4924	0.0374	0.1541	76%	0.5076	93%
2	0.2435	0.1520	0.3380	55%	0.7565	80%
3	0.5914	0.0201	0.0520	61%	0.4086	95%
4	0.3058	0.1379	0.1453	5%	0.6942	80%
5	0.4533	0.0798	0.2294	65%	0.5467	85%

Table 1. In the second column the volume of the model is shown. In columns 3,4, and 6 we show how much volume there is left for milling using our method, the convex hull, and the bounding box, respectively. All volumes are normalised with respect to the volume of the bounding box. In columns 5 and 7 we show how much more volume our method removes compared to the convex hull and the bounding box, respectively.

4. Conclusion and Future Work

We have described a novel method that finds a one-sided approximation of a free-form surface by a piecewise ruled surface. The method guarantees that no extension of the rulings cut through the original surface. This allows us to use the method and wire cutting as a pre-processing step for milling.

Compared to using the convex hull as a pre-processing step, we obtain an improvement from 5 to 75% and typically around 50%.

For simplicity we have limited ourselves to cuts parallel to the coordinate planes, but relaxing this conditions and allowing any cutting direction will improve the result.

We represent the piecewise ruled surface as a tensor product spline surface of degree one in one direction and the planarity condition of the rulings restrict the flexibility of the piecewise ruled surface. To overcome this, we can choose another representation where we explicitly move a planar polygon through space.

In this work we have assumed that the full architectural model has been segmented into block-sized portions, and we have only considered the piece inside a single block. An interesting possibility is to use our method to aid in the segmentation. If we allow for non-convex polygons in the optimisation, we will obtain a better fit and we could use the concave vertices to guide the segmentation.

Our system can be incorporated into a design workflow in one of several ways. The most obvious use is to simply run the method in order to reduce the amount of milling that is required to produce a given piece. However, the designer might also wish to create surfaces whose final design will change little when our method runs. In this case, we simply ensure that a design which is close to cuttable will, in fact, be so. Finally, our method seems to provides a distinct design expression with a plurality of intersection curves if cuts are made from several directions. We surmise that this could be used explicitly in some cases.

Acknowledgements

References

Burry, M. 2016. "Robots at the Sagrada Familia Basilica: A Brief History of Robotised Stone-Cutting." In *Robotic Fabrication in Architecture, Art and Design 2016.* Springer, 2–15.

Cao, L., and L. Dong. 2015. "Positioning Method of a Cylindrical Cutter for Ruled Surface Machining Based on Minimizing One-Sided Hausdorff Distance." *Chinese Journal of Aeronautics 28*, 5: 1564–1573.

Chiou, J. C. 2004. "Accurate Tool Position for Five-Axis Ruled Surface Machining by Swept Envelope Approach." *Computer-Aided Design 36*, 10: 967–974.

Chu, C.-H., and J.-T. Chen. 2006. "Tool Path Planning for Five-Axis Flank Milling with Developable Surface Approximation." *The International Journal of Advanced Manufacturing Technology 29*, 7-8: 707–713.

Elber, G., and R. Fish. 1997. "5-Axis Freeform Surface Milling Using Piecewise Ruled Surface Approximation." *Journal of Manufacturing Science and Engineering 119*, 3: 383–387.

Elber, G. 1995. "Model Fabrication Using Surface Layout Projection." *Computer-Aided Design 27*, 4: 283–291.

Feringa, J., and A. Søndergaard. 2015. "Fabricating architectural volume." *Fabricate*.

Feringa, J. 2014. "Entrepreneurship in Architectural Robotics: The Simultaneity of Craft, Economics and Design." *Architectural Design 84*, 3: 60–65.

Flöry, S., and H. Pottmann. 2010. "Ruled Surfaces for Rationalization and Design in Architecture. *LIFE in: formation. On Responsive Information and Variations in Architecture*, 103–109.

Flöry, S., Y. Nagai, F. Isvoranu, H. Pottmann, and J. Wallner. 2012. "Ruled Free Forms." In *Advances in Architectural Geometry*.

Jiang, C., C. Tang, M. Tomicic, J. Wallner, and H. Pottmann, H. 2014. "Interactive Modeling of Architectural Freeform Structures – Combining Geometry with Fabrication and Statics." In *Advances in Architectural Geometry*, ed. by P. Block, J. Knippers, and W. Wang. Springer.

Julius, D., V. Kraevoy, and A. Sheffer. 2005. "D-Charts: Quasi-Developable Mesh Segmentation." *Computer Graphics Forum.*

Kohler, F., M. Gramazio, and J. Willmann. 2014. "The robotic touch: How robots change architecture."

Mcgee, W., J. Feringa, and A. Søndergaard. 2013. *Rob – Arch 2012: Robotic Fabrication in Architecture, Art, and Design.* Vienna: Springer, ch. Processes for an Architecture of Volume, 62–71.

Piegl, L., and W. Tiller. 2012. *The NURBS book.* Springer Science & Business Media.

Søndergaard, A., J. Feringa, T. Nørbjerg, K. Steenstrup, D. Brander, J. Graversen, S. Markvorsen, A. Bærentzen, K. Petkov, J. Hattel, J., et al. 2016. "Robotic Hot-Blade Cutting." In *Robotic Fabrication in Architecture, Art and Design 2016*. Springer, 150–164.

Søndergaard, A. 2014. "Odico Formwork Robotics." *Architectural Design 84*, 3: 66–67.

Steinhagen, G., J. Braumann, J. Brüninghaus, M. Neuhaus, S. Brell-Cokcan, and B. Kuhlenkötter, B. 2016. „Path Planning for Robotic Artistic Stone Surface Production." In *Robotic Fabrication in Architecture, Art and Design 2016*. Springer, 122–135.

Wächter, A., and T. L. Biegler. 2005. "On the Implementation of an Interior-Point Filter Line-Search Algorithm for Large-Scale Nonlinear Programming." *Mathematical Programming 106*, 1: 25–57.

Wang, C. C., and G. Elber. 2014. "Multi-Dimensional Dynamic Programming in Ruled Surface Fitting." *Computer-Aided Design 51*: 39–49.