

Seamless Reconstruction of Part-Based High-Relief Models from Hand-Drawn Images

Marek Dvorožňák
Czech Technical University in Prague,
Faculty of Electrical Engineering
dvoromar@fel.cvut.cz

Saman Sepehri Nejad
University of Utah
samansepehri@gmail.com

Ondřej Jamriška
Czech Technical University in Prague,
Faculty of Electrical Engineering
jamriond@fel.cvut.cz

Alec Jacobson
University of Toronto
jacobson@cs.toronto.edu

Ladislav Kavan
University of Utah
ladislav.kavan@gmail.com

Daniel Sýkora
Czech Technical University in Prague,
Faculty of Electrical Engineering
sykorad@fel.cvut.cz




Figure 1: A comparison of our approach with the current state-of-the-art: the original input drawing (a); the result of Entem et al. [2015] (b) in contrast to the result of our technique (c) that produces a more natural transition between individual parts; the result of Sýkora et al. [2014] suffers from visible seams between individual parts (d) whereas our approach delivers smooth transition (e). (Images (a) and (b) come from [Entem et al. 2015].)

ABSTRACT

We present a new approach to reconstruction of *high-relief* surface models from hand-made drawings. Our method is tailored to an interactive modeling scenario where the input drawing can be separated into a set of semantically meaningful parts of which relative depth order is known beforehand. For this kind of input, our technique allows inflating individual components to have a semi-elliptical profile, positioning them to satisfy prescribed depth order, and providing their seamless interconnection. Compared to previous methods, our approach is the first that formulates this reconstruction process as a single non-linear optimization problem. Because its direct optimization is computationally challenging, we propose an approximate solution which delivers comparable results orders of magnitude faster enabling an interactive user workflow. We evaluate our approach on various hand-made drawings and demonstrate that it provides state-of-the-art quality in comparison with previous methods which require comparable user intervention.

CCS CONCEPTS

• **Computing methodologies** → **Mesh models; Reconstruction; Non-photorealistic rendering;**

KEYWORDS

sketch-based 3D modeling, high-relief models, hand-drawn images

ACM Reference Format:

Marek Dvorožňák, Saman Sepehri Nejad, Ondřej Jamriška, Alec Jacobson, Ladislav Kavan, and Daniel Sýkora. 2018. Seamless Reconstruction of Part-Based High-Relief Models from Hand-Drawn Images. In *Expressive '18: The Joint Symposium on Computational Aesthetics and Sketch Based Interfaces and Modeling and Non-Photorealistic Animation and Rendering, August 17–19, 2018, Victoria, BC, Canada*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3229147.3229153>

1 INTRODUCTION

Recent advances in interactive 3D modeling from a single image [Entem et al. 2015; Feng et al. 2016; Li et al. 2017; Sýkora et al. 2014; Xu et al. 2014; Yeh et al. 2017] make the creation of 3D models less demanding for artists and also more accessible for novice users who do not have sufficient experience with professional 3D modeling tools. Such tools need complex manipulation with geometric primitives in 3D space which requires working with multiple 2D projections. A key advantage of staying in the 2D domain is that it allows the user to remain entirely focused on the creative process and not be distracted by resolving consistency in depth. This important aspect also explains why 2D sketches are often used as concept art or in early stages of 3D model design.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Expressive '18, August 17–19, 2018, Victoria, BC, Canada

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5892-7/18/08...\$15.00

<https://doi.org/10.1145/3229147.3229153>

In this paper, we focus on a branch of single image modeling methods that are suitable for organic structures composed of several rounded parts that are positioned in depth and attached or smoothly connected to each other [Entem et al. 2015; Feng et al. 2016; Sýkora et al. 2014; Yeh et al. 2017]. In the original image, these parts are usually delineated by outlines or have distinct boundaries that can be easily extracted. Thanks to the 2.5D layered structure, this kind of input usually requires only little user intervention while the resulting meshes still have a certain level of complexity which would be more difficult to achieve using standard 3D modeling tools. The desired result is akin to a *high-relief* sculpture, defined in classic sculpture as a relief where more than half of the shape projects from the background at full depth (see, e.g., [Read 1961]).

With previous methods, however, the reconstruction process is usually separated into a set of individual sub-problems which are solved sequentially, e.g., the input regions are first inflated and then shifted to preserve the relative depth. Finally, already inflated and shifted components are stitched together or smoothly interconnected. Due to this sequential process, the quality of the resulting mesh often suffers from the lack of smoothness in the areas where individual parts were stitched together. In this paper, we formulate a single optimization problem that unifies all of the above-mentioned sub-problems within a single energy minimization framework, delivering seamless organic shapes that seem like sculpted from a single block of material.

The contributions of our work are as follows: (1) We formulate the reconstruction of high-relief models from a single hand-drawn image as a minimization of a unified non-linear energy functional. Thanks to this joint formulation, our technique naturally produces meshes where the individual parts are interconnected smoothly and seamlessly. (2) We propose an efficient approximate method to our non-linear solution which enables interactive modeling workflow.

2 RELATED WORK

Igarashi et al. [1999] introduced the concept of modeling by inflation. They add volume to a 2D shape procedurally by triangulating the shape and setting vertex heights proportional to chordal axis distance. This concept was later extended by others using convolutional surfaces [Tai et al. 2004], sweeping 2D template scalar field [Schmidt et al. 2005], using mass-spring system [Karpenko and Hughes 2006], non-linear optimization [Nealen et al. 2007], generalized cylinders [Borosán et al. 2012; Zeng et al. 2015], surfaces of revolution [Bessmeltsev et al. 2015], level set method [Levi and Gotsman 2013], and finally by an implicit surface which is defined by a skeleton of the inflated region’s shape and its radius function [Entem et al. 2015].

Other methods provide an extension that allows specification of cross-sectional functions for individual components [Olsen and Samavati 2010] or define a set of primitives that can be used to approximate their shape [Chen et al. 2013; Gingold et al. 2009; Shtof et al. 2013]. Sýkora et al. [2014] and Feng et al. [2016] obtain inflated shapes with semi-elliptical profiles by solving Poisson equation of the squared height (recovering the height by taking the square root). Yeh et al. [2017] and Jayaraman et al. [2018] utilize user-specified curvature annotations to infer gradient fields and obtain inflated shapes with parabolic profiles by solving Poisson equation




Figure 2: Comparison of inflation with a parabolic (a) and a semi-elliptical (b) profile for a frontal (top) and a side (bottom) view. Notice how the semi-elliptical inflation is steeper at boundaries and more evenly rounded which gives the frontal render (top right) a more natural appearance.

of the (unsquared) height (see Fig. 2). Li et al. [2017] use an iterative process guided by several types of user-provided curves.

As the inflation process is usually applied only to a single 2D region, the resulting 3D object has only limited structural complexity. To produce more complex 3D objects, individual components need to be inflated separately and then joined together. In this “piecewise” workflow, correct absolute depth values need to be specified in order to preserve the overall 3D structure and avoid potential penetration of the individual parts. A typical approach how to resolve this problem is to let the user to view the object from side-views and specify absolute depths manually [Borosán et al. 2012; Feng et al. 2016; Igarashi et al. 1999; Nealen et al. 2007]. Gingold et al. [2009] use manually constructed intersection curves to guide the positioning of parts in depth, and Bessmeltsev et al. [2015] delegate depth specification to an underlying 3D skeleton positioned by the user. Finally, Yeh et al. [2017] allows the user to specify a set of sparse slope cues which locally define surface increase in depth.

Other approaches try to infer the missing depth information using various cues and perform the positioning in an automatic fashion. Sýkora et al. [2010] use a set of sparse depth (in)equalities given by the user. Sýkora et al. [2014] utilize illusory surfaces to predict depth (in)equalities. Liu et al. [2013] use relative depth cues represented by T-junctions while Zeng et al. [2015] combines them also with ground contact cues. In [Yeh et al. 2015], angles at junctions and region overlaps are used as a layering metric while Shtof et al. [2013], as well as Chen et al. [2013], utilize geosemantic constraints as relative depth cues. Finally, Sýkora et al. [2014] use quadratic programming to automatically find smooth surfaces that translate and deform parts along axis perpendicular to the image so that the relative depth ordering is satisfied.

After determining the relative depth positioning of the individual components, the remaining task is to join the components together, ideally in a smooth and seamless way. Igarashi et al. [1999] lets the user specify a region for subsequent surface fairing by low-pass filtering [Taubin 1995]. In a follow-up work [Igarashi and Hughes




Figure 3: Original drawings (top) and the input to our method (bottom): pre-tesselated regions with completed occluded parts have their relative depth ordering visualized in grayscale (the lighter is the closer). Equality and inequality constraints are visualized using cyan and magenta, respectively, and interconnection lines using blue color. The Robin boundary constraints (green-to-red gradient) may be automatically determined by our system. (Original drawings: wolf, bunny, farmer © Anifilm; unicorn, elephant come from [Entem et al. 2015]; lamb © Marek Dvorožňák.)

2003], simple filtering is replaced by a non-linear method [Schneider and Kobbelt 2001]. A similar approach is used by Nealen et al. [2007], utilizing a fairing interpolation of surfaces defined by control curves. Borosán et al. [2012] use Laplacian smoothing around intersection loops, and Levi and Gotsman [2013] utilize a heuristic where a boolean union of spheres and reconstructed parts at joints is followed by bi-Laplacian smoothing. Blending of implicit primitives based on Ricci’s operator [Ricci 1973] is used in Schmidt et al. [2005] as well as in Entem et al. [2015]. This approach allows control over the smoothness of each joint. To produce smooth joints, Sýkora et al. [2014] perform additional smoothing step by performing biharmonic interpolation in regions corresponding to connections of the individual components.

Despite the progress made in the above-mentioned work, the modeling process remains decoupled into separate steps. In this paper, we *unify* inflation, positioning, and seamless joining of individual components. Due to this unified formulation, hand-drawn images can be converted into high-quality meshes with minimal user intervention.

In this section, we have described methods that deal with the reconstruction of 3D shapes from sketches. The results are either full 3D models or some of their approximations, like high-relief models. There are also different techniques [Arpa et al. 2015; Cignoni et al. 1997; Schüller et al. 2014; Weyrich et al. 2007] that aim at base or high-relief generation out of a full 3D model. However, those deal with depth compression of the model as opposed to the model reconstruction.

3 OUR APPROACH

The aim of our approach is to reconstruct a high-relief model from a hand-drawn image. As an input to our method, we expect a set of semantically meaningful regions with completed occluded parts

of which relative depth ordering is known. In addition to that, we assume (in)equality constraints for region boundaries and boundary vertices where two regions should merge (see Fig. 3). All this can be obtained using a semi-automatic process described in [Sýkora et al. 2014].

We would like the final high-relief model to satisfy the following requirements:

- regions should be inflated in a way so that the resulting shapes have semi-elliptical profiles,
- they should be shifted in depth to satisfy the prescribed relative depth ordering,
- interconnection of regions should be seamless, and
- the resulting model should closely match contours of the input 2D drawing when using orthographic projection.

Although the inflation with parabolic profile is solvable using linear system [Sýkora et al. 2014], non-linear semi-elliptical inflation is more desirable because it produces shapes that are steeper at boundaries and more evenly rounded. This is especially important for organic models such as cartoon characters (see Fig. 2 for comparison).

To restrict the high-relief model to stay within the boundaries given by the outlines of the original drawing, we restrict the inflation and shifting of individual parts to take place only in a direction that is perpendicular to the original image plane, i.e., we will optimize only z coordinates of the final 3D mesh as in [Sýkora et al. 2014].

In the rest of this section, we first show how to inflate a single region to have a desired semi-elliptical profile by formulating a non-linear inflation functional ($E_{inflation}$). Then, we consider joint inflation of multiple regions and satisfaction of their relative depth order. For this, we combine $E_{inflation}$ with a shifting functional (E_{shift}) which gives us the final non-linear energy E that expresses

the whole reconstruction problem. Finally, we show how to linearize E and get an approximate solution which can be solved using a quadratic program. This enables substantial speedup while retaining similar quality as the original non-linear solution.

3.1 Non-linear Formulation

Inflation of a single region. An initial planar region Ω_i can be inflated to a semi-elliptical profile by finding a function $f(\mathbf{x}) : \Omega_i \rightarrow \mathbb{R}$ that minimizes the following energy functional

$$E_{inflation}^s = \int_{int(\Omega_i)} (\Delta f^2(\mathbf{x}) - c)^2 dx \quad (1)$$

subject to

$$f(\mathbf{x}) = 0 \quad \forall \mathbf{x} \in \partial\Omega_i \quad (2)$$

where $int(\Omega_i)$ and $\partial\Omega_i$ is the interior and the boundary of Ω_i , respectively, Δ is the Laplacian operator and c is a scalar specifying a user-controllable amount of inflation. The energy is non-linear because the Laplacian is applied to the square of f and therefore, the result is not invariant to translation of the boundary conditions. To be able to move the parts in the z direction, which may be necessary to meet the relative depth order, we introduce a separate shifting function.

Simultaneous inflation and stitching of multiple regions. Given n initially planar regions Ω_i , we topologically interconnect those individual overlapping components at areas where seamless connection is desired (see Fig. 3). This enables us to achieve seamless transitions among them. Now, the task is to find an inflation and a shifting function $f(\mathbf{x}) : \Omega \rightarrow \mathbb{R}$ and $g(\mathbf{x}) : \Omega \rightarrow \mathbb{R}$, respectively, that minimize the following energy functional:

$$E = E_{inflation} + \lambda_{shift} E_{shift} \quad (3)$$

where

$$E_{inflation} = \int_{int(\Omega)} (\Delta f^2(\mathbf{x}) - c)^2 dx + \lambda_{bnd} \int_B (f(\mathbf{x}))^2 dx, \quad (4)$$

$$E_{shift} = \int_{\Omega} \|\nabla g(\mathbf{x})\|^2 dx, \quad (5)$$

$\Omega = \Omega_1 \cup \dots \cup \Omega_n$ is a unifying region that contains all the regions Ω_i and λ_{shift} is a regularization parameter controlling the balance between the inflation and the shifting of parts in the optimization.

As compared to the inflation of a single region where the boundary of the region is fixed at the plane $z = 0$ (Formula 1 and 2), the inflation energy $E_{inflation}$ is extended by a term (controlled by the parameter λ_{bnd}) that allows movement of the boundary of f on a subset $B \subseteq \partial\Omega$. This relaxation proves useful in alleviating unwanted reconstruction artifacts visually resembling *pits*. The influence of λ_{bnd} is visualized in Fig. 4 and also in our supplementary video. Fig. 9 (a) and (c) shows how the results with and without these artifacts look like when rendered.

The aim of the shifting energy E_{shift} , where ∇ stands for the gradient operator, is to find a function g that deforms the inflated shape f in a way that the result $h = f + g$ satisfies the relative depth conditions while encouraging g to be as flat as possible. The energy

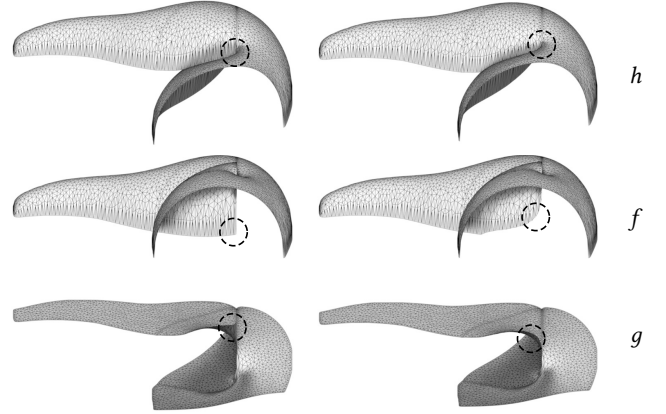


Figure 4: Converged results of the non-linear optimization with visualization of the inflation (f) and the shifting (g) function for the final result $h = f + g$, and also of the influence of the inflation boundary relaxation term λ_{bnd} . The result in the right column has its boundary more flexible ($\lambda_{bnd} = 0.01$) than the result in the left column ($\lambda_{bnd} = 100$) which mitigate formation of pits (depicted inside circles). See accompanying video for an animation.

E is minimized subject to the following relative depth conditions:

$$\begin{aligned} f_i(\mathbf{x}) + g_i(\mathbf{x}) &= f_j(\mathbf{x}) + g_j(\mathbf{x}) \quad \forall \mathbf{x} \in C_{i,j}^=, \\ f_i(\mathbf{x}) + g_i(\mathbf{x}) &\leq f_j(\mathbf{x}) + g_j(\mathbf{x}) \quad \forall \mathbf{x} \in C_{i,j}^{\leq}, \\ f_i(\mathbf{x}) + g_i(\mathbf{x}) &\geq f_j(\mathbf{x}) + g_j(\mathbf{x}) \quad \forall \mathbf{x} \in C_{i,j}^{\geq}, \end{aligned} \quad (6)$$

where $C_{i,j}^=, C_{i,j}^{\leq}, C_{i,j}^{\geq} \subseteq \Omega_i \cap \partial\Omega_j$ are sets of points that specify relative depth order for two overlapping regions Ω_i and Ω_j . The resulting function h is then simply $h = f + g$.

3.2 Efficient Reformulation of Non-linear Energy

Although the non-linear solution that we have described fulfills our reconstruction requirements, depending on the mesh resolution and chosen numerical method, the convergence of the non-linear optimization may be relatively slow. In our implementation, the optimization often lasts hours for input with moderate complexity (see more details in Section 4). Even though more sophisticated numerical methods could offer higher performance, in this section we propose a different approach based on the observation that for a suitable λ_{shift} , when the optimization converges, the inflation energy $E_{inflation}$ is minimized and f is therefore completely inflated. To obtain a solution that is orders of magnitude faster while producing results that are comparable in quality, we build on this observation and separate the problem into the two subsequent steps: inflation and shifting.

As shown by Sýkora et al. [2014], the inflation with a semi-elliptical profile that corresponds to the Formula 1 can be obtained efficiently by solving the Poisson equation

$$\Delta \tilde{f}(\mathbf{x}) = c \quad (7)$$

which produces a shape \tilde{f} with a parabolic profile, and then the following cross-section function is used to change the shape of f to a semi-elliptical profile:

$$f(\mathbf{x}) = d\sqrt{\tilde{f}(\mathbf{x})} \quad (8)$$

where d is a scaling factor allowing to obtain flatter profile or reverse inflation.

To allow for greater modeling flexibility, we extend this inflation method to support transition between fixed and free boundaries via Robin boundary conditions for Equation 7:

$$\alpha(\mathbf{x})f(\mathbf{x}) + (1 - \alpha(\mathbf{x}))\frac{\partial f}{\partial \mathbf{n}}(\mathbf{x}) = 0 \quad \forall \mathbf{x} \in \partial\Omega_i \quad (9)$$

where $\alpha(\mathbf{x}) \in [0, 1]$ specifies the interpolation between Dirichlet and Neumann boundary constraints. This extension allows us to mimic the behavior of non-linear solution where the boundary can be shifted (Formula 4).

To obtain the final surface h that satisfies the specified relative depth order, we minimize:

$$\int_{\Omega} \|\nabla h(\mathbf{x}) - \nabla f(\mathbf{x})\|^2 d\mathbf{x} \quad (10)$$

subject to (in)equality constraints representing the relative depth conditions:

$$\begin{aligned} h_i(\mathbf{x}) &= h_j(\mathbf{x}) & \forall \mathbf{x} \in C_{i,j}^=, \\ h_i(\mathbf{x}) &\leq h_j(\mathbf{x}) & \forall \mathbf{x} \in C_{i,j}^{\leq}, \\ h_i(\mathbf{x}) &\geq h_j(\mathbf{x}) & \forall \mathbf{x} \in C_{i,j}^{\geq}. \end{aligned} \quad (11)$$

This formulation is mathematically equivalent to the one used in [Sýkora et al. 2014]. However, it allows us to directly optimize for the final mesh h as opposed to the two-step procedure in [Sýkora et al. 2014].

3.3 Implementation Details

We discretize our continuous formulation using the finite element method. We assume that each planar region Ω_i is converted into a triangular mesh with additional interior vertices for boundary vertices of each region Ω_j that overlaps with Ω_i . This instantly gives us pairs of corresponding vertices that are used to satisfy relative depth ordering, i.e., the sets $C_{i,j}^=$, $C_{i,j}^{\leq}$ and $C_{i,j}^{\geq}$. Please refer to [Sýkora et al. 2014] for more details about the procedure. Then, we reconnect the meshed regions at vertices where seamless transition is expected. These are visualized using blue lines in Fig. 3.

Non-linear formulation details. The inflation and shifting non-linear energies (Formula 4 and 5) are discretized as follows

$$\begin{aligned} E_{inflation} &\approx (\mathbf{M}_{in}^{-1}\mathbf{L}_{in}\mathbf{f}_{in}^2 - \mathbf{c})^T \mathbf{M}_{in} (\mathbf{M}_{in}^{-1}\mathbf{L}_{in}\mathbf{f}_{in}^2 - \mathbf{c}) \\ &\quad + \lambda_{bnd}\mathbf{f}_{bnd}^T \mathbf{M}_{bnd}\mathbf{f}_{bnd}, \\ E_{shift} &\approx (\mathbf{G}\mathbf{g})^T \mathbf{T} (\mathbf{G}\mathbf{g}), \end{aligned}$$

where \mathbf{c} is a column vector of scalars c (see Formula 1), \mathbf{T} , \mathbf{M} and \mathbf{M}^{-1} are diagonal matrices representing areas of mesh triangles, the mass matrix and its inverse, \mathbf{G} and \mathbf{L} are sparse matrices representing discretization of the gradient operator and the usual cotangent discretization of the Laplacian operator [Meyer et al. 2003]. The square of a vector is understood as element-wise operation and




Figure 5: Example of boundary conditions used for mitigation of pits. Movable boundary utilized in our non-linear solution is visualized using yellow color (left) and Robin boundary conditions that we employ in our approximate solution using green-to-red gradient (right). User input is visualized using green and red points which specify the range of the specified conditions.

the subscripts *in* and *bnd* denote a part of a matrix or a vector corresponding to interior and boundary vertices, respectively.

Efficient reformulation details. The Poisson equation for obtaining a parabolic inflation (Formula 7) is discretized as

$$\mathbf{L}\tilde{\mathbf{f}} = \mathbf{M}\mathbf{c}$$

and the minimization of functional in Formula 10 with (in)equalities (Formula 11) is discretized as a quadratic program, i.e., we minimize:

$$\frac{1}{2}\mathbf{h}^T \mathbf{L}\mathbf{h} - \mathbf{h}^T \mathbf{L}\mathbf{f} \quad (12)$$

subject to (in)equality constraints:

$$\begin{aligned} \mathbf{h}_i(p) &= \mathbf{h}_j(p) & \forall p \in C_{i,j}^=, \\ \mathbf{h}_i(p) &\leq \mathbf{h}_j(p) & \forall p \in C_{i,j}^{\leq}, \\ \mathbf{h}_i(p) &\geq \mathbf{h}_j(p) & \forall p \in C_{i,j}^{\geq}. \end{aligned} \quad (13)$$

4 RESULTS

We implemented both the original non-linear method and the linearized approximation in C++. Our implementation relies on the Eigen library [Guennebaud et al. 2010] for linear algebra routines, libigl [Jacobson et al. 2013] for discrete differential operators, solvers of quadratic problems and programs, and L-BFGS non-linear solver [Liu and Nocedal 1989]. To compute the gradient of the energy used in the non-linear solver, we used reverse-mode automatic differentiation from the Stan Math Library [Carpenter et al. 2015].

For all results included in this paper, we use the following parameters: $c = 4$ to obtain an inflation with hemispherical profile for circular regions and a semi-elliptical profile for regions with different shapes, $\lambda_{shift} = 1$ to equally balance inflation with shifting, $\lambda_{bnd} = 100$ for results that contain pits, and $\lambda_{bnd} = 0.01$ for results that mitigate pits.

According to our experiments, including the entire boundary of the function f into the subset B in Formula 4 may result in unwanted shifting of parts of the boundary that should stay fixed. This is caused by the competition between the individual terms in energy E (Formula 3). We resolve this by restricting B as follows: For each two overlapping regions that are interconnected and one is supposed to be above the other one, we include only boundary




Figure 6: Comparison of results produced by Entem et al. [2015] (top) with our approach (bottom); for the unicorn and the elephant example. Our results more closely reproduce outlines of the original input drawings and contain seamlessly merged body parts without bulges. (Images located at the top come from [Entem et al. 2015].)

vertices of the top region that are overlapping with the bottom region into B (an example of the resulting subset B is depicted using yellow color in Fig. 5, left). For top regions that are entirely surrounded by the bottom components (e.g., elephant’s ear), a user intervention may be needed to obtain satisfactory result. For these cases, we provide a simple two-click interface to specify the range manually.

We use the same technique for specification of Robin boundary conditions for the top region (Formula 9) in our approximate solution. We set $\alpha(\mathbf{x}) \in [0, 0.2]$ for boundary points that are adjacent to the interconnection line (blue line in Fig. 5), $\alpha(\mathbf{x}) = 1$ for points that are at the borderline of the bottom region and linearly interpolate α between the points. For all other boundary points, we assume $\alpha(\mathbf{x}) = 1$. For obtaining a reverse inflation (e.g., rabbit’s ears or farmer’s shovel in Fig. 7), the parameter d was set to -1 , otherwise we used $d = 1$.

We ran our implementation on a quad-core CPU (Core i7, 2.7 GHz, 16 GB RAM). Summarized running times for our approximate solution are presented in Table 1.

Our implementation of non-linear solution often required hours to converge (see the supplementary material for a time-lapse video which illustrates convergence of the non-linear optimization). In contrast, the approximate version is significantly faster while the results look nearly identical (see Fig. 9 for a comparison).

We also compared the results produced by our method with the results produced by the two most closely related previous works: Ink-and-Ray [Sýkora et al. 2014] and a method by Entem et al. [2015].

The results of Ink-and-Ray with visible seams at connections of regions are shown in Fig. 7 and are even more pronounced when the lighting conditions changes, see Fig. 8 and the supplementary material. Those seams are caused by bi-Laplacian smoothing in a user-specified area around the connection that is performed in a post-processing grafting phase. Since we reconnect the overlapping regions before the reconstruction phase, our results naturally reproduce seamless connections that respect the specified inflation profile.

Table 1: Running times of our approximate solution.

model	# vertices	# faces	time
wolf	32 k	52 k	10.1 s
bunny	19 k	33 k	2.7 s
farmer	33 k	55 k	1.5 s
unicorn	22 k	36 k	3.8 s
elephant	39 k	68 k	12.4 s
lamb	30 k	50 k	7.1 s

The method of Entem et al. [2015] represents each part of a reconstructed 3D model by a skeleton-based convolution surface. These parts are then positioned in depth based on thickness of their 3D reconstruction and then smoothly blended together by simple summing operation. The quality of this blend strongly depends on the shape of completed regions as well as on their positions in depth. Due to this reason, unwanted bulges may appear in the final solution (see Fig. 6) as opposed to our method which guarantees to produce seamless connections between individual parts.

In addition to the rendered results presented in this paper, we include the resulting meshes for all models present in Fig. 7 and 9 in our supplementary material.

5 LIMITATIONS AND FUTURE WORK

Our method enables reconstruction of smooth high-relief models for a variety of different input drawings. However, we would like to point out some limitations of our approach which can serve as motivation for future work.

In some configurations where there are two regions smoothly interconnected and one is assumed to be above the other one, our solution may pull the top region down a bit and produce unwanted deformation of the lower region (see Fig. 10 for an example). As a future work, we plan to incorporate additional user-assisted compensation for such kind of configurations.

Our method assumes that the resulting model consists of a set of rounded shapes which have semi-elliptical profiles. Although this assumption is realistic for most organic shapes, there can still be situations which would require local modifications of the shape profile. For those, one may incorporate, e.g., the concept of curvature cues used in [Yeh et al. 2017] and modify the ∇f in Formula 10.

Although we provide a quick user-assisted mitigation of pits, an interesting avenue for future work would be an automatic approach such as determining suitable $\alpha(x)$ without user intervention. As future work, we plan to extend our technique from high-relief models to full 3D models by taking into account shape mirroring extension as used in [Feng et al. 2016].




Figure 7: Comparison of results produced by Ink-and-Ray system [Sýkora et al. 2014] (top row) and the results of our method (bottom row). The differences in smoothness are pointed out with arrows. The reverse inflation of rabbit’s ears and mouth, and farmer’s shovel may be obtained by setting the parameter d to -1 for these parts.




Figure 8: Comparison between a sequence of light variation on results produced by Ink-and-Ray method [Sýkora et al. 2014] (top) and our method (bottom). See accompanying video for an animation.

6 CONCLUSION

We presented a method to reconstruct high-relief part-based models from a single hand-drawn image. In contrast to previous techniques where the modeling process was subdivided into several independent steps, we proposed a unified non-linear energy minimization formulation which enables joint inflation and shifting of individual parts. In addition, we also proposed an efficient approximate method which delivers comparable solution as the original non-linear formulation but is notably faster. This enables us to create an interactive 3D modeling framework that enables production of high-quality meshes where individual parts are interconnected seamlessly. We confirmed the improvement in quality by comparing renderings of our resulting models and those obtained with the current state-of-the-art techniques.

ACKNOWLEDGMENTS

We would like to thank all anonymous reviewers for their fruitful comments and suggestions. We gratefully acknowledge the support of Activision, Adobe, and hardware donation from NVIDIA Corp. This research has been supported by the Technology Agency of the Czech Republic under research program TE01020415 (V3C – Visual Computing Competence Center), the Grant Agency of the Czech Technical University in Prague, grant No. SGS16/237/OHK3/3T/13 (Research of Modern Computer Graphics Methods), Research Center for Informatics No. CZ.02.1.01/0.0/0.0/16_019/0000765, the Fulbright Commission in the Czech Republic, the NSERC Discovery Grants (RGPIN-2017-05235 & RGPAS-2017-507938), a Canada Research Chair award, the Connaught Fund, and the National Science Foundation under Grant Numbers IIS-1617172 and IIS-1622360. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.




Figure 9: Comparison between the non-linear solution (a and c) and the approximate solution (b and d)—(a) results of the non-linear solution with pits, (b) corresponding approximations, (c) results of the non-linear solution with mitigated pits, (d) corresponding approximations (our final results).




Figure 10: Limitation of our method: A top region (nose) which is interconnected with a bottom region (head) may pull the bottom region down a bit which may produce unwanted deformation. This deformation is more evident when rendered from a side view (right).

REFERENCES

Sami Arpa, Sabine Süsstrunk, and Roger D. Hersch. 2015. High Reliefs from 3D Scenes. *Computer Graphics Forum* 34, 2 (2015), 253–263.

Mikhail Bessmeltsev, Will Chang, Nicholas Vining, Alla Sheffer, and Karan Singh. 2015. Modeling Character Canvases from Cartoon Drawings. *ACM Transactions on Graphics* 34, 5 (2015), 162.

Peter Borosán, Ming Jin, Doug DeCarlo, Yotam Gingold, and Andrew Nealen. 2012. RigMesh: Automatic Rigging for Part-Based Shape Modeling and Deformation. *ACM Transactions on Graphics* 31, 6 (2012), 198.

Bob Carpenter, Matthew D. Hoffman, Marcus Brubaker, Daniel Lee, Peter Li, and Michael Betancourt. 2015. The Stan Math Library: Reverse-Mode Automatic Differentiation in C++. *CoRR* abs/1509.07164 (2015).

Tao Chen, Zhe Zhu, Ariel Shamir, Shi-Min Hu, and Daniel Cohen-Or. 2013. 3-Sweep: Extracting Editable Objects from a Single Photo. *ACM Transactions on Graphics* 32, 6 (2013), 195.

Paolo Cignoni, Claudio Montani, and Roberto Scopigno. 1997. Computer-Assisted Generation of Bas- and High-Reliefs. *Journal of Graphics Tools* 2, 3 (1997), 15–28.

Even Entem, Loïc Barthe, Marie-Paule Cani, Frederic Cordier, and Michiel van de Panne. 2015. Modeling 3D animals from a side-view sketch. *Computers & Graphics* 46 (2015), 221–230.

Lele Feng, Xubo Yang, Shuangjiu Xiao, and Fan Jiang. 2016. An Interactive 2D-to-3D Cartoon Modeling System. In *Proceedings of International Conference on Technologies for E-Learning and Digital Entertainment*. 193–204.

Yotam Gingold, Takeo Igarashi, and Denis Zorin. 2009. Structured Annotations for 2D-to-3D Modeling. *ACM Transactions on Graphics* 28, 5 (2009), 148.

Gaël Guennebaud, Benoit Jacob, et al. 2010. Eigen v3. <http://eigen.tuxfamily.org>. (2010).

Takeo Igarashi and John F. Hughes. 2003. Smooth Meshes for Sketch-based Freeform Modeling. In *Proceedings of Symposium on Interactive 3D Graphics*. 139–142.

- Takeo Igarashi, Satoshi Matsuoka, and Hidehiko Tanaka. 1999. Teddy: A Sketching Interface for 3D Freeform Design. In *SIGGRAPH Conference Proceedings*. 409–416.
- Alec Jacobson, Daniele Panozzo, et al. 2013. libigl: A simple C++ geometry processing library. (2013).
- Pradeep Kumar Jayaraman, Chi-Wing Fu, Jianmin Zheng, Xueting Liu, and Tien-Tsin Wong. 2018. Globally Consistent Wrinkle-Aware Shading of Line Drawings. *IEEE Transactions on Visualization and Computer Graphics* (2018).
- Olga A. Karpenko and John F. Hughes. 2006. SmoothSketch: 3D free-form shapes from complex sketches. *ACM Transactions on Graphics* 25, 3 (2006), 589–598.
- Zohar Levi and Craig Gotsman. 2013. ArtiSketch: A System for Articulated Sketch Modeling. *Computer Graphics Forum* 32, 2pt2 (2013), 235–244.
- Chang-Jian Li, Hao Pan, Yang Liu, Xin Tong, Alla Sheffer, and Wenping Wang. 2017. BendSketch: Modeling Freeform Surfaces Through 2D Sketching. *ACM Transactions on Graphics* 36, 4 (2017), 125.
- D. C. Liu and J. Nocedal. 1989. On the Limited Memory BFGS Method for Large Scale Optimization. *Mathematical Programming* 45, 3 (1989), 503–528.
- Xueting Liu, Xiangyu Mao, Xuan Yang, Linling Zhang, and Tien-Tsin Wong. 2013. Stereoscopizing Cel Animations. *ACM Transactions on Graphics* 32, 6 (2013), 223.
- Mark Meyer, Mathieu Desbrun, Peter Schröder, and Alan H. Barr. 2003. Discrete Differential-Geometry Operators for Triangulated 2-Manifolds. In *Visualization and Mathematics III*, Hans-Christian Hege and Konrad Polthier (Eds.), 35–57.
- Andrew Nealen, Takeo Igarashi, Olga Sorkine, and Marc Alexa. 2007. FiberMesh: Designing Freeform Surfaces with 3D Curves. *ACM Transactions on Graphics* 26, 3 (2007), 41.
- Luke Olsen and Faramarz F. Samavati. 2010. Image-assisted Modeling from Sketches. In *Proceedings of Graphics Interface*. 225–232.
- H Read. 1961. *The Art of Sculpture*.
- A. Ricci. 1973. A constructive geometry for computer graphics. *Comput. J.* 16, 2 (1973), 157–160.
- R. Schmidt, B. Wyvill, M. C. Sousa, and J. A. Jorge. 2005. ShapeShop: Sketch-based Solid Modeling with BlobTrees. In *Eurographics Workshop on Sketch-Based Interfaces and Modeling*. 53–62.
- Robert Schneider and Leif Kobbelt. 2001. Geometric fairing of irregular meshes for free-form surface design. *Computer Aided Geometric Design* 18, 4 (2001), 359–379.
- Christian Schüller, Daniele Panozzo, and Olga Sorkine-Hornung. 2014. Appearance-Mimicking Surfaces. *ACM Transactions on Graphics* 33, 6 (2014), 216:1–216:10.
- Alex Shtof, Alexander Agathos, Yotam Gingold, Ariel Shamir, and Daniel Cohen-Or. 2013. Geosemantic Snapping for Sketch-Based Modeling. *Computer Graphics Forum* 32, 2 (2013), 245–253.
- Daniel Šykora, Ladislav Kavan, Martin Čadík, Ondřej Jamriška, Alec Jacobson, Brian Whited, Maryann Simmons, and Olga Sorkine-Hornung. 2014. Ink-and-Ray: Bas-Relief Meshes for Adding Global Illumination Effects to Hand-Drawn Characters. *ACM Transactions on Graphics* 33, 2 (2014), 16.
- Daniel Šykora, David Sedláček, Sun Jinchao, John Dingliana, and Steven Collins. 2010. Adding Depth to Cartoons Using Sparse Depth (In)equalities. *Computer Graphics Forum* 29, 2 (2010), 615–623.
- Chiew-Lan Tai, Hongxin Zhang, and Jacky Chun-Kin Fong. 2004. Prototype modeling from sketched silhouettes based on convolution surfaces. *Computer Graphics Forum* 23, 1 (2004), 71–83.
- Gabriel Taubin. 1995. A Signal Processing Approach to Fair Surface Design. In *SIGGRAPH Conference Proceedings*. 351–358.
- Tim Weyrich, Jia Deng, Connelly Barnes, Szymon Rusinkiewicz, and Adam Finkelstein. 2007. Digital Bas-Relief from 3D Scenes. *ACM Transactions on Graphics* 26, 3 (2007), 32.
- Baoxuan Xu, William Chang, Alla Sheffer, Adrien Bousseau, James McCrae, and Karan Singh. 2014. True2Form: 3D Curve Networks from 2D Sketches via Selective Regularization. *ACM Transactions on Graphics* 33, 4 (2014), 131.
- Chih-Kuo Yeh, Shi-Yang Huang, Pradeep Kumar Jayaraman, Chi-Wing Fu, and Tong-Yee Lee. 2017. Interactive High-Relief Reconstruction for Organic and Double-Sided Objects from a Photo. *IEEE Transactions on Visualization and Computer Graphics* 23, 7 (2017), 1796–1808.
- Chih-Kuo Yeh, Pradeep Kumar Jayaraman, Xiaopei Liu, Chi-Wing Fu, and Tong-Yee Lee. 2015. 2.5D Cartoon Hair Modeling and Manipulation. *IEEE Transactions on Visualization and Computer Graphics* 21, 3 (2015), 304–314.
- Qiong Zeng, Wenzheng Chen, Huan Wang, Changhe Tu, Daniel Cohen-Or, Dani Lischinski, and Baoquan Chen. 2015. Hallucinating Stereoscopy from a Single Image. *Computer Graphics Forum* 34, 2 (2015), 1–12.