

Stress Relief: Improving Structural Strength of 3D Printable Objects

Ondrej Stava¹ Juraj Vanek¹ Bedrich Benes¹ Nathan Carr² Radomír Měch²

¹ Purdue University ² Adobe Systems Incorporated




Figure 1: A model of a dragon was first hollowed to reduce stress caused by its weight if held by the head (a). The stress decreased, but the neck had to be still thickened (b,c). The object was still too front heavy causing a twist deformation on the legs (c), that was eliminated by fixing the model to the pedestal by a strut (red) (d). These steps were done automatically by our system.

Abstract

The use of 3D printing has rapidly expanded in the past couple of years. It is now possible to produce 3D-printed objects with exceptionally high fidelity and precision. However, although the quality of 3D printing has improved, both the time to print and the material costs have remained high. Moreover, there is no guarantee that a printed model is structurally sound. The printed product often does not survive cleaning, transportation, or handling, or it may even collapse under its own weight. We present a system that addresses this issue by providing automatic detection and correction of the problematic cases. The structural problems are detected by combining a lightweight structural analysis solver with 3D medial axis approximations. After areas with high structural stress are found, the model is corrected by combining three approaches: hollowing, thickening, and strut insertion. Both detection and correction steps are repeated until the problems have been eliminated. Our process is designed to create a model that is visually similar to the original model but possessing greater structural integrity.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically based modeling; I.3.8 [Computer Graphics]: Applications;

Keywords: 3D printing, structural analysis, physics-based modeling

Links: DL PDF WEB

ACM Reference Format

Stava, O., Vanek, J., Benes, B., Carr, N., Měch, R. 2012. Stress Relief: Improving Structural Strength of 3D Printable Objects. *ACM Trans. Graph.* 31 4, Article 48 (July 2012), 11 pages. DOI = 10.1145/2185520.2185544 <http://doi.acm.org/10.1145/2185520.2185544>.

Copyright Notice

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701, fax +1 (212) 869-0481, or permissions@acm.org.
© 2012 ACM 0730-0301/2012/08-ART48 \$15.00 DOI 10.1145/2185520.2185544
<http://doi.acm.org/10.1145/2185520.2185544>

1 Introduction and Related Work

3D printing enables the physical realization of 3D shapes designed using a computer. Originally, 3D printing referred to powder deposition printing by Z Corporation, but since then it has been used in connection with any additive manufacturing process [Shapeways 2011a], and the various techniques have become an important part of rapid prototyping processes (see the list at [Engineering Handbook 2011]). A 3D printed object is created layer by layer, similar to how a 2D image may be printed line by line. Fused deposition modeling and jetting liquid polymer methods deposit or jet new material for each layer of the printed object. Selective laser sintering, electron beam melting, and powder deposition printing melt or bind powdered material to produce the layers. Many manufacturers produce 3D printers, and their prices are rapidly dropping, though the precision, printable color selections, and material types are increasing, making them more accessible to the general public.

The promise of seamlessly moving creations from a virtual space into reality is truly tantalizing, and its applications go far beyond basic manufacturing and rapid prototyping. Unfortunately, many obstacles remain for 3D printing to be practical and commonplace. Although the quality of 3D printing has increased, both the time to print and the material costs have remained high. For a model to be 3D printable, it must meet a set of topological requirements to ensure that its interior and exterior volumes are clearly defined. This is usually achieved by making the input mesh a watertight manifold surface. Producing watertight manifold models from, for example, point clouds or surface representations has been well studied and is available in existing systems [Hornung and Kobbelt 2006; Liu et al. 2007]. Another common requirement is that the input mesh has no parts thinner than the printer can print. However, these requirements do not reflect the structural stability of the model. The manufactured objects are commonly so fragile that they do not survive transportation, cannot be handled, and even collapse under their own weight.

Whereas the structural properties are a well-studied subject in other disciplines such as mechanical engineering, printability of virtual objects for 3D were addressed only in a limited extent within the context of Computer Graphics. In this context, procedural methods have been coupled with physics-based analysis for masonry buildings [Whiting et al. 2009] and structural analysis has been used for

3D virtual growth of plants and environmental response [Jirasek et al. 2000; Hart et al. 2003]. In general, a structurally stable object can be obtained using methods based on shape optimization [Allaire 2002; Haslinger and Mäkinen 2003], where the geometry of a 3D model is modified to minimize certain objective function. However, this approach is computationally very expensive for even moderately complex input models (see Section 4.5 for more details). Companies providing 3D printing usually publish a set of constraints that any 3D model should satisfy before it is printed [Shapeways 2011b; Z Corporation 2011] and users are thus responsible to manually check and correct the model before printing. This process can be very cumbersome and time-consuming even for experienced users. Also, these constraints define only local properties, such as minimal thickness, and they do not capture all possible structural problems. The first automatic method devised to detect structural issues of printed objects was proposed by [Telea and Jalba 2011]. The method detects thin and thick parts and then uses a set of basic geometric rules to determine whether a given thin part can support its attached parts. Although the method detects some problematic cases, it lacks the structural analysis component necessary to accurately approximate stress in more complex topological forms. Furthermore, this work does not address the fundamental goal of correcting the detected structural issues. Even in the presence of perfect analysis, it may not be obvious or easy for the user to know where to start making corrections. In our work, we underscore the importance of having both computer-assisted correction and detection combined in a tightly coupled system.

We propose a first step in the direction of physically based automatic solutions for the detection and correction of major structural problems in 3D models before they are printed, while attempting to minimally alter their appearance as shown in Fig. 1. We first detect major printability issues of the input model through printability assessments. In our work, we focus on structural problems resulting from forces that can occur when the object is standing or when it is held at probable locations, such as pedestals or appropriately sized handles. We then correct the model by using three different approaches: local thickening, strut adding, and hollowing. Local thickening increases the strength of the object's thin parts, and it preserves the form and the surface detail of the object. Struts can be inserted to prevent nonrigid deformations of the object, and they are automatically placed at locations that are least visible to the user. Hollowing eliminates material inside the object, reducing the weight of certain parts to alleviate the stress at other locations. Hollowing may also lower the cost by reducing the amount of used material. The correction technique is chosen based on the visual impact on the result, and it can be further adjusted by user parameters. After a correction is made, the process repeats until all structural problems are eliminated. We show our technique on various objects, such as game characters or technical models. We measured the stress on original and corrected objects, and we experimentally verified that the corrected models are more robust. Our paper continues with a description of the overall schema of the system in Section 2.

2 System Pipeline

Figure 2 shows the overall schema of our system. The input model is a triangular watertight mesh representing a 3D manifold that is analyzed to detect possible structural problems. We first detect issues resulting from gravity acting on a standing object and from humans manipulating it by pressing the object at various handling points. The model is then iteratively evaluated and corrected. At each iteration the structural analysis module reports whether the stress is above the desired level, and the best model correction option is evaluated and chosen. The options for the structural model

corrections are (1) adding a strut (2) thickening some parts, or (3) hollowing some part of the model. Care is taken not to modify the model in a way that would cause a significant visual alteration. The result is a modified triangular mesh with improved structural properties. In the following text, we discuss individual modules of the pipeline in detail. The structural problems of general objects are very complex because they depend on the material, printer, and the object itself; thus the complete solution is not within the scope of this paper. We present a general pipeline and a solution that can be applied to a wide variety of input models printed on an arbitrary 3D printer. Please note that the implementation of individual modules can be adjusted to provide better accuracy or to meet requirements that are specific to the used 3D printing technology or to the type of the printed object.




Figure 2: The overall schema of our method. The potential structural loads are generated and the object is repeatedly analyzed for structural problems that are corrected by one of the three proposed methods. The analysis stops when the stress is under a user-specified level.

3 Printability Analysis

Printing a 3D object is a complex process. Each printer manufacturer provides its own printer drivers that address some issues. For example, to print an object with concave or topologically disconnected parts, a support medium is introduced. Once the model has been printed, the support must be removed by blowing, melting, or washing. A successfully printed model must not only survive printing, but also subsequent handling during the cleanup phase and regular use.

To perform printability analysis, we first determine the most probable structural loads that the object may be subject to. Each load is defined by a set of forces and constraints that are used for a physically based structural analysis of the object.

3.1 Determining Structural Loads

It would be difficult to determine all possible structural loads that a given object must withstand; thus we are looking only for the set of highly probable cases. We divide the loads into two categories: permanent and imposed. Permanent loads occur when the object is placed in a resting position on a flat surface. It is deformed by gravity while its movement is constrained by the flat plane where it rests. Imposed loads capture situations when the object is handled and is thus subject to both gravity and external forces.

Permanent Loads To determine the permanent loads we find a set of orientations for which the object can be stably placed onto a flat surface. Each orientation is tested to determine whether the model can support its own weight. We assume that the model's *default upright orientation* is known. However, it may be automatically determined by the algorithm of Fu et al. [2008]. We find a set of additional valid orientations for the object placement by closely




Figure 3: The convex hull and the center of mass are used to find the most probable rest orientations of the object.

following concepts detailed in [Fu et al. 2008]. We compute the convex hull of the object along with its center of mass (see Fig. 3). Each face of the convex hull represents a candidate orientation for placing the object on a flat surface. We test each face by orthogonally projecting the center of mass onto the face plane. If the projected center falls outside the face, the configuration is unstable. We sort the stable faces by their decreasing area and then move the default orientation to the head of the list. The area criterion provides a simple heuristic to prioritize the more likely orientations if the object were to be randomly placed down. The prioritized ordering allows test cases to be truncated into a smaller set when computing resources are limited. We note that the set of cases may be reduced further by exploiting symmetries.

Imposed Loads An object can be manipulated in many different ways. Because of the physical size of objects that are usually printed, we restrict ourselves to the analysis of loads generated by a *pinch grip*, where the object is grasped by two fingers. Two-finger grips represent only a limited sub-set of possible handling configurations, but they are especially important for printability analysis because the forces applied to the model are higher when only two fingers are used.

To further reduce the search space of grip configurations, we first create a set of plausible grip positions that are then evaluated and from which the best ones are selected. To determine plausible grip positions, we first mark each face center of the input mesh as a possible grasp site P_1 for the first finger. For each site P_1 , we define a set of grip directions that determine the axis of a single pinch grip. The set of directions should represent the most natural ways by which a person would grasp the given object.

As shown in [Balasubramanian et al. 2010], people tend to instinctively hold objects in such ways that the center of the mass is located either in the grip axis itself or somewhere on the plane defined by the grip axis and the gravity vector. To find such configurations for every site P_1 , we define one grip direction that connects the site to the center of the object’s mass and one direction that points to the center of mass projected to a plane that is parallel to the ground level. Another common case is holding the object by some local feature, such as the handle of a cup. We capture these cases by creating a grip direction that points to the closest point on the medial axis to P_1 , and we create one ray for each of these directions. We then find all intersections of the ray with the object, and for each one we create a second grasp site P_2 that, together with the given site P_1 , forms a grip configuration G_{P_1, P_2} .

The above process results in a large number of possible grip configurations that is lessened by analyzing geometric properties of the grip configurations and checking how accessible they are for human

fingers. The value of each configuration Λ_G is then computed as

$$\Lambda_G = \prod_{\mathbf{p} \in \{P_1, P_2\}} (\mathbf{n}_{A_{\mathbf{p}}} \cdot \mathbf{l}_{G, \mathbf{p}}) O_G(\mathbf{p}), \quad (1)$$

where we compare the grip axis direction $\mathbf{l}_{G, \mathbf{p}}$ to the average normal direction $\mathbf{n}_{A_{\mathbf{p}}}$ on the grasped surface $A_{\mathbf{p}}$. The size of the surface $A_{\mathbf{p}}$ is defined by the average size of a human finger and its center is the position \mathbf{p} of a given grasp site P . This ensures that we preferably choose grip configurations where the grasped surface is flat and aligned with the grip axis. The accessibility of each site is then estimated using ambient occlusion $O_G(\mathbf{p})$.

We select N_G grip configurations with highest Λ while keeping them significantly different from each other. Equation (1) would lead to similar Λ values for closely located grip configurations. For each selected grip configuration G_s , we multiply the value of Λ for all remaining configurations with a penalty term $\sqrt{d_{G_i, G_s}}$, where d_{G_i, G_s} is the average geodesic distance between grip sites.

The force F_{grip} of the grip is evaluated using an empirical model of [Kinoshita et al. 1997] that predicts the grip force based on the tangential force F_t and the tangential torque T_z on the fingers holding the object:

$$F_{grip} = \frac{F_t}{\mu_i} + \frac{|T_z|}{\mu_{rot}} - \frac{0.011}{\mu_i} F_t |T_z|, \quad (2)$$

where μ_i and μ_{rot} are the coefficients of linear and rotational friction, respectively, and their values are determined by the properties of the material. The tangential force is $F_t = F_g (\mathbf{l}_G \times \mathbf{g})$, where \mathbf{g} is the direction of gravity and F_g is the gravitational force. The value of tangential torque depends on the relative positions of the grasp sites \mathbf{p}_1 and \mathbf{p}_2 to the center of mass of the object \mathbf{c}_o as $T_z = F_g \sum_{i=1,2} \mathbf{l}_G \cdot ((\mathbf{p}_i - \mathbf{c}_o) \times \mathbf{g})$.

Note that if the grip axis \mathbf{l}_G is parallel to \mathbf{g} , both F_t and T_z are zero and the given grip requires no force. However, it is still important to test these configurations because the effects of gravity alone might be significant, especially when \mathbf{l}_G is far away from the center of mass of the object. An example of the first five detected pinch grips for a model of a kitten is shown in Fig. 4.




Figure 4: Pinch grip positions. Each grip pair has the same color.

3.2 Structural Analysis

The structural stability of a printed model depends on the computed load cases and on the printing technology and material that are used. To make our approach applicable to a wide variety of 3D printing settings, we compute the stability from the solution of the linear elasticity problem considering a homogenous material. The accuracy of the structural analysis can be improved by using more advanced methods targeted at specific 3D printers and materials.

The input object is converted into a tetrahedral mesh using method in [Si 2011], and the stress is computed from the solution of the discrete linear elasticity problem using the Finite Element Method (FEM) with quadratic tetrahedral elements [Hughes 1987]:

$$\mathbf{K}d = \mathbf{F}, \quad (3)$$

where \mathbf{K} is the stiffness matrix and d is the deformation of the mesh caused by external and internal forces represented by \mathbf{F} .

The tetrahedral mesh and the material properties are used to compute the stiffness matrix \mathbf{K} of the linear system that is reused for all load scenarios. Each load yields the boundary conditions and the right-hand side \mathbf{F} as described in Section 3.2.1. The solution to the linear elasticity problem is a vector of deformations d of the input mesh under the given input conditions. The deformations are used to compute the strain and stress in the input mesh using Cauchy's linear strain tensor and Hooke's law, respectively. To determine the structural problems, we use the von Mises yield criterion [von Mises 1913] that specifies the critical value of von Mises.

3.2.1 Analysis of Loads

To solve the system in Eq. (3) for a given load scenario, we need to define the boundary conditions and the right-hand side \mathbf{F} . The boundary conditions specify fixed vertices of the tetrahedral mesh that have a prescribed deformation d_v equal to zero. The right-hand side \mathbf{F} is computed from the body forces and external forces that affect the printed object, and it can be decomposed into the sum of elemental contributions $\mathbf{F} = \sum \mathbf{f}^e$ that are defined as

$$\mathbf{f}^e = \int_{\Omega_e} \mathbf{N}^T \mathbf{b} \, d\Omega + \int_{\Gamma_e} \mathbf{N}^T \mathbf{t} \, d\Gamma, \quad (4)$$

where Ω_e is a tetrahedral element e , Γ_e is one face of the tetrahedron, \mathbf{N} is the matrix of quadratic basis functions, \mathbf{b} is the body force on the element (in our case only gravity), and \mathbf{t} is the surface load at the given face of an element.

Permanent loads are caused by gravity acting on an object standing on a flat surface; therefore the fixed vertices can be identified by locating all vertices that lie within some distance to the flat surface. Also, the only force affecting the object is gravity (a body force \mathbf{b} in Eq. (4)). The surface load $\mathbf{t} = 0$ everywhere on the object. To compute the right-hand side \mathbf{F} we need to integrate the gravity over each element Ω_e . Since we use quadratic basis functions, an accurate integral can be computed using numerical integration based on the Gaussian quadrature rule [Golub and Welsch 1967].

Imposed loads are caused by two fingers holding the object, subjecting the object to both gravity and external forces generated by the grip itself. We decompose the problem into two different cases: one to account for gravity without any external forces and one to analyze the pressure caused by the grip while ignoring the effects of gravity. The decomposition is necessary because the structural analysis method does not support friction, which forces us to fix all grasped sites when testing the effects of gravity. The fixed site then defines new boundary conditions of our system while the right-hand side is computed in the same way as for the permanent loads.

When testing the pressure caused by the grip, we fix all points that are grasped by one finger while we apply the entire required grip force vector $\mathbf{F}_{grip} = F_{grip} \mathbf{I}_G$ to all affected mesh faces Γ_G that are located under the second grasp site. The load \mathbf{t}_τ that affects a given mesh face τ is computed from its size and orientation:

$$\mathbf{t}_\tau = \mathbf{F}_{grip} \frac{\max(0, -\mathbf{n}_\tau \cdot \mathbf{F}_{grip})}{\sum_{\gamma \in \Gamma_g} A_\gamma \max(0, -\mathbf{n}_\gamma \cdot \mathbf{F}_{grip})}, \quad (5)$$

where A is the area of a given face and \mathbf{n} is its normal.

3.2.2 Stress Analysis

Once all boundary conditions and forces are defined, the system from Eq. (3) is solved and the deformation of the mesh is computed. The deformation is then used to solve the linear strain tensor ε and stress tensor σ :

$$\varepsilon = \frac{1}{2}((\nabla \mathbf{d})^T + \nabla \mathbf{d}) \quad (6)$$

$$\sigma = \mathbf{C}\varepsilon, \quad (7)$$

where \mathbf{C} is the elasticity matrix that contains properties of the printed material.

To compute the critical level of the stress, we use the von Mises yield criterion [von Mises 1913] that compares the von Mises stress σ_{vm} with a material-specific yield strength S_m . The value of von Mises stress is

$$\sigma_{vm}^2 = \frac{(\Delta\sigma_{1,2})^2 + (\Delta\sigma_{2,3})^2 + (\Delta\sigma_{3,1})^2 + 6(\sigma_{12}^2 + \sigma_{23}^2 + \sigma_{31}^2)}{2}, \quad (8)$$

where $\Delta\sigma_{i,j} = \sigma_{ii} - \sigma_{jj}$.

The critical level of the stress σ_c is then defined as $\sigma_c = \lambda S_m$, where λ is a constant $0 < \lambda < 1$ that specifies the maximum allowed value of stress with respect to the yield strength of the material S_m . By decreasing the value of λ , we increase the sensitivity of our method so that it detects more printability problems but can lead to a generation of false positives. From our experiments, a value $\lambda = 0.2$ provides reasonable results. At the end, we compare the maximum value of the von Mises stress $\max \sigma_v$ with the maximum allowed critical stress σ_c . If the value of the von Mises stress is higher than the critical value, we say that the input object has a printability problem.

4 Printability Corrections

A detected printability problem indicates that it may be unsafe to print the object in its current state. To remove these problem, we reduce the stress in the input object so that it is below a given critical value σ_c that specifies a potential structural failure of the used material by using local thickening, strut adding, and hollowing. The particular correction is automatically selected based on its impact on the visual appearance of the model and its effectiveness.

First, we pick the structural load that is associated with the highest stress in the tested model. We then apply each of our model correction methods and generate a list of all possible corrections. There could be more than one suggested correction for each of the three methods. Every correction from the list is evaluated for its effectiveness δ_ψ^e and for its impact on the appearance δ_ψ^v . These factors are then used to compute the final cost of each correction γ as

$$\gamma = w_\psi (w_e \delta_\psi^e + w_v \delta_\psi^v), \quad (9)$$

where w_ψ is a user-defined weight of a given type of printability correction, and w_e and w_v are weights defining the importance of effectiveness and visual appearance, respectively. A correction with the lowest cost is then applied. For most of the analyzed objects, we used an equal weight w_ψ for all corrections, while the effectiveness and visual importance were set to $w_e = 0.7$ and $w_v = 0.3$.

4.1 Medial Geometry Analysis

The results of the structural analysis are stored in a 3D tetrahedral grid that represents the object. In general, the structure of 3D objects can be complicated; therefore we reduce the degree of complexity of the problem by analyzing the computed stress on the medial axis of the object rather than on the object itself.




Figure 5: The interior medial axis simplifies the analysis of the results. Red indicates areas of high stress.




Figure 6: It is not sufficient to apply thickening locally on highly stressed parts of the model (in red) (left) because the stress would be simply moved somewhere else (right).

We first construct an approximate medial axis for both the interior and exterior of the object following the method described in Dey and Zhao [2003]. The medial axis is prone to include minor surface perturbations, which may result in many extraneous branches that often correspond to insignificant surface details. We remove these branches using a topology-preserving pruning based on a separation angle strategy as described in Liu et al. [2010]. The value of the separation angle can be used to define the size of the surface detail that we want to preserve.

Lastly, we compute the maximal stress for every medial point on the interior medial axis by gathering the maximum stress value from its associated nearest tetrahedrons (see Fig. 5).

4.2 Local Thickening

Many of the structural problems are caused by excessive stress in thin parts of the printed object that can easily break or fracture. The goal of local thickening is to identify and fix the problematic thin parts of the model.

Component Analysis We first identify candidate parts for thickening. The excessive stress is often concentrated only in a limited volume, but in general it is not sufficient to thicken the mesh only around the stressed areas, because such an operation would only move the stress somewhere else (see Fig. 6). To solve this problem, we must identify the entire thin component that should be thickened as a whole. When such a component is thickened, the stress is redistributed evenly across its entire volume, which effectively reduces the stress below the critical value. Also, a small thickening of a large part of the model has a smaller visual impact than excessive modification of a local part.

To determine which part of the mesh should be considered a thin component, we need to take into account the value of the stress that we want to reduce. When the stress is high, we need to thicken larger parts of the model than when the stress is less. We use combined information about the stress and about the local thickness and topology that is stored on the computed interior medial axis, as described in Section 4.1.

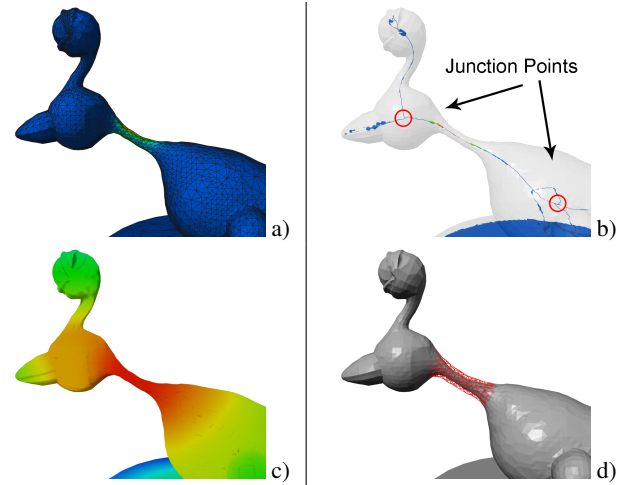


Figure 7: To reduce the stress in the object (a), we identify the linear chains of the medial axis with an excessive stress. (b). The part of the component that must be thickened is determined by the value of the stress and by the distance of individual vertices from the medial points of the component (c). The thickness of the component together with the stress is used to compute a thickening correction (d).

We limit the length of components that can be thinned by non-manifold junctions in the medial axis as shown in Fig. 7 and by the threshold thickness that separates a given thin part from thick parts. To compute the threshold thickness, we take the location of the maximal stress σ_m and find its closest medial point. Based on the local thickness t_m associated with the medial point, we can then compute a new thickness t_d that would reduce the stress to the allowed critical value σ_c . Since most of the problematic components resemble thin tubes, we simplify the situation by assuming that the component has a beam shape with a circular cross-section. If needed, this approximation can be improved by detecting the actual cross-section of the component at the point with the highest stress. Using the second modulus of a cross-section, it can be shown that the stress σ caused by bending forces on a beam is directly related to the thickness t of the cross-section as

$$\sigma = c_b/t^3, \quad (10)$$

where c_b is a constant value specific for a given position at the beam that is independent to the local thickness t . Please note that this relation is valid only when the beam is subject to bending forces, and it does not hold for general cases. However, most cases that will be encountered when handling a printed object would result either in bending forces or axial forces. Axial forces usually result in lower stress than bending forces, thus thickening the surface enough to reduce the bending forces sufficiently reduces stress to also cover axial forces.

We can then use Eq. (10) to determine the target thickness t_d as

$$t_d = t_m \sqrt[3]{\sigma_m/\sigma_c}. \quad (11)$$

The threshold thickness t_t that separates the thin part from the thick parts is defined as $t_t = c_t t_d$, where c_t is a predefined constant (we set $c_t = 1.1$ in all our examples).

We then compute the set of points on the medial axis that are bounded by the threshold thickness. The surface region corresponding to these medial points defines the thin components of the mesh.

Mesh Displacement To make a given thin component thicker, we displace the original mesh away from the medial axis. We displace only vertices whose distance t_v to the nearest medial point on the thin component is less than t_t . The amount of displacement of the given vertex is computed as $d_v = (t_t - t_v)/2$, where t_t is the threshold thickness of the given component.

The target offset surface can be generated using many existing techniques, such as level sets [Cohen et al. 1996; Peng et al. 2004]. However, these methods do not preserve small surface details that are important for the visual appearance of the models. To solve this issue, we displace individual vertices away from their nearest medial points. To avoid self-intersections, we displace the vertices along a smooth distance field that is computed as a solution to Laplace’s equation on a tetrahedral grid that discretizes space within the exterior medial axis of the object. The grid contains a single point for each medial vertex and for each surface vertex of the mesh. All of these vertices are then used to define the boundary conditions of Laplace’s equation, using their distance to the nearest internal medial point. The system is then solved using the finite element method with quadratic tetrahedral elements. The displacement is equal to traversal of the resulting distance field, which is implemented using the Runge-Kutta integration method [Butcher 2008]. The traversal ends when a given vertex reaches the distance equal to $\frac{t_v}{2} + d_v$.

Thickening Evaluation The input model may contain several thin components that are subject to excessive stress. Although we might thicken the components one by one, but that might not always lead to the best structural and visual results. If we have several problematic thin components, and if we thicken only one, the thickened component might gather some stress from the remaining others. This can lead to uneven thickening of different components, even if their initial parameters are the same. Therefore when we evaluate the thickening corrections, we compute the cost of all combinations of nonoverlapping thickenings.

The effectiveness of a thickening correction δ_i^e that consists of n_t nonoverlapping thickenings is computed as

$$\delta_i^e = w_\sigma^t \prod_{i < n_t} \frac{\sigma_c}{\sigma_{t_i}} + w_V^t \frac{1}{n_t} \sum_{i < n_t} \frac{V_{t_i}}{V_{c_i}}, \quad (12)$$

where σ_c is the critical stress as defined in Section 3.2.2. The value of σ_{t_i} is the current maximal stress on the thickened component i , and V_{t_i} is the new volume added to the original volume of the component V_{c_i} . w_σ^t and w_V^t are user-defined weights that determine the importance of stress reduction compared to the new volume of the thickening. In our examples, we used $w_\sigma^t = 0.8$ and $w_V^t = 0.2$.

The visual impact of a thickening on the appearance of the object can be determined from the changes in the visual profile of the thickened component. We define the visual profile as the average visible area from several predefined directions. To compute the profile, we render the scene from a set of n_v view directions, and we measure the number of pixels that were rendered for the given component. For the evaluation, we measure the visual profile of the original object r_o , of the thin component r_c , and of the thickened component $r_{c'}$. The visual impact of one thickening is then computed as

$$\delta_i^v = \frac{1}{n_v} \sum_{i < n_v} \left(w_{r_c}^t \frac{r_{c_i}}{r_o} + w_{r_{c'}}^t \frac{r_{c'_i}}{r_{c_i}} \right), \quad (13)$$

where weights $w_{r_c}^t$ and $w_{r_{c'}}^t$ control the importance of choosing a visually small part for thickening compared to the size of the whole object and not locally changing the size too much, respectively.




Figure 8: An imposed pressure load caused an excessive stress in the material (left). The search space for possible struts is limited by a part of exterior medial axis (purple color) that automatically identifies all concave parts of the input model for possible connections by a strut (right).

For our experiments, we used $n_v = 5$ different directions sampled over the hemisphere defined by the upright position of the object: one direction pointing up and the remaining four uniformly distributed along a circle defined by a polar angle $\pi/4$. We found that the default setting of $w_{r_c}^t = 0.3$ and $w_{r_{c'}}^t = 0.7$ produced good results.

4.3 Adding Struts

The local thickening can be effectively used to solve many of the structural problems of the printed objects, but in certain cases, the thin components that cause the printability issues are essential for the object’s visual style. When they are made noticeably thicker, the objects might be visually altered too much. Also, the local thickening is sometimes unfeasible from a practical point of view, such as when the printed object consists of many interconnected thin components, and to remove the stress, we would need to thicken the whole object. Our second model correction technique attempts to solve these problems by adding supporting structures, that is, struts. Unlike local thickening that thickens affected components, the struts are designed to remove the stress by addressing its cause, which is a nonrigid deformation in the object. Our method automatically places struts between parts that are deformed, which moves the stress from the problematic parts into the struts.

Strut Location Detection We can usually find many different struts that would effectively reduce the stress in the input object. The suitability of each strut depends on several factors, including its visibility, thickness, length, and many others (see below for more details). To find a good location for a strut, we first detect a set of all plausible strut locations and we evaluate them to select the best ones.

Struts are essentially new, straight connections between two parts of the object, which means that they can be constructed only in exterior concave regions of the input objects; otherwise, drawing non-intersecting line between end points of a strut would be impossible. The concave regions can be directly identified from the exterior medial axis, as described in Section 4.1. The exterior medial axis defines sheets that every strut must intersect, which greatly reduces the search space for possible struts (Fig. 8).

The struts can intersect the medial axis at any point and angle. To obtain a finite number of struts, we find a subset of medial points using a uniform sampling of the medial axis. Each point from this

subset defines a set of struts that pass through the point at a few predefined angles. The angles are computed from the normal vector on the medial sheet and from the normal and deformation vectors of the closest mesh vertices. Some of these struts may not be valid, since they can point outside of the object. Therefore for every medial point and every angle, we shoot a ray in both directions and we check whether it intersects the input object. The intersection points are the end points of a strut. In our examples we generated nine uniformly distributed directions for every medial point. We generate additional struts by shooting rays from the end points of the struts that are passing the medial points. The direction of these rays corresponds to the direction of the deformation of the mesh surface at the given point that is obtained from the structural analysis as explained in Section 3.2. The struts obtained from these rays act, in general, against the normal deformations of the object while they minimize the effects of skew deformations, which are unwanted because they can easily break the struts.

Strut Thickness The thickness of a strut depends on the strength and orientation of forces that act on its end points. Unfortunately these forces are unknown and cannot be directly computed from the known displacement of the input object. Therefore we use an approximate solution that usually leads to good results, noting that if the computed thickness is insufficient, it would be thickened by the local thickening correction later on.

To compute the thickness, we first find the most stressed medial point that lies on a geodesic path along the medial axis between the two end points of the strut. The stress value σ_m and the thickness at the selected medial point is then used to compute a new desired thickness of the component using Eq. (11). We use this thickness as the thickness of the strut, which usually leads to a structurally stable correction no matter the nature of forces affecting the ends of the strut.




Figure 9: The most effective strut to reduce the stress of the object from Fig. 8 (a). The geodesic curve connecting the strut end points is marked in green, and the stress is significantly reduced across the model (c). A strut that prefers the visual impact over efficiency (b) transfers a portion of the stress from the left to the right; however, it accumulates stress at the strut end points (d).

Strut Evaluation We estimate the thickness of a strut independent of the forces that affect its end points, but in reality the efficiency of a strut is heavily influenced by its orientation and position. In general, struts that are subject to axial forces are more stable compared to similar struts that must withstand many bending forces. Therefore to evaluate the efficiency of a strut, we take into

consideration the direction of forces acting on the end points of the strut. Although these forces are unknown, we can still estimate their directions from the computed displacement \mathbf{d} of the analyzed object. To evaluate the locations \mathbf{p}_{Fp} and directions \mathbf{F}_{dp} of forces affecting a given strut end point \mathbf{p} , we compute a weighted sum of all displacements of the medial points that are within some threshold distance ε_d to the medial point associated with the point \mathbf{p} . The volume of the object that is considered during evaluation of the effectiveness of the strut is determined by ε_d . The amount of supported volume mostly depends on the geodesic distance ν_s between the two end points of the strut; therefore we set $\varepsilon_d = \nu_s/4$. \mathbf{F}_{dp} is then decomposed into normal component \mathbf{F}_{np} and tangential component \mathbf{F}_{tp} with respect to the direction of the strut. Furthermore if the normal force is eccentric with respect to the position of the strut, we need to compute the bending moment $\mathbf{M}_{bp} = (\mathbf{p}_c - \mathbf{p}_{Fp}) \times \mathbf{F}_{np}$, where \mathbf{p}_c is the midpoint of the strut. The total strength of normal force that affects the strut is then $F_{ns} = \|\mathbf{F}_{np0} - \mathbf{F}_{np1}\|$, the strength of tangential force is $F_{ts} = \|\mathbf{F}_{tp0} - \mathbf{F}_{tp1}\|$, and the total moment is $M_{bs} = \|\mathbf{M}_{bp0} + \mathbf{M}_{bp1}\|$.

We then evaluate the efficiency of a given strut as

$$\delta_s^e = w_\sigma^s \frac{\sigma_c}{\sigma_m} \frac{M_{bs} + F_{ts}}{F_{ns}} + w_V^s \frac{V_s}{V_o}, \quad (14)$$

where σ_m is the maximal stress detected between two end-points of the strut, V_s is the volume of the strut, and V_o is the volume of the printed object. Both factors can be weighted by users, similar to the Eq. (12). In our examples we used $w_\sigma^s = 0.8$ while the effect of the volume was only $w_V^s = 0.2$.

The visual impact of a strut is defined as

$$\delta_s^v = w_{vi}^s O_s \frac{V_s}{V_o}, \quad (15)$$

where O_s is the visibility of a given strut. We approximate the visibility by computing the ambient occlusion on several sample points on a surface of each strut. Because struts are visually more disturbing than other corrections, we set the weight of the visual importance to $w_{vi}^s = 2$. An example of selecting different struts based on varying weights is shown in Fig. 9.

Strut Geometry When a strut is selected to correct the input object, its geometric representation must be created, and it must then be attached to the original mesh. The geometry of a strut is a tube with both ends open and skewed according to the normal vector at the target surface. Each open end of the tube is then projected to the surface, and faces within a predefined geodesic distance to the projected area are removed from the mesh. This creates an open hole in the original mesh that must be connected to the end of the tube by generating a new triangulation of the open space between the two open holes. To do so, we first embed both holes into a circle as described in [Schmidt and Singh 2010]. Both circles are then projected to a single plane, and a new triangulation is created using the algorithm proposed by [Shewchuk 2001]. Vertices belonging to the open holes are then projected back to the global space. The remaining new interior vertices of the connections are then smoothed by solving a linear bi-Laplacian system [Andrews et al. 2011] that is constrained by the positions and tangent vectors of the vertices that belong to the open holes.

4.4 Hollowing

Should the excessive stress emerge during testing of the gravity effects of both permanent and imposed loads, we know that the source of the stress must be the gravity force. Because the strength of the gravity force depends on the mass of the object, we try to reduce its effects by reducing the mass of the printed object.

Problems caused by gravity usually emerge when we have a thick component that is supported by thin components that cannot hold its weight. Therefore we first use the maximal value of stress in the model to identify thin components, applying the same method as described in Section 4.2. All the remaining parts of the model are then considered thick components.




Figure 10: Different loads might lead to different hollowing solutions. If we hold the object by the top, only the bottom components are hollowed (left). On the other hand, if the object sits on a surface all spherical components are hollowed while the pedestal remains intact (right).

To learn which parts should be hollowed, we select the thin component with the highest value of stress and traverse the graph of thin and thick components, adding all visited components to a marked list. We do not allow traversal over any other thin components that are subject to an excessive level of stress. If some marked component reaches a fixed part of the tested load (floor surface or hold positions), we remove all connected marked components from the list. The remaining marked thick components are then hollowed (see Fig. 10). The amount of hollowing is defined on every medial point v_m of the marked thick components by their hollowing thickness $t_{hollow}(v_m)$. The thickness is determined by the amount of volume that we want to remove from the thick components. Since in general the stress scales linearly with the amount of applied force, we can compute the hollowing thickness as

$$t_{hollow}(v_m) = \min \left(t(v_m) - t_{min}, t(v_m) \sqrt[3]{1 - \frac{\sigma_c}{\sigma_m}} \right), \quad (16)$$

where $t(v_m)$ is the current thickness on the given medial point, t_{min} is some minimal allowed thickness of the printed material, and σ_m is the maximal stress in the object.

The hollowing itself follows a modified approach that was used for mesh displacement as explained in Section 4.2. Similarly, we compute a distance field using Laplace's equation, but with different boundary conditions. For hollowing, the interior medial points that belong to the marked thick components bound the Laplace's equation by their thickness, and all mesh vertices define a zero thickness

boundary constraint. Once the distance field is computed, we normalize the computed values based on t_{hollow} of the nearest medial point to the given grid point. Lastly, we construct a new mesh that represents the hole using a marching tetrahedron method for an iso-surface with a value equal to 1.




Figure 11: To create a hollowing correction, we compute a distance field on the object's interior using Laplace's equation bounded by the medial points of the hollowed component and the vertices of the mesh (a). Marching tetrahedrons creates a new triangular mesh that represents the hole (b), and the mesh is simplified and applied to the object (c.)

Note that most 3D printing technologies use an additional material to support the layers during printing of the object, and it is quite likely that there would be some support material inside the hollowed part. We need to add a small hole or holes into the structure so that the support material can be cleaned or melted. Such holes can be added manually before the support material is cleaned.

Lastly, hollowing can significantly affect the location of the center of mass of the printed object. If the hollowing makes the object unstable in its upright position, we automatically discard it and proceed with the next best model correction.

Hollowing Evaluation Hollowing has no visual impact on the appearance of the printed model. We evaluated hollowing as follows:

$$\delta_h^e = w_\sigma^h \frac{\sigma_c}{\sigma_m} \frac{V_t}{V_h}, \quad (17)$$

where V_h is the actual volume of removed material, and V_t is the target volume that we wanted to remove to reduce the stress below the critical stress σ_c . The target volume is directly proportional to the amount of stress that we want to reduce as $V_t = V_o \frac{\sigma_c}{\sigma_m}$. Since hollowing is a preferred solution to most gravity-related issues, we set its importance as $w_\sigma^h = 3$.

4.5 Shape Optimization

The above-mentioned model correction techniques are essentially heuristics that aim to reduce the amount of excessive stress under the critical value while maintaining the visual appearance of the input model. A general solution to this problem could be achieved by using so-called shape optimization techniques. The goal of these methods is to find the optimal geometry and topology of an input shape that minimizes certain objective function such as compliance [Allaire 2004] or the total stress in a shape [Allaire and Jouve 2008]. In our case, the goal would be to optimize the shape so that it minimizes the amount of excessive stress under the given load conditions while minimizing the visual difference from the original input model. As shown in [Allaire 2004], shape optimization problems can be solved iteratively using the gradient descent algorithm, where at each iteration the shape of the input model is deformed along the derivative of the objective function. The output of the algorithm is then an optimally corrected model that can be safely printed.

The main limitation of the shape optimization methods is their computational cost. To accurately evaluate the gradient of the objective function, we need to solve the linear elasticity problem whenever the shape of the model is changed. To do so we need to recompute the tetrahedral mesh, build a new stiffness matrix, and solve the resulting system. All of these steps are computationally expensive and would need to be performed during every iteration of the optimization process. This would make the method too slow for even moderately complex input models. Furthermore, shape optimization methods can lead to topology changes that are difficult to handle using our mesh-based approach. Allaire et al. [2004] proposed a new shape optimization approach that relies on level-set deformations of the initial shape [Osher and Sethian 1988] that automatically handles changing topologies. This approach also provides a performance benefit because the simulation mesh for linear elasticity remains constant, as defined by the underlying Eulerian grid. Unfortunately, to accurately represent the detail that can be produced by the modern 3D printers, a very fine resolution grid with millions of elements would need to be used. Our approach is a compromise between the speed and the output quality. We leave the research of a feasible shape optimization for printability corrections for future work. However, if a better or different solution is developed, it can be readily added into our pipeline.

5 Results

We evaluated our methods practically by printing our objects using two 3D printers: the Objet™ Alaris30 desktop printer and the Objet™ Eden350V professional printer. Both printers use the same printing pattern: the object is built by assembling thin layers of material (thickness of one layer can be from 16 to 28 microns). Average printing time was 18 hours with three objects, and tray size was 300x200x150mm.

The building process uses two kinds of material: primary and support. As a primary material, we used VeroWhite™ rigid opaque material and FullCure 720™ transparent one. Both materials are based on liquid polymers that solidify when exposed to UV light. After solidification they have plastic-like structure with tensile strength equal to 50MPa. Support material has a gel-like structure and helps to support the incomplete object during printing. After printing, the support material is washed away using a high pressure (up to 12MPa) water jet nozzle. The cleaning process proved to be dangerous to printed models with thin parts because the high pressure of the water jet can easily break them, as happened, for example, in our banana model (Fig. 12), whose left arm broke during cleaning.

When we used our system to improve our input models, all three methods were used: struts, thickening and hollowing. In all cases, structural strength of the models improved significantly, and printed models were more resistant to permanent and imposed loads.

Figure 12 shows an object that is unstable after printing and can be easily broken because of very thin legs. Structural analysis showed too much stress on the legs, and our system suggested adding a strut at the back of the model. Subsequent analysis showed that the strut was sufficient to alleviate the stress. We were able to confirm this by printing a new version of the model, which proved to be stable.

Another example of adding struts is in Fig. 13. Although the model of a soccer cup is able to stand still, it deforms easily when grabbed by the supports that hold the ball. To relieve the stress, the system iteratively added five interconnecting struts that created bridges between individual supports. Note that the struts are introduced in areas where they have the least visual impact. Figure 14 shows a model created using the Spore™ video game that has a weak spot on the neck. This causes the head of the model to break off when the object is grabbed by the head and tail. Our system detected this




Figure 12: A banana model printed without optimizations is quite likely to break because of very thin legs. An optimized model with an added strut shows much more structural strength.

issue and suggested the structure be locally thickened along this weak spot. The thicker neck was then able to withstand the pressure caused by the grip, the correction did not leave any significant impact on the visual appearance of the model.

The example in Fig. 15 illustrates hollowing. A hollowed model of a bunny consumed significantly less material than the full version (56%) and was also lighter. To maintain structural strength, our system also added six inner struts, as can be seen in Fig. 15 left. Hollowing has no additional effect on visual appearance.

In conclusion, to correct a model of a dragon (Fig. 1), our system used all three types of corrections. First, the system detected that the neck cannot withstand the weight of the rest of the body when the model is held by the head. The first fix was to hollow the main body, but even though the stress decreased by approximately 25%, the remaining stress was still above the allowed limit. The stress on the neck was fully removed after thickening was applied. Then another problem emerged when the model was tested standing on a flat surface. The object is too front-heavy, causing twisting and bending deformation on the legs, and this resulted in dangerous stress levels. The problem was automatically solved by fixing the body of the dragon to the pedestal with a strut connected to the dragon's tail.

Table 1 lists the models we have used, the number of polygons representing them, the types of correction, the maximum and average stress of the model before and after correction, and the weight of used material, also before and after the correction. The reduction of the maximum stress indicates how well we fixed the problematic area, and the average stress indicates whether we did not increase the overall stress, for example, because of hollowing. We are less concerned with average stress, since we focus on making the weakest parts of the model stronger.

6 Limitations and Future Work

Our system for automatic detection and correction of structural issues in 3D printed objects has its limitations, but it could be extended in many directions. To provide fast solutions, we have used many simplifications. One area for future research is a more thor-

Model	Polys (before/after)	Corrections	Max. stress [MPa]	Avg. stress [KPa]	Material [g]
Banana (Fig. 12)	19,352 / 19,540	1s	30.334 / 5.872 (81%)	3.195 / 0.538 (83%)	74 / 75 (1.4%)
Cup (Fig 13)	17,062 / 26,114	5s	22.967 / 7.590 (66%)	22.543 / 8.217 (64%)	161 / 161 (0%)
Bunny (Fig 15)	83644 / 88530	1h,6s	4.526 / 2.328 (49%)	9.205 / 7.233 (21%)	455 / 198 (56%)
Shell (Fig 10)	12,660 / 15,508	2s	5.541 / 1.884 (64%)	233.786 / 66.316 (72%)	46 / 49 (6.5%)
Molecule (see video)	3,816 / 3,816	5t	11.1 / 4.083 (63%)	18.035 / 14.166 (21%)	137 / 147 (7.3%)
Spore TM (Fig 10)	11,180 / 11,082	1t	17.817 / 4.881 (73%)	4.985 / 4.543 (9%)	103 / 104 (1%)
Dragon (Fig 1)	20572 / 9984	1t,1h,1s	29.092 / 6.048 (79%)	6.139 / 5.283 (14%)	124 / 105 (15%)

Table 1: Algorithm performance with respect to stress relief and material use. The applied number and types of corrections for each model is provided with *s*, *t*, and *h* to denote strut, thickening, and hollowing, respectively. The other values are formatted as before / after (change%). Note that when measuring the stress on the bunny, we compared the stress in the hollowed bunny with the stress in the reinforced one.

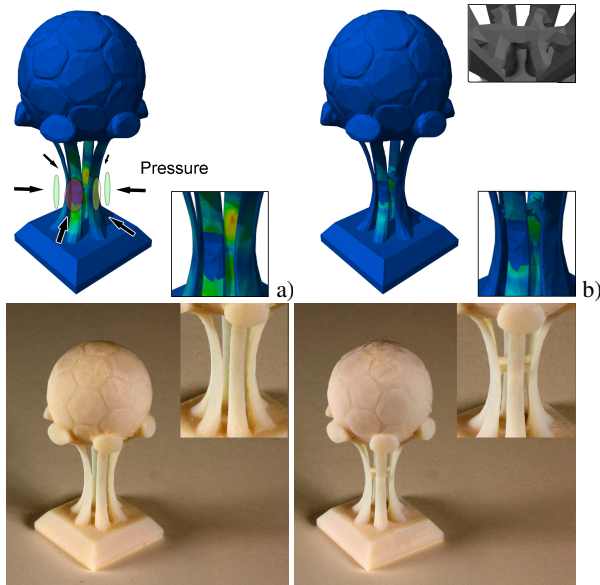


Figure 13: The handle of the soccer cup can break under pressure of the grip (a). Our method automatically connects all parts of the handle together by means of five different struts (b).

ough evaluation of grip positions on the object that can vary for different objects and users, among other things. Another area of improvement is structural analysis. The structural analysis presented in this paper works accurately only when the used material is homogenous and rigid. In reality, many of the used materials are plastic with nonhomogeneous properties caused by the layering of the material during printing. New printers can also print multiple materials at the same time, which is not directly supported by our method. Such situations require more advanced structural analysis as well as model specification. Improving the accuracy and precision with a better physical simulation can enhance our results by providing greater predictability.

Next, although we try to limit the visual impact of the model correction methods, we note that some of the corrections, especially struts, might alter the visual appearance of the printed object more than desired. The visual impact of struts can be reduced if their geometry and shape follow the visual style of the printed object. A detailed user study should be done to determine which structural changes are acceptable for which class of objects. We can speculate that adding struts to objects representing living objects will be less tolerated than doing the same for art objects.

Although our method generates good results for most tested cases, some scenarios remain that are not captured efficiently. First of all, we do not detect structural problems that might arise when the




Figure 14: The thin neck of a SporeTM creature can snap easily if the printed model is held by its head and tail. Our local thickening correction strengthens the neck without a significant effect on the visual appearance of the model.

model is cleaned after printing. If the cleaning is done by melting the support material away in hot water, for example, there is no additional stress. However, some cleaning involves the use of a water stream that may add more stress to some parts. Evaluation of such stresses should be further investigated.

Lastly, our current system is limited to the analysis and correction of *static* models. Many CAD-type models may have movable parts such as hinges, joints, wheels, bearings, and cogs. It is important not only to evaluate stress in the various configurations, but also to ensure that any corrections do not constrain the intended motion of the model. Especially important is taking into account a model's de-




Figure 15: Example of hollowing. A hole is created in the object to save material and to lighten the whole object. Inner struts are added to maintain structural strength.

sired function as part of a printability analysis. It might sometimes be possible to deduce an object's function from its form without need of additional user intervention.

7 Conclusions

We have detailed an automatic system for both detecting and correcting 3D printability issues in static models. To our knowledge, our work is the first to combine both physically based analysis with automatic geometric corrections in a single unified system for improving printability. We believe that our system aids both novices and experts in resolving issues prior to costly printing (both in time and money). Such technology is increasingly important as 3D printers become available for consumers. Although our work provides a strong foundation, we identified several challenging future research directions in the area of 3D printability.

Acknowledgements

We would like to thank Purdue CGVLAB for using their 3D printer.

References

- ALLAIRE, G., AND JOUVE, F. 2008. Minimum stress optimal design with the level set method. *Engineering Analysis with Boundary Elements* 32, 11, 909–918.
- ALLAIRE. 2002. A level-set method for shape optimization. *Comptes Rendus Mathematique* 334, 12, 1125–1130.
- ALLAIRE, G. 2004. Structural optimization using sensitivity analysis and a level-set method. *Journal of Computational Physics* 194, 1, 363–393.
- ANDREWS, J., JOSHI, P., AND CARR, N. 2011. A linear variational system for modelling from curves. *Comput. Graph. Forum*.
- BALASUBRAMANIAN, R., XU, L., BROOK, P. D., SMITH, J. R., AND MATSUOKA, Y. 2010. Human-guided grasp measures improve grasp robustness on physical robot. In *ICRA*, 2294–2301.
- BUTCHER, J. C. 2008. *Numerical Methods for Ordinary Differential Equations*. John Wiley and Sons, Ltd.
- COHEN, J., VARSHNEY, A., MANOCHA, D., TURK, G., WEBER, H., AGARWAL, P., BROOKS, F., AND WRIGHT, W. 1996. Simplification envelopes. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, SIGGRAPH '96, 119–128.
- DEY, T. K., AND ZHAO, W. 2003. Approximating the medial axis from the voronoi diagram with a convergence guarantee. *Algorithmica* 38 (October), 179–200.
- ENGINEERING HANDBOOK, 2011. Rapid Prototyping – Processes. <http://engineershandbook.com/RapidPrototyping/rpprocesses.htm>.
- FU, H., COHEN-OR, D., DROR, G., AND SHEFFER, A. 2008. Upright orientation of man-made objects. *ACM Trans. Graph.* 27 (August), 42:1–42:7.
- GOLUB, G. H., AND WELSCH, J. H. 1967. Calculation of gauss quadrature rules. Tech. rep., Stanford, CA, USA.
- HART, J. C., BAKER, B., AND MICHAELRAJ, J. 2003. Structural simulation of tree growth and response. *The Visual Computer* 19, 2-3, 151–163.
- HASLINGER, J., AND MÄKINEN, R. A. E. 2003. *Introduction to Shape Optimization: Theory, Approximation, and Computation*. SIAM, Philadelphia.
- HORNUNG, A., AND KOBBELT, L. 2006. Robust reconstruction of watertight 3d models from non-uniformly sampled point clouds without normal information. In *Proc. of Symposium on Geometry Processing*, Eurographics Association, 41–50.
- HUGHES, T. J. R. 1987. *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*. Prentice-Hall.
- JIRASEK, C., PRUSINKIEWICZ, P., AND MOULIA, B. 2000. Integrating biomechanics into developmental plant models expressed using l-systems. In *Proc. Plant Biomechanics*, 615–624.
- KINOSHITA, H., BÄCKSTRÖM, L., FLANAGAN, J. R., AND JOHANSSON, R. S. 1997. Tangential torque effects on the control of grip forces when holding objects with a precision grip. *Journal of Neurophysiology* 78, 3, 1619–1630.
- LIU, R., BURSCHKA, D., AND HIRZINGER, G. 2007. On the way to water-tight mesh. In *Proc. of ISPRS*.
- LIU, L., CHAMBERS, E. W., LETSCHER, D., AND JU, T. 2010. A simple and robust thinning algorithm on cell complexes. *Comput. Graph. Forum* 29, 7, 2253–2260.
- OSHER, S., AND SETHIAN, J. A. 1988. Fronts propagating with curvature dependent speed. *Journal of Computational Physics* 79, 12–49.
- PENG, J., KRISTJANSSON, D., AND ZORIN, D. 2004. Interactive modeling of topologically complex geometric detail. In *ACM SIGGRAPH 2004 Papers*, ACM, New York, NY, USA, SIGGRAPH '04, 635–643.
- SCHMIDT, R., AND SINGH, K. 2010. meshmixer: an interface for rapid mesh composition. In *ACM SIGGRAPH 2010 Talks*, SIGGRAPH '10, ACM, 6:1–6:1.
- SHAPEWAYS, 2011. 3D printing in 4 simple steps. http://www.shapeways.com/about/how_does_it_work.
- SHAPEWAYS, 2011. Things to keep in mind when designing for 3D printing. <http://www.shapeways.com/tutorials/things-to-keep-in-mind>.
- SHEWCHUK, J. R. 2001. Delaunay refinement algorithms for triangular mesh generation. *Computational Geometry: Theory and Applications* 22, 1–3.
- SI, H., 2011. Tetgen: A quality tetrahedral mesh generator and a 3D delaunay triangulator.
- TELEA, A., AND JALBA, A. 2011. Voxel-based assessment of printability of 3d shapes. In *Proc. of Mathematical morphology and its applications to image and signal processing*, Springer-Verlag, Berlin, Heidelberg, ISMM, 393–404.
- VON MISES, R. 1913. Mechanik der festen körper im plastisch-deformablen zustand. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse* 1913, 582–592.
- WHITING, E., OCHSENDORF, J., AND DURAND, F. 2009. Procedural modeling of structurally-sound masonry buildings. *ACM Trans. Graph.* 28, 5, 112.
- Z CORPORATION, 2011. Architectural Design Guide - Printing 3D Architectural Models. <http://crl.ap.buffalo.edu/digitalworkshop/zcorp/pdfs/archdesguide.pdf>.