

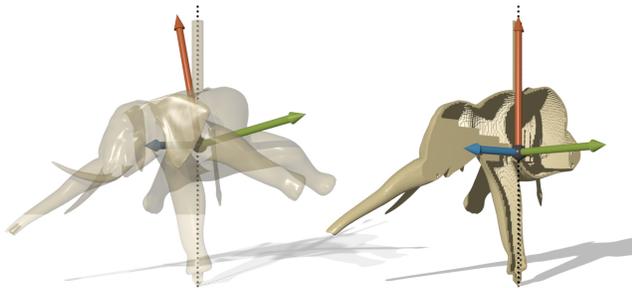
Spin-It: Optimizing Moment of Inertia for Spinnable Objects

Moritz Bächer
Disney Research Zurich

Emily Whiting
ETH Zurich

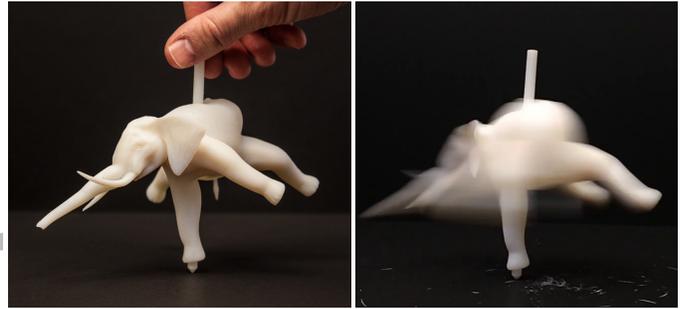
Bernd Bickel
Disney Research Zurich

Olga Sorkine-Hornung
ETH Zurich



(a) unstable input

(b) hollowed, optimized model



(c) our spinning top design

(d) elephant in motion

Figure 1: We introduce an algorithm for the design of spinning tops and yo-yos. Our method optimizes the inertia tensor of an input model by changing its mass distribution, allowing long and stable spins even for complex, asymmetric shapes.

Abstract

Spinning tops and yo-yos have long fascinated cultures around the world with their unexpected, graceful motions that seemingly elude gravity. We present an algorithm to generate designs for spinning objects by optimizing rotational dynamics properties. As input, the user provides a solid 3D model and a desired axis of rotation. Our approach then modifies the mass distribution such that the principal directions of the moment of inertia align with the target rotation frame. We augment the model by creating voids inside its volume, with interior fill represented by an adaptive multi-resolution voxelization. The discrete voxel fill values are optimized using a continuous, nonlinear formulation. Further, we optimize for rotational stability by maximizing the dominant principal moment. We extend our technique to incorporate deformation and multiple materials for cases where internal voids alone are insufficient. Our method is well-suited for a variety of 3D printed models, ranging from characters to abstract shapes. We demonstrate tops and yo-yos that spin surprisingly stably despite their asymmetric appearance.

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

Keywords: fabrication, moment of inertia, shape optimization

Links: [DL](#) [PDF](#)

1 Introduction

Spinning toys have existed since antiquity as playful objects that capture the imagination. Invented independently all over the world, spinning tops are referenced in ancient Greek literature [Gould

1975], and evidence of clay tops has been found in ancient cities dating as early as 3500 B.C. Similarly, while yo-yos are believed to have been invented in China, there are many historical references, including in Mozart’s *The Marriage of Figaro* where a yo-yo is spun to relieve stress [Malko 1978]. Despite the long tradition of these toys, until today creating new designs has been a trial-and-error process, calling on the intuition and meticulous patience of artists and hobbyists. Moreover, there has been little departure from rotationally symmetric designs.

Much attention has been devoted in the field of classical mechanics to *explaining* the motion of spinning objects, however, the focus has been primarily on analysis [Crabtree 1909; Perry 1957; Provatidis 2012; Cross 2013] rather than design. In this paper, we investigate the unique geometric properties of shapes that spin, with an eye on digital modeling and free-form design. A stable spin has requirements on rotational inertia, including precise positioning of the center of mass and correct alignment of the primary axes of the body. We propose an algorithm to optimize for these inertial properties, for example to design a spinning top that rotates smoothly and stably and can be fabricated using 3D printing.

In our approach, users provide an initial design for a spinning model, specified as a 3D surface mesh. Along with the input geometry, the user may specify the desired axis of spinning and the contact point with the support. The mass distribution is then optimized to ensure that the primary axis for the moment of inertia aligns with the desired axis of rotation. Since the moment of inertia depends on the entire volume, rather than just on the surface geometry, we preserve the appearance of the input design by prioritizing changes to the internal mass distribution. The algorithm may also deform the model to ensure correct rotational properties in cases where internal modifications are not sufficient. Alternatively, we can optimize dual-material models that compensate unfavorable configurations with a higher density fill, avoiding changes to the external shape.

Overall, we present a novel technique to design objects with special dynamic properties, and make the following contributions: We formulate a nonlinear functional that measures the *spinnability* of a solid shape about a user-defined axis. Using this measure, we devise constraint optimization problems that align the principal axes for moment of inertia with user-specified rotation axes. To this end, we maximize the ratio of principal moments in the dominant and lateral directions and place the center of mass on the rotation axis.

For our tops, we further improve stability by lowering the center of mass, simultaneously reducing the mass.

Our approach is effective on a wide range of models, from characters to abstract geometric forms. We employ an adaptive octree for discretizing the fill volume of our input shapes and a cage-based scheme to parameterize their deformations. We validate our results by fabricating the optimized shapes; as can be seen in the accompanying video, the objects can be stably spun despite their complex, asymmetric exterior appearance.

2 Related work

Fabrication-oriented design has recently gained increasing interest from the computer graphics community, triggered by advances in 3D manufacturing technology. Various physical properties have been explored in this shape modeling context, such as reproducible material appearance (see [Hullin et al. 2013] for a survey), deformation properties [Bickel et al. 2010; Skouras et al. 2013], articulation behavior [Bächer et al. 2012; Cali et al. 2012; Zhou et al. 2014], and kinematic structures [Zhu et al. 2012; Coros et al. 2013; Ceylan et al. 2013]. Complementary to these works, we focus on rigid models and dynamic properties resulting from mass distribution.

Prévost et al. [2013] proposed an approach for balancing static models at rest, that applies a combination of voxel carving and deformation to control the center of mass. Our work addresses a more general problem of stability under rotational motion, involving both center of mass and moment of inertia. While Prévost et al. [2013] use a plane sweeping heuristic for carving, we solve our constrained combinatorial problems by recasting them as sequential linear-quadratic programs using relaxation on the fill variables. As discussed later, our optimization can be used for static balancing, tending to find more stable solutions.

To assess printability of digital models, structural strength analysis under various loading conditions has been proposed, e.g. [Umetani and Schmidt 2013; Zhou et al. 2013]. Given a model with structural weaknesses, Stava et al. [2012] and Wang et al. [2013] create support structures and modify surface thicknesses to improve robustness. We select surface thickness based on printer resolution, but further robustness measures could also be integrated.

Rotational dynamics design. Furuta et al. [2010] combine a geometric modeling interface and a rigid body simulator for the design of kinetic art, providing real-time previews of the resulting motion during the design process. While restricted to forward simulations, this approach allows the user to quickly explore many trial-and-error experiments. We avoid trial-and-error and simulation, directly estimating models from user-specified geometries. Hirose et al. [2011] enforce symmetries along with additional geometric constraints to create sphericons. In contrast, we do not require a feasible starting solution and do not incorporate constraints other than the ones prescribed by the desired physical properties, enabling free-form design. To the best of our knowledge, we are the first to study the computational design of spinning toys with asymmetric appearance.

Moment of inertia is a fundamental property of rigid bodies. It specifies the required torque needed for a change in angular velocity and is, for example, an essential component in physics-based animation for rigid body simulations [Eberly 2003] or dynamics and control of characters [Macchietto et al. 2009]. Design for moment of inertia has been investigated in mechanical engineering, for example, in reducing inertial resistance of car rims [König and Wintermantel 2004]. However, the methods and objectives used differ significantly: the domain is restricted to regular, low-

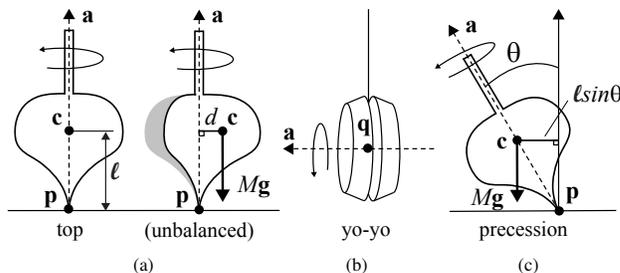


Figure 2: Spinning Yo-yos and Tops stably: For spinning tops, the center of mass must lie on the user-specified spinning axis \mathbf{a} , otherwise it will cause an unbalanced external torque $|\tau| = Mg d$ relative to \mathbf{p} (a). For slower angular velocities, the precession angle θ between rotational (vertical axis) and spinning axes becomes larger (c). For smaller ℓ , the gravitational torque $|\tau| = Mg \ell \sin \theta$ is smaller for equal precession, resulting in a longer spin. For yo-yos, we require the center of mass to coincide with \mathbf{q} (b).

resolution grids in combination with an approximate inertia formulation [Kang et al. 2009]; evolutionary algorithms are employed as an optimization strategy [Proos et al. 2001; König and Wintermantel 2004]. In contrast, our approach is generalizable to free-form shapes and we formulate an exact energy and derivatives.

Topology optimization methods solve engineering problems of distributing a limited amount of material in a design space [Bendsøe and Sigmund 2012]. While our adaptive voxel discretization shares similarities with the method proposed by DeRose and Díaz [1996], spinnability properties have not been considered by prior work in this field. Furthermore, we propose a joint hollowing and deformation optimization that allows deformation of the grid.

3 Fundamentals and overview

Given a 3D shape, we aim to generate spinnable models by altering their mass properties, while keeping the appearance as close to the original as possible. In the following sections we describe the user input, fundamental mass properties, and spinnability metrics needed to optimize the input toward a stably rotating object.

3.1 User input

The user provides the surface of a solid 3D shape, along with the desired spinning axis \mathbf{a} . The axis origin is set to the contact point \mathbf{p} as shown in Fig. 2 (a), which can be user-defined or chosen as the lowest point on the model w.r.t. the up-direction \mathbf{a} . For yo-yo designs, the shape is partitioned into two parts and connected with an axle that aligns with \mathbf{a} , to allow string coiling. The user selects a point \mathbf{q} on the axle to define the coiling location (Fig. 2 (b)).

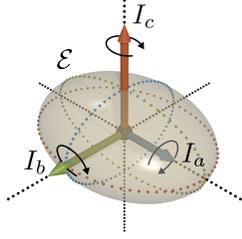
3.2 Mass properties and constraints

Center of mass. We denote by M the mass of our object and by \mathbf{c} the center of mass. If we assume a frictionless spin, the only external torque acting on a spinning top relative to \mathbf{p} , is the gravitational torque with magnitude $|\tau| = Mg d$, where g is Earth’s gravity and d is the distance from \mathbf{c} to the spinning axis (Fig. 2 (a)). We constrain the center of mass to lie on the spinning axis so that the net torque on the model around the ground contact point is zero.

Refer to Fig. 2 (c): during the spinning motion, the precession angle θ between the rotational (vertical) and spinning axes increases if the angular velocity ω becomes smaller. We can express the gravitational torque as $|\tau| = Mg \ell \sin \theta$, where ℓ is the height of the center of mass. Hence, we expect a longer, more stable spin for smaller values of ℓ and M .

For yo-yos, the gravitational torque remains zero throughout the spin if we neglect the effect of an uneven coiling of the string.

Moment of inertia is the analog of mass for rotational motion and measures the resistance to rotations about a given axis. Euler’s equations from classical mechanics (see, e.g., [Goldstein et al. 2001]) conveniently describe the rotating motion of a rigid body in its body frame, whose axes are the three principal axes of inertia and the origin is \mathbf{c} . Since there is no external torque acting on the body (for \mathbf{c} on the spinning axis), we can only spin about an axis with constant angular velocity if it is a principal axis of inertia.



For an arbitrary rigid body, there exists an equivalent ellipsoid with the same inertial properties. We can discuss the preferable axis using an ellipsoid \mathcal{E} with half-axes $\mathbf{h}_a, \mathbf{h}_b, \mathbf{h}_c$ ($\|\mathbf{h}_c\| \leq \|\mathbf{h}_b\| \leq \|\mathbf{h}_a\|$). Due to symmetry, \mathcal{E} ’s principal axes of inertia are parallel to its half-axes, and the corresponding moments I_a, I_b , and I_c each equal the sum of squares of the two other half-

axes’ lengths (omitting a common scale factor), as illustrated in the inset. Hence, the maximal principal axis of inertia corresponds to the shortest axis \mathbf{h}_c , and we have $I_c \geq I_b \geq I_a$. If we spin the ellipsoid \mathcal{E} with a constant angular velocity ω about a principal axis of inertia, the kinetic energy K in our system is $\frac{1}{2}I\omega^2$, $I \in \{I_a, I_b, I_c\}$. Since K is proportional to I , we can expect a longer spin for $I = I_c$.

Rotational stability refers to a body’s behavior under small disturbances to its angular velocity ω due to, e.g., frictional forces. Given three distinct values for the principal moments of inertia, $I_c > I_b > I_a$, rotation is stable under small perturbations only about the largest and the smallest axis [Goldstein et al. 2001]. In the case of two axes having identical principal moments, the rotation is stable only around the distinct axis. For $I_c = I_b = I_a$, no axis is stable, neglecting contact friction. We can observe this effect when trying to spin a marble in place: the orientation of the body changes over time.

For an asymmetric shape whose maximal principal axis of inertia aligns with the spin (and gravitational) axis and whose moments are distinct $I_c > I_b > I_a$, the top spins stably under the condition [Lewis et al. 1992]:

$$\omega^2 > \frac{Mg\ell}{I_c - I_b}. \quad (1)$$

From this relation we can see that the stability limit depends on the height of the center of mass ℓ and the mass M itself: the lower the centroid and the smaller the mass, the less angular velocity ω is required for a stable spin, confirming our conclusion from the above discussion on precession. Similarly, we need a smaller ω the higher the absolute difference between the largest moment I_c and the mid-moment I_b .

In summary, in order to spin stably, four basic requirements on the mass distribution of the model must be met:

1. The center of mass \mathbf{c} must lie on axis \mathbf{a} for spinning tops, or coincide with the axle center \mathbf{q} for yo-yos.
2. The center of mass \mathbf{c} should be closer to contact point \mathbf{p} and the mass M minimal for our tops.
3. The axis \mathbf{a} should be parallel to the maximal principal axis of inertia.
4. The magnitude of the largest principal moment of inertia should dominate over lateral moments to ensure the stability of the spin.

3.3 Measuring spin quality

To distill the above analysis of spinning properties into a spin quality measure, we formulate energy functionals for our yo-yos and tops. Provided that the basic constraints from Section 3.2 are fulfilled, our functionals assign a spin quality score to a given model \mathcal{M} based on the stability criterion (1). Note that while Eq. (1) suggests that a comparison of the mid- and largest moments is sufficient for tops, we consider all moments in our quality measures because the ordering of mid- and smallest axes might flip during our dynamic balancing optimization (see Section 4).

Yo-yos. We measure the spin quality of a yo-yo by summing the squared ratios of the dominant to lateral principal moments of inertia:

$$f_{\text{yo-yo}} = \gamma_{\text{I}} \left[\left(\frac{I_a}{I_c} \right)^2 + \left(\frac{I_b}{I_c} \right)^2 \right], \quad (2)$$

assuming that I_c corresponds to the given spin axis and the center of mass \mathbf{c} equals the axle center \mathbf{q} . The function $f_{\text{yo-yo}}$ is our *yo-yo energy functional*; small values correspond to longer, more stable spins.

Tops. To measure the quality of a spinning top, we add a term that penalizes the distance ℓ between the center of mass \mathbf{c} (which is constrained to lie on the axis \mathbf{a}) and the contact point \mathbf{p} and minimizes the mass M , yielding the *top energy functional*:

$$f_{\text{top}} = \gamma_{\text{c}} (\ell M)^2 + f_{\text{yo-yo}}. \quad (3)$$

The two weights γ_{c} and γ_{I} allow calibrating the relative contributions of the center of mass, inertia and the regularization terms of the parameterizations that follow (Sections 5–7).

3.4 Optimizing tops and yo-yos

We turn models into spinnable objects by altering their mass properties while keeping the appearance as close to the original as possible. Our primary focus is redistributing mass by hollowing the interior with precisely shaped voids. This method has significant effect on the inertia tensor, with the added benefit of preserving appearance. However, due to non-negligible material on the object’s shell, voids cannot always accomplish a stable spin. Consequently, we introduce deformation and dual-material optimization as extensions to our approach.

Hollowing. We first optimize the shape’s mass distribution by introducing inner voids. We adopt a multi-resolution octree to discretize the interior volume of the object. To generate the voids, we optimize for voxel fill values using a continuous, nonlinear formulation as discussed in more detail in Section 5. We maximize stability through the energy functionals $f_{\text{yo-yo}}$ (2) and f_{top} (3), respectively.

Cage-based deformation. While hollowing is effective for many models, some special cases over-extend our stability requirements. In these instances, we further manipulate mass by introducing deformation optimization. We automatically extract a cage from our octree and deform both the surface and the interior voids. We are able to simultaneously optimize for hollowing, using a unified formulation described in Section 6.

Dual-density optimization. As an alternative to deformation, we modify our hollowing optimization to incorporate multiple densities in our material specifications. A heavier material is used in the interior to compensate for highly non-optimal geometry of the model’s shell. While this technique presents an additional fabrication effort, its benefit is avoiding changes to the exterior shape. We describe the method in Section 7.

4 Optimizing dynamic balance

To evaluate our two quality measures f_{yo-yo} and f_{top} on a model \mathcal{M} made of a homogeneous material, we need to express its mass properties M , \mathbf{c} , and the 3×3 symmetric inertia tensor \mathbf{I} . Assume that the surface \mathcal{M} encloses a region $\Omega \in \mathbb{R}^3$ that corresponds to a solid object with constant density ρ . We express the above quantities using the ten integrals of the monomials of degree ≤ 2 over Ω , collected in a 10-vector:

$$\mathbf{s}_\Omega(\rho) = [s_1, s_x, s_y, s_z, s_{xy}, s_{yz}, s_{xz}, s_{x^2}, s_{y^2}, s_{z^2}]^T, \quad (4)$$

where $s_t = \rho \int_\Omega t dV$, e.g., $s_{xy} = \rho \int_\Omega xy dV$.

We obtain the following expressions for the mass and center of mass:

$$M = s_1 \quad \text{and} \quad \mathbf{c} = \frac{1}{M} [s_x, s_y, s_z]^T,$$

and \mathcal{M} 's inertia tensor:

$$\mathbf{I} = \begin{bmatrix} s_{y^2} + s_{z^2} & -s_{xy} & -s_{xz} \\ -s_{xy} & s_{x^2} + s_{z^2} & -s_{yz} \\ -s_{xz} & -s_{yz} & s_{x^2} + s_{y^2} \end{bmatrix}.$$

Note that we can reduce the volume integrals in \mathbf{s}_Ω to surface integrals $\mathbf{s}_{\partial\Omega}$ using the Divergence Theorem, resulting in analytical expressions for a triangulated surface $\partial\Omega$; see supplemental material for the derivations and pseudocode.

Coordinate frame for yo-yos. As evident from the formulas above, \mathbf{c} and \mathbf{I} are expressed w.r.t. a coordinate frame. For our yo-yos, the most convenient frame has its origin at the user-provided spin point \mathbf{q} and one of the three axes, say z , points in the direction of the desired spin axis \mathbf{a} , as illustrated in the inset. For this choice of frame, the model can only be spun about \mathbf{a} if the center of mass components s_x , s_y , and s_z , and also the off-diagonal elements $-s_{xz}$, $-s_{yz}$ of \mathbf{I} equal zero. Otherwise, \mathbf{c} does not equal \mathbf{q} or the z -axis is not a principal axis of inertia of \mathcal{M} . Provided \mathcal{M} fulfills these constraints, $I_z = s_{x^2} + s_{y^2}$ takes on the role of I_c in our functional f_{yo-yo} . Because our evaluation is independent of the choice of the x and y axes, we parameterize the 2×2 upper block of \mathbf{I} by an orientation angle ϕ using a Givens rotation:

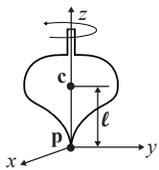
$$\mathbf{R} \begin{bmatrix} s_{y^2} + s_{z^2} & -s_{xy} \\ -s_{xy} & s_{x^2} + s_{z^2} \end{bmatrix} \mathbf{R}^T = \begin{bmatrix} I_x & 0 \\ 0 & I_y \end{bmatrix},$$

$$\mathbf{R} = \begin{bmatrix} \cos \phi & -\sin \phi \\ \sin \phi & \cos \phi \end{bmatrix}.$$

Optimizing yo-yos. To turn an arbitrary model \mathcal{M} into a yo-yo, we therefore need to minimize f_{yo-yo} with $I_a := I_x$, $I_b := I_y$, and $I_c := I_z$, with the constraints

$$s_t = 0, \quad \forall t \in \{x, y, z, xz, yz\} \quad \text{and} \quad (5)$$

$$\cos \phi \sin \phi (s_{x^2} - s_{y^2}) + (\cos^2 \phi - \sin^2 \phi) s_{xy} = 0. \quad (6)$$



Parallel Axis Theorem. The body frame centered at \mathbf{c} is not an ideal coordinate system for our tops because the center of mass can move freely along the axis \mathbf{a} . A better-suited frame is centered at the contact point \mathbf{p} , with the z -axis aligned with \mathbf{a} (see inset). Within this frame, the center of mass \mathbf{c} lies at height ℓ on the z -axis, so

that the inertia tensor \mathbf{I} is computed w.r.t. a frame shifted by ℓ w.r.t. our body frame. To evaluate f_{top} , we use the Parallel Axis Theorem, which states that if the axes of two frames are parallel, we can determine the new inertia tensor using the translation vector between the two origins and the body's mass:

$$\mathbf{I}_{CoM} = \mathbf{I} + M (\mathbf{c}\mathbf{c}^T - \mathbf{c}^T \mathbf{c} \mathbf{E}),$$

where \mathbf{E} is the identity matrix. For our choice of frame, where the center of mass is at $[0, 0, \ell]^T$, the theorem simplifies to

$$\mathbf{I}_{CoM} = \mathbf{I} - \frac{s_z^2}{s_1} \text{diag}(1, 1, 0).$$

Optimizing tops. For our tops, we minimize f_{top} , where I_a and I_b now refer to the rotated components of the shifted inertia tensor \mathbf{I}_{CoM} , and $I_c = s_{x^2} + s_{y^2}$ as before. Unlike the yo-yo case, \mathbf{c} can move freely on the z -axis. Hence, we relax the equality constraint $s_z = 0$, instead substituting $\ell M = s_z$ in the objective f_{top} (refer to Eq. (3)). The constraints to the optimization are then

$$s_t = 0, \quad \forall t \in \{x, y, xz, yz\} \quad \text{and} \quad (7)$$

$$\cos \phi \sin \phi (s_{x^2} - s_{y^2}) + (\cos^2 \phi - \sin^2 \phi) s_{xy} = 0. \quad (8)$$

Optimizing static balance. Interestingly, the problem of balancing a model at rest is a relaxed version of the top optimization:

$$\text{minimize } f_{\text{static}} = \gamma c \ell^2 \quad \text{subject to } s_x = s_y = 0$$

where we remove the mass term M in f_{static} because only the lowering of \mathbf{c} improves the balance at rest.

5 Hollowing

The most nonintrusive way to compensate for unfavorable mass distributions in a model is to introduce voids in the interior, as illustrated in Fig. 3. The idea of carving the shape's interior by sweeping a plane through a uniform voxel grid was explored in [Prévost et al. 2013] for static balancing. We propose a different optimization approach that addresses the inertia tensor in addition to the center of mass, uses a spatially-adaptive discretization, and avoids heuristics.

As explained in the previous sections, we aim to minimize $f_{top}(\Omega)$ subject to the constraints (7), (8), or $f_{yo-yo}(\Omega)$ subject to the constraints (5), (6). The variable in the optimization is the spatial mass distribution inside the shape, as detailed below. Recall that the functionals and the constraints are expressed in terms of the integrals \mathbf{s}_Ω ; we explain how these integrals depend on our unknowns.

5.1 Fabrication considerations

We enforce a minimal wall thickness to ensure that the resulting models can be fabricated. As shown in Fig. 3 left, we partition the region Ω into a boundary shell Ω_b and the interior Ω_i , restricting the hollowing to Ω_i . To account for a hollowed region $\Omega' \subseteq \Omega_i$ in our cost functionals, we adjust the volume integrals in Eq. (4):

$$\mathbf{s}_{\Omega - \Omega'} = \mathbf{s}_\Omega - \mathbf{s}_{\Omega'}.$$

Recall, given an axis, the contribution of a mass element to the moment of inertia is proportional to its squared distance from this axis. Mass on the boundary Ω_b has a high influence on the moment of inertia since it is far from the axis. Therefore, it is desirable for the wall to be as thin as possible within fabrication limits.

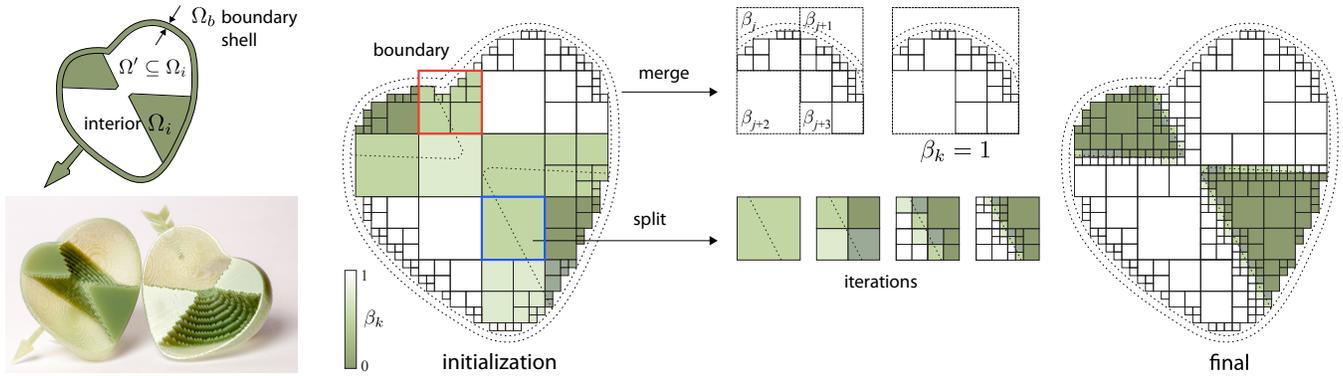


Figure 3: Hollowing: (Left) Our input encloses a volume Ω . By introducing voids Ω' , we can compensate for an unfavorable mass distribution. (Right) To reduce the number of variables and overall time complexity for our voids optimization, we summarize contributions of octree leaf cells in a partition of larger cells shown here for a boundary and an interior cell.

5.2 Voxelization

We discretize the interior Ω_i into mass elements Ω_k and optimize a binary fill variable $\beta_k \in \{0, 1\}$ for each, where a value of 1 means that we hollow that element, and 0 means we keep it filled. To handle free-form surfaces in our input and provide sufficient degrees of freedom for interior voids, we require our discretization to support fine enough mass elements. One possibility would be to use a high-resolution uniform voxel grid. However, we observe that finest-resolution voxels are only required at the surface separating the void space from the fill and external surface (see, e.g., the interior mass distribution of the Heart in Fig. 3 left, bottom). We therefore employ a multi-resolution voxelization based on an adaptive octree, thereby significantly reducing the number of fill variables. Our discretized volume integrals then become

$$s_{\Omega - \Omega'} = s_{\Omega} - \sum_k \beta_k s_{\Omega_k}$$

where $\Omega_i = \bigcup_k \Omega_k$ is a partitioning of the interior into octree cells Ω_k . The void space Ω' consists of all cells Ω_k for which $\beta_k = 1$.

5.3 Optimization approach

Given our adaptive voxel discretization, since the fill values are binary, the resulting minimization problem would be combinatorial. In order to take advantage of continuous optimization techniques, we propose a relaxation approach that allows β_k to take on a continuous value in the interval $[0, 1]$.

The goal of the optimization eventually is to assign binary fill values to each voxel. In practice, we observed that fill variables β_k with a fractional value only occur on the boundary between voids and solid regions. Hence, we sample these regions at a high resolution, ensuring final fractional values correspond to finest resolution cells only (compare with Fig. 3 final). Values are then rounded to binary numbers after convergence of the optimization.

This motivates the following optimization algorithm using adaptive refinement (refer to Fig. 3 right):

Initialization. We initialize the octree to a mid-level refinement (blue in Fig. 3) as a compromise between number of variables and resolution of the initial partitioning. For each cell, we compute s_{Ω_k} . For cells which overlap the boundary Ω_b (red), we only take the contribution from the volume in Ω_i into account.

Optimization step. We then optimize the fill variables β_k for all cells k as explained in detail below.

Split-and-merge. All cells k whose fill values are not binary ($\beta_k \in [\varepsilon, 1 - \varepsilon]$) after minimization, are split one level lower if they are not yet at the maximum resolution (see split branch). Conversely, cells with fill values within ε of 0 or 1 are candidates for merging. We merge neighboring cells with the same values into as coarse cells as possible (see merge branch). This gives us a new set of cells Ω_k for which we update s_{Ω_k} .

Convergence. After each optimization step, and split-and-merge, we check whether all fill values $\beta_k \in [\varepsilon, 1 - \varepsilon]$ correspond to cells Ω_k at the maximum resolution. If so, we terminate the optimization.

Our functionals f_{top} and $f_{y_0-y_0}$ are nonlinear in the fill variables β_k . To prevent an under-determined minimization problem, we use a uniform symmetric Laplacian \mathbf{L} as a regularizer, constructed over neighboring cells. This results in the following optimization:

$$\min_{\beta} f(\beta) + \gamma_{\mathbf{L}} \frac{1}{2} \beta^T \mathbf{L} \beta,$$

where β is a vector containing all β_k , and $f(\beta)$ refers either to $f_{\text{top}}(\beta)$ or $f_{y_0-y_0}(\beta)$.

5.4 Implementation details

Cells overlapping the boundary need special handling. We represent the content of such cells with a tree itself, rooted at the cell's level, refined to highest resolution in close proximity to the boundary of \mathcal{M} (red cell in Fig. 3). We illustrate cell splitting in the inset on the left. Cells corresponding to fill variables are marked in red, their subtrees in black.

Note that the example cell overlaps both internal void and the external model boundary.

To optimize the above regularized functionals, we use an active set algorithm with sequential linear-quadratic programming (SLQP) [Nocedal and Wright 2000; Byrd et al. 2006]. We restrict the fill values to the unit interval using box constraints and incorporate the center of mass and inertia as equality constraints. As the Hessian is dense, in our experiments we experienced better time performance when using LBFGS [Fletcher 1987], a memory-efficient approximation of the Hessian.

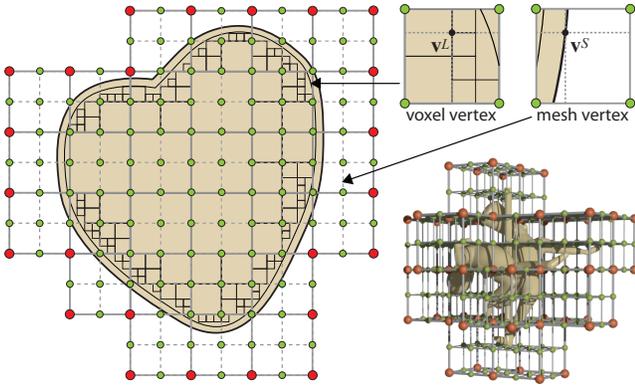


Figure 4: Octree-cage extraction: We first identify all octree cells (solid grey) intersecting with the objects’ volume. The exterior surface of these octree cells defines our cage and the cell corners on the surface are the handles (red). We then precompute bounded biharmonic weights [Jacobson et al. 2011] for all grid vertices (green) and tri-linearly interpolate them for mesh and voxel vertices.

6 Cage-based deformation

If our void optimization does not lead to a sufficiently stable spinning axis, we additionally allow the optimization to deform the surface and the interior cells, by using a cage-based scheme. We extract the cage from our octree by identifying all cells at a specific refinement level that intersect with the object’s volume as illustrated in Fig. 4. The m vertices on the outer boundary of these cells become the cage vertices, so-called *handles* (red). We then introduce a uniform grid (dotted and solid grey lines) at a finer octree level and precompute the bounded biharmonic weights [Jacobson et al. 2011] w_{ij} for each grid vertex \mathbf{v}_i (green and red) and handle j . For surface vertices of \mathcal{M} , and initial and newly created cell vertices (after splits), we tri-linearly interpolate the weights.

Similar to Prévost et al. [2013], we parametrize the handle transformations by uniform scales $u_j \in \mathbb{R}$ and translations $\mathbf{t}_j \in \mathbb{R}^3$, such that the overall deformation at a vertex \mathbf{v}_i has the form

$$\mathbf{v}'_i = \sum_{j=1}^m \begin{bmatrix} w_{ij} \mathbf{v}_i^T & w_{ij} \end{bmatrix} \begin{bmatrix} u_j \\ \mathbf{t}_j \end{bmatrix}.$$

In contrast to their work, however, the handles are automatically extracted without any 3D input from the user.

Collecting all original locations of the n triangle vertices \mathbf{v}_i^S and current leaf cell corners \mathbf{v}_i^L , with their tri-linearly interpolated weights in a $3n \times 4m$ weight matrix \mathbf{W} and, similarly, collecting the parameters u_j, \mathbf{t}_j of the m handles in a $4m$ column vector \mathbf{h} , we get the matrix version of the Linear Blend Skinning (LBS) deformation:

$$\mathbf{v}' = \mathbf{W}\mathbf{h}, \text{ where } \mathbf{v}' = \begin{bmatrix} \mathbf{v}'^S \\ \mathbf{v}'^L \end{bmatrix}, \mathbf{W} = \begin{bmatrix} \mathbf{W}^S \\ \mathbf{W}^L \end{bmatrix}.$$

The $3n$ -vector \mathbf{v}' collects all the deformed vertex and grid point positions. As individual points are typically only influenced by a partial set of handles, \mathbf{W} is relatively sparse.

Allowing deformation means that the integrals \mathbf{s} involved in our optimization formulation become functions of the handle transformations \mathbf{h} , in addition to the fill variables β :

$$s_D(\beta, \mathbf{h}) = s_\Omega(\mathbf{h}) - \sum_k \beta_k s_{\Omega_k}(\mathbf{h}).$$

6.1 Optimization

For the combined optimization of fill values and deformations, we add two additional regularizers that keep the deformation of the surface within reasonable limits. Since we only expect small rotations, we employ Sorkine et al. [2004]’s bi-Laplacian $\mathbf{L}_D^T \mathbf{L}_D$, which accounts for linearized rotations and favors smooth deformations. Additionally, we penalize deformations far from the original surface:

$$\gamma_{D,L} (\mathbf{v}'^S)^T \mathbf{L}_D^T \mathbf{L}_D \mathbf{v}'^S + \gamma_{D,v} \|\mathbf{v}'^S - \mathbf{v}^S\|_2^2.$$

As before, we optimize the resulting constraint minimization using active sets, with an analytical gradient, and an LBFGS approximation for the Hessian. The gradient with respect to handle transformations is straightforward if we assume the derivatives of our s -integrals with respect to vertex positions to be known (see our supplemental material for derivations and pseudocode),

$$\frac{\partial}{\partial \mathbf{h}} s_D(\beta, \mathbf{h}) = \left[\frac{\partial s_\Omega}{\partial \mathbf{v}'} - \sum_k \beta_k \frac{\partial s_{\Omega_k}}{\partial \mathbf{v}'} \right] \frac{\partial \mathbf{v}'}{\partial \mathbf{h}}$$

where $\partial \mathbf{v}' / \partial \mathbf{h} = \mathbf{W}$. Note that the cells’ volume integrals s_{Ω_k} with their respective gradients are all independent, allowing a massively parallel update after a deformation step.

We can easily replace our volume deformation with any other cage-based scheme and handle parametrizations by switching the LBS formulation and gradient $\partial \mathbf{v}' / \partial \mathbf{h}$.

7 Optimizing dual-material models

As previously motivated in Section 5.1, the masses of the elements in the boundary shell Ω_b generally have the most influence on a model’s inertial properties. So far we considered two balancing strategies: hollowing, and the combination of hollowing and deformation. While these strategies are well-suited for physical fabrication using single material 3D printers, an artist might not be satisfied with a deformed spinning toy. We can further increase our feasible set for the hollowing-only strategy by allowing for dual material optimizations. We use a material with high density ρ_i on the interior, where $\rho_i \gg \rho_b$. Given the same interior volume Ω_i , the heavier material makes it possible to achieve larger changes in moment of inertia and avoid the need for deformation. Our discretized volume integrals are then a function of the material densities:

$$s_{\Omega-\Omega'} = (s_{\Omega_b}(\rho_b) + s_{\Omega_i}(\rho_i)) - \sum_k \beta_k s_{\Omega_k}(\rho_i).$$

where the first term represents the solid model ($s_\Omega = s_{\Omega_b} + s_{\Omega_i}$) with different densities for boundary and interior. Fig. 5 shows an example fabricated result, where the model’s boundary and interior are made of low- and high-density materials, respectively.



Figure 5: Dual-material optimization: (Left) dual-material model, (right) 3D printed mold and resulting metal cast.

8 Results

Fabrication. All our models were printed on an Objet Connex 350 with an ABS-like plastic (green surface finish) and Objet’s “Vero White” material (white finish). The printer has a resolution of 600 and 1600 DPI on the two horizontal and vertical axes, respectively. The Connex 350 – like most other 3D printers – builds models layer-by-layer in a bottom-up manner, requiring a supporting structure for fabricating overhanging parts. Because we cannot remove any support from the interior without introducing holes in the models’ shells, we cut them prior to printing and glue them afterward.

Spinning tops. We validated our approach by designing and fabricating a variety of spinning tops, ranging from posed characters and abstract shapes to household objects. For the models presented in Figs. 6, 7, 8, and 9, we use an adaptive octree with a maximum refinement level of nine during the optimization. On a standard desktop computer with 3.2 GHz and 8 cores, the complete processing time for each takes less than a minute. This includes loading the input mesh, initializing the octree, performing hollowing optimization, and writing the output mesh. The hollowing optimization itself takes approximately 10 seconds.

In the figures below we illustrate the before-and-after body frames with black spheres for the center of mass, and red, green, and blue arrows for the maximal, mid-, and minimal principal axes of inertia (see, e.g., Fig. 7): the Ellipsoid in Fig. 6 demonstrates how we can turn asymmetric models, whose principal axes are far off the user-specified rotation axis, into dynamically balanced models that spin stably.



Figure 6: Asymmetric “Ellipsoid”: (Left) Unstable input design with misaligned principal axes. (Middle) Cross-section of optimized result after hollowing. The dominant principal axis (red) aligns with the spin axis. (Right) Fabricated result with hollowing.

Similar to the Ellipsoid, the input model for the Heart in Fig. 7 has poor mass properties, with the maximal principal axis extremely misaligned from the desired rotation axle (cupid’s arrow). Our optimization fixes the axis’ orientation and produces a very stable spin, as shown in the accompanying video.

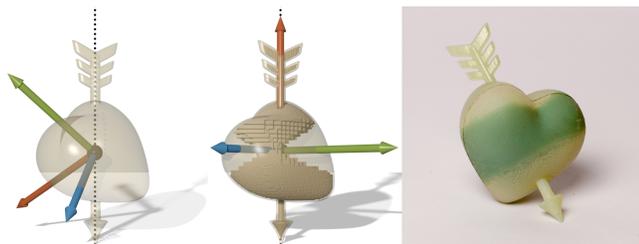


Figure 7: “Heart”: The initial design (left) has very poor alignment of the dominant principal axis with the spin axis. (Middle) Optimized result after hollowing, showing the interior mass distribution. Opaque surfaces indicate the boundary of the void space. (Right) 3D printed top.

Finally, two break-dancing Armadillos are shown in Fig. 8, one spinning on his back shell, one on the tip of his finger. Our hollowing successfully aligns the maximal principal axis of inertia with the user-specified one, even if it is far off as for the Armadillo spinning on his shell (compare left and right visualizations). Both Armadillos “dance” very stably around \mathbf{a} , as we demonstrate in our video.

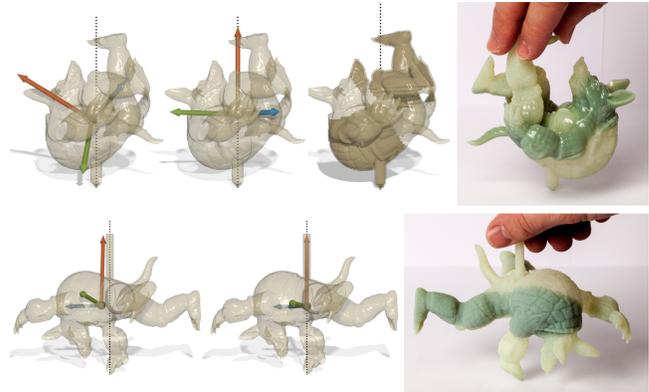


Figure 8: “Break-dancing Armadillos”: Through our hollowing optimization, the Armadillos can perform spinning dance moves. For each design, the unstable input (left), and the optimized stable output (right) are shown. The Armadillo on its shell is particularly badly aligned in the initial model.

Rotational stability. For the Teapot model (inset), the center of mass is reasonably close to the central spinning axis and the maximal principal axis of inertia is parallel to \mathbf{a} .



However, as we observe in the accompanying video, the solid model does not spin when actuated by hand. In accordance to the rotational stability criterion 1, a large angular velocity is required for a stable spin since the moments of inertia are similar. Our hollowing maximizes the ratio of I_c over lateral moments and allows us to reduce the angular velocity by a factor 1.56 (see Fig. 9 left, inertia only: $f_{\text{top}} = f_{y_0-y_0}$), while a simultaneous lowering of the center of mass allows for a reduction by a factor 1.60 as illustrated in Fig. 9 middle (lowering only: $f_{\text{top}} = \gamma_c \ell^2 + f_{y_0-y_0}$). We can achieve an even higher reduction of ω if we include mass M (see Fig. 9 right), resulting in a factor 1.68. Interestingly, the lowering only strategy shifts the mass distribution towards the contact point (compare left with middle cross-sections), while the simultaneous mass reduction lowers the center of mass less but reduces the mass inward out (compare middle with right cross-sections).

Yo-yos. We designed and fabricated two yo-yo examples. The Cuboid in Fig. 10 top is a case where the initial principal axes of the inertia tensor are far from the user specifications. Even with the highly non-optimal starting shape, the optimized output model spins stably. In our Woven Ring example (Fig. 10 bottom), the hollowing procedure successfully aligned the maximal principal axis despite complex surface geometry.

Cage-based deformation and dual-material models. There are physical limitations on how much the inertial properties of an object can be changed by hollowing if a minimum shell thickness has to be maintained. In this case, we propose two strategies. For the Dancing Couple model, we apply a cage-based deformation, as shown in Fig. 11. Refer to Fig. 4 lower, right for the cage.



Figure 9: “Teapot”: (Left) Hollowed result showing voxelized interior mass and aligned axes using $f_{top} = f_{yo-yo}$. (Middle) Lowering of the center of mass shifts the mass distribution closer to the contact point. If we include mass reduction (right), mass is reduced inward out, resulting in the design with highest rotational stability.

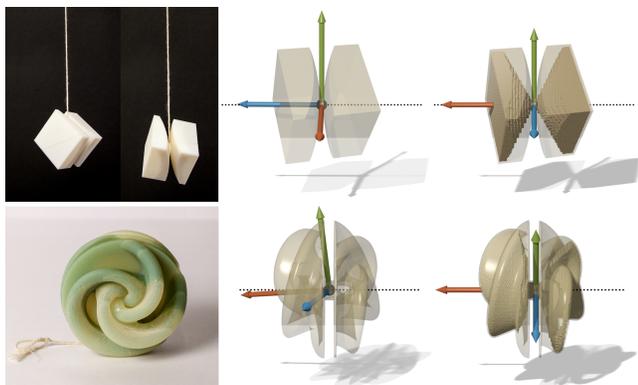


Figure 10: Yo-yo designs: (Left to right) 3D print; input model; optimized output model after hollowing. (Top) “Cuboid”: Our optimization rotates the original principal axes frame about the mid-magnitude axis. (Bottom) “Woven Ring”: The axis of dominant principal moment is precisely aligned to the spin direction.

Alternatively, we also handle multiple densities. The interior of the Sumone model (Fig. 5) consists of tin-solder material with significantly higher density ($\rho = 8.1 \text{ g/cm}^3$) compared to our printer material ($\rho = 1.17 \text{ g/cm}^3$). While dual material optimization extends the design space of spinning tops, it comes at the cost of a two-step fabrication process involving casting from a mold. Where lower densities are sufficient, 3D printing of multi-material content is possible [Chen et al. 2013; Vidimčič et al. 2013].

Static balancing is an inherent part of our optimization approach. In Fig. 12, we compare our balancing to the voxel-based sweep plane heuristic by Prévost and colleagues [Prévost et al. 2013]. For a fair comparison, we use voxel sizes that match our finest cells of a level 9 octree. In addition to static balancing, our method is capable of lowering the center of mass as we demonstrate in Fig. 12 top, left: while our center is 42% of the character’s height, Prévost et al.’s method places it at 56%. Furthermore, in contrast to Prévost et al., our method precisely places the center of mass at the center of the support polygon. This improves stable balance, as shown in the tilting plane test (Fig. 12 bottom). While our “T-Rex” keeps its balance up to a tilting angle of 8 degrees, Prévost et al.’s output already topples over at 1 degree.

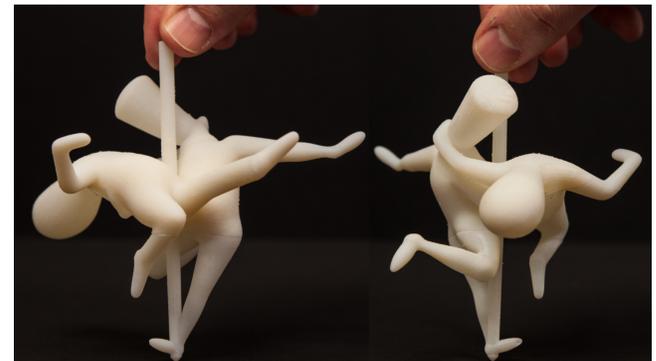
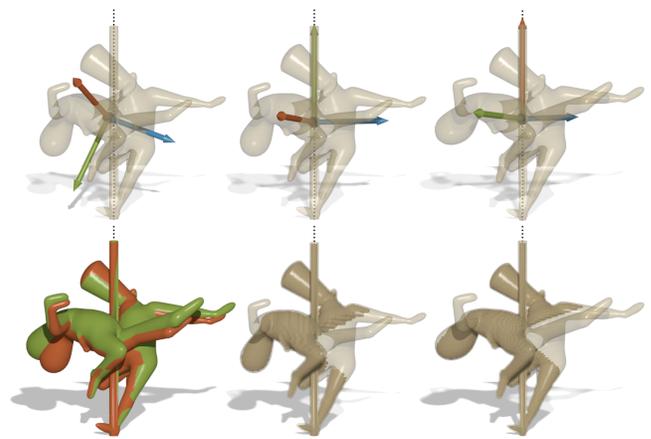


Figure 11: “Dancing Couple”. (Top: left to right) Initial design with principal axes rotated away from spin frame; after hollowing, the dominant primary axis is still not aligned; optimized result after deformation. (Middle: left to right) Initial (red) and deformed (green) models; voxelization after hollowing; voxelization with deformation optimization. (Bottom) The 3D printed result.

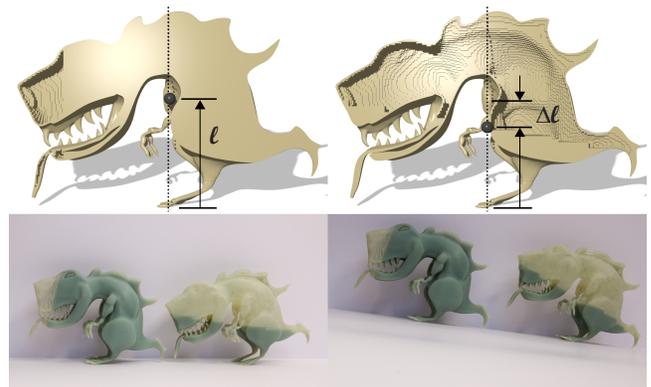


Figure 12: Statically balancing “T-Rex”: Our method handles static balancing. Compared to Prévost et al. (top-left), our hollowing result (top-right) has a lower center of mass, Δl . (Bottom) Inclined-plane stability test: the model by Prévost et al. loses balance significantly earlier (1°) than our optimized model (8°).

Cutting and voids. Due to the mathematical properties of moment of inertia, we can expect a small number of interior void spaces: among all our demo models (see Table 1), the Armadillo spinning on his shell had the largest number (5) of void spaces (see Fig. 8 left). However, merely 2 planar cuts were sufficient to access all voids. The highest number of planar cuts necessary (5)

for removing support material was for our Dancing Couple, even though there is only 1 void space. For powder-based printing, a single cut should be sufficient. We placed cuts manually, but could incorporate automated partitioning techniques in the future, for example, as an extension of Luo et al. [2012].

Model	# voids	# cuts
Cuboid, Ellipsoid, Heart, T-Rex	1	1
Armadillo (hand)	1	3
Woven Ring	1	4
Dancing Couple	1	5
Teapot	2	1
Elephant	3	2
Armadillo (shell)	5	2

Table 1: Cuts & voids: Models sorted by increasing # of voids.

Limitations. Our method is concerned with the concept of stability under perfect contact conditions with the support, and neglecting effects from air drag. However, frictional forces can lead to interesting phenomena on spinning objects. For example, the ‘tippe top’ is designed to flip vertically during its spin and relies on friction with the spin surface [Cohen 1977]. Similarly, a hard-boiled egg changes its spinning axis by 90 degrees. However, both examples have a particular shape, which likely imposes restrictions on the design space and limits free-form design. Further, simulation of air drag can also be significant for designs with complex surface geometry.

Our method is subject to practical limitations in scale. While larger models are easier to optimize, since minimum printable thickness is constant, models with high mass are difficult to spin by hand.

Lastly, to increase the value of our method as a design tool, it would be advantageous to integrate a selection of user controls. We would like to replace the fully automatic pipeline with user-guidance, affording exploration of the top and yo-yo design space.

9 Discussion

We have presented a technique to take arbitrary 3D models and create spinnable toys with surprising dynamic properties. Our algorithm generates novel designs for tops and yo-yos by optimizing rotational dynamics properties. We augment the model by creating voids inside its volume, with interior fill represented by an adaptive multi-resolution voxelization. The discrete voxel fill values are optimized using a continuous, nonlinear formulation. Further, we optimize for rotational stability by maximizing the dominant principal moment. We extend our technique to incorporate deformation and multiple materials with different densities. We have shown through theoretical results and physical validation that our method is well-suited for a range of geometries, from characters to abstract shapes.

Moment of inertia is a physical property fundamental to mechanical systems. As their computational design becomes increasingly popular, control over their inertial properties is an important feature. Our spinning toy application serves as empirical evidence that our energy terms are meaningful and intuitive. However, our energy formulation and solution strategies are generally applicable. Our work could inspire new inertia control techniques, e.g. in design of mechanical structures [Zhu et al. 2012; Coros et al. 2013; Ceylan et al. 2013], animatronics, and robotics: our method could be adopted to control inertial properties of individual parts, thereby

minimizing the system’s overall inertial resistance. This can allow for low-power actuators, reducing energy consumption and cost, or facilitate the design of passive dynamic systems. Further, while our models are passive objects with intricate shapes, Gajamohan et al. [2012] present a cube with actively controlled reaction wheels for balancing and even controlled jumping and falling. Embedding active control mechanisms is out of scope for this paper but might be an interesting avenue for future work.

Acknowledgments

We thank the anonymous reviewers for their helpful comments; Maurizio Nitti for model design; Romain Prévost for Make-It-Stand comparisons; Alexander Sorkine-Hornung, Kaan Yücer and Changil Kim for video and photo assistance; Ronnie Gänsli for metal casting; Alec Jacobson for the posed Elephant model, and skinning demo [Jacobson et al. 2012] for Armadillo poses; Romain Prévost and Amit Bermanto for print preparation. Model sources include: Woven Ring: generated by ‘Sculpture Generator 1’ by Carlo H. Séquin, UC Berkeley; Elephant: De Espona model library, courtesy of Robert Sumner; T-Rex: TurboSquid; Armadillo: Stanford Computer Graphics Laboratory; and Utah Teapot: Martin Newell, University of Utah. This project was supported in part by the ERC Starting Grant iModel (StG-2012-306877). Emily Whiting is supported by the ETH Zurich / Marie Curie COFUND Postdoctoral Fellowship.

References

- BÄCHER, M., BICKEL, B., JAMES, D. L., AND PFISTER, H. 2012. Fabricating articulated characters from skinned meshes. *ACM Trans. Graph.* 31, 4, 47:1–47:9.
- BENDSØE, M., AND SIGMUND, O. 2012. *Topology Optimization: Theory, Methods and Applications*. Engineering online library. Springer.
- BICKEL, B., BÄCHER, M., OTADUY, M. A., LEE, H. R., PFISTER, H., GROSS, M., AND MATUSIK, W. 2010. Design and fabrication of materials with desired deformation behavior. *ACM Trans. Graph.* 29, 4.
- BYRD, R. H., NOCEDAL, J., AND WALTZ, R. A. 2006. Knitro: An integrated package for nonlinear optimization. In *Large-scale nonlinear optimization*. Springer, 35–59.
- CALÌ, J., CALIAN, D. A., AMATI, C., KLEINBERGER, R., STEED, A., KAUTZ, J., AND WEYRICH, T. 2012. 3D-printing of non-assembly, articulated models. *ACM Trans. Graph.* 31, 6, 130:1–130:8.
- CEYLAN, D., LI, W., MITRA, N., AGRAWALA, M., AND PAULY, M. 2013. Designing and fabricating mechanical automata from mocap sequences. *ACM Trans. Graph.* 32, 6.
- CHEN, D., LEVIN, D. I. W., DIDYK, P., SITTHI-AMORN, P., AND MATUSIK, W. 2013. Spec2Fab: A reducer-tuner model for translating specifications to 3D prints. *ACM Trans. Graph.* 32, 4, 135:1–135:10.
- COHEN, R. J. 1977. The tippe top revisited. *American Journal of Physics* 45, 1, 12–17.
- COROS, S., THOMASZEWSKI, B., NORIS, G., SUEDA, S., FORBERG, M., SUMNER, R. W., MATUSIK, W., AND BICKEL, B. 2013. Computational design of mechanical characters. *ACM Trans. Graph.* 32, 4, 83:1–83:12.

- CRABTREE, H. 1909. *An Elementary Treatment of the Theory of Spinning Tops*. Longmans, Green and Co.
- CROSS, R. 2013. The rise and fall of spinning tops. *American Journal of Physics* 81, 280.
- DEROSE, G. C. J., AND DÍAZ, A. R. 1996. Hierarchical solution of large-scale three-dimensional topology optimization problems. In *ASME Design Engineering Technical Conferences and Computers in Engineering Conference*.
- EBERLY, D. H. 2003. *Game Physics*. Elsevier Science Inc.
- FLETCHER, R. 1987. *Practical Methods of Optimization; (2Nd Ed.)*. Wiley-Interscience, New York, NY, USA.
- FURUTA, Y., MITANI, J., IGARASHI, T., AND FUKUI, Y. 2010. Kinetic art design system comprising rigid body simulation. *Computer-Aided Design and Applications* 7, 4, 533–546.
- GAJAMOHAN, M., MERZ, M., THOMMEN, I., AND D’ANDREA, R. 2012. The Cubli: A cube that can jump up and balance. In *Proc. IROS, IEEE*, 3722–3727.
- GOLDSTEIN, H., POOLE, C., AND SAFKO, J. 2001. *Classical Mechanics*, 3rd ed. Addison Wesley.
- GOULD, D. 1975. *The Top: Universal Toy Enduring Pastime*. Bailey Brothers and Swinfen Ltd.
- HIROSE, M., MITANI, J., KANAMORI, Y., AND FUKUI, Y. 2011. An interactive design system for sphericon-based geometric toys using conical voxels. In *Proc. International Conference on Smart Graphics*, 37–47.
- HULLIN, M. B., IHRKE, I., HEIDRICH, W., WEYRICH, T., DAMBERG, G., AND FUCHS, M. 2013. Computational fabrication and display of material appearance. In *Eurographics State-of-the-Art Reports (STAR)*, 17 pages.
- JACOBSON, A., BARAN, I., POPOVIĆ, J., AND SORKINE, O. 2011. Bounded biharmonic weights for real-time deformation. *ACM Trans. Graph.* 30, 4, 78:1–78:8.
- JACOBSON, A., BARAN, I., KAVAN, L., POPOVIĆ, J., AND SORKINE, O. 2012. Fast automatic skinning transformations. *ACM Transactions on Graphics (proceedings of ACM SIGGRAPH)* 31, 4, 77:1–77:10.
- KANG, Z., WANG, X., AND WANG, R. 2009. Topology optimization of space vehicle structures considering attitude control effort. *Finite Elements in Analysis and Design* 45, 431–438.
- KÖNIG, O., AND WINTERMANTEL, M., 2004. CAD-based evolutionary design optimization with CATIA V5. 1st Weimar Optimization and Stochastic Days.
- LEWIS, D., RATIU, T., SIMO, J. C., AND MARSDEN, J. E. 1992. The heavy top: a geometric treatment. *Nonlinearity* 5, 1, 1.
- LUO, L., BARAN, I., RUSINKIEWICZ, S., AND MATUSIK, W. 2012. Chopper: partitioning models into 3D-printable parts. *ACM Trans. Graph.* 31, 6, 129:1–129:9.
- MACCHIETTO, A., ZORDAN, V., AND SHELTON, C. R. 2009. Momentum control for balance. *ACM Trans. Graph.* 28, 3.
- MALKO, G. 1978. *The One and Only Yo-Yo Book*. Avon.
- NOCEDAL, J., AND WRIGHT, S. J. 2000. *Numerical Optimization*. Springer.
- PERRY, J. 1957. *Spinning Tops and Gyroscopic Motion*. Dover Publications.
- PRÉVOST, R., WHITING, E., LEFEBVRE, S., AND SORKINE-HORNUNG, O. 2013. Make It Stand: Balancing shapes for 3D fabrication. *ACM Trans. Graph.* 32, 4, 81:1–81:10.
- PROOS, K., STEVEN, G., QUERIN, O., AND XIE, Y. 2001. Stiffness and inertia multicriteria evolutionary structural optimisation. *Engineering Computations* 18, 7, 1031 – 1054.
- PROVATIDIS, C. G. 2012. Revisiting the spinning top. *International Journal of Material and Mechanical Engineering (IJMME)* 1, 4, 71–88.
- SKOURAS, M., THOMASZEWSKI, B., COROS, S., BICKEL, B., AND GROSS, M. 2013. Computational design of actuated deformable characters. *ACM Trans. Graph.* 32, 4, 82:1–82:10.
- SORKINE, O., COHEN-OR, D., LIPMAN, Y., ALEXA, M., RÖSSL, C., AND SEIDEL, H.-P. 2004. Laplacian surface editing. In *Proc. Symposium on Geometry Processing*, 179–188.
- STAVA, O., VANEK, J., BENES, B., CARR, N., AND MĚCH, R. 2012. Stress relief: Improving structural strength of 3D printable objects. *ACM Trans. Graph.* 31, 4, 48:1–48:11.
- UMETANI, N., AND SCHMIDT, R. 2013. Cross-sectional structural analysis for 3D printing optimization. In *SIGGRAPH Asia 2013 Technical Briefs*, 5:1–5:4.
- VIDIMČE, K., WANG, S.-P., RAGAN-KELLEY, J., AND MATUSIK, W. 2013. OpenFab: A programmable pipeline for multi-material fabrication. *ACM Trans. Graph.* 32, 4, 136:1–136:12.
- WANG, W., WANG, T. Y., YANG, Z., LIU, L., TONG, X., TONG, W., DENG, J., CHEN, F., AND LIU, X. 2013. Cost-effective printing of 3D objects with skin-frame structures. *ACM Trans. Graph.* 32, 6, 177:1–177:10.
- ZHOU, Q., PANETTA, J., AND ZORIN, D. 2013. Worst-case structural analysis. *ACM Trans. Graph.* 32, 4, 137:1–137:12.
- ZHOU, Y., SUEDA, S., MATUSIK, W., AND SHAMIR, A. 2014. Boxelization: Folding 3d objects into boxes. *ACM Trans. Graph.* 33, 4 (Jul), (to appear).
- ZHU, L., XU, W., SNYDER, J., LIU, Y., WANG, G., AND GUO, B. 2012. Motion-guided mechanical toy modeling. *ACM Trans. Graph.* 31, 6, 127.