

Spin-It: Optimizing Moment of Inertia for Spinnable Objects

By Moritz Bächer, Bernd Bickel, Emily Whiting, and Olga Sorkine-Hornung

Abstract

Spinning tops and yo-yos have long fascinated cultures around the world with their unexpected, graceful motions that seemingly elude gravity. Yet, due to the exceeding difficulty of creating stably spinning objects of asymmetric shape in a manual trial-and-error process, there has been little departure from rotationally symmetric designs. With modern 3D printing technologies, however, we can manufacture shapes of almost unbounded complexity at the press of a button, shifting this design complexity toward computation.

In this article, we describe an algorithm to generate designs for spinning objects by optimizing their mass distribution: as input, the user provides a solid 3D model and a desired axis of rotation. Our approach then modifies the interior mass distribution such that the principal directions of the moment of inertia align with the target rotation frame. To create voids inside the model, we represent its volume with an adaptive multiresolution voxelization and optimize the discrete voxel fill values using a continuous, nonlinear formulation. We further optimize for rotational stability by maximizing the dominant principal moment. 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.

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,¹² 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cities dating as early as 3500 B.C. Similarly, while yo-yos are believed to have been invented in China, there are many historical references, including in Mozart's *The Marriage of Figaro* where a yo-yo is spun to relieve stress.¹⁷ Despite the long tradition of these toys, until today creating new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.

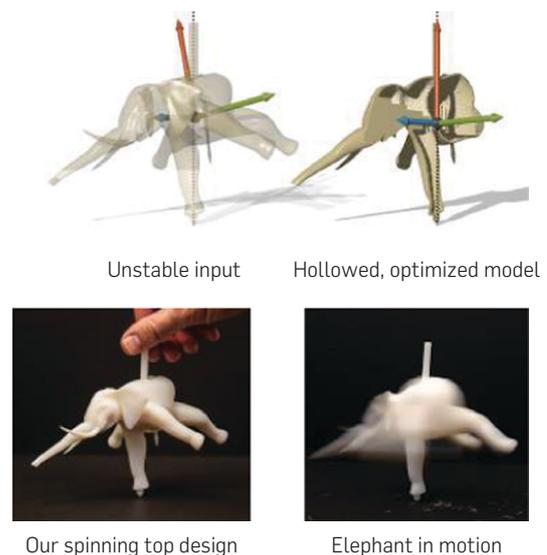
Much attention has been devoted in the field of classical mechanics to *explaining* the motion of spinning objects; however, the focus has been primarily on analysis^{8, 9, 19, 21} rather than design. In this article, we investigate the unique geometric properties of shapes that spin, with an eye on 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 changing the internal mass distribution as we illustrated in Figure 1 on an elephant top.

We first formulate a nonlinear functional that measures the *spinnability* of a solid shape about a user-defined axis. Using this measure, we then devise constrained 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.

Figure 1. We describe 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.



The original version of this paper was published in *Proceedings of SIGGRAPH'14*, August 2014, ACM.

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 validate our results by fabricating the optimized shapes; the objects can be stably spun despite their complex, asymmetric exterior appearance.

2. RELATED WORK

Fabrication-oriented design. Fabrication-oriented design has 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 including deformation properties,^{4, 22} articulation behavior,^{1, 5} structural strength,^{23, 24} and kinematic structures.^{6, 7, 25}

Most related to our effort is the work by Prévost et al.²⁰: they 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.²⁰ 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.

Rotational dynamics design. Furuta et al.¹⁰ 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.¹³ 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. 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 or dynamics and control of characters.¹⁶ Design for moment of inertia has been investigated in mechanical engineering. However, the methods and objectives used differ significantly. Our approach further generalizes to free-form shapes and we formulate an exact energy and derivatives.

Topology optimization. Topology optimization methods solve engineering problems of distributing a limited amount of material in a design space.³ While our adaptive voxel discretization shares similarities with solution techniques common in this field, spinnability properties have not been considered by prior work.

3. FUNDAMENTALS AND OVERVIEW

Given a 3D shape, we aim to generate spinnable models by altering their mass distribution, 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 Figure 2a, 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 (Figure 2b).

3.2. Mass properties and constraints

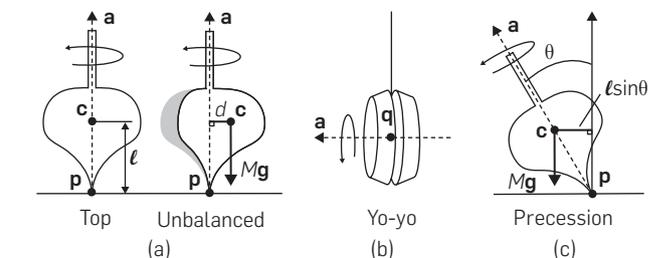
Center of mass. We denote by M the mass of our object and by \mathbf{c} the center of mass. If we assume a frictionless spin, the only external torque acting on a spinning top relative to \mathbf{p} , is the gravitational torque with magnitude $|\tau| = Mg d$, where g is Earth's gravity and d is the distance from \mathbf{c} to the spinning axis (Figure 2a). 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 Figure 2c: 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.

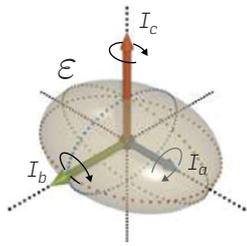
Moment of inertia. Moment of inertia is the analog of mass for rotational motion and measures the resistance to rotations about a given axis. Euler's equations from classical

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).



mechanics (see, e.g., Ref.¹¹) 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. Rotational stability refers to a body's behavior under small disturbances to its angular velocity ω due to, for example, 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.¹¹ 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. As long as the condition $I_c > I_b \geq I_a$ holds, we call I_c the dominant and I_b and I_a the lateral principal moments of inertia.

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 condition¹⁴:

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

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 Equation (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:

$$f_{\text{yo-yo}} = \gamma_1 \left[\left(\frac{I_a}{I_c} \right)^2 + \left(\frac{I_b}{I_c} \right)^2 \right], \quad (2)$$

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*:

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

The two weights γ_c and γ_1 allow calibrating the relative contributions of the center of mass, inertia and the regularization term of the parameterization that follows (see Section 5).

3.4. Optimizing tops and yo-yos

We turn models into spinnable objects by altering their mass distribution while keeping their appearance unchanged. To this end, we redistribute mass by hollowing the interior with precisely shaped voids. We adopt a multiresolution 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 we discuss in more detail in Section 5. We maximize stability through the energy functionals $f_{\text{yo-yo}}$ (2) and f_{top} (3), respectively.

While hollowing is effective for many models, some special cases over-extend our stability requirements and voids alone cannot accomplish a stable spin. This is due to the non-negligible material on the object's shell. In the original version of this article,² we introduce extensions to our approach, further manipulating mass by either deforming the surface and interior voids or compensating for highly nonoptimal mass distributions with a heavier material in the interior.

4. OPTIMIZING DYNAMIC BALANCE

Before we explain our multiresolution discretization of the interior mass distribution, we formalize our quality measures and requirements from the previous section, casting them as optimization problems in a discretization-independent manner. 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:

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

$$\text{where } s_i = \rho \int_\Omega t^i dV, \quad \text{e.g., } s_{xy} = \rho \int_\Omega xy dV.$$

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 accompanying the original version of this article Ref.²

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} . If I_x and I_y denote the eigenvalues of the 2×2 upper block

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

we recall that the trace of $\bar{\mathbf{I}}$ is the sum of the squared eigenvalues $I_x^2 + I_y^2$, leading to an elegant reformulation of our yo-yo functional f_{yo-yo}

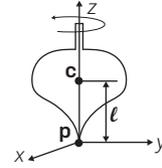
$$\gamma_1 \frac{\text{tr}(\bar{\mathbf{I}}^2)}{I_z^2} = \gamma_1 \frac{(s_{y^2} + s_{z^2})^2 + 2s_{xy}^2 + (s_{x^2} + s_{z^2})^2}{(s_{x^2} + s_{y^2})^2}, \quad (5)$$

where tr is the trace operator and I_x and I_y take on the roles of I_a and I_b , respectively.

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

$$s_i = 0, \quad \forall t \in \{x, y, z, xz, yz\}. \quad (6)$$

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 are now the eigenvalues of the upper 2×2 block of the 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 Equation 3).

The constraints to the optimization are then

$$s_i = 0, \quad \forall t \in \{x, y, xz, yz\}. \quad (7)$$

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

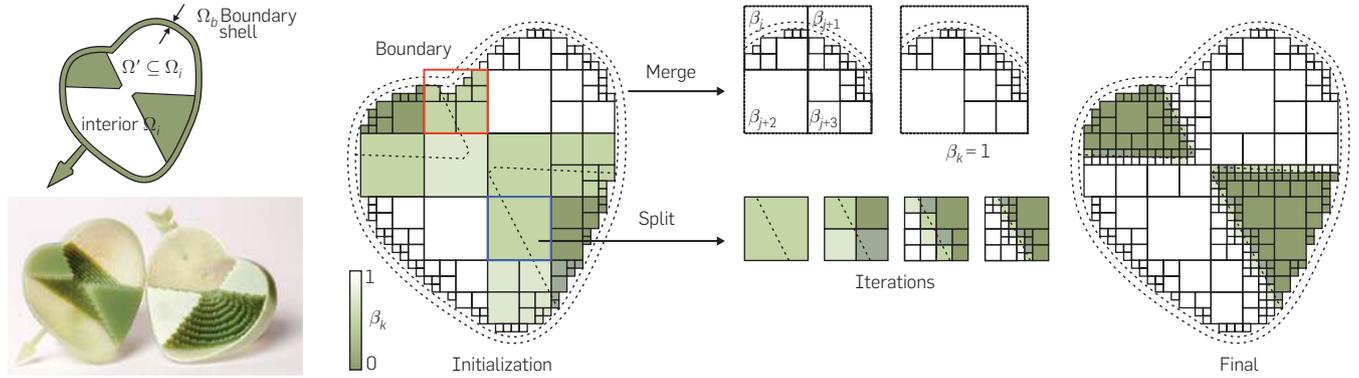
$$\text{minimize } f_{\text{static}} = \gamma_c \ell^2 \quad \text{subject to } s_x = s_y = 0,$$

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 Figure 3. The idea of carving the shape's interior by sweeping a plane through a uniform voxel grid was explored in Prévost et al.²⁰ 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.

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.



As explained in the previous sections, we aim to minimize $f_{\text{top}}(\Omega)$ subject to the constraints (7), or $f_{y_0-y_0}(\Omega)$ subject to the constraints (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 Figure 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 Equation (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.

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 Figure 3, left, bottom). We therefore employ a multiresolution voxelization based on an adaptive octree, thereby significantly reducing the number of fill variables. Our discretized volume integrals then become

$$\mathbf{s}_{\Omega-\Omega'} = \mathbf{s}_\Omega - \sum_k \beta_k \mathbf{s}_{\Omega_k},$$

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 Figure 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 Figure 3, right):

Initialization. We initialize the octree to a mid-level refinement (blue in Figure 3) as a compromise between number of variables and resolution of the initial partitioning. For each cell, we compute \mathbf{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 $f_{\text{yo-yo}}$ are nonlinear in the fill variables β_k . To prevent an underdetermined minimization problem, we penalize differences between fill variables using a uniform symmetric Laplacian L , constructed over neighboring cells. This results in the following regularized optimization problem:

$$\min_{\beta} f(\beta) + \gamma_L \frac{1}{2} \beta^T L \beta \quad \text{subject to } s_i = 0,$$

where β is a vector containing all β_k , and $f(\beta)$ refers either to $f_{\text{yo-yo}}(\beta)$ or $f_{\text{top}}(\beta)$, and s_i denote the respective linear equality constraints (6) or (7).

To optimize the above regularized functionals, we use an active set algorithm with sequential linear-quadratic programming.¹⁸ We further restrict the fill values to the unit interval using box constraints. As the Hessian is dense, in our experiments we experienced better time performance when using LBFGS,¹⁸ a memory-efficient approximation of the Hessian.

6. 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.

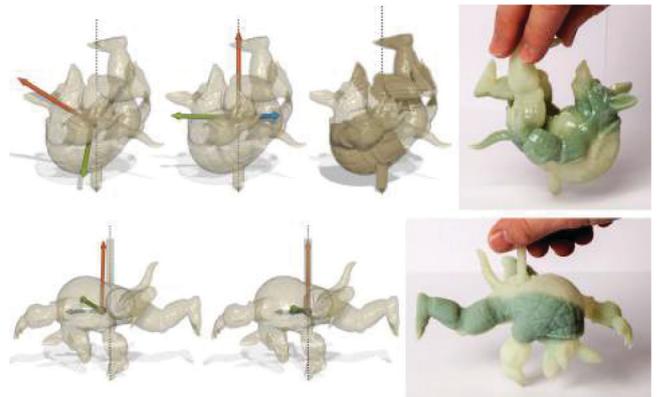
Spinning tops. We validated our approach by designing and fabricating a variety of spinning tops, ranging from posed characters and abstract shapes to household objects. For the models presented in Figures 1, 4, and 5, we use an adaptive octree with a maximum refinement level of nine during the optimization. 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 s.

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., Figure 4): the Ellipsoid in Figure 4 (top) 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 4. Asymmetric “Ellipsoid” and “Heart”: (Left) Unstable input designs with misaligned principal axes. (Middle) Optimized results after hollowing: for the “Ellipsoid”, a cross-section is shown. The dominant principal axis (red) aligns with the spin axis. Opaque surfaces indicate the boundary of the void space. (Right) Fabricated results with hollowing.



Figure 5. “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.



Similar to the Ellipsoid, the input model for the Heart in Figure 4 (bottom) has a poor mass distribution, leading to a principal axis far off the desired rotation axle (cupid’s arrow). Our optimization fixes the axis’ orientation and produces a very stable spin.

Finally, two break-dancing Armadillos are shown in Figure 5, 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} .

Rotational stability. For the Teapot model (inset), the center of mass is reasonably close to the central spinning

Figure 6. “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.



axis and the maximal principal axis of inertia is parallel to \mathbf{a} . However, 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

Figure 6, left, inertia only: $f_{top} = f_{yo-yo}$), while a simultaneous lowering of the center of mass allows for a reduction by a factor 1.60 as illustrated in Figure 6 (middle; lowering only: $f_{top} = \gamma_c \ell^2 + f_{yo-yo}$). We can achieve an even higher reduction of ω if we include mass M (see Figure 6, right), resulting in a factor 1.68. Interestingly, the lowering only strategy shifts the mass distribution toward 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 Figure 7 (top) is a case where the initial principal axes of the inertia tensor are far from the user specifications. Even with the highly nonoptimal starting shape, the optimized output model spins stably. In our Woven Ring example (Figure 7, bottom), the hollowing procedure successfully aligned the maximal principal axis despite complex surface geometry.

Static balancing. Static balancing is an inherent part of our optimization approach. In Figure 8, we compare our balancing to the voxel-based sweep plane heuristic by Prévost et al.²⁰ 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 Figure 8 (top-left): while our center is 42% of the character’s height, Prévost et al.’s method places it at 56%. Furthermore, in contrast

Figure 7. 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.

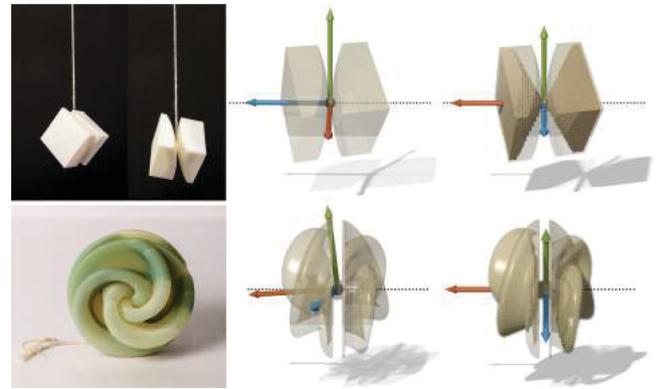
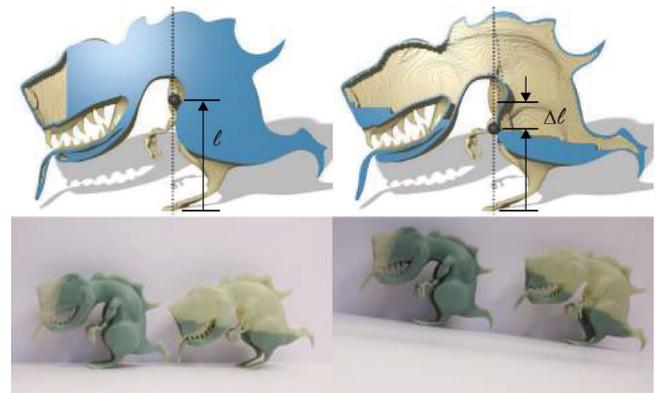


Figure 8. Statically balancing “T-Rex”: Compared to Prévost et al.²⁰ (top-left), our hollowing result (top-right) has a lower center of mass, $\Delta \ell$. Cross-sections are shown in blue. (Bottom) Inclined-plane stability test: the model by Prévost et al. loses balance significantly earlier (1°) than our optimized model (8°).



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 (Figure 8, bottom). While our “T-Rex” keeps its balance up to a tilting angle of 8°, Prévost et al.’s output already topples over at 1°.

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, the Armadillo spinning on his shell had the largest number (5) of void spaces (see Figure 5, left). However, merely two planar cuts were sufficient to access all voids. 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.¹⁵

Limitations. Our method is concerned with the concept of stability under perfect contact conditions with the

support, and neglecting effects from air drag. However, simulation of air drag can be significant for designs with complex surface geometry. Our method is further 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.

7. DISCUSSION

Spinning tops and yo-yos have existed since millennia and we have witnessed only very limited departure from symmetric designs. Utilizing the shift in design complexity from manufacturing toward computation, we have presented a technique to take arbitrary, asymmetric 3D models and turn them into stably spinning toys with previously unseen and surprising dynamic properties. While we have not considered friction in our modeling, 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. Similarly, a hard-boiled egg changes its spinning axis by 90°. However, both examples have a particular shape, which likely imposes restrictions on the design space and limits free-form design.

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 strategy are generally applicable. Our work could inspire new inertia control techniques, for example, in design of mechanical structures,^{6, 7, 25} 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.

Acknowledgments

First and foremost, we would like to thank our editor Steve Marschner for his invaluable feedback. We were fortunate to get further help from 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 and Armadillo models, and Romain Prévost and Amit Bermanno 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 was supported by the ETH Zurich/Marie Curie COFUND Postdoctoral Fellowship.

References

- Bächer, M., Bickel, B., James, D.L., Pfister, H. Fabricating articulated characters from skinned meshes. *ACM Trans. Graph.* 31, 4 (2012), 47:1–47:9.
- Bächer, M., Whiting, E., Bickel, B., Sorkine-Hornung, O. Spin-it: Optimizing moment of inertia for spinnable objects. *ACM Trans. Graph.* 33, 4 (July 2014), 96:1–96:10.
- Bendsøe, M., Sigmund, O. *Topology Optimization: Theory, Methods and Applications*. Engineering online library. Springer, Berlin, Germany, 2012.
- Bickel, B., Bächer, M., Otaduy, M.A., Lee, H.R., Pfister, H., Gross, M., Matusik, W. Design and fabrication of materials with desired deformation behavior. *ACM Trans. Graph.* 29, 4 (2010), 63:1–63:10.
- Calì, J., Calian, D.A., Amati, C., Kleinberger, R., Steed, A., Kautz, J., Weyrich, T. 3D-printing of non-assembly, articulated models. *ACM Trans. Graph.* 31, 6 (2012), 130:1–130:8.
- Ceylan, D., Li, W., Mitra, N., Agrawala, M., Pauly, M. Designing and fabricating mechanical automata from mocap sequences. *ACM Trans. Graph.* 32, 6 (2013), 186:1–186:11.
- Coros, S., Thomaszewski, B., Noris, G., Sueda, S., Forberg, M., Sumner, R.W., Matusik, W., Bickel, B. Computational design of mechanical characters. *ACM Trans. Graph.* 32, 4 (2013), 83:1–83:12.
- Crabtree, H. *An Elementary Treatment of the Theory of Spinning Tops*. Longmans, Green and Co., London, UK, 1909.
- Cross, R. The rise and fall of spinning tops. *Am. J. Phys.* 81 (2013) 81:280.
- Furuta, Y., Mitani, J., Igarashi, T., Fukui, Y. Kinetic art design system comprising rigid body simulation. *Comput. Aided Des. Appl.* 7, 4 (2010), 533–546.
- Goldstein, H., Poole, C., Safko, J. *Classical Mechanics*, 3rd edn. Addison Wesley, Boston, MA, USA, 2001.
- Gould, D. *The Top: Universal Toy Enduring Pastime*. Bailey Brothers and Swinfen Ltd, London, UK, 1975.
- Hirose, M., Mitani, J., Kanamori, Y., Fukui, Y. An interactive design system for sphericon-based geometric toys using conical voxels. In *Proceedings of the International Conference on Smart Graphics* (2011). Springer, Berlin/Heidelberg, Germany, 37–47.
- Lewis, D., Ratiu, T., Simo, J.C., Marsden, J.E. The heavy top: A geometric treatment. *Nonlinearity* 5, 1 (1992), 1.
- Luo, L., Baran, I., Rusinkiewicz, S., Matusik, W. Chopper: Partitioning models into 3D-printable parts. *ACM Trans. Graph.* 31, 6 (2012), 129:1–129:9.
- Macchietto, A., Zordan, V., Shelton, C.R. Momentum control for balance. *ACM Trans. Graph.* 28, 3 (2009), 80:1–80:8.
- Malko, G. *The One and Only Yo-Yo Book*. Avon, New York, NY, USA, 1978.
- Nocedal, J., Wright, S.J. *Numerical Optimization*. Springer, New York, NY, USA, 2000.
- Perry, J. *Spinning Tops and Gyroscopic Motion*. Dover Publications, Whitefish, MT, USA, 1957.
- Prévost, R., Whiting, E., Lefebvre, S., Sorkine-Hornung, O. Make it stand: Balancing shapes for 3D fabrication. *ACM Trans. Graph.* 32, 4 (2013), 81:1–81:10.
- Provatidis, C.G. Revisiting the spinning top. *Int. J. Mater. Mech. Eng.* 1, 4 (2012), 71–88.
- Skouras, M., Thomaszewski, B., Coros, S., Bickel, B., Gross, M. Computational design of actuated deformable characters. *ACM Trans. Graph.* 32, 4 (2013), 82:1–82:10.
- Stava, O., Vanek, J., Benes, B., Carr, N., Mèch, R. Stress relief: Improving structural strength of 3D printable objects. *ACM Trans. Graph.* 31, 4 (2012), 48:1–48:11.
- Zhou, Q., Panetta, J., Zorin, D. Worst-case structural analysis. *ACM Trans. Graph.* 32, 4 (2013), 137:1–137:12.
- Zhu, L., Xu, W., Snyder, J., Liu, Y., Wang, G., Guo, B. Motion-guided mechanical toy modeling. *ACM Trans. Graph.* 31, 6 (2012), 127.

Moritz Bächer (moritz.baecher@disneyresearch.com), Disney Research Zurich, Switzerland.

Bernd Bickel (bernd.bickel@ist.ac.at), Institute of Science and Technology Austria, Klosterneuburg, Austria.

Emily Whiting (emily@cs.dartmouth.edu), Dartmouth College, Hanover, NH.

Olga Sorkine-Hornung (sorkine@inf.ethz.ch), ETH Zurich, CS Department, Zurich, Switzerland.