Perception of physical stability and center of mass of 3-D objects

Steven A. Cholewiak

Department of Psychology and Center for Cognitive Science, Rutgers University, New Brunswick, NJ, USA

Roland W. Fleming

Department of Psychology, Justus-Liebig-University Giessen, Giessen, Germany

Manish Singh

Department of Psychology and Center for Cognitive Science, Rutgers University, New Brunswick, NJ, USA

Humans can judge from vision alone whether an object is physically stable or not. Such judgments allow observers to predict the physical behavior of objects, and hence to guide their motor actions. We investigated the visual estimation of physical stability of 3-D objects (shown in stereoscopically viewed rendered scenes) and how it relates to visual estimates of their center of mass (COM). In Experiment 1, observers viewed an object near the edge of a table and adjusted its tilt to the perceived critical angle, i.e., the tilt angle at which the object was seen as equally likely to fall or return to its upright stable position. In Experiment 2, observers visually localized the COM of the same set of objects. In both experiments, observers' settings were compared to physical predictions based on the objects' geometry. In both tasks, deviations from physical predictions were, on average, relatively small. More detailed analyses of individual observers' settings in the two tasks, however, revealed mutual inconsistencies between observers' critical-angle and COM settings. The results suggest that observers did not use their COM estimates in a physically correct manner when making visual judgments of physical stability.

Introduction

Research on visual perception generally focuses on the representation of overtly visible surface properties. In addition to estimating such properties—i.e., representing what objects in view are and where they are located (e.g., Marr, 1982)—an important goal of the visual system is to *predict* how visible objects are likely to behave in the near future. Predicting the physical behavior of objects is, among other things, crucial for the perceptual guidance of motor actions. Consider, for example, the visual guidance of motor actions aimed at intercepting an object in motion, or at catching a precariously balanced object that is about to fall. Predicting the physical behavior of objects in these and other situations requires observers to infer hidden forces acting on objects—e.g., gravity, support, friction—and often to do so from vision alone. Visually estimating the physical stability of objects involves an inference of unseen forces that requires integrating shape information across the entire 3-D object to accurately estimate the object's physical parameters, such as its center of mass. An open question has been how humans make these types of judgments and how their estimates relate to the physical dynamics of the world.

Traditional research on intuitive physics has shown that people often hold erroneous physical intuitions (McCloskey, 1983a, 1983b; McCloskey, Caramazza, & Green, 1980). For example, many people expect that a ball being swung at the end of a string will, if the string breaks, continue moving in a curved trajectory; or that an object dropped from a flying airplane will fall vertically straight down (McCloskey et al., 1980). On the other hand, our visuomotor interactions with objects in everyday life suggest that we have a good comprehension of physical attributes such as gravity, friction, and support relations. Indeed, subsequent research has shown that people are much more sensitive to violations of physical laws when they view real-time dynamic displays than when they are explicitly asked about their intuitions (Kaiser, Proffitt, & Anderson, 1985; Proffitt & Gilden, 1989). In visuomotor interactions involving catching a falling ball, for instance, McIntyre, Zago, Berthoz, and Lacquaniti (2001) have
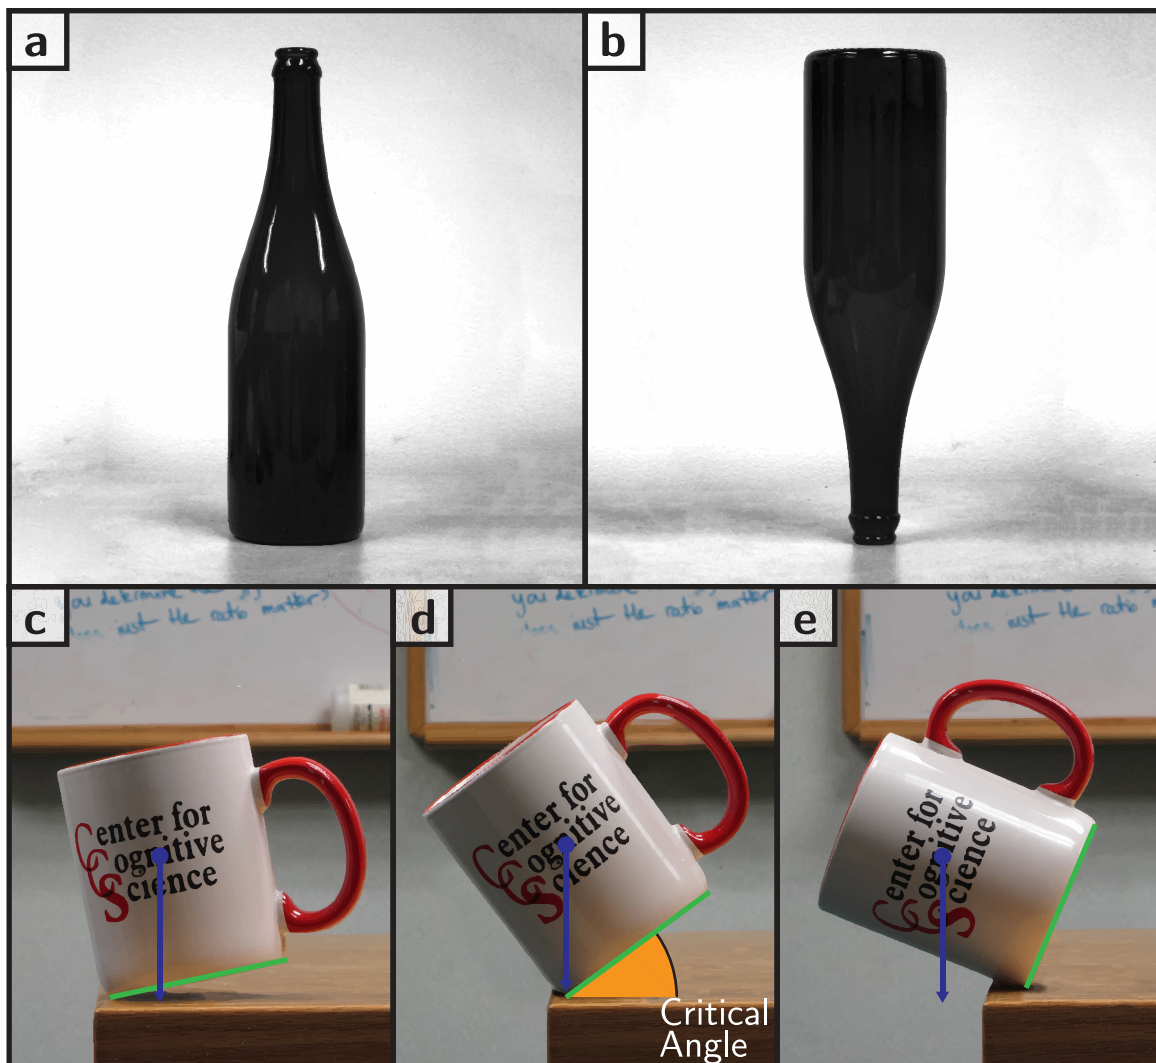
Figure 1. Examples illustrating the perceived stability of real objects. The top photographs show a bottle on a table, in one configuration that is very perceptually stable (a) and another that is very perceptually unstable (b). The bottom photographs show a coffee cup with its center of mass—represented by the blue circle and the blue line representing the gravity vector—vertically above the base (green; c), directly over the contact point when the object is at its critical angle (d), and outside the base (e).

shown that, in timing their hand movements, observers take into account acceleration due to gravity in a manner that is consistent with Newton's laws of motion.

Perhaps even more impressive than perceptual predictions involving moving objects are cases where observers can infer the action of underlying forces from a static scene or image. Infants as young as 6.5 months implicitly understand the influence of gravity and expect that objects that are not supported will fall down (Baillargeon, Needham, & DeVos, 1992). By 8 months they can also judge to some extent whether or not a cuboidal object is *adequately* supported (Baillargeon & Hanko-Summers, 1990; Baillargeon et al., 1992). Such judgments about how an object is likely to behave have important implications for judgments of future object behavior. Freyd, Pantzer, and Cheng (1988) found that

adult subjects who were shown a static image of an unsupported object that was previously shown to be supported had a systematic memory distortion when tasked with a same/different judgment—consistent with how the depicted object would behave if its support were in fact physically removed (e.g., in the case of gravity, they misremembered it as being lower than it actually was in the image). Based on this and other evidence, the authors argue that the representation of static scenes includes not just a kinematic component but a dynamic one as well—in other words, a representation of underlying forces.

An ecologically important judgment that relies on the implicit inference of underlying forces is the perceptual estimation of an object's physical stability. Consider the two bottles depicted in Figure 1: We can readily judge from vision alone that the bottle in Figure

1a is more physically stable than the one in Figure 1b. In other words, we naturally expect that the bottle in Figure 1a would be more resistant to the action of perturbing forces.

Similarly, we can judge in a quick glance that in Figure 1c, the cup is likely to return to its vertically upright position, whereas the same cup in Figure 1e will fall over. These expectations make physical sense. All forces acting on an object in a uniform gravitational field can be summarized by a single net force that acts on its center of mass (COM). Hence, when the gravity-projected COM lies within the supporting "base" surface (as in Figure 1c), the net torque acting on the object causes it to return to its upright position. However, with a large enough tilt—once the gravity-projected COM lies outside the support area, as in Figure 1e—the object topples over. A natural way to quantify the difference in stability is in terms of the critical angle, i.e., the angle through which an object in a given state of stable equilibrium can be tilted before it will topple over (Cholewiak, Fleming, & Singh, 2013). The critical angle corresponds to a state of unstable equilibrium in which the object is equally likely to fall over or to return to its upright position. As can be seen in Figure 1d, the critical angle is a function of both the width of the object's base and the height of its COM. The bottle in Figure 1a is more physically stable because it has a larger critical angle than the bottle in Figure 1b.

Samuel and Kerzel (2011) examined the perceptual estimation of balance of 2-D polygonal objects. In one experiment, the shapes consisted of two polygonal parts and rested on a vertex. Subjects adjusted the orientation of the object until it was perceived to be equally likely to fall to the right or to the left—which, physically speaking, requires that its (COM) be vertically aligned with its supporting vertex. The results showed that although subjects could perform this task, they were overly influenced by the eccentricity of the top part of the object, which led to errors in their judgments. In a second experiment, Samuel and Kerzel used polygonal planar objects sitting on a supporting edge, with varying degrees of equilibrium states, and showed that observers' stability judgments exhibited a conservative/anticipatory bias—namely, a bias in the direction of perceiving an object to be unstable, even though physically it would maintain its upright posture. This bias remained even after taking into account subjects' own perceptual estimates of the objects' COMs. The authors propose that subjects' responses may have been guided by a tendency to stay "on the safe side."

Recent work by Battaglia, Hamrick, and Tenenbaum (2013) presented a framework for evaluating physical stability in more complex scenes using a rigid-body dynamics simulation; they compared their model's performance to observers' judgments of stability. In their experiments, observers viewed scenes that contained towers of stacked cuboidal objects and were asked to evaluate whether the objects would fall (in scenes that were either physically stable or unstable) or in which direction they would fall (in scenes where the towers would always fall). The authors then used a computational physics simulation—built upon the Open Dynamics Engine and described by them as the Intuitive Physics Engine (IPE)—that incorporated the effect of gravity on the towers to model the scenes. There were three parameters manipulated in the IPE model that controlled state uncertainty (e.g., uncertainly about the object positions), mass densities of the objects, and latent forces (e.g., bumps and vibrations that may have been applied to the scene). Using their simulation, the authors could query the last state of the scene to determine how many of the blocks had fallen, the direction in which they fell, and the distances they fell from their starting point. The predictions from their model had good correspondence with observers' judgments of whether the towers were stable and in which direction unstable towers would fall. However, their model is based upon a rigid-body dynamics simulation engine that, by extension, assumes that the brain has a representation of the full physical state of a scene and is "inverting the physics of the scene" when judging the stability. Therefore, although it provides a framework to describe human performance for scenes with multiple interacting objects, the physics simulation by itself does not directly address the question of how humans make stability judgments based upon vision alone (e.g., based on the shape of a 3-D object).

Barnett-Cowan, Fleming, Singh, and Bülthoff (2011) investigated whether visual judgments of physical stability incorporate multisensory information and how the perceived stability of a single object can be affected by changes in the observer's gravitational frame of reference. They had observers either sit upright or lie on their left or right side when judging the critical angles of a series of objects. The objects were generated surfaces of revolution with a protrusion that was shifted up or down to alter the center of mass. The authors found a small but statistically significant effect of the shape manipulation, but interestingly, observers made stability judgments that were biased towards the subjective visual vertical—which was biased by their body posture—a finding that suggests that the visual system's estimates of physical stability incorporate multisensory information and may utilize biased internal representations of the gravitational frame.

Cholewiak et al. (2013) further investigated the influence of 3-D shape on visual estimates of stability by examining how well observers could track the

critical angle of an asymmetric 3-D object as a function of the direction in which it is tilted. In addition to a task involving the visual estimation of critical angle in different directions, they used asymmetric matching of object stability to determine what attributes of the asymmetric objects were used when judging overall physical stability: the average critical angle around the circumference or the minimum. They found that observers could track the critical angle as a function of tilt quite accurately and that their stability judgments in the asymmetric matching task were better explained by the asymmetric object's minimum critical angle (i.e., in the direction in which the object was least stable). The results suggested that physical stability is likely represented along a unitary dimension, where objects can be visually judged as more or less stable, despite variations along various shape dimensions.

The current study investigates the visual estimation of physical stability of rotationally symmetric 3-D objects, and its dependence on shape attributes. In the first experiment, we measured observers' perceived critical angles and compared these against the corresponding physical critical angles. As summarized earlier, this method has been successfully employed in recent work to measure perceived stability (Barnett-Cowan et al., 2011; Cholewiak et al., 2013). Experiment 1 examined the perception of physical stability by manipulating the objects' aspect ratios and overall volumes. Changing the aspect ratio provides a simple manipulation of the object's shape, whereas changing the volume allows us to test for size invariance—whether, consistent with physical predictions, visual estimates of critical angle are a function of intrinsic shape.[1]

As is clear from the definition of the critical angle, the physical notion of stability involves an object's COM in a fundamental way. Previous research has shown that observers use the COM to visually localize a shape—both in perceptual tasks, such as estimating the separation between two dot clusters (Morgan, Hole, & Glennerster, 1990), and in saccadic localization (Vishwanath & Kowler, 2003, 2004). Some systematic errors in perceptual estimation of the COM have also been reported in specific contexts, however (Baud-Bovy & Gentaz, 2004; Baud-Bovy & Soechting, 2001; Bingham & Muchisky, 1993).[2] Thus, if the estimation of perceived critical angle is found to deviate from the corresponding physical prediction, one must consider the possibility that perhaps a perceptual mislocalization of the COM is at least partly responsible. In Experiment 2, we therefore used a task involving the perceptual localization of COM for the same set of shapes. We examine whether any misperception in the physical stability of an object (Experiment 1) can be explained in terms of a corresponding misperception its COM (Experiment 2).

Experiment 1

We examined how accurately naïve observers could estimate the critical angle—the angle through which an object in a given state of stable equilibrium can be tilted before it will topple over (Cholewiak et al., 2013)—for different 3-D objects.

Methods

Observers

Fifteen observers at Rutgers University participated in the experiment. All were naïve as to the purpose of the experiment and received course credit for their participation.

Apparatus

The stimuli were generated in MATLAB using Psychtoolbox-3 (Brainard, 1997; Kleiner, Brainard, & Pelli, 2007; Pelli, 1997) running on an HP desktop computer and presented on a Sony Trinitron 20-in. CRT with a 1024 × 768 pixel resolution at a refresh rate of 120 Hz. Within Psychtoolbox-3, experimental scenes were rendered using the MATLAB OpenGL (MOGL) toolbox and were presented stereoscopically using an NVIDIA Quadro 4000 graphics card, CrystalEyes LCD shutter glasses, and an E-2 IR emitter by StereoGraphics (60-Hz binocular update rate).[3] Observers were comfortably seated with a chin rest supporting their head 80 cm from the screen.

COM locations and critical angles were analytically computed using Mathematica 8. The experimental scenes were illuminated with OpenGL's per-vertex lighting model using a fixed function pipeline with specular highlights applied to textured objects, and with the specular reflection angles calculated assuming a constant view direction parallel to the direction of the z-axis. There were two light sources, both with white ambient, diffuse, and specular components, and the objects' surfaces had ambient, diffuse, and specular reflectances. Objects were textured with 2-D planar textures mapped to their surfaces.

Stimuli

The experiment used the method of adjustment to measure the perceived critical angle for objects in a 4 (aspect ratios) × 2 (volumes) factorial design. Stimuli were rendered 3-D scenes containing a table surface with a single object placed near its left edge (see top image in Figure 2 for an example). Each object was an upright conical frustum, rendered with a wood-grain texture to mimic the visual appearance of a solid block
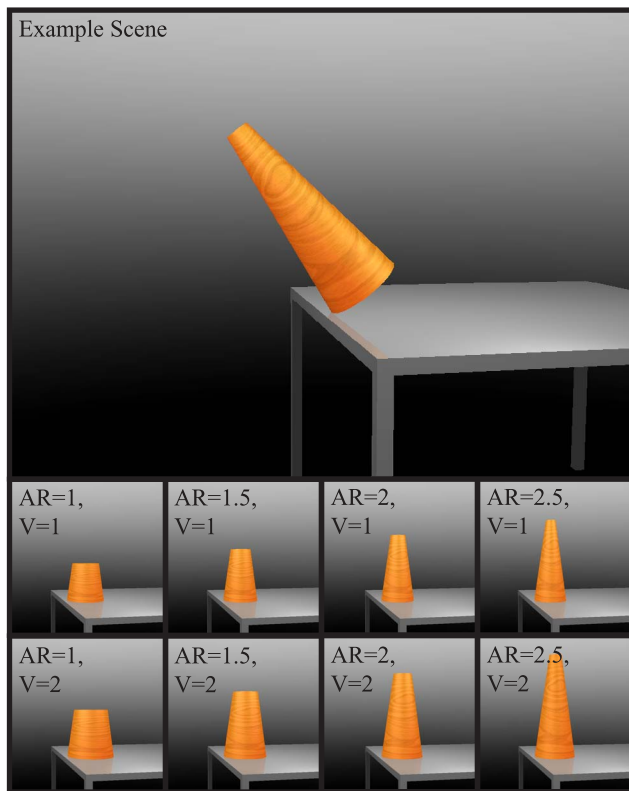
Figure 2. At the top, an example scene from Experiment 1 with the conical frustum (aspect ratio = 1.5, volume = 2) tilted so that it is at its critical angle (38.90°). Below are cropped examples from Experiment 1 showing the four aspect ratios of the conical frustums (1, 1.5, 2, and 2.5) and the two volumes (1 and 2).

of wood. On any given trial, the presented object had one of four aspect ratios (height/base diameter): 1, 1.5, 2, or 2.5 (see bottom rows of Figure 2). Aspect ratio was manipulated in such a way that the overall volumes of the frustums were preserved. The four aspect ratios were repeated with two different volumes that differed by a factor of 2. The heights of the frustums subtended approximately 4.1° of visual angle for the shortest frustum with volume = 1 to 11.4° of visual angle for the tallest with volume = 2. The taper angle of the conical frustums was fixed at 82°.

Procedure

The observers' task was to adjust the tilt of the object until it was perceived to be equally likely to fall off the table or to return to its upright position on the table. We refer to this angle of tilt as the perceived critical angle. Pressing and holding the Shift key while pressing the left or right arrows allowed observers to adjust the tilt angle at a faster rate (2°/press instead of 0.2°/press). The object's motion was constrained to be orthogonal to the edge of the table, with the axis of rotation at the

point on the base of the object closest to the table's edge.

On each trial, the conical frustum was shown with an initial tilt angle of either 0° (vertically upright with base fully on the table) or 90° (tilted so that the entire object was over the precipice).

Each observer performed a total of 112 adjustments: 14 for each of the 4 × 2 combinations of aspect ratio and volume. Half of these had an initial tilt angle of 0°, and the other half had an initial tilt angle of 90°.

Results

Because we used inexperienced observers from an undergraduate subject pool, we established a criterion to exclude observers with highly noisy data—specifically, that the overall standard deviation of their settings (collapsed using root mean square across all conditions) should not exceed 10°. Data from two observers were excluded based on this criterion.[4] Figure 3a shows the group data collapsed across the remaining 13 observers: Average critical-angle settings are plotted as a function of frustum aspect ratio, for the two volumes and two initial tilt angles. The black dotted curve shows the physical predictions for the critical angle. Observers were able to track a change in the stability as the aspect ratio was manipulated. However, on average, the critical angle—and hence object stability—tended to be somewhat underestimated for small aspect ratios (short and wide shapes) and overestimated for large aspect ratios (long and narrow shapes).

A repeated-measures ANOVA revealed a highly significant effect of aspect ratio, $F(3, 36) = 50.89$, $p < 0.01$. All 13 observers exhibited a significant influence of aspect ratio in their individual data. A small but statistically reliable effect of initial tilt angle was also seen in the group data, possibly due to hysteresis, $F(1, 12) = 7.36$, $p < 0.05$. The effect was quite small, however. Only one of the 13 observers, S10, exhibited an average difference of more than 5°. Finally, observers' settings showed no reliable dependence on the overall volume of the conical frustum, $F(1, 12) = 1.48$, $p > 0.05$. This lack of dependence on volume was evident both in the group data and in individual observers' data. This suggests that the perception of object stability is unaffected by the overall size of the object and—consistent with the physical definition of critical angle—is a function of its intrinsic shape only.

Modeling

We next wished to model observers' settings of object stability as a function of object shape (here,
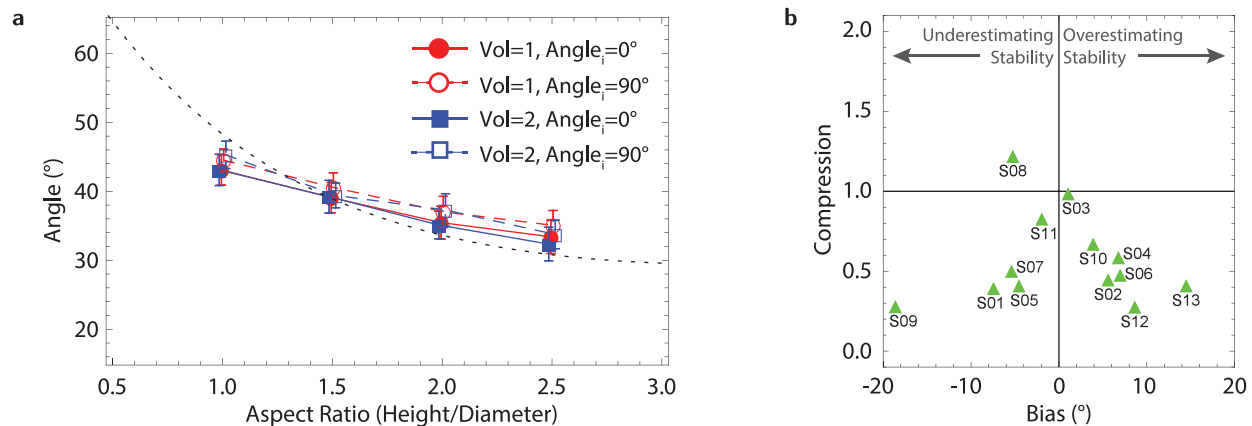
Figure 3. (a) Average performance across all observers as an effect of aspect ratio (x-axis), volume (red circles = 1, blue squares = 2), and initial tilt angle (solid = 0°, dashed = 90°). The physical critical angle is the black dotted line, plotted as a function of the aspect ratio. Note that for the lowest aspect ratio, observers underestimated stability, while for the highest aspect ratios, they overestimated stability. (b) Computational results comparing observers' judgments relative to the physical predictions. Overall bias is a measure of an observer's over/underestimation bias, while compression is a measure of the influence of the aspect ratio on judgments. Each green triangle represents a single observer (labeled by ID number). Note that there is no consistent overestimation or underestimation bias across observers and that the observers' responses—as an effect of aspect ratio—were compressed relative to the physical prediction (the majority are <1).

aspect ratio). An examination of the individual observer data suggested that the group pattern of underestimating stability for small aspect ratios and overestimating it for large aspect ratios might be the result of two factors: (a) the fact that most observers' data curves exhibit a shallower slope than the physical-prediction curve (i.e., a smaller dependency on aspect ratio than physically predicted), and (b) the fact that there is no consistent bias across observers (despite the fact that different individual observers exhibit small biases to either underestimate or overestimate object stability). In order to quantify these differences across observers, we fitted a model to each individual observer's data, of the form

$$y = \beta_0 + \beta_1 F(x), \tag{1}$$

where x is the aspect ratio, y is the perceived critical angle, β_0 and β_1 are free parameters, and $F(x)$ specifies the physical critical angle as a function of the conical frustum aspect ratio—i.e., the function plotted in the dotted black curve in Figure 3a.[5] The fits of this model thus capture two ways in which individual observers' settings can deviate from the physical predictions. We summarized each observer's stability settings in terms of two variables of interest:

(a) Bias B, or the average magnitude of underestimation or overestimation in physical stability. We measure bias as the mean difference between the fitted model to an observer's data and the corresponding physical prediction over the four levels of aspect ratio (Equation 2). Positive biases ($B > 0$) indicate an overall overestimation of the critical

angle and, therefore, overestimation of the stability. Likewise, negative biases ($B < 0$) correspond to underestimation of the critical angle and object stability.

$$B = \frac{\sum_{i=1}^{4} [\hat{y}(x_i) - F(x_i)]}{4} \tag{2}$$

(b) Compression C, or the influence of aspect ratio on perceived critical angle relative to the corresponding influence of aspect ratio on the *physical* critical angle. The compression is therefore a measure of the relative magnitude of the effect of aspect ratio on the observers' judgments. This is simply the fitted parameter β_1 (hence $C = \beta_1$). When $\beta_1 > 1$, an observer's stability settings are more influenced by aspect ratio than is physical stability. When $\beta_1 < 1$, an observer's stability settings are less influenced by aspect ratio than is physical stability.

Figure 3b summarizes each individual observer's data as a point within the 2-D space defined by these two fitted parameters. Two trends are evident in this plot: First, although individual observers exhibit different biases, observers who overestimate stability (positive bias) and those who underestimate stability (negative bias) are roughly balanced. In other words, we see no systematic tendency to underestimate stability in this experiment (i.e., no systematic conservative/anticipatory bias), as was observed in the Samuel and Kerzel (2011) study. Second, the compression parameter C is consistently smaller than 1 for

all but one of the 13 observers, indicating that most observers' critical-angle settings are less influenced by aspect ratio than are the corresponding physical predictions.

Discussion

In Experiment 1, we found that object stability—measured in terms of critical angle—was, on average, underestimated for short, wide shapes with lower aspect ratios and overestimated for tall, narrow shapes with higher aspect ratios. There was an effect of aspect ratio, meaning that as the aspect ratio increased, observers changed their settings to reflect the change in the physical stability. There was no effect of the volume manipulation on the perceived critical angle, suggesting size/volume invariance. This size invariance is consistent with the prediction that critical-angle judgments should be dependent upon geometric relations between the COM and the supporting base of the object, and hence should stay constant as the volume is increased.

As the individual results illustrate (Figure 3b), some observers had overestimation or underestimation biases, but there was no consistent bias across all observers. In addition, the compression in the observers' responses illustrates that although there was an influence of aspect ratio on their critical-angle judgments, it was not as large as the physical prediction for the majority of the observers.

Experiment 2

The pattern of settings in the visual estimation of object stability in Experiment 1 can be summarized as follows: (a) We saw no systematic tendency across observers to overestimate or underestimate physical stability, and (b) observers' settings of object stability were less affected by aspect ratio than physically predicted. Previous empirical work on the visual estimation of COM has pointed to judgment biases (e.g., Baud-Bovy & Gentaz, 2004; Baud-Bovy & Soechting, 2001; Bingham & Muchisky, 1993). So a natural question is whether the observed pattern of biases in the estimation of object stability might reflect a corresponding misperception of the objects' COM. Specifically: (a) Are perceptual COM estimates *also* less influenced by aspect ratio than are physical COMs and (b) Do individual observers who tend to underestimate object stability when judging the critical angle correspondingly perceive the COM of the frustum objects to be *higher* than it actually is (and vice versa)? Experiment 2 addresses these questions by obtaining

visual estimates of COM on the same set of 3-D objects from the same group of observers.

Methods
Observers

The same group of observers participated as in Experiment 1.

Apparatus

The apparatus was the same as in Experiment 1.

Stimuli

Observers participated in Experiment 2 after completing Experiment 1 and taking a short break. The method of adjustment was used to estimate their perceived COM for the same set of conical frustum shapes used in Experiment 1. Unlike the opaque surfaces in Experiment 1, however, the conical frustums in Experiment 2 were rendered with a transparent surface through which the axis of symmetry was visible—shown as a solid gray line (see top image in Figure 4). Along this line, the vertical height of a small red dot—technically a small sphere in the 3-D rendered scene—was adjusted by the observer.

Each observer made 10 adjustments of COM height for each combination of aspect ratio and volume, for a total of 80 trials. The initial location of the probe dot, at the start of each trial, was at the top of the object on half the trials and at the bottom on the other half.

In the instructions to the observers, the concept of the COM of the object was explained by describing the balancing point of a serving tray, and noting the existence of a central point where the tray could be balanced on a single finger. Observers generally found this concept quite intuitive. They were then instructed that their task was to similarly localize the balance point of the conical frustums.

To make their COM settings, observers used the up and down arrow keys to adjust the dot's vertical position along the axis of symmetry. They were instructed to adjust the height of the probe dot to localize the perceived COM of the object. As in Experiment 1, pressing the Shift key allowed them to make faster adjustments.

Results

Since our goal was to examine the degree of mutual consistency between settings of object stability and settings of COM, we only analyze data from the 13 observers whose stability data were reported in
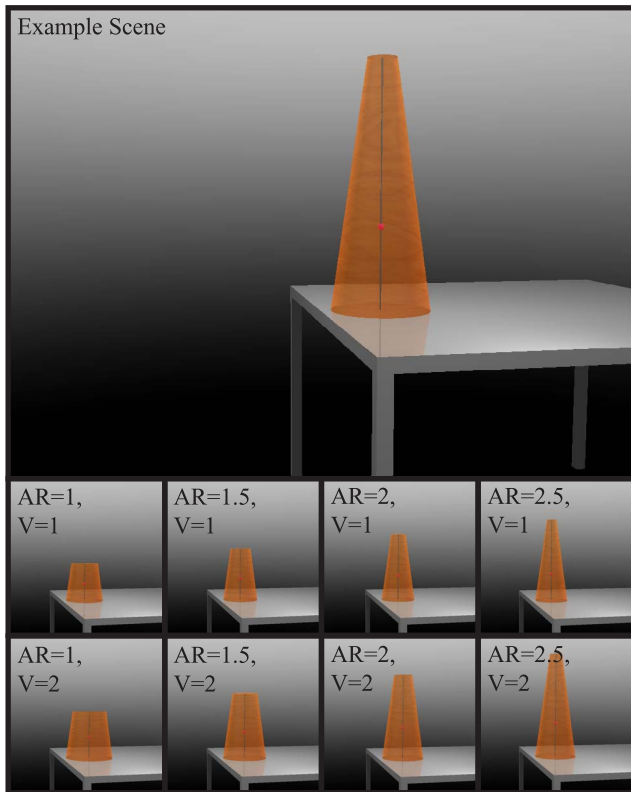
Figure 4. At the top, an example scene from Experiment 2 with a marker placed at the conical frustum's center of mass (aspect ratio = 2, volume = 2). Below are cropped examples from Experiment 2 showing the four aspect ratios of the conical frustums (1, 1.5, 2, and 2.5) and the two volumes (1 and 2). Note that the observer's judgment of center of mass (normally rendered as a red sphere inside the object) was constrained to the vertical axis of symmetry (solid gray line).

Experiment 1.[6] Figure 5a shows the group data across the 13 observers: The average COM height settings are plotted as a function of frustum aspect ratio, for the two volumes and two initial probe locations. The two dashed curves show the physical predictions of the COM height for the two different volumes. Unlike the critical angle, the COM heights are of course *not* invariant to the size of the object. Although the deviations from the physical COM predictions are not large, for most aspect ratios the observers tend to exhibit a slight tendency to underestimate the height of the COM (i.e., to perceive the COMs as being located slightly lower within the frustums than they actually are).

A repeated-measures ANOVA on the group data revealed a highly significant effect of aspect ratio, $F(3, 36) = 275.7$, $p < 0.001$. All 13 observers exhibited a highly significant influence of aspect ratio in their individual data. Similarly, there was a highly significant influence of volume on observers' COM settings, $F(1, 12) = 235.5$, $p < 0.001$, with all observers exhibiting a

significant influence in their individual data. The effect of the initial location of the probe dot was not significant, $F(1, 12) = 2.54$, $p > 0.05$. None of the interactions were significant.

Modeling

In order to compare observers' COM settings to their estimates of object stability from Experiment 1, we converted their observed COM settings into the critical angles they *would* generate, were the observers using these COM estimates in the physically correct manner to estimate the critical angles. This transformation allows for direct comparison of the results from the two experiments, to determine if observers are using the COM in a physically consistent manner when judging an object's critical angle. These "implied critical angles" can simply be calculated using the radius of the object's base (R_{base}) and the height of the observer's COM setting (h_{COM}):

$$\Theta_{CA} = \tan^{-1}\left(\frac{R_{base}}{h_{COM}}\right)\frac{180}{\pi}. \tag{3}$$

The group data for these implied critical angles are plotted in Figure 5b. In these plots, there is only one physical prediction curve because the physical critical angles are invariant to object volume. Consistent with this prediction, a repeated-measures ANOVA on these transformed data revealed that there was no longer a main effect of volume, $F(1, 12) = 2.02$, $p > 0.05$.

The group COM data replotted in terms of implied critical angle now exhibit, for all aspect ratios but the highest, a systematic tendency to overestimate the critical angle.[7] Were observers to use their perceptually estimated COMs in the physically correct manner to estimate the critical angle, their stability estimates from Experiment 1 should have matched those in Figure 5b. However, comparing Figure 5b with observers' critical-angle settings in Experiment 1 (shown in Figure 3a), we see that performance in the critical-angle stability task does not appear to be consistent with a strategy of deriving critical angles from perceived COMs—at least not in the physically correct manner.

We next examine the degree of consistency between the two tasks at the level of individual observers. As in Experiment 1, we summarize individual observers' performance using the two parameters of compression (from Equation 1) and bias (from Equation 2) to characterize their implied critical angles derived from their COM settings. Figure 6 summarizes each individual observer's data as points (orange downward-facing triangles) within the 2-D space defined by these two variables. A comparison with the corresponding data from Experiment 1 (green upward-facing triangles) indicates large and systematic differences between
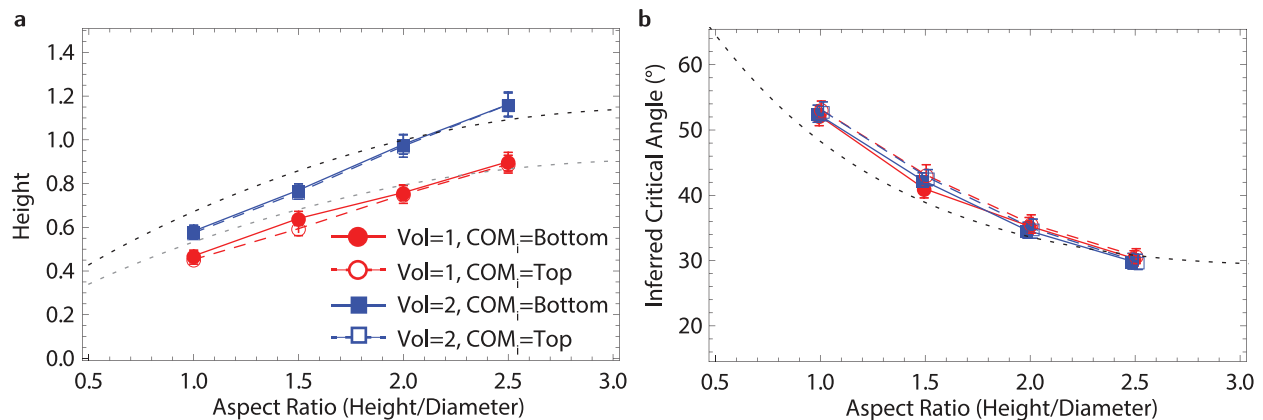
Figure 5. (a) Average judgments of center of mass (COM) across all observers as a function of aspect ratio (x-axis), volume (red circles = 1, blue squares = 2), and initial COM start location (solid lines = estimates that began at the bottom of the object, dashed = started at the top). The physical COM locations for the two volumes are the dotted lines. Note that observers underestimated the COM height for low aspect ratios and overestimated it for high aspect ratios. (b) Average COM judgments, as implied critical angles plotted as a function of aspect ratio.

the two. There was a large reduction in bias for the majority of observers (i.e., observers were more accurate) when they were asked to judge the objects' COMs. Additionally, individual observers' COM estimates were not systematically biased in the same direction as their biases in the critical-angle task: Observers who overestimated stability in Experiment 1 were not necessarily biased to make their COM judgments closer to the base of the object in Experiment 2, and visa versa. Thus, the majority of observers were much closer to the physical prediction in the COM task of Experiment 2 than in the critical-angle task of Experiment 1.

The conclusion that observers were not using the perceived COM as measured in Experiment 2—at least not in a physically correct manner—to make their stability judgments was confirmed by independence tests for the bias and compression parameters for the

two experiments. The bias in Experiment 1 was not predictive of the bias in Experiment 2, Spearman's $\rho = 0.033$, $p = 0.915$, and neither was the compression in Experiment 1 predictive of the compression in Experiment 2, Spearman's $\rho = -0.352$, $p = 0.239$.

Discussion

In Experiment 2, we found that the same observers from Experiment 1, whose responses were compressed relative to the physical prediction, now had responses that were much closer to the physical prediction. For all 13 observers, the compression in their responses was greater than 1, indicating that they put a larger emphasis on the aspect ratio when making their COM judgments than would be predicted physically. This is in contrast to their responses on the critical-angle task, where the compression factor was less than 1.

In addition, the magnitude of bias was consistently smaller in Experiment 2 (i.e., accuracy was better in the COM task than the stability task), and the sign of bias was not balanced across all observers in Experiment 2; unlike in Experiment 1, the bias tended to be more in the positive direction.
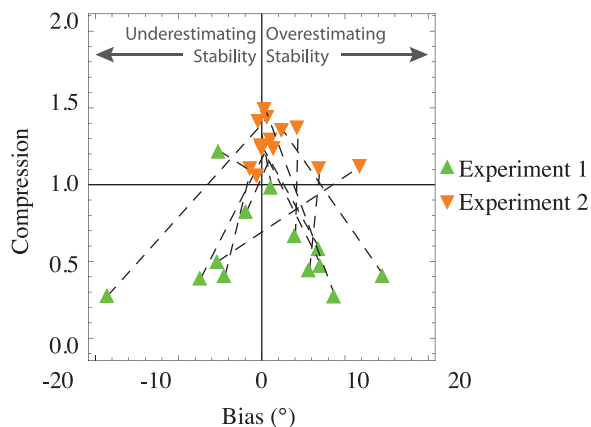

Figure 6. Compression (y-axis) and bias (x-axis) for the model fits in Experiments 1 (green triangles) and 2 (orange triangles). The dashed lines connect individual observers' model fits.

Conclusions

The experiments reported here show that people can use purely visual information to infer the forces acting upon 3-D objects in a scene. Using these inferences, observers were able to make stability judgments using a critical-angle task (Experiment 1) and could use the

visual representation of shape to estimate an object's COM (Experiment 2).

The results of Experiment 1 showed that, on average, observers are quite accurate when making their judgments of object stability using a critical-angle task. Their visual judgments of stability were invariant to size/volume, and appeared to be a function of only shape. In the group data, there was a bias to underestimate the stability of low-aspect-ratio shapes and overestimate the stability of high-aspect-ratio shapes in the average data; however, this was due to (i) individual variability in bias (with some observers underestimating stability, others overestimating it) and (ii) a compression in the data curves relative to the physical predictions (i.e., a smaller perceptual influence of aspect ratio than physically predicted).

These results contrast with those of Samuel and Kerzel (2011), where observers' stability judgments were biased towards underestimating the stability of their shapes. However, it is important to note that the shapes used by Samuel and Kerzel were 2-D polygons, while the objects used in these experiments were more naturalistic, 3-D stimuli presented stereoscopically in rendered scenes. When judging the critical angle of 3-D shapes in our experiments, observers were *not* universally underestimating the stability of their judgments, and therefore we conclude that stability judgments are not always judged on the safe side.

The critical-angle results contrasted with the COM estimates in Experiment 2, showing that observers were much closer to the physical predictions when they were asked to judge the center of mass of the same objects. There was much less bias (observers were more accurate), and the compression observed in Experiment 1—the shallower slopes in observers' stability estimates as a function of aspect ratio—was no longer present. The Intuitive Physics Engine proposed by Battaglia et al. (2013) could potentially be used to explain some of the observed compression, given that the model assumes that observers integrate a noisy representation of the scene and the objects therein; however, it is unclear how it could explain the individual biases present in the observer data. Observers who systematically underestimated the stability of the objects in the critical-angle task were *not* more likely to overestimate the COM height in the second experiment, so performance in the critical-angle task is not predictive of their performance in the COM localization task.

Observers' settings of center of mass were thus inconsistent with their settings of critical angle. Although observers are able to localize the center of mass, it does not appear to explain performance in a visual balancing task. Therefore, comparing the two sets of results, observers do not appear to use their perceived COM—at least not in the physically correct manner—when making judgments of object stability.

Keywords: 3-D shape, perceived object stability, critical angle, center of mass

Footnotes

[1] Physical stability depends on the mass distribution within an object, which is a function of (i) the object's shape and (ii) variations in the object's density. In the absence of specific information to the contrary, observers presumably make the default assumption of uniform density. In that case, the critical angle is fully determined by the object's 3-D shape.

[2] The term "center of gravity" (COG) is commonly used in this literature. For all practical purposes, the concepts of COG and COM are identical. Physically speaking, the two become distinguishable only in situations where the gravitational field is nonuniform.

[3] The binocular views were rendered using a fixed camera separation of 6 cm.

[4] We confirmed, however, that including the data from these two observers does not alter any of our conclusions.

[5] The functional form of the physical critical-angle curve (in radians) for conical frustums with taper angles of 82° is as follows:

$$F(x) = \tan^{-1}\left[\frac{3 - 6x\cot(82\frac{\pi}{180}) + 4x^2\cot(82\frac{\pi}{180})^2}{x(3 - 8x\cot(82\frac{\pi}{180}) + 6x^2\cot(82\frac{\pi}{180})^2)}\right]$$

[6] We note, moreover, that the two observers whose data were excluded from the Experiment 1 analysis also produced very noisy data in the COM task.

[7] By definition, an underestimation of an object's COM height implies an overestimation of its critical angle.

References

Baillargeon, R., & Hanko-Summers, S. (1990). Is the top object adequately supported by the bottom object? Young infants' understanding of support relations. *Cognitive Development, 5,* 29–53, doi:10.1016/0885-2014(90)90011-H.

Baillargeon, R., Needham, A., & DeVos, J. (1992). The development of young infants' intuitions about support. *Early Development and Parenting, 1*(2), 69–78, doi:10.1002/edp.2430010203.

Barnett-Cowan, M., Fleming, R. W., Singh, M., & Bülthoff, H. H. (2011). Perceived object stability depends on multisensory estimates of gravity. *PLoS ONE, 6,* 1–5, doi:10.1371/journal.pone.0019289.

Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences, USA, 110*(45), 18327–18332, doi:10.1073/pnas.1306572110.

Baud-Bovy, G., & Gentaz, E. (2004). The visual localization of the centre of triangles in young children and adults. *Current Psychology Letters, 13,* 2–11.

Baud-Bovy, G., & Soechting, J. (2001). Visual localization of the center of mass of compact, asymmetric, two-dimensional shapes. *Journal of Experimental Psychology: Human Perception and Performance, 27*(3), 692–706, doi:10.1037/0096-1523.27.3.692.

Bingham, G. P., & Muchisky, M. M. (1993). Center of mass perception and inertial frames of reference. *Perception & Psychophysics, 54*(5), 617–632, doi:10.3758/BF03211785.

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision, 10*(4), 433–436, doi:10.1163/156856897X00357.

Cholewiak, S. A., Fleming, R. W., & Singh, M. (2013). Visual perception of the physical stability of asymmetric three-dimensional objects. *Journal of Vision, 13*(4):12, 1–13, http://www.journalofvision.org/content/13/4/12, doi:10.1167/13.4.12. [PubMed] [Article]

Freyd, J. J., Pantzer, T. M., & Cheng, J. L. (1988). Representing statics as forces in equilibrium. *Journal of Experimental Psychology: General, 117*(4), 395–407, doi:10.1037/0096-3445.117.4.395.

Kaiser, M. K., Proffitt, D. R., & Anderson, K. (1985). Judgments of natural and anomalous trajectories in the presence and absence of motion. *Learning, Memory, and Cognition, 11*(4), 795–803, doi:10.1037/0278-7393.11.1-4.795.

Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in Psychtoolbox-3? *Perception,* 36, ECVP Abstract Supplement.

Marr, D. (1982). *Vision*. Cambridge, MA: MIT Press, doi:10.7551/mitpress/9780262514620.001.0001.

McCloskey, M. (1983a, April). Intuitive physics. *Scientific American, 248,* 122–130, doi:10.1038/scientificamerican0483-122.

McCloskey, M. (1983b). Naive theories of motion. In D. Gentner & A. L. Stevens (Eds.), *Mental models* (pp. 299–324). Hillsdale, NJ: Erlbaum.

McCloskey, M., Caramazza, A., & Green, B. (1980, December). Curvilinear motion in the absence of external forces: Naive beliefs about the motion of objects. *Science, 210*(4474), 1139–1141, doi:10.1126/science.210.4474.1139.

McIntyre, J., Zago, M., Berthoz, A., & Lacquaniti, F. (2001). Does the brain model Newton's laws? *Nature Neuroscience, 4,* 693–694, doi:10.1038/89477.

Morgan, M. J., Hole, G. J., & Glennerster, A. (1990). Biases and sensitivities in geometrical illusions. *Vision Research, 30*(11), 1793–1810, doi:10.1016/0042-6989(90)90160-M.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10*(4), 437–442, doi:10.1163/156856897X00366.

Proffitt, D. R., & Gilden, D. L. (1989). Understanding natural dynamics. *Journal of Experimental Psychology: Human Perception and Performance, 15*(2), 384–393, doi:10.1037/0096-1523.15.2.384.

Samuel, F., & Kerzel, D. (2011). Is this object balanced or unbalanced? Judgments are on the safe side. *Journal of Experimental Psychology: Human Perception and Performance, 37*(2), 529–538, doi:10.1037/a0018732.

Vishwanath, D., & Kowler, E. (2003). Localization of shapes: Eye movements and perception compared. *Vision Research, 43*(15), 1637–1653, doi:10.1016/S0042-6989(03)00168-8.

Vishwanath, D., & Kowler, E. (2004). Saccadic localization in the presence of cues to three-dimensional shape. *Journal of Vision, 4*(6):4, 445–458, http://www.journalofvision.org/content/4/6/4, doi:10.1167/4.6.4. [PubMed] [Article]