Expectations affect the perception of material properties

L. M. Alley, A. C. Schmid, K. Doerschner
Authors:

Lorilei M. Alley*[1], Alexandra C. Schmid[*+1], Katja Doerschner[1,2]

[1] Justus Liebig University Giessen

[2] Bilkent University

[*] Equal contribution

[+] Corresponding author

13.12.2019

ABSTRACT

Many objects that we encounter have 'typical' material qualities: spoons are hard, pillows are soft, and jell-O dessert is wobbly. Over a lifetime of experiences, strong associations between an object and its typical material properties may be formed, and these associations not only include how glossy, rough or pink an object is, but also how it behaves under force: we expect knocked over vases to shatter, popped bike tires to deflate, and gooey grilled cheese to hang between two slices of bread when pulled apart. Here we ask how such rich visual priors affect the visual perception of material qualities and present a particularly striking example of expectation violation. In a cue conflict design, we pair computer-rendered familiar objects with surprising material behaviors (a linen curtain shattering, a porcelain teacup wrinkling, etc.) and find that material qualities are not solely estimated from the object's kinematics (i.e. its physical (atypical) motion while shattering, wrinkling, wobbling etc.); rather, material appearance is sometimes "pulled" towards the "native" motion, shape, and optical properties that are associated with this object. Our results, in addition to patterns we find in response time data, suggest that visual priors about materials can set up high-level expectations about complex future states of an object and show how these priors modulate material appearance.

INTRODUCTION

The material an object is made of endows it with a certain utility or function: chairs are usually rigid, in order to afford sitting, spoons are hard to afford eating, and towels are soft and absorbent to afford drying. The material choices for many types of man-made objects (keys, cups, cushions) tend to be restricted, and thus many objects that we encounter have a 'typical' material: bookshelves are made from wood, pillows from cloth and down, and keys from metal. Over a lifetime of interacting with objects, we learn not only what they look like, but also how their material 'behaves' under different types of forces. For example, we know that porcelain vases shatter when knocked over, that soft jell-O wobbles when poked, and that rubber balls bounce when thrown at the wall. Strong associations are formed between the object's shape and its optical and mechanical material properties.

Observers rely on these associations when estimating material qualities, which can lead to paradoxical findings. For example, physically, the softness of an object (compliance) is independent of its optical properties (color, gloss or translucency), and with sophisticated production techniques, it is possible to manufacture soft objects with any set of optical properties. Yet, when judging the softness of an object in static images, research has shown that that optical characteristics of the material affected observers' judgements, e.g. a velvety cube was perceived as softer than a chrome cube of the same shape and size (Paulun, Schmidt, van Assen, & Fleming, 2017; Schmidt, Paulun, van Assen & Fleming, 2017). To account for these results, Paulun et al., (2017) propose an indirect route to perception, where image cues activate the memory of a particular material along with its learned associations, including its typical mechanical properties (e.g. 'this looks like honey - so it's probably quite runny.'). They suggest that this route operates in parallel with a more direct one, where image cues directly convey something about the properties of the material (e.g. Paulun et al., 2017, Paulun, Kawabe, Nishida, & Fleming, 2015, van Assen, Wijntjes, & Pont, 2016, van Assen, Barla, & Fleming, 2018, Marlow & Anderson, 2016, see Figure 1 for an illustration).

Learned associations about material properties are not only evoked by the optical characteristics of an object but may also be elicited by its shape: Experience with soft materials or liquids appears to create strong associations between shape deformations and perceived mechanical material qualities (e.g. Paulun et al., 2017, Schmidt et al., 2017, Paulun et al., 2015, van Assen et al., 2016, van Assen et al., 2018, Schmid & Doerschner, 2018). Similarly, if a shape reminds the observer of a specific object, its typical mechanical material properties might also become activated. Consequently, recognizing an object should cause observers to activate strong predictions about the object's material properties, and how it should behave under

physical forces. Here we ask whether these predictions in turn affect how we perceive material properties?

Figure 1. Contribution of prior associations and image cues on perceived material qualities. *The role of predictions or associative mechanisms in material perception is not well understood. A. The perception of material qualities (such as gelatinousness) can be influenced by prior associations between dynamic optics, shape, and motion properties.* ***A.*** *Watching the green (left) object deform may evoke an association with green Jell-O (****B****.), and may therefore be perceived as wobblier and more gelatinous than the matte object, despite both objects wobbling in identical ways (as shown in Supplementary Movies S1 and S2).* ***C.*** *Alternatively, the green object may be perceived as wobblier due to larger image differences between frames, and potentially higher motion energy, as illustrated on the right* (Doerschner et al., 2011, Doerschner, Kersten, & Schrater, 2011)*). A combination of associative and modulatory mechanisms is also possible. The difference in motion energy in greyscale images of the translucent object in panel C. is about seven (6.8) times larger than that of the matte one, purely due to the difference in optical properties between these two objects.*

The human brain uses prior knowledge to continuously generate predictions about visual input in order to make quick decisions and guide our actions (see Kveraga, Ghuman, & Bar, 2007 for a review). Predictions about material properties should be no exception to this: To avoid small daily disasters, we need to be able to predict how bendable and heavy a cup is before picking it up. How prior knowledge is combined with visual input has been investigated in several areas of vision. For example, researchers have studied how color memory for objects influences perceived color (e.g. Hansen, Olkkonen, Walter, & Gegenfurtner, 2006; Olkkonen & Allred, 2014), or how memory modulates the perception of motion direction or dynamics of binocular rivalry (Scocchia, Valsecchi, Gegenfurtner, & Triesch, 2013 & Scocchia, Valsecchi, & Triesch, 2014, Chang & Pearson, 2018). Within the framework of Bayesian models (e.g. Landy, Maloney, Johnston, & Young, 1995; Ernst & Banks, 2002; Hillis, Ernst, Banks, & Landy, 2002; Kersten, Mamassian, & Yuille, 2004) the influence of prior knowledge on the integration of different types of sensory input has been looked at more formally. In particular,

cue conflict scenarios have proven extremely useful to generate insights about the complex interplay of prior selection and the weighting of sensory input in the perception of object properties (e.g. Knill, 2007). We will use an experimental paradigm analogous to cue conflict by juxtaposing indirect (prior) and direct (sensory) information in the perception of material properties, to formally test whether expectations about an object's mechanical properties are generated.

Violating the expected mechanical properties of materials necessarily involves image motion. Figure 2 illustrates this. Shown below are three frames from an animation by Florent Porta. The first panel sets up the viewer's expectations about the objects material properties, (i.e. the balloon will pop when it comes into contact with the spines of the cactus). As the movie proceeds, our expectations about the event-to-unfold are violated, and the viewer is quite surprised when the cactus pops like the balloon would. In the present study we directly test how such expectations affect the perception of material properties, by compare material perception for falling objects that deform in surprising and expected ways.

Figure 2. Three frames from "Preposterous" by Florent Porta*. Artists have played with our expectation of how objects and their materials should behave. In this study, we compare material perception for falling objects that deform in surprising and unsurprising (i.e. expected) ways. Retrieved from https://vimeo.com/191444383.*

We anticipate our results using the Bayesian framework for analogy. Suppose the task of an observer, watching the movie in Figure 2, was to rate the cactus' rigidness. If the observer recognizes the object as a cactus, it is probably safe to assume that observers have a strong prior belief about how rigid cacti are, but it is also possible that other, weaker, prior beliefs exist. When confronted with a popping cactus, the visual system may either veto all cues that previously suggested that the cactus was rigid (e.g. Landy, Maloney, Johnston, & Young, 1995)

and judge the cactus as very soft, or it may be that the visual system down-weighed cues to rigidness (Knill, 2007). If cue vetoing occured, we would expect to find no difference in the perception of mechanical material properties (e.g. perceived rigidness in our example) between objects that deform in surprising (popping cactus) and expected (popping balloon) ways– because in either case the strongest cue is used (i.e. the visual sensory input which shows the cactus is popping and thus really not that rigid). If, however, a down-weighing of the visual input occurred, the visual prior (on the typical rigidness of cactii) would exert an influence when observers rated the rigidness of the popping cactus and we would find they are rated as somewhat more rigid than popping baloons. The Bayesian framework predicts another result:when expectations are violated, as in the exploding cactus scenario, the visual system needs to update its internal model of the world (the generative model, e.g. see Kersten et al., 2004), in order to minimize the prediction error for future tasks. This updating (or prediction error correction) is thought to be a reiterative process, which may take time to complete. In fact, a recent study by (Urgen & Boyaci, 2019) models an individuation task and shows that for surprising trials, the model predicted longer duration thresholds. By analogy, we expect to find that it takes observers longer to perform perceptual tasks when judging material attributes of surprisingly deforming objects (e.g. popping cactii) than ones that deform as expected (e.g. popping balloons).

In this study, we use a violation of expectation paradigm to investigate how predictions about object deformations based on object knowledge influence how we perceive the material of an object. In order to manipulate object familiarity, we rendered two types of objects: Familiar objects, for which there exist strong predictions about their mechanical material properties, and Novel, unfamiliar objects, for which no strong predictions should exist. For each of these object types, we rendered two motion sequences that showed how objects deformed when being dropped on the floor: an ‘Expected’ sequence, in which objects’ deformations were consistent with the observer’s prior beliefs about the mechanical properties of the material, and a ‘Surprising’ sequence in which they were not. The Novel objects inherited their optical and motion properties from a corresponding Familiar object. Participants rated the objects in these movies on various material attributes. While we expected to find differences in ratings and response times between Expected and Surprising events for Familiar objects (see above), we also expected to find differences - although attenuated - for Novel shapes. Note that we expected *some* rating differences between Expected and Surprising deformations also to occur for the unfamiliar, Novel shapes due to their optical cues alone and/or the fact that they all were

bounded 3D objects, which might have elicited a prior expectation (e.g. a rigidity prior) about these objects' mechanical material properties. We included two additional experimental conditions in which observers rated the same attributes on static images that showed each familiar and novel object during the first and last frame of the motion sequence, in order to explicitly assess for each object (familiar and novel) existing priors on material properties, as well as the influence of shape recognizability (i.e. after the object had deformed) on ratings, respectively.

METHODS

Stimuli

Objects. We used two types of objects: Familiar and Novel. Figure 3 gives an overview of all objects used in the experiment. In order to create a stimulus set with broad range of typical material classes, we choose 15 Familiar objects belonging to one of 5 material mechanics: non-deforming (wooden chair, metal key, metal spoon), wobbling (red and green jell-O, custard), shattering (wineglass, terracotta pot, porcelain teacup), wrinkling (linen, velvet, and silk curtain) and splashing (milk, honey, and water drops). All objects were rendered (Blender 2.77a, Stichting Blender Foundation, Amsterdam, NL) with their typical optical material properties, e.g. metallic-looking key, green transparent jell-O, or a silky appearing curtain). Objects were located in a room that had brown walls and a hardwood, polished floor, and were illuminated by the Campus" environment map (Debevec, 1998). The 3D meshes for of the Familiar objects were obtained or adapted from TF3DM (www.tf3dm.com) and TurboSquid (www.turbosquid.com); the remaining were created by hand. Unfamiliar shapes ('Glavens', Phillips, 2004) inherited their optical and mechanical properties from the corresponding Familiar object.

Figure 3. Stimuli Expected Motion. *Shown are all 15 familiar (top) and corresponding novel objects (bottom) used in the experiments. Objects are organized according to their typical material kinematics, i.e. how they deform under force. Note that individual scenes are scaled to maximize the view of the object (First frame, left columns), or in order to give a good impression of the material kinematics (Last frame, right columns).*

Deformations. For each object, we rendered short movies that showed an object falling from a height and interacting with the ground . In order to manipulate surprise in our experiment, an object could either behave as expected, e.g. a glass would shatter, or it could inherit the mechanical material behavior from another object, e.g. a milk drops would stay rigid upon impact (Figure 4B shows two more examples of surprising material 'behaviors', also see Supplementary Figure 1). We created corresponding Expected and Surprising movies for Novel objects (Figure 3, Figure 4C & Supplementary Figure 1). Table 1 shows which objects inherited which material mechanics in the Surprising condition. All experimental movies can be downloaded at https://doi.org/10.5281/zenodo.2542577.

Table 1. Overview of objects and conditions in the experiment. *Columns show the five categories of typical material mechanics (material class) in the experiment; rows show which objects occurred in the Expected and Surprising conditions. Every object in the Expected condition would also appear exactly once in the Surprising condition, where it would be rendered with very atypical mechanical properties, e.g. the silk curtain would splash (see*

Figure 3 for corresponding renderings). Novel objects (Glavens) had no familiar shape but had the same optical and mechanical material properties as the Familiar objects in this table.

	Non-deforming	Wobbling	Shattering	Wrinkling	Splashing
Expected	Painted chair	Green jell-O	Clay pot	Linen curtain	Honey droplets
	Brass key	Red jell-O	Wine glass	Silk curtain	Milk droplets
	Metal spoon	Custard	Teacup	Velvet curtain	Water droplets
Surprising	Milk drops	Wine glass	Custard	Green jell-O	Painted chair
	Red jell-O	Honey droplets	Linen curtain	Metal spoon	Clay pot
	Velvet curtain	Brass key	Water droplets	Teacup	Silk curtain

Animations. Each movie consisted of 48 frames, depicting an object suspended in air, which then fell to the ground. Impact occurred exactly at the at the 11th frame for all objects. The largest extent of the objects in the first frame varied between 6.91 (clay pot) and 12.6 (spoon) degrees visual angle. The largest extent of the objects in the last frame depended on the deformation, but varied between 48.85 degrees visual angle for shattering/splashing items and 4.29 degrees visual angle for rigidly falling items. Object deformations were simulated using the Rigid Body, Cloth, and Particles System physics engines in Blender. For technical specifications about the Particle System simulations, we refer the reader to the parameters listed in (Schmid & Doerschner, 2018).

Apparatus

The experiment was coded in MATLAB 2015a (MathWorks, Natick, MA) using the Psychophysics Toolbox extension (version 3.8.5, Brainard, 1997; Pelli, 1997; Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, 2007), and presented on a 24 5/8" PVM-2541 Sony (Sony Corporation, Minato, Tokyo, Japan) 10-bit OLED monitor, with a resolution of 1024 x 768 and a refresh rate of 60 Hz. Videos were played at a rate of 24 frames per second. The participants were seated approximately 60 cm from the screen.

Task and Procedure

Main (Motion) Experiment: Observers were asked to watch a short video clip to the end and then to rate the object they saw on one of four attributes, as quickly (but as accurately) as they could. We choose the attributes such that they would capture some aspect of the mechanical material qualities of the objects. For example, a splashing object is likely to be rated as very *liquid*, and a non-deforming object not, a wiggling object is likely to be rated as very *wobbly* but a shattering object not. In order to familiarize observers with the rating task, the use

of the slider bar and the keypresses, they completed four practice trials with two objects that did not occur in the actual experiment and with two rating adjectives that also did not occur in the experiment (e.g. rate how shiny this object is).

The experiment was organized into four blocks, one block per attribute. Before the block started, the observer was familiarized with the rating question and then proceeded with a button press to start the trials. On every trial, a reminder of the question of this block remained at the top of the screen, e.g. '*hard*' (for 'How hard is the object?'), together with the first frame of a movie, which was held static for three seconds before the movie was played to the end. After this the movie clip repeated 2 more times (without the hold at the beginning), if necessary. Participants were asked to first watch the video until it finished (i.e. the first play through) and then to rate the object as quickly as they could (while still maintaining accuracy). They indicated their rating by using the mouse to adjust the height of a slider bar placed on the right side of the screen (Figure 4A). A zero setting indicated the absence of an attribute, e.g. not *gelatinous* at all, while a maximum setting would correspond to the subjective maximum value of an attribute. The trial was completed when the observer pressed the space key on the keyboard, after which time the next trial would immediately begin. Response time was measured from the beginning to the end of a trial (between spacebar presses). The slider position of the previous trial was carried over the new trial in order to give the experimental interface a more natural feel to it.

Participants completed 240 trials in total (2 surprise conditions (Expected, Surprising) × 2 object familiarity (Familiar, Novel) × 4 attributes × 15 objects). Surprising condition and object type were the two relevant manipulations in the experiment. While the order of blocks was the same for all observers (*hard*, *gelatinous*, *heavy*, *liquid*), the trial order in each block was randomized.

First Frame/Last Frame Experiment: The tasks, setup, and procedures were identical to that of the motion experiment. In contrast to the Motion experiment, in these experiments, the first/last frame of each of the videos was held on-screen for the same duration as a single presentation of the video.

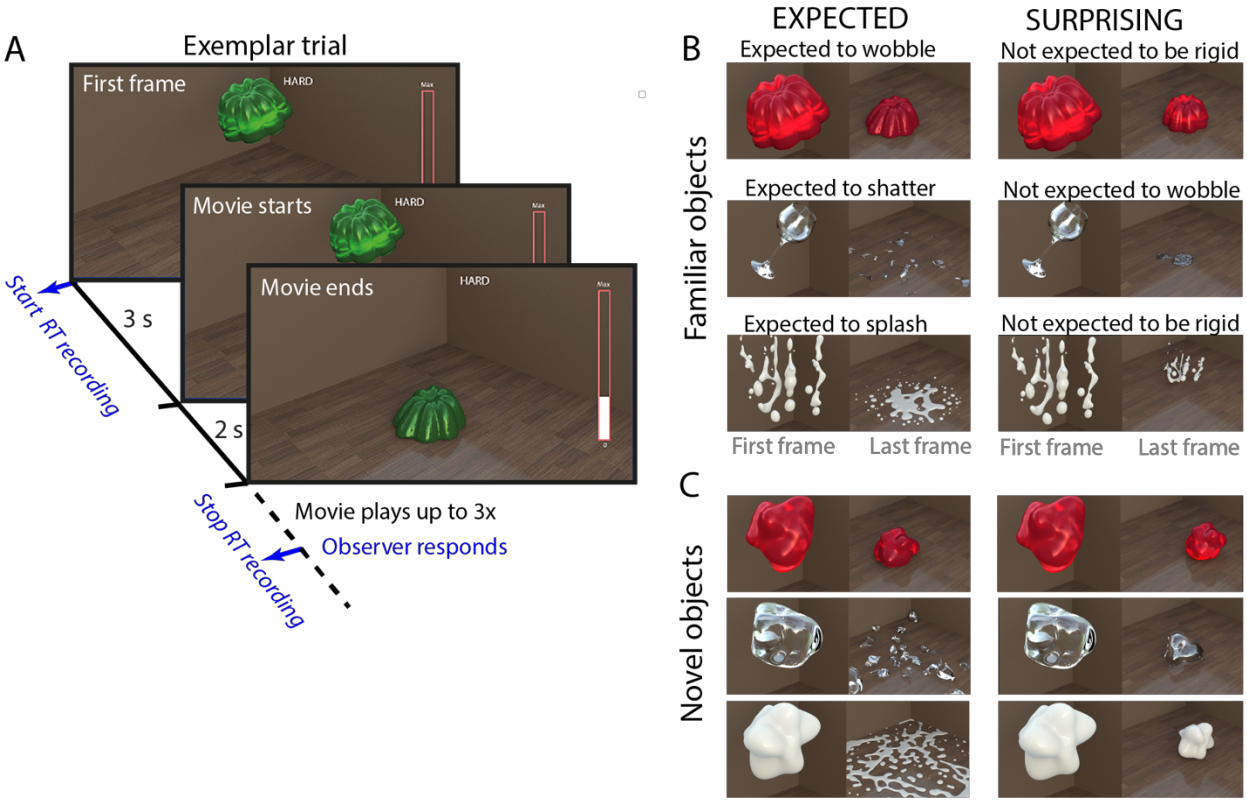

Figure 4. Trial and stimuli. A. *An exemplar trial.* ***B.-C.*** *A subset of familiar objects (B) and corresponding novel objects (C) used in the experiments. Familiar objects could either behave as expected, or in a surprising manner. Note that this distinction (expected vs surprising) is only meaningful for familiar objects. Note that individual scenes are scaled to maximize the view of the object (First frame), or to give a good impression of the material kinematics (Last frame). Figures 3 (Expected condition) and Supplementary Figure 1 (Surprising condition) show corresponding views for the entire stimulus set.*

Participants

Motion Experiment: 25 participants (mean age 24.8; 18 female) participated in the experiment; 23 were right-handed and all had self-reported normal or corrected-to-normal vision.

First Frame Experiment: 15 participants (mean age 26.40; 13 female) participated in this experiment; 13 participants were right-handed and all had self-reported normal or corrected-to-normal vision.

Last Frame Experiment: 14 participants (mean age 26.85; 10 female) participated in the 'Static Last Frame' condition of Experiment 1; 12 were right-handed. All participants had self-reported normal or corrected-to-normal vision.

All participants were native German speakers, and the experiment was given entirely in German. The experiment followed the guidelines set forth by the Declaration of Helsinki, and participants were naïve as to the purpose of the experiment. All participants provided written informed consent and were reimbursed at a rate of €8 per hour.

Analysis.

Rating differences: Expected versus Surprising

To measure the influence of object knowledge on perceived material properties, we computed rating differences between Expected and Surprising conditions (Figure 4B & C). For each object type (Familiar, Novel), each attribute (*how hard, how gelatinous, how heavy, how liquid*), and each material class (Splashing, Shattering, Non-deforming, Wobbling, Wrinkling), rating differences were calculated between the three objects with Expected outcomes and the three (different) objects with Surprising outcomes (ratings were averaged over the three objects before difference scores were computed). Thus, *how* an object deformed was the same, e.g. it would splash, but *which* object (Familiar or Novel) would do the (e.g.) splashing is different in Expected (honey, milk, water) and surprising conditions (chair, pot, curtain).

Initially we were agnostic about the direction in which the rating will change (e.g. will a Splashing chair look more *liquid*, or *harder*, or *heavier* than Splashing milk?), so we took the absolute value of this difference score as an index of the effect of expectation on ratings (Effect of Expectation Index, $\boldsymbol{\epsilon}$):

$$\boldsymbol{\epsilon} = |\ \textbf{av. ratings Expected} - \textbf{av. ratings Surprising}\ | \qquad \textit{Equation 1}$$

This resulted in 40 difference scores for each observer, 20 in the Familiar object condition and 20 in the Novel object condition. If object knowledge influences perceived material properties, then ratings should differ between Familiar objects that behave as expected, and Familiar objects that behave in a surprising way, i.e. $\boldsymbol{\epsilon}$ for Familiar objects should be larger than $\boldsymbol{\epsilon}$ for Novel objects ($\boldsymbol{\epsilon}_{\text{Familiar object}} > \boldsymbol{\epsilon}_{\text{Novel object}}$). A binomial sign test was used to compute the likelihood of obtaining *k* or more instances in which $\boldsymbol{\epsilon}$ (averaged across observers) was greater for Familiar versus Novel objects.

Prior Pull

The four attributes that observers rated in our experiments are related to the kinematic properties of objects. Such properties can be best estimated when we observe how object

interactins with another one, e.g. if two objects collide and one deforms more that the other we can tell that one object is softer. Conversely, it should be difficult to judge such properties (hardness gelatinousness etc) in static frames – unless we rely on our previous experiences observing how object interact. Therefore, one could use the First Frame ratings of Familiar objects as a measure of *prior knowledge* about the material qualities from familiar shape and optics associations. Ratings of moving Novel objects, on the other hand, could be used as a measure of how much the image motion, generated by the kinematics of the material ('sensory' route), influences the rating (no influence of shape, equating for the effect of familiar optical properties). These two conditions make similar predictions (i.e. yield similar ratings) when the material behavior is expected but make *different* predictions (i.e. yield different ratings) when the material behavior is surprising. We operationally define and measure the "prior pull" as the distance between Familiar and Novel object motion ratings in the direction of First Frame ratings. We assess whether prior pull occurs more frequently in the Surprising motion condition compared to the Expected motion condition.

Linear Regression Models

We developed two models with the aim to account for the differences (magnitude and direction) we observed in ratings of familiar and novel moving objects, e.g. differences between the green jell-O Familiar Object and the corresponding green jell-O Novel Object - in both expected and surprising conditions. Low- and high-level models had each exactly 3 predictors, for each of which we provide the rationale and computation below.

High-level Model: Each predictor in the high-level model was multiplied by a weight G that took into account the *predictability* of the "behavior" of a given stimulus (computed as the difference first frame and motion ratings in the Expected condition) and the *extremeness* of the rating (computed as the twice the difference between first frame ratings in the Expected condition and .5). Thus, predictors in the high-level model were:

X1 (=X1*G), where G= *predictability*extremeness*: the *optics & bounded shape prior* (computed as the difference between First Frame ratings of Novel objects in the Expected condition and moving stimuli ratings of Novel objects in the Expected condition). <u>Rationale:</u> This predictor captures to what extent the reflectance properties and the fact that observers see a bounded shape will lead to prior expectations about the material kinematics.

X2 (=X2*(1-G)): the *prior-pull* (computed as the overall differences between First frame ratings of Familiar objects and corresponding ratings of moving Novel objects (across both expected and surprising conditions). <u>Rationale</u>: This predictor captures to what extent

expectations generated by the familiar shape of the object can explain the rating differences, not controlling for the effects of optics and bounded shape (which are captured by X1).

X3 (=X3*(1-G)): *last frame shape recognizability* (computed as the overall differences between Last frame ratings of Familiar and Novel objects (across both, expected and surprising conditions). Rationale: The reason for including this predictor is our hypothesis that the conflict between prior and surprising motion should be larger, if the object would be still recognizable after falling to the ground (as in the wobbling or rigid motion conditions, see Table 1.)

The prediction success of this high-level model was compared with a **Low-level Model** that had the following predictors:

X1: *motion energy differe*nce (this was computed in two steps): 1. Take the sum of the absolute value of all consecutive image differences, starting with the impact frame, e.g. sum(abs(f11-f12), abs(f12-f13), …, abs(f39-f40)), (Tschacher, Rees, & Ramseyer, 2014), for all experimental conditions/stimuli (60 in total) and normalize these values by object size (number of pixels corresponding to the object on frame 11 ME_n), 2. For each object in the expected and surprising conditions, compute the differences of $ME_{Familiar}$ and ME_{Novel} conditions. Rationale: Familiar and Novel objects were often different in 3D volume and shape, thus differences in ratings might be attributable to resulting differences in motion energy when these two classes of objects deform.

X2: *first frame object size* (computed as number of pixels in the first frame corresponding to the object). Rationale: Differences in the area taken up by Familiar and Novel objects in the first frame might account for rating differences, e.g. if the Familiar object was smaller, it might have generated predictions to be less heavy than the corresponding Novel object (irrespective of the familiar shape of the object).

X3: *last frame object siz*e (computed as numbers of pixels in the last frame corresponding to the object). Rationale: this predictor is related to X3 in the high-level model. Differences in the area taken up by Familiar and Novel objects in the last frame might account for rating differences, e.g. if the novel object 'splashed' more (thus took up more area in the last frame) it might have been rated more liquid simply due to this fact.

Response times

Time taken to make each judgment (response time) was measured. We reasoned that rating the material properties of materials that behave surprisingly might involve the reiterative correction of a prediction error by the visual system, and this error correction might be associated with an increase in response time when rating objects that behave in a surprising

way. Before computing the difference in response time between Expected and Surprising conditions, we pre-processed response time data as follows: we subtracted the time to impact (3 seconds static first frame + 0.45 seconds to impact) from the raw response times so that a response time of zero would now indicate time of impact. Data points that were faster than 0.75 seconds after impact (fastest possible button press) were excluded. Response latencies that were longer than 2 standard deviations above the mean were also excluded. Following these exclusions, approximately 6% of the data were excluded for response times that were too fast or too slow according to this criterium (~1.3% too fast, ~4.7% too slow).

Exclusions

For the first 10 participants in the motion experiment, data points for the Novel object matched to the Surprising key object were excluded due to a stimulus presentation error that showed the rigid version of the object instead of the wobbling one, (10 subjects × 4 attributes = 40 data points excluded). For the remaining 15 subjects this error was fixed.

Data points that were faster than 0.75 seconds after impact (fastest possible button press) were excluded. Response latencies that were longer than 2 standard deviations above the mean were also excluded. Following these exclusions, approximately 6% of the data were excluded for response times that were too fast or too slow according to this criterium (~1.3% too fast, ~4.7% too slow, see Supplementary Figure 2 for a breakdown of exclusions per experimental condition).

RESULTS

Ratings and the Effect of Expectation (ϵ)

Overall, observers made broadly intuitive responses in that the ratings of observers varied systematically with the mechanical material properties of Familiar and Novel objects: Figure 5, for example, shows that non-deforming (yellow bars) and shattering (orange) objects were rated as hard (top plot), and not *gelatinous* (bottom plot). Unsurprisingly, objects that splashed (yellow bars), wobbled (green bars), and wrinkled (purple bars) were rated less hard. Finally, objects that wobbled (green bars) were rated as very *gelatinous*. Overall, Familiar and Novel objects tended to be rated similarly.

There were a few deviations from this overall pattern (also see Supplementary Figure 3 for all results), i.e. that ratings for individual Familiar objects were either substantially above or below their group mean: e.g. in the Surprising condition, *hardness* and *gelatinousness* ratings

for the non-deforming red jell-O tended to be lower and higher than the group average, respectively. Similarly, ratings for the wobbling glass and the wobbling key appear to be outliers. We follow this up in the regression analysis below, allowing for different high and low-level factors to explain rating patterns.

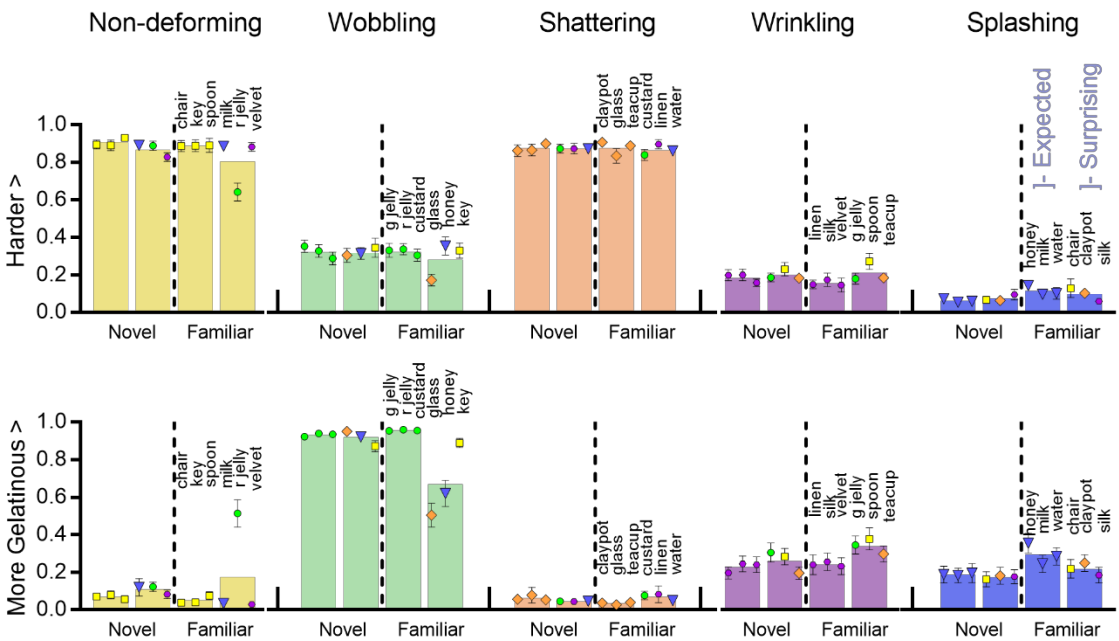

Figure 5. Ratings of two attributes: how hard (top) and how gelatinous (bottom) averaged across participants. *A higher value on the y-axis indicates a higher rating for that attribute. The different colors show the results for each material class (i.e. how the objects behaved on impact with the ground). Within each group, Novel object ratings are shown on the left of the dotted line, and Familiar object ratings (labelled) are shown on the right. Bars show the average rating of the three objects in each condition, and symbols show ratings for each object. The symbol color and shape refer to the outcome that would be expected for that object (which differ from the bar color in the Surprising conditions): yellow squares would be expected not to deform, green circles would be expected to wobble, orange diamonds would be expected to shatter, purple hexagons would be expected to wrinkle, and blue triangles would be expected to splash. Error bars are one standard error of the mean. See Supplementary Figure 3 for all results.*

Figure 6 plots the Effect of Expectation Index (ϵ) for each attribute and each material class, averaged across subjects. In all but one condition (19/20), ϵ was greater for Familiar objects than Novel objects, sign test: $p<0.001$. Results from a paired *t*-test corroborate the finding that on average, ϵ was greater for Familiar objects than Novel objects, $t(24)=6.70$, $p<.001$, which

also verifies the effect at an individual level. This supports our hypothesis that judgments of material qualities are not based purely on the observed material mechanics, but are also affected by prior knowledge about the typical material mechanics of a (familiar) object. Note that the effect of expectation index for Novel objects was not zero, which suggests that observers might generate some predictions about the mechanical properties based on the estimated optical properties of these objects, and/or based on associations they have with the general material properties of a bounded convex shape (e.g. a rigidity prior (Ullman, 1976; Grzywacz & Hildreth, 1987; but also see Jain & Zaidi, 2011). Below we break down these specific influences on material judgements and based on this develop a linear regression model that tries to account for the observed rating differences in Expected and Surprising conditions.

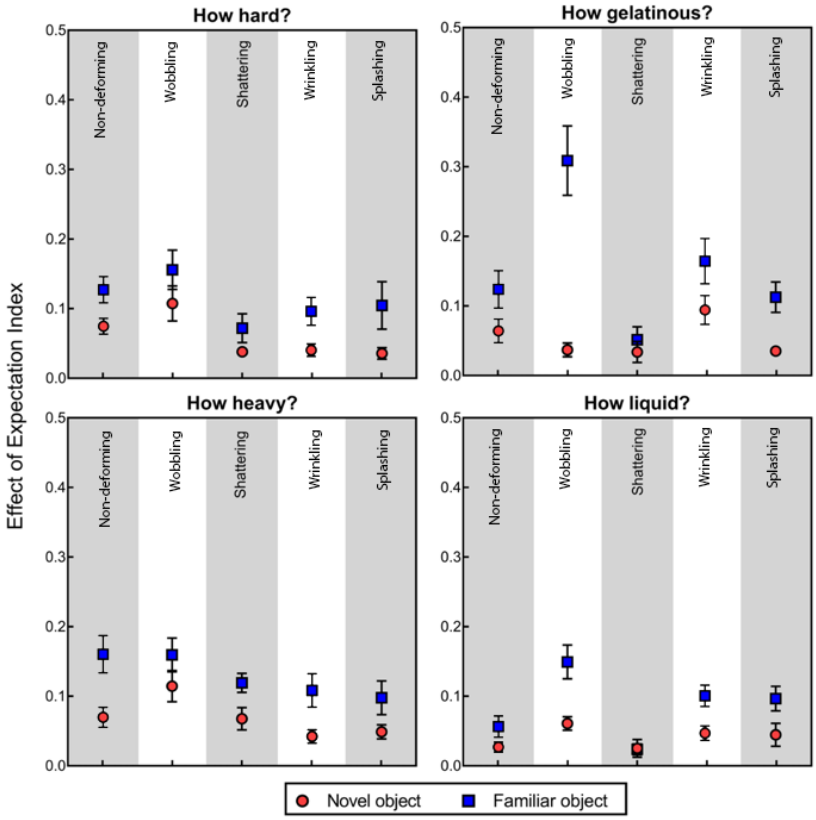

Figure 6. Effect of Expectation Index (ϵ) averaged across subjects. ϵ *was calculated as the absolute difference between ratings in the Expected and Surprising conditions (see analysis section, Equation 1). Error bars are one standard error of the mean.*

Prior Pull

In some cases, the directionality of the rating differences (Expected vs Surprising) is directly interpretable. For example, the jell-O, spoon and teacup that wrinkled surprisingly (Figure 5, purple bars) were rated as *harder* (top plot) than the linen, silk, and velvet curtains that wrinkled

expectedly. Prior knowledge about spoons and teacups (and even jell-O relative to cloth) being hard seems to have led to increased ratings of *hardness* compared to their soft curtain counterparts, despite all of these objects wrinkling. Thus, our results suggest that prior object knowledge about *hardness* "pulled" ratings of *hardness* towards this expectation. We visualize this 'prior pull' in the data in Figure 7.

Even though we plot data from expected (Figure 7A) and surprising motion conditions (Figure 7B) by our experimental design, a meaningful prior 'pull' occurs only when the ratings of atypically-moving (Surprising) Familiar (light green dots) and Novel objects (small, dark green dots) do not overlap and when ratings of atypically-moving (Surprising) familiar are pulled in the direction of First Frame ratings (significant cases highlighted in yellow; also see Table S1). If, instead, ratings of moving Familiar and Novel objects overlap completely, the object shape prior did not exert any significant influence on the rating (see Figure 7C for an illustration). Prior pull can only occur if ratings on a given quality in First frame (large green dots) and Surprising motion experiments differ substantially, as indicated by the length of the dark gray lines between these two types of data points in Figure 7. We would expect to find this occurring more frequently in the Surprising motion condition, and our results confirm this: First Frame ratings differed significantly from Novel object motion ratings for 48 out of 60 (80%) conditions in the Surprising motion condition, and only 24 out of 60 (40%) conditions in the Expected motion condition (also see Table S1). Overall, the more the familiar object remained intact in the Surprising motion condition (Figure 7B), the more likely the prior exerted an influence over the material appearance.

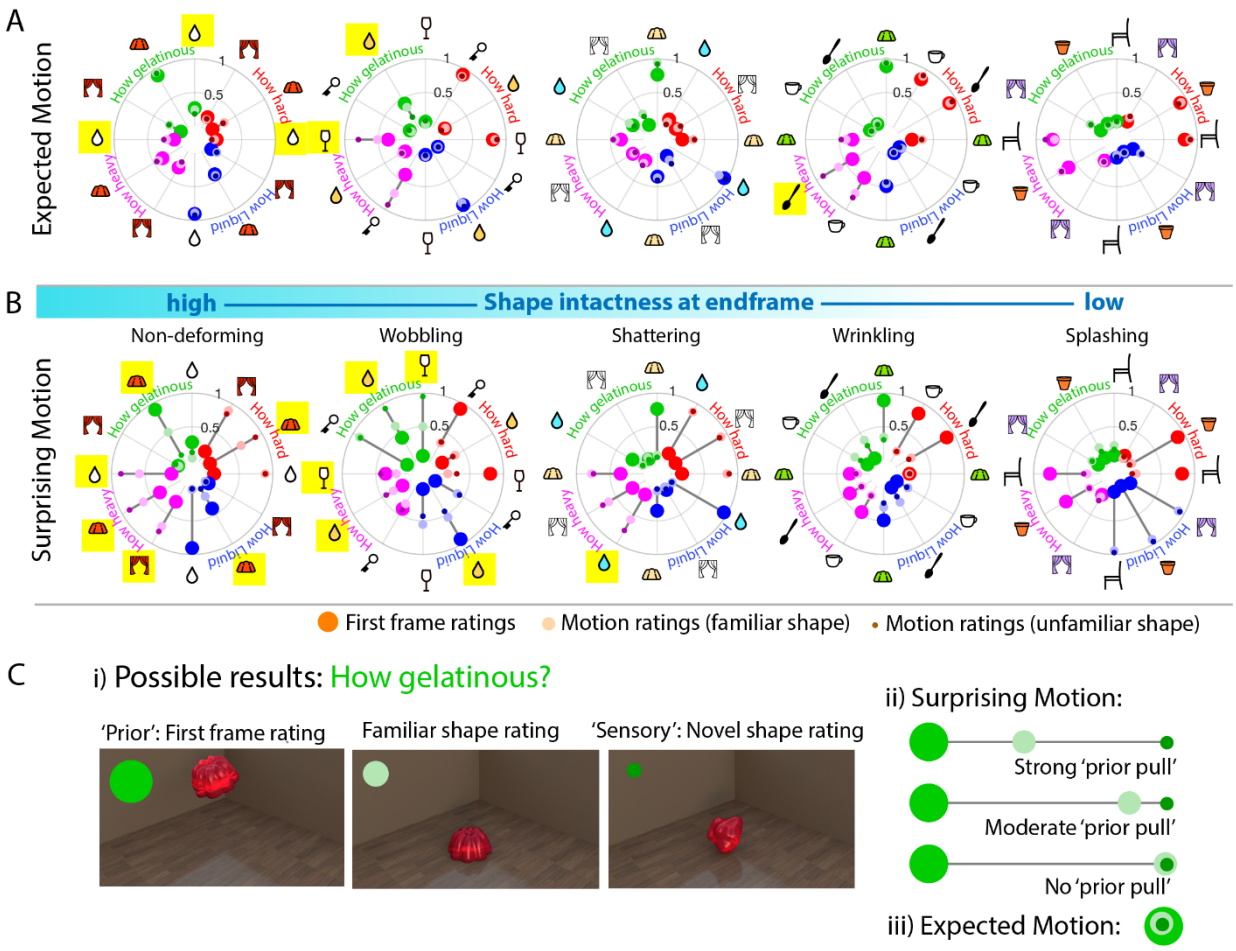

Figure 7. Material quality ratings and prior attraction. A. *Average observer ratings for different questions about material qualities e.g. how hard, liquid, heavy or gelatinous an object appears. Icons symbolize individual Familiar objects (chair, key, cup, pot, glass, spoon, blue droplet: water, yellow droplet: honey, white droplet: milk, violet curtain: silk, red curtain: velvet, white curtain: linen; yellow custard, red and green jell-O). Each question and corresponding data are coded in the same color (red: hard, blue: liquid, purple: heavy, and green: gelatinous). Ratings could vary between 0 (lowest) and 1 (highest). On average, ratings from all 3 conditions (i.e. First Frame Familiar objects (large dots), typically-behaving Familiar objects (medium dots) and corresponding (moving) Novel objects (small dots)) tended to overlap. Organization of objects follows that of B.* ***B.*** *Same as A, but here, ratings of atypically-behaving familiar objects are plotted as medium-sized desaturated dots (now organized by motion), and ratings of corresponding Novel objects - i.e. unfamiliar shapes-- which inherit their optical and kinematic qualities from a familiar object - as small dark dots.* ***C. i)*** *Illustrates how we measure how much the rating of an atypically-moving Familiar object (middle image) overlaps with the rating of a material-matched moving Novel object (right image), or conversely, how much it is pulled towards ratings of a static view of the Familiar object (left image).* ***ii)*** *Shows possible results: For example, seeing an image of red jell-O in*

its classical shape, observers tend to expect that it is quite gelatinous. When they see an object with the same optical properties that falls and does not wobble when it hits the floor, they rate it - unsurprisingly - as very non-gelatinous. When a classically-shaped red jell-O falls on the floor and doesn't wobble, observers could either rate it similar to the novel object -after all it doesn't wobble at all - or it could be rated as somewhat more gelatinous, despite the sensory input, possibly because prior experience influences the appearance, making observers perceive wobble when there isn't. ***iii)*** *When the familiar object moves exactly as expected, and when there is no strong influence of shape familiarity on material judgements, all three ratings will overlap.*

That only a subset of our objects exhibits a "pull" towards the prior is consistent with literature that shows when sensory input is unambiguous, it will dominate the percept, and prior knowledge does not have an effect (Summerfield & De Lange, 2014). In our case, particular combinations of shape and optics (often involving "shape recognizability") at the end of the animation may lead observers to essentially "ignore" sensory information. For example, the rigid jell-O is perceived as gelatinous despite not wobbling. Interestingly, prior associations with specific combinations of motion, shape, and optics at the end of the movie may "enhance" differences between Familiar and Novel object ratings (e.g. in the case of the wobbling glass, the particular combination of translucent optics forming what looks like a puddle might evoke strong representations of liquidity, so despite the wobble, it is not considered gelatinous).

One might argue that the reason why we see relatively few instances of significant prior pull is that Novel objects might too elicit expectations (after all, they are bounded shapes with particular optical properties). The correlations in Figure 8 between First Frame and motion ratings, however, suggests that, for the most part, this is not the case: while Familiar object first frame ratings predicted Expected motion ratings extremely well (Figure 8A, R^2=0.9), Novel object first frame ratings predict motion ratings very poorly (Figure 8B, R^2=0.07). These results are also supported by the average interobserver correlations in Figure 8C. We can see that this correlation is much lower for novel objects in the First frame experiment than for familiar objects suggesting that static images of Novel objects do not elicit consistent (typical) expectations about kinematic material properties across observers. Thus, for the most part, Familiar object first frame ratings are a good measure of observers' prior expectations about the material qualities of these objects. In contrast, Novel object motion ratings are an ideal measure of sensory input because the motion from the material behavior includes information

about both the mechanical deformations, as well as any additional effect of image motion from optics (Doerschner et al., 2011).

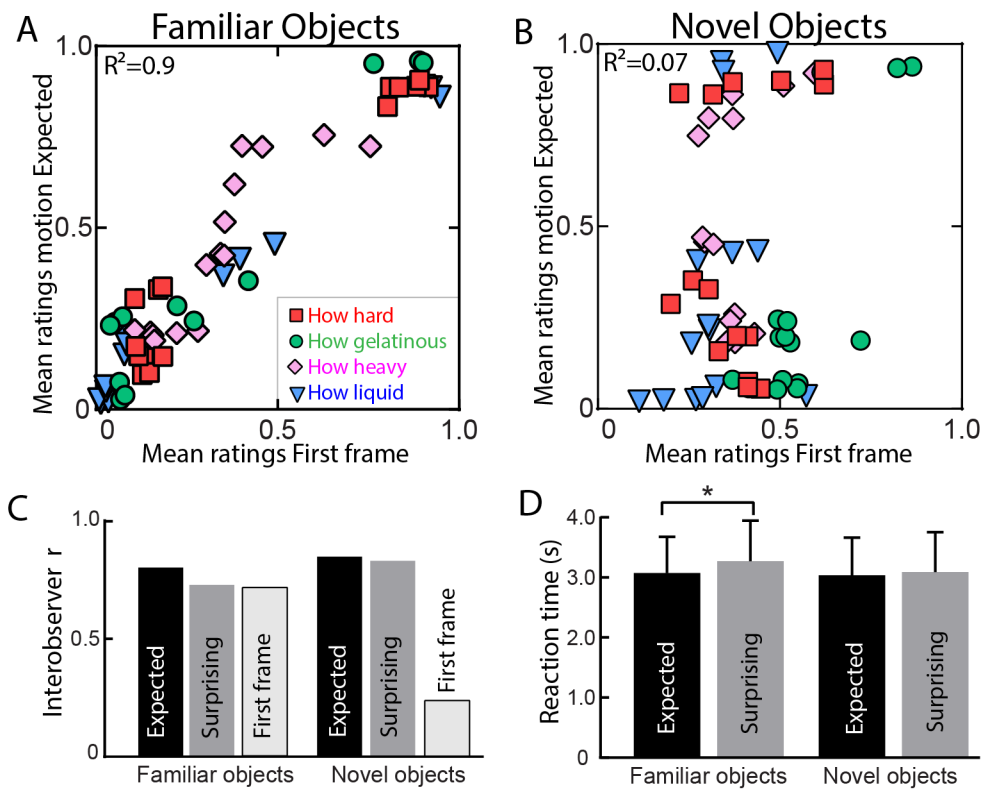

Figure 8. Prediction strength response time differences and interobserver correlations. A. *Correlation between mean First frame ratings and mean Expected motion ratings for Familiar and Novel objects (**B.**) A high correlation indicates that the First frames (still images) of objects are highly predictive of the objects' kinematic properties, and thus are in good agreement with ratings in the Expected motion condition, where objects fall and deform according to their typical material kinematics. This is clearly not the case for Novel objects, suggesting that these objects do not elicit strong prior expectations about how an object will deform.* ***C.*** *Average interobserver correlation for Expected and Surprising motion trials, as well as the First frame experiment. Note that only for Novel objects, this latter correlation was quite low, suggesting that still images of unfamiliar objects do not elicit a strong prior in observers about the material qualities measured in this experiment.* ***D.*** *Response time data averaged across all observers for Expected (black) and Surprising trials (medium gray). Stars indicate significant differences $p<0.005$.*

Linear Regression Models

Given that there are a few cases where Novel objects do seem to generate correct predictions about the material outcome (those at the bottom left and top right of Figure 8B), and since some of the magnitude of the prior pull may be explained by shape recognizability at

the end of the movie, we tested a linear regression model that predicted the direction and magnitude of Familiar and Novel object rating differences from prior pull from optics, familiar shape, and shape recognizability at the end of the animation (Figure 9). We can model the data moderately well with this *high-level* model (see analysis, R^2 between 0.266 and 0.59 depending on the question), but not perfectly, potentially due to specific motion-shape-optics interactions (Schmid & Doerschner, 2018). We also tested whether *low-level* image differences between Familiar and Novel objects could be driving differences in ratings. To this end we modelled the data using differences in size between Familiar and Novel objects in the first and last frames (pixel area difference, see analysis section) and calculated differences in motion energy. Such a model performs extremely poorly ($R^2 = 0.025$, $p>0.05$, Figure 9). This suggests rating differences are not caused by differences in object size or image motion.

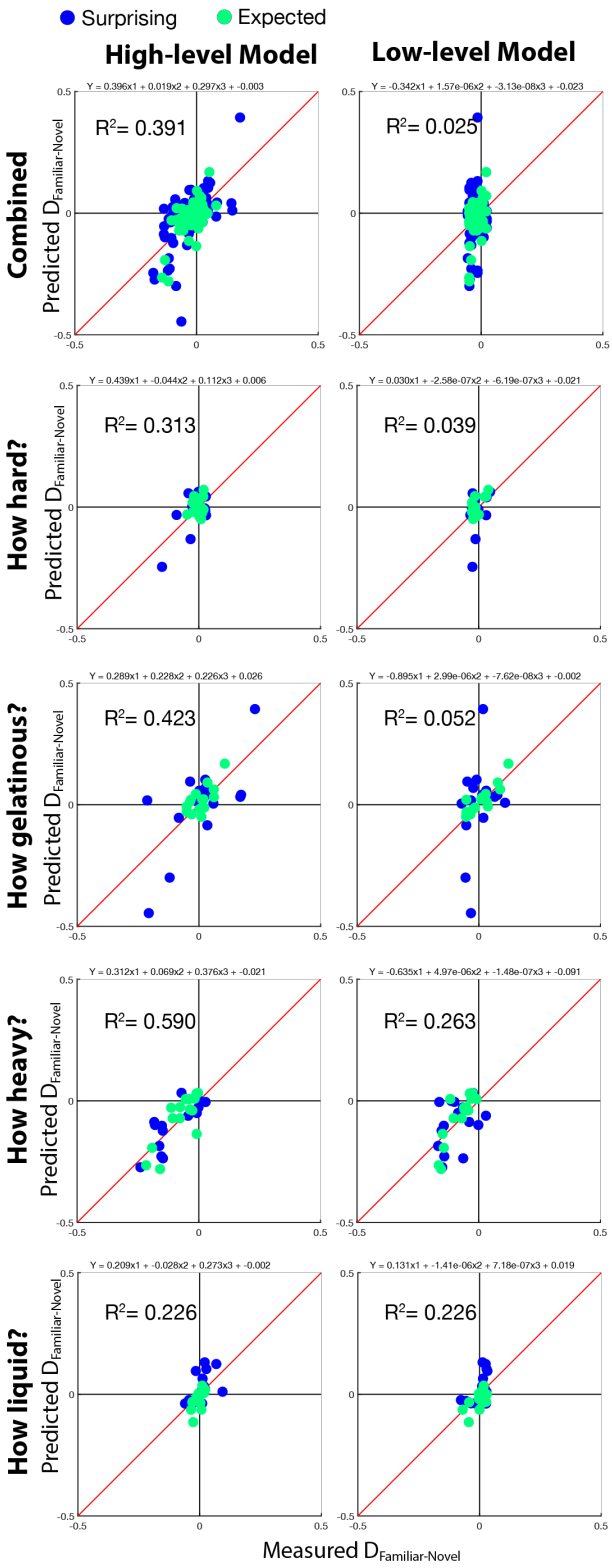

Figure 9. Low and high level linear regression models predicting the difference between ratings of moving familiar and control objects $D_{familiar\text{-}novel}$. *We developed two models with the aim to account for the differences we observed in ratings of familiar and novel moving objects. Overall, the high-level model (left side) was more successful in predicting this difference than the low-level model (right side). However, this varied as a function of questions.*

The high-level model performed best for ratings of gelatinousness and hardness, whereas the low-level model performed as good as the high-level model for ratings of liquidness. This makes sense since ratings of the latter might be dominated by how much a substance physically spreads in the image, whereas the former requires more high-level inference. Low- and high-level models had each exactly 3 predictors.

Response times

We find a small but significant increase in RT in the Familiar object surprising condition - which is the condition that most strongly juxtaposes prior expectation with sensory evidence (Figure 8D). There was no significant response time difference between expected and surprising Novel objects conditions.

DISCUSSION

Visual perception is not a one-way (bottom-up) road; how we process visual input is influenced by expectations about the sensory environment, which develop from our previous experience and learning about existing regularities in the world, i.e. associating things or events that co-occur. Expectations have been shown to facilitate visual processing in the case of priming, to modulate the frequency of a particular percept in bi-stable stimuli, and to change our interpretation of ambiguous stimuli (see Kveraga et al., 2007, or Panichello, Cheung, & Bar, 2013 for a review). However, the stimuli used in these experiments have been fairly simple (static images of objects), and it has been shown that learning associations can also include fairly complex phenomena. For example, recently, Bates, Yildirim, Tenenbaum, & Battaglia, (2015) (also Kubricht et al., 2016; Battaglia, Hamrick, & Tenenbaum, 2013), showed that humans can learn to predict how different liquids flow around solid obstacles (also see other examples for predicting of motion trajectories of rigid objects, e.g. (Flombaum, Scholl, & Santos, 2009; Gao & Scholl, 2010, or Soechting, Juveli, & Rao, 2009). While the authors attributed human performance to an ability to "reason" about fluid dynamics, here we explicitly test whether existing *perceptual expectations* about material properties can set up rather complex predictions about future states, and whether – and to what extent – these expectations influence material appearance. We show that the qualities of "surprising" materials (Itti & Baldi, 2009a) are perceived different to expected ones that behave the same (Figure 7), and that surprise leads to increases in processing time of the stimuli (Figure 8D, in line with Itti & Baldi, 2009b; Baldi & Itti, 2010; or recently Urgen & Boyaci, 2019). Our method provides a general technique to differentiate the extent to which material qualities are directly estimated from

material kinematics versus being modulated by prior associations from familiar shape and optical properties, and contributes thus an essential piece to the puzzle of how the human visual system accomplishes material perception (Adelson, 2001; Maloney & Brainard, 2010; Anderson, 2011; Fleming, Nishida, & Gegenfurtner, 2015; Fleming, 2014; Komatsu & Goda, 2018)

Knowledge affects (material) perception

There are countless demonstrations showing that knowledge affects how we perceive the world: from detecting the Dalmatian amidst black and white blotches or an animal in the scene (Thorpe, Fize, & Marlot, 1996), deciding on the identity of a blob by (Oliva & Torralba, 2007), to being a greeble or bird expert (Gauthier, James, Curby, & Tarr, 2003). Object knowledge does not just facilitate categorical judgments, it also affects the estimation of visual properties such as color or motion (Hansen et al., 2006; Olkkonen & Allred, 2014; Scocchia et al., 2013). How exactly knowledge alters and facilitates neural processes in visual perception is a topic of ongoing research (e.g. Rahman & Sommer, 2008; Gauthier, Skudlarski, Gore, & Anderson, 2000; or Kveraga et al., 2007).

In contrast, the role of predictions or associative mechanisms in material perception is not well understood (van Assen et al., 2018; Schmidt et al., 2017; Paulun et al., 2017). Knowledge about materials entails several dimensions and can include taxonomic relations: gold is a metal, metals are elements with physical properties, metals are usually malleable and ductile, as well as perceptual regularities: gold looks yellowish, often has a very shiny, polished and smooth appearance, feels cool to the touch etc. Our experimental results suggest that identifying a material (i.e. knowing what it is) not only co-activates its typical optical qualities, but also elicits strong predictions about the typical kinematic properties and resulting material 'behaviors'. For example, liquids are not only translucent or transparent, they also tend to run down, splash, or ooze. Importantly, we seem to have quite specific ideas of what running down, splashing, or oozing should look like (e.g. Dövencioglu, van Doorn, Koenderink, & Doerschner, 2018 probed such 'ideas' explicitly), supposedly because we have ample visual (but also haptic) experiences with liquids, and thus opportunities to learn the regularities (statistical or other) associated with a specific material category. Our results show that these specific ideas, or 'priors', about 'material behaviors interfere with the bottom-up processing of visual information, leading to predictable differences in ratings of material properties between Expected and Surprising conditions. Thus, we come to a somewhat different conclusion than Sharan, Rosenholtz & Adelson (2013). While these authors acknowledged that 'the relationship

between objects and materials is not arbitrary' (Sharan et al. 2013), they later suggest that 'material categorization is not merely derived from high-level object recognition' (Sharan et al. 2013). The basis for this conclusion is an apparent dissociation of object vs material categorization performance at short stimulus presentation times. We believe our experimental paradigm might be more suitable to directly measure how object knowledge affects material perception. Interestingly, Sharan et al. (2013) find a substantial decrease for the material categorization tasks at very short presentation times. Given our results that show response time increases in the surprise condition for Familiar objects, it would be interesting to know if this decrease is primarily driven by the surprising stimuli in their experiment, i.e. 'fake' ones – since for those, the object shape/identity (a natural/edible thing) was in conflict with the material of the object (plastic/inedible).

Prior pull

The yellow highlighted cases in Figure 7 show that "prior pull" occurred twice as much in the Surprising than the Expected condition, and more in conditions where the object was still intact and recognizable at the end of the movie (objects that behaved rigidly or wobbled). Prior pull in the Expected condition also occurred where estimation (sensory input) and associative (prior knowledge) accounts made different predictions (gelatinousness of the honey, heaviness of the wine glass, Figure 7A). In these cases, the 'expected' cases were not so expected – this may be related to shape properties, or the size of the splashing of the liquids. Although we controlled for the effects of image motion from optics (e.g. specular highlights), perhaps other low-level image differences exist between Familiar and Novel objects that could be driving differences in ratings. However, we were able to rule this out, since a linear regression model that incorporates these low-level effects as predictors performs extremely poorly. This suggests rating differences were not caused by differences in object size or image motion.

That only a subset of our objects exhibits a "pull" towards the prior is consistent with literature that shows when sensory input is unambiguous, it will dominate the percept, and prior knowledge does not have an effect (Summerfield & De Lange, 2014). In our case, particular combinations of shape and optics (often involving "shape recognizability") at the end of the animation may lead observers to essentially "ignore" sensory information. For example, the rigid jell-O is perceived as gelatinous despite not wobbling. Interestingly, prior associations with specific combinations of motion, shape, and optics at the end of the movie may "enhance" differences between Familiar and Novel object ratings (e.g. in the case of the wobbling glass,

the particular combination of translucent optics forming what looks like a puddle might evoke strong representations of liquidity, so despite the wobble, it is not considered gelatinous).

We did not aim to test an exhaustive list of material attributes, but to determine whether effects of prior associations on visual input might depend on the type of material attribute judged, and on how the object behaves under external forces. We found that "mechanical" qualities like hardness and liquidity appear to be more directly estimated from material kinematics in "shape destroying" conditions (splashing, shattering, wrinkling), but prior associations play a modulatory role (to the extent where material kinematics can even be ignored (e.g. red and green jell-O) when shape remains somewhat intact. These latter conditions seem to create more of a cue conflict, and are more ambiguous. On the other hand, qualities like gelatinousness and heaviness (which are much more difficult to estimate directly from mechanical deformations) were more affected by familiar shape and optics associations.

Alternative explanations for rating differences

One might argue that the 'prior pull' we demonstrated here is not perceptual but in fact is due to a particular 'cognitive strategy' of some observers (i.e. explicitly ignoring the motion information and thus rating material qualities of atypically moving familiar objects as they rated objects on the first frame, while other observers' ratings were 100% identical to novel object ratings). This would have resulted in bimodal rating distributions and/or low interobserver correlation in the object motion condition, neither of which we found (Figure S4 and Figure 8C, respectively).

Another argument against the cognitive strategy approach is supported by response time (RT) patterns in our experiments. A small but significant increase in RT in the familiar object surprising condition - which is the condition that most strongly juxtaposes prior expectation with sensory evidence - would be consistent with the idea of recurrent prediction error correction (Urgen & Boyaci, 2019, Figure 8D). Importantly, we do not find evidence for a response time advantage in expected familiar objects condition (compared to novel, also consistent with (Urgen & Boyaci, 2019), which suggests that this increase in RT cannot simply be due to the fact that observers positioned the slider in advance to the 'wrong' (expected) position. If observers adopted such a strategy, we should have also seen faster RTs for expectedly moving familiar objects.

It is quite striking that the same material deformations were rated differently in the Expected and Surprising conditions. Toscani, Valsecchi, & Gegenfurtner (2013) showed that depending on the task, observers pointed their gaze at specific points at the stimulus, e.g. near

682 the brightest regions on an object for lightness judgements. One possibility could thus be that
683 the task, object knowledge, and expectations about the material 'behavior' guided eye
684 movements of observers in our experiment to specific locations on the stimulus. While in
685 Expected conditions fixation patterns might have been optimal with respect to the task, e.g.
686 observers correctly anticipated how the object would deform (or shatter, splash, wobble etc.),
687 it is possible that in the Surprising conditions the 'wrong' expectation guided eye movements
688 to the 'wrong' locations on the object, which in turn lead to a different sampling of information
689 and ultimately influenced their judgements of material qualities. We are currently investigating
690 this possibility directly.
691
692 **Bayes**
693 We believe that the results of this study fit well within the Bayesian framework, which offers
694 an account of how prior knowledge is integrated with sensory input. Our experiment constitutes
695 a situation not unlike classical cue-conflict experiments (e.g. Ernst & Banks, 2002; Knill, 2007),
696 where the sensory cues may be in conflict with one another and/or the prior belief. While we
697 do not aim to model our results formally, we still believe that this analogy is useful in
698 interpreting our findings. We will first focus on the rating differences that we found in the
699 Expected and Surprising conditions. In the latter, in many instances, ratings were 'pulled'
700 towards the expected material property, not the signaled one. For example, a wrinkling spoon
701 or teacup was rated *harder* than any of the wrinkling cloths. Here, prior knowledge about
702 spoons and teacups being *hard* seem to have led to increased ratings of *hardness* compared to
703 their soft curtain counterparts, despite all of these objects wrinkling. This outcome would be
704 best explained by a so-called down-weighing of the cues (Knill, 2007) to *hardness*, which
705 suggests that the visual system entertains multiple priors (strong and weaker ones) about the
706 state of the world, and that, depending on the sensory input, it adjusts the weights of these
707 priors. This implies that in a cue-conflict situation, the unlikely interpretation of the input does
708 not simply get vetoed down (Landy et al., 1995), but that instead it would factor into the final
709 percept – just as we observed it in our results.
710 When faced with violated expectations, as in our Surprising condition, the visual system
711 needs to update the generative model in order to minimize the prediction error i.e. the error
712 between the expected state and 'measured' state of the world, in order to perform future tasks
713 (Urgen & Boyaci, 2019). Because this updating is a reiterative process, we reasoned that it
714 would take observers longer to perform perceptual tasks when judging material attributes of
715 surprisingly deforming objects. This is exactly what we found. One might criticize that response

times for ratings were much longer than times measured in classic reaction times studies (e.g see a review by Eckstein, 2011 on visual search). However, it is not all that uncommon to consider response times of 2 seconds and longer, as in categorical color perception (Boynton & Olson, 1990; or Okazawa, Koida, & Komatsu, 2011). Note that we treated response time data as conservatively as classic reaction time studies, e.g. by removing data points that were two standard deviations above the mean.

Estimation vs Association

Our work shows that previously acquired object-material associations play a central role in material perception and are much more sophisticated than previously appreciated. Our results also offer an explanation to the seemingly conflicting findings in research that investigated the perception of material qualities of non-rigid objects. While some work proposes that perceived material qualities like softness are strongly influenced by motion and shape cues, which completely dominate optical cues (Paulun et al., 2017), other work showed that *both* optical and mechanical cues affect estimates of viscosity (van Assen & Fleming, 2016) and yet other research concludes that optical properties dominates over image motion and shape cues when judging the stiffness of cloth (Aliaga, O'Sullivan, Gutierrez, & Tamstorf, 2015). Our results suggest that the prior on the material category (e.g. rubber, liquids, cloth) determines how image cues (optical, motion, shape) are weighted and integrated with existing material knowledge to yield a specific material percept. Thus, the conclusions of previous research can be reconciled when the role of familiar shape priors is considered. This study not only shows that existing *perceptual expectations* about material properties can set up rather complex predictions about future states of materials, it also extends a growing theme in the material perception literature that studying the perception of kinematic material qualities can serve as a tool to guide investigations of the neural mechanisms about material properties, as it provides insight into components (high and low level) that make up material perception as a whole (Schmid & Doerschner, 2019).

CONCLUSION

This work shows that the visual system can predict the future states of rigid and non-rigid materials. Such predictions can be activated by the shape of an object and – to a lesser extent - also by the optical qualities of a surface. Understanding how high-level expectations are integrated with incoming sensory evidence is an essential step towards understanding how the human visual system accomplishes material perception.

CONTRIBUTIONS

LA, KD, and AS conceptualized and developed the experiments in this paper. AS created the stimuli and programmed the experiments. AS, LA, and KD developed and performed the analyses. KD, AS, and LA wrote the paper.

ACKNOWLEDGEMENTS

This work has been supported by a Sofja Kovalevskaja Award endowed by the German Federal Ministry of Education. We thank Roland Fleming, Huseyin Boyaci, and Karl Gegenfurtner for helpful discussions and feedback on earlier versions of this manuscript, and Flip Phillips for providing the 'Glaven' 3D meshes.

References

Rahman, A., & Sommer, W. (2008). Seeing what we know and understand: How knowledge shapes perception. *Psychonomic Bulletin & Review*, *15*(6), 1055–1063. https://doi.org/10.3758/PBR.15.6.1055

Adelson, E. H. (2001). *Human Vision and Electronic Imaging VI. Proceedings of the. SPIE* (Vol. 4299). Retrieved from http://persci.mit.edu/pub_pdfs/adelson_spie_01.pdf

Aliaga, C., O'Sullivan, C., Gutierrez, D., & Tamstorf, R. (2015). Sackcloth or silk? In *Proceedings of the ACM SIGGRAPH Symposium on Applied Perception - SAP '15* (pp. 41–46). New York, New York, USA: ACM Press. https://doi.org/10.1145/2804408.2804412

Anderson, B. L. (2011). Visual perception of materials and surfaces. *Current Biology*, *21*(24), R978–R983. https://doi.org/10.1016/J.CUB.2011.11.022

Baldi, P., & Itti, L. (2010). Of bits and wows: A Bayesian theory of surprise with applications to attention. *Neural Networks*, *23*(5), 649–666. https://doi.org/10.1016/j.neunet.2009.12.007

Bates, C. J., Yildirim, I., Tenenbaum, J. B., & Battaglia, P. W. (2015). Humans predict liquid dynamics using probabilistic simulation. In *Cognitive Science Society* (pp. 172–177). Retrieved from https://pdfs.semanticscholar.org/1867/112a8deb73d72d5e08553e7d065346b7308f.pdf?_ga=2.153744048.96345254.1562848276-1034217185.1562848276

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*, *110*(45), 18327–18332. https://doi.org/10.1073/pnas.1306572110

Boynton, R. M., & Olson, C. X. (1990). Salience of chromatic basic color terms confirmed by three measures. *Vision Research*, *30*(9), 1311–1317. https://doi.org/10.1016/0042-6989(90)90005-6

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*(4), 433–436. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/9176952

Chang, S., & Pearson, J. (2018). The functional effects of prior motion imagery and motion perception. *Cortex*, *105*, 83–96. https://doi.org/10.1016/j.cortex.2017.08.036

Debevec, P. (1998). Rendering synthetic objects into real scenes. In *Proceedings of the 25th*

annual conference on Computer graphics and interactive techniques - SIGGRAPH '98 (pp. 189–198). New York, New York, USA: ACM Press. https://doi.org/10.1145/280814.280864

Doerschner, K., Kersten, D., & Schrater, P. R. (2011). Rapid classification of specular and diffuse reflection from image velocities. In *Pattern Recognition* (Vol. 44, pp. 1874–1884). Pergamon. https://doi.org/10.1016/j.patcog.2010.09.007

Doerschner, K., Fleming, R. W., Yilmaz, O., Schrater, P. R., Hartung, B., & Kersten, D. (2011). Visual motion and the perception of surface material. *Current Biology*, *21*(23), 2010–2016. https://doi.org/10.1016/j.cub.2011.10.036

Dövencioglu, D. N., van Doorn, A., Koenderink, J., & Doerschner, K. (2018). Seeing through transparent layers. *Journal of Vision*, *18*(9), 25. https://doi.org/10.1167/18.9.25

Eckstein, M. P. (2011). Visual search: A retrospective. *Journal of Vision*, *11*(5), 14–14. https://doi.org/10.1167/11.5.14

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*(6870), 429–433. https://doi.org/10.1038/415429a

Fleming, R. W. (2014, January 1). Visual perception of materials and their properties. *Vision Research*. Pergamon. https://doi.org/10.1016/j.visres.2013.11.004

Fleming, R. W., Nishida, S., & Gegenfurtner, K. R. (2015). Perception of material properties. *Vision Research*, *115*, 157–162. https://doi.org/10.1016/J.VISRES.2015.08.006

Flombaum, J. I., Scholl, B. J., & Santos, L. R. (2009). Spatiotemporal priority as a fundamental principle of object persistence. In *The Origins of Object Knowledge* (pp. 135–164). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199216895.003.0006

Gao, T., & Scholl, B. J. (2010). Are objects required for object-files? Roles of segmentation and spatiotemporal continuity in computing object persistence. *Visual Cognition*, *18*(1), 82–109. https://doi.org/10.1080/13506280802614966

Gauthier, I., James, T. W., Curby, K. M., & Tarr, M. J. (2003). THE INFLUENCE OF CONCEPTUAL KNOWLEDGE ON VISUAL DISCRIMINATION. *Cognitive Neuropsychology*, *20*(3–6), 507–523. https://doi.org/10.1080/02643290244000275

Gauthier, I., Skudlarski, P., Gore, J. C., & Anderson, A. W. (2000). Expertise for cars and birds recruits brain areas involved in face recognition. *Nature Neuroscience*, *3*, 191.

Retrieved from https://doi.org/10.1038/72140

Grzywacz, N. M., & Hildreth, E. C. (1987). Incremental rigidity scheme for recovering structure from motion: position-based versus velocity-based formulations. *Journal of the Optical Society of America. A, Optics and Image Science*, *4*(3), 503–518. https://doi.org/10.1364/josaa.4.000503

Hansen, T., Olkkonen, M., Walter, S., & Gegenfurtner, K. R. (2006). Memory modulates color appearance. *Nature Neuroscience*, *9*(11), 1367–1368. https://doi.org/10.1038/nn1794

Hillis, J. H., Ernst, M. O., Banks, M. S., & Landy, M. S. (2002). Combining sensory information: Mandatory fusion within, but not between, senses. *Science*, *298*(5598), 1627–1630. https://doi.org/10.1126/science.1075396

Itti, L., & Baldi, P. (2009a). Bayesian surprise attracts human attention. *Vision Research*, *49*(10), 1295–1306. https://doi.org/https://doi.org/10.1016/j.visres.2008.09.007

Itti, L., & Baldi, P. (2009b). Bayesian surprise attracts human attention. *Vision Research*, *49*(10), 1295–1306. https://doi.org/10.1016/j.visres.2008.09.007

Jain, A., & Zaidi, Q. (2011). Discerning nonrigid 3D shapes from motion cues. *Proceedings of the National Academy of Sciences*, *108*(4), 1663 LP – 1668. https://doi.org/10.1073/pnas.1016211108

Kersten, D., Mamassian, P., & Yuille, A. (2004). Object Perception as Bayesian Inference. *Annual Review of Psychology*, *55*(1), 271–304. https://doi.org/10.1146/annurev.psych.55.090902.142005

Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in psychtoolbox-3. *Perception*, *36*(14), 1–16.

Knill, D. C. (2007). Learning Bayesian priors for depth perception. *Journal of Vision*, *7*(8), 13. https://doi.org/10.1167/7.8.13

Komatsu, H., & Goda, N. (2018, November 10). Neural Mechanisms of Material Perception: Quest on Shitsukan. *Neuroscience*. https://doi.org/10.1016/j.neuroscience.2018.09.001

Kubricht, J., Jiang, C., Zhu, Y., Zhu, S.-C., Terzopoulos, D., & Lu, H. (2016). Probabilistic Simulation Predicts Human Performance on Viscous Fluid-Pouring Problem. In *Proceedings of the 38th Annual Meeting of the Cognitive Science Society* (pp. 1805–1810). Retrieved from https://www.seas.upenn.edu/~cffjiang/research/cogsci16/nips16.pdf

Kveraga, K., Ghuman, A. S., & Bar, M. (2007, November). Top-down predictions in the cognitive brain. *Brain and Cognition*. NIH Public Access. https://doi.org/10.1016/j.bandc.2007.06.007

Landy, M. S., Maloney, L. T., Johnston, E. B., & Young, M. (1995). Measurement and modeling of depth cue combination: in defense of weak fusion. *Vision Research*, *35*(3), 389–412. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/7892735

Maloney, L. T., & Brainard, D. H. (2010). Color and material perception: Achievements and challenges. *Journal of Vision*, *10*(9), 19–19. https://doi.org/10.1167/10.9.19

Marlow, P. J., & Anderson, B. L. (2016). Motion and texture shape cues modulate perceived material properties. *Journal of Vision*, *16*(1), 1–14. https://doi.org/10.1167/16.1.5

Okazawa, G., Koida, K., & Komatsu, H. (2011). Categorical properties of the color term "GOLD" *Journal of Vision*, *11*(8), 4–4. https://doi.org/10.1167/11.8.4

Oliva, A., & Torralba, A. (2007). The role of context in object recognition. *Trends in Cognitive Sciences*, *11*(12), 520–527. https://doi.org/https://doi.org/10.1016/j.tics.2007.09.009

Olkkonen, M., & Allred, S. R. (2014). Short-Term Memory Affects Color Perception in Context. *PLoS ONE*, *9*(1), e86488. https://doi.org/10.1371/journal.pone.0086488

Panichello, M. F., Cheung, O. S., & Bar, M. (2013, January 21). Predictive feedback and conscious visual experience. *Frontiers in Psychology*. Frontiers. https://doi.org/10.3389/fpsyg.2012.00620

Paulun, V. C., Kawabe, T., Nishida, S., & Fleming, R. W. (2015). Seeing liquids from static snapshots. *Vision Research*, *115*, 163–174. https://doi.org/10.1016/J.VISRES.2015.01.023

Paulun, V. C., Schmidt, F., van Assen, J. J. R., & Fleming, R. W. (2017). Shape, motion, and optical cues to stiffness of elastic objects. *Journal of Vision*, *17*(1), 20. https://doi.org/10.1167/17.1.20

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, *10*(4), 437–442. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/9176953

Phillips, F. (2004). Creating Noisy Stimuli. *Perception*, *33*(7), 837–854. https://doi.org/10.1068/p5141

Schmid, A. C., & Doerschner, K. (2018). Shatter and splatter: The contribution of mechanical and optical properties to the perception of soft and hard breaking materials. *Journal of Vision*, *18*(1), 14. https://doi.org/10.1167/18.1.14

Schmid, A. C., & Doerschner, K. (2019). Representing stuff in the human brain. *Current Opinion in Behavioral Sciences*, *30*, 178–185. https://doi.org/10.1016/J.COBEHA.2019.10.007

Schmidt, F., Hegele, M., & Fleming, R. W. (2017). Perceiving animacy from shape. *Journal of Vision*, *17*(11), 10. https://doi.org/10.1167/17.11.10

Schmidt F., Paulun V.C., van Assen J.J.R. & Fleming, R.W. (2017). Inferring the stiffness of unfamiliar objects from optical, shape, and motion cues. *Journal of Vision*, 17(3):18. doi: 10.1167/17.3.18

Scocchia, L., Valsecchi, M., Gegenfurtner, K. R., & Triesch, J. (2013). Visual Working Memory Contents Bias Ambiguous Structure from Motion Perception. *PLoS ONE*, *8*(3), e59217. https://doi.org/10.1371/journal.pone.0059217

Scocchia, L., Valsecchi, M., & Triesch, J. (2014). Top-down influences on ambiguous perception: the role of stable and transient states of the observer. *Frontiers in Human Neuroscience*, *8*, 979. https://doi.org/10.3389/fnhum.2014.00979

Soechting, J. F., Juveli, J. Z., & Rao, H. M. (2009). Models for the Extrapolation of Target Motion for Manual Interception. *Journal of Neurophysiology*, *102*(3), 1491–1502. https://doi.org/10.1152/jn.00398.2009

Sharan, L., Rosenholtz, R. & Adelson, E. H. (2013). Accuracy and speed of material categorization in real-world images. *Journal of Vision*, 14(9):12, 1–24, doi: 10.1167/14.9.12

Summerfield, C., & De Lange, F. P. (2014, November 15). Expectation in perceptual decision making: Neural and computational mechanisms. *Nature Reviews Neuroscience*. https://doi.org/10.1038/nrn3838

Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature*, *381*(6582), 520–522. https://doi.org/10.1038/381520a0

Toscani, M., Valsecchi, M., & Gegenfurtner, K. R. (2013). Optimal sampling of visual information for lightness judgments. *Proceedings of the National Academy of Sciences*, *110*(27), 11163–11168. https://doi.org/10.1073/pnas.1216954110

Tschacher, W., Rees, G. M., & Ramseyer, F. (2014). Nonverbal synchrony and affect in

dyadic interactions. *Frontiers in Psychology*, *5*, 1323. https://doi.org/10.3389/fpsyg.2014.01323

Ullman, S. (1976). The interpretation of structure from motion. *Proceedings of the Royal Society of London: B. Biological Sciences*, *203*, 405–426.

Urgen, B. M., & Boyaci, H. (2019). When expectations are not met: unraveling the computational mechanisms underlying the effect of expectation on perceptual thresholds. *BioRxiv*, 545244. https://doi.org/10.1101/545244

van Assen, J. J. R., Barla, P., & Fleming, R. W. (2018). Visual Features in the Perception of Liquids. *Current Biology*, *28*(3), 452-458.e4. https://doi.org/10.1016/j.cub.2017.12.037

van Assen, J. J. R., & Fleming, R. W. (2016). Influence of optical material properties on the perception of liquids. *Journal of Vision*, *16*(15), 12. https://doi.org/10.1167/16.15.12

van Assen, J. J. R., Wijntjes, M. W. A., & Pont, S. C. (2016). Highlight shapes and perception of gloss for real and photographed objects. *Journal of Vision*, *16*(6), 6. https://doi.org/10.1167/16.6.6