
Visual affects: Linking curiosity, Aha-Erlebnis, and memory through information gain

Sander Van de Cruys^{1*}, Claudia Damiano¹, Yannick Boddez^{2,3}, Magdalena Król⁴, Lore Goetschalckx¹, Johan Wagemans¹

Accepted for publication in Cognition

1 Laboratory of Experimental Psychology, KU Leuven

2 Department of Experimental Clinical and Health Psychology, Ghent University

3 Centre for the Psychology of Learning and Experimental Psychopathology, KU Leuven

4 Institute of Psychology, SWPS University of Social Sciences and Humanities

Abstract

Current theories propose that our sense of curiosity is determined by the learning progress or information gain that our cognitive system expects to make. However, few studies have explicitly tried to quantify subjective information gain and link it to measures of curiosity. Here, we asked people to report their curiosity about the intrinsically engaging perceptual ‘puzzles’ known as Mooney images, and to report on the strength of their aha experience upon revealing the solution image (curiosity relief). We also asked our participants (279) to make a guess concerning the solution of the image, and used the distribution of these guesses to compute the crowdsourced semantic entropy (or ambiguity) of the images, as a measure of the potential for information gain. Our results confirm that curiosity and, even more so, aha experience is substantially associated with this semantic information gain measure. These findings support the expected information gain theory of curiosity and suggest that the aha experience or intrinsic reward is driven by the actual information gain. In an unannounced memory part, we also established that the often reported influence of curiosity on memory is fully mediated by the aha experience or curiosity relief. We discuss the implications of our results for the burgeoning fields of curiosity and psychoaesthetics.

Keywords: *curiosity, aha-Erlebnis, memory, aesthetic appreciation, information gain, uncertainty, intrinsic motivation, Mooney images, semantic ambiguity.*

* Correspondence to: Sander Van de Cruys, Brain & Cognition, Laboratory of Experimental Psychology, KU Leuven, 3000 Leuven, Belgium. E-mail: sander.vandecruys@kuleuven.be
<https://orcid.org/0000-0003-4831-7800>

Introduction

Why do we need a sense of curiosity? One plausible answer emerging from the recent revival of curiosity research (Gottlieb & Oudeyer, 2018; Kidd & Hayden, 2015) is that learning is costly. We need some sense of whether we will be able to make progress in learning the structure of our world, or we will just be wasting valuable computational resources. We need to navigate to activities or environments that will reveal learnable regularities and avoid spending limited resources on input variability that is either due to mere noise or to regularities that are just too complex given the mental models (competence) we currently possess. This is the role of curiosity defined as a type of motivation that is inherent within information processing, no matter the immediate adaptive value of the processed information (Hunt, 1981; Livson, 1967).

Indeed, from very early on in development, infants show a keen sensitivity for this. For example, 17-month-old infants intuitively attend more to learnable than to unlearnable artificial grammars, hence avoiding to labor in vain on inputs without learnable patterns (Gerken et al., 2011). Similarly, in the so-called Goldilocks effect, very young infants selectively pay attention to visual or auditory materials of intermediate predictability (Kidd et al., 2012, 2014), again preventing a waste of computational resources on inputs that are already known (too simple) or unknowable (too unpredictable). These early signs of curiosity, measured by looking time, align with measures of interest and curiosity later in life, often showing a similar focus on materials (be it artworks or artificial stimuli) of medium complexity (Berlyne, 1966; Day, 1981).

These strands of evidence, together with work in developmental robotics and computational neuroscience (Gottlieb et al., 2013; Schmidhuber, 2009), have converged on a concept of curiosity as expected learning progress or information gain. It casts curiosity as a metacognitive feeling based on specific information-theoretic principles and directing us to the best opportunities for learning. It is metacognitive because it is rooted in an evaluation of whether there is sufficient ground to (continue trying to) learn particular materials (Metcalfe et al., 2020). In other words, it “indicates when there’s a match between the presented learning material and the learner’s readiness to encode it” (Wade & Kidd, 2019), ensuring that we remain as much as possible in the so-called “zone of proximal development” (Vygotsky, 1962), the optimal region of learning. This concept of curiosity goes beyond the influential “information gap” theory (Loewenstein, 1994). Indeed, it is not sufficient to notice a gap in your knowledge (uncertainty) to become curious. In addition, curiosity requires a meta-cognitive expectation that the gap is bridgeable with current

capacities, in other words, that the uncertainty is resolvable, so we can ‘cope’ with it (see also, Silvia, 2005).

An account based on learning progress or information gain can explain why we often are most curious about medium uncertainty. Indeed, we can usually expect to make progress here, because medium uncertainty often implies we have at least some mental models for this domain. However, crucially, such an account does not enforce an inverted U-shape relation between uncertainty and curiosity. Indeed, contrary to the seminal works of Berlyne and Loewenstein, an information gain account predicts a monotonically increasing relation between uncertainty and curiosity, as long as the uncertainty in question is *expected to be reducible*. Evidently, the potential for learning progress increases with uncertainty.

The evidence for intermediate uncertainty being most curiosity-inducing is indeed mixed. While some studies have found support for medium uncertainty or confidence about a solution (Baranes et al., 2015; Kang et al., 2009; Marvin & Shohamy, 2016), usually in the context of curiosity about trivia questions, others showed that curiosity or exploration monotonically increases with uncertainty (van Lieshout et al., 2018). These inconsistencies may be due in part to a lack of consideration for the expected reducibility of the uncertainty studied (subjective appraisal of one’s coping potential). However, another major difficulty is the quantification of subjective uncertainty (and expected information gain) as such. Self-reported confidence ratings or objective stimulus-based uncertainty measures at best crudely approximate and at worst seriously misestimate that parameter.

Here, we use a new way of quantifying subjective uncertainty and information gain, for stimuli that have not been used before in curiosity research, namely Mooney images. Mooney images (Mooney & Ferguson, 1951), the most famous example of which is the ‘camouflaged’ Dalmatian, are constructed by blurring and thresholding natural grayscale images to arrive at irregular black-and-white patchworks, often impossible to recognize without extra cues (see **Figure 1**). These images have already given us many insights into our visual system (e.g., Dolan et al., 1997; Gorlin et al., 2012; Hegdé & Kersten, 2010) and they have a couple of features that make them extremely suited for curiosity research as well. Mooney images are naturally engaging perceptual ‘puzzles’ for participants. They can elicit a strong tip-of-the-tongue feeling (the feeling of being on the brink of resolution), associated with curiosity (Metcalf et al., 2020). In other words, they often give a sense that the uncertainty is reducible with continued sampling (eye fixations) and mental effort. With the best instances of Mooney images, perceivers also experience a strong phenomenological shift when eventually they autonomously discover (or are shown) the solution: One cannot ‘unsee’ it when again confronted with the same Mooney image (cf. one-shot learning; Giovannelli et al., 2010; Ishikawa & Mogi, 2011). The disambiguation is usually

accompanied by an Aha-Erlebnis (Kounios & Beeman, 2014), a positive feeling of insight and relief.



Figure 1: Mooney or two-tone image (left) and its grayscale source (solution) image (right).

The extent to which the experience of these images follows that prototypical Mooney character, i.e., the instant ‘click’ of recognition and the strong aha experience, depends on both individual factors (e.g., previous experience), as well as image-based factors (notably the level of ‘support’ for the solution). However, very little is known about the specific factors determining the positive aha experience. One tempting hypothesis is that where curiosity gauges *expected* information gain, the aha marks *actual* information gain: the intrinsic reward of *curiosity relief*. Information gain quantifies the reduction in uncertainty after the solution is known (it is also known as relative entropy, with entropy being the information-theoretic measure of uncertainty; Cover & Thomas, 1991). Theoretically, the idea is that a perceiver’s visual system has a particular probability distribution over candidate hypotheses or possible “hidden causes” for a given Mooney image. However, at this stage, its meaning is still ambiguous. Once the solution is known, all probability mass is concentrated on one best-supported hidden cause that explains the image features well (virtually zero uncertainty). In other words, the shift in the probability distribution over the hypothesis space should be proportional to the aha experience. Assuming that the posterior entropy is indeed zero, the initial entropy can be used as an estimate of the information gain or reduction in entropy (the shift). Still, this entropy cannot be directly measured for any particular subject and image combination (i.e., we do not know yet how such distributions are encoded in neural activity). Here, we approximated this semantic entropy of an image

by ‘crowdsourcing’ the distribution based on the counts of the guesses made by our total sample of participants. Under the assumption that curiosity measures *expected* uncertainty reduction (or information gain) and the relief (aha) is proportional to the *actual* uncertainty reduction, we asked to what extent people’s curiosity before the reveal accurately predicts their post-reveal aha experience, and whether both indeed correlate with semantic entropy.

Another outstanding question in curiosity research concerns the relation of curiosity with memory. Given curiosity’s above-mentioned role in making learning more efficient, a positive influence of curiosity on memory is to be expected. Indeed, several studies have gathered evidence that items for which participants reported higher curiosity, were also better remembered, sometimes up to weeks later (Gruber et al., 2014; Kang et al., 2009). However, a recent study suggests this effect might not be quite as strong as one would expect (Wade & Kidd, 2019). Importantly, most of these studies did not measure the participants’ reaction upon revealing the answer or solution for the item. Specifically, the intrinsic reward of curiosity relief could be a crucial mediator of the encoding strength in memory. While one study found a positive correlation between amygdala activity at reveal and solution memory for Mooney images (Ludmer et al., 2011), no previous study related the strength of the positive aha experience to memory. To sum up, we use Mooney images to investigate whether curiosity and curiosity relief are driven by subjective uncertainty and (potential) information gain, and whether curiosity and curiosity relief have independent effects on memory.

Methods

Participants

280 first-year psychology students (242 women) completed the study for course credits, 194 from collective testing sessions, and 86 from individual web-based sessions. The procedure was exactly the same for both. One participant was removed because she responded with the exact same level of curiosity on all images (making it impossible to compute z-scores).

This study was approved by the Social and Societal Ethical Committee of the KU Leuven and all participants gave explicit informed consent.

Stimuli

The prototypical Mooney images cannot be recognized without cues or sufficient time and effort, but become instantly recognizable as soon as the grayscale source image has been revealed once. Upon recognition, a strong phenomenological shift takes place when the observer is confronted with the Mooney again. In the prototypical cases, one cannot ‘unsee’

it (even after considerable delay), and a positive Aha-Erlebnis is felt with this perceptual ‘insight’. However, there is considerable individual variability in these characteristics (i.e., different people will experience the Mooney-effect for different images, to a different extent). For all these reasons, there is no automated way to create such images and not all grayscale source images are equally suited to generate Mooney versions (depending on lighting, object segmentation, background texture, etc.). Hence, we used a combination of automatized Mooney generation, hand selection, and piloting to create and narrow down the stimulus set.

First, a very large set of Mooney images was created from grayscale images (‘solutions’) by the procedure and code modified from Imamoglu et al. (2012), using source images from the Caltech 256 (Griffin et al., 2007) and MemCat (Goetschalckx & Wagemans, 2019) image databases. The former is a validated set of at least 80 different images for 256 everyday object categories, which has been used to benchmark object recognition in AI systems. The latter is an image set consisting of five broader semantic categories (animal, sports, food, landscapes, vehicles), with 2K exemplars each, further divided into different subcategories (e.g., bear, pigeon, cat, etc. for the animal category). The actual “Mooneyfication” process consists of grayscaling, low-pass filtering, and thresholding an image such that grayscale values below a certain value become white and those above this value become black (hence two tones). The optimal threshold is determined by Otsu’s method, which maximizes the variance between the two classes of pixels which are separated by the threshold (equivalently, it minimizes intra-class variance) (van der Walt et al., 2014).

To reduce the set size and obtain good Mooney candidates, we first eliminated images 1) that contained less or more than one identifiable foreground object, 2) that had little or no discernible structure after Mooneyfication (i.e. excluding images that had few, very large patches of black or white), or 3) in which the object was still overly obvious after Mooneyfication. The remaining 755 candidate Mooneys were presented (in two one-hour sessions) to 8 motivated participants. In this pretest, we only asked people whether (y/n) they recognized the object and to what extent seeing the solution induced an aha-experience (on a 7-point scale). Based on the results, we removed those images that were recognized without help by most participants (images with $>.8$ recognition rate) and kept those that generated at least some aha experience (images with >3.7 aha strength). Only Mooney images with low initial recognition rate and high post-solution aha were selected for the current study. The images all had a width and height between 250 pixels and 750 pixels. All stimulus creation and selection procedures are documented in the Open Science Framework (see Analysis section).

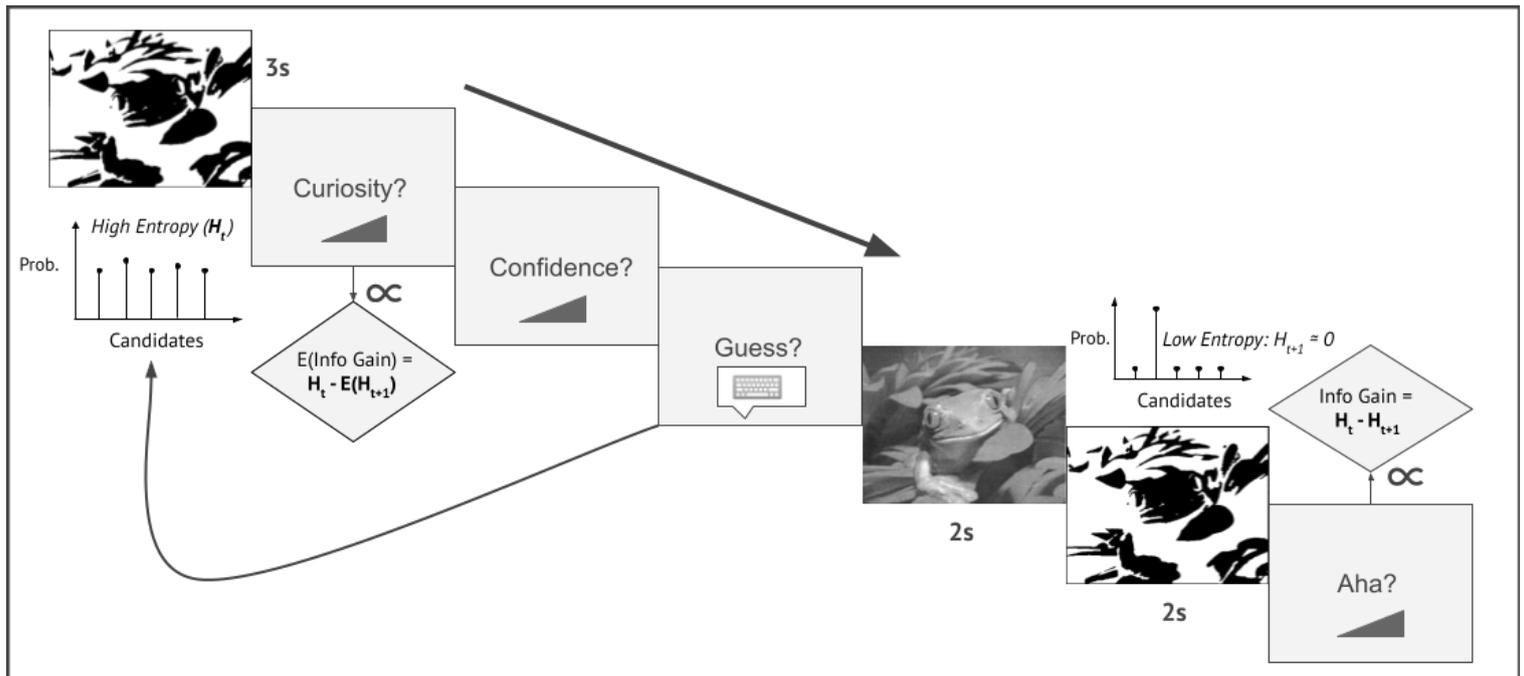


Figure 2: Flow chart of the procedure of part 1, with our measurements in rectangles and the corresponding theoretical concepts (and equations) in lozenges. The probability distribution plots represent semantic entropy before (H_t) and after (H_{t+1}) the solution has been shown (Prob.=Probability; E=Expected).

Procedure

In the first part of the study (Figure 2), participants were shown a fixation cross (200ms) and then a Mooney image for 3s followed by two 7-point rating scales. On one scale they indicated the curiosity they felt when seeing the image (“indifferent” to “very curious”), and on the other how confident they were about the solution of the image (“no clue at all” to “very certain about my answer”). On the next screen, participants were encouraged to make their guess concerning the content of the Mooney with a typed response. Participants were informed in the instructions that all images depicted objects from one of the following broad-level categories: animals, inanimate objects, plants, sports, vehicles, and food. Hence, a good guess means knowing more than these broad labels for a given image. We used examples (e.g., when a parrot is depicted, “parrot” or “bird” is correct, but not “animal”) to indicate that we were looking for basic- or subordinate-level categories. After making their guess, the solution (grayscale image) was shown (2s), followed by the Mooney image again (2s). This is the moment the participant will (or will not) experience the phenomenological shift. Immediately afterward, they were required to indicate the strength of their aha

experience, again on a 7-point rating scale (from “absent” to “very intense”). The aha experience is further described in the instructions as the “positive feeling of the ‘click’ you sometimes experience when the pieces of the ‘puzzle’ fall together and you suddenly have ‘insight’ in the image”. Each participant received a random set of 80 Mooney images (never more than 3 from the same category) out of the total set of 203.

In the second part, participants saw all images again (for 2s), supplemented by 20 new Mooney images, all in random order. Now, the task was to remember (y/n) whether they had seen the Mooney image already in part one (familiarity) and to remember the solution (another open response). After this unannounced memory part, participants filled out four questionnaires: two about different dimensions of trait curiosity, one about autism traits (AQ-short; Hoekstra et al., 2011), and the Need for Closure scale (Roets & Van Hiel, 2011). The latter measures a tendency to want cognitive closure and avoid ambiguity or confusion, which we reasoned may influence one’s response to ambiguous Mooney images. The AQ-short measures (subclinical) autism-like traits, which some studies suggest may alter Mooney perception (Król & Król, 2019; Loth et al., 2010). The Five-Dimensional Curiosity Scale Revised (5DCR; Kashdan et al., 2020) is a validated questionnaire assessing five dimensions of curiosity (joyous exploration, deprivation sensitivity, stress tolerance, thrill-seeking, and social curiosity). However, it does not include a perceptual curiosity dimension, so we included the 10 items of the Perceptual Curiosity Scale (PCS; Litman et al., 2005). Although trait curiosity might be a domain-general concept, we suspected trait perceptual curiosity would be more related to the type of state curiosity induced by our Mooney task.

Analysis

The open responses (first guesses pre-reveal and solution memory responses in the second phase) were corrected programmatically with a fuzzy matching algorithm comparing responses with pre-generated lists of valid responses, as described in Van de Cruys et al. (2018), accounting for variations in spelling or typos. Our sampling of 80 images per participant (out of the set of 203) gave us an average of 110 observations per image. One participant had no variation at all on the curiosity measure, and was therefore excluded from further analyses. All analyses except for the participant-based ones (relating to the questionnaires) were done on raw scale scores as well as (participant-based) z-transformed ones to account for individual differences in the use of the scales. Unless noted otherwise, all correlations are Pearson correlations on the z-transformed scores. We clearly indicate when analyses are post hoc, so corrected alpha-values apply. In terms of statistical tests, we used Generalized Estimating Equations (McNeish et al., 2016) as implemented in the Python

Statsmodels library (Seabold & Perktold, 2010) to correct for clustering of data points within participants. These have the advantage of making a smaller number of assumptions than hierarchical (generalized) linear models and the resulting beta's are interpreted identically to conventional (logistic) regression as the slope connecting the predictor to the dependent variable.

The semantic entropy measure was computed as follows. Each guess on a trial (pre-reveal guesses only) is compared to the list of used labels. If the guess is not in the list of labels (or if the list is still empty), a new possible label is created for this image, to which all other guesses for this image (by other participants) are compared (always using fuzzy-matching, see Methods). Frequencies of all labels in the list are used to compute proportions of the label in the total number of guesses (i.e. summed up frequencies of all the different guesses for a given image). Note that empty responses are also tallied towards the total. These 'probabilities' were then used to compute the information entropy for an image as Shannon defined it:

$$H(X) = - \sum P(x_i) \log(P(x_i))$$

Where x_i is the i -th guess in the list of used guesses. In this way, we attain a crowdsourced subjective uncertainty of an image. Note that this analysis only uses the raw guesses and totally disregards the ground truth for the image (guess accuracy).

In addition to the semantic entropy measure, we also assessed several measures of perceptual entropy calculated directly from the Mooney and grayscale images to investigate whether any low-level image statistics influence curiosity and aha ratings or memory accuracy. We included five measures in total – pixel entropy, edge density, PHOG complexity, anisotropy, and self-similarity – which have been found to relate to subjective ratings of complexity, interest, and pleasure (Greibenkina et al., 2018; Lyssenko et al., 2016; Van Geert & Wagemans, 2020). We refer to the supplementary materials for a brief explanation of how these measures are computed. We also calculated the structural similarity between the Mooney image and its grayscale solution as a measure of the physical match between the grayscale image and its Mooney counterpart. We hypothesized that the strength of the match would be related to the obviousness of the Mooney's solution, thus perhaps playing a role in the accuracy of the initial recognition of the Mooney or the strength of the aha experience. Structural similarity was calculated using the "ssim" function in Matlab (Wang et al., 2004).

All data and code for running the experiments, for deriving the image-based and semantic entropy measures, and for all analyses and plots, are available at osf.io/hm2kb. At the same

location, our Mooney stimulus set is made available for future studies, as well as an interactive data explorer. The latter allows the reader to experience the Mooney images and their resolution across the dimensions we measured. Information on inter-rater reliability can be found in the supplementary materials.

Results

Semantic entropy

Our first hypothesis concerned the effect of semantic entropy on curiosity and aha experience, tested in an analysis that disregards the accuracy of guesses. Indeed, most people made a guess on most trials (on average 76% of trials), but mean accuracy was only .22 (in the pre-reveal part). We created the semantic entropy measure using the number of different guesses and their frequency (across participants) for a given image (see *Analysis*). It is a proxy for the entropy of the subjective probability distribution over the hypothesis space for an image, possibly capturing the guesses that may be implicit in many participants but were not given (or even not conscious). We hypothesized that such an entropy measure would be proportional to curiosity, but it is usually hard to compute for a particular individual being confronted with a non-artificial stimulus (see also, Holm, 2017; Nicki, 1970). Of course, since our probabilities are based on the ‘votes’ from the whole group, our crowdsourced entropy measure should not be considered a direct measure of subjective entropy either, but rather an approximation of it. Related to that, the correlation between the semantic entropy of an image and the average confidence is only $-.28$ ($p < .001$). This shows these are different measures of uncertainty: for one guess versus a potential distribution. Indeed medium confidence may mean one sees different partial solutions (with similar confidence in each guess) or sees one solution for which the support in the image is not high.

We found that semantic entropy correlated substantially with curiosity ($r = .38$, $p < .001$) and aha ($r = .55$, $p < .001$; **Figure 3**), similar to confidence. However, confidence correlates more with curiosity ($r = .78$, $p < .001$) than with aha ($r = .57$, $p < .001$), while the reverse is true for semantic entropy (test of correlation difference: $z = -3.38$, $p = .001$; see Diedenhofen & Musch, 2015). This is consistent with aha being proportional to actual information gain, while curiosity tracks only an imperfect expectation of information gain. Indeed, (semantic) entropy represents a potential for learning progress, in that a highly spread out distribution of hypotheses (high entropy) will, after reveal, collapse into a distribution with concentrated probability mass on one particular hypothesis (i.e., posterior entropy is necessarily zero). It is

this shift in distribution that represents information gain and that we have identified as a determinant of the strength of the positive aha experience.

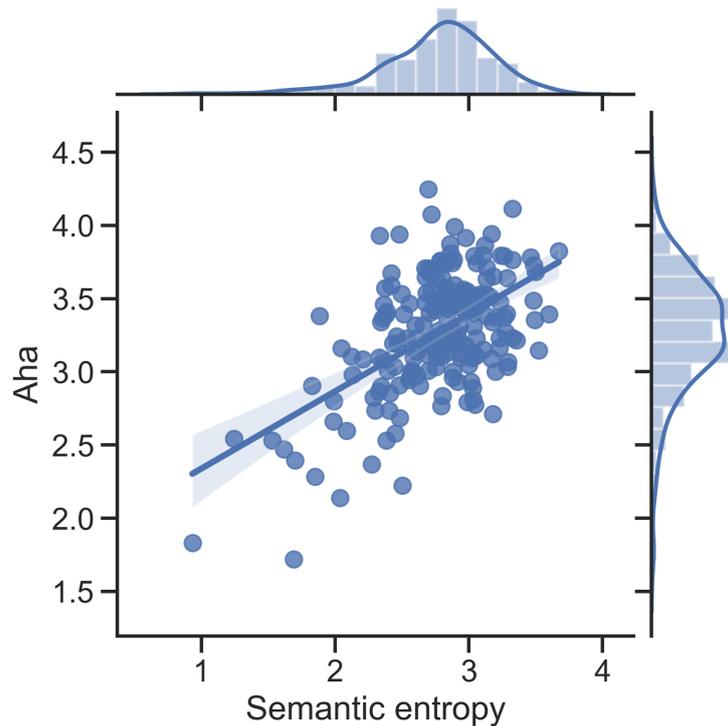


Figure 3: Average aha plotted as a function of the semantic entropy of a Mooney image. Dots are individual images, the solid line is the linear best fit. The shaded area is a 95% confidence interval for the regression.

In addition, the potential of multiple plausible hypotheses for a Mooney suggests it is an engaging image (a measure of its ‘poly-interpretability’), so semantic entropy may be a better basis to select prototypical Mooney images. It might be even better than average aha itself, because of the large individual differences in aha (see supplementary materials).

Notably, this measure of uncertainty at the ‘meaningful object candidate’ level turns out to have no correlations whatsoever with the low-level image complexity measures (all correlations lower than .05). Although it is sometimes implied that low-level measures capture the general complexity of an image, our findings show these measures really capture different things. The image cues that give rise to (multiple) interpretations seem to still escape us.

Curiosity, confidence, and aha

Next, we looked at the relation between confidence or certainty about the guess and curiosity. We find no evidence for a U-shaped relation as often reported in studies using trivia questions (verified with the “two lines” procedure; Simonsohn, 2018), but rather a significant monotonously decreasing relation, with people experiencing more curiosity the less confident they were about their guess ($r = -.33$; $p < .001$; see **Figure 4**). We also looked at what people are curious about: “the contents of the perceived information gap, or simply the confirmation of whether or not their best guesses are correct” (Wade & Kidd, 2019). Clearly, people were also, in fact even more, curious when they did not make a guess ($t(8409.55) = 16.91$; $p < .001$; see also supplementary figure 1). As a sanity check, we also could confirm that the more confident people were about their guess, the more likely they were indeed accurate ($r = .5$; $p < .001$).

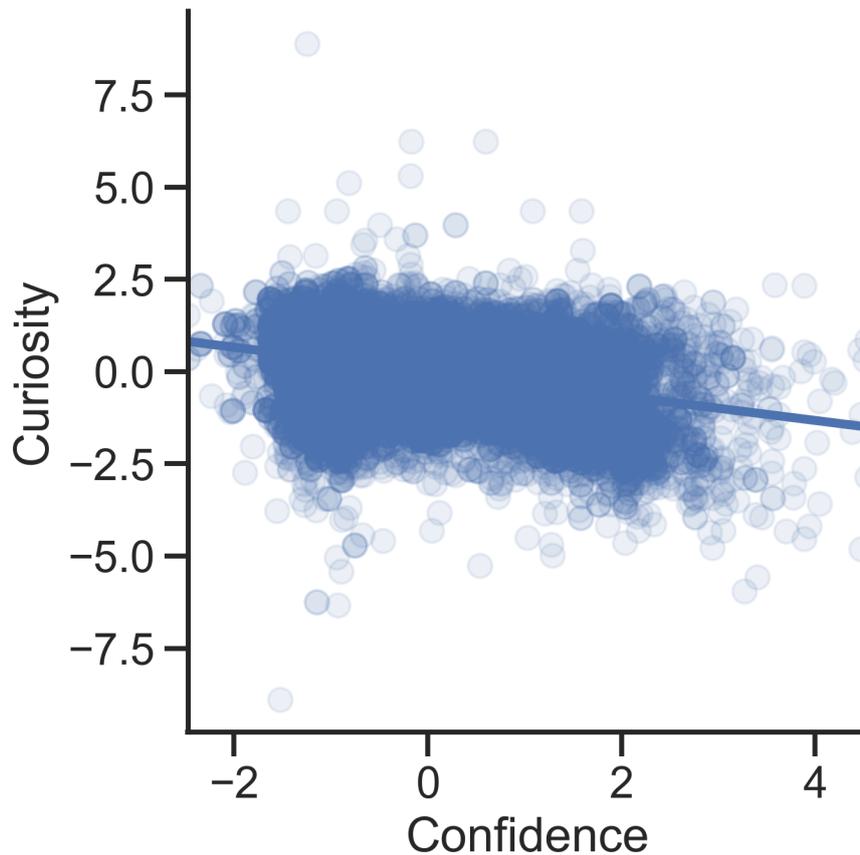


Figure 4: Standardized curiosity rating as a function of standardized confidence rating. Dots are individual trials, the solid line is the linear best fit. The shaded area (too small to see here) is a 95% confidence interval for the regression.

We could also confirm that people's curiosity (pre-reveal) predicts the strength of their (post-reveal) aha-experience ($r = .28$; $p < .001$, **Figure 5**). Assuming curiosity is indeed a measure of expected learning gain, while the aha experience marks the actual gains one has made, it seems that, through their curiosity, people indeed have some (imperfect) sense of future learning progress. This was even more obvious on a per-image analysis: images that on average elicited higher curiosity were also the ones that tended to elicit higher aha experiences ($r = .68$; $p < .001$, **Figure 5**). Note that these findings are modulated by accuracy, in the sense that, quite logically, accurate, high confidence trials give little or no curiosity or aha. Indeed, with regard to the relation between confidence and aha experience, we can see that only for accurate guesses, aha really goes down with confidence (for incorrect guesses: $r = -.02$, $p = .008$; for correct guesses: $r = -.31$; $p < .001$; see **Figure 6**), meaning that if you were highly confident of your guess and it turned out to be right, you have decreased aha. In those cases, the solution was likely too obvious, and you did not gain any new knowledge. If your guess was inaccurate, your confidence about it does not matter much in the aha experience, i.e., if you were highly confident of an inaccurate guess, there is no cost or benefit in strength of aha.

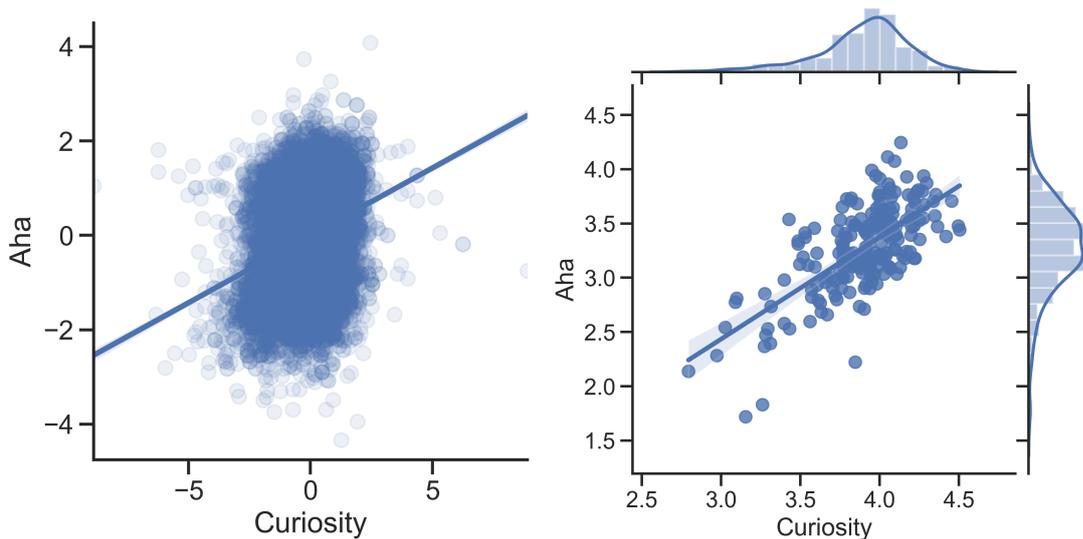


Figure 5: Standardized aha rating as a function of standardized curiosity rating (left panel). Dots are individual trials, the solid line is the linear best fit. The shaded area is a 95% confidence interval for the regression. In the right panel, each dot is an image, and average curiosity of this image is plotted by its average aha rating.

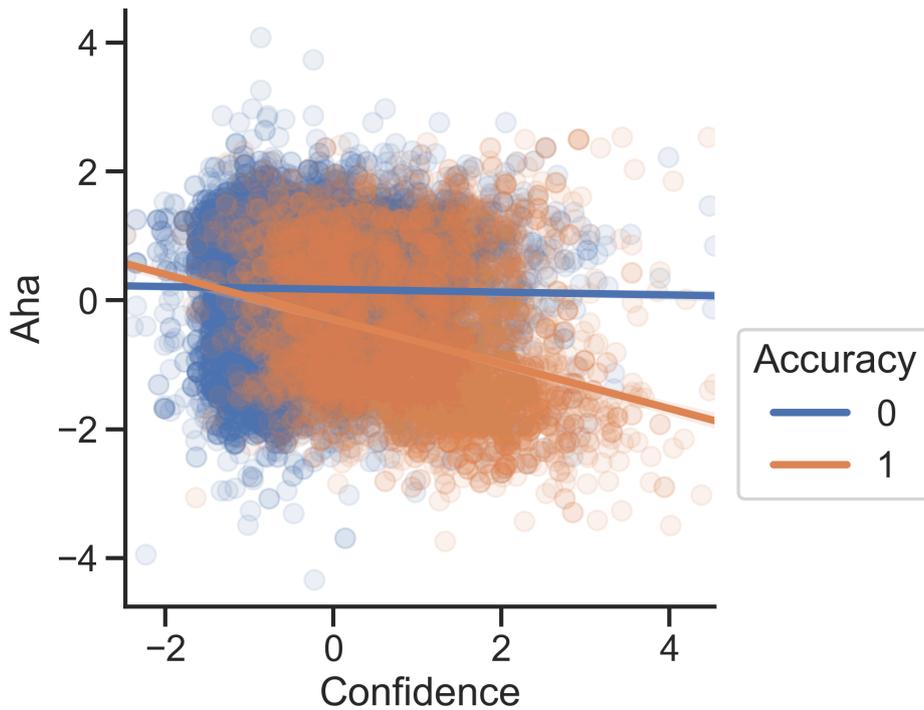


Figure 6: Standardized aha rating as a function of standardized confidence rating (left). Dots are individual trials, the solid line is the linear best fit. The shaded area is a 95% confidence interval for the regression. Data is plotted separately for accurate (red) and inaccurate trials (blue).

Curiosity, aha, and memory

Our next question is whether people indeed have better memory for images that induced greater curiosity and aha. In general, accuracy increased on average by 35 percentage points in the memory phase, compared to the pre-reveal phase (mean accuracy from .22 to .57). We found indeed that aha is predictive of better solution memory accuracy ($r = .12$; $p < .001$; see **Figure 7**) and that this is also the case, but less so, for curiosity ($r = .04$; $p < .01$). Note that for this analysis, we obviously removed the trials of the images for which participants

were already accurate the first time around. We confirmed this with a model including every measure from the first phase (accuracy, confidence, curiosity, and aha) as a predictor for dependent variable solution memory. While the main effects of aha ($B = .25$; $z = 11.66$; $p < .001$), confidence ($B = .17$; $z = 8.82$; $p < .001$), accuracy ($B = 1.64$; $z = 31.23$; $p < .001$), and the interaction between aha and accuracy ($B = -.34$; $z = -6.22$; $p < .001$) were significant, the effect of curiosity was not significant ($B = .03$; $z = 1.77$; $p = .07$). However, when aha was dropped from the model, the effect of curiosity did become significant ($B = .07$; $z = 4.33$; $p < .001$). Finally, the influence of aha on solution memory was confirmed using per-image data, again only including incorrect trials from Part 1, showing that those images that were (on average) better remembered, were the ones that elicited greater aha experiences in part 1 ($r = .31$; $p < .001$). Finally, participants with higher average aha scores (across images) tended to also have better memory, though this correlation was smaller ($r = .13$; $p = .03$). It suggests that most of the effect of aha on memory takes place within individuals, and between images, as one might expect. We will not report the findings on Mooney familiarity (old vs. new Mooney) memory in detail as correlations with curiosity and aha were systematically smaller and/or insignificant, suggesting particularly the solution memory is influenced by our predictors (these results can be found online at osf.io/hm2kb).

In a post hoc analysis, we added a computed variable representing the deviation of aha rating compared to a local baseline (moving window average of aha ratings in the previous 10 trials). If we control for aha, confidence and curiosity, only this aha deviation measure remains significant ($B = .1$; $z = 4.07$; $p < .001$) with regard to predicting solution memory. It seems then that if aha is a measure of intrinsic reward associated with the relief of curiosity (with information gain), particularly the (information) reward prediction error compared to the currently expected information reward (which seems to be continuously updated) influences memory (see also Marvin & Shohamy, 2016).

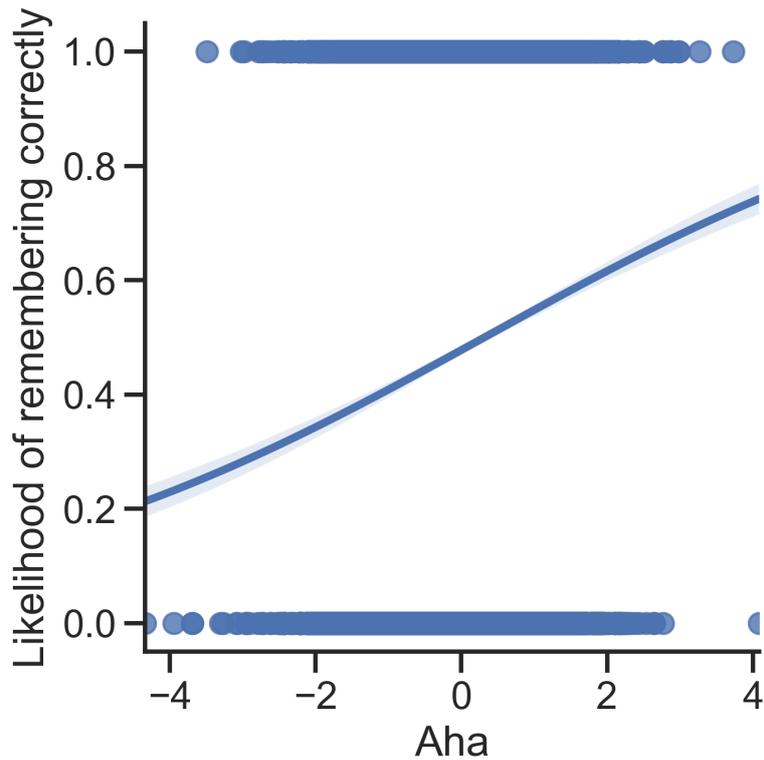


Figure 7: The likelihood of remembering an image correctly (given the Mooney version in part two) as a function of standardized aha rating (from part one). Dots are individual trials, the solid line is the logistic fit.

The shaded area is a 95% confidence interval for the regression.

In sum, our findings so far suggest that the greater their aha experience, the more strongly people encoded the image. The Mooney images that elicit a strong aha may be the “best” in terms of information gain, giving us the largest perceptual shift and the best solution memory. However, the impact of curiosity on memory was not very large. In fact, a mediation analysis (Imai et al., 2010) revealed that the effect of curiosity on solution memory is significantly mediated by aha experience (causal mediation effect: $B = .01$, $p < .001$) such that the direct effect of curiosity on memory is negligible ($B = -.003$, $p = .29$). This suggests that previous studies finding an effect of curiosity level of items on memory for those items might have missed this mediation, if they did not measure the experience (of appreciation or aha) upon revealing the missing information. Still, our study also confirms curiosity predicts this appreciation, so arguably curiosity is a good measure of how cognitively engaged one is or will be with the item, which in turn predicts memory (cf. the elaboration effect of memory; Craik & Tulving, 1975). Related to this, if we only consider inaccurate first guesses, we find higher strengths of aha experience when one has actually made a guess

($t(9221.23) = -6.04, p < .001$), suggesting that more engagement, in addition to expectation violation (wrong guess) has considerable influence on aha (and so memory). However, contrary to Brod and Breitwieser (2019) who report that making a prediction increases curiosity, we did not find a positive effect of engagement (making a guess) on curiosity (cf. supra). In fact, we find the opposite effect: people are more curious when they do not make a prediction. This could be due to differences in baseline engagement (interest) in the stimulus materials (numerical trivia facts in Brod and Breitwieser) or, as highlighted in the introduction, with a strong expectation of solvability (uncertainty reduction) because solutions were always provided.

State and trait curiosity

We were also interested in whether a person's average state curiosity as measured in our task is correlated with a person's trait curiosity as measured by validated questionnaires on trait curiosity. We indeed find that one's average state curiosity is positively predictive of one's score on the 5DCR, but only very mildly so ($r = .13; p = .03$). Surprisingly, this correlation seems to be driven largely by a relation with the social curiosity subscales, as the correlations with the other subscales were all insignificant (overt social curiosity $r = .177; p = .003$; covert social curiosity $r = .14; p = .016$; all other subscales: $p > .15$). However, also the separate trait perceptual curiosity scale correlated significantly, though again very mildly, with average state curiosity ($r = .15, p = .01$), as expected for a perceptual task like ours. Given those small correlations, we can conclude that our form of induced state curiosity is not captured well by the trait curiosity measures. Interestingly, participants that had higher average curiosity (or aha) did not have a higher memory accuracy ($r = -.003, p = .91$) which implies that the effects of curiosity or aha on memory take place within persons, as the information gain account would predict. Finally, neither the Need for Closure Scale nor the Autism Quotient questionnaire (AQ-28) correlated with task curiosity.

The contribution of low-level image statistics

Finally, we explored the effect of those low-level image complexity measures on the strength of curiosity, aha, and in (memory) accuracy. We computed multiple measures of image complexity (see Analysis) for both the Mooney image and the corresponding grayscale image. In addition, we computed the structural similarity between the Mooney image and its corresponding grayscale image, which we expected would be able to capture part of the ease of resolution or the obviousness of the solution in the Mooney. Structural similarity had no effect on first guess accuracy ($r = -.02; p = .83$), but it was related to post-reveal solution memory ($r = .19, p = .007$; first part inaccurate guesses only; **Figure 8**),

suggesting that it is not so much a measure of obviousness of meaningful structure, but that it may only help to make similar features become actual cues after you have already seen the solution. However, structural similarity was not correlated with aha experience ($r = .06$, $p = .39$), nor with curiosity or confidence, implying that it does not capture the processing characteristics that determine aha experience. Top-down influences (not captured by low-level measures such as structural similarity) may have a more important role here. With regard to our complexity measures, our analysis should be considered exploratory. Although relations have been found between preference and interest and various computed complexity measures for (abstract) artworks (Lyssenko et al., 2016) or arrays of objects (Van Geert & Wagemans, 2020), our analyses did not reveal (Bonferroni-corrected) significant correlations between curiosity or aha and any of our complexity measures. The only correlations that survived corrections were between complexity measures (primarily: anisotropy of the Mooney, anisotropy of the grayscale, edge entropy of the grayscale), and familiarity memory (old vs. new Mooney), such that familiarity memory was more accurate for less complex images. Interestingly, measures concerned both the Mooney and grayscale versions, meaning that Mooney familiarity accuracy varied with characteristics of the corresponding solution. Correlations with solution memory accuracy pointed in the same direction but did not survive our strict correction. In sum, low-level image statistics explain little of the variance in aha, curiosity, or memory.

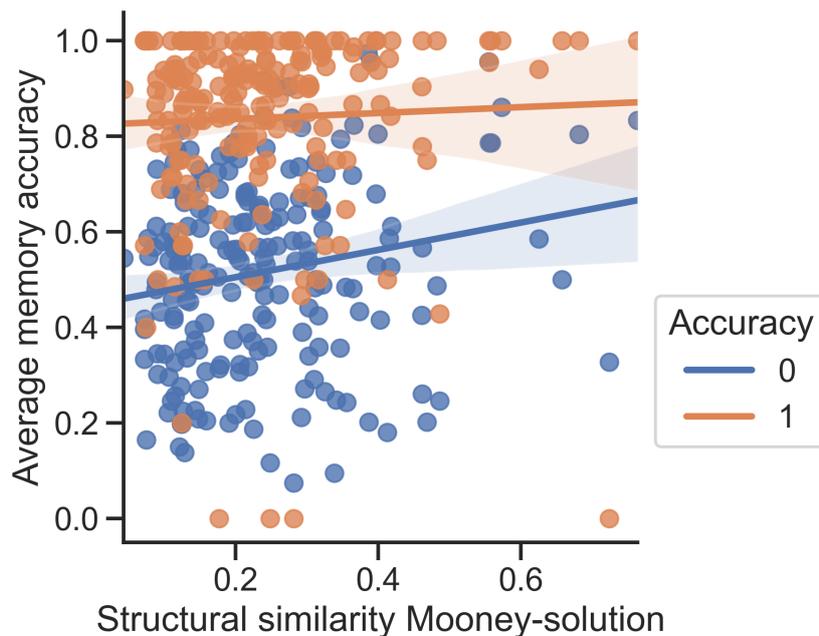


Figure 8: Average memory accuracy plotted as a function of the structural similarity between Mooney image and grayscale solution (source) image. Dots are individual images, the solid line is the linear best fit. The shaded area is a 95% confidence interval for the regression. Data is plotted separately for accurate (red) and inaccurate trials (blue).

Discussion

The key contributions of our study can be summed up in five points. First, we found clear evidence that curiosity (relief) tracks semantic information gain, consistent with current theories of curiosity based on expected learning progress. Second, our analyses indicate that the effect of curiosity on (incidental) solution memory is mediated by aha or curiosity relief, which may imply that curiosity facilitates memory, but only if the relief (actual gain) aligns with or exceeds curiosity (expected gain). Third, we show that the strength of ‘state’ perceptual curiosity reflects a domain-general individual trait curiosity only to a very limited extent, as measured by established curiosity questionnaires. Fourth, low-level image-based statistics do not explain much of the variance in curiosity, aha, memory, or semantic interpretability (semantic entropy). This suggests that these image statistics have limited importance in understanding complex evaluations such as curiosity and (aesthetic) appreciation, which has been likened to the aha experience (Muth & Carbon, 2013). Fifth and finally, we provide a new set of stimuli for research on perception and curiosity, with reference data on multiple dimensions (most notably semantic entropy) that can be used in future studies into the neural or psychophysiological correlates of curiosity and curiosity relief.

Any experimental ‘task’ on curiosity will technically be a bit contrived because a ‘task’ is imposed, while curiosity provides a sense of what ‘tasks’ to engage in to begin with (but see, Geana et al., 2016). This implies an unavoidable degree of motivational impurity in studies on curiosity or intrinsic motivation, since “any instruction intended to entice subjects to work on a task carries with it an implication that the experimenter will be pleased if the subject does so and displeased if the subject does not” (Walker, 1981). For this reason, we started out with intrinsically engaging visual stimuli, that may be more suited for studying curiosity than the oft-used trivia questions. Two earlier studies used image materials but in blurred (low-pass filtered) versions, instead of the thresholded (Mooney) images that we used. Nicki (1970) found that medium blur leads to maximum uncertainty about object identity, using guesses and confidence ratings of one individual. Still, consistent with our findings, he further reports that curiosity, as measured by the preference to see the unblurred

object (rather than an unrelated but comparable clear image) after being presented with the blurred object, follows an inverted U-shaped function of blurredness. Jepma et al. (2012) similarly used images with an intermediate degree of blur, but did not make trial by trial measurements of curiosity (or relief). They did report that resolving blurred images led to improved incidental memory, consistent with our findings.

Van Lieshout et al. (2018) studied curiosity in a lottery task, in which participants were confronted with a vase of marbles of two different colors and in different proportions, to manipulate uncertainty. Participants could win monetary rewards based on the outcome of the draw from the vase, and their curiosity about the outcome was measured either through self-report or through their willingness to wait to see the outcome. As in the current study, their results showed that curiosity was a monotonically increasing function of uncertainty, computed as the outcome uncertainty for a given vase (as well as independently varied reward probability). Note that this is a purely passive observation task because participants could not influence outcomes in any way. This is slightly different from our task, given that it was possible to influence outcomes using eye movements and self-generating possible solutions (at least for the time the Mooney image was on-screen). If anything this should increase curiosity.

Indeed, the availability of actions to gather more information (i.e., exploration) is crucial, because, as indicated in the introduction, it heightens the sense of (expected) reducibility of uncertainty that we identified as crucial in curiosity. In line with the recent predictive processing (also called active inference) accounts of the brain (Clark, 2015; Hohwy, 2013), this implies a proactive stance of the brain (agent) not only inferring the hidden causes of the current input (e.g. the object that ‘generated’ the Mooney image) but also predicting the uncertainty that is to come, as well as whether that uncertainty is expected to be resolvable. The latter is called epistemic value (or information gain) and is evaluated in order to select future actions so as to minimize (expected) uncertainty (Friston et al., 2017; Schwartenbeck et al., 2013; Seth et al., 2020). Without going into the computational technicalities of this account of curiosity, it brings two important corrections on classical theories of curiosity (Berlyne, 1966; Loewenstein, 1994). The latter were drawing on classical information theory centered on the idea of passive receivers, while active inference emphasizes the active contributions of agents. First, this means that it is not some unspecific or objective uncertainty that matters for curiosity but rather the subjective uncertainty, which is always relative to a particular model. Consequently, expected learning progress is not only based on current sensory evidence but also on prior knowledge formed by similar experiences (Wade & Kidd, 2019). As van Lieshout et al. (2018) remark this may explain why studies based on trivia questions found moderate uncertainty to be most curiosity-inducing: Participants

might just not have had mental models for many of the topics of the questions (so no interest). Second, the active inference account makes apparent that curiosity is not just a function of perception but also of the potential for action. It is captured by expectations on the resolvability of uncertainty, and it is another reason why the maxim of moderate uncertainty may not hold. Future studies using our Mooney stimuli could systematically address whether the possibility for action or so-called ‘epistemic foraging’ (exploration to learn about the structure of sensory inputs, see e.g., Clark, 2018), modulates curiosity. For example, participants could actively control partial revealing of information, by gradually blending in the solution with the Mooney image. Another limitation of the current study is that we did not ask participants to explicitly judge the expected solvability (reducibility of the uncertainty) of the Mooney images, because it might have influenced the curiosity measure. Future studies could select a subset of high curiosity-inducing and low curiosity-inducing images to verify whether these indeed differ reliably in the extent to which people estimate that they would be able to solve them given unlimited viewing time. Given that we eventually provided the solution for all of the images, participants might have judged reducibility to be similar for all images. However, given that information gain explains only part of the variance in curiosity, it is conceivable that differences in reducibility did play a role here.

Let us now turn to the interpretation of the aha-related results. In fact, the aha experience is often connected to aesthetic appreciation and preference. For example, Muth and Carbon (2013) found that having the insight or Aha-Erlebnis for a Mooney image increases subsequent liking of the image. And indeed, the literature on preference or (aesthetic) appreciation and curiosity or interest is highly overlapping, both in theoretical ideas (cf. inverted U-curve) and empirical measurements (e.g., preference is used to measure both curiosity and appreciation). This has historically always been the case, and for good reasons. Our conceptualization in terms of expected (curiosity) and actual (aha or appreciation) information gain solidifies this connection. Hence, we discuss aha as a form of intrinsic reward causing the appreciation of visual images, while acknowledging that the aha phenomenon may only explain part of what makes us like images (aesthetically).

The inventive early work on appreciation or liking focused on geometric patterns for which objective entropy (complexity) could still be quantified (e.g., Terwilliger, 1963). Consistent with the reasoning on learning progress, the results show a clear influence of learning and a preference for patterns that deviate systematically from what one is used to (and has presumably learned already). The objective entropy that is most pleasurable might indeed be a moving target because it is subjective entropy that counts. Interestingly, Terwilliger (1963) already noted that 1) even for these ‘simple’ patterns, there are multiple ways to

quantify ‘objective’ complexity depending on subjective choice of the coding scheme (for an instructive formal treatment, see Feldman, 2004), and, 2) that proper comparisons were clouded by familiarity and meaning emerging to different degrees along the entropy continuum, sometimes in idiosyncratic ways. More recent research has tried to link objective low-level image entropy of actual artworks or more ecological arrays of objects with appreciation or interest for these stimuli but has all-in-all led to only modest correlations (Graham, 2019; Lyssenko et al., 2016; Van Geert & Wagemans, 2020), consistent with our own findings. Although our stimuli are obviously not actual artworks, the different low-level entropy measures indeed explain very little variance in meaning and appreciation (or curiosity). Maybe this should not surprise us, given that the perceptual features that make something a potential bearer of meaning may require integration on a more global scale, something that is not adequately captured by the low-level characteristics we used. Furthermore, the very task we used (identifying the content of images) may have biased participants to the semantic level, thereby curtailing the effect of lower-level factors on aha (and curiosity). That said, there is evidence for a default semantics-first processing of (visual) stimuli (Peterson, 1994; Pinna, 2010), suggesting that our results may generalize beyond our specific task or materials.

Indeed, in contrast to pixel entropy, semantic entropy of images did strongly predict appreciation, which matches with findings of semantic instability, ambiguity, or indeterminacy being conducive to (aesthetic) appreciation (Ishai et al., 2007; Muth et al., 2016; Muth & Carbon, 2016; Nicki et al., 1979; Pepperell, 2006; Zeki, 2004). However, because we measured appreciation as the aha experience at a time when the stimulus is already perfectly disambiguated, our findings throw new light on the reason why semantically unstable or indeterminate images are often liked more. This is, arguably, the case because these poly-interpretable, indeterminate images have a greater potential for information gain. Indeed, we have attributed the strong aha to large shifts in the distribution of beliefs (in the Bayesian sense) technically known as information gain, relative entropy, or Kullback–Leibler divergence (see also Itti & Baldi, 2009; Tanner & Itti, 2017). Liking a stimulus is a function not of the stimulus per se, but of the subjective *process* of going from a state of high uncertainty to a state of lower uncertainty (Van de Cruys, 2017; Van de Cruys & Wagemans, 2011).

That said, information gain at the level of semantic entropy clearly does not capture all factors involved in the aha experience. Specifically, it does not fully account for the quality of the end solution: The support in the image for the particular solution. Knowing the solution compresses our mental representation of the Mooney image: It becomes predictable so you know which features, patches, or edges in the Mooney image are relevant, which

belong together and which are irrelevant (e.g. background, shadows,...). If all cues in an image line up, it is unlikely to be created by a random process, increasing the “eureka” feeling (Feldman, 2004). This is an additional way in which we reduce entropy or make learning progress here better (but equivalently) described as predictive or compression progress (Schmidhuber, 2009). Images vary on this aspect as well, so that the sensory ‘evidence’ in the image will be explained better or worse by the hidden object. One might conjecture that, if there are definite image cues that lend themselves to multiple interpretations (cf. high semantic entropy), that also implies these are well-explained by the actual solution, i.e. can serve well as evidence or ‘support’ for it. However, whether this form of compression gain can be a source of the positive aha experience would need to be tested explicitly (e.g. using the paradigm by Król & Król, 2018).

Finally, our findings concerning the modest relations of average state curiosity to questionnaires about trait curiosity, indicate that the latter clearly need to be validated more thoroughly based on actual experimentally-induced curiosity (or exploration), instead of merely using other questionnaires measuring similar constructs. Only then will we be able to establish whether curiosity is a domain-general trait. Note as well that the information gain account of curiosity casts serious doubts on whether a domain-general factor would account for much of the variability. The importance of a variable sense of curiosity sensitive to actual good learning opportunities is mostly apparent within individuals (as our results confirm).

Along the same lines, our findings suggest hacking our sense of curiosity to improve memory (e.g., in education), would only work if curiosity relief (aha) is boosted accordingly. Marvin and Shohamy (2016) measured both curiosity about and satisfaction with the answers to different trivia questions and computed the information prediction error as the discrepancy between the anticipated (the curiosity rating) versus the received information reward (the satisfaction rating). They found that this information prediction error predicted subsequent likelihood of remembering the answers. We could not confirm this in our data (using the difference between curiosity and aha, instead of satisfaction). Marvin and Shohamy also found that curiosity as such predicts memory too, but do not report the relation of satisfaction with memory separately, nor whether curiosity’s effects on memory could be completely mediated by satisfaction, as we found for aha. However, analogous to Marvin and Shohamy’s (2016) information prediction error findings, our results seemed to show that memory was a function of the deviation of aha (information reward) from a recent (running average) baseline of aha. This suggests that people build up an expected information gain from their experienced gains in the past trials and that the ‘information gain prediction error’ of the current trial predicts their memory (more than

‘raw’ aha).

While we hypothesize that information gain is a key factor in the memory improvements that we see with increasing curiosity and aha, generic processes may mediate this memory facilitation. We could not confirm that engagement as measured by whether participants made a guess or not had an effect on memory, but more granular measures of engagement (such as eye movement patterns) may reveal such an effect. Given that specifically aha strength is related to memory, it is telling that the aha usually appears when presented with the original Mooney image again and it ‘clicks’ after having seen the solution. This may point to a role for the well-established generation effect (Slamecka & Graf, 1978): the finding that materials that you yourself inferred or (re)constructed are better retained in memory.

In conclusion, we seem to be curious about and like experiences which allow the greatest information gain, or, equivalently, have the most potential to influence our model of the given perceptual inputs. Causal connections must remain tentative, because our study is purely correlational, even though semantic entropy and curiosity precede aha by design. Our findings also shed new light on the interplay between curiosity, information gain, aha, and memory, showing the usefulness of our stimulus set (and the collected norming data) for this field.

Supplementary material

Our raw data, stimuli, and experiment/analysis code can be found at <https://osf.io/hm2kb/>

Acknowledgments

We would like to thank Pieter Moors and Eline van Geert for advice on the ICC analyses, and Jo Bervoets for inspiring discussions. This work is supported by a Methusalem grant from the Flemish Government (METH/14/02), awarded to JW. MK is supported by the National Science Centre in Poland (Narodowe Centrum Nauki; 2017/27/B/HS6/00169) and YB is supported by a Ghent University grant (BOF16/MET_V/002).

References

Baranes, A., Oudeyer, P.-Y., & Gottlieb, J. (2015). Eye movements reveal epistemic curiosity in

- human observers. *Vision Research*, 117, 81–90. <https://doi.org/10.1016/j.visres.2015.10.009>
- Berlyne, D. E. (1966). Curiosity and Exploration. *Science*, 153(3731), 25–33.
<https://doi.org/10.1126/science.153.3731.25>
- Brod, G., & Breitwieser, J. (2019). Lighting the wick in the candle of learning: generating a prediction stimulates curiosity. *Npj Science of Learning*, 4(1), 1–7.
<https://doi.org/10.1038/s41539-019-0056-y>
- Clark, A. (2015). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press. <https://market.android.com/details?id=book-TnqECgAAQBAJ>
- Clark, A. (2018). A nice surprise? Predictive processing and the active pursuit of novelty. *Phenomenology and the Cognitive Sciences*, 17(3), 521–534.
<https://doi.org/10.1007/s11097-017-9525-z>
- Cover, T. M., & Thomas, J. A. (1991). Entropy, relative entropy and mutual information. *Elements of Information Theory*, 2(1), 12–13.
<http://www.cs.columbia.edu/~vh/courses/LexicalSemantics/Association/Cover&Thomas-Ch2.pdf>
- Craik, F. I. M., & Tulving, E. (1975). Depth of processing and the retention of words in episodic memory. *Journal of Experimental Psychology. General*, 104(3), 268–294.
<https://doi.org/10.1037/0096-3445.104.3.268>
- Day, H. I. (1981). *Advances in Intrinsic Motivation and Aesthetics*. Springer Science & Business Media.
- Diedenhofen, B., & Musch, J. (2015). cocor: a comprehensive solution for the statistical comparison of correlations. *PloS One*, 10(3), e0121945. <https://doi.org/10.1371/journal.pone.0121945>
- Dolan, R. J., Fink, G. R., Rolls, E., Booth, M., Holmes, A., Frackowiak, R. S. J., & Friston, K. J. (1997). How the brain learns to see objects and faces in an impoverished context. *Nature*,

389(6651), 596–599. <https://doi.org/10.1038/39309>

- Feldman, J. (2004). How surprising is a simple pattern? Quantifying “Eureka!” *Cognition*, *93*, 199–224.
- Friston, K. J., Lin, M., Frith, C. D., Pezzulo, G., Hobson, J. A., & Ondobaka, S. (2017). Active Inference, Curiosity and Insight. *Neural Computation*, *29*(10), 2633–2683. https://doi.org/10.1162/neco_a_00999
- Geana, A., Wilson, R., Daw, N. D., & Cohen, J. D. (2016). Boredom, Information-Seeking and Exploration. *CogSci*. <https://cogsci.mindmodeling.org/2016/papers/0307/paper0307.pdf>
- Gerken, L., Balcomb, F. K., & Minton, J. L. (2011). Infants avoid “labouring in vain” by attending more to learnable than unlearnable linguistic patterns. *Developmental Science*, *14*(5), 972–979. <https://doi.org/10.1111/j.1467-7687.2011.01046.x>
- Giovannelli, F., Silingardi, D., Borgheresi, A., Feurra, M., Amati, G., Pizzorusso, T., Viggiano, M. P., Zaccara, G., Berardi, N., & Cincotta, M. (2010). Involvement of the parietal cortex in perceptual learning (Eureka effect): an interference approach using rTMS. *Neuropsychologia*, *48*(6), 1807–1812. <https://doi.org/10.1016/j.neuropsychologia.2010.02.031>
- Goetschalckx, L., & Wagemans, J. (2019). MemCat: a new category-based image set quantified on memorability. *PeerJ*, *7*, e8169. <https://doi.org/10.7717/peerj.8169>
- Gorlin, S., Meng, M., Sharma, J., Sugihara, H., Sur, M., & Sinha, P. (2012). Imaging prior information in the brain. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(20), 7935–7940. <https://doi.org/10.1073/pnas.1111224109>
- Gottlieb, J., & Oudeyer, P.-Y. (2018). Towards a neuroscience of active sampling and curiosity. *Nature Reviews. Neuroscience*, *19*(12), 758–770. <https://doi.org/10.1038/s41583-018-0078-0>
- Gottlieb, J., Oudeyer, P.-Y., Lopes, M., & Baranes, A. (2013). Information-seeking, curiosity, and

- attention: computational and neural mechanisms. *Trends in Cognitive Sciences*, 17(11), 585–593.
<https://doi.org/10.1016/j.tics.2013.09.001>
- Graham, D. (2019). The Use of Visual Statistical Features in Empirical Aesthetics. In Marcos Nadal And (Ed.), *The Oxford Handbook of Empirical Aesthetics*. Oxford University Press.
<https://doi.org/10.1093/oxfordhb/9780198824350.013.19>
- Grebenkina, M., Brachmann, A., Bertamini, M., Kaduhm, A., & Redies, C. (2018). Edge–Orientation Entropy Predicts Preference for Diverse Types of Man–Made Images. *Frontiers in Neuroscience*, 12, 678. <https://doi.org/10.3389/fnins.2018.00678>
- Griffin, G., Holub, A., & Perona, P. (2007). *Caltech-256 Object Category Dataset*. 20.
<https://authors.library.caltech.edu/7694>
- Gruber, M. J., Gelman, B. D., & Ranganath, C. (2014). States of curiosity modulate hippocampus–dependent learning via the dopaminergic circuit. *Neuron*, 84(2), 486–496.
<https://doi.org/10.1016/j.neuron.2014.08.060>
- Hegd , J., & Kersten, D. (2010). A link between visual disambiguation and visual memory. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 30(45), 15124–15133.
<https://doi.org/10.1523/JNEUROSCI.4415-09.2010>
- Hoekstra, R. A., Vinkhuyzen, A. A. E., Wheelwright, S., Bartels, M., Boomsma, D. I., Baron-Cohen, S., Posthuma, D., & van der Sluis, S. (2011). The construction and validation of an abridged version of the autism–spectrum quotient (AQ–Short). *Journal of Autism and Developmental Disorders*, 41(5), 589–596. <https://doi.org/10.1007/s10803-010-1073-0>
- Hohwy, J. (2013). *The Predictive Mind*. Oxford University Press.
<http://books.google.be/books?id=3m8nAgAAQBAJ>
- Holm, L. (2017). Curiosity and expected information gain in word learning. *Editors: Anders*

Arweström Jansson Uppsala University Anders. Arwestrom.Jansson@It.Uu. Se Anton Axelsson Uppsala University Anton. Axelsson@It.Uu. Se Rebecca Andreasson Uppsala University Rebecca. Andreasson@It.Uu. Se, 43.

<https://www.diva-portal.org/smash/get/diva2:1156189/FULLTEXT01.pdf#page=51>

Hunt, J. M. (1981). Experiential Roots of Intention, Initiative, and Trust. In H. I. Day (Ed.), *Advances in Intrinsic Motivation and Aesthetics* (pp. 169–202). Springer US.

https://doi.org/10.1007/978-1-4613-3195-7_8

Imai, K., Keele, L., & Tingley, D. (2010). A general approach to causal mediation analysis.

Psychological Methods, 15(4), 309–334. <https://doi.org/10.1037/a0020761>

Imamoglu, F., Kahnt, T., Koch, C., & Haynes, J.-D. (2012). Changes in functional connectivity support conscious object recognition. *NeuroImage*, 63(4), 1909–1917.

<https://doi.org/10.1016/j.neuroimage.2012.07.056>

Ishai, A., Fairhall, S. L., & Pepperell, R. (2007). Perception, memory and aesthetics of indeterminate art. *Brain Research Bulletin*, 73(4–6), 319–324. <https://doi.org/10.1016/j.brainresbull.2007.04.009>

Ishikawa, T., & Mogi, K. (2011). Visual one-shot learning as an “anti-camouflage device”: a novel morphing paradigm. *Cognitive Neurodynamics*, 5(3), 231–239.

<https://doi.org/10.1007/s11571-011-9171-z>

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

Jepma, M., Verdonschot, R. G., van Steenbergen, H., Rombouts, S. A. R. B., & Nieuwenhuis, S. (2012). Neural mechanisms underlying the induction and relief of perceptual curiosity. *Frontiers in Behavioral Neuroscience*, 6. <https://doi.org/10.3389/fnbeh.2012.00005>

Kang, M. J., Hsu, M., Krajbich, I. M., Loewenstein, G., McClure, S. M., Wang, J. T.-Y., &

- Camerer, C. F. (2009). The wick in the candle of learning: epistemic curiosity activates reward circuitry and enhances memory. *Psychological Science, 20*(8), 963–973.
<https://doi.org/10.1111/j.1467-9280.2009.02402.x>
- Kashdan, T. B., Disabato, D. J., Goodman, F. R., & McKnight, P. E. (2020). The Five-Dimensional Curiosity Scale Revised (5DCR): Briefer subscales while separating overt and covert social curiosity. *Personality and Individual Differences, 157*, 109836.
<https://doi.org/10.1016/j.paid.2020.109836>
- Kidd, C., & Hayden, B. Y. (2015). The Psychology and Neuroscience of Curiosity. *Neuron, 88*(3), 449–460. <https://doi.org/10.1016/j.neuron.2015.09.010>
- Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2012). The Goldilocks Effect: Human Infants Allocate Attention to Visual Sequences That Are Neither Too Simple Nor Too Complex. *PloS One, 7*(5), e36399. <https://doi.org/10.1371/journal.pone.0036399>
- Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2014). The Goldilocks effect in infant auditory attention. *Child Development, 85*(5), 1795–1804. <https://doi.org/10.1111/cdev.12263>
- Kounios, J., & Beeman, M. (2014). The cognitive neuroscience of insight. *Annual Review of Psychology, 65*, 71–93. <https://doi.org/10.1146/annurev-psych-010213-115154>
- Król, M., & Król, M. (2018). “Economies of Experience”—Disambiguation of Degraded Stimuli Leads to a Decreased Dispersion of Eye–Movement Patterns. *Cognitive Science*.
<https://onlinelibrary.wiley.com/doi/abs/10.1111/cogs.12566>
- Król, M., & Król, M. (2019). The world as we know it and the world as it is: Eye–movement patterns reveal decreased use of prior knowledge in individuals with autism. *Autism Research: Official Journal of the International Society for Autism Research, 12*(9), 1386–1398.
<https://doi.org/10.1002/aur.2133>

- Litman, J. A., Collins, R. P., & Spielberger, C. D. (2005). The nature and measurement of sensory curiosity. *Personality and Individual Differences, 39*(6), 1123–1133.
<https://doi.org/10.1016/j.paid.2005.05.001>
- Livson, N. (1967). Towards a differentiated construct of curiosity. *The Journal of Genetic Psychology, 111*(1st Half), 73–84. <https://doi.org/10.1080/00221325.1967.10533749>
- Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin, 116*(1), 75. <http://psycnet.apa.org/record/1994-41058-001>
- Loth, E., Gómez, J. C., & Happé, F. (2010). When seeing depends on knowing: adults with Autism Spectrum Conditions show diminished top-down processes in the visual perception of degraded faces but not degraded objects. *Neuropsychologia, 48*(5), 1227–1236.
<https://doi.org/10.1016/j.neuropsychologia.2009.12.023>
- Ludmer, R., Dudai, Y., & Rubin, N. (2011). Uncovering Camouflage: Amygdala Activation Predicts Long-Term Memory of Induced Perceptual Insight. *Neuron, 69*(5), 1002–1014.
<https://doi.org/10.1016/j.neuron.2011.02.013>
- Lyssenko, N., Redies, C., & Hayn-Leichsenring, G. U. (2016). Evaluating Abstract Art: Relation between Term Usage, Subjective Ratings, Image Properties and Personality Traits. *Frontiers in Psychology, 7*, 973. <https://doi.org/10.3389/fpsyg.2016.00973>
- Marvin, C. B., & Shohamy, D. (2016). Curiosity and reward: Valence predicts choice and information prediction errors enhance learning. *Journal of Experimental Psychology. General, 145*(3), 266–272. <https://doi.org/10.1037/xge0000140>
- McNeish, D., Stapleton, L. M., & Silverman, R. D. (2016). On the Unnecessary Ubiquity of Hierarchical Linear Modeling. *Psychological Methods, 21*(1), 1–17. <https://doi.org/10.1037/met0000078>
- Metcalfé, J., Schwartz, B. L., & Eich, T. S. (2020). Epistemic curiosity and the region of proximal

- learning. *Current Opinion in Behavioral Sciences*, 35, 40–47.
<https://doi.org/10.1016/j.cobeha.2020.06.007>
- Mooney, C. M., & Ferguson, G. A. (1951). A new closure test. *Canadian Journal of Psychology*, 5(3), 129–133. <https://www.ncbi.nlm.nih.gov/pubmed/14870072>
- Muth, C., & Carbon, C.-C. (2013). The Aesthetic Aha: On the pleasure of having insights into Gestalt. *Acta Psychologica*, 144(1), 25–30. <https://doi.org/10.1016/j.actpsy.2013.05.001>
- Muth, C., & Carbon, C.-C. (2016). SeIns: Semantic Instability in Art. *Art & Perception*, 4(1-2), 145–184. <https://doi.org/10.1163/22134913-00002049>
- Muth, C., Raab, M. H., & Carbon, C.-C. (2016). Semantic Stability is More Pleasurable in Unstable Episodic Contexts. On the Relevance of Perceptual Challenge in Art Appreciation. *Frontiers in Human Neuroscience*, 10, 43. <https://doi.org/10.3389/fnhum.2016.00043>
- Nicki, R. M. (1970). The reinforcing effect of uncertainty reduction on a human operant. *Canadian Journal of Psychology/Revue Canadienne de Psychologie*, 24(6), 389–400.
<https://doi.org/10.1037/h0082875>
- Nicki, R. M., Forestell, P., & Short, P. (1979). Uncertainty and preference for “ambiguous” figures, “impossible” figures and the drawings of M. C. Escher1. *Scandinavian Journal of Psychology*, 20(1), 277–281. <https://doi.org/10.1111/j.1467-9450.1979.tb00709.x>
- Pepperell, R. (2006). Seeing without Objects: Visual Indeterminacy and Art. *Leonardo*, 39(5), 394–400. <https://doi.org/10.1162/leon.2006.39.5.394>
- Peterson, M. A. (1994). Object Recognition Processes Can and Do Operate Before Figure–Ground Organization. *Current Directions in Psychological Science*, 3(4), 105–111.
<https://doi.org/10.1111/j.1467-8721.1994.tb00156.x>
- Pinna, B. (2010). New Gestalt principles of perceptual organization: an extension from grouping to

- shape and meaning. *Gestalt Theory*, 32, 1–67.
- Roets, A., & Van Hiel, A. (2011). Item selection and validation of a brief, 15-item version of the Need for Closure Scale. *Personality and Individual Differences*, 50(1), 90–94.
<https://doi.org/10.1016/j.paid.2010.09.004>
- Schmidhuber, J. (2009). Driven by compression progress: A simple principle explains essential aspects of subjective beauty, novelty, surprise, interestingness, attention, curiosity, creativity, art, science, music, jokes. *Anticipatory Behavior in Adaptive Learning Systems*, 48–76.
<http://www.springerlink.com/index/p450676545184762.pdf>
- Schwartenbeck, P., Fitzgerald, T., Dolan, R. J., & Friston, K. (2013). Exploration, novelty, surprise, and free energy minimization. *Frontiers in Psychology*, 4, 710.
<https://doi.org/10.3389/fpsyg.2013.00710>
- Seabold, S., & Perktold, J. (2010). Statsmodels: Econometric and statistical modeling with python. *Proceedings of the 9th Python in Science Conference*, 57, 61.
<https://pdfs.semanticscholar.org/3a27/6417e5350e29cb6bf04ea5a4785601d5a215.pdf>
- Seth, A. K., Millidge, B., Buckley, C. L., & Tschantz, A. (2020). Curious Inferences: Reply to Sun and Firestone on the Dark Room Problem [Review of *Curious Inferences: Reply to Sun and Firestone on the Dark Room Problem*]. *Trends in Cognitive Sciences*, 24(9), 681–683. cell.com.
<https://doi.org/10.1016/j.tics.2020.05.011>
- Silvia, P. J. (2005). Emotional responses to art: From collation and arousal to cognition and emotion. *Review of General Psychology: Journal of Division 1, of the American Psychological Association*, 9(4), 342. <https://pdfs.semanticscholar.org/a87a/95da82793dfe51b6f0a6d9d52c1c5d428f7c.pdf>
- Simonsohn, U. (2018). Two Lines: A Valid Alternative to the Invalid Testing of U-Shaped Relationships With Quadratic Regressions. *Advances in Methods and Practices in Psychological*

- Science*, 1(4), 538–555. <https://doi.org/10.1177/2515245918805755>
- Slamecka, N. J., & Graf, P. (1978). The generation effect: Delineation of a phenomenon. *Journal of Experimental Psychology. Human Learning and Memory*, 4(6), 592–604.
<https://doi.org/10.1037/0278-7393.4.6.592>
- Tanner, J., & Itti, L. (2017). Goal relevance as a quantitative model of human task relevance. *Psychological Review*, 124(2), 168–178. <https://doi.org/10.1037/rev0000053>
- Terwilliger, R. F. (1963). Pattern complexity and affective arousal. *Perceptual and Motor Skills*, 17, 387–395. <https://doi.org/10.2466/pms.1963.17.2.387>
- Van de Cruys, S. (2017). Affective Value in the Predictive Mind. In T. K. Metzinger & W. Wiese (Eds.), *Philosophy and Predictive Processing*. MIND Group.
<https://doi.org/10.15502/9783958573253>
- Van de Cruys, S., Vanmarcke, S., Van de Put, I., & Wagemans, J. (2018). The Use of Prior Knowledge for Perceptual Inference Is Preserved in ASD. *Clinical Psychological Science*, 6(3), 382–393. <https://doi.org/10.1177/2167702617740955>
- Van de Cruys, S., & Wagemans, J. (2011). Putting reward in art: A tentative prediction error account of visual art. *I-Perception*, 2(9), 1035–1062. <https://doi.org/10.1068/i0466aap>
- van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., Boulogne, F., Warner, J. D., Yager, N., Gouillart, E., Yu, T., & scikit-image contributors. (2014). scikit-image: image processing in Python. *PeerJ*, 2, e453. <https://doi.org/10.7717/peerj.453>
- Van Geert, E., & Wagemans, J. (2020). Order, complexity, and aesthetic appreciation. *Psychology of Aesthetics, Creativity, and the Arts*, 14(2), 135. <https://psycnet.apa.org/journals/aca/14/2/135/>
- van Lieshout, L. L. F., Vandenbroucke, A. R. E., Müller, N. C. J., Cools, R., & de Lange, F. P. (2018). Induction and relief of curiosity elicit parietal and frontal activity. *The Journal of*

Neuroscience: The Official Journal of the Society for Neuroscience.

<https://doi.org/10.1523/JNEUROSCI.2816-17.2018>

- Vygotsky, L. S. (1962). *Thought and word*. <https://psycnet.apa.org/record/2006-10268-007>
- Wade, S., & Kidd, C. (2019). The role of prior knowledge and curiosity in learning. *Psychonomic Bulletin & Review*, 26(4), 1377–1387. <https://doi.org/10.3758/s13423-019-01598-6>
- Walker, E. L. (1981). The quest for the inverted U. In *Advances in intrinsic motivation and aesthetics* (pp. 39–70). Springer. https://link.springer.com/chapter/10.1007/978-1-4613-3195-7_3
- Wang, Z., Bovik, A. C., Sheikh, H. R., & Simoncelli, E. P. (2004). Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing: A Publication of the IEEE Signal Processing Society*, 13(4), 600–612. <https://doi.org/10.1109/tip.2003.819861>
- Zeki, S. (2004). The neurology of ambiguity. *Consciousness and Cognition*, 13(1), 173–196. <https://doi.org/10.1016/j.concog.2003.10.003>