To cite this chapter:

Van de Cruys, S., Bervoets, J., & Moors, A. (2022). Preferences need inferences: Learning, valuation, and curiosity in aesthetic experience. In M. Nadal & M. Skow (Eds.), *The Routledge International Handbook of Neuroaesthetics* (pp. 475-506). Routledge. https://doi.org/10.4324/9781003008675-28

26

PREFERENCES NEED INFERENCES

Learning, valuation, and curiosity in aesthetic experience

Sander Van de Cruys, Jo Bervoets and Agnes Moors

More than 40 years ago, pioneering social psychologist Robert Zajonc (1980) published his seminal work titled "Preferences Need No Inferences," in which he argued for the primacy of affect over cognition. Affective evaluation (the preference) comes first, he claimed, and only then do cognitive processes (the inferences) kick in. The central piece of his evidence for this was the *mere exposure* effect: the finding that the mere repeated presentation of a stimulus increases its liking, no matter whether the stimulus is (consciously) perceived or categorized. Bracketing the discussion of the evidence for this effect for now (we will get to that), watch how deceptive the mere exposure concept is in light of what we have come to understand about the perceptual system in the past decades. "Mere exposure" perpetuates the myth that there is some raw sense data to expose ("merely") to the agent and to be "picked up" passively by the brain. The empirical evidence, however, carried by the Bayesian view of perception, abundantly shows that perception is an active, (re)constructive process, biased from the very start by (implicit) hypotheses and expectations. Perceptual illusions are obvious examples here, but they only illustrate the general principle, which is that the visual system makes sense of ambiguous, imprecise sensory information by combining it with priors or probabilistic hypotheses. If "mere perception" already consists of inferential processes, it is impossible for preferences to "need no inferences," as Zajonc had it. Moreover, his thesis is fundamentally at odds with the recent idea that the brain performs *all* its functions through (approximate) Bayesian inference. This influential theory is known as *predictive processing* (also called the *Bayesian brain* or *active inference*). It holds that perception and learning are inference and, perhaps counterintuitively, valuation too.

In this chapter, we will tie these three—inference, learning and valuation—together to find answers to the classical questions of what we look for (curiosity), what we appreciate (aesthetically) and why we prefer what we prefer. We bracket "aesthetically" because we consider art an activity continuous with our more mundane sensorimotor activities and experiences, so we should not expect separate answers with respect to the motivational and affective principles that govern it. Our emphasis throughout will be on valuation—understood as the process of how we come to value, prefer or appreciate—as a function of learning and inference. Importantly, we will focus on so-called non-reinforced preference, meaning preferences that are not due to pairing of stimuli with rewards or punishments (e.g., Boddez et al., 2019). Indeed, one of the counterintuitive aspects of the predictive processing theory is that it has no conventional concept of rewards, goals or values. It eschews the classical sharp schism between the epistemic (beliefs, representations, cognitions) and the conative (desires, preferences, motivations). However, what it puts in its stead is key to tackle traditional

DOI: 10.4324/9781003008675-28

conundrums in the science of aesthetic experience. It forces us to radically rethink how value emerges and to reinterpret conventional strands of thinking on (aesthetic) valuation.

In what follows, we will set up our problem by way of a brief review of the major theories concerning appreciation and curiosity in psychology. Curiosity and appreciation seem an inseparable pair when trying to understand aesthetic experience, and yet, we still miss good integrated accounts. As we will argue in the subsequent section, this is due to the lack of a language for articulating the *beholder's share*, the precise kind of active involvement of the subject with the artwork. Next, we give a brief overview of the state of the art in the predictive processing theory and propose that it provides such a language. The subsequent section on mechanisms of valuation unpacks this and provides the core of our argument. In the final part of the chapter, we will illustrate this with a new take on the findings on mere exposure as well as (aesthetic) valuation in general. The final sections try to broach the most gripping but ineffable capacities of art with the same tack.

A brief history of appreciation and curiosity in psychology

Appreciation

The Gestalt psychologists were among the first to put the question of what we prefer to look at and why we have those perceptual preferences on the scientific agenda (Wagemans et al., 2012). Their answer can be summarized as the *perceptual parsimony* thesis, where "parsimony" comprised such characteristics as order, simplicity and symmetry (e.g., Koffka, 1935). The Gestalt psychologists argued that we tend to organize visual inputs in the simplest, most orderly way, and this is also what we prefer perceptually. While the Gestalt tradition considered parsimony largely determined by the input properties, much later the emphasis was shifted toward parsimony of processing, in the so-called *processing fluency* tradition (e.g., Reber et al., 2004). Many of the same characteristics (order, symmetry, etc.) are thought to increase the ease or fluency of processing, in addition to several factors that did not enter the equation for Gestalt psychologists, such as familiarity (repetition) and prototypicality. This marked an important move towards a subject-dependent definition of parsimony, away from purely stimulus-bound characteristics. Appreciation now is very much dependent on individual processing and learning.

Of course, the processing fluency theory was proposed in large part to explain the swath of studies on the mere exposure effect (Bornstein, 1989; Zajonc, 1968), under the assumption that, with repeated encounters, the processing of a stimulus happens more fluently. To further explain the fluency phenomenon, and to reconcile it with the work on reinforcement learning, the positive affective mark of fluency is sometimes seen as a safety signal, a quelling of our assumed innate fear of the unknown with repeated "uneventful" experiences (i.e., fluency as learned inhibition of fear; e.g., Winkielman et al., 2003). Another explanation of the positive effect of fluency considers it a signal that one has been able to (cognitively) deal with the stimulus in the past. Here, liking would be the result of misattributing a characteristic of the subjective processing (i.e., the ease) to the external stimulus. It is the stimulus that is liked, even though what is monitored and appraised is a property of processing it. The liking of fluency is then a metacognitive signal of processing quality (Alter & Oppenheimer, 2009): how successful I am (or have I been) in processing the current inputs. We will return to this important idea later on.

Empirical findings have cast serious doubt on the parsimony thesis, whether in the "objective" form proposed by the Gestalt tradition or in the "subjective" fluency/mere exposure tradition. People often seem to prefer or appreciate medium complexity or medium orderliness (termed the *Goldilocks principle*) rather than the most fluently processed, ordered or simple stimuli. This is of course apparent in art (Figure 26.1), where certain violations of familiarity or order are often conducive to aesthetic appreciation, but it also holds in everyday life and controlled experiments. For example, a large meta-analysis of mere exposure studies (Montoya et al., 2017) found that while appreciation increases with repeated presentations, most studies also found an

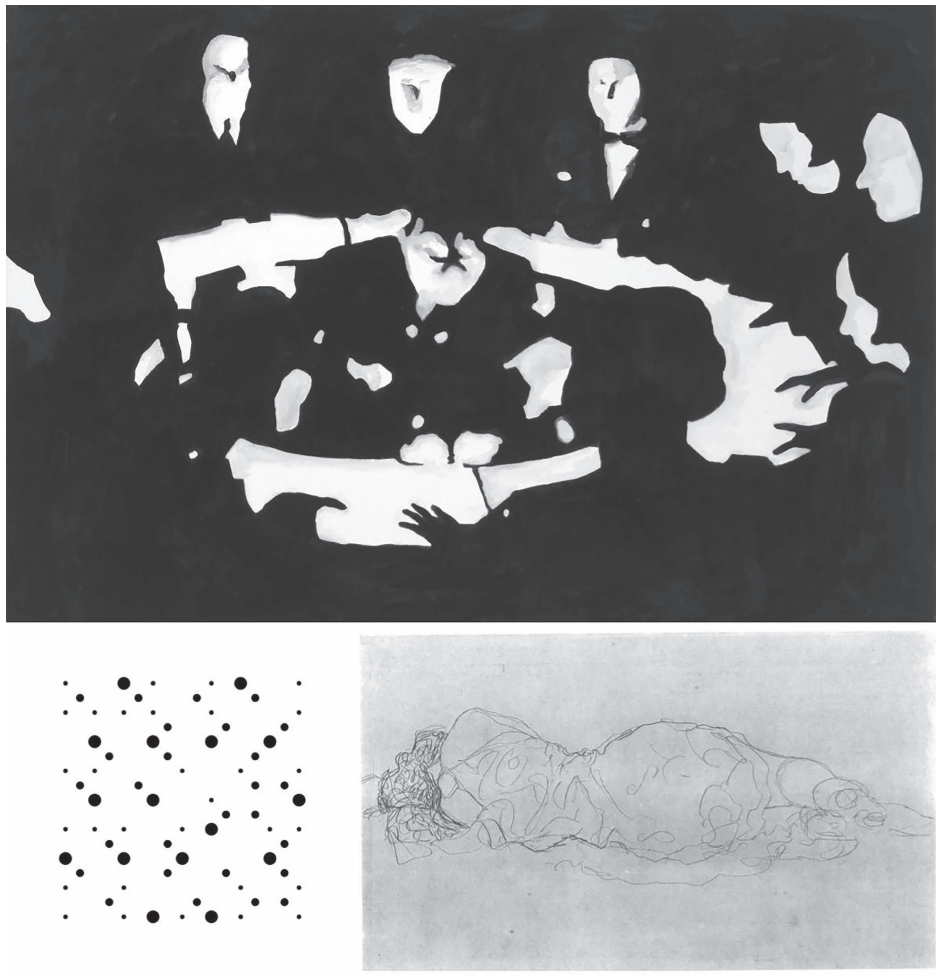

Figure 26.1 Art, a prime example of appreciated stimuli, often breaks order or simplicity. For example, in these works by Jan Vanriet (top; *Séance*, oil on canvas) and Gustav Klimt (bottom right; *Reclining Woman*) familiar shapes are not depicted in their simplest, most recognizable forms. Similarly, the computer rendition by Neil Dodgson (2009) inspired by Bridget Riley's op-art (bottom left), includes clear violations of symmetry. The work by Vanriet is reprinted with permission from the artist. The drawing by Klimt is reproduced from Wikipedia Commons (https://commons.wikimedia.org/wiki/File:Klimt_-_Ausgestreckte_Frau_001.jpg). The work by Dodgson is reprinted with permission from the artist.

inverted U-curve function, indicating that there is an optimum in liking for a medium number of repetitions, after which liking goes down again (sometimes described as a boredom effect; Bornstein et al., 1990). Of course, this directly links to Berlyne's (1960) work on preference for stimuli of intermediate complexity (see also Chmiel & Schubert, 2017).

Walker (1981, p. 40) summarizes Berlyne's theory in two basic postulates: "1. There is an optimal level of psychological complexity for a psychological event that will be preferred to either simpler or more complex events. 2. Repeated experience of an event will lead to progressive simplification of that event," thereby making the connection between repeated exposure and complexity (simplicity; cf. Gestalt school) or fluency

explicit. Walker goes on to review the evidence for both postulates, concluding that there is reasonable support for both but that many methodological challenges hamper the empirical identification of an inverted U-curve (provided there is one). For example, stimulus sets or contexts usually have a strong effect, where appreciation (and complexity) ratings greatly depend on the particular number, order (serial dependency) and types of stimuli in the test set. This problem is compounded by possible nonlinearities in both the complexity and the liking dimension and by the fact that people do not even seem to process all information in complex stimuli. Together these factors bias how much of the complexity continuum participants have experienced. Even more problematic, participants may be evaluating different things (from the intended dimensions), possibly depending on the language used by the experimenter when describing the task of rating complexity or appreciation. One counterintuitive consequence of all these confounds is that an (partial or shifted) inverted U-curve pattern could be present in all of the participants but not (or even a regular, non-inverted U-curve) in the shape of the average data, or vice versa (see Güçlütürk et al., 2016; Spehar et al., 2016; Walker, 1981).

Underlying these problems, is, as so often in psychology, the *mere measurement effect* clouding possible U-curve relations. That is, the very fact of presenting, processing and measuring will substantially influence the measurement. Every repeat of a stimulus with a certain level of complexity will in fact be a different event, with a different psychological complexity, so averaging, even across the same stimulus (level), may bias the shape of the curve. Indeed, the second postulate—progressive simplification with experience—urges us to pay attention to *subjective* complexity/simplicity for a particular individual with her particular learning history. Nonetheless, many early studies (e.g., Terwilliger, 1963) on the relation between complexity and liking relied on objective complexity (measured for example as the number of angles in geometric shapes).

More recent elegant experiments in young infants managed to improve on this by quantifying (subjective) complexity based on the presentation history of stimuli of a given participant (Kidd et al., 2012, 2014). With this more individualized, computational approach, the authors could confirm a preference for intermediate complexity (or predictability) in visual and auditory sequences. In parallel, studies on semantic-level complexity also underline the need to account for subjective complexity that is not computable based on stimulus features alone. For example, Nicki (1970) found that people prefer photographs of objects with intermediate blur rather than lower or higher blur but that this is likely due to the fact that images with intermediate blur elicit a higher number of guesses about the content of those images, with even confidence about the guesses (see also Van de Cruys et al., 2021). In other words, while intermediate uncertainty or blur seemed to be preferred, this was actually the consequence of a preference for the *most* (semantically) unpredictable or uncertain images. Indeed, research in artworks shows that low-level image statistics of complexity or uncertainty explain little of the variance in appreciation (Van Geert & Wagemans, 2020) but that semantic uncertainty (ambiguity) is a crucial factor in appreciation (Muth et al., 2016; Wang et al., 2020).

In recent years, the importance of taking into account the individual learning history has been bolstered by studies showing that particularly *relative* fluency or exposure matters for liking rather than the absolute, momentary fluency of a stimulus or trial. For example, Wänke and Hansen (2015) have shown that the typical mere exposure effect could be replicated only when the old stimuli were mixed with new stimuli. Also emphasizing individual learning, Forster et al. (2015) showed that an effect of ease of processing could only be found in within-participant (not between-subject) comparisons. Similarly, relative exposure, that is, being presented more often than other stimuli in the set, increases liking more than absolute exposure (Mrkva & Van Boven, 2020). Together, such findings suggest that participants implicitly track the expected fluency across the recent history of experiences and attribute value to deviations from this expectation, a notion that will become important later on.

An emphasis on learning and processing dynamics may hold the key to a reconciliation of the "simplicity-fluency" camp and the "intermediate complexity" camp as well. In particular, one could describe what perceivers like in terms of the subjective progress made in dealing with the stimulus (or activity), in other words, a reduction in disfluency (Graf & Landwehr, 2015; Muth et al., 2015; Schmidhuber, 2009; Van de Cruys &

Wagemans, 2011) or an increase in the system's compression of a given stimulus, going from unpredictable (or complex) to predictable (or simple). If it seems as if humans like simple stimuli, we are confusing a static end product with the crucial process. If it seems as if we prefer medium complexity, we are confusing a static starting point with the crucial process (Van de Cruys, 2018) because the chance of learning progress is the highest for environments of medium subjective complexity. Under those conditions, we have some idea as to what type of regularities rule the situation or environment, but a level of active (mental) engagement or learning is required to resolve the remaining uncertainty.

Curiosity

While curiosity and intrinsic motivation clearly direct our engagement with and shape our appreciation of art, few behavioural studies have explicitly examined the interplay of curiosity and appreciation in the aesthetic experience. This is even more remarkable given that discussions on appreciation are partly mirrored in those about curiosity, especially with regard to the role of complexity and uncertainty. This may in part be due to the fact that the same dependent measures, namely choice or preference behaviour, have been used to infer participants' appreciation in some studies and their curiosity in other studies, so they are hard to disentangle in practice, even though they are cognitively and affectively clearly distinct. For good measure, we understand curiosity as "a tendency or motive to acquire or transform information under circumstances that offer no immediate adaptive value for such activity" (Livson, 1967, p. 75). It is a form of motivation that is "inherent within information processing and action" (Hunt, 1981). Broadly speaking, there are four (types of) theories of curiosity. *Optimal level theories* (e.g., Berlyne, 1978) argue, in parallel with the inverted U-curve discussion for appreciation, that medium complexity or uncertainty elicits the most curiosity. Although this is a descriptive theory, Berlyne proposed that medium complexity corresponds to an optimal, medium arousal that animals would gravitate towards in a homeostatic sense. However, it remains a puzzle why a system would strive to be in a medium arousal state in the first place, and, indeed, subsequent work has largely discredited the connection with arousal in Berlyne's thesis (Silvia, 2005; Walker, 1981).

The second and currently most popular theory of curiosity is Loewenstein's (1994) account of curiosity as our sense of a gap in our knowledge, quantified using information theory (as the entropy, i.e., the number of alternative options and their probability). The gap theory clearly captures something important, namely that curiosity can be an increasing function of uncertainty *relative* to one's current knowledge and expectations rather than reaching a peak for intermediate uncertainty, something that later empirical studies have also confirmed (e.g., Van de Cruys et al., 2021; van Lieshout et al., 2018). However, the computational tools of classical information theory (entropy) are limited in truly shedding light on curiosity. For example, entropy is only defined once particular predictions (or questions) are formulated, but Loewenstein provides few clues on how our system generates those. Another problem in Loewenstein's theory stems from his framing of the curiosity gap as a classical drive: something that needs to be "removed" by action. Because people want to resolve it, Loewenstein (and many others after him) infer curiosity must be aversive. In the same contorted way, Freud counted sexual arousal among the aversive drives because the orgasm terminates the arousal (Hunt, 1981). Obviously, curiosity, like sex, is usually pleasurable and appetitive, even though the uncertainty at its core (the "gap") is indeed often a cause or catalyst of anxiety (Hirsh et al., 2012). So what is missing here? When is uncertainty not a bad thing?

The missing element may very well be provided by a third theory on curiosity that is based on the appraisal theory of emotions. Silvia (2005) proposed that curiosity (or "interest," as he prefers to call the emotion associated with intrinsic motivation and learning) is the result of a combination of two different implicit cognitive appraisals. The first is the novelty appraisal, which, analogously to Berlyne and Loewenstein, is an evaluation of whether something is new, unfamiliar, uncertain, complex or unexpected (i.e., the "gap"). The second crucial ingredient is an appraisal of coping potential. It involves an estimation of one's ability to

understand or deal with the new, unexpected event, in the sense of rendering it predictable or meaningful again. Like (relative) fluency previously, this is a metacognitive evaluation, but here it is about the extent to which one expects to be able to reduce uncertainty. One might describe "promised insight" (Muth et al., 2015) as an *expected* instead of an actual fluency. As long as there is this component to curiosity, uncertainty can be a source of joy and "approach" behaviours, instead of just anxiety. Although we may be more likely to be able to reduce uncertainty for medium levels of uncertainty, there is nothing that strictly binds curiosity to just this level of uncertainty.

A fourth and final strand of theorizing on curiosity (Keller & Voss, 1983; Livson, 1967; McReynolds, 1971) dates from the 60s but has been largely forgotten, at least in comparison to Berlyne's influential work from the same era. Based on animal research, McReynolds (1971) developed a view of curiosity as "an *expected rate* of cognitive structuring that an individual tries to *maintain* by exploration (or lack thereof)." *Cognitive structuring* is a somewhat dated term, but McReynolds used it to refer to the process of assimilation by which new percepts are made to fit with the existing mental schemata or models. Like Berlyne's theory, this view of curiosity has a strong adaptive, homeostatic flavour to it, not so much in the physiological domain (cf. Berlyne's arousal) but in the information processing domain. Originally, McReynolds used the term *perceptualization rate* to denote the expected rate of structuring, inspired by Glanzer (1958) who described the organism as "an information processing system that requires certain amounts of information per unit time." However, later on, McReynolds perceptively remarks that the essential variable to be tracked and regulated is not the stimulation or some "raw" information rate but rather "the rate of cognitive structural change" (McReynolds, 1971, p. 37). Hence, he preferred the concept of *innovation rate* centring on the assimilation or learning process. The system's core concern is to minimize "unassimilated material"—the uncertain, unexpected inputs—and to optimize "innovation rate." In more modern terms, curiosity is driven by the expected rate of updating one's mental models (the expectation of uncertainty reduction relative to one's mental models). In this way, the expected rate-based explanation of curiosity further specifies the previous account centred on coping potential, or expected reducibility of uncertainty (i.e., a positive rate).

Using examples from rodent research, McReynolds reasoned that both long-term experience (e.g., "rearing environment") and short-term experience (cf. deprivation studies) could contribute to forming those expected or "accustomed" rates of processing. If organisms form adaptive expectations on the rate of "structuring" or uncertainty reduction that they try to maintain, this would require active exploration (as the expression of curiosity) to fulfil: "an active seeking for new data to be digested" (McReynolds, 1971). It is easy to see our creative, artistic explorations as (one of) the system's efforts to upregulate the rate to match the expected one. A concrete recent example was provided by the COVID-19 pandemic lockdown, during which people apparently en masse took up learning to play an instrument (Hill, 2021). At the other end of the spectrum, some environments may raise uncertainty levels too quickly too much, which also prevents expected structuring rates from being re-established. Instead of exploration, here one would see (anxious) avoidance or repetitive, rigidly structured (stereotypical) forms of self-stimulation to reinstate the expected rates of uncertainty reduction (as also found in clinical disorders; e.g., Van de Cruys et al., 2014). The parsimony of McReynolds' view of anxiety and curiosity (and their behavioural expressions) as two sides of the same coin is attractive and intuitively plausible and will become important for aesthetics as well.

Let us wrap up this historical overview by asking why we need a sense of curiosity in the first place. A plausible answer is that we need it because learning is costly and fallible, so being able to sensitively direct resources to where the best learning progress can be made has a considerable advantage for the system. It entails learning not only the "subject matter" but also (meta-)learning what is learnable based on how your mental models performed in similar environments (Gottlieb et al., 2013). Indeed, even young infants (in artificial grammar learning) seem to already have a sense of where learnable inputs are situated rather than just noise (Gerken et al., 2011), here investigated through a looking times paradigm. Hence, curiosity and exploration can ensure that we remain in what Vygotsky (1962) has called the zone of proximal development

(Metcalfe et al., 2020; Oudeyer et al., 2007), the optimal region for learning just above one's current mental models or abilities. Specifically, it allows the agent to avoid wasting resources on inputs that are either mere noise (nothing to learn), too easy (all regularities are already learned) or too complex (for the current mental models). One interesting way to implement this is through contextually adjusted expected rates of uncertainty reduction, as McReynolds envisaged.

The quest for the beholder's share

However, an underlying reason for the lack of work on the interplay of curiosity and appreciation is a lack of a conceptual language to bring this to the fore. Most accounts of appreciation have been applications of classical information theory and so adopted the vernacular of a passive receiver of a static artwork (the sender's "message"), leaving little room for the proactive involvement of the subject. While information theory has clearly led to crucial insights (cf. Loewenstein), it also comes with serious limitations in articulating the beholder's share (Gombrich, 1963; Seth, 2019). It misses a formal account of how the expectations are built through learning and how they are brought to bear in experience. Indeed, if our goal is to explain *aesthetic* experience, our analysis suggests those expectations are formed not only about the content level but also about the meta-level of the *content-generating* processes (e.g., expected fluency or rates of uncertainty reduction).

It must be said that the active involvement of the subject and the role of internal dynamics in the aesthetic experience have not been overlooked by philosophers and psychologists working on the subject (e.g., Dewey, 2005; Hume, 2008). Notably, in 1942 Eysenck published his "law of aesthetic appreciation," with great ambition. The law states that "The pleasure derived from a percept as such is directly proportional to the decrease of energy capable of doing work in the total nervous system, as compared with the original state of the whole system" (Eysenck, 1942). There are three notable things in Eysenck's law. First, aesthetic pleasure is about the frugal use of energetic resources. Second, aesthetic appreciation is *subjective* in that the energy changes are measured against a given system's prior state. Third, the law is decidedly dynamic in nature. Eysenck was clearly inspired by the early Gestalt psychologists like Köhler here, who had a pretty dynamic view on perceptual experience already (using the language of fields and equilibrium processes; Köhler, 1920; Wagemans, 2015). However, it failed to gain traction, as was also the case for Eysenck's law. Indeed, as we sketched earlier, the mistaken picture that survived from the Gestalt tradition is one centred on *static* order and simplicity.

It is tempting to blame this disregard for internal dynamics on the cognitive revolution born in the slipstream of the rise of information theory and heralding the computer metaphor of the mind. Indeed, some authors (e.g., Cupchik & Heinrichs, 1981) lament that the static, atomistic approach of tracking objective probabilities of occurrence of singular elements or events in an artistic "stimulus" (as per information theory) has limited import in understanding (aesthetic) appreciation. For example, Cupchik and Heinrichs (1981) complain that "The notion that dabs of colour have a quantifiable uncertainty of appearing next to each other does not reflect the structure and creative origin of the work" and that "information as a quantity is different and independent from the meaning of a work of art." On the face of it, this critique crushes the hope that one could come up with a broadly applicable, formal account of aesthetic appreciation. However, in the next sections, we will discuss the ways in which the predictive processing account helps to bridge the distance from mere "atomistic" information to meaning. Information theory will remain central, although on its own it indeed falls short in capturing the relevant processes.

That brings us to the actual reason that the dynamic view of Köhler, Eysenck and others (see also Arnheim, 1974; Pepperell, 2018) did not catch on at the time: They failed to articulate how their metaphoric theorizing on forces, fields or energy dynamics mapped onto tractable psychological or neurophysiological variables. For example, Eysenck (1942) writes: those external stimuli will be judged the most beautiful which

are most in agreement with the *internal forces* of perception" (our emphasis). But what are those internal forces of perception? Have we made any progress in the last decades that could give us another shot at those richer accounts of aesthetics of the days before the cognitive turn? We still do not have a good grasp on the coding scheme, the "internal forces," of the brain, which would be necessary to determine subjective (brain-based) energetic or processing gains across time. However, one plausible, well-elaborated candidate for this coding scheme is provided by hierarchical predictive processing.

Predictive processing in brief

Producing future

According to poet and philosopher Paul Valéry, the purpose of the brain is to produce future. In recent decades, that very idea has been computationally fleshed out in the theory of predictive processing, known under different guises such as the *Bayesian brain* or *active inference* (Clark, 2013; Friston, 2010; Hohwy, 2020). To "produce future" efficiently means not just to respond to stimuli but to proactively model and predict sensory inputs and the opportunities and challenges they represent. In its most basic sense, being (alive) means creating those conditions that allow one to persist into the future and not dissipate under the second law of thermodynamics. In information-theoretic terms, we can say that there is a limited set of expected or "target" states that is consistent with continued existence and so has to be maintained (cf. homeostasis). More concretely, an organism has become equipped through evolution with particular interoceptive expectations (e.g., blood glucose or temperature at a particular level), which evolution has "discovered" to be consistent with survival. For the most part, it cannot fulfil those interoceptive expectations on its own but only by going via its environment (e.g., finding a food source) (Pezzulo et al., 2015).

To do so, the organism must possess or learn a (minimal) generative model of its environment. It is called "generative" because it is a model about how the expected (interoceptive) observations can be generated or recreated. We only have the sensory effects of the processes in the world around us to go on, but from the regularities in those sensory inputs, we can infer their likely hidden causes. So we will experience a world populated with objects such as rocks, animals and clouds. The pattern of inputs created by the cloud on our senses is more alike than that of not-cloud, for example, that pattern follows a similar motion on the retina (cf. grouping by "common fate"). The cloud "object" explains such input patterns, as well as the exact pattern of darkness and temperature over time (and space) that I sense when the cloud covers the sun, and so on. Ultimately, we see clouds and other hidden causes because it helps us predict exteroceptive (e.g., when the sky turns dark grey, it's likely to rain soon) and interoceptive states (e.g., I will get cold).

To turn events to our advantage and actually realize our target interoceptive states, we need to learn not only about our surroundings but also about ourselves as a hidden cause. We are just another inferred hidden cause in the world, albeit one with which we have very intimate and rich experience. By the kind of stable correlated patterns in our different senses that our actions tend to cause, we discover who we are and how our own actions can accomplish expected interoceptive observations (e.g., when putting on a raincoat, I will not get cold). The same principle that allows us to see clouds and other objects—stable patterns of sensory inputs can be used to successfully predict our surroundings—readily extends to the intentions and goals that we perceive in ourselves and others: These inferred mental causes are as real as clouds or other objects we perceive (Dennett, 1991), though they are situated at a hierarchically higher (or "deeper") level summarizing larger patterns in inputs. This does not mean that social perception is any less direct. We will perceive coordinated motions of limbs immediately as intentions, just because of the predictive power of such inferred causes that abstract away the concrete minor variations in the constituent motions.

We experience the world as populated with objects and intentions, because, irrespective of their veridicality, this is the interface that works best in navigating our world and fulfilling our interoceptive expectations.

By continually predicting incoming inputs based on current context, we actively construct our interface to the world around us. Only in our failures, the mismatch between our constructed pattern and the observed inputs, do we meet reality (von Glasersfeld, 1995). It is those prediction errors that keep our constructs in check, so they continue to "work." Input activity that is correctly predicted is suppressed (explained away), while the errors are sent upwards to update predictions. This allows the system to reserve its resources for unexpected observations (prediction error) while building increasingly efficient reconstructions of those observations. By minimizing prediction errors, iterated across levels in the brain, we infer latent causes that best explain the regularities in impinging inputs. Mathematically, the process is a form of gradient descent optimization scheme, often described as descending an error "landscape" to find its minima (for computational details, see Aitchison & Lengyel, 2017; Gershman, 2019; G. B. Keller & Mrsic-Flogel, 2018; Smith et al., 2021; Spratling, 2017).

Both perception (inference) and learning can be captured with this (approximate Bayesian) predictive updating scheme. The speed and automaticity with which the visual system uses its generative model to disambiguate the "hidden causes" of the sensory inputs and settles on the best explanation (e.g., this is figure, that is ground) hides the underlying inferential process. We have to severely distort images, as we do to create so-called Mooney images (Figure 26.2), to obstruct this process and block the use of our learned hierarchical models. Even then, a very brief priming of the right semantic model is often sufficient to cause a dramatic shift in how we see the image and as well as a near inability to unsee the right solution. The newly inferred cause seems so strongly supported by image cues that we cannot but assume it was always there in the image and not just in our experience.

The error-based updating scheme is repeated at each level of the cortical processing hierarchy, with each higher level trying to predict activity in the region below. On a local level, "predicting" should be understood in the technical sense of reconstructing input activity (from the lower region) quasi-simultaneously, broader than the everyday notion of prediction as anticipation. However, across the hierarchy, predictions are formed capturing statistical regularities that span more and more space and time. Take the example of reading a book. While reading, we form expectations on event sequences across time in a story, which in turn



Figure 26.2 Two-tone or so-called Mooney images are created by blurring and thresholding greyscale photographs (see the source photograph in Figure 26.3). They are examples of one-shot learning: Once you find or are confronted with the solution, you cannot unsee it. "Discovery" of the familiar structure in the image usually gives a positive feeling of insight or *Aha-Erlebnis*. Notice the similarity of this type of image degradation with the technique used in the artworks in Figure 26.1.

creates expectations on high-level meanings that we will likely encounter, as well as the kind of words that may appear. Reading a sentence, we have particular expectations about which words will likely appear in which locations (based on learned syntax and semantics), which can further be unpacked in expected letters and their compositing features (oriented lines). The hierarchy forms a cascade of interdependent predictions going all the way down to the peripheral level where predictions can be "answered" at the level of the retina (or receptors of other senses). Recent evidence supports the existence of such a predictive hierarchy in the brain (Heilbron et al., 2021; Rohe & Noppeney, 2015; Wacongne et al., 2011). Using these predictions, we exploit the redundancies in reading materials, which allows us to be more selective and efficient in our sampling of the inputs. Indeed, eye movement patterns show we jump from word part to word part in a sentence, without having to "read" each letter (Karl J. Friston et al., 2018). The whole process happens largely flawlessly, though typos may be missed (especially when we edit our own text, because of strong and accurate predictions) and unexpected "turns of sentences" might need to be read twice.

Part of the act

So far, we have only covered one way of reducing prediction errors, namely by changing your predictions or our models, known as perceptual inference and learning. The other, complementary way to minimize prediction errors is by changing the things predicted: by acting to bring the world (our observations) closer to your models of it. Hence, adaptations to improve the fit between world and mind can have two different directions: *world-to-mind* (assimilation) or *mind-to-world* (accommodation). This closes the perception-action cycle at the core of this framework, using "fit" as quantified by prediction errors as the unified optimization criterion.

Actions can be treated with the same predictive machinery described previously, when they are conceptualized not as motor commands but rather as their expected sensory consequences, in proprioception (the state of our muscles and tendons) and exteroception (e.g., feeling the cup in my hand). This idea goes back to classical work on ideomotor theory (James, 1890) and the view of behaviour as the control of perception (Powers, 1973). Like in perception, action models are hierarchically organized with abstract goals and intentions at the top to be unpacked into policies (action sequences) and further into specific motor programs represented as the exteroceptive and proprioceptive outcomes that they are expected to realize. At the lowest,



Figure 26.3 A photograph of a frog that has been used to create Figure 26.2.

peripheral levels, proprioceptive predictions are confronted with the current state (before any movement) of muscle and tendon receptors to form prediction errors. Instead of leading to model updates, here prediction errors trigger classical reflex arcs executing the movement (K. J. Friston et al., 2010). The generative action models can be exploited for action *execution*, that is, inferring the action sequences that accomplish intended effects. However, in line with the mirror neuron system (Kilner et al., 2007), the same model can be applied for action *inference* as well, that is, inferring your own or others' intentions from sensory effects of actions—a framework that is now referred to as *active inference*.

One important role for actions is foraging for more information (conversely, more information enriches our action capabilities). Indeed, predicting well also means predicting when one *cannot* reliably predict something—when more sensory sampling is needed (e.g., through head and/or eye movements). We continuously fine-tune our rhythm and direction of sampling based on the current context in the service of better prediction error minimization. For example, when reading a difficult text (or one in a foreign language), our predictions do not provide a lot of support, so reading becomes effortful (i.e., with frequent model updates) and sampling intensive. But to actively direct our sampling of the world in an optimal way, we need to have a way of gauging the quality of the evidence and of our predictions. Those "quality predictions" are called expected *precisions*, which, in essence, are estimates of uncertainty or reliability attached to observations and our first-order predictions. Since precisions are (meta-)predictions as well (predictions of prediction errors), they can be updated in the same way as the "regular" predictions. The role of precision in inference is to regulate the relative influence of new evidence (prediction errors) versus the predictions, for example, to allow top-down predictions to dominate our percept when sensory information is noisy (unreliable). Estimating the quality (precision) of upcoming information also allows us to direct actions to seek out informative (precise) prediction errors: Those future prediction errors that are relevant and reducible through actions or model updates.

Here, we quickly need to dispel the criticism often raised at predictive processing that it would lead to immobilization instead of action: An organism driven by prediction error minimization would retreat from the world into a dark room where all its errors will eternally be minimized (if you just predict darkness) (Seth et al., 2020; Sun & Firestone, 2020; Van de Cruys et al., 2020). It should be clear by now, however, that minimizing prediction errors on those lowest sensory levels does nothing to resolve the multilevel and multimodal expectations that an organism holds. We have seen that organisms often have to build a world of higher-level constructs such as objects and intentions to be able to fulfil their interoceptive target states. Staying in a dark room will typically not remain "expected" for very long, given the particular characteristic states (the phenotype) of an organism (Friston et al., 2012). This implies that to ensure prediction error minimization in the longer run, we need to seek out prediction errors in the here and now, especially the most informative (high-precision) ones with regard to our preferred states.

The "dark room" confusion is born out of the misconception of a mental model as a purely epistemic thing, a search for correspondence to some truth out in the world. Contrary to this conception, remember that the generative model as used in predictive processing doubles as a specification of our preferences. To be viable, our model has to be equipped with some prior expectations about the states we will tend to end up in. That means our needs and goals will be described as an expectation (technically a probabilistic prior distribution) on observations; for example, we strongly (i.e., with high precision) expect to sense a temperature between certain parameter values. In other words, our model has to be *optimistically biased* (Sharot, 2011), in that it might not represent the current state of the world but rather a state that can, fallibly, be attained by one's actions. Indeed, pure optimistic hallucination of more favourable conditions without a generative model of how to reliably accomplish those conditions would quickly be the death of a system. Similarly, a perfectly accurate, exhaustive representation without concern for (the intrinsic regularities that reproduce the limited resources of) the very system that does the representing would drive the system to self-destruction. Any perceptual system can only exist if it translates sensory inputs into a format that says

something about the further persistence of that perceiving system. This format need not necessarily be completely veridical (Hoffman et al., 2015; Tschantz et al., 2020), as long as it works: If it infers those hidden causes and actions that allow it to efficiently "generate" its preferred or expected states and so minimize prediction errors. This makes room for a *controlled optimism* (Van de Cruys et al., 2020) that is continually "negotiated" with the world.

With that in mind, we can see that the possibility to act in the world truly urges the organism to become future-oriented. Being able to choose actions shifts the focus from just the current, momentary prediction errors to *expected prediction errors* in the future and how our actions can be chosen optimally to resolve them. Here, we take our models offline as it were, and we create "artificial" prediction errors by confronting our preferred outcomes with the expected outcomes based on our action models. We use temporally extended ("deep") counterfactual models to ask "What if I do this, instead of that, how well does that reduce my expected prediction errors (uncertainty)?" (Friston et al., 2021). We play out different models against each other (largely implicitly), and we let our models or hypotheses die in our stead (Dennett, 2017). We choose our actions such that they minimize the divergence (i.e., prediction errors, under some assumptions) between two probability distributions (models): one describing our expected outcomes if we were to follow a course of action and another describing our desired future (our prior or preferred outcomes).

To do this, we of course first need to resolve uncertainty or confidence about which actions lead to which outcomes. So in choosing actions, the balance that an organism needs to strike entails consideration of both epistemic value and pragmatic value ("rewards" or prior preferences) on *equal footing* (Friston et al., 2015). Maximizing epistemic value (or information gain) means exploring the world to disclose its structure, to reduce future prediction errors. Once this structure is known, one can exploit it to realize one's prior preferences or expected observations (i.e., pragmatic value). Indeed, this intuition is vindicated by mathematical treatments of active inference which show that minimizing uncertainty (expected prediction errors) relative to the generative model can be decomposed into maximizing epistemic and pragmatic (goal-fulfilling or reward-maximizing) value (Friston et al., 2015).

Psychologists in the 1970s have done experiments where people had to choose between tasks with a range of different difficulty levels (Schneider & Heckhausen, 1981; Trope, 1975). It turned out people choose intermediate difficulty levels (somewhat biased in the optimistic direction, i.e., expecting success levels slightly higher than their objective performance level), where subjective uncertainty with respect to task outcome is greatest: I am sure I will win when I pick low difficulty and I will lose when I pick high difficulty. Winning did not gain them anything in this task, so they chose the most uncertain option. Indeed, when (pragmatic) value is equivocal, epistemic value determines action, according to predictive processing. You first learn (to reduce uncertainty about) the structure of the world; then you will be able to exploit that structure to generate the states of your strongest priors (preferences).

A sceptic might object that it is overly reductionist to posit that the mind (and all of our behaviour) is governed by the minimization of uncertainty. It defies our intuitions, as well as the history of psychology, where conative constructs like needs, desires and goals have always taken a central role. However, as we saw, conative constructs are not made obsolete by predictive processing; they are merely absorbed in the model that the organism embodies (Clark, 2020). In this view, what we colloquially call goals can be recast as stable expectations in the sense that they usually have a longer temporal horizon, and a higher precision, meaning that they are more robust to negating evidence, so they will be realized by actions (using precise action models) instead of prediction updating. To put it more positively, minimizing uncertainty is equivalent to maximizing the evidence for the model that defines us, often described as *self-evidencing* (Hohwy, 2014). Generative models are not purely representational but bent towards desired states of the organism. It is a record of representations (hidden states) that best explained past inputs, as well as a target to be realized in future inputs with our actions (Hafner et al., 2020). The same model is used for both inference and action.

Framed like this, uncertainty minimization is "merely" the underlying mechanism by which we accomplish our needs and goals.

Mechanisms of valuation

Beyond mere exposure

With this brief overview of the predictive processing account, we can begin to consider why it is particularly well equipped to address our questions about appreciation and curiosity. In the past 10 years, several authors made important advances in articulating the beholder's share using the framework, as well as in exploring the mechanisms of affect generation in aesthetic experience (in music and visual art) (Forster, 2019; Kesner, 2014; Koelsch et al., 2019; Van de Cruys & Wagemans, 2011). In the remainder of this chapter, we sketch the main lines of reasoning and zoom in on the mechanisms of valuation as they follow from predictive processing.

As prefigured in the introduction, "mere exposure" theory is rendered meaningless in a predictive or inferential view of perception. Instead, with more experience, we form better representations in the sense that our generative models will explain (similar) inputs better. Capturing regularities in the inputs better means that those inputs become subjectively simpler. Thanks to our adaptive, predictive models, commonly occurring inputs also become least metabolically costly (Sengupta et al., 2013), simply because they gradually elicit fewer prediction errors. So increased predictive matching of a stimulus naturally goes along with progressive simplification or increasingly efficient (compressed) coding of what appears often in the agent's environment. In this way, the account ties together predictability, (subjective) complexity and frugality.

So when we find that people like what is likely ("mere exposure"), typical or average, we should say that they like the sensory inputs that their generative models can easily reconstruct (Brielmann & Dayan, 2022). This pleasantness of "generative success" can get quite literal: Chamberlain et al. (2021) showed that drawings created by natural, human-like movement dynamics are preferred, at least by viewers with drawing experience (and thus the necessary generative models). In a similar vein, we tend to prefer stimuli with the same low-level statistical regularities (e.g., with respect to colour composition or 1/f spatial frequency spectrum) as natural images (Graham & Redies, 2010; Nascimento et al., 2021), but this principle presumably generalizes it to regularities on any level of abstraction (in the generative model). Additional support for this comes from the finding that stimuli that are gazed upon more tend to be preferred, even though gaze is manipulated by the experimenter, and there is no differential reward history whatsoever (Schonberg & Katz, 2020). Of course, all of this is consistent with processing fluency ideas, but we now have a mechanistic way to conceptualize the beholder's share (and how it evolves).

Since actions are naturally part of the predictive models, even regular action patterns (habits in layman's terms) can become intrinsically preferred patterns (as they indeed often do) following the same rationale. Indeed, an unexpected *in*ability to perform one's habits (prediction error) can arouse intensely negative emotions. But in isolation, habits do not normally come with intensely positive appreciations, similar to how the most frequent or predictable stimuli are usually not liked much, except in special cases like faces (Ryali et al., 2020) or wallpapers. The preference for regular wallpapers and average faces may have more to do with the fact that we just want those to be the clean canvases for the actually interesting matter: social expressions or interactions. This is similar to how habits just need to become the neutral background against which our more burning deliberations can play out (e.g., wearing the same type of clothes every day to free up cognitive space for more important decisions).

And yet we often take pleasure in returning to our familiar routines, for example, after a long, trying day, even though there is nothing "objectively" rewarding about them. This "warm glow of familiarity" is not

necessarily just an expression of conservative taste. Remember that our whole generative model is built up to fulfil our visceral predictions. It should not surprise us then that enacting the regimes inscribed with the most precision in our generative model can regulate our homeostasis best. In some cases the force of habit is so strong that even "objectively" painful stimuli can become attractive because they have become part of one's entrenched behavioural patterns.

The more general lesson here is that once we describe our preferences as prior expectations for particular observations, they no longer need to be pre-specified genetically (indeed they cannot) but can be learned as well: Our idiosyncratic expected patterns become attractors in themselves, irrespective of reward gains. Instead of just learning what *is* in the world, we also learn what to want (Bem, 1972; Srivastava & Schrater, 2015). As social creatures, our parents, caregivers and peers are of course the main sources of familiar patterns we are exposed to early in life. This type of learning usually happens in a behaviour-first fashion: Children first merely ape the sensorimotor "rituals" and only later infer the goals and preferences (as hidden causes) behind them. At that point, the sensory "evidence" of oneself behaving in a certain fashion is already so irrefutably strong, that the best explanation to adopt will be "I have a preference to do X." We see this principle of putting the practice (ritual) before the ideology ("sacred values") in acculturation of religions as well (Heylighen et al., 2018).

Another telling example of how people often infer their preferences based on (past) behaviour is the finding that our preferences will change when we have approached or chosen a particular stimulus before (Kawakami et al., 2007; Schonberg & Katz, 2020; Van Dessel et al., 2019). The mere act of choosing something (even if the choice is in fact cunningly manipulated by an experimenter), similar to the mere act of perceiving something, is shown to increase subsequent preference for it. Again, people seem to gravitate towards perceptual or behavioural patterns that are already part of their generative model. Here, a preference appears as just another hidden (high-level) cause explaining your sensorimotor interactions (minimizing prediction errors). In our predictive nets, a preference is shorthand for the complex composite of multimodal sensory cues and its associated reactive dispositions (Clark, 2019). All the things that make me feel X (and do Y) have a hidden cause we call "value." Similarly, we can learn from experience that when we have "desires," we tend to feel and do particular things, and that when we have similar desires/values, we end up doing similar things. Again, preferences, desires and values are not eliminated in the predictive processing view; they are just constructed. Preferences as high-level "empirical priors"—learned predictions of regular "packages" of sensorimotor flow—can be readily recognized in new instances and, through verbal labelling, allow for efficient communication and coordinated action (Clark, 2019).

This certainly does not imply that the usual, intuitive causal relation—going from preference to (choice) behaviour—has no part to play. It merely says that the inferences go both ways. Often when we need to determine value or preference, we rely on experience sampling; that is, we infer value based on past interactions and behavioural dispositions (Gershman & Daw, 2012). We continually infer our own appreciation, just like we infer other people's preferences from their choice behaviour (except that we have additional interoceptive data to infer our own preferences). This reminds us not to take preferences and values as "essences" but rather as fluctuating and context-dependent constructs. That will come as no surprise to researchers in aesthetic appreciation.

Once people can estimate (expect) the "value" (including the interoceptive observations it entails) that they are going to experience, they can also form models on what factors generate or modulate that experience. The abstract construct of value is then just another tool to help us in our self-evidencing (so continued existence), although animals without this level of abstracted causes can perfectly survive. Regardless, the underlying mechanics remains that of uncertainty minimization. If I have a set of high precision expected (interoceptive) states (a.k.a. the ones that "feel good"), and I have a model of how to generate those (e.g., through getting things that I value), active inference will get me to them. In the same way, I can have a model of how to reach pleasure from a film or a book, namely by actually watching or reading it, instead of

by just hearing the denouement. So we dislike spoilers even though that would make the whole experience more predictable (in our colloquial sense of the word). Of course, our internal model of pleasure generation might be wrong here. Research has shown that spoiling does not, in fact, substantially diminish the pleasure of watching (Leavitt & Christenfeld, 2011). Indeed, people commonly reread the same book or rewatch the same film.

To sum up, what we claim is not that people necessarily prefer the most predictable stimuli, person or situation, by our colloquial sense of "predictable," but rather that the way we become aware of our preferences (as well as our distastes) is as products of a thoroughly constructive, predictive process. Our preferences are phenomenally transparent in Metzinger's (2007) sense that we experience them as direct or given, with no access to preference forming (valuation) processes. "The phenomenologically expressed preferences of a conscious agent do not bear any evidence on the nature of the underlying brain processes that produce such conscious preferences" (Weaver, personal communication).

The value of obstacles

Perhaps the reader still harbours some doubts about the picture of value so far. Perhaps an agent driven by the proposed principles still too often seems to find value in echo chambers instead of museum halls. And maybe the view of preference as just another inference falls short in capturing the specific phenomenology of affective experience: the thing that determines the valence of value. What can our approach offer to ease those worries?

For this we have to return to the way predictive processing puts value on epistemic concerns. Under the account, an agent's search to understand its surroundings is a crucial part of biology, not just limited to human beings. Traditional scholars of motivation and psychological aesthetics have found themselves being forced to postulate a separate "need for cognition" (e.g., Kagan, 1972) (known under many different guises such as the need for "knowledge," for "structure" or for "closure") to adequately explain human behaviour. In predictive processing, on the other hand, the tendency to explore and learn the structure of one's environment directly follows the principle of (expected) prediction error minimization. Bear in mind that this is not a quest for some absolute truth (though it may have enabled one in humans): It operates on a pragmatic principle of increasing coherence between expected and input patterns (and in the service of the body) rather than towards increasing correspondence to some external ground truth.

The way that curious and explorative actions are driven by expected predictive progress (i.e., uncertainty that is expected to be reducible) also aligns with recent theoretical, computational and empirical work casting curiosity as expected learning progress, sometimes denoted as expected information or compression gains (Gottlieb et al., 2013; Holm, 2017; Kidd & Hayden, 2015; Schmidhuber, 2009; Félix Schoeller et al., 2018; Van de Cruys et al., 2021). What this means is that we are tuned to find regions in the input space that, given our current mental models, hold the highest potential for—the best slopes of—prediction error (uncertainty) reduction. As we saw in our overview of curiosity theories, we can say that we navigate the world guided by learned (meta-)expectations on "cognitive structuring" rates, or expected rates of prediction error minimization. The hypothesis is that people have a (imperfect) sense of where predictive progress can be made, which we experience as curiosity. Actually *making* this predictive progress is experienced as pleasure or appreciation. Consequently, positive affect is determined by the rate of prediction error minimization, a change in uncertainty over time, and not by the momentary absence of prediction errors (cf. familiarity previously).

This readily applies to our experience with art. Several seminal works in psychological aesthetics (e.g., Huron, 2006; Meyer, 1961) have stressed the importance of expectation in aesthetic experience (especially music). Artists exploit learned (culturally dominant) predictions or establish new predictions in their work (e.g., a rhythm or motif in music). These predictions can be based on lower-level regularities (e.g., perceptual symmetry) or on higher-level, semantic associations (e.g., story archetypes). Crucially, however,

they intuitively add prediction errors in their works to allow predictive progress in their audience: an active experience of recovering structure or meaning, sometimes across levels of abstraction (Van de Cruys & Wagemans, 2011). For example, Picasso's perceptually fragmented face in his famous work *Weeping Woman* may be broken because it expresses sadness. Or Munch's human figure in his painting *Separation* loses its perceptual boundary with the background because it may express loss. But these simplistic examples betray the complexity and uncertainties of the process. In practice, artworks will invite many cycles of curious anticipation of prediction error resolution alternated by (local) instances of actual resolution, often without reaching a final resolution, whatever that may mean (Muth & Carbon, 2016). A common way to put this is that artists use errors to delay understanding, to allow their public to experience things anew, as if it is *their first encounter* with the perceptual world (Shklovsky, 1917), when the biggest learning progress could still be made, but always with the risk of leaving their audience in the dark. Indeed, artists mostly do not deliberately search for the sweet spot of their audience (the result would be a gimmick instead of an artwork). An artist may not be particularly popular in their own time. They often push the limits in terms of unexpectedness and then have to rely on a cultural learning process to bring unpredictability to a level that is appreciated in terms of perceived structure and reducibility of prediction errors. In fact, one could say that our account is infused with the realization that, once one would elevate something (some predictable "rule") to "art," artists would break this rule (introduce uncertainty). Indeed, they must do so, to be able to touch their audiences again.

Precise prediction errors can create curiosity in the sense that they allow expectations on error reduction to grow, which spurs engagement through epistemic actions (e.g., eye movements on a painting). However, as those prediction errors and the mental effort spent on matching them increase, the expectation of error reduction decreases. In the limit, this may lead to anxiety, irritation or just looking away, which, in keeping with our approach, should be understood as using overt action to reduce prediction errors, namely by avoidance behaviour (instead of prediction updating). We shift attention to a more predictable setting, in which expectations on rates of prediction error minimization can again be fulfilled. Just before this point of abandonment, however, there is an opportunity for a *higher than expected* rate of error minimization. This characterizes the intensely positive experience of sudden insight ("eureka") or "aha." Indeed, empirical work has recently shown that aha moments arise when we are able to resolve a problem faster than expected (Dubey et al., 2021).

The role of obstacles is also clear in the so-called *generation effect* in memory (Slamecka & Graf, 1978): the finding that we remember things better if we can generate or infer them ourselves. Marketers exploit this phenomenon to make their brands stick. They would give their potential customers only just enough (or even slightly distorted) information to make the inference or generate the meaning themselves rather than providing complete information. Again, this is a kind of anti-simplicity or disfluency principle that causes people to remember and, we would say, appreciate a stimulus more (whether a painting or a brand; see Figure 26.4).

Notice that predictive processing provides an account of the process here, as well as of the directionality: One needs the overarching systemic principle of minimization of uncertainty (self-evidencing) to be able to attach valence (so emotion) to it. More concretely, changes towards more uncertainty (e.g., dissolving figures) will be less pleasant than the reverse operation, unless the figures (unexpectedly) morph into some new structure instead of mere disorder. The directionality is important but change is *conditio sine qua non*. Hence the captivating power of ambiguous images that invite change (Figure 26.4).

As before, this analysis naturally extends to action. The action analogue of art is games. In playing games we adopt rules, consisting of goals (expected outcomes) as well as "unnecessary *obstacles* [to our goals] to make possible the activity of overcoming them" (Nguyen, 2017). The adopted ends are often arbitrary and only adhered to for the duration of the game. But the "game" of increasing and decreasing uncertainties, in this case about whether your actions can accomplish goals, is the same. Nguyen therefore calls games a veritable "art of agency." The aesthetic experience is associated with the increasing fit or harmony between

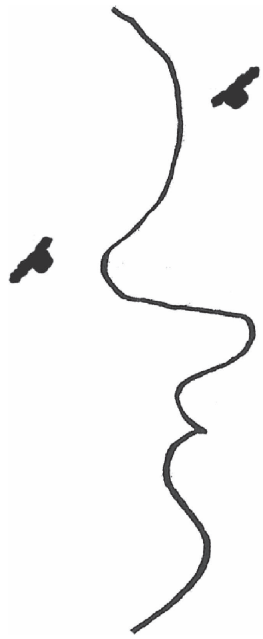

Figure 26.4 Ambiguous figure adapted from Hebb (1949). Why do figures like this appeal? It is arguably not because of the simplicity of the input per se (contra Gestaltists). Absences can become meaningful (as obstacles or prediction errors) in a predictive processing account. The line is of course simple as such, which might make you expect little meaning in it. But that changes once you look a fraction longer, probably because your visual system registers that the line deviates (prediction error) from what you expect an average/randomly drawn line looks like. This prediction error suggests an intention of the drawer, which in turn creates an expected reducibility of the prediction errors, which is fulfilled (after a brief search) with the discovery of a face. But this leaves some remaining errors: The one face has an odd contour, which leads to the discovery of a second face. In what seem to be unassuming beginnings, there are already micro-cycles of curiosity, discovery and appreciation. These would make up the "goodness" of a figure according to the current account.

one's abilities and the challenges of the situation (Nguyen, 2017). In art and games (as in life), the struggle is inseparable from the joy. Of course, the dynamics of overcoming obstacles also provides the plot for much of the (cultural) narratives that we enjoy so much, as many scholars and artists such as Kurt Vonnegut have noticed (Felix Schoeller & Perlovsky, 2016; Singh, 2021). It is by incurring informational (and metabolic) costs in the form of divergence between one's models and the sensed world (Zénon et al., 2019) that a fallible but pleasurable process of (re)harmonization is enabled.

Note that we carefully steer away from identifying aesthetic appreciation with just a resolved end product (regularity) or mere effort/energy spent (obstacles or prediction error). Artworks (even static ones) that are ambiguous or indeterminate are never free from regularities, however unstable they may be (and so regularities are usually only temporarily, mentally reachable). Our mere failures in the encounter with art are not the art, but artists are very creative in finding patterns in your failures, turning failure into art. Absolute noise does not capture attention or please, except if there's some story or rule (read: predictability) to it or if it merges dynamically into something (identifiable). There is something pretty trivial to this idea, in the sense that to be "identifiable" just means having a regularity. An aesthetic experience can definitely consist of just lingering on one's failing (prediction errors), but there has to be some curiosity still and so an expected

reducibility of (some of) the "errors." This in turn means that you have indeed been able to reduce them in your past experience (in similar situations, in other art contexts, in other artworks from a particular artist or the same artwork just moments ago). The artist cannot keep you hanging indefinitely or she loses your engagement, precluding any further experience, aesthetic or not (of course her mileage may vary in the next observer). In that sense, any experience is indeed aesthetic (else we would not have it), as Dewey (2005) realized.

Learning dynamics and valuation

The general idea that takes shape here is that affective experience reflects the dynamics of inference and learning (i.e., dynamics in prediction errors or uncertainty). This is a comprehensive hypothesis in the sense that it subsumes curiosity and (aesthetic) appreciation, as well as our more conventional strivings: the progress (or regress) in reducing discrepancies relative to our goals as expected outcomes for the self. Positive (vs. negative) rates of discrepancy reduction relative to one's goals have long been proposed to be "what it's like" to experience positive (vs. negative) affect (Carver & Scheier, 1990; Kivetz et al., 2006). Since goals are just one important type of (higher-level) expectations, we can readily generalize this view and situate valence in the increasing versus decreasing rates of prediction error minimization. Perception and learning, as well as discrepancy reduction by action, always take place against the backdrop of the model that includes our goals and homeostatic expected states, so it is never divorced from valuation: Indeed, it is the very *core* of valuation. With Skov and Nadal (2020), we emphasize that the basic mechanics of valuation are not specific to our engagement with art. However, our particular hypothesis on the nature of valuation and affective experience is relatively new, so its computational and psychological implications have only just begun to be explored (Hesp et al., 2021; Joffily & Coricelli, 2013; Kiverstein et al., 2019; Solms, 2018; Van de Cruys, 2017; S. Wilkinson et al., 2019).

To see why the hypothesis is plausible, we first need to acknowledge that the notion of reward is not self-evident or objectively determinable. So it cannot serve as the root cause of valence it is often presumed to be. Reward (value) is not objectively given by the environment but instead always intrinsically determined or inferred by the organism. Contrary to most reinforcement learning settings, there is no external "teacher" that can configure the organism's reward function (Juechems & Summerfield, 2019). Evolution is often thought to take on this role by, as we saw, equipping an organism with desired or "expected" observations. However, phenomena like satiation and the hedonic treadmill (the tendency to quickly return to a stable level of pleasure; Brickman & Campbell, 1971) indicate that what is valuable is always crucially and very contextually dependent on internal states and learning. It is impossible for evolution to pre-specify a fixed reward function for organisms that, like ourselves, are so dependent on learning for survival (we can become hunter-gatherers or wage-earners-supermarket-goers). As we discussed, predictive processing naturally accommodates learning of what is valuable (e.g., new goals) during our lifetime. What this theory suggests is that the progress on goal attainment (including homeostatic set-points), can be evaluated by the *change* over time in the distance between expected state and current state (change in prediction error). In other words, no (lack of) momentary discrepancy or prediction error should amount to affective value but rather the dynamics in prediction errors: Are we in the process of decreasing or increasing them? This formulation reveals that the epistemic or aesthetic emotions are not separate from the conventional "existential" or homeostatic ones but rather rely on similar computational mechanisms.

Perhaps counterintuitively, the system is optimizing for learning instead of for rewards: It searches for regions in the input space (the environment) of large gradients of prediction errors. Finding the pattern is the actual reward, not getting a particular stimulus or substance. This isn't to say that we engage with art *because* we want (have a conscious goal) to experience (the positive feeling of) uncertainty resolution. The system just uses inference as uncertainty reduction mechanisms, so the appreciation could as well

be epiphenomenal. However, the fact that we consciously experience these uncertainty dynamics not as "content" but as affect means that the system can represent (predict) the changes in uncertainty per se, as well as inferring the hidden causes not of inputs but of its *own processing dynamics*. It can build a generative model of those changes as expected states. In other words, it represents "value" (and a generative model of value) as a more abstract, universal currency, enabling comparisons among diverse situations or modalities. In the previous sections, we described value as a hidden cause summarizing a regular pattern of sensory consequences and behavioural reactions that is shared among instances of value. Here, we can add that a core part of the "value complex" will be an expected dynamic in internal processes: the uncertainty gradients.

In this view, people will seek out art because they have a model of how this type of environment provides them with good prediction error slopes. This is just a complicated way of saying they engage with art because of the sense-making (regularity-revealing) value it generates (which as we saw crucially involves sense-breaking). But for any specific artwork, additional experiences will reduce future uncertainty resolving potential. This does not mean an artwork or piece of music will necessarily lose its emotional potential once we are acquainted with it. For example, repeated experience with Picasso's cubist faces will not overwrite our overlearned beliefs on the proper structure of faces (lest it would break our everyday perception), so prediction errors (and hence the potential for resolution) will continue to be present despite repeated exposure. However, frequent engagement with a piece of art or music can detract from its emotionality under this learning-based view, suggesting that we are often chasing an experience we cannot fully relive when revisiting the same artwork.

Affective experience arises when we are able to build meta-models that monitor (model) how our modelling efforts on the lower levels are doing. Hence, (first-order) learning is phylogenetically prior to affect (Ginsberg & Jablonka, 2019). The need for meta-models may only have emerged from a powerful capacity to learn and model our world in unprecedented temporal (hierarchical) depth. This capacity required equally powerful feedback mechanisms on how well the system is doing in modelling, that is, coping with, its environment. Indeed, in this conception, emotions track the fluctuations in uncertainty or confidence in our models on longer timescales (spanning multiple glances or actions; Hesp et al., 2021). As an explanation of emotions, this may have an overly intellectualistic or cognitivist ring to it, but in fact it merely reinforces ideas that were always central in theories of emotion. Different prominent scholars describe emotions as a form of nonconceptual monitoring of coping performance in our interaction with the world (Frijda, 2006; Reisenzein, 2009). We just propose that prediction error dynamics are the form of feedback on the system's own functioning required to spell out those accounts in computational specificity (Hesp et al., 2021; Joffily & Coricelli, 2013; Schillaci et al., 2020).

If we turn to the *direct* evidence for this view, we must conclude it is scarce so far. Of course, the view is still very young, and the challenges to operationalize it are considerable. At its centre are unobservable, hidden states and their dynamics which intrinsically depend on individual learning history. Advances in disentangling and tracking neural activity associated with predictions versus prediction errors across time and along the cortical hierarchy are modest so far (e.g., Corlett et al., 2021; Issa et al., 2018). Meanwhile, in silico computational modelling implementations of these ideas provide an interesting proof of concept (Hesp et al., 2021) but cannot make the hypothesized connections to the experiential dimensions of affective value. Still, the idea that positive affect is linked *not* to conventional reward (magnitudes) per se but to learning dynamics—specifically one's success in learning the structure of the environment—received support in a recent study that managed to cleanly disentangle those factors (Blain & Rutledge, 2020). More fundamentally, the reward value of "mere" information gain (uncertainty resolution) is also well supported (Bromberg-Martin & Monosov, 2020; Fiorillo et al., 2003; Fortes et al., 2016) and shown to be associated with activity in the same dopaminergic brain regions as conventional rewards (Bromberg-Martin & Hikosaka, 2011). Indirect evidence for the link between positive affective experiences and learning comes from

studies reporting that a pleasurable aha experience is associated with better memory (Danek & Wiley, 2020; Kizilirmak et al., 2016; see also Sarasso et al., 2021; Van de Cruys et al., 2021).

In the domain of psychological aesthetics, several recent findings nicely fit with a view that attaches pleasure to learning (uncertainty) dynamics. For example, Cheung et al. (2019) showed that pleasure in music is associated with the surprise (prediction error) of hearing a chord in a chord progression, when one actually had a strong (low uncertainty) prediction about the chord that would follow (but didn't). The authors managed to quantify uncertainty and surprise using a large corpus of chord progressions in popular music (see also Gold et al., 2019; Pearce, 2018). We assume that the unexpected increase in prediction error (of surprising tones, beats or chords in music) is what allows the pleasurable resolution of uncertainty when the pattern is reinstated (in music). Similarly, the studies reporting a preference for intermediate predictability in auditory or visual sequences (e.g., Delplanque et al., 2019; Witek et al., 2015) can be marshalled in support of the proposed view, although the precise mechanisms usually remain beyond reach in those studies. In a study using video art, Muth et al. (2015) found that dynamics in (self-reported) semantic uncertainty while viewing animated sketches influence appreciation. Finally, the set effects discussed earlier (relative fluency effect; Wänke & Hansen, 2015) also sit comfortably with our reasoning, because they suggest that affective experience (appreciation) reflects a type of feedback or monitoring of one's inferential processes across multiple stimuli or trials. Such studies underscore that even though we are asked to appreciate a single stimulus and we ourselves also intuitively (mis)attribute our appreciation to a singular work or stimulus, in reality our affective experience is a function of our own way of processing, in light of the more extensive context. Similar to how it is hard to shake your confidence that the content of a perceptual experience is not just out there in the world but the result of perceptual inference, the immediateness of the affective experience hides the underlying (meta)inferences.

The touch of art

Attuning generative models

What is true for all perception holds for perceiving art: We build up a generative model of the artwork. Because we know the artwork is human-made, the hidden causes that we infer will not only include the objects portrayed or the medium used (e.g., brush strokes as latent causes of particular patterns of paint) but also the emotions and intentions that the painter might have had (Freedberg & Gallese, 2007; Seth, 2019). For example, when we infer aggression as the hidden cause for an odd (unexpected) trace of a brushstroke, we turn that sensory effect (brushstroke) into an expression. Of course, our inferences do not necessarily reflect the meaning or emotion of the artist, but we can't help but approach the work in this (re)constructive sense (as we do for all sensory inputs), trying to see potential sense in the artist's apparent sense-breaking. Articulating these processes makes it seem a deliberate matter. In truth, we automatically infer and readily *perceive* the causal history of visual stimuli, making irregularities regular, as experiments have shown (Chen & Scholl, 2016; H. Leder et al., 2012; Pinna, 2010). Sometimes our inferences rely on intuitive (possibly innate) mappings from "generative" emotions to auditory or visual features, like when slow rhythmic and melodic contours express tranquillity in lullabies. Other times a learned "language" of expressions is needed to "get" the artwork (Jackendoff & Lerdahl, 2006). The basic process is not unlike empathy (Wilkinson et al., 2021), but with a work of (personified) art instead of with an actual person. We enter into a sort of dialogue in which we probe the work with predictions that we "test" using targeted sampling—for example, with eye movements to regions with high expected informativeness (precision). And the work will "respond" with inputs that become clues under our probing (Koenderink, 2010).

When reading Aristotle on the origins of aesthetic experience (in his *Poetics*), one is struck by how much of the current theorizing on inference and learning he already foreshadowed. Tracy (1946) summarizes Aristotle's line of argument as follows:

> in the case of a work of art, the observer is establishing a significant connection between the presentation he sees (picture, play, etc.) and some original of which he has knowledge from his own experience; the inferences drawn by an observer of a work of art have to do with, and are conditioned by, the necessarily imperfect degree of adequacy it achieves; i.e., some effort on the part of the observer is required to get him en rapport with the artist; satisfaction comes from the successful integration between the artist's way of presenting a given situation or object and the observer's power to interpret the artist's procedures.

The key point here is the pleasure created by the increasing alignment (by surpassing obstacles; cf. effort) between generative models of the work (or implicit artist) and the observer. Aristotle links this to the emotional impact of tragedies or more generally the pleasure we can derive from negative emotions in art (see also Menninghaus et al., 2017). For example, reading a good novel can provide an unexpected attunement with core but largely unarticulated dynamics of the self (models), even (or especially) if those dynamics concern negative emotions. A piece of art can provide validation or external evidence for the type of regularities in our own mental life which we consider very personal and idiosyncratic (Van de Cruys et al., 2017) and in this way supports self-evidencing at a very fundamental (and uniquely human) level. It is as if we can briefly breach the epistemic and conative boundary between ourselves and the outside world (other agents). This brings home the point that a predictive processing view of the aesthetic experience is about more than mere "epistemic" problem-solving or striving for predictability. It captures the existential consolation of art as well. Moreover, art puts in motion a self-reinforcing cycle, in the sense that validation of your models will provide you with the confidence and safety—read: give you an expectation of a good rate of prediction error minimization in a given context—to be curious again and to explore prediction errors that afford new meaning. Those cycles may demarcate aesthetic experience from other experiences, more so than any of its component processes separately.

Dislodge and roam free

With the great power of our modelling capacities—particularly the increased hierarchical *depth* of networks of inferred causes—comes great vulnerability as well. Specifically, the (Bayesian) beliefs that make up our models are at risk of becoming overly convoluted or even free-floating, impervious to evidence. Indeed, in daily life, we often have only very sparse, biased or indirect data to speak to and constrain our predictive models (Hoel, 2020). On top of that, because of the hierarchical model structure, rich auxiliary hypotheses can be called upon to take the blame and explain away current prediction errors, thereby saving our "preferred" (high confidence) models from revision. In this way, we end up with overfitted models: models that try to capture too much of a given set of sense data and as a result will not generalize to new situations (hence creating prediction errors in the longer term). These models are too specific and inefficient in that they have to use too many parameters (high complexity) in order to match every new situation encountered instead of latching on to stable regularities across samples. Overfitted models capture incidental variability (noise) in the "training" samples as putative regularities. Interestingly, a deliberate injection with noise inputs is a proven technique in artificial intelligence to *improve* the generalizability of learned predictive models (see also cybernetic phenomenon of "order from noise"; Von Foerster, 1960). Intuitively, it shakes the system out of its habits that were too tuned to the particular features of its limited, biased sample of experiences.

Here, art may provide a safe, "offline" testing ground for our models: It can be a way of safely introducing "chaos" to make our models more robust in the long term (Hoel, 2020). It may seem disrespectful to describe art as mere noise, but it aligns with the freedom with which artists can upset our predictions, intentionally creating ambiguities and fictions to shake up our categories (Hoel, 2019), often allowing our models to settle down again, finding a more global and robust minimum in the prediction error landscape. Artists intuitively direct their "noise" not only to hone existing models—reducing their complexity while retaining their predictive power—but also to open new pathways of future prediction error minimization.

At this point, we encounter the artist's capability that seems diametrically opposed to the predictive processing account with its sole focus on uncertainty minimization. Keats (Poetry Foundation, 2021b) called it *negative capability*, that is, "when a man is capable of being in uncertainties, mysteries, doubts, without any irritable reaching after fact and reason." According to Keats and many after him, great artists have this aptitude that allows them to "bury self-consciousness, dwell in a state of openness to all experience, and identify with the object contemplated" (Poetry Foundation, 2021a). It is the capacity to dwell in uncertainty—to "stay with the trouble" (Haraway, 2016)—despite the existential threat it holds and despite the fact that our whole being, per predictive processing, is fundamentally oriented towards uncertainty minimization. It is a true capability because it is the openness for uncertainty or unpredictability that gives the opportunity for radically new and different models to take hold. Instead of merely triggering a shift in probabilities of the currently available hypotheses (prediction updates), it opens new hypothesis spaces. It is a capability that cannot be approached as a capability, hence the qualifier "negative."

The paradox is that expecting uncertainty and inviting chaos (what we could call *radical curiosity*) can lead one to perceive new layers of regularities in reality. Such discoveries may rely on a recycling of existing regularities (models) for new domains (cf. metaphors), but it crucially requires relinquishing our default, prepotent models for that other domain, in other words, allowing uncertainty to rise first. Sometimes this is just a matter of enlarging (the salience) the unexplained residue of our default models, as artists sometimes do. However this "breaking of new ground" is done—what type of generative models allow for it (Williams, 2020)—it seems it is only possible in a very roundabout way. Like greedy reward optimization (exploitation instead of exploration) makes you miss out on opportunities for better rewards, a focus on greedy information gain or predictive progress would foreclose the discovery of new, more efficient regularities and so better means for uncertainty minimization. We need to relinquish our focus on information gain relative to current models to radically reconfigure the hypothesis space of our models and try on different (self-)models. No artist (improviser) has to be convinced of this kind of generative power of errors and uncertainty, but it is hard to form a model and plan for this generativity. However, it is something one recognizes when it happens. It is marked by intensely positive emotion: *Unexpected* gains, or faster than expected rates of prediction error reduction, feel better than mere gains. It is impossible to use information gain as a guide towards unexpected gains, except by deliberately staying with or increasing one's uncertainty. There is interesting work ahead in exploring how this translates into the artist's concrete (epistemic) practices.

A parallel can be drawn with evolution in single-celled organisms (Freddolino & Tavazoie, 2012). In harsh conditions, these organisms upregulate mutation rate (randomness or uncertainty) to increase the chances for the population to discover new useful regularities (hypotheses) and hence survive. In more complex organisms, individuals are less expendable, so they have developed ways to ensure the organism does not need to die with its hypotheses. However, as in single-celled organisms, the risk of introducing uncertainty seems to be offset by the social group, albeit in a different way. A child takes more risks in the presence of the parent that it trusts to resolve its uncertainty. Its freedom in making errors begets novel structure. Similarly, a context of art, because it allows atypical amounts of certainty or randomness to be safely introduced, may be an accelerator of novel predictive models. However, trust seems to be a precondition for this. Trust in oneself as the metacognitive belief that one's generative model is up to dealing with whatever is thrown at it (e.g., a jazz musician who knows that, through her skills, she can deal with unexpected events or "errors"). And trust in

the social context of art, understood here as our basic expectation that prediction errors are reducible when one takes the effort to engage with it. I know there is a model; I just don't know the model (yet). This trust is rooted in the responsiveness of our sociocultural environment: The way it, from early on in development, folds to our anticipatory activities (Hunt, 1981), thereby kickstarting our predictive senses and providing a unique type of reliable validation of one's models. If that trust or confidence has percolated into hierarchically deeper layers of our models, it seems to create room for the "letting be" of uncertainties on lower-level branches of models, instigating novel, flexible model reconfigurations at those levels.

Allowing uncertainty to surge in an organism that owes its very existence to an uncertainty minimizing process seems a perilous and even self-contradictory (literally and figuratively) undertaking. One way to respond to this would be to say that uncertainty minimization is not a sufficient condition for existence; it can only be used in retrospect to say that, if an organism is alive now, it must have been minimizing uncertainty (Constant, 2021). This reasoning is reminiscent of evolutionary theory, where the emphasis on adaptation as necessary for survival, risks obscuring the equally crucial role of non-adaptive variability (random mutations). Weick (1979) remarks that "adaption precludes adaptability": A built-in specialization (adaptation) for a particular "goal," while possibly more efficient in the short term, will get you trapped in a dead end when the environment changes. Variability or uncertainty is the raw material for adaptive change and novel regularities to be picked up. The example of single-celled organisms changing their mutation rates illustrates the power of channelling uncertainty in volatile or harsh conditions. It is the active planning for uncertainty in order to better minimize uncertainty (as the population), similar to what we saw with precision expectations in predictive processing (at the individual level). We can talk about "planning for uncertainty," however contradictory that sounds, because expected uncertainties become part of the model and the system benefits from making uncertainty selective (e.g., limiting the increase of mutations to those parts of genetic material most likely to solve unexpected environmental challenges but not to others). One could compare this to an immune system that only works well when it has been challenged properly during its development (see also the idea of *hormesis* and *antifragility*; Taleb, 2013). An even more plastic system like the brain similarly needs those challenges, and so good uncertainty minimization slopes are what it should expect.

Humans, and artists all the more, seem specialized in selectively channelling and compartmentalizing uncertainty in order to improve and better support their models. What enables this and allows us to find novelty within an individual instead of a population is our "dissociative" capacity to accumulate evidence for a new model without immediately having to relinquish existing models (Kelly, 1964). Art is a product of our capacity to turn our models (including self-models) into objects themselves that can be taken "offline" and explicitly questioned with fictive "data" (Miller et al., 2020) instead of just applied and updated in "online" behaviour. This is what art does: asking and following through "what if" stories. Such counterfactual thinking softens the threat of model invalidations. We can dry-run different models without (yet) committing. It is the sand-boxing of a (self-)model to circumvent the logic of the law of the excluded middle, which says that we can't simultaneously affirm and deny a (hypo)thesis (Kelly, 1958). Indeed, given the constructivist and pragmatist inspiration of our predictive modelling, nothing prevents us from violating that law of rationality if it is to our benefit in coping with our world. The fascinating questions that have yet to be confronted by predictive processing are precisely about these "dissociative" capacities and the competition between fractionated (self)models that emerges. Some of these fractionated models are inferentially shielded so do not have any way to cross-talk. The result is food for psychotherapists or for . . . artists.

Conclusions

Throughout this chapter, we developed the view that when we "artify"—to focus on the activity instead of the products—we are, implicitly, optimizing for learning. We are curiously guided by expected predictive progress, we pleasurably make such progress, attune our models to the (social) world, and in the

process of introducing and sustaining (barely) controlled levels of uncertainty, we prevent overfitting and facilitate the unmooring of old models and the discovery of new ones. Our view retells accounts of aesthetic appreciation as old as Aristotle's and extends recent predictive processing-based accounts of aesthetic experience (Pelowski et al., 2017; Sarasso et al., 2020; Felix Schoeller et al., 2015; Van de Cruys & Wagemans, 2011).

Our artistic sense emerges as a neotenous trait with roots in playful behaviour in children and even non-human animals (Bekoff, 2015). A relaxed, playful dog picks out the most unwieldy stick, the one that fights back most. A child chooses the most unfamiliar, inappropriate object to play with (Andersen & Roepstorff, 2021; Bonawitz et al., 2012). At least in origin, this is not driven by self-conscious showmanship, let alone a display of prowess or the "costly signaling" of one's fitness, an evolutionary rationale that has been invoked to explain art as well (Dissanayake, 2007; Helmut Leder & Nadal, 2014). Rather, the child is intrinsically motivated to handle the more challenging objects, because the obstacles allow for predictive progress or advances in the "grip" it has on its environment. So it goes in art as well, where artists throw up barriers and constraints for their own performances to keep it genuinely *interesting* for themselves and their observers, in other words: to make room for new learning.

The emphasis on learning and inference in this chapter may give the impression that art is a mere epistemic matter, serving the "world-disclosing" goal of uncertainty and complexity reduction (preventing overfitting, i.e., future prediction errors). This may align with the Kantian "disinterestedness" as precondition for aesthetic judgment. However, as we saw, perceiving and learning is always dependent on "concepts" (predictions) and is already valued in predictive processing; value is not something added to it. There is no "veridical perception plus the values"; behavioural control and perception are one (against conventional modularity ideas). Value functions are absorbed into our probabilistic models (priors), so utilities and probabilities are not separately represented. This may explain why a positive aesthetic appreciation often goes together with perceived truth (Reber et al., 2004) (incidentally, it may also explain the difficulty in separating fact from value in everyday and ideological discussions). Of course, we humans can build explicit conceptual models of value as something construed to be strictly divorced from probabilities or expectancies. But the underlying processes do not honour that strict dissociation. On the one hand, the view reinforces the classical adagium from Western philosophy that value (beauty) and truth coincide, if only in experience. On the other hand, it overthrows the strict separation of cognition (probabilities, inferences) and emotion (values, preferences), equally deeply rooted in Western thinking. Still, the approach does not explain away value but reveals the mechanisms of valuation within an uncertainty minimizing agent.

It follows that we implicitly reconfigure value functions by exploring different models, taking on different predictive "sets" or roles and moulding our models to better fit internal (self) and external sensorium. In art, just as in play behaviour, our inferred models become targets in themselves. It is the context par excellence for us to suspend the normal utility functions absorbed in our routine models and explore not just how to attain our goals and values but rather *what can be valuable* in the first place. We do so by adopting arbitrary goals (or expectations), "assigning arbitrary rewards and accepting unnecessary costs" (Chu & Schulz, 2020), like we do in playing. We do so by relaxing the tight grip of prediction errors in the here and now, as well as relaxing optimization of information gain relative to the current predictive set, to ultimately—and fallibly—finding better long-term prediction error minimizing models. Far from being disinterested, it is about exploring a surplus of interest.

Finally, the approach also allows us to find a middle ground between two classical, orthogonal views of aesthetic appreciation. According to the "idiosyncratic" view, aesthetic appreciation is just a function of personal taste and the vagaries of arbitrary individual emotions, while to the nomothetic view, our aesthetic sense reflects some universal, objective (and human-independent) values. Philosophically, the latter view is of course rooted in classical Platonic theories of aesthetic value (Hart, 1971; Zangwill, 2021). But the dispute still plays out in current discussions in experimental aesthetics, with some studies reporting that aesthetic

appreciation is largely dependent on individual rather than shared taste (A. A. Brielmann & Pelli, 2019), while others try to establish objective stimulus features that determine appreciation (Graham & Redies, 2010). The view espoused here follows the idiosyncratic view in that aesthetic value emerges in the relation between a subject (with her particular models) and an object, much in line with the work of Dewey. Values are dynamically inferred rather than objective and static, and dynamics in inferences bring about affective values. However, as we have tried to show in this chapter, the same move gives us a window on the nomothetic principles governing (aesthetic) valuation as well.

Acknowledgements

I (SVdC) thank my twin daughters Martha and Jacoba for reminding me of the value of radical curiosity. We also thank Eline van Geert, Pieter van Dessel, David Gijbels, Pietro Sarasso, Félix Schoeller, Pablo Fernandez Velasco and Jacopo Frascaroli for helpful feedback on earlier versions of this chapter. This work is partly supported by a Methusalem grant from the Flemish government (METH/14/02). JB is supported by the NeuroEpigenEthics project that received funding from European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 804881). AM receives funding from the Research Foundation–Flanders (FWO, G073317N) and the Research Fund of KU Leuven (C14/17/047).

References

Aitchison, L., & Lengyel, M. (2017). With or without you: Predictive coding and Bayesian inference in the brain. *Current Opinion in Neurobiology*, *46*, 219–227. https://doi.org/10.1016/j.conb.2017.08.010

Alter, A. L., & Oppenheimer, D. M. (2009). Uniting the tribes of fluency to form a metacognitive nation. *Personality and Social Psychology Review*, *13*(3), 219–235. https://doi.org/10.1177/1088868309341564

Andersen, M. M., & Roepstorff, A. (2021). *Play in predictive minds: A cognitive theory of play*. https://doi.org/10.31234/osf.io/u86qy.

Arnheim, R. (1974). *Entropy and art: An essay on disorder and order*. University of California Press. http://books.google.com/books?id=kL75af7hbXsC

Bekoff, M. (2015). Playful fun in dogs. *Current Biology: CB*, *25*(1), R4–R7. https://doi.org/10.1016/j.cub.2014.09.007

Bem, D. J. (1972). Self-perception theory. *Advances in Experimental Social Psychology*, *6*(1), 1–62. http://healthyinfluence.com/wordpress/wp-content/uploads/2011/05/SP-Theory-Bem-Advances.pdf. https://doi.org/10.1016/S0065-2601(08)60024-6

Berlyne, D. E. (1960). *Conflict, arousal, and curiosity*. http://psycnet.apa.org/psycinfo/2006-09643-000/

Berlyne, D. E. (1978). Curiosity and learning. *Motivation and Emotion*, *2*(2), 97–175. https://doi.org/10.1007/BF00993037

Blain, B., & Rutledge, R. B. (2020). Momentary subjective well-being depends on learning and not reward. *eLife*, *9*. https://doi.org/10.7554/eLife.57977

Boddez, Y., Descheemaeker, M., Mertens, G., Truyts, A., & Van de Cruys, S. (2019). Like what you see: Generalization of social learning determines art appreciation. *Acta Psychologica*, *196*, 18–25. https://doi.org/10.1016/j.actpsy.2019.04.001

Bonawitz, E. B., van Schijndel, T. J. P., Friel, D., & Schulz, L. (2012). Children balance theories and evidence in exploration, explanation, and learning. *Cognitive Psychology*, *64*(4), 215–234. https://doi.org/10.1016/j.cogpsych.2011.12.002

Bornstein, R. F. (1989). Exposure and affect: Overview and meta-analysis of research, 1968–1987. *Psychological Bulletin*, *106*(2), 265–289. https://doi.org/10.1037/0033-2909.106.2.265

Bornstein, R. F., Kale, A. R., & Cornell, K. R. (1990). Boredom as a limiting condition on the mere exposure effect. *Journal of Personality and Social Psychology*, *58*(5), 791–800. https://doi.org/10.1037/0022-3514.58.5.791

Brickman, P., & Campbell, D. T. (1971). Hedonic relativism and planning the good science. In M. H. Appley (Ed.), *Adaptation level theory: A symposium* (pp. 287–302). Academic Press.

Brielmann, A. A., & Pelli, D. G. (2019). Intense beauty requires intense pleasure. *Frontiers in Psychology*, *10*, 2420. https://doi.org/10.3389/fpsyg.2019.02420

Brielmann, A. A., & Dayan, P. (2022). *Introducing a computational model of aesthetic value*. https://doi.org/10.31234/osf.io/eaqkc.

Bromberg-Martin, E. S., & Hikosaka, O. (2011). Lateral habenula neurons signal errors in the prediction of reward information. *Nature Neuroscience*, *14*(9), 1209–1216. https://doi.org/10.1038/nn.2902

Bromberg-Martin, E. S., & Monosov, I. E. (2020). Neural circuitry of information seeking. *Current Opinion in Behavioral Sciences*, *35*, 62–70. https://doi.org/10.1016/j.cobeha.2020.07.006

Carver, C. S., & Scheier, M. F. (1990). Origins and functions of positive and negative affect: A control-process view. *Psychological Review*, *97*(1), 19–35. https://doi.org/10.1037/0033-295X.97.1.19

Chamberlain, R., Berio, D., Mayer, V., Chana, K., Leymarie, F. F., & Orgs, G. (2021). A dot that went for a walk: People prefer lines drawn with human-like kinematics. *British Journal of Psychology*, *113*(1), 105–130. https://doi.org/10.1111/bjop.12527

Chen, Y. C., & Scholl, B. J. (2016). The perception of history: Seeing causal history in static shapes induces illusory motion perception. *Psychological Science*, *27*(6), 923–930. https://doi.org/10.1177/0956797616628525

Cheung, V. K. M., Harrison, P. M. C., Meyer, L., Pearce, M. T., Haynes, J. D., & Koelsch, S. (2019). Uncertainty and surprise jointly predict musical pleasure and amygdala, hippocampus, and auditory cortex activity. *Current Biology: CB*, *29*(23), 4084–4092.e4. https://doi.org/10.1016/j.cub.2019.09.067

Chmiel, A., & Schubert, E. (2017). Back to the inverted-U for music preference: A review of the literature. *Psychology of Music*, *45*(6), 886–909. https://doi.org/10.1177/0305735617697507

Chu, J., & Schulz, L. E. (2020). Play, curiosity, and cognition. *Annual Review of Developmental Psychology*, *2*(1), 317–343. https://doi.org/10.1146/annurev-devpsych-070120-014806

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, *36*(3), 181–204. https://doi.org/10.1017/S0140525X12000477

Clark, A. (2019). Consciousness as generative entanglement. *Journal of Philosophy*, *116*(12), 645–662. https://doi.org/10.5840/jphil20191161241

Clark, A. (2020). Beyond desire? Agency, choice, and the predictive mind. *Australasian Journal of Philosophy*, *98*(1), 1–15. https://doi.org/10.1080/00048402.2019.1602661

Constant, A. (2021). The free energy principle: It's not about what it takes, it's about what took you there. In *Biology and Philosophy*, *36*(2). https://doi.org/10.1007/s10539-021-09787-1

Corlett, P. R., Mollick, J. A., & Kober, H. (2021). *Substrates of human prediction error for incentives, perception, cognition, and action*. https://doi.org/10.31234/osf.io/pf89k.

Cupchik, G. C., & Heinrichs, R. W. (1981). Toward an integrated theory of aesthetic perception in the visual arts. In H. I. Day (Ed.), *Advances in intrinsic motivation and aesthetics* (pp. 463–485). Springer. https://doi.org/10.1007/978-1-4613-3195-7_19

Danek, A. H., & Wiley, J. (2020). What causes the insight memory advantage? *Cognition*, *205*, 104411. https://doi.org/10.1016/j.cognition.2020.104411

Delplanque, J., De Loof, E., Janssens, C., & Verguts, T. (2019). The sound of beauty: How complexity determines aesthetic preference. *Acta Psychologica*, *192*, 146–152. https://doi.org/10.1016/j.actpsy.2018.11.011

Dennett, D. C. (1991). Real patterns. *Journal of Philosophy*, *88*(1), 27–51. https://doi.org/10.2307/2027085

Dennett, D. C. (2017). *From bacteria to Bach and back: The evolution of minds*. W. W. Norton, and Company. https://play.google.com/store/books/details?id=PEp8DAAAQBAJ

Dewey, J. (2005). *Art as experience*. Penguin. http://books.google.com/books?id=aAbqAGo5MwwC

Dissanayake, E. (2007). What art is and what art does: An overview of contemporary evolutionary hypotheses. In *Evolutionary and neurocognitive approaches to aesthetics, creativity, and the arts* (pp. 1–14). https://books.google.be/books?hl=nl&lr=&id=np6RDwAAQBAJ&oi=fnd&pg=PT8&dq=artistic+aesthetic+evolutionary+adaptation+&ots=9epYxgUzVm&sig=k7gulj_FA7I96RDmg8oaH8CRafU

Dubey, R., Ho, M. K., Mehta, H., & Griffiths, T. (2021). Aha! moments correspond to meta-cognitive prediction errors. PsyArXiv. Retrieved June 22, from https://psyarxiv.com/c5v42/download/?format=pdf

Eysenck, H. J. (1942). The experimental study of the "good Gestalt"—A new approach. *Psychological Review*, *49*(4), 344–364. https://doi.org/10.1037/h0057013

Fiorillo, C. D., Tobler, P. N., & Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. *Science*, *299*(5614), 1898–1902. https://doi.org/10.1126/science.1077349

Forster, M. (2019). Processing fluency. In M. N. A. Vartanian (Ed.), *The Oxford handbook of empirical aesthetics*. Oxford University Press. https://doi.org/10.1093/oxfordhb/9780198824350.013.21

Forster, M., Gerger, G., & Leder, H. (2015). Everything's relative? Relative differences in processing fluency and the effects on liking. *Everything's Relative. PLoS One*, *10*(8), e0135944. https://doi.org/10.1371/journal.pone.0135944

Fortes, I., Vasconcelos, M., & Machado, A. (2016). Testing the boundaries of "paradoxical" predictions: Pigeons do disregard bad news. *Journal of Experimental Psychology. Animal Learning and Cognition*, *42*(4), 336–346. https://doi.org/10.1037/xan0000114

Freddolino, P. L., & Tavazoie, S. (2012). Beyond homeostasis: A predictive-dynamic framework for understanding cellular behaviour. *Annual Review of Cell and Developmental Biology*, *28*, 363–384. https://doi.org/10.1146/annurev-cellbio-092910-154129

Freedberg, D., & Gallese, V. (2007). Motion, emotion and empathy in esthetic experience. *Trends in Cognitive Sciences*, *11*(5), 197–203. https://doi.org/10.1016/j.tics.2007.02.003

Frijda, N. H. (2006). *The laws of emotion*. Erlbaum.

Friston, K. J. (2010). The free-energy principle: A unified brain theory? *Nature Reviews. Neuroscience*, *11*(2), 127–138. https://doi.org/10.1038/nrn2787

Friston, K. J., Da Costa, L., Hafner, D., Hesp, C., & Parr, T. (2021). Sophisticated inference. *Neural Computation*, *33*(3), 713–763. https://doi.org/10.1162/neco_a_01351

Friston, K. J., Daunizeau, J., Kilner, J., & Kiebel, S. J. (2010). Action and behaviour: A free-energy formulation. *Biological Cybernetics*, *102*(3), 227–260. https://doi.org/10.1007/s00422-010-0364-z

Friston, K. J., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., & Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive Neuroscience*, *6*(4), 187–214. https://doi.org/10.1080/17588928.2015.1020053

Friston, K. J., Rosch, R., Parr, T., Price, C., & Bowman, H. (2018). Deep temporal models and active inference. *Neuroscience and Biobehavioral Reviews*, *90*, 486–501. https://doi.org/10.1016/j.neubiorev.2018.04.004

Friston, K. J., Thornton, C., & Clark, A. (2012). Free-energy minimization and the dark-room problem. *Frontiers in Psychology*, *3*, 130. https://doi.org/10.3389/fpsyg.2012.00130

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

Gershman, S. J. (2019). What does the free energy principle tell us about the brain? In arXiv [q-bio.NC]. arXiv. http://arxiv.org/abs/1901.07945

Gershman, S. J., & Daw, N. D. (2012). Perception, action and utility: The tangled skein. Principles of brain dynamics: Global state. *Interactions*, 293–312. http://books.google.be/books?hl=nl&lr=&id=KOWZ2sZNlWQC&oi=fnd&pg=PA293&dq=Perception,+Action,+and+Utility:+The+Tangled+Skein+gershman&ots=e-7WaGMx4V&sig=CO2ffgFSyMpRkWXR5l08YHTnGrc

Ginsberg, S., & Jablonka, E. (2019). *The evolution of the sensitive soul*. MIT Press.

Glanzer, M. (1958). Curiosity, exploratory drive, and stimulus satiation. *Psychological Bulletin*, *55*(5), 302–315. https://doi.org/10.1037/h0044731, PubMed: 13591450

Gold, B. P., Pearce, M. T., Mas-Herrero, E., Dagher, A., & Zatorre, R. J. (2019). Predictability and uncertainty in the pleasure of music: A reward for learning? *Journal of Neuroscience: the Official Journal of the Society for Neuroscience*, *39*(47), 9397–9409. https://doi.org/10.1523/JNEUROSCI.0428-19.2019

Gombrich, E. H. (1963). *Meditations on a hobby horse: And other essays on the theory of art*. Phaidon Press.

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

Graf, L. K. M., & Landwehr, J. R. (2015). A dual-process perspective on fluency-based aesthetics: The pleasure-interest model of aesthetic liking. *Personality and Social Psychology Review*, *19*(4), 395–410. https://doi.org/10.1177/1088868315574978

Graham, D. J., & Redies, C. (2010). Statistical regularities in art: Relations with visual coding and perception. *Vision Research*, *50*(16), 1503–1509. https://doi.org/10.1016/j.visres.2010.05.002

Güçlütürk, Y., Jacobs, R. H. A. H., & van Lier, R. (2016). Liking versus complexity: Decomposing the inverted U-curve. *Frontiers in Human Neuroscience*, *10*, 112. https://doi.org/10.3389/fnhum.2016.00112

Hafner, D., Ortega, P. A., Ba, J., Parr, T., Friston, K., & Heess, N. (2020). Action and perception as divergence minimization. In arXiv [cs.AI]. arXiv. http://arxiv.org/abs/2009.01791

Haraway, D. J. (2016). *Staying with the trouble: Making kin in the Chthulucene*. Duke University Press. https://play.google.com/store/books/details?id=ZvDgDAAAQBAJ

Hart, S. L. (1971). Axiology—Theory of values. *Philosophy and Phenomenological Research*, *32*(1), 29–41. https://doi.org/10.2307/2105883

Hebb, D. O. (1949). *The organization of behaviour*. John Wiley.

Heilbron, M., Armeni, K., Schoffelen, J.-M., Hagoort, P., & de Lange, F. P. (2021). *A hierarchy of linguistic predictions during natural language comprehension*. In bioRxiv (p. 2020.12.03.410399). https://doi.org/10.1101/2020.12.03.410399

Hesp, C., Smith, R., Parr, T., Allen, M., Friston, K. J., & Ramstead, M. J. D. (2021). Deeply felt affect: The emergence of valence in deep active inference. *Neural Computation*, *33*(2), 398–446. https://doi.org/10.1162/neco_a_01341

Heylighen, F., Kingsbury, K., Lenartowicz, M., Harmsen, T., & Beigi, S. (2018). Social systems programming: Behavioural and emotional mechanisms co-opted for social control. Manuscript Submitted for Publication. http://pespmc1. *Vub*. Ac. be/Papers/SSP2mechanisms. Pdf. [SVdC]. http://pespmc1.vub.ac.be/Papers/SSP2mechanisms.pdf

Hill, A. (2021, March 23). 'A new obsession': The people who learned to play instruments during lockdown. *The Guardian*. www.theguardian.com/world/2021/mar/23/a-new-obsession-the-people-who-learnt-to-play-instruments-during-lockdown

Hirsh, J. B., Mar, R. A., & Peterson, J. B. (2012). Psychological entropy: A framework for understanding uncertainty-related anxiety. *Psychological Review*, *119*(2), 304–320. https://doi.org/10.1037/a0026767

Hoel, E. (2019). Enter the supersensorium. *The Baffler*, *45*, 118–133. www.jstor.org/stable/26639757?casa_token=sqMDDLFrA5cAAAAA:ANKYuvUQ34o8hrXXXMNkU-nj74nTaZevqPv6h1_kB6q-TGNg8eto6pxmbGTaF_Qz1D84e_05UhB2aae-kBFmEYYljoI17XQTfffvH7U5q-l4ePsVlEZi

Hoel, E. (2020). The Overfitted Brain: Dreams evolved to assist generalization. arXiv Preprint arXiv:2007.09560. http://arxiv.org/abs/2007.09560

Hoffman, D. D., Singh, M., & Prakash, C. (2015). The interface theory of perception. *Psychonomic Bulletin and Review*, *22*(6), 1480–1506. https://doi.org/10.3758/s13423-015-0890-8

Hohwy, J. (2014). The self-evidencing brain. *Noûs*, *50*(2), 259–285. https://doi.org/10.1111/nous.12062

Hohwy, J. (2020). New directions in predictive processing. *Mind and Language*, *35*(2), 209–223. https://onlinelibrary.wiley.com/doi/abs/10.1111/mila.12281

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

Hume, D. (2008). *An enquiry concerning human understanding*. https://doi.org/10.1093/owc/9780199549900.001.0001

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

Huron, D. (2006). *Sweet anticipation: Music and the psychology of expectation* (1st ed). MIT Press. http://www.amazon.com/dp/0262083450

Issa, E. B., Cadieu, C. F., & DiCarlo, J. J. (2018). Neural dynamics at successive stages of the ventral visual stream are consistent with hierarchical error signals. *eLife*, 7. https://doi.org/10.7554/eLife.42870

Jackendoff, R., & Lerdahl, F. (2006). The capacity for music: What is it, and what's special about it? *Cognition*, *100*(1), 33–72. https://doi.org/10.1016/j.cognition.2005.11.005

James, W. (1890). *The principles of psychology*. Harvard University Press. www.wwww4.com/read_book/_the_principles_of_psychology_1541318.pdf

Joffily, M., & Coricelli, G. (2013). Emotional valence and the free-energy principle. *PLOS Computational Biology*, *9*(6), e1003094. https://doi.org/10.1371/journal.pcbi.1003094

Juechems, K., & Summerfield, C. (2019). Where does value come from? *Trends in Cognitive Sciences*, *23*(10), 836–850. https://doi.org/10.1016/j.tics.2019.07.012

Kagan, J. (1972). Motives and development. *Journal of Personality and Social Psychology*, *22*(1), 51–66. https://doi.org/10.1037/h0032356

Kawakami, K., Phills, C. E., Steele, J. R., & Dovidio, J. F. (2007). (Close) distance makes the heart grow fonder: Improving implicit racial attitudes and interracial interactions through approach behaviours. *Journal of Personality and Social Psychology*, *92*(6), 957–971. https://doi.org/10.1037/0022-3514.92.6.957

Keller, G. B., & Mrsic-Flogel, T. D. (2018). Predictive processing: A canonical cortical computation. *Neuron*, *100*(2), 424–435. https://doi.org/10.1016/j.neuron.2018.10.003

Keller, H., & Voss, H.-G. (1983). *Curiosity and exploration: Theories and results*. Academic Press.

Kelly, G. A. (1958). Man's construction of his alternatives. *Assessment of Human Motives*, 33–64.

Kelly, G. A. (1964). The language of hypothesis: Man's psychological instrument. *Journal of Individual Psychology*, *20*(2), 137–152. www.aippc.it/wp-content/uploads/2019/04/2014.01.005.015.pdf

Kesner, L. (2014). The predictive mind and the experience of visual art work. *Frontiers in Psychology*, *5*, 1417. https://doi.org/10.3389/fpsyg.2014.01417

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

Kilner, J. M., Friston, K. J., & Frith, C. D. (2007). Predictive coding: An account of the mirror neuron system. *Cognitive Processing*, *8*(3), 159–166. https://doi.org/10.1007/s10339-007-0170-2

Kiverstein, J., Miller, M., & Rietveld, E. (2019). The feeling of grip: Novelty, error dynamics, and the predictive brain. *Synthese*, *196*(7), 2847–2869. https://doi.org/10.1007/s11229-017-1583-9

Kivetz, R., Urminsky, O., & Zheng, Y. (2006). The goal-gradient hypothesis resurrected: Purchase acceleration, illusionary goal progress, and customer retention. *Journal of Marketing Research*, *43*(1), 39–58. https://doi.org/10.1509/jmkr.43.1.39

Kizilirmak, J. M., Galvao Gomes da Silva, J., Imamoglu, F., & Richardson-Klavehn, A. (2016). Generation and the subjective feeling of "aha!" are independently related to learning from insight. *Psychological Research*, *80*(6), 1059–1074. https://doi.org/10.1007/s00426-015-0697-2

Koelsch, S., Vuust, P., & Friston, K. J. (2019). Predictive processes and the peculiar case of music. *Trends in Cognitive Sciences*, *23*(1), 63–77. https://doi.org/10.1016/j.tics.2018.10.006

Koenderink, J. (2010). Vision and information. In L. Albertazzi, G. J. Van Tonder, & D. Vishwanath (Eds.), *Perception beyond inference: The information content of visual processes* (p. 27). MIT Press. http://books.google.be/books?hl=nl&lr=&id=Kaw100wfsBYC&oi=fnd&pg=PA27&dq=koenderink+vision+information&ots=_4TRc59NZY&sig=ItPvAKWBOuI09r-nV-rUt8OcpVI

Koffka, K. (1935). *Principles of Gestalt psychology*. Brace.

Köhler, W. (1920). *Die physischen Gestalten in Ruhe und im stationären Zustand*. Vieweg.

Leavitt, J. D., & Christenfeld, N. J. S. (2011). Story spoilers don't spoil stories. *Psychological Science*, *22*(9), 1152–1154. https://doi.org/10.1177/0956797611417007

Leder, H., Bär, S., & Topolinski, S. (2012). Covert painting simulations influence aesthetic appreciation of artworks. *Psychological Science*, *23*(12), 1479–1481. https://doi.org/10.1177/0956797612452866

Leder, H., & Nadal, M. (2014). Ten years of a model of aesthetic appreciation and aesthetic judgments: The aesthetic episode—Developments and challenges in empirical aesthetics. *British Journal of Psychology*, *105*(4), 443–464. https://doi.org/10.1111/bjop.12084

Livson, N. (1967). Towards a differentiated construct of curiosity. *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), 75–98. http://psycnet.apa.org/record/1994-41058-001. https://doi.org/10.1037/0033-2909.116.1.75

McReynolds, P. (1971). The three faces of cognitive motivation. In H. I. Day, D. E. Berlyne & D. E. Hunt (Eds.), *Intrinsic motivation: A new direction in education* (pp. 33–45). Holt, Rinehart and Winston of Canada. https://market.android.com/details?id=book-o27GQgAACAAJ.

Menninghaus, W., Wagner, V., Hanich, J., Wassiliwizky, E., Jacobsen, T., & Koelsch, S. (2017). The distancing-embracing model of the enjoyment of negative emotions in art reception. *Behavioral and Brain Sciences*, *40*, e347. https://doi.org/10.1017/S0140525X17000309

Metcalfe, 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

Metzinger, T. (2007). Self models. *Scholarpedia*, *2*(10), 4174. www.scholarpedia.org/article/Self_models?&sa=U&ei=Y8UNVPSgLMK_ygPWgYLADw&ved=0CFMQFjAO&usg=AFQjCNEiwQEhiPd-fSDqXm0_kTQEYkMnUg

Meyer. (1961). *Emotion and meaning in music*. University of Chicago Press. www.amazon.co.uk/dp/0226521397

Miller, M., Nave, K., & Clark, G. D. (2020, September 25). Use uncertainty to leverage the power of your predictive brain—Mark Miller, Kathryn Nave. *George Deane & Andy Clark*. *Aeon*. https://aeon.co/essays/use-uncertainty-to-leverage-the-power-of-your-predictive-brain

Montoya, R. M., Horton, R. S., Vevea, J. L., Citkowicz, M., & Lauber, E. A. (2017). A re-examination of the mere exposure effect: The influence of repeated exposure on recognition, familiarity, and liking. *Psychological Bulletin*, *143*(5), 459–498. https://doi.org/10.1037/bul0000085

Mrkva, K., & Van Boven, L. (2020). Salience theory of mere exposure: Relative exposure increases liking, extremity, and emotional intensity. *Journal of Personality and Social Psychology*, *118*(6), 1118–1145. https://doi.org/10.1037/pspa0000184

Muth, C., & Carbon, C.-C. (2016). SeIns: Semantic instability in art. *Art and Perception*, *4*(1–2), 145–184. https://doi.org/10.1163/22134913-00002049

Muth, C., Raab, M. H., & Carbon, C. C. (2015). The stream of experience when watching artistic movies. Dynamic aesthetic effects revealed by the continuous evaluation procedure (CEP). *Frontiers in Psychology*, *6*, 365. https://doi.org/10.3389/fpsyg.2015.00365

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

Nascimento, S. M. C., Marit Albers, A., & Gegenfurtner, K. R. (2021). Naturalness and aesthetics of colors—Preference for color compositions perceived as natural. *Vision Research*, *185*, 98–110. https://doi.org/10.1016/j.visres.2021.03.010

Nguyen, C. T. (2017). Philosophy of games. *Philosophy Compass*, *12*(8), e12426. https://doi.org/10.1111/phc3.12426

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

Oudeyer, P.-Y., Kaplan, F., & Hafner, V. V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation*, *11*(2), 265–286. https://doi.org/10.1109/TEVC.2006.890271

Pearce, M. T. (2018). Statistical learning and probabilistic prediction in music cognition: Mechanisms of stylistic enculturation. *Annals of the New York Academy of Sciences*. https://doi.org/10.1111/nyas.13654

Pelowski, M., Markey, P. S., Forster, M., Gerger, G., & Leder, H. (2017). Move me, astonish me . . . delight my eyes and brain: The Vienna integrated model of top-down and bottom-up processes in art perception (VIMAP) and corresponding affective, evaluative, and neurophysiological correlates. *Physics of Life Reviews*, *21*, 80–125. www.sciencedirect.com/science/article/pii/S1571064517300325

Pepperell, R. (2018). Chapter 19. Art, energy, and the brain. In J. F. Christensen & A. Gomila (Eds.), *Progress in brain research* (Vol. 237, pp. 417–435). Elsevier. https://doi.org/10.1016/bs.pbr.2018.03.022

Pezzulo, G., Rigoli, F., & Friston, K. (2015). Active Inference, homeostatic regulation and adaptive behavioural control. *Progress in Neurobiology*, *134*, 17–35. https://doi.org/10.1016/j.pneurobio.2015.09.001

Pinna, B. (2010). New Gestalt principles of perceptual organization: An extension from grouping to shape and meaning. *Gestalt Theory*, *32*, 1–67.

Poetry Foundation. (2021a). *Negative capability*. www.poetryfoundation.org/learn/glossary-terms/negative-capability

Poetry Foundation. (2021b). *Selections from Keats's letters by John Keats*. www.poetryfoundation.org/articles/69384/selections-from-keatss-letters

Powers, W. T. (1973). *Behaviour: The control of perception*, *xi*. Aldine.

Reber, R., Schwarz, N., & Winkielman, P. (2004). Processing fluency and aesthetic pleasure: Is beauty in the perceiver's processing experience? *Personality and Social Psychology Review*, *8*(4), 364–382. https://doi.org/10.1207/s15327957pspr0804_3

Reisenzein, R. (2009). Emotional experience in the computational belief–desire theory of emotion. *Emotion Review*, *1*(3), 214–222. https://doi.org/10.1177/1754073909103589

Rohe, T., & Noppeney, U. (2015). Cortical hierarchies perform Bayesian causal inference in multisensory perception. *PLOS Biology*, *13*(2), e1002073. https://doi.org/10.1371/journal.pbio.1002073

Ryali, C. K., Goffin, S., Winkielman, P., & Yu, A. J. (2020). From likely to likable: The role of statistical typicality in human social assessment of faces. *Proceedings of the National Academy of Sciences of the United States of America*, *117*(47), 29371–29380. https://doi.org/10.1073/pnas.1912343117

Sarasso, P., Neppi-Modona, M., Sacco, K., & Ronga, I. (2020). 'Stopping for knowledge': The sense of beauty in the perception-action cycle. *Neuroscience and Biobehavioral Reviews*, *118*, 723–738. https://doi.org/10.1016/j.neubiorev.2020.09.004

Sarasso, P., Perna, P., Barbieri, P., Neppi-Modona, M., Sacco, K., & Ronga, I. (2021). Memorisation and implicit perceptual learning are enhanced for preferred musical intervals and chords. *Psychonomic Bulletin and Review*, *28*(5), 1623–1637. https://doi.org/10.3758/s13423-021-01922-z

Schillaci, G., Ciria, A., & Lara, B. (2020). Tracking emotions: Intrinsic motivation grounded on multi-level prediction error dynamics. In *Joint IEEE 10th international conference on development and learning and epigenetic robotics (ICDL-EpiRob)* (pp. 1–8). https://doi.org/10.1109/ICDL-EpiRob48136.2020.9278106

Schmidhuber, J. (2009). Simple algorithmic theory of subjective beauty, novelty, surprise, interestingness, attention, curiosity, creativity, art, science, music, jokes. *Journal of SICE*, *48*(1). www.idsia.ch/~juergen/sice2009.pdf.

Schneider, K., & Heckhausen, H. (1981). Subjective uncertainty and task preference. In H. I. Day (Ed.), *Advances in intrinsic motivation and aesthetics* (pp. 149–167). Springer. https://doi.org/10.1007/978-1-4613-3195-7_7

Schoeller, F., & Perlovsky, L. (2016). Aesthetic chills: Knowledge-acquisition, meaning-making, and aesthetic emotions. *Frontiers in Psychology*, 7, 1093. https://doi.org/10.3389/fpsyg.2016.01093

Schoeller, F., Perlovsky, L., & Arseniev, D. (2018). Physics of mind: Experimental confirmations of theoretical predictions. *Physics of Life Reviews*, *25*, 45–68. https://doi.org/10.1016/j.plrev.2017.11.021

Schoeller, F., Perlovsky, L., et al. (2015). Great expectations—Narratives and the elicitation of aesthetic chills. *Psychology*, *6*(16), 2098. www.scirp.org/html/5-6901655_62308.htm

Schonberg, T., & Katz, L. N. (2020). A neural pathway for nonreinforced preference change. *Trends in Cognitive Sciences*, *24*(7), 504–514. https://doi.org/10.1016/j.tics.2020.04.002

Sengupta, B., Stemmler, M. B., & Friston, K. J. (2013). Information and efficiency in the nervous system—A synthesis. *PLOS Computational Biology*, *9*(7), e1003157. https://doi.org/10.1371/journal.pcbi.1003157

Seth, A. K. (2019). From unconscious inference to the beholder's share: Predictive perception and human experience. *European Review*, *27*(3), 378–410. https://doi.org/10.1017/S1062798719000061

Seth, A. K., Millidge, B., Buckley, C. L., & Tschantz, A. (2020). 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

Sharot, T. (2011). The optimism bias. *Current Biology: CB*, *21*(23), R941–R945. https://doi.org/10.1016/j.cub.2011.10.030

Shklovsky, V. (1917). Art as technique. *Literary Theory: An Anthology*, *3*. https://books.google.be/books?hl=en&lr=&id=XIL5DQAAQBAJ&oi=fnd&pg=RA1-PA8&dq=shklovsky+technique&ots=kldegdtA5O&sig=ktxgstenMqqOypFXcoMbIk577Yg.

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

Singh, M. (2021). The sympathetic plot, its psychological origins, and implications for the evolution of fiction. *Emotion Review*, *13*(3), 183–198. https://doi.org/10.1177/17540739211022824

Skov, M., & Nadal, M. (2020). A farewell to art: Aesthetics as a topic in psychology and neuroscience. *Perspectives on Psychological Science*, *15*(3), 630–642. https://doi.org/10.1177/1745691619897963

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

Smith, R., Friston, K., & Whyte, C. (2021). A step-by-step tutorial on active inference and its application to empirical data. https://doi.org/10.31234/osf.io/b4jm6

Solms, M. (2018). The hard problem of consciousness and the free energy principle. *Frontiers in Psychology*, *9*, 2714. https://doi.org/10.3389/fpsyg.2018.02714

Spehar, B., Walker, N., & Taylor, R. P. (2016). Taxonomy of individual variations in aesthetic responses to fractal patterns. *Frontiers in Human Neuroscience*, *10*, 350. https://doi.org/10.3389/fnhum.2016.00350

Spratling, M. W. (2017). A review of predictive coding algorithms. *Brain and Cognition*, *112*, 92–97. https://doi.org/10.1016/j.bandc.2015.11.003

Srivastava, N., & Schrater, P. (2015). Learning what to want: Context-sensitive preference learning. *PLoS One*, *10*(10), e0141129. https://doi.org/10.1371/journal.pone.0141129

Sun, Z., & Firestone, C. (2020). The dark room problem. *Trends in Cognitive Sciences*, *24*(5), 346–348. https://doi.org/10.1016/j.tics.2020.02.006

Taleb, N. N. (2013). 'Antifragility'as a mathematical idea. *Nature*, *494*(7438), 430–430. www.nature.com/articles/494430e

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

Tracy, H. L. (1946). Aristotle on aesthetic pleasure. *Classical Philology*, *41*(1), 43–46. https://doi.org/10.1086/362919

Trope, Y. (1975). Seeking information about one's ability as a determinant of choice among tasks. *Journal of Personality and Social Psychology*, *32*(6), 1004–1013. https://doi.org/10.1037/0022-3514.32.6.1004

Tschantz, A., Seth, A. K., & Buckley, C. L. (2020). Learning action-oriented models through active inference. *PLOS Computational Biology*, *16*(4), e1007805. https://doi.org/10.1371/journal.pcbi.1007805

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. (2018). Upgrading Gestalt psychology with variational neuroethology: The case of perceptual pleasures: Comment on "Answering Schrödinger's question: A free-energy formulation" by M.J. Desormeau Ramstead et al. *Physics of Life Reviews*, *24*(1), 21–23. https://lirias.kuleuven.be/retrieve/503242

Van de Cruys, S., Chamberlain, R., & Wagemans, J. (2017). Tuning in to art: A predictive processing account of negative emotion in art. *Behavioral and Brain Sciences*, *40*, e377. https://doi.org/10.1017/S0140525X17001868

Van de Cruys, S., Damiano, C., Boddez, Y., Król, M., Goetschalckx, L., & Wagemans, J. (2021). Visual affects: Linking curiosity, Aha-Erlebnis, and memory through information gain. *Cognition*, *212*, 104698. https://doi.org/10.1016/j.cognition.2021.104698

Van de Cruys, S., Evers, K., Van der Hallen, R., Van Eylen, L., Boets, B., de-Wit, L., & Wagemans, J. (2014). Precise minds in uncertain worlds: Predictive coding in autism. *Psychological Review*, *121*(4), 649–675. https://doi.org/10.1037/a0037665

Van de Cruys, S., Friston, K. J., & Clark, A. (2020). Controlled optimism: Reply to Sun and Firestone on the dark room problem. *Trends in Cognitive Sciences*, *24*(9), 680–681. https://doi.org/10.1016/j.tics.2020.05.012

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 Dessel, P., Hughes, S., & De Houwer, J. (2019). How do actions influence attitudes? An inferential account of the impact of action performance on stimulus evaluation. *Personality and Social Psychology Review*, *23*(3), 267–284. https://doi.org/10.1177/1088868318795730

Van Geert, E., & Wagemans, J. (2020). Order, complexity, and aesthetic appreciation. *Psychology of Aesthetics, Creativity, and the Arts*, *14*(2), 135–154. 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. *Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *38*(10), 2579–2588. https://doi.org/10.1523/JNEUROSCI.2816-17.2018

Von Foerster, H. (1960). *On self-organizing systems and their environments* (pp. 31–50). Self-Organizing Systems. www.researchgate.net/profile/Mariana-Broens/publication/271732162_Complexity_and_information_technologies_an_ethical_inquiry_into_human_autonomous_action/links/564a83ca08ae127ff986aa22/Complexity-and-information-technologies-an-ethical-inquiry-into-human-autonomous-action.pdf

von Glasersfeld, E. (1995). Radical constructivism: A way of knowing and learning. Studies In Mathematics Education Series: 6. *Falmer press, Taylor and Francis Inc*. http://eric.ed.gov/?id=ED381352

Vygotsky, L. S. (1962). *Thought and word*. https://psycnet.apa.org/record/2006-10268-007

Wacongne, C., Labyt, E., van Wassenhove, V., Bekinschtein, T., Naccache, L., & Dehaene, S. (2011). Evidence for a hierarchy of predictions and prediction errors in human cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(51), 20754–20759. https://doi.org/10.1073/pnas.1117807108

Wagemans, J. (2015). Historical and conceptual background: Gestalt theory. The Oxford handbook of perceptual. *Organization*, 3–20. www.gestaltrevision.be/pdfs/oxford/Wagemans-Historical_and_conceptual_background_Gestalt_theory.pdf

Wagemans, J., Elder, J. H., Kubovy, M., Palmer, S. E., Peterson, M. A., Singh, M., & Von der, H. R. (2012). A century of Gestalt psychology in visual perception: I. Perceptual grouping and figure-ground organization. *Psychological Bulletin*, *138*(6), 1172–1217. https://doi.org/10.1037/a0029333

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, X., Bylinskii, Z., Hertzmann, A., & Pepperell, R. (2020). Toward quantifying ambiguities in artistic images. *ACM Transactions on Applied Perception*, *17*(4), 1–10. https://doi.org/10.1145/3418054

Wänke, M., & Hansen, J. (2015). Relative processing fluency. *Current Directions in Psychological Science*, *24*(3), 195–199. https://doi.org/10.1177/0963721414561766

Weick, K. E. (1979). The social psychology of organizing. Random House. https://play.google.com/store/books/details?id=gHJPAAAAMAAJ

Wilkinson, S., Deane, G., Nave, K., & Clark, A. (2019). Getting warmer: Predictive processing and the nature of emotion. In L. Candiotto (Ed.), *The value of emotions for knowledge* (pp. 101–119). Springer International Publishing. https://doi.org/10.1007/978-3-030-15667-1_5

Wilkinson, Z., Cunningham, R., & Elliott, M. A. (2021). The influence of empathy on the perceptual response to visual art. *Psychology of Aesthetics, Creativity, and the Arts*. https://doi.org/10.1037/aca0000418

Williams, D. (2020). Predictive coding and thought. *Synthese*, *197*(4), 1749–1775. https://doi.org/10.1007/s11229-018-1768-x

Winkielman, P., Schwarz, N., Fazendeiro, T., Reber, R., et al. (2003). The hedonic marking of processing fluency: Implications for evaluative judgment. *Psychology of Evaluation: Affective Processes in Cognition and Emotion*, *189*, 217. https://books.google.com/books?hl=en&lr=&id=t1h6AgAAQBAJ&oi=fnd&pg=PA195&dq=fazendeiro+winkielman&ots=bDBy66NH8_&sig=lve9qpJ1K5569bs3U3cGksDne58

Witek, M. A. G., Clarke, E. F., Wallentin, M., Kringelbach, M. L., & Vuust, P. (2015). Correction: Syncopation, body-movement and pleasure in groove music. *PLoS One*, *10*(9), e0139409. https://doi.org/10.1371/journal.pone.0139409

Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology*, *9*(2, Pt.2), 1–27. https://doi.org/10.1037/h0025848

Zajonc, R. B. (1980). Feeling and thinking: Preferences need no inferences. *American Psychologist*, *35*(2), 151–175. https://doi.org/10.1037/0003-066X.35.2.151

Zangwill, N. (2021). Aesthetic judgment. In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy*. Metaphysics Research Laboratory, Stanford University. https://plato.stanford.edu/archives/spr2021/entries/aesthetic-judgment/

Zénon, A., Solopchuk, O., & Pezzulo, G. (2019). An information-theoretic perspective on the costs of cognition. *Neuropsychologia*, *123*, 5–18. https://doi.org/10.1016/j.neuropsychologia.2018.09.013