Beyond individual features: Identifying computational principles in perceiving visual beauty

Mustafa Alperen Ekinci[1,2,3], Philipp Flieger[1,3], and Daniel Kaiser[1,2,3,4]

[1] Neural Computation Group, Justus Liebig University Giessen, 35394 Giessen, Germany
[2] Center for Mind, Brain and Behavior (CMBB), Universities of Giessen, Marburg, and Darmstadt, 35032 Marburg, Germany
[3] Center for Applied Computer Science and Data Science (ZAD), Justus Liebig University Giessen, 35392 Giessen, Germany
[4] Cluster of Excellence "The Adaptive Mind", Universities of Giessen, Marburg, and Darmstadt, 35394 Giessen, Germany

Abstract:

Prior research in empirical aesthetics has related beauty to visual features like curvature, color, or shape. Although these feature-based accounts have identified reliable visual correlates of beauty, the literature on feature preferences is highly fragmented, with individual features studied for individual stimulus domains. Here, we argue that making further progress hinges on complementing feature-based approaches with approaches aimed at distilling computational principles that underlie the perception of beauty. Recent advances in computational neuroscience can drive this new line of research by linking visual beauty to the coding of visual information in artificial neural networks and the biological brain. We highlight two examples of how this approach has identified overarching computational principles. First, metabolic processing costs, both in artificial and biological vision, are inversely related to perceived beauty, suggesting that beauty may reflect energy-efficient, fluent visual coding. Second, integrative processing in artificial and biological vision captures the efficient spatial representation of image elements across space and predicts beauty ratings. Synthesizing these findings, we derive a neuro-computational framework for assessing computational principles in the perception of beauty, which uses artificial neural networks for exploration at scale and then validates candidate principles using human behavior and brain data. We discuss candidate principles that could be explored using this approach, like spatiotemporal predictability, and delineate how the new insights could be reconciled with feature-based approaches. In conclusion, complementing feature-based studies with studies targeting computational principles promises a comprehensive and mechanistic understanding of why certain visual stimuli are perceived as beautiful, and others are not.

Keywords: Neuroaesthetics, computational principles, neural coding, processing fluency, artificial neural networks

Introduction

In our everyday lives, visual beauty can be found all around us, in artworks and landscapes, and in objects and scenes. One can experience beauty unexpectedly while cycling to work or deliberately in a museum. In both cases, encountering beauty enriches our lives and brings us joy.

How can we explain visual beauty, and which factors determine a stimulus's aesthetic appeal? These questions have been asked since Fechner's first empirical investigations on image preferences in the nineteenth century (Fechner, 1876; Bader, 2018). We have come a long way from these classical studies. Based on decades of empirical research, we now have different theoretical models that characterize the biological basis of visual aesthetics (Leder, 2004; Palmer et al., 2013; Chatterjee & Vartanian, 2014; Redies, 2015). Although these models differ in proposed mechanisms and stress different cognitive processes, they all converge on the notion that sensory processing is the critical backbone: Models assume "perception" (Leder, 2004), "sensory-motor" (Chatterjee & Vartanian, 2014), or "sensory coding" (Redies, 2015) steps, where sensory features are the first stepping stone towards aesthetic appeal.

The importance of visual features in aesthetic appreciation is well established. Studies in empirical aesthetics have identified a range of low- and mid-level visual features associated with aesthetic appeal, including curvature (Corradi et al., 2019; Vartanian et al., 2024), color (Graham et al., 2016; Schloss & Palmer, 2011; Swartz et al., 2024), complexity (Van Geert & Wagemans, 2020), and symmetry (Rhodes et al., 1998; Huang et al., 2018), among others. Visual features can reliably explain domain-specific beauty (e.g., symmetric vs asymmetric faces) and unveil shared aesthetic preferences (e.g., preferences for curved rather than angular objects). More recently, studies have shown that combining multiple features into predictive models can yield powerful predictions of aesthetic judgements for static (Iigaya et al., 2021) and dynamic (Ekinci et al., 2026) stimuli. This shows that multiple visual features are reliably and systematically linked to perceived beauty.

Yet, even though we have learned a lot from these studies, answers to a more general question remain elusive: *Why* are these specific features preferred? For instance, why is a curved object preferred over a rectilinear one? Why is a symmetric face preferred over an asymmetric one? Though for each of these examples, individual explanations have been proposed (Rhodes, 2006; Bar & Neta, 2006), we argue that there may be overarching *computational principles*, which are fundamentally ingrained in human visual processing, and which may explain perceived beauty across individual features and across stimulus domains. In other words, they may ultimately at the same time explain why we prefer curved objects and symmetric faces.

In the following, we use computational principles to refer to neural coding motifs in sensory systems. Such neural coding motifs form specific and recurring patterns of neural information processing that characterize or constrain the processing of visual information in the brain. Critically, we ask whether and how we can identify

computational principles that are indicative of perceived beauty - without necessarily being tied to individual features or stimulus domains.

A useful starting point for thinking about such computational principles is to ask why the visual system should represent the beauty of visual inputs in the first place. A tractable proposal is offered by *processing fluency* theories (Reber et al., 2004; Graf & Landwehr, 2015), which posit that beauty arises from enhanced processing efficiency in sensory systems. This idea offers concrete suggestions for candidate computational principles: Neural motifs that facilitate efficient coding in resource-limited sensory systems should be linked to perceived beauty. Computational principles consistent with such efficient coding may include neural motifs that reduce neural competition through integration (Röth et al., 2021), increase efficiency by sharpening representations (Pouget et al., 2000), integrate sensory inputs with predictions based on prior expectations (Rao & Ballard, 1999), or facilitate neural communication through oscillatory dynamics (Fries, 2015; Stecher et al., 2025). While these candidate principles are an arbitrary subset, they can all operate across multiple levels of the visual hierarchy and thus transcend individual feature or stimulus preferences.

As these examples also illustrate, there are numerous candidate principles that facilitate efficient representation in the visual brain, and there are potentially several different neural indicators for each of them. Establishing reliable links between a computational principle and perceived beauty therefore requires tedious, time-consuming, and expensive experimentation. As an alternative, however, modern artificial neural networks (ANNs) offer cheap and high-throughput in silico testbeds, where computational principles can be examined at scale and in turn linked to human beauty ratings (Figure 1A). Once a candidate principle has been identified, it can be computationally benchmarked in an ANN architecture using a set of suitable measures. Variations in these measures can then be linked to variations in beauty ratings across stimuli. After establishing links between a computational principle in an ANN and perceived beauty in human observers, this association can be biologically validated using targeted human neuroimaging. This validation is important, given that ANNs, although in many cases suitable for modelling human vision (Kietzmann et al., 2017; Cichy & Kaiser, 2019; Doerig et al., 2023), can show systematic divergences from human perception and neural processing (Xu & Vaziri-Pashkam, 2021; Bowers et al., 2022).

In the following, we show the feasibility of this framework by highlighting two recent studies that used ANNs, and specifically deep neural networks (DNNs), in conjunction with brain imaging data to identify, benchmark, and validate computational principles that underlie the perception of beauty. Specifically, Tang et al. (2025) investigated *metabolic cost* (i.e., the relative energy expenditure in artificial and biological vision associated with processing a stimulus), and Nara and Kaiser (2024) investigated *integrative processing* (i.e., the relative degree of spatial integration among image elements).

Metabolic cost

Inspired by fluency-based theories, Tang et al. (2025) hypothesized that sharper neural tuning and its associated reduction in energy expenditure are indicative of perceived beauty. Consequently, they proposed an inverse relationship between the metabolic cost of neural activation and aesthetic preferences for visual stimuli. Akin to the framework outlined in Fig. 1A, they first benchmarked the proposed association in a DNN model and then validated the association using biological brain data.

Computational benchmarking. For computational benchmarking, the authors used a VGG19 model pre-trained on object recognition (Simonyan & Zisserman, 2014). The total number of active filters in the DNN was used as a measure for in silico metabolic cost. This measure was extracted for 4,916 natural images (from BOLD5000; Chang et al., 2019). Across images, metabolic cost in the DNN was negatively correlated with beauty ratings from an independent group of human observers (Fig. 1B), confirming the hypothesis that beauty is linked to the energy efficiency of stimulus processing. An examination of layer-wise metabolic cost across the DNN hierarchy revealed stronger predictions when moving from shallower to deeper layers, suggesting that metabolic costs associated with high-level feature coding are most indicative of visual beauty.

Biological validation. To validate the inverse relationship between metabolic cost and beauty using biological brain data, the authors used fMRI data from four participants who viewed the same 4,916 images shown to the DNN (Chang et al., 2019). The sum of the beta values for each image and for each region in the brain served as a measure of biological metabolic cost. Confirming the computational benchmarking, beauty ratings across images were negatively correlated with metabolic costs in the scene-selective parahippocampal place area (PPA) and occipital place area (OPA), and the face-selective fusiform face area (FFA) (Fig. 1C). These findings suggest that lower energy expenditure in the human visual system, similarly to lower processing costs in visual DNNs, is indicative of greater perceived beauty. As in the DNN, this association was strongest for complex feature representations in category-selective regions of high-level visual cortex.

To sum up, this study first hypothesized metabolic cost to be a candidate principle tied to visual beauty, then used DNN benchmarking as a testbed for this association and in turn validated the association using human neuroimaging data. The resulting association between metabolic cost and perceived beauty not only shows the feasibility of the proposed framework. It also highlights a computational principle (energy efficiency) that can operate across multiple stages of the visual hierarchy and thus promises exploratory power beyond individual features and stimulus domains. Further support for this principle comes from a related study that investigated the sparsity of DNN representations for faces and paintings (Dibot et al., 2023).

Integrative processing

Starting from fluency-based considerations, too, Nara and Kaiser (2024) investigated a complementary computational principle that increases coding efficiency: integrative processing. Inspired by human neuroimaging work (MacMains & Kastern, 2011; Kaiser et al. 2014; Kaiser & Peelen, 2018), they hypothesized that the efficient integration of image elements across space would reduce processing demands via a compression of information into holistic representational formats. The consequently lower processing demands should lend greater beauty to images that are more readily integrated across space. The authors thus hypothesized a positive correlation between the degree of spatial integration and perceived beauty across images. The hypothesized association was again benchmarked in a computational model and then validated using human brain data.

Computational benchmarking. For computational benchmarking, the authors used a VGG16 model pre-trained on scene recognition (Simonyan & Zisserman, 2014; Zhou et al., 2017). The degree of spatial integration for each image was measured as the sign-inverted correlation of DNN activity patterns for the whole image and an average of the activity patterns for complementary image halves (generated from opposing checkerboard patterns overlaid on the image). Lower correlations were taken as indicative of more integration (i.e., the whole is differently represented from the halves), whereas higher correlations were indicative of less integration (i.e., the whole is similarly represented to the halves). Across images, the degree of integration in the DNN, and particularly in intermediate to late layers, was positively correlated with beauty ratings from an independent group of human observers (Fig. 1D), confirming the hypothesis that perceived beauty is linked to how efficiently the brain integrates information across whole images. This relationship remained robust across four experiments with different sets of natural images, from landscapes to varied real-world scenes.

Biological validation. To validate the relationship between integrative processing and beauty using biological brain data, the authors conducted an fMRI experiment where observers passively viewed a subset of natural scene images and image halves. Akin to the DNN analysis, correlations of fMRI response patterns across the whole scene and the two complementary halves were used to quantify integrative processing, with lower whole-to-halves correlations indicating more integration and greater whole-to-halves correlations indicating less integration. Confirming the computational benchmarking, integrative processing was positively correlated with perceived beauty across the visual cortex, with a significant correlation obtained in scene-selective PPA (Fig. 1E). This suggests that the efficient spatial integration of scene features in high-level visual cortex is linked to perceived beauty.

To sum up, this study applied the proposed research framework to distill integrative processing as a computational principle connected to the perception of beauty. Integrative processing provides another great example of a computational principle that can operate across feature hierarchies and across stimulus domains. In fact, a more

recent study has already shown that integrative processing is similarly, albeit somewhat more weakly, associated with preferences for visual artworks, too (Lin et al., 2026).

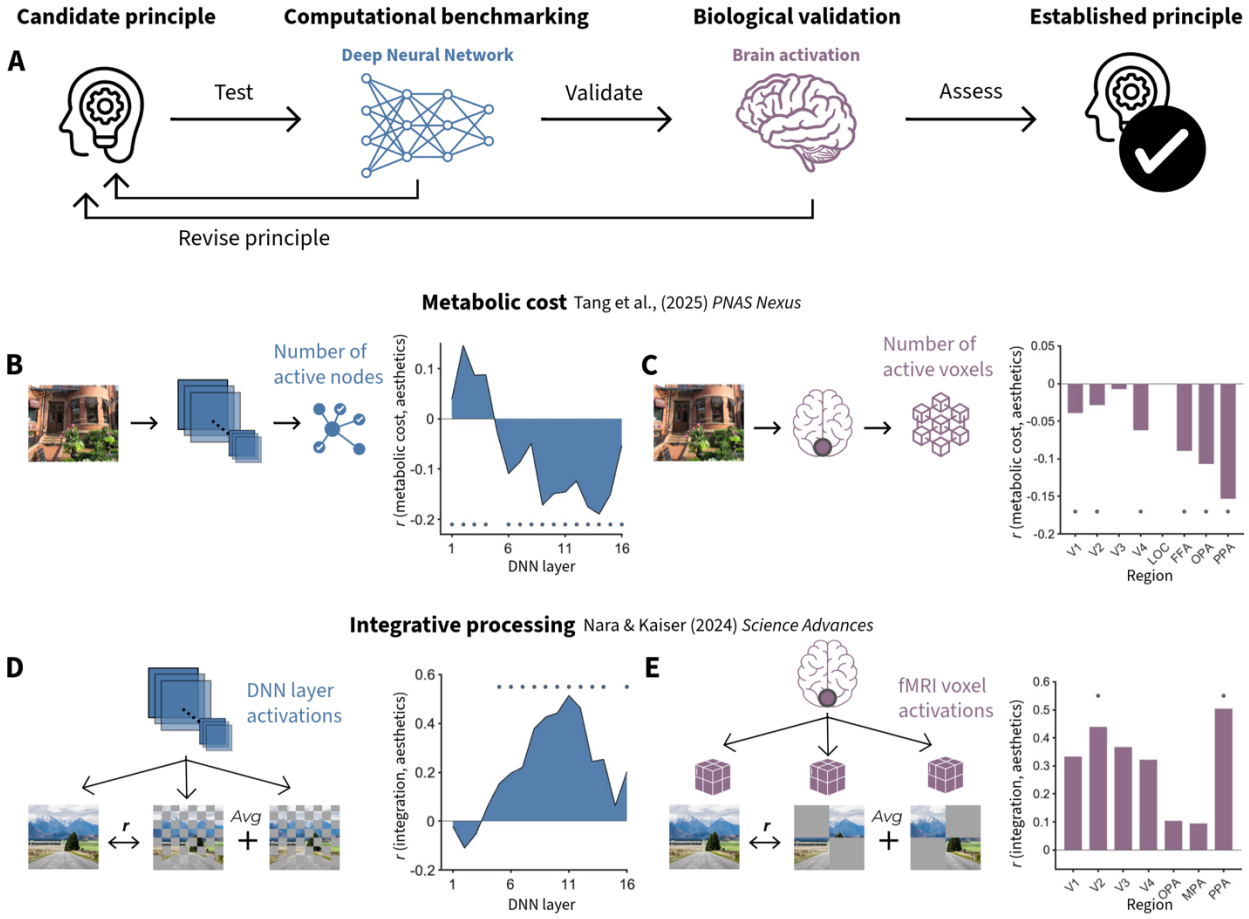


Figure 1. Outline of the framework and application in two example studies. (A) The diagram shows the process of testing and validating a candidate computational principle. Candidate principles are first tested through computational benchmarking, continuously revising the computational implementation in the process. If successful prediction is observed, the idea is validated biologically in human neuroimaging. Finally, the principle is established as a predictor of perceived beauty or is revised again. (B) Example 1: Metabolic cost. Here, it is hypothesized that lower metabolic processing cost is linked to increased perceived beauty. The association between decreased energy cost and higher visual beauty was first tested on a DNN, and (C) the same approach was applied to human brain data, in which active voxels were used as a measure of energy cost. (D) Example 2: Integrative processing. Here, it is hypothesized that a greater extent of spatial integration across images, resulting in more compressed and efficient representations, is linked to increased perceived beauty. Whether integrative processing predicts visual beauty was also first tested on a DNN, and (E) the relationship was then validated on the brain data, applying the same analysis approach as in the DNN. Images are licensed under CC BY. Photographers, in order: Tom Hall and Jeremy T. Hetzel. Images were obtained through Openverse.

A neuro-computational framework for identifying overarching principles

Both highlighted studies follow a very similar framework to test and validate candidate neuro-computational principles, combining modern artificial intelligence (AI) technologies for benchmarking and neuroscientific experiments for validation. Their approach exemplifies a new line of research with a hypothesis-driven neuro-computational framework (Fig. 1A). They also highlight the feasibility of this framework for establishing overarching principles in the neural coding of beauty, which transcend the feature preferences identified in previous work.

Now, how can we further build on this framework? There are already other candidate principles on the horizon that could be benchmarked and validated using the proposed framework. Spatiotemporal redundancy seems promising in this regard. We know that, already at the photoreceptor level, neuronal signals are reformatted into a decorrelated representational format to decrease redundancy and increase efficiency (Dan et al., 1996). Similar processes operating across the whole visual hierarchy can make use of the predictability of information across space and time. This redundancy reduction could be related to perceived beauty, as decreasing redundancy will yield more fluent processing of the compressed inputs. While this constitutes just one promising example of a candidate principle, many others are readily testable using the framework we highlight here. These could include principles such as representational sharpening (Pouget et al., 2000), prediction-error minimization (Van de Cruys & Wagemans, 2011), and an optimal balance between feedforward and feedback communication (Herz et al., 2020), each constituting a distinct computational mechanism underlying aesthetic processing.

While many of these principles are directly derived from fluency-based theories, other theoretical views may posit different candidate principles. For instance, reward-based models of aesthetic value also emphasize that current and expected future learning predict aesthetic experiences (Brielmann & Dayan, 2022). Other candidates could be derived from pioneer theories of aesthetics, such as Berlyne's psychobiological theory (1971), in which complexity and novelty are related to beauty in a U-shaped curve, and Martindale's cognitive theory (1988), which suggests prototypicality and meaningfulness are determinants of beauty, yielding a fresh look at classically proposed mechanisms.

Moving beyond prediction

When using ANNs in service of distilling computational principles that are biologically derived and interpretable, we can find new explanations for *why* certain stimuli are preferred over others. This approach is complementary to using ANNs as prediction machines in the context of maximizing accuracy in predicting whether or not a stimulus is preferred. ANNs, like DNNs, naturally excel at this task, with their high-dimensional representations offering powerful predictions of human beauty judgments when directly regressed onto them. For instance, Conwell et al. (2025) showed that a set of 180 different ANN architectures trained on visual recognition could robustly predict human

beauty ratings. By now, similar findings have been obtained across a range of studies and domains (Lu et al., 2014; Seresinhe et al., 2017; Jin et al., 2019; Nara & Kaiser, 2024). However, such studies largely rely on “brute-force” approaches, where many features are jointly exploited to enable high-accuracy predictions. Unlike the approach proposed here (Fig. 1A), this essentially leaves the underlying mechanisms hidden in a black box.

This does not necessarily mean that such approaches cannot yield insights further down the line, where opening the black box and examining the high-dimensional representations that drive predictions can yield further insights. For instance, a recent study by Soydaner and Wagemans (2026) predicted aesthetic evaluations using human-rated image attributes by employing different machine learning models trained on those attributes. To clarify the importance of each feature, they used an explainable AI technique, Shapley Additive Explanations, by measuring how the prediction performance changes when excluding and including each feature across different combinations of other features. This technique allowed them to assess the contribution of specific features to predicting aesthetic evaluations while exploring whether there is a nonlinear relationship between aesthetic evaluations and network features. In the future, further research along these lines may unify approaches based on first principles and approaches based on maximizing predictive power.

Finally, using ANNs to benchmark candidate computational principles also highlights a critical feature of ANNs in cognitive neuroscience: The approach leverages ANNs as scientific tools for exploring hypothesis spaces (Cichy & Kaiser, 2019; Wichmann & Geirhos, 2023). Here, the power of ANNs as reasonably good models of human visual processing enables cheap and high-throughput testing of hypothesized associations, which in turn can be validated in bespoke human experiments that use targeted paradigms and stimuli. In the future, this computational exploration of putative effects can even be fully automatized into an iterative hypothesis generation and testing loop (Lu et al., 2024; Musslick et al., 2025; Kabra et al., 2026) that automatically produces a set of computational principles that can then be validated.

Closing the loop: relating computational principles and feature preferences

At this point, it is important to stress that we do not advocate for an approach that replaces studies on feature preferences. Such studies are critically important for characterizing the building blocks of visual preferences within and across domains. We rather advocate for a complementary approach that coexists with studies of feature preferences, providing a new starting point for adding explanations for feature preferences on a more abstract level of description.

But how do we go forward and bring these approaches together to distill such explanations? After identifying neuro-computational principles (e.g., metabolic cost or integrative processing), we can use these new insights and directly test whether established feature preferences can be derived from these first principles. Taking the example of metabolic cost, the common reason curved lines are preferred over angular ones, and symmetric faces are preferred over non-symmetric ones may be that they are

both coded more fluently, requiring less energy expenditure in neural circuits. Following this example, one could utilize previous stimuli in experiments relating curvature and beauty (or symmetry and beauty), compute the desired measure in an ANN (here: metabolic cost), and test whether it is systematically related to curvature and whether the shared variance in turn explains human beauty ratings. Work along these lines thus has the potential to unite approaches based on computational principles and studies focusing on feature preferences. Some first steps in this direction have been made. For instance, Nara and Kaiser (2024), in their spatial integration study, tested which visual features contributed most to spatial integration in DNNs and found that the integration of high spatial frequency contours was most prominently linked to the prediction of perceived beauty. This was achieved by manipulating the visual feature content of the images while repeating the in silico benchmarking in their DNN models. Moving forward, however, it may turn out to be relatively tricky to relate individual visual features and a stimulus optimally "fulfilling" the computational principles in question: Most everyday stimuli are defined by a myriad of visual features and relevant dimensions, and a stimulus may reach an optimal point in this feature space through intricate interactions between its many features.

Identifying overarching computational principles also promises to reconcile the literature on feature preferences in cases where it does not fully converge. A good example is the relationship between complexity and perceived beauty. While some studies found that complexity relates to beauty following an inverted U-shape (Berlyne, 1971; Nadal et al., 2010), others reported negative or positive linear relationships (Tinio & Leder, 2009; Nadal et al., 2010; Delplangue et al., 2019). If we can identify computational principles that explain why complexity is related to beauty, we can then scrutinize under which conditions certain stimuli and domains produce different types of relationships between complexity and beauty: The different relationships observed in different studies may directly follow from a differential alignment with the more overarching coding motif that relates complexity and beauty in the first place.

Principles of sensory coding and what cognition does with them

Here, our focus has been strictly on perceptual processing and how it relates to beauty. However, it is possible that the framework proposed here offers a fruitful starting point for investigating the cognitive processing related to beauty, too. For one, computations in cognitive systems may also systematically link to perceived beauty. For instance, fluency is not constrained to sensory systems. The fluency of processing at different levels of information processing can also predict cognitive judgments (Oppenheimer, 2008), such as judging fluent statements to have higher truth (Reber & Schwarz, 1999) or believing fluent-named companies in the stock market to have better values (Alter & Oppenheimer, 2006). From an aesthetic perspective, Brielmann and Dayan (2022) suggested that decision-making plays a critical role in our aesthetic preferences, and applying fluency considerations to decision-making processes could be a promising avenue for characterizing computational principles on a cognitive level.

Additionally, cognitive systems may read out sensory beauty by evaluating perceptual representations in line with their overarching properties, such as how compressed and sparse they are. It has been shown that higher-level neural correlates of beauty operate on re-weighted sensory representations (Iigaya et al., 2023). It is conceivable that cognitive systems do not read out visual feature codes directly when assessing beauty in the world. They may rather read out the quality of a representation and whether coding is in line with principles that are shared across sensory inputs that have been associated with beauty in the past (Walther, 2026). Understanding what these principles are would consequently open new ways of thinking about the nature of such high-level representations.

Conclusion

In conclusion, the framework proposed here offers a new methodology for understanding the nature of visual preferences. First, the framework focuses on establishing broader principles of aesthetic processing rather than limiting the scope of the explanations to single perceptual features using a single type of content. Second, through recent advancements in ANN modelling, it offers new opportunities to overcome critical challenges in empirical aesthetic research by (i) enabling a fast testbed for computational hypotheses and (ii) distilling computational principles that can be compared with those observed in the human brain. We believe that this approach has the potential to kickstart a fruitful new line of investigations into the neural correlates of beauty, which may ultimately yield explanations for *why* we assign beauty to the world around us.

Author Contributions: M.A.E., P.F., and D.K. wrote the paper.

Competing Interest Statement: The authors declare no competing interests.

Acknowledgements

This work was supported by the Deutsche Forschungsgemeinschaft (DFG), grant KA4683/6-1 (project no. 536053998), and under Germany's Excellence Strategy (EXC 3066/1, "The Adaptive Mind", project no. 533717223). It was further supported by a European Research Council (ERC) Starting Grant (PEP, ERC-2022-STG 101076057). Views and opinions expressed are those of the authors only and do not necessarily reflect those of the funders. Neither the funders nor the granting authority can be held responsible for them.

References

Alter, A. L., & Oppenheimer, D. M. (2006). Predicting short-term stock fluctuations by using processing fluency. *Proceedings of the National Academy of Sciences*, *103*(24), 9369-9372. doi: https://doi.org/10.1073/pnas.0601071103

Bader, L. (2018). Artists versus art historians? Conflicting interpretations in the Holbein controversy. *Journal of Art Historiography*, *19* (December). doi: https://doi.org/10.11588/artdok.00008481

Bar, M., & Neta, M. (2006). Humans prefer curved visual objects. *Psychological Science*, *17*(8), 645-648. doi: https://doi.org/10.1111/j.1467-9280.2006.01759.x

Berlyne, D. E. (1971). Aesthetics and psychobiology.

Bowers, J. S., Malhotra, G., Dujmović, M., Montero, M. L., Tsvetkov, C., Biscione, V., ... & Blything, R. (2022). Deep problems with neural network models of human vision. *Behavioral and Brain Sciences*, *46*. doi: https://doi.org/10.1017/S0140525X22002813

Brielmann A.A. & Dayan P. (2022). A computational model of aesthetic value. *Psychol. Rev.* 129, 1319. doi: https://doi.org/10.1037/rev0000337

Chang, N., Pyles, J. A., Marcus, A., Gupta, A., Tarr, M. J., & Aminoff, E. M. (2019). BOLD5000, a public fMRI dataset while viewing 5000 visual images. *Scientific data*, *6*(1), 49. doi: https://doi.org/10.1038/s41597-019-0052-3

Chatterjee, A., & Vartanian, O. (2016). Neuroscience of aesthetics. *Ann. N. Y. Acad. Sci.* 1369, 172–194. doi: https://doi.org/10.1111/nyas.13035

Cichy, R. M., & Kaiser, D. (2019). Deep neural networks as scientific models. *Trends in Cognitive Sciences*, *23*(4), 305-317. doi: 10.1016/j.tics.2019.01.009

Conwell, C., Graham, D., Boccagno, C., & Vessel, E. A. (2025). The perceptual primacy of feeling: Affectless visual machines explain a majority of variance in human visually evoked affect. *Proceedings of the National Academy of Sciences*, *122*(4), e2306025121. doi: https://doi.org/10.1073/pnas.2306025121

Corradi, G., Belman, M., Currò, T., Chuquichambi, E. G., Rey, C., & Nadal, M. (2019). Aesthetic sensitivity to curvature in real objects and abstract designs. *Acta Psychol.* 197, 124–130. doi: 10.1016/j.actpsy.2019.05.012

Dan, Y., Atick, J. J., & Reid, R. C. (1996). Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *The Journal of Neuroscience*, *16*(10), 3351-3362. doi: https://doi.org/10.1523/JNEUROSCI.16-10-03351.1996

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

Dibot, N. M., Tieo, S., Mendelson, T. C., Puech, W., & Renoult, J. P. (2023). Sparsity in an artificial neural network predicts beauty: Towards a model of processing-based

aesthetics. *PLoS Computational Biology*, *19*(12), e1011703. doi: https://doi.org/10.1371/journal.pcbi.1011703

Doerig, A., Sommers, R. P., Seeliger, K., Richards, B., Ismael, J., Lindsay, G. W., ... & Kietzmann, T. C. (2023). The neuroconnectionist research programme. *Nature Reviews Neuroscience*, *24*(7), 431-450. doi: https://doi.org/10.1038/s41583-023-00705-w

Ekinci, M. A., Buhlmann, N., & Kaiser, D. (2026). Visual features explain dynamic aesthetic experiences across distinct movie content. *Communications Psychology*, *4*(1), 127. doi: https://doi.org/10.1038/s44271-026-00531-7

Fechner, G. T. (1876). *Vorschule der Ästhetik* (Vol. 1). Breitkopf & Härtel.

Fries, P. (2015). Rhythms for cognition: communication through coherence. *Neuron*, *88*(1), 220-235. doi: 10.1016/j.neuron.2015.09.034

Graf, L. K., & 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. doi: https://doi.org/10.1177/1088868315574978

Graham, D., Schwarz, B., Chatterjee, A., & Leder, H. (2016). Preference for luminance histogram regularities in natural scenes. *Vision Research*, *120*, 11-21. doi: 10.1016/j.visres.2015.03.018

Herz, N., Baror, S., & Bar, M. (2020). Overarching states of mind. *Trends in Cognitive Sciences*, *24*(3), 184-199. doi: 10.1016/j.tics.2019.12.015

Huang, Y., Xue, X., Spelke, E., Huang, L., Zheng, W., & Peng, K. (2018).The aesthetic preference for symmetry dissociates from early-emerging attention to symmetry. *Sci. Rep.* **8**, 6263. doi: https://doi.org/10.1038/s41598-018-24558-x

Iigaya, K., Yi, S., Wahle, I. A., Tanwisuth, K., & O'Doherty, J. P. (2021). Aesthetic preference for art can be predicted from a mixture of low-and high-level visual features. *Nature Human Behaviour*, *5*(6), 743-755. doi: https://doi.org/10.1038/s41562-021-01124-6

Iigaya, K., Yi, S., Wahle, I. A., Tanwisuth, S., Cross, L., & O'Doherty, J. P. (2023). Neural mechanisms underlying the hierarchical construction of perceived aesthetic value. *Nature Communications*, *14*(1), 127. doi: https://doi.org/10.1038/s41467-022-35654-y

Jin, X., Wu, L., Zhao, G., Li, X., Zhang, X., Ge, S., ... & Zhou, X. (2019). Aesthetic attributes assessment of images. In *Proceedings of the 27th ACM International Conference on Multimedia* (pp. 311-319). doi: https://doi.org/10.1145/3343031.3350970

Kabra, S., Abhyankar, N., Desai, S., Iyer, P., & Reddy, C. K. (2026). LLM-AutoSciLab: Closed-Loop Scientific Discovery via Active Experimentation with LLMs. *arXiv preprint arXiv:2605.24043*. doi: https://doi.org/10.48550/arXiv.2605.24043

Kaiser, D., Stein, T., & Peelen, M. V. (2014). Object grouping based on real-world regularities facilitates perception by reducing competitive interactions in visual cortex. *Proceedings of the National Academy of Sciences*, *111*(30), 11217-11222. doi: https://doi.org/10.1073/pnas.1400559111

Kaiser, D., & Peelen, M. V. (2018). Transformation from independent to integrative coding of multi-object arrangements in human visual cortex. *NeuroImage*, *169*, 334-341. doi: 10.1016/j.neuroimage.2017.12.065

Kietzmann, T. C., McClure, P., & Kriegeskorte, N. (2017). Deep neural networks in computational neuroscience. *BioRxiv*, 133504. doi: https://doi.org/10.1101/133504

Leder, H., Belke, B., Oeberst, A., & Augustin, D. (2004). A model of aesthetic appreciation and aesthetic judgments. *British journal of psychology*, *95*(4), 489-508. doi: https://doi.org/10.1348/0007126042369811

Lin, Y., Verlinde, L., Wagemans, J., & de Beeck, H. O. (2026). Tracing aesthetic experience from perception and conception to appraisal using deep convolutional neural networks. *Iscience*, *29*(6). doi: 10.1016/j.isci.2026.116232

Lu, X., Lin, Z., Jin, H., Yang, J., & Wang, J. Z. (2014). Rapid: Rating pictorial aesthetics using deep learning. In *Proceedings of the 22nd ACM International Conference on Multimedia* (pp. 457-466). doi: https://doi.org/10.1145/2647868.2654927

Lu, C., Lu, C., Lange, R. T., Foerster, J., Clune, J., & Ha, D. (2024). The AI scientist: Towards fully automated open-ended scientific discovery. *arXiv preprint arXiv:2408.06292*. doi: https://doi.org/10.48550/arXiv.2408.06292

McMains, S., & Kastner, S. (2011). Interactions of top-down and bottom-up mechanisms in human visual cortex. *Journal of Neuroscience*, *31*(2), 587-597. doi: https://doi.org/10.1523/JNEUROSCI.3766-10.2011

Martindale, C., Moore, K., & West, A. (1988). Relationship of preference judgments to typicality, novelty, and mere exposure. *Empirical studies of the arts*, *6*(1), 79-96. doi: https://doi.org/10.2190/MCAJ-0GQT-DJTL-LNQD

Musslick, S., Bartlett, L. K., Chandramouli, S. H., Dubova, M., Gobet, F., Griffiths, T. L., ... & Holmes, W. R. (2025). Automating the practice of science: Opportunities, challenges, and implications. *Proceedings of the National Academy of Sciences*, *122*(5), e2401238121. doi: https://doi.org/10.1073/pnas.2401238121

Nadal, M., Munar, E., Marty, G., & Cela-Conde, C. J. (2010). Visual complexity and beauty appreciation: Explaining the divergence of results. *Empirical Studies of the Arts*, *28*(2), 173-191. doi: https://doi.org/10.2190/EM.28.2.d

Nara, S., & Kaiser, D. (2024). Integrative processing in artificial and biological vision predicts the perceived beauty of natural images. *Science Advances*, *10*(9), eadi9294. doi: 10.1126/sciadv.adi9294

Oppenheimer, D. M. (2008). The secret life of fluency. *Trends in Cognitive Sciences*, *12*(6), 237-241. doi: 10.1016/j.tics.2008.02.014

Palmer, S. E., Schloss, K. B., & Sammartino, J. (2013). Visual aesthetics and human preference. *Annual review of psychology*, *64*(1), 77-107. doi: https://doi.org/10.1146/annurev-psych-120710-100504

Pouget, A., Dayan, P., & Zemel, R. (2000). Information processing with population codes. *Nature Reviews Neuroscience*, *1*(2), 125-132. doi: https://doi.org/10.1038/35039062

Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*(1), 79-87. doi: https://doi.org/10.1038/4580

Reber, R., & Schwarz, N. (1999). Effects of perceptual fluency on judgments of truth. *Consciousness and Cognition*, *8*(3), 338-342. doi: 10.1006/ccog.1999.0386

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. doi: https://doi.org/10.1207/s15327957pspr0804_3

Redies C. (2015). Combining universal beauty and cultural context in a unifying model of visual aesthetic experience. *Front. Hum. Neurosci.* 9, 218. doi: https://doi.org/10.3389/fnhum.2015.00218

Rhodes, G., Proffitt, F., Grady, J. M., & Sumich, A. (1998). Facial symmetry and the perception of beauty. *Psychon. Bull. Rev.* 5, 659–669. doi: https://doi.org/10.3758/BF03208842

Rhodes, G. (2006). The evolutionary psychology of facial beauty. *Annu. Rev. Psychol.*, *57*(1), 199-226. doi: https://doi.org/10.1146/annurev.psych.57.102904.190208

Röth, K., Shao, S., & Gjorgjieva, J. (2021). Efficient population coding depends on stimulus convergence and source of noise. *PLoS Computational Biology*, *17*(4), e1008897. doi: https://doi.org/10.1371/journal.pcbi.1008897

Schloss, K. B., & Palmer, S. E. (2011). Aesthetic response to color combinations: preference, harmony, and similarity. *Atten. Percept. Psychophys.* 73, 551–571. doi: https://doi.org/10.3758/s13414-010-0027-0

Seresinhe, C. I., Preis, T., & Moat, H. S. (2017). Using deep learning to quantify the beauty of outdoor places. *Royal Society Open Science*, *4*(7), 170170. doi: 10.1098/rsos.170170

Simonyan, K., & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*. doi: https://doi.org/10.48550/arXiv.1409.1556

Soydaner, D., & Wagemans, J. (2026). Unveiling the factors of aesthetic preferences with explainable AI. *British Journal of Psychology*, *117*(2), 444-478. doi: https://doi.org/10.1111/bjop.12707

Stecher, R., Cichy, R. M., & Kaiser, D. (2025). Decoding the rhythmic representation and communication of visual contents. *Trends in Neurosciences*, *48*(3), 178-188. doi: 10.1016/j.tins.2024.12.005

Swartz, A., Skelton, A. E., Mather, G., Bosten, J. M., Maule, J., & Franklin, A. (2024). The perceived beauty of art is not strongly calibrated to the statistical regularities of real-world scenes. *Sci. Rep.* 14, 19368. doi: https://doi.org/10.1038/s41598-024-69689-6

Tang, Y., Cunningham, W. A., & Walther, D. B. (2025). Less is more: aesthetic liking is inversely related to metabolic expense by the visual system. *PNAS Nexus* **4**, pgaf347. doi: https://doi.org/10.1093/pnasnexus/pgaf347

Tinio, P. P., & Leder, H. (2009). Just how stable are stable aesthetic features? Symmetry, complexity, and the jaws of massive familiarization. *Acta Psychologica*, *130*(3), 241-250. doi: 10.1016/j.actpsy.2009.01.001

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. doi: https://doi.org/10.1068/i0466aap

Van Geert, E., & Wagemans, J. (2020). Order, complexity, and aesthetic appreciation. *Psychol. Aesthet. Creat. Arts* 14, 135–154. doi: https://doi.org/10.1037/aca0000224

Vartanian, O., Farzanfar, D., Munar, E., Skov, M., Hayn-Leichsenring, G., Ho, P. K., & Walther, D. B. (2024). Neural dissociation between computational and perceived measures of curvature. *Sci. Rep.* 14, 26529. doi: https://doi.org/10.1038/s41598-024-76931-8

Walther, D. B. (2026). The Sweet Spot of Seeing: A Categorization-Based Account of Visual Aesthetic Pleasure. *Current Opinion in Psychology*, 102351. doi: 10.1016/j.copsyc.2026.102351

Wichmann, F. A., & Geirhos, R. (2023). Are deep neural networks adequate behavioral models of human visual perception?. *Annual review of vision science*, *9*(1), 501-524. doi: https://doi.org/10.1146/annurev-vision-120522-031739

Xu, Y., & Vaziri-Pashkam, M. (2021). Limits to visual representational correspondence between convolutional neural networks and the human brain. *Nature communications*, *12*(1), 2065. doi: https://doi.org/10.1038/s41467-021-22244-7

Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., & Torralba, A. (2017). Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, *40*(6), 1452-1464. doi: 10.1109/TPAMI.2017.2723009