

From decisions to aesthetic values in cognitive neuroscience

Pradyumna Sepúlveda

1. Italian Academy for Advanced Studies in America, Columbia University
2. Department of Psychiatry, Columbia University

In cognitive neuroscience, the study of the decision process has shown that neurons in an extensive network of brain regions can track subjective *value*, an internal representation of how good an option is. In these experiments, animals and humans learn the associations between simple stimuli, such as light or sounds, and potential rewards and punishments. These experiments, supported by mechanistic computational models, have allowed us to find these hidden value signals in the depths of the brain cortex. However, it is still an open question how the brain actually transforms real and complex stimuli into a simple subjective value metric. Even more, we know that humans can express preferences for completely novel items, which suggests that value judgments can be constructed on the fly. The value signal can be highly flexible and affected by current contexts and goals, which is crucial to adapt to ever-changing environments.

From ancient cave paintings to millions of Instagram pictures, aesthetic experience has always been a fundamental constituent of human lives. The question of how people make aesthetic judgments has been investigated in multiple disciplines for centuries. Yet, how the human brain computes aesthetic value from complex visual stimuli is not fully understood. The approach taken in cognitive neuroscience (in conjunction with novel computational methods) could open a new avenue for a *mechanistic* understanding of aesthetic judgments.

Moreover, the study of aesthetic value can be an excellent case study to address how we generate value in all our decisions. Visual artworks can be complex and highly varied, offering many possibilities regarding the potential stimulus to probe. At the same time, it is very intuitive to appreciate a painting: even if we are not experts and it is our first time seeing an artwork, we can quickly tell if it pleases our taste. Experimentally, extracting objective metrics from the raw visual information (e.g., luminosity, color) is possible, which is an advantage in the design of controlled experiments. The process of visual judgment is dynamic; our eyes move, exploring each corner of the canvas or screen, which can reveal important insights about sampling information to generate subjective reports. Furthermore, art generates a multifaceted experience, which could be affected by contextual and other internal factors, which we can access through explicit reports.

In this pre-paper, I briefly introduce some background on the study of decision-making research in neuroscience, highlighting relevant aspects for our study of aesthetic values. Then, I present a computational model that describes how visual aesthetic judgments could be constructed by integrating multiple visual information pieces.

Perceptual decisions

Even the most trivial decisions, like choosing coffee or tea, are incredibly intricate, involving obscure and complex interactions of multiple factors. For this reason, initial neuroscientific efforts to formalize a methodology to study decisions relied mainly on analyzing perceptual choices, i.e., psychophysics (Shadlen and Kiani, 2013). In simple perceptual decisions, subjects use sensory stimuli, such as the contrast of lines, shades of color, or dot motion, to make their choices. For

example, subjects may be asked to look at two circles on a screen and instructed to select the option with higher luminosity. This setup allows the experimenter a high degree of control over the evidence used in the decisions. These perceptual studies imply that a vast amount of sensory information is integrated into “internal representations” that extract relevant dimensions, e.g., a signal in the brain that tracks variations in luminosity. This representation is eventually used to make a choice. Further work has shown that this information could indeed be reflected in specific activity in the brain: the study of the visual system in the brain has been pivotal in understanding how sensory information is transformed into neural signals. For example, monkeys performing demanding visual discrimination tasks have shown that neurons in area MT/V5, a region in the posterior part of the brain associated with visual motion processing, can track decision variability (Newsome et al., 1989). In this way, neuroscientists have described a sophisticated map in which areas of the brain process visual information from basic lines and contrast to more complex representations such as objects and faces (Figure 1).




Figure 1. Basic visual processing in the brain. Left panel: Information processing starts at the retina, proceeds to the lateral geniculate nucleus (LGN) in the thalamus, and then to primary visual areas, V1, V2, V4, and inferior temporal cortex (IT). Center panel: the size of receptive fields (the area of visual space where a neuron is responsive) increases along the occipital cortex. Neurons in higher areas gradually increase their receptive field sizes, integrating information over larger and larger regions of the visual field. Right panel: Increasing levels of complexity in the processing found along the visual areas generate a hierarchy of features— for instance, V1 codes for basic features such as edges and lines. Higher-level neurons integrate information over multiple low-level neurons with smaller receptive fields and code for more complex features. Reproduced from Herzog and Clarke (2014).

Value-based decisions

On the other hand, for value-based decisions, the subjective assessment of the alternatives does not rely entirely on the sensory aspects but on internal signals. A monkey choosing juice type or humans trading in the stock market are considered to be making decisions in this category (Rangel et al., 2008). In neuroscience, insights from economics, biology, animal learning, psychology, and computer science have informed the notion of value. At the core is the idea that a *value* signal is a representation that guides choices to maximize future rewards. In other words,

to make appropriate decisions, these values should be reliable predictors of the benefits that could result from each option.

This value concept allows us to compare objects not commensurable in straight sensory or quantitative terms. For example, in an economic decision, comparing the value of vastly distinct objects such as La Gioconda, the International Space Station, and a banana is possible. This abstract nature of value has sometimes been denominated as having a “common currency” in the brain (Levy and Glimcher, 2012). To understand how these values can be learned, animal and machine learning knowledge has been integrated into the “reward hypothesis”. This proposal states that the environment provides reinforcement signals that indicate the probable costs or benefits of states and actions in the environment. For example, a dog can learn to associate sitting with getting a treat. After many repetitions, the animal learns that the “sitting” action leads to a reward, i.e., this action will have a high value. Animals can learn to act in ways that maximize their utility over the short or long term (Juechems and Summerfield, 2019). This hypothesis has been developed and formalized in computational models such as the reinforcement learning (RL) model, a pivotal computational formalization in computer science and neuroscience.

The presence of valuation systems has been generally accepted, although their exact neural implementation and characteristics are an ongoing area of research. Some standard proposals for valuation systems include the *Pavlovian system*, which assigns prominence to a restricted set of evolutionarily relevant actions, such as preparatory behaviors for food or avoiding aversive stimuli, such as heat or electric shocks. The behavioral repertoire of Pavlovian responses is rather limited and inflexible. In contrast, the *habitual valuation system* can cover a more comprehensive set of actions, assuming they are repeatedly experienced. In the habitual system, stimulus-response associations are learned, relying on past experiences through a trial-and-error training process. This makes habitual responses a slow learning process but fast (or automatic) to deploy once learning is achieved. For example, this category includes a rat learning to press a lever for liquids in response to a sound or a smoker’s desire for a cigarette after a meal. Finally, in the *goal-directed valuation system*, values depend on the association between states/actions and outcomes. Since each action can lead to multiple results, it is assumed that some mapping between actions and outcomes exists, which can be more flexible, not relying entirely on repeated training of automatic responses. For example, suppose we need to get to Columbia University, and the 1 subway line is broken. In that case, we can harness our knowledge of the New York City transit system to look for alternative routes. However, given a more detailed assessment of the possibilities, they may require a higher investment of cognitive resources. The separation between these three systems is not necessarily clean-cut since they can interact and be modulated by additional factors, such as risk, uncertainty, and time.

Although most of the research in value-based decision-making has been related to obtaining hedonic rewards, recent work has shown that the value could represent a more general and flexible construct, which aligns with a goal-directed valuation. For example, a couple of loose matches could be an inconsequential object in the context of your kitchen, but they could become your most prized possession if you are stranded on an empty island in the middle of the Pacific Ocean. Recent work has shown that the brain network tracking reward value can indeed contain information about the “usefulness” of an object in experiments with multiple goals (Castegnetti et al., 2021). Indeed, the prefrontal cortex activity (usually associated with value computation) is reshaped by goals, indicating that the mapping of the value of state/actions can be quickly repurposed depending on task demands (e.g., Frömer 2019, Sepulveda et al., 2020; De Martino and Cortese, 2023).

Computational frameworks in decision making

While we can probe the brain's anatomy, this static approach does not fully capture its dynamic flow of information. A method to understand how the brain could operate as an information-processing machine is to use mechanistic modeling of neural and behavioral processes, which has been primarily developed in computational neuroscience. David Marr, one of the founders of computational neuroscience, proposed that the brain can be understood as an information-processing system. To understand the system, we can ask questions at three distinct levels: 1) What computational problem does the system solve? 2) What computational algorithms (the steps or code to solve the problem) does the system use? 3) How does the system implement the algorithm? (Marr, 2010).




Figure 2. Illustration of Marr's three levels to understand decision-making. A person tries to choose the best snack to eat (goal, computational level). One way of doing this is by assessing different taste dimensions and picking the most satisfactory combination (algorithmic realization). In the brain, this process could be implemented through various neural interactions and circuits (physical implementation). There are many ways in which the goal can be achieved, which opens the possibility of multiple algorithms with distinct implementations in the brain. This approach also allows for the discarding of algorithms that do not have biological support.

For example, we could focus on the problem of selecting a snack to eat during breakfast (**Figure 2**). This is the computational level. To solve this problem, we must define the specific steps to fulfill the goal or the algorithmic level. For example, one possible algorithm could be: 1) identify available snacks, 2) assess benefits

(e.g., the best taste, best nutrients), 3) account for costs or drawbacks (e.g., digestive difficulties), 4) integrate all pros and cons and generate value for each item, 4) pick the option with the highest value. We could take an alternative algorithm: 1) remember what I ate yesterday, 2) pick the same snack. Both algorithms yield a decision and are possible candidates to explain why a participant makes specific choices. Mathematical models formalize the algorithms *quantitatively* and make predictions about future behavior. We can use these predictions to assess the precision of the algorithms and identify the options that are more likely to underlie the cognitive mechanisms.

Once a candidate algorithm is identified, we need to find in the brain how each step could be implemented in a biologically compatible manner, the implementation level. At this level, we try to identify brain activity that reflects each one of the steps of the algorithm, e.g., the neural activity of features of the food item. Following the example, by combining computational models and neuroimaging (O'Doherty et al., 2007), researchers have identified that information associated with nutritious attributes of food items, such as proteins or fat composition, could be found in the orbitofrontal cortex, the lower frontal part of the brain (Suzuki et al., 2016) when participants assess various snacks.

Further, functional magnetic resonance imaging (fMRI), a technique that indirectly tracks brain activity and measures changes in blood flow properties, has been important in showing relevant value processing areas in the human brain. Estimations of participants' internal values in decision tasks are reflected in the activity in brain areas, such as the ventromedial prefrontal cortex (vmPFC) (Cortese et al., 2012) (**Figure 3**). Recent studies using this approach have shown that value can be constructed by many features, as in the aforementioned work on food items (Suzuki et al., 2016) and others (Hare et al., 2011; Lim et al., 2013). One caveat of this approach is that the features to describe an object could be unlimited, easily overwhelming the ability of biological and artificial agents to process information. Further studies integrating human behavior and learning in multidimensional environments have shown that attention plays a critical role in selecting the relevant features and discarding irrelevant information, fostering the generation of efficient value representations (Niv et al., 2015; Leong et al., 2017; Cortese et al., 2021).




Figure 3. Model-based fMRI. Activity associated with the anticipated value extracted from a computational (RL) model in a decision-making task. vmPFC is an area recurrently associated with value tracking in the brain. Reproduced from Cortese et al. (2021).

What about aesthetic value?

Aesthetic valuation is difficult to categorize, with many aspects of the decision-making process converging on it. After this overview,

we hint at many theoretical and methodological insights that could help understand how humans make aesthetic judgments. The neural basis of sensory processing can define some fundamental features used to construct aesthetic experiences. From value-based decisions, prior experience with a stimulus or object can influence value judgments, making previous memories and the history of associations relevant to estimating idiosyncratic aesthetic preferences. On the other hand, value decisions are also flexible and highly adaptable, which in aesthetic judgments is reflected in the capacity to express preferences for completely novel stimuli, somehow generating “value” almost on the fly, e.g., I can instantly know if I like a painting, even if I have never seen it before. This flexibility could rely on humans having a “baseline” value mapping, an internal model of visual aesthetic preference, which can be used over novel (and familiar) objects.

Approaching this big question, previous work in our research group has focused on aesthetic judgments as an example of how the brain can transform realistically complex stimuli into a simple subjective value (Iigaya et al., 2021). In this study, a high-dimensional input, a visual artwork, is assessed by participants to generate a one-dimensional output, a rating of how much the painting is liked. The dimensionality reduction process (or how people can identify task-relevant features) is crucial for humans to generalize their knowledge to novel situations, and it has been studied in decision-making literature, usually based on simplified visual inputs (e.g., Mach et al., 2020; Cortese et al., 2021). Yet, little is known about how this process could operate in the context of aesthetic judgments. From the perspective of cognitive models, the diversity of visual art presents an interesting and challenging study case. Even in paintings alone, there is an overwhelmingly broad range of elements, themes, and styles. Therefore, artworks offer a good test for understanding the computation mechanisms the brain uses to generalize and generate values across diverse sets of stimuli.

It is important to note the field of neuroaesthetics has already used tools from cognitive neuroscience to study the aesthetic experience. For example, neuroimaging studies have shown how some brain areas increase their activity in the presence of stimuli with higher aesthetic values (Cela-Conde et al., 2004; Kawabata and Zeki, 2004). While this approach has been useful in assessing the relevance of some brain regions to aesthetic judgments, it is limited in describing how the brain *computes* aesthetic value from visual stimuli in the first place. Using computational methods can bring some additional insights into potential algorithms and implementation.

A model of aesthetic value construction

The intuition behind the model is that the brain actively constructs the value of a stimulus by integrating its simple attributes or features. These features describe some basic characteristics of a painting and are general enough to be found in any visual stimulus, i.e., they can be used across all images, including new ones. For instance, in any work of visual art, we can quantify different colors, intensities, contrasts, textures, shapes, etc. We can also consider more complex features, such as if the painting is abstract or figurative. All these features can be used to characterize a painting. Using computational methods, the model captures the relevance of each one of these features in the generation of the value preferences, which is implemented in an algorithm with three fundamental parts (**Figure 4A**):

- 1) *Feature metrics for novel images (input)*: Each “new” image to be judged is decomposed into its basic feature metrics. For example, the agent quantifies the novel painting’s global luminosity, hue, symmetry, etc. While the features are always the same, the feature content will differ depending on the paintings.
- 2) *Feature weights*: The model considers that each agent (a human or artificial observer) can have a unique set of weights that characterize each one of the features. These weights reflect an individual’s subjective judgment about how much a particular feature should count toward the overall value of a stimulus. For example, one person may emphasize symmetry in their aesthetic preference, while another may consider color contrasts more important for their subjective report. These weights are assumed to be stable for each individual during the judgments, allowing the model to define an idiosyncratic *feature space*.
- 3) *Integration of features (output)*: For the novel painting, the value of the features and their respective weights are used to compute the overall value of the image, which is reported

as an aesthetic judgment. This integration is assumed to follow a linear function (i.e., adding up the weighted metrics for all the features), although it could be implemented using non-linear functions. This stage generates the *output* of the model, the subjective value.

As you probably noted, selecting the features to be included in the model is crucial for its effectiveness. In the current implementation, two main sets of features are extracted: low-level and high-level. Low-level features can be generated by multiple methods of computer vision, which rely entirely on global characteristics of the image (e.g., color distributions, brightness effects, hue, saturation, etc.) or local features extracted from automatic segmentation. High-level features are generated based on more abstract dimensions, not easily computed by a simple algorithm such as concreteness, dynamics, or emotional valence value for each image (Chatterjee et al., 2010; Vaidya et al., 2018). In total, the model can include over 80 visual features.

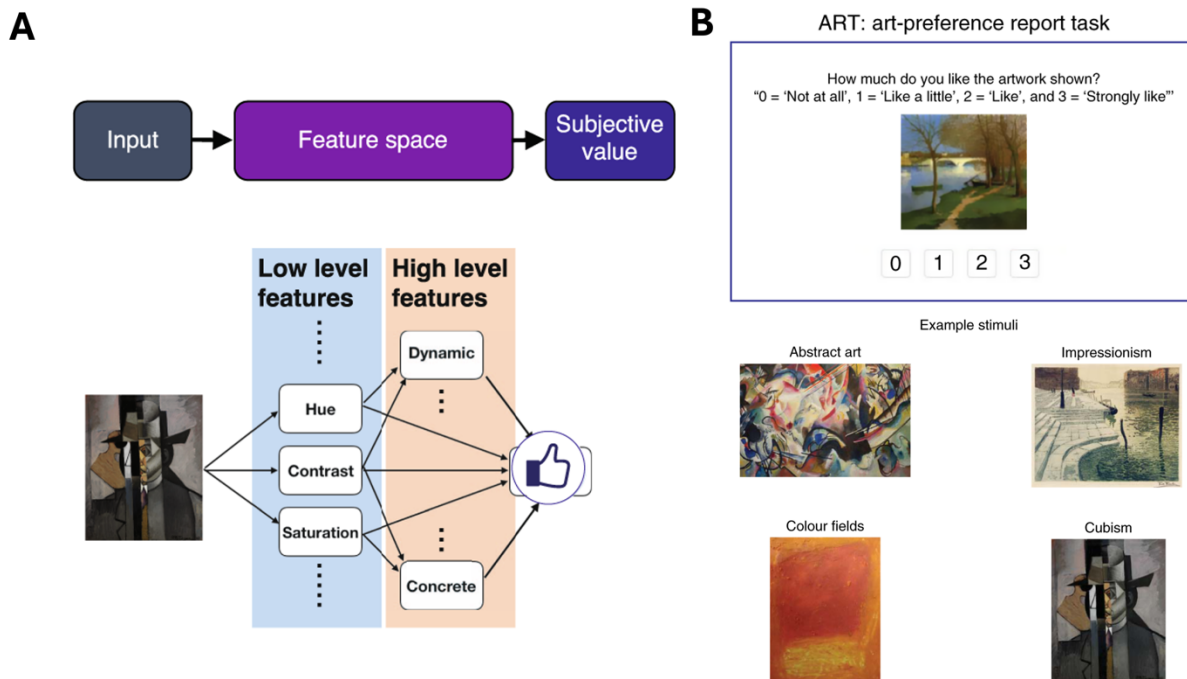


Figure 4. The model for prediction of aesthetic value. A) The idea of value construction captured by the model implies that input is projected into a feature space, in which the subjective value judgment is performed. Importantly, the feature space is shared across images, enabling this mechanism to generalize across various paintings and photos, including novel ones. In the model, a visual stimulus (e.g., artwork) is decomposed into multiple low-level features (e.g., mean hue, mean contrast) and high-level features (e.g., concreteness, dynamics). One hypothesis of the model is that high-level features can be constructed from low-level features. Subjective value is constructed from a linear combination of all features (low and high levels). **(B)** The task (ART: art-liking rating task). Participants were asked to report how much they liked a stimulus (a piece of artwork) shown on the screen using a four-point Likert rating ranging from 0 to 3. Example of some of the paintings presented to participants. Some styles presented to participants were Cubism, Impressionism, abstract art, and color fields. Reproduced from ligaya et al. (2021).

To validate the model, participants with no expertise in art reported their aesthetic ratings for hundreds of images of different styles (e.g., color field, impressionism, cubism) during in-person and online experiments (**Figure 4B**). Naturally, each participant had unique preferences regarding the type of artwork they preferred. The model can account for this variability since it is possible to find the best set of weights that allows the model to capture individual participant responses (**Figure 5**). The aesthetic value construction model generates predictions that are significantly correlated with participants' responses, even when it is used to predict the aesthetic preferences for non-painting visual stimuli, such as photography. These observations suggest that a non-negligible proportion of the variance in participants' aesthetic ratings can be captured by how they assess simple visual features (although certainly there is room to improve the model). Some features were found to be consistently relevant across participants (e.g., concreteness), while in other cases, features presented varied relevance across participants (**Figure 5**).

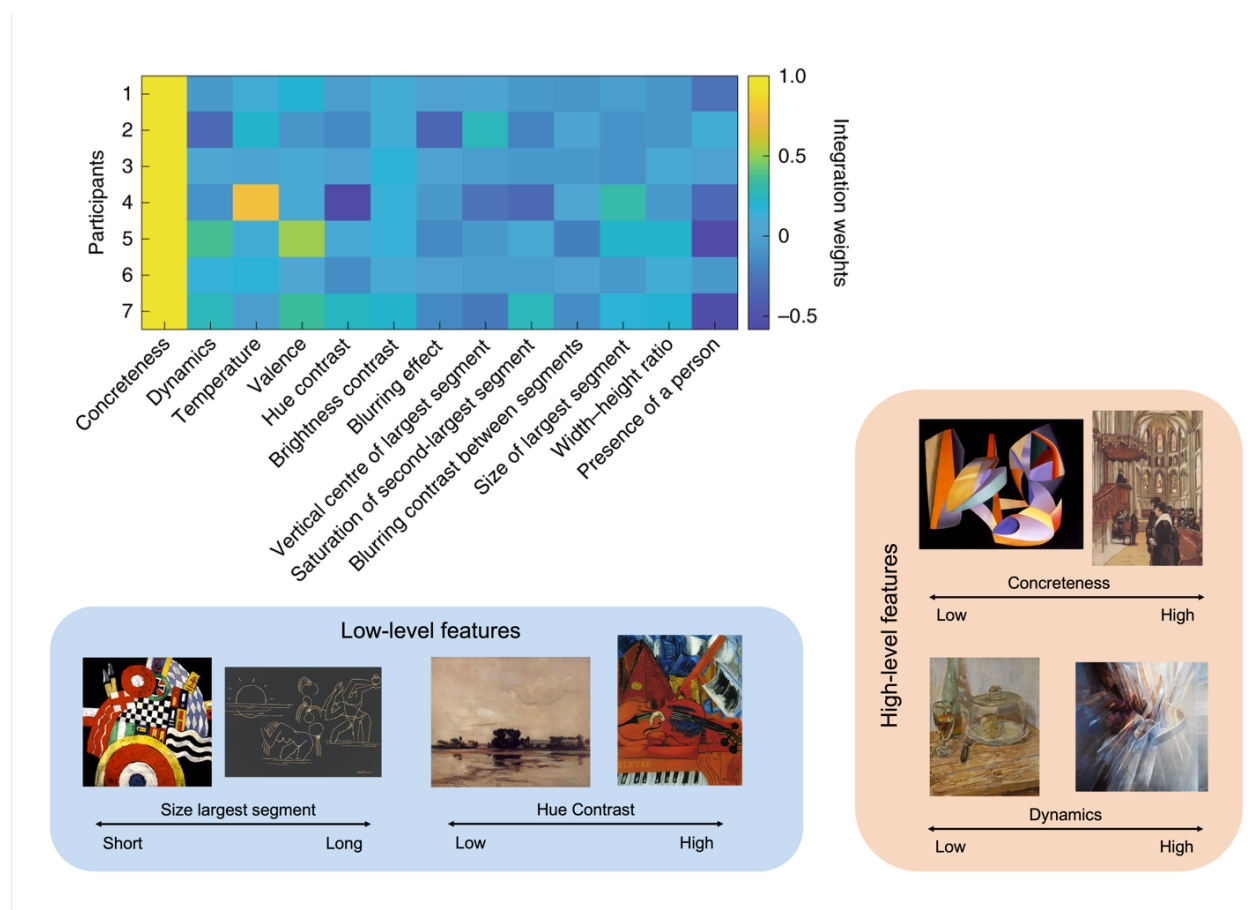


Figure 5. Value of feature weights for seven in-person participants. Weights were obtained by estimating the model for each participant. A reduced model, with some of the most relevant features, is presented here. The color bar in the figure indicates the value of feature weights for each participant. Positive weights (yellow) indicate that a higher feature metric will increase the preference for the images (e.g., higher concreteness predicts a higher subjective value for this group of participants). Negative weights (blue) indicate a reduction in subjective value rating depending on the feature. Examples of low and high-level features are presented. Selected paintings are exemplars of images representing the feature metrics. Reproduced from ligaya et al. (2021).

The model alone suggests a potential algorithm used to make aesthetic judgments. However, to elucidate the actual implementation of this process in humans, we need to include the brain in the analysis. Participants performed an equivalent artwork rating task inside the MR scanner, and a model-based fMRI analysis was used to characterize the potential neural substrates for the different steps of the model (Iigaya et al., 2023). Brain decoding shows that areas in the visual ventral stream (which encompasses part of the occipital and temporal cortex, similar to the diagram in **Figure 1**) presented a hierarchical representation of model features (**Figure 6A**). This means that primary visual areas (such as V1) contain information mainly associated with low-level features (blue bars in **Figure 6A**). On the other hand, more anterior areas down the visual stream (e.g., posterior parahippocampal cortex, PHC) represented mostly high-level features of the model. These observations align with previous decision and vision studies showing sensory features are strongly represented in these regions (e.g., Cortese et al., 2021; Castegnetti et al., 2021). Following the hierarchy, parietal and frontal areas of the brain show mixed representations of the low- and high-level features. Even more, moving forward in the brain towards the prefrontal cortex (PFC), we see stronger representations of subjective aesthetic valuation (increasing red bars in **Figure 6B**), in line with previous reports (e.g., Hare et al., 2011, Lim et al., 2013). Additional analysis shows that brain areas representing features could work together with those tracking subjective value, as observed from the activity coupling when people make aesthetic judgments.

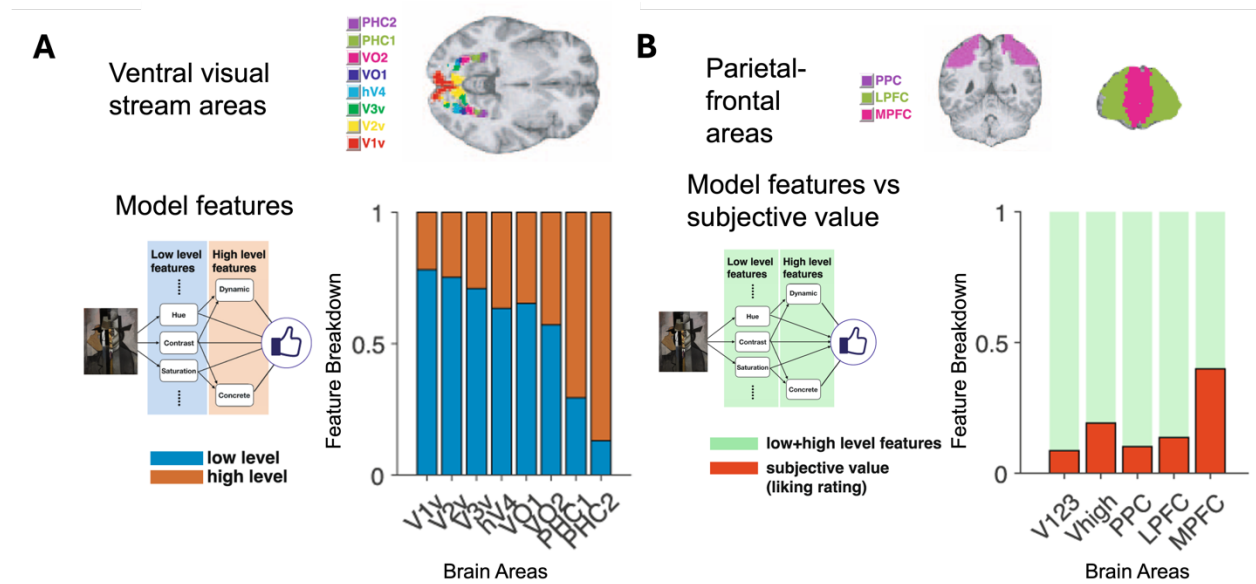


Figure 6. The visual ventral-temporal stream encodes low (blue) and high-level (orange) features in a graded hierarchical manner. Deeper areas in the ventral stream show a preferential encoding for high-level features, as expected from the visual hierarchy processing. Brain parcellation with different colors for the areas in the visual stream is presented. **B**) Encoding of low- and high-level features (green) and liking ratings (red) across brain regions in frontal areas. Frontal areas (e.g., MPFC) tend to represent the subjective value more strongly than posterior regions of the brain. Note that the ROIs for

the visual areas are now grouped as V1-2-3 (V1, V2, and V3) and V-high (Visual areas higher than V3). Brain parcellation with different colors for the posterior parietal cortex (PPC), lateral prefrontal cortex (LPFC), and medial prefrontal cortex (MPFC) is presented. Reproduced from Iigaya et al. (2023).

An essential aspect of the model is that weights are assumed to be stable at the participant level during aesthetic judgments. It is an open question if aesthetic judgments can also be flexible, as shown for value in previous value-based decision research. Alternative goals and contexts, such as where we appreciate a painting (e.g., museum vs. phone), mood, learning, and expertise, could impact the *feature space* participants use to construct their aesthetic judgments. Further, the model describes specific weights for each participant, allowing the uncovering of individual variability in how people assess artworks. How this integration of information for aesthetic judgment reflects personal traits is an open question that this framework could help to understand.

Final remarks

The objective of this overview was to connect the study of decision-making from the perspective of cognitive neuroscience with the complex question of aesthetic evaluation. I introduced a model that captures some aspects of constructing these aesthetic values. During my time at the Italian Academy, I worked to expand this framework, taking the presented computational model as a starting point to assess how aesthetic judgments could be affected by contexts and goals. Furthermore, the construction of each value judgment is highly dynamic, and how people explore the artworks and pay attention to features before making a judgment could be informative of this active process. Finally, this model allows us to understand part of the natural variation in how people make aesthetic judgments. I will explore some of these points during my presentation.

References

- Bandettini, P. A. (2012). Twenty years of functional MRI: The science and the stories. *NeuroImage*, 62(2), 575–588. <https://doi.org/10.1016/j.neuroimage.2012.04.026>
- Briellmann, A. A., & Dayan, P. (2022). A computational model of aesthetic value. *Psychological Review*, 129(6), 1319–1337. <https://doi.org/10.1037/rev0000337>
- Castegnetti, G., Zurita, M., & De Martino, B. (2021). How usefulness shapes neural representations during goal-directed behavior. *Science Advances*, 7(15), eabd5363. <https://doi.org/10.1126/sciadv.abd5363>
- Cela-Conde, C. J., Marty, G., Maestú, F., Ortiz, T., Munar, E., Fernández, A., Roca, M., Rosselló, J., & Quesney, F. (2004). Activation of the prefrontal cortex in the human visual aesthetic perception. *Proceedings of the National Academy of Sciences*, 101(16), 6321–6325. <https://doi.org/10.1073/pnas.0401427101>
- Chatterjee, A., Widick, P., Sternschein, R., Smith, W. B., & Bromberger, B. (2010). The Assessment of Art Attributes. *Empirical Studies of the Arts*, 28(2), 207–222. <https://doi.org/10.2190/EM.28.2.f>
- Cortese, A., Yamamoto, A., Hashemzadeh, M., Sepulveda, P., Kawato, M., & De Martino, B. (2021). Value signals guide abstraction during learning. *eLife*, 10, e68943. <https://doi.org/10.7554/eLife.68943>
- De Martino, B., & Cortese, A. (2023). Goals, usefulness and abstraction in value-based choice. *Trends in Cognitive Sciences*, 27(1), 65–80. <https://doi.org/10.1016/j.tics.2022.11.001>
- Frömer, R., Dean Wolf, C. K., & Shenhav, A. (2019). Goal congruency dominates reward value in accounting for behavioral and neural correlates of value-based decision-making. *Nature Communications*, 10(1), 4926. <https://doi.org/10.1038/s41467-019-12931-x>
- Gore, J. C. (2003). Principles and practice of functional MRI of the human brain. *Journal of Clinical Investigation*, 112(1), 4–9. <https://doi.org/10.1172/JCI200319010>
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics* (Vol. 1, pp. 1969-2012). New York: Wiley.
- Hare, T. A., Malmaud, J., & Rangel, A. (2011). Focusing Attention on the Health Aspects of Foods Changes Value Signals in vmPFC and Improves Dietary Choice. *Journal of Neuroscience*, 31(30), 11077–11087. <https://doi.org/10.1523/JNEUROSCI.6383-10.2011>
- Herzog, M. H., & Clarke, A. M. (2014). Why vision is not both hierarchical and feedforward. *Frontiers in computational neuroscience*, 8, 135.
- Iigaya, K., Yi, S., Wahle, I. A., Tanwisuth, K., & O'Doherty, J. P. (2020). Aesthetic preference for art emerges from a weighted integration over hierarchically structured visual features in the brain. <https://doi.org/10.1101/2020.02.09.940353>
- 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. <https://doi.org/10.1038/s41562-021-01124-6>
- O'Doherty, J. P., Hampton, A., & Kim, H. (2007). Model-based fMRI and its application to reward learning and decision making. *Annals of the New York Academy of sciences*, 1104(1), 35-53.
- 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>
- Kawabata, H., & Zeki, S. (2004). Neural Correlates of Beauty. *Journal of Neurophysiology*, 91(4), 1699–1705. <https://doi.org/10.1152/jn.00696.2003>
- Leong, Y. C., Radulescu, A., Daniel, R., DeWoskin, V., & Niv, Y. (2017). Dynamic Interaction between Reinforcement Learning and Attention in Multidimensional Environments. *Neuron*, 93(2), 451–463. <https://doi.org/10.1016/j.neuron.2016.12.040>

- Levy, D. J., & Glimcher, P. W. (2012). The root of all value: A neural common currency for choice. *Current Opinion in Neurobiology*, 22(6), 1027–1038. <https://doi.org/10.1016/j.conb.2012.06.001>
- Lim, S.-L., O'Doherty, J. P., & Rangel, A. (2013). Stimulus Value Signals in Ventromedial PFC Reflect the Integration of Attribute Value Signals Computed in Fusiform Gyrus and Posterior Superior Temporal Gyrus. *The Journal of Neuroscience*, 33(20), 8729–8741. <https://doi.org/10.1523/JNEUROSCI.4809-12.2013>
- Mack, M. L., Preston, A. R., & Love, B. C. (2020). Ventromedial prefrontal cortex compression during concept learning. *Nature Communications*, 11(1), 46. <https://doi.org/10.1038/s41467-019-13930-8>
- Marr, D. (2010). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. The MIT Press. <https://doi.org/10.7551/mitpress/9780262514620.001.0001>
- Newsome, W. T., Britten, K. H., & Movshon, J. A. (1989). Neuronal correlates of a perceptual decision. *Nature*, 341(6237), 52–54. <https://doi.org/10.1038/341052a0>
- Niv, Y., Daniel, R., Geana, A., Gershman, S. J., Leong, Y. C., Radulescu, A., & Wilson, R. C. (2015). Reinforcement Learning in Multidimensional Environments Relies on Attention Mechanisms. *The Journal of Neuroscience*, 35(21), 8145–8157. <https://doi.org/10.1523/JNEUROSCI.2978-14.2015>
- O'Doherty, J. P., Hampton, A., & Kim, H. (2007). Model-Based fMRI and Its Application to Reward Learning and Decision Making. *Annals of the New York Academy of Sciences*, 1104(1), 35–53. <https://doi.org/10.1196/annals.1390.022>
- Pagnoni, G. Z., Caroline F.; Montague, P. Read; Berns, Gregory S. (2002). Activity in human ventral striatum locked to errors of reward prediction. *Nature Neuroscience*, 5(2), 97–98. <https://doi.org/10.1038/nm802>
- Pessiglione, M. S., Ben; Flandin, Guillaume; Dolan, Raymond J.; Frith, Chris D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature*, 442(7106), 1042–1045. <https://doi.org/10.1038/nature05051>
- Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, 9(7), 545–556. <https://doi.org/10.1038/nrn2357>
- Schultz, W. (1998). Predictive Reward Signal of Dopamine Neurons. *Journal of Neurophysiology*, 80(1), 1–27. <https://doi.org/10.1152/jn.1998.80.1.1>
- Sepulveda, P., Usher, M., Davies, N., Benson, A. A., Ortoleva, P., & De Martino, B. (2020). Visual attention modulates the integration of goal-relevant evidence and not value. *eLife*, 9, e60705. <https://doi.org/10.7554/eLife.60705>
- Seymour, B. O., John P.; Dayan, Peter; Koltzenburg, Martin; Jones, Anthony K. P.; Dolan, Raymond J.; Friston, Karl J.; Frackowiak, Richard S. J. (2004). Temporal difference models describe higher-order learning in humans. *Nature*, 429(6992), 664–667. <https://doi.org/10.1038/nature02581>
- Shadlen, M. N., & Kiani, R. (2013). Decision Making as a Window on Cognition. *Neuron*, 80(3), 791–806. <https://doi.org/10.1016/j.neuron.2013.10.047>
- Penny, W. D., Friston, K. J., Ashburner, J. T., Kiebel, S. J., & Nichols, T. E. (Eds.). *Statistical Parametric Mapping*. (2007). Elsevier. <https://doi.org/10.1016/B978-0-12-372560-8.X5000-1>
- Suzuki, S., Cross, L., & O'Doherty, J. P. (2017). Elucidating the underlying components of food valuation in the human orbitofrontal cortex. *Nature Neuroscience*, 20(12), 1780–1786. <https://doi.org/10.1038/s41593-017-0008-x>
- Vaidya, A. R., Sefranek, M., & Fellows, L. K. (2018). Ventromedial Frontal Lobe Damage Alters how Specific Attributes are Weighed in Subjective Valuation. *Cerebral Cortex*, 28(11), 3857–3867. <https://doi.org/10.1093/cercor/bhx246>