RIFL (2024) SFL: 222-236
ISSN 2036-6728 DOI: 10.4396/2025SFL06

The role of concepts in the perceptual learning of aesthetic properties

Marta Benenti[1]
Universidad de Murcia
marta.benenti@um.es

Abstract This paper explores the role played by concepts in the perceptual learning of aesthetic properties. After introducing the phenomenon of perceptual learning, I suggest that some aesthetic properties count as perceptual properties we learn to detect, and that conceptual feedback involved in learning processes can account for the evaluative nature of the resulting experiences. I then examine how labeling shapes perceptual categorization in infants to suggest that, although both top-down and bottom-up explanations can account for aesthetic experiences that involve perception, the latter can shed light on a gray area of aesthetic learning that, while involving concepts, does yield fully reasons-responsive experiences.

Keywords: perceptual learning, cognitive permeation, labeling, aesthetic experience, evaluative properties

Received 05/05/2025; accepted 14/10/2025.

0. Introduzione

De gustibus non est disputandum, that is, *about taste there's no arguing*, a well-known saying goes, because it doesn't seem possible to establish a fixed, universal standard for what we should aesthetically appreciate. Yet we are constantly arguing about what deserves aesthetic appreciation. Sometimes, aesthetic arguments stemming from disagreement take place for the sake of making important collective decisions, such as those regarding artistic and cultural heritage, urban planning, or the conservation of natural ecosystems. Some other discussions touch upon everyday issues, like when we rate the latest HBO tv show, or decide what colors to wear.

In spite of the pervasiveness of aesthetic arguments, the processes that govern the acquisition and development of those aesthetic preferences we defend or counter remain largely obscure to both philosophers and psychologists. In particular, it is unclear whether these processes should be explained by perceptual and emotional mechanisms, or in terms

[1] This project has been funded by the European Union's Horizon Marie Skłodowska-Curie Action N° 10.3030/101062570.
The earliest version of this paper was originally presented and received helpful feedback at the 29th Conference of the Italian Society for the Philosophy of Language in Rome (2024). I am also grateful to the members of the ARESMUR research group at the University of Murcia for providing insightful comments on a previous version of this paper, and to two anonymous reviewers of this journal. My heartful gratitude to Mara Floris for generously sharing and discussing her work with me.

of higher-level cognitive functions such as beliefs, imagination, or memory – or else by an interaction among multiple functions. Yet, setting aside some exceptions, most experiences that we are willing to consider aesthetically significant (such as those involving artworks) engage our senses, thus pointing at a perceptual access to the appreciated properties (Levinson 2005). A graceful, balanced, elegant dance move; a forceful or unsettling musical gesture; a grotesque painting composition; an intense cinematic sequence-all seem to display properties that we grasp primarily through perception. In other, more technical words, the *phenomenal character* of those experiences that these properties feature in is perceptual.

At the same time, however, such aesthetic experiences are taken to provide epistemic grounds for evaluating the experienced object as meriting, or not, our aesthetic appreciation. Thus, while perception is paradigmatically non-evaluative, for it picks-up features of the world that do not depend on our preferences or tastes, explaining aesthetic perception requires that we introduce and account for the way in which aesthetic perceptions can not only justify aesthetic evaluations, but be themselves justified or unjustified. That is, not only we need to account for the way in which perceiving a dance performance as elegant justifies our evaluation of it as a beautiful performance, but also for the fact that we are justified or unjustified in perceiving it as elegant. One way to go is by allowing for some of our perceptual experiences to be permeable to, shaped, or modified by concepts.

As it is well known, in philosophy of mind and epistemology there is an unsettled debate about the possibility that we perceive evaluative properties, that is, properties that are capable of grounding our judgements in virtue of their evaluative nature (Bergqvist & Cowan 2018). But given the breadth of this debate, in this paper I will only consider the possible contribution of concepts to our perception of such properties. And since recent discussions in analytic philosophy of mind have reignited the discussion, the focus of this paper will be especially on the latest insights about perceptual learning and the cognitive permeability of perceptual contents.

While I am open to the idea that our perceptual experiences can, on some occasions, be shaped and modified by concepts, I contend that we need more clarity about the way in which concepts are meant to perform this permeating task. That is, on how we mobilize concepts and apply them to perceptual contents so that the resulting perceptual experiences are conceptually structured and capable as such of grounding our judgements. One way to address this issue is by focusing on how concepts help us improve our perceptual skills. In this paper, I will do so with the aim of suggesting a possible explanation of the role of concepts in our aesthetic appreciation.

Before proceeding further, a preliminary elucidation about aesthetic concepts is needed. Aesthetic concepts, that is, concepts that we use in aesthetic discourses like judgements and classifications, constitute a heterogeneous class. According to a seminal taxonomy proposed by Frank Sibley (1959), they can be purely descriptive (like *impressionist*), purely evaluative (like *beautiful*), or mixed (like *elegant*). While each type of concept might require a specific theoretical and experimental treatment, in what follows, I will concentrate on the way in which these terms are originally offered to us in language games – and so partake in our perceptual training – and assume that they do so roughly in the same way, that is, through labeling. As such, these concepts can be seen as contributing to those perceptual classifications that are constitutive of diverse aesthetic experiences including seeing a work as belonging to an art movement, perceiving it as balanced, and also as, for example, mediocre.[2]

[2] Thanks to an anonymous reviewer of this journal for urging me to clarify this premise.

I will thus begin by introducing the phenomenon of perceptual learning, thereby endorsing the hypothesis that we can learn to detect certain perceptual properties. After suggesting that some aesthetic properties can be counted among perceptual properties we can learn to detect, I zoom on the role of concepts in this process. The conceptual feedback provided by verbal labels to such learning processes, I will argue, can account for the perceptual learning of certain aesthetic properties in at least two distinct fashions. In particular, I will suggest that by helping to shape our perceptions "bottom-up", verbal labels promise to acknowledge a gray zone of aesthetic learning that, while involving conceptual feedback is not yet fully reasons-responsive.

1. Perceptual learning

The intuition that perception is subject to change is old and widespread. A rough genealogy may trace it back at least to the 3rd century, when Diogenes Laertius refers to the "trained eye" of a sculptor who sees the same statue differently than an ordinary person (1925: 161). In the 18th century, Thomas Reid spoke at length about "acquired perceptions," referring to perceptual abilities that can be gained through repeated experiences and intermodal correlations (EIP II xxi: 239). In the following century, William James wrote that an expert sommelier might be able to distinguish between the top and bottom halves of a bottle of a particular wine relying solely on taste (1898: 509), and famously John Dewey (1922) emphasized that the ability to have "distinct sensations" is a sign of training, skill, and habit. More recently, psychologist Eleanor Gibson defined perceptual learning as "Any relatively permanent and consistent change in the perception of a stimulus array, following practice or experience with this array" (Gibson, 1963: 29), and empirical evidence in neuroscience has since then confirmed that at least some changes in perceptual discrimination depend on perceptual learning (Goldstone 2003).

In line with the just mentioned examples, perceptual learning results for the most part in creating new perceptual units or – as some have put it – in *chunking reality* in different ways (Stokes 2020; Jenkin 2022). For example, after years of practice, a professional photographer will immediately detect the degree of exposure of an image. A beginner, on the other hand, will simply observe the respective darkness or brightness of the same image without detecting the correct lighting. This is a case of *differentiation* resulting from perceptual learning, that is, the acquired ability to distinguish between properties that, before training, seemed indistinguishable. Another possible outcome of perceptual learning is *unification*, that is, the acquired capacity to perceive as a single property what was previously perceived as two or more distinct properties. For instance, before learning to perform a given movement, a dancer will have to learn its different motor components, i.e. how to move an arm, how to coordinate it with legs, how to bend their back, how long to stand still. While at the beginning the dancer will perceive and execute the movements disconnected from each other, at the end of the training the movement will appear to the dancer's proprioception as unified. Similarly, when learning to taste wine, one may begin by detecting some specific gustatory indicators of that wine being a *Beaujolais*, like its being fruity and juicy, and then become capable of recognizing it as being a *Beaujolais* without perceptually separating the fruity and the juicy notes. As it has been noticed, all these cases might be accounted for as depending on *attentional weighting*. Through practice or experience, one comes to systematically pay attention to certain objects and properties and move away from others, thereby unifying or differentiating among them. Painting historians who become capable of attributing paintings to their authors by looking at the works' features acquire the capacity to *automatically* allocate their attention to certain particulars while overlooking other, less relevant ones. The automaticity of this process is crucial for distinguishing it from standard attention

allocation. That is: one can be told to look at certain features over others and learn to do this every time one sees a painting. Yet until this process is automatic it would not count as perceptual learning. Automaticity – it is said – *is* the product of the learning process (Stokes 2020).

Clearly, not all perceptual modifications count as the result of perceptual learning. Consistently with Gibson's seminal definition, the typically accepted criteria for identifying aesthetic learning and differentiating it from other related phenomena are that changes must be long-lasting as opposed to temporary, that they result from training and experiences rather than from physiological modifications, and that they are genuinely perceptual as opposed to explicable in terms of other mechanisms (Connolly 2019).

First, by requiring that a perceptual change be long-lasting in order to be considered the effect of learning, one can exclude short-term perceptual changes from the category of perceptual learning. Suppose we enter a dark room after having been in blinding sunlight. Adapting to the new lighting conditions may take some time, so our perceptual experience may undergo some changes, but this will last only for a few seconds. Although genuinely perceptual, these changes are neither long-lasting, nor do they result from any practice, and therefore would not count as the result of perceptual learning. Secondly, the requirement of training, or practice, as a necessary condition for a change to be considered a result of perceptual learning means that the effects of physical modifications such as injuries to the perceptual apparatus, cannot count as perceptual learning.

But while the first two requirements are quite uncontroversial, the latter is more debatable. Requiring the changes to be *genuinely perceptual* allows to retain the distinction between perceptual learning and post-perceptual modifications, such as changes in beliefs that leave the perceptual content unaltered.[3] Nevertheless, the debate about perceptual learning inevitably intersects with the discussion about whether perception is permeable to concepts. Examples of cognitive permeation[4] range from knowledge influencing our perception of colors (Macpherson 2012), to facial emotional expressions (Marchi & Newen 2015), to aesthetic expertise (Stokes 2014). Those who argue that perception is encapsulated and paradigmatically non-conceptual explain such examples in terms of inferential knowledge based on perception, or of allocation of attention (Pylyshyn 1999), while those who insist that concepts can enter and shape perceptual contents appeal to cognitive permeation. Although, typically, cases of cognitive permeation do not meet the long-term (or permanence) requirement for a perceptual experience to result from perceptual learning (if I see an ambiguous face as angry because of my occurrent contextual beliefs, this won't necessarily last in time nor be true of further occurrences of similar experiences), instances of potentially overlapping phenomena have been discussed (Arstila 2016, Ransom 2020). While I will not take side here on the debate about cognitive permeation (see Block 2023, Ch. 9-10 for recent extensive discussion), I shall maintain the distinction between it and perceptual learning. Whereas the former applies, paradigmatically, to cases in which concepts keep driving our perceptual exploration of stimuli, the latter applies, paradigmatically, to perceptual outcomes that do not recruit concepts "online".

[3] A purely phenomenological perspective that appeals to introspection to establish whether one's recognitional capacity is the result of perceptual learning or rather the application of updated beliefs to unchanged perceptual contents, is doomed to uncertainty. In such cases, empirical tests can determine the extent to which first-person changes align with significant neural modifications in perception.

[4] I follow Becko Copenhaver (personal conversation), Madeleine Ransom (2020) and Kevin Connolly (2019) in using 'cognitive permeation' instead of the still widespread but slightly disturbing 'cognitive penetration'.

RIFL (2024) SFL: 222-236
ISSN 2036-6728 DOI: 10.4396/2025SFL06

2. Perceptual training without and with concepts

Now that I have illustrated some of the possible results of perceptual learning and the general criteria for distinguishing it from other phenomena, it remains to describe the training processes that make it possible.

It is commonly agreed that perceptual learning requires repeated exposure to perceptual stimuli. In the domain of ecological psychology, which claims a direct descent from the Gibsonian framework, it is maintained that perceptual learning is an improvement in the capacity to detect perceptual information given the agent's purpose. This view is supported by the assumption that we perceive our environment directly, without any conceptual mediation. Direct perception implies that the perceptual information that reach the organism are already rich, rather than in need for enrichment by means of other cognitive sources, so what the perceiver can do to improve their performance is become better at selecting those pieces of information that better serve their goals (Raja 2018). To this aim, repeated interaction with the stimuli of an informational space is required – and sufficient. The organism receives perceptual (and, in particular, sensory-motor) feedback that creates a resonance between the perceiver and the perceptual environment. In a progressive process of adjustment, the perceiving organism improves its selection skills by weighing the perceptual variables to which one has to pay attention in order to obtain their goal. Thus, consistently with the generally accepted features of perceptual learning, also in the ecological framework much of the perceptual training depends on the allocation of attention.

Outside the ecological framework, perceptual learning without conceptual integration has been extensively explored appealing to the so-called "mere exposure effect" (Zajonc 1968). In a series of experiments, psychologist Robert Zajonc presented participants with various stimuli – including photographs, Chinese characters, and meaningless words – via screen projection at varying frequencies and then asked them to rate the stimuli using a liking scale ranging from good to bad. It turned out that the more the participants were made (passively) familiar with the stimulus, the more they liked it. According to a later meta-analysis of experiments testing exposure in several domains (language, food, social and political attitudes), effects of mere exposure are robust and reliable (Bornstein 1989). In the domain of aesthetics, exposure effects have been tested for musical and for abstract and figurative painting preferences (See Meskin et al. 2013 and Nanay 2017 about the theoretical implications of mere exposure effects in aesthetics).

Notably, however, both assessments of perceptual learning in terms of purely perceptual training are open to the possibility of cognitive enrichments, modifications, and integrations. Although in ecological psychology the focus is on training processes that do not require any conceptual intervention, but rather draw on the direct relationship between the organism and the environment, there is an ongoing discussion addressing the role of *norms* in perceptual learning. Understood as criteria of correctness that enable evaluation of our own or others' behaviors, such norms do not present themselves as material, perceivable stimuli in the environment, but are rather conceived as *semantic abilities* (Zipoli-Caiani 2022) or as resulting from socially educated intentions and expectations (Segundo-Ortín 2024). As to the mere exposure effect, further studies have shown that the phenomenon also applies to stimuli that, although novel, comply with some internalized system of *rules* (Gordon & Holyoak 1983; Manza & Bornstein 1995). In addition, the recent reassessment of the phenomenon of mere exposure in the predictive processing framework highlighted the compatibility of learning by mere exposure with forms of inferential learning (Van de Cruys et al. 2022).[5] More in general,

[5] I am grateful to an anonymous reviewer of this journal for bringing this study to my attention.

it has been argued that perceptual "expert performance" might require a combination of repeated and systematic exposure to stimuli accompanied by conceptual categorization which functions as feedback and as a guide for further perceptual experiences (Stokes 2020).

Several experimental examples back this claim. Birdwatchers, for instance, become experts in recognizing birds if they receive systematic feedback about the categories the detected animals belong to (Tanaka and Taylor 1991). Radiologists, who are trained in diagnostics, seem to perform better than x-ray technologists in perceiving the relevant saliences in radiographic images, although the latter are more exposed to the images (Nodine et al., 1999). Chess masters can reconstruct chessboards after viewing them for five seconds thanks to their capacity to unitize the pieces based on the moves they afford, that is, the conceptual rules of the game (Jenkin 2022). Therefore, without denying that the capacity for perceptual detection can be improved via exposure and interactions that do not appeal to concepts and, more broadly, higher level cognitive interventions (e.g. from imagination, memory, desires), we can legitimately focus on training processes that depend on conceptual integrations.[6]

One particularly interesting outcome of perceptual learning that rely on conceptual training is that it allows to treat the resulting perceptual experiences as "reasons-responsive" (Jenkin 2022). While mere perception is hardly conceived as responding to reasons (no matter how good reasons we have to think that the two lines of the Hering illusion are straight, our perception will keep presenting them as curve), mental states like beliefs typically are. Thus, if concepts constituting our beliefs play a role in our related perceptual experiences, then such experiences could turn out to be, at least partly capable of responding to reasons, that is, justifiable. For instance, chess players may argue in favor (or against) their visual experience of the chessboard as affording a given move, based on the conceptual rules that have governed their learning processes. As such, perceptual learning integrated by conceptual feedback can be seen as one mechanism that makes our perceptual experiences reasons-responsive.

If we apply to the aesthetic case what we have seen works in other instances of concept-supported perceptual learning, we can hypothesize that the ability to apply aesthetic concepts to perceivable properties instantiated in our aesthetic experiences is part of the training that enhances our ability to perceptually recognize aesthetic qualities. At least some aesthetic properties are perceivable properties, so that the perceptual training that results in the detection of such properties likely makes use of aesthetic concepts. Think, for example, of art historians who learn to recognize at a glance stylistic features that allow them to attribute artworks to periods, artistic movements, or to alleged authors or to musicians who learn to produce and reproduce sound patterns that meet the criteria not only of a correct performance, but above all of a performance that is balanced, elegant, passionate, engaging, virtuosic, and so on.

As mentioned at the beginning, perceivable aesthetic properties distinguish themselves from standard perceptual properties in virtue of their evaluative nature. Perceiving something aesthetically amounts, among other things, to evaluating it along some kind of evaluative scale. And if we are not prone to reduce evaluation to purely hedonic criteria, then we may want to conceive of such a scale as somehow *responding to reasons*, i.e. reasons-

[6] Other, more explored ways of addressing this and related issues is by focusing on how we detect or learn to detect gestalt (Landers 2021) or grouping properties (Voltolini 2023) in aesthetic experiences. Although I am fully sympathetic with these approaches to aesthetic perception, I will limit the scope of my contribution to the much narrower yet not uninteresting role of *verbal labels* in the involved learning processes.

responsive.[7] Evaluative aesthetic properties can accordingly be conceived as justifiable and capable of providing us with reasons to justify our judgements. Accounting for the way in which aesthetic concepts take part in perceptual training processes should help explain how perceivable aesthetic properties can be reason-responsive and thus ground our aesthetic judgments.

3. Conceptual feedback as *labeling*

How do the concepts with which categorize the objects and properties of our aesthetic experiences interact with those properties during processes of training and learning? As the above-mentioned examples of perceptual training show, perceptual learning results in an increased capacity to categorize perceptual stimuli such as objects or their features – and to do so perceptually. The conceptual intervention in these processes amounts to *categorization*, which is the capacity to use concepts to produce judgements about the membership of an item in a category (Machery 2009).[8] For conceptual categorization to contribute to perceptual training – rather that consisting in a completely separate process which is at best grounded in perception – it should provide one with a modified perceptual experience.

Given this broad conception, it is admittedly difficult to assess the impact of concepts on perceptual categorization. In the psychological literature there is an ongoing debate about how *linguistic labels* contribute to our capacity to categorize what we perceive, especially as newborns and young children (see Floris 2022 for a recent review). The advantage of narrowing the scope of the research to this specific mode of conceptual intervention in our experiences is that it allows us to focus on a concrete praxis, the language game of labeling, that is pervasive in our linguistic practices, particularly in our relationship to infants. While not capturing the totality of conceptual interventions on perception, verbal labeling – particularly in aesthetic discourses – seems a good entry point to a problem that is clearly larger and more complex. Therefore, in what follows, I will briefly introduce the two broad theoretical trends that can be identified within this research domain, namely the so-called *Top-down* and *Bottom-up* approaches to labeling, to see whether and how they promise to explain the way conceptual feedback facilitates perceptual learning in aesthetic experiences.

According to the first family of theories, categorization processes involve the use of labels that direct and guide perception by virtue of their referential nature-that is, their capacity to refer to certain features of objects. Psychologist Gary Lupyan has introduced the *Label*

[7] Here I remain neutral over the nature of aesthetic reasons. Accordingly, I take the notion of "reason-responsiveness" as broad enough to include justifications in terms of true beliefs about perceptual properties and justifications in terms of persuasive reasons for experiencing a given object or feature under a certain aspect (see Alcaraz León 2008 for clear analysis of justification in aesthetic discourse).

[8] Talking about *concepts* is admittedly vague. Not only philosophy and psychology tend to conceive of them in different ways, but also within philosophy there is no unitary conception of what concepts amount to. On the widespread assumption that concepts are the building blocks of thoughts, the class of concepts includes, at least, *prototypes, exemplars,* and *theories*. Each kind of concept has a different structure that allows categorization, drawing inductions, understanding the meaning of the words, and so on (Macherie 2009). It is therefore difficult to account for the conceptual feedback in perceptual learning regardless of the sort of concept that is involved in such a process. For the sake of this contribution, however, I will leverage the "heterogeneity hypothesis" according to which a given referent (a category, a substance, or kind of events) is referred to by several distinct coreferential concepts belonging to the just mentioned fundamental kinds (Macherie 2009). According to this proposal, the concept of, say, "*pointilliste*" encompasses prototypical instances, exemplars and theories that co-refer to pointilliste works. The same, I suspect, could be said for aesthetic evaluative and mixed concepts such as *beauty*, *balance*, *elegance*, *vividness*, and so on. On this background, "labels" must be seen as verbal applications of concepts, regardless of the fundamental kind they could belong to.

Feedback Hypothesis (Lupyan 2012) claiming that our labels can manipulate perceptions "online," that is, they can intervene directly in the perceptual process itself. More specifically, naming objects makes salient those perceptual features that are diagnostic of the named category (Floris 2022: 46), thereby resulting in new ways of perceiving the encountered objects. Such intervention in the perceptual process would be so rapid and automatic that it results in a perceptual categorization. This top-down proposal appears to capture an intuition that is particularly strong in the aesthetic domain. Philosopher Malcolm Budd seems to suggest just this when he invites us to imagine a situation in which we experience a work of art but struggle to find the right word to describe it, that is, the appropriate label. Then someone offers us a term that suddenly seems to perfectly match the character of the artwork. Once we (or someone else) have found the word that applies to that work or one of its aspects, our experience becomes one of perceptually recognizing a quality that fits that concept:

> When she now looks at, reads, or listens to the work, regarding this characterization as being well-suited to convey the work's character, without anything else needing to happen, she thereby experiences the work as having that character. Her experience of the work has changed: previously it was inchoate, the character being obscure; now it is distinct, the character apparent (Budd 2008: 150).

Top-down labeling can thus be seen as one way in which linguistic categorization provides us with conceptual feedback that influences our perceptual experiences of artworks and, more broadly, of objects of aesthetic appreciation. The perceptual detection of features can, at times, be integrated and improved by labeling, that is, by applying names to those features that we perceive yet in a somehow unsatisfying fashion. According to Lupyan, the ability of concepts to modify perceptions is explained by the phenomenon of cognitive permeability. Verbal labeling permeates the contents of our perception manipulating them while we are perceiving the objects' features (Lupyan 2015).

An alternative way to account for the conceptual intervention in perceptual learning processes is the *bottom-up* approach, which understands the labels we use to categorize objects and features on par with other perceptual saliences and regularities, such as colors, surface orientations, rhythm, and so on (Sloutsky & Fisher 2012; Plunkett et al., 2008). Defenders of this stance argue that labeling functions similarly to the perceptual constants one is presented with in repeated perceptual experiences. By being exposed to perceptual saliences of objects *and* to verbal labels, we improve our capacity to perceptually categorize (i.e. perceptually grouping or distinguishing) items.

From this perspective, labeling is neither more central nor more effective than other saliences in categorization processes, and above all, it does not act through the referential power of concepts, but rather contributes to these processes in a purely *quantitative* way. In other words, the tendency to see certain labels associated with certain perceptual stimuli increases the likelihood that we will categorize further occurrences of those saliences with those labels. For instance, some experimental findings attest that, if asked to assess the similarities between an item and a target, the fact that the two share the same label does not result in an "all-or-nothing" categorization as one may expect if labels worked referentially, but rather contributes to categorization like other, genuinely perceptual similarities in a way that can be statistically measured (Sloutsky & Lo 1999; Sloutsky et al. 2001). Moreover, it has been shown that labels contribute to category formation in early infancy before children understand their referential use, and thus in a purely causal manner (Plunkett et al. 2008).

The bottom-up approach allows us to treat the concepts we use to linguistically classify perceptual properties as perceptual components of the learning processes. Accordingly,

labels help refine our recognitional abilities through processes of repeated exposure rather than through their ability to refer to something. As such, conceptual feedback participates in perceptual learning in a way that does not force us to appeal to cognitive permeation, yet in principle could account for the improvement of perceptual abilities. Notably, this conception of conceptual feedback is consistent with the idea that someone could remain "ignorant of any relevant theory, term, or sophisticated concept concerning the art form", but nonetheless acquire:

> a capacity to discriminate graceful from non-graceful performances. In this way, they would have acquired the ability to see certain actions as graceful, organized in the typical way of 'being graceful,' but without possessing the relevant terms. This would be a rich perceptual experience of an aesthetic property, enabled by perceptual learning (Stokes 2018: 49).

A scenario in which we only deal with artworks and artistic performances in a purely perceptual way, without applying any labels is implausible. However, such applications are not always referential, in the sense that they can accompany our experiences of artworks without necessarily describe – or being understood as describing – them. The bottom-up approach would accommodate labeling practices with non-referential uses of concepts – whose format would instead clash with the apparently purely perceptual nature of the experience. This view is tempting, especially because it can nicely be integrated with perceptual learning without resorting to cognitive permeation. And still, to argue this is to say that, ultimately, there is nothing specific about concept-guided perceptual learning that distinguishes it from perceptual learning guided by repeated exposure to other relevant perceptual saliences. If this is the case, then one may question the capacity of bottom-up theories of labeling to account for the reason-responsiveness of the resulting perceptual experiences. In other words: how do bottom-up labelling shape our perceptual experiences in such a way that makes them justifiable on the ground of their partly conceptual nature, if their contribution to perceptual learning is non-referential?

One way to answer this question is to consider bottom-up and top-down views of labelling as being compatible with one another. Although these proposals are typically presented as alternatives, they should not be conceived as mutually exclusive. In particular, seen from a developmental perspective, it is plausible that verbal labels we use to categorize initially partake in perceptual learning in a bottom-up fashion, and later assume a distinct "supervisory" role, more consistent with top-down views (Floris 2022). In addition, in adults, the two roles could both survive. If this is possible, then the richness and, according to some, the primitiveness of aesthetic phenomena and the learning practices associated with them allows us, at least at a purely speculative level, to have our cake and eat it too: to explain certain learning processes using the Label Feedback Hypothesis, and others by relying on a bottom-up approach.

4. Outline of an aesthetic proposal

So far, I have introduced two distinct yet interconnected psychological mechanisms, namely perceptual learning and conceptual feedback in categorization processes. By focusing on the effects of labeling on perception, I have suggested that we might account for perceptual expertise in a way that acknowledges the role of concepts in perceptual training while securing the perceptual output of the learning process. In this conclusive paragraph I would like to outline a composite explanation of the perceptual learning of aesthetic properties, one that draws on both the top-down and the bottom-up approaches

to labeling, in order to account for possible interactions between perceptual contents and aesthetic concepts.

Elaborating on the question whether it is necessary that someone possess musical concepts for them to experience a musical work with understanding, Malcolm Budd argues that, on the one hand,

> neither the lack of a certain concept of a particular phenomenon nor the inability to recognize instances of the phenomenon as falling under the concept prevents a person from being sensitive to the presence of the phenomenon in a work of art and alive to the aesthetic or artistic function of the phenomenon in the work (Budd 2008: 139)

This is, in Budd's view, explained by the distinction between possessing a concept and being able to apply it on the basis of perceptual saliences. As he notices, we do not need to have the concept of a dominant seventh chord in order to "have a full sensitivity" to its "harmonic implications" in a given work (*ibidem*). We might, on the other hand, learn that a dominant seventh chord is "a major triad plus the note which forms with the root of the triad the interval of a minor seventh" but at the same time be unable to apply the concept to the relevant auditory phenomenon. Nonetheless, he continues, the mastery of musical concept is essential "to understand and engage in the practice of musical analysis and criticism" (Budd 2008: 141), that is, to adequately explain the relationship holding among the various musical phenomena that constitute a given work.

Thus, to fully account for the role of concepts in aesthetic perception, it seems that we must resort to both psychological mechanisms that I have just described. Top-down views of conceptual feedback, and especially the Label Feedback Hypothesis, are well suited to explain those learning processes that involve the ability to recruit concepts and orient perceptual experiences in the light of those concepts. This ability is well exemplified by the phenomenon known in the philosophical literature as *seeing as* (Wittgenstein 1953; Gombrich 1960). We see a certain pattern as a representation, hear a musical passage as a lament, see a painting as an expression of inner states. Learning to recognize something in light of a concept involves the recruitment and application of aesthetic concepts to perceptual profiles, which allows us to recognize them as belonging to a certain aesthetic category-such as depictions of rabbits, emotional expressions, dynamic objects, unsettling, grotesque, or elegant forms. In such cases, it is hard to claim that the apparent change in our perception results from a training that permanently alters its content, that is, we cannot explain them in terms of perceptual learning. Instead, what we seem to learn and perhaps improve at, is applying and reapplying certain labels in such a way that they guide our perceptual experience.

Whether these conceptual applications permeate the contents of our perceptual experiences "online" or rather leave them unchanged is debatable. Those sympathetic with cognitive permeation may want to insist that aesthetic concepts change our perception online, while those who are skeptical towards the existence of this phenomenon will be happy with a less ambitious explanation of seeing-as experiences. My intuition is that, even in this case, we should better answer this question on a case-by-case basis: while once we are told how to look at bi-stable figures to see what they depict, we have hard times in changing our perspective on them, we may well manage to shift from hearing a funeral march as solemn to hear it as grotesque – if we are given reasons for that. So, while we may account for the former case in terms of cognitive permeation (e.g. Voltolini 2013), the latter case could be better explained in terms of change of perspective towards an untouched perceptual content.

If cases of perceptual uncertainty or ambiguity in aesthetic experiences and their responsiveness to reasons can be explained in terms of top-down cognitive feedback on perception, they hardly count as perceptual learning. This is not problematic per se. However, there seems to be a relevant portion of aesthetic learning that, although implies categorization and is accompanied by language, does not fully rely on its referential or properly semantic function. I am thinking, on one hand, of the learning processes in very young children, and on the other, of the gradual formation of aesthetic habits.

It seems reasonable to suggest that children develop the ability to recognize aesthetic features through conceptual feedback which, at least at first, does not work because of its referential function – which young children might not be able to grasp yet – but because aesthetic labels are presented regularly in association with certain sensory stimuli. Some musical motifs will thus be perceived and gradually recognized as cheerful, and others as agitated, partly because of those labels (and likely gestures and behaviors) that systematically accompany those specific sounds. Consistently with experimental results of bottom-up approaches, the conceptual feedback in these cases might partake in the learning process that, however, won't be determined by the way in which labels refer to perceptual patterns, but rather by repeated exposure to those stimuli enhanced by labeling. Consider, furthermore, exposure phenomena that shape our aesthetic preferences as adults – and the impact they have on our individual and collective lives, from the formation of trends in pop culture and social media, to effective advertising. A theory of conceptual feedback that treats labels as properties of stimuli seems to provide an explanation both for the low responsiveness to reasons in such phenomena, and for their pervasiveness. As to the former, if verbal labels sometimes function as perceptual-like stimuli in perceptual learning (rather than in virtue of their referential power), then their capacity to respond to reasons in argumentative practices will be deflated, so that we won't be able to explain why this season we are particularly inclined to wear cerulean[9]. Although we may be able to verbally articulate that shade of blue, compare it with others, describe its relationships with different colors, our preference will turn out as something similar to a matter of fact, hard to justify. Still, it won't count purely as the result of mere exposure[10] fully deprived of conceptual interventions, because our encounter with instances of cerulean will have been mediated by accompanying descriptions (advertising, chatter, media commentaries…). And, I submit, being this way of encountering features and objects of aesthetic evaluation so widespread, bottom-up approaches promise to account for a large part of our aesthetic learning, one that is *almost* inchoate, *almost* non-conceptual, yet already available for further conceptually refined categorization and evaluation.

5. Conclusions

The goal of this contribution was to offer a plausible account of how concepts affect processes of perceptual learning in aesthetic experiences. After suggesting that we learn to detect at least some aesthetic properties via perceptual training, I tried to answer the question of how conceptual feedback plays a role in such training processes. To do so, I appealed to the debate about *labeling* and presented two ways in which labels have an impact on perceptual categorization, one top-down and the other bottom-up. In the final paragraph I outlined a compatibilist way to apply these views to aesthetic learning, such

[9] Those who have watched David Frankel' 2006 *The devil wears Prada* should be familiar with the example.
[10] Cutting (2003) suggests that mere exposure effects can explain the formation and promotion of artistic canons. However, the role of concepts is not the focus of his experimental research.

that preserves both the perceptual nature of most aesthetic experiences and their expected responsiveness to reasons.

Everything said so far suggests that the mentioned approaches in psychology and in philosophy of mind promise to clarify certain aspects of a way more complex phenomenon. Applying them to aesthetics requires further work though. If conceptual feedback is based on labels' capacity to refer to certain properties, then we must ask how exactly aesthetic labels refer to aesthetic properties. Answering that question requires, for example, assessing the metaphorical nature of most aesthetic attributions. If instead conceptual feedback is just one of many ways in which we categorize perceptual stimuli, then the empirical question we must ask concerns the statistically measurable ways in which aesthetic labels accompany and impact our perceptual experiences in the training practices we actually undergo. In this case, the evaluative component of aesthetic attributions must be weighed accordingly.

References

Alcaraz León, María José (2008), «The Rational Justification of Aesthetic Judgments» in *Journal of Aesthetics and Art Criticism*, vol. 66, n. 3, pp. 291-300.

Arstila, Valtteri (2016), «Perceptual Learning Explains Two Candidates for Cognitive Penetration» in *Erkenntnis*, vol. 81, n. 6, pp. 1151-1172.

Bergqvist, Anna, Cowan, Robert (2018) (eds.), *Evaluative Perception*, Oxford, Oxford University Press.

Block, Ned (2023), *The Border Between Seeing and Thinking*, New York, Oxford University Press.

Bornstein, Robert (1989), «Exposure and Affect: Overview and Meta-analysis of Research, 1968-1987» in *Psychological Bulletin*, vol. 106, pp. 265-289.

Budd, Malcolm (2008), *Aesthetic Essays*, Oxford, Oxford University Press.

Connolly, Kevin (2019), *Perceptual Learning*, Oxford, Oxford University Press.

Cutting, James (2003), «Gustave Caillebotte, French Impressionism, and Mere Exposure» in *Psychonomic Bulletin & Review*, vol. 10, pp. 319-343. https://doi.org/10.3758/BF03196493.

Dewey, John (1922), *Human Nature and Conduct: An Introduction to Social Psychology*, New York, Henry Holt.

Floris, Mara (2022), *Language and Categorisation in Infancy*, Rome, Aracne.

Gibson, Eleanor J. (1963), «Perceptual Learning» in *Annual Review of Psychology*, vol. 14.

Goldstone, Robert (2003), «Learning to Perceive While Perceiving to Learn», in Kimchi, Ruth, Behrmann, Marlene & Olson, Carl R. (eds.) (2003), *Perceptual Organization in Vision: Behavioral and Neural Perspectives*, Mahwah (NJ), Lawrence Erlbaum Associates, pp. 233-278.

Gombrich, Ernst H. (1960), *Art and Illusion*, London, Phaidon.

Gordon, Peter C., Holyoak, Keith J. (1983), «Implicit Learning and Generalization of the "Mere Exposure" Effect» in *Journal of Personality and Social Psychology*, vol. 45, n. 3, pp. 492-500. https://doi.org/10.1037/0022-3514.45.3.492

James, William (1890/1950), *Principles of Psychology*, New York, Dover Publications.

Jenkin, Zoe (2022), «Perceptual Learning and Reasons-Responsiveness» in *Noûs*, pp. 1-28. https://doi.org/10.1111/nous.12425

Laertius, Diogenes (1925), *Lives of Eminent Philosophers*, Vol. 2, translated by R. D. Hicks, Cambridge, Harvard University Press (Loeb Classical Library).

Landers, Casey (2021), «Specialized Visual Experiences» in *Philosophical Quarterly*, vol. 71, n. 1, pp. 74-98.

Levinson, Jerrold (2005), «What Are Aesthetic Properties?» in *Aristotelian Society Supplementary Volume*, vol. 79, pp. 191-227.

Lupyan, Gary (2012), «Linguistically Modulated Perception and Cognition: The Label-Feedback Hypothesis» in *Frontiers in Psychology*, vol. 3, n. 54, pp. 1-13.

Lupyan, Gary (2015a), «Cognitive Penetrability of Perception in the Age of Prediction: Predictive Systems Are Penetrable Systems» in *Review of Philosophy and Psychology*, vol. 6, n. 4, pp. 547-569.

Machery, Edouard (2009), *Doing Without Concepts*, New York, Oxford University Press.

Macpherson, Fiona (2012), «Cognitive Penetration of Colour Experience: Rethinking the Issue in Light of an Indirect Mechanism» in *Philosophy and Phenomenological Research*, vol. 84, pp. 24-62.

Manza, Louis, Bornstein, Robert F. (1995), «Affective Discrimination and the Implicit Learning Process» in *Consciousness and Cognition: An International Journal*, vol. 4, n. 4, pp. 399-409. https://doi.org/10.1006/ccog.1995.1047.

Marchi, Filippo, Newen, Albert (2015), «Cognitive Penetrability and Emotion Recognition in Human Facial Expression» in *Frontiers in Psychology*, vol. 6, n. 828, https://doi.org/10.3389/fpsyg.2015.00828.

Meskin, Aaron, Phelan, Mark, Moore, Michael, Kieran, Matthew (2013), «Mere Exposure to Bad Art» in *British Journal of Aesthetics*, vol. 53, n. 2, pp. 139-164.

Nanay, Bence (2017), «Perceptual Learning, the Mere Exposure Effect and Aesthetic Antirealism» in *Leonardo*, vol. 50, n. 1, pp. 58-63.

Nodine, Calvin F., Kundel, Harold L., Mello-Thoms, Claudia, Weinstein, Susan P., Orel, Suzanne G., Sullivan, Diane C., Conant, Elizabeth F. (1999), «How Experience and Training Influence Mammography Expertise» in *Academic Radiology*, vol. 6, n. 10, pp. 575-585, https://doi.org/10.1016/s1076-6332(99)80252-9.

Plunkett, Kim, Hu, Jianfeng, Cohen, Leslie B. (2008), «Labels Can Override Perceptual Categories in Early Infancy» in *Cognition*, vol. 106, n. 2, pp. 665-681.

Pylyshyn, Zenon (1999), «Is Vision Continuous with Cognition? The Case for Cognitive Impenetrability of Visual Perception» in *Behavioral and Brain Sciences*, vol. 22, pp. 341-365.

Raja, Vicente (2018), «A Theory of Resonance: Towards an Ecological Cognitive Architecture» in *Minds and Machines*, vol. 28, n. 1, pp. 29-51.

Ransom, Madeleine (2020), «Expert Knowledge by Perception» in *Philosophy*, vol. 95, n. 3, pp. 309-335.

Reid, Thomas (1764/1997), *An Inquiry into the Human Mind on the Principles of Common Sense*, Edinburgh, Edinburgh University Press.

Sibley, Frank (1959), «Aesthetic Concepts» in *The Philosophical Review*, vol. 68, n. 4, pp. 421-450, https://doi.org/10.2307/2182490.

Segundo-Ortín, Miguel (2022), «Socio-cultural Norms in Ecological Psychology: The Education of Intention» in *Phenomenology and Cognitive Sciences*, vol. 23, pp. 1-19.

Sloutsky, Vladimir M., Fisher, Anna V. (2012), «Linguistic Labels: Conceptual Markers or Object Features?» in *Journal of Experimental Child Psychology*, vol. 111, n. 1, pp. 65-86.

Sloutsky, Vladimir M., Lo, Y.-F. (1999), «How Much Does a Shared Name Make Things Similar? Part 1. Linguistic Labels and the Development of Similarity Judgment» in *Developmental Psychology*, vol. 35, n. 6, pp. 1478-1492.

Sloutsky, Vladimir M., Lo, Y.-F., Fisher, Anna V. (2001), «How Much Does a Shared Name Make Things Similar? Linguistic Labels, Similarity, and the Development of Inductive Inference» in *Child Development*, vol. 72, n. 6, pp. 1695-1709.

Stokes, Dustin (2013), «Cognitive Penetrability of Perception» in *Philosophy Compass*, vol. 8, n. 7, pp. 646-663.

Stokes, Dustin (2014), «Cognitive Penetration and the Perception of Art» in *Dialectica*, vol. 68, n. 1, pp. 1-34.

Stokes, Dustin (2018), *Rich Perceptual Content and Aesthetic Properties*, in Bergqvist, Anna & Cowan, Robert (eds.) (2018), *Evaluative Perception*, Oxford, Oxford University Press.

Stokes, Dustin (2020), «On Perceptual Expertise» in *Mind and Language*, vol. 36, n. 2, pp. 241-263.

Tanaka, James, Taylor, Marjorie (1991), «Object Categories and Expertise: Is the Basic Level in the Eye of the Beholder?» in *Cognitive Psychology*, vol. 23, pp. 457-482. https://doi.org/10.1016/0010-0285(91)90016-H.

Van de Cruys, Sander, Bervoets, Jo, Moors, Agnes (2022), «Preferences Need Inferences: Learning, Valuation, and Curiosity in Aesthetic Experience», in *The Routledge International Handbook of Neuroaesthetics*, London, Routledge, pp. 475-506.

Voltolini, Alberto (2013), «The Content of a Seeing-As Experience» in *Aisthesis*, vol. 6, n. 1, pp. 215-237.

Voltolini, Alberto (2023), «Perceiving Aesthetic Properties» in *The British Journal of Aesthetics*, vol. 63, n. 3, pp. 417-434. https://doi.org/10.1093/aesthj/ayac039.

Wittgenstein, Ludwig (1953), *Philosophical Investigations*, London, Macmillan.

Zajonc, Robert B. (1968), «Attitudinal Effects of Mere Exposure» in *Journal of Personality and Social Psychology*, vol. 9, pp. 1-27.

Zipoli Caiani, Silvano (2022), «Intelligence Involves Intensionality: An Explanatory Issue for Radical Enactivism (Again)» in *Synthese*, vol. 200, n. 132. https://doi.org/10.1007/s11229-022-03527-y.