



HAL
open science

Fifty Shades of Affective Colouring of Perception

Frédérique de Vignemont

► **To cite this version:**

Frédérique de Vignemont. Fifty Shades of Affective Colouring of Perception. *Australasian Journal of Philosophy*, 2021, pp.1-15. 10.1080/00048402.2021.1965176 . hal-03854382

HAL Id: hal-03854382

<https://hal.science/hal-03854382v1>

Submitted on 15 Nov 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fifty shades of affective colouring of perception

Frédérique de Vignemont

Abstract:

Recent evidence in cognitive neuroscience indicates that the visual system is influenced by the outcome of an early appraisal mechanism that automatically evaluates what is seen as being harmful or beneficial for the organism. This indicates that there could be valence in perception. But what could it mean for one to see something positively or negatively? Although most theories of emotions accept that valence involves being related to values, the nature of this relation remains highly debated. Some explain valence in terms of *evaluative content*, others in terms of *evaluative attitude*. Here I shall argue that an account of affective perception in terms of attitude has more chance to succeed. To do so, I will first highlight the difficulties that a content-based approach faces, considering the many forms it might take. I will conclude that seeing the world positively or negatively involves more than a positive or negative content; it involves a distinctive attitude, but which one? Should it be conceived in imperative or evaluative terms? And what makes this attitude distinct from a proper emotion?

Introduction

Affective colouring of sensory perception is pervasive (Fulkerson, 2020). By that, I simply mean that many sensory experiences *somehow* involve an affective dimension in addition to their sensory dimension. Typically, French cheeses do not only smell like dirty socks. They smell awful. Hearing footsteps behind you in a dark alley is not simply hearing a soft sound.

It is hearing something bad approaching. And it can feel so horrible for some people to see blood that they prefer to turn their head to avoid looking at it. The crucial question is at what level the affective dimension kicks in. Taste, smell and affective touch, which are sometimes called ‘affective senses’, resemble pain and have thus been analysed on its model. But what about other types of sensory experiences that do not intuitively strike us as affective, such as visual experiences?

According to a conservative account, visual experiences can give rise to emotions, desires, and evaluative judgments, but they themselves cannot be affectively loaded. However, it has been recently proposed that there may be more to visual experiences than what the conservatives assume, including evaluative, affective and motivational components (Stokes, 2018; Audi, 2013; Noordhof, 2018; Mendelovici, 2013; Fulkerson, 2020). This richer conception of perception is made plausible by recent evidence in cognitive neuroscience showing that the visual system is influenced by the outcome of an early appraisal mechanism that automatically evaluates what is seen as being harmful or beneficial for the organism (e.g. Barbot and Carrasco, 2018). It may then be that some visual experiences can be endowed with features that are normally distinctive of emotional experiences. But how, then, to differentiate those sensory experiences that are affectively-loaded (affective perception) from emotional experiences that are grounded in sensory experiences (perceptual emotion) (Fulkerson, 2020)? Seeing a storm can fill you with awe, for instance. But is your visual experience itself affectively-loaded or is it merely the cognitive base of a subsequent emotion? It is not even clear that you can have one without the other: in most cases at least, affective perception should give rise to perceptual emotion. In such a situation, can the subject herself distinguish between the two? For many, it may seem doubtful, especially if they defend a perceptual theory of emotion, which conceives of emotions as a non-sensory form of perceptual

awareness (Prinz, 2004; Tappolet, 2016; Döring, 2007).¹ To find conceptual space for the notion of affective perception is thus a delicate enterprise.

In this paper I shall examine what it would be required for visual experiences to be affectively-loaded. For sake of simplicity, I shall take valence as the key signature of affectivity (Carruthers, 2018).² I will further assume that it makes a phenomenological difference to be valenced, that visual experiences can have an affective phenomenal character. But what could it mean for one to see something positively or negatively? Although most theories of emotions accept that valence involves being related to values, the nature of this relation remains highly debated. Some explain valence in terms of *evaluative content* (e.g. Carruthers, 2018; Tappolet, 2016; Mendelovici, 2013; Bain, 2013), others in terms of *evaluative attitude* (e.g. Deonna and Teroni, 2012; Mitchell, forthcoming; Jacobson, 2019). Here I shall argue that an account of affective perception in terms of attitude has more chance to succeed. To do so, I will first highlight the difficulties that a content-based approach faces, considering the many forms it might take. I will conclude that seeing the world positively or negatively involves more than a positive or negative content; it involves a distinctive attitude, but which one? Should it be conceived in imperative or evaluative terms? And what makes this attitude distinct from a proper emotion?

¹ To avoid any confusion, I will systematically keep the term ‘affective perception’ to refer to affectively-loaded *sensory* perception, and use ‘quasi-perception’ for emotion defined in perceptual terms.

² Hedonic tone and motivational force often follow from valence but I shall leave these standard attributes of affective experiences aside and start with the most minimalist criterion of valence. I shall thus not address the question of sensory pleasure. Though the boundary between affective perception and sensory pleasure is not clean-cut or simple, they refer to two distinct phenomena (Aydede and Fulkerson, 2013). Sensory pleasures refer to sensory sensations that feel pleasant whereas in affective perception, what is affectively tinted is the object of perception, and not the sensation itself.

1. Tidy and Scruffy

Imagine that two flatmates, Tidy and Scruffy, find a rotten sandwich in the fridge. Tidy is very orderly and keeps everything clean while Scruffy is more relaxed about cleanliness and hygiene, to say the least. One can easily grant that it does not feel the same for the two flatmates to look at the rotten food. But what is the difference exactly? One can simply claim that Tidy feels disgusted by the mold on the sandwich, whereas Scruffy does not care. But is that all there is? Recent evidence in cognitive neuroscience indicates that the difference between Tidy and Scruffy can also be at the visual level. More specifically, it has been found that (i) there is an automatic appraisal mechanism that occurs very early on and (ii) this specialized mechanism influences visual processing, boosting performance and altering the way things subjectively appear. Here I will simply summarize some of the main findings without going into much experimental details (for review, see Maratos and Pessoa, 2019; Pourtois et al., 2013).

By default, the amygdala automatically monitors the environment for objects and events relevant for the well-being of the organism, whether one pays attention to them or not, whether one is consciously aware of them or not. This appraisal system promotes self-preservation (defensive circuit) and propagation opportunities (nourishment, procreation, and nurturance). One may then say that the amygdala responds to values. According to what Jacobson (2019) calls the “relationalist biological view”, values are mind-independent properties to harm or benefit a given organism in specific circumstances.³ Because of their importance for survival, the brain evolved a specialized system specifically tuned to these values that operates as early as 80 ms after the presentation of the stimulus. It discriminates between neutral stimuli and valenced stimuli (mainly threats such as snakes and angry faces, but sometimes also appetitive stimuli, such as food and smiling faces), and it does so even

³ It is thus possible for an individual to be mistaken about what is dangerous or beneficial for her given the situation.

when the visual information is weak and irrelevant for the ongoing task. It relies on two feedforward sweeps of inputs, one throughout the visual pathway and one through a subcortical route, which has been described as being “quick and dirty”. Both routes allow for rapid evaluative processing based on limited and rudimentary visual information. Once the stimulus is evaluated as relevant for the organism, the amygdala increases the neural response in the visual cortex. As a consequence, participants are better able to attend and to track affectively salient visual images than neutral ones. It has also been found that performance is better or faster in visual search, attentional blink, and spatial orienting tasks. Interestingly, a recent study shows that the sudden appearance of a fearful face enhances contrast appearance of a subsequent Gabor patch (Barbot and Carrasco, 2018). This effect is larger in participants with anxiety personality traits. This shows that seeing negative stimuli not only improves performance but also modifies the way they subjectively appear. This is just a small sample of a vast array of studies, which are summarized as follows by the neuroscientist Patrick Vuilleumier, one of the experts in the field:

Thus, affective processes do not only contribute ‘subjective’ components to perception, by bringing ‘colors’ or ‘flavors’ to it, but may actually determine and shape the contents of perception itself through multiple effects on sensory pathways. (Vuilleumier, 2015, p. 30).

Though fascinating, these results have two limits. First, the only stimuli that have been used are those that are potentially harmful to the individual (and more rarely beneficial). They are not exclusively stimuli for which one is innately hard-wired to react since early appraisal can be influenced by learning and expertise, but it remains an open empirical question whether one can generalize the notion of affective perception to non-biological values, to aesthetic

perception, for instance, or to gustatory likes and dislikes.⁴ Hereafter, I shall thus limit myself to a relatively minimal notion of affective perception, restricted to biological values. The second limit is that the results show that early appraisal gives salience to stimuli that have relevance for the well-being of the organism, tagging them as important and potentiating attentional effects. But salience and valence should not be confused. The fact that the visual system processes more effectively the rotten food that is evaluated as a potential threat does not suffice to demonstrate that there is valence in perception. Nonetheless, sensitivity to biological values at the level of the visual system makes *prima facie* plausible the hypothesis that visual experiences can be affectively charged. The working assumption from now on will thus be that Tidy's visual experience is negatively charged, whereas Scruffy's visual experience is not. One can further claim that Tidy's experience is appropriate to what is potentially harmful in their fridge, unlike Scruffy's experience. The crucial question now is to determine how to best account for the affective load of visual experiences, whether we should do it in terms of content or in terms of attitude.

2. Content and attitude

I assume here a relatively minimal representationalist theory, according to which sensory experiences have intentional content. Many have argued that intentional content exhausts phenomenal character, appealing to the argument from transparency (Harman, 1990). When one turns one's introspection to one's visual experience of a cube, one is typically aware of the properties of the cube only, and not of the visual experience itself. It has thus been argued that the properties of the cube that are represented suffice to account for what it is like to see it. How, then, to explain that it feels different when one touches the same cube? According to

⁴ For instance, I shall not include within the scope of affective perception cases such as Tidy seeing aubergines in her fridge, which she dislikes. Early appraisal is not concerned by flimsy like and dislike but rather by allergy, phobia, and possibly addiction. My explanandum thus differs from Jacobson's (2019).

what Chalmers (2004) calls pure representationalism, there is no phenomenal difference between the two experiences that cannot be explained by differences in content (colour, for instance, being represented only in the visual content). Alternatively, one may be more modest and defend only impure representationalism: the phenomenal character is determined not only by the particular intentional content, but also by the attitude-type (Crane, 2003) or what has been also called mental paint, which refers to the properties of the vehicle of the experience (Block, 1996). The objective here is clearly not to settle this major debate, but only to apply it to the issue of affective perception. One way to conceive of affective perception is indeed as a distinct modality of sensory perception, an affective one. We can then ask whether one can account for the affective colouring of visual experiences by appealing exclusively to intentional content (pure representationalism), or also to mental paint (impure representationalism).⁵

Interestingly, the debate between a content-based approach and an attitude-based approach has been at the core of recent discussions on emotions.⁶ Most theories assume that emotions are affective representations with non-conceptual evaluative content (Tappolet, 2016; Carruthers, 2018; Tye, 2008). On this view, emotions are defined by their content, which is non-conceptual and which represents evaluative properties. For instance, what makes an emotion an episode of fear or anger is that it represents danger or offensiveness. However, the content-based approach to emotions has been recently criticized in favour of an attitude-based approach (Deonna and Teroni, 2012; Mitchell, forthcoming; Jacobson, 2019). According to

⁵ Here I shall stay within the representationalist framework and determine which version has the most chance to account for affective perception, pure or impure. I will not consider the possible contribution of non-representational properties, what Block (1996) calls mental latex.

⁶ The following taxonomy does not aim to be exhaustive. In particular, it focuses on philosophical theories, leaving aside emotion theories described in cognitive neuroscience.

Deonna and Teroni (2012), for instance, the phenomenal character of the emotion must be explained in terms of the specific attitude that the subject experiences, an attitude that consists in a holistic pattern of action-readiness (the subject experiences her body as being poised to act in various ways towards the target of her emotion). Mitchell (forthcoming) defends a more moderate version of the attitude-based approach. Like advocates of the content-based approach, he argues that emotions have non-conceptual evaluative content, which represents the formal object distinctive of each emotion, but he also argues that the phenomenal character of emotions outruns the evaluative content and that it is also determined by a specific evaluative attitude. For example, when you feel afraid of the dog, you disfavour the dog that you represent as dangerous.

Both the content-based and attitude-based accounts have been criticized (e.g., Dokic and Lemaire, 2013; Dokic and Lemaire, 2015), but it is beside the point which theory offers the most promising account of *emotion*. Our concern is only how to best explain the affective charge of visual experiences. All we need at this stage is to state that there are two main theoretical paths for visual experiences to be affectively-loaded, through intentional content and through intentional mode.

3. Evaluative content

Insofar as most theories have explained emotions by appealing to evaluative content, this may seem the most promising path to pursue for affective perception. This brings affective perception close to aesthetic perception and to moral perception, which have both recently been described in terms of evaluative content too. In this view, when seeing a ballet or a murder, the gracefulness of one or the evilness of the other is said to be represented at the perceptual level (Stokes, 2018; Audi, 2013; Noordhoff, 2018). Likewise, one may claim, when Tidy sees the rotten sandwich, its badness is represented by her visual experience. More

generally, one can propose that affective perception involves a certain kind of aspectual sensory experience: one sees the world under a specific aspect, which is determined by one's biological cares and concerns, and the sensory content represents thin evaluative properties, such as being good or bad.

This hypothesis goes against a sparse conception of perception, according to which only low-level sensory features can be represented. According to the sparse interpretation, Tidy's and Scruffy's visual experiences represent the same properties: a greenish layer at the top of a yellowish blob. Tidy can then *judge* this to be bad, but she cannot visually experience it because thin evaluative properties are not admissible candidates for visual content. The sparse conception, however, has been challenged in favour of a richer conception of perceptual content, which can represent some type properties, including natural kinds, artefacts, and causation (Bayne 2009; Siegel 2010). It has been recently extended to include evaluative properties (Stokes, 2018).⁷ One can then propose the following model:

Rich-content model: A subject's affectively-charged visual experience consists in her undergoing an experience whose visual content represents thin evaluative properties along with shape, colour and motion.

A now classic argument in favour of a rich conception of perception appeals to the method of phenomenal contrast but in the case of *evaluative phenomenal contrast*, it is difficult to rule out alternative explanations (Stokes, 2018; Werner, 2018; Logue, 2018). Let us return to the contrast between Tidy and Scruffy. First, one has to show that it follows from differences in visual phenomenology *per se*. As said earlier, there may also be emotional differences, such

⁷ For some, evaluative properties such as aesthetic properties are actually more admissible for perceptual content than natural kinds (Stokes, 2018; Logue, 2018). Whether a ballet is graceful or not is a matter of visual appearance, whereas whether I see a tomato (instead of a fake plastic toy that looks like a tomato) is a matter of internal structure, which cannot be seen.

as Tidy feeling disgusted. Furthermore, even if one grants that Tidy and Scruffy have different visual experiences of the food, those differences can result from mere attentional effects and it is hard to prove that Tidy's visual experience represents what she sees as bad on the only basis of the method of phenomenal contrast. Even if we leave aside these difficulties, the rich-content model faces another major problem. Thin evaluative properties have no unique sensory look, no unique visual appearance. There can be a sensory look for natural and artefactual kinds. There can be a sensory look for causation (involving spatio-temporal proximity between two events, for instance). There can also be a sensory look for aesthetic properties, such as gracefulness. In all these cases, one can claim that high-level properties supervene on a conjunction of low-level properties (Stokes, 2018). The difficulty with valence, however, comes from the variety of what can be appraised as being good or bad. For instance, the visual appearance of what can be evaluated as bad goes from the void below your feet to the snake in front of you, from the angry face of the man that you offended to the rotten food in your fridge. Valence does not look like anything, or it looks like too many things. Furthermore, the instantiation of the conjunction of some non-evaluative properties of the object does not suffice to predict whether the object is good or bad. For instance, peanuts are normally good unless one is allergic. As said earlier, valence tracks biological values, which are properties to harm or benefit a given organism in specific circumstances but the properties of the organism (such as its allergy) cannot be visually experienced. Values are subject-dependent (what is beneficial for one may be harmful for another), and context-dependent (what is good for one person now may be bad for her later if the circumstances change). It is difficult to make sense of a visual look for such properties.

Because of the differences between sensory and evaluative properties, some proponents of the content-based approach argue that they cannot be represented in the same way, and thus that goodness and badness cannot be part of the visual content. Nonetheless, they also claim,

the phenomenal content of perceptual experiences can include more than sensory content: “Perceptibility is not exhausted by perceptuality” (Audi, 2013: 41). Perceptual experiences can include a further layer of content, which is non-sensory and which represents evaluative properties (Audi, 2013; Noordhof, 2018). One may then phrase the following new interpretation of the content-based approach to affective perception:

Layered-content model: A subject’s affectively-charged visual experience consists in her undergoing an experience that integrates a visual content, which represents shape, colour and motion, and a non-sensory content, which represents thin evaluative properties.

A good example of the layered-content model can be found in the evaluativist theory of pain, according to which pain has both somatosensory content (sensory-discriminative component) and evaluative content (affective component):

A subject’s being in unpleasant pain consists in his (i) undergoing an experience (the pain) that represents a disturbance of a certain sort, and (ii) that same experience additionally representing the disturbance as bad for him in the bodily sense (Bain, 2013: S82)

Likewise, Fulkerson (2020) proposes that affective perception includes two contents:

Perception thus seems to have a kind of dual content: there is a component that tracks objective sensible features (intrinsic and/or relational) and another that tracks their emotional features (their affective-motivational import). (Fulkerson, 2020)⁸

An additional clause could be a feeling of integration between the two contents. Audi (2013), for instance, insists that it is essential to moral perception for the sensory and the evaluative

⁸ Fulkerson defends a psychofunctionalist account but his view is compatible with evaluativism.

contents to be bound together. The notion of integration is generally left underspecified but we can suggest that it works along the same principles as perceptual binding. We are constantly confronted with visual scenes full of objects with their distinctive shapes and colours. The visual system has then the task to bind the right colour with the right shape for each object. This is known as the binding problem in the literature: to select and segregate the relevant properties to integrate together from those pertaining to other objects. One might then suggest that binding operates in the same way for valence as for colour and shape.

With this conception, part of the worries that we have had so far disappears insofar as evaluative properties no longer need to be included in the sensory content. However, one still needs to give a satisfying account of the non-sensory content and of its link to the sensory one. The evaluative content is often conceived as amodal, cognitively penetrable, and possibly conceptual (but not necessarily). For instance, according to Audi (2013), it is part of what he calls non-basic perception, which represents non-observable properties and which is non-inferentially grounded in basic perception (which represents observable properties). Lyons (2018) goes a cognitive step further. He claims that evaluative properties can be represented in what he calls perceptual seemings (or intuitions). Perceptual seemings are not beliefs: one does not endorse their content to be true; things just *look* to be good or bad. According to Lyons, perceptual seemings qualify as perceptual for the sole reason that they result from perceptual processes but one may wonder whether this is sufficient. By making the evaluative content extrinsic to the sensory content, there is a strong risk that valence becomes only an external addendum coming from further cognitive or emotional states.⁹

⁹ One solution is to clearly differentiate the affective layer from an emotion. Fulkerson (2020), for instance, proposes that in affective perception only, one experiences the affective features as response-dependent properties of the perceived object. The problem, however, is what Fulkerson claims to be specific to affective

To recapitulate, the content-based approach to affective perception can proceed in several ways, from the strong hypothesis that the sensory content itself can be about evaluative properties to the weak hypothesis that perceptual experiences can include two types of content, sensory and evaluative, to the even weaker hypothesis that evaluative properties figure only at the level of perceptual seemings. Though the layered conception is more promising, it also opens crucial questions. Indeed, if the evaluative content is too high-level and if it is not appropriately bound to the sensory content, there is the risk that it is no longer legitimate to conceive of it as being part of visual experiences.

4. Hot perception

We have just seen that to explain how visual content can represent thin evaluative properties is already a theoretical challenge. However, even if we managed to do so, this would still not be sufficient for visual experiences to be affectively-loaded. What seems to be missing is the distinctive phenomenal character of affective states. This may be described as a form of engagement toward the world.

Seeing the utterly specific ways in which a situation, animal or person is appealing or repellent requires an appropriate *affective engagement* with the situation, animal or person. (Johnston, 2001: 181, my underline)

It is difficult, however, to precisely grasp what the notion of affective engagement entails. We can start by giving a negative definition. Consider Goldie's (2000) example of the rally driver who sees a sudden incoming car but reacts in a cold-blooded manner: he remains detached from what he sees. He pays focused attention but this is not sufficient for his visual

perception is actually the description of emotions given by quasi-perceptual theories (see Tappolet, 2016, for instance).

experiences to be 'hot'. Affective engagement is not simply attentional engagement. Nor should it be reduced to motor engagement. The rally driver's motor responses are indeed quick and in direct relation with his perception. This is precisely what makes him such a good driver. A proponent of the content-based approach might then propose that what is missing is the representation of the situation as being bad or dangerous. On this interpretation, his visual state only represents: <car in a collision course>. But would it change anything if it visually represented the situation as bad? It seems that the rally driver could still remain cold-blooded. Advocates of the content-based approach to emotions themselves acknowledge that the notion of evaluative content can fail to capture affective phenomenology when they criticize cognitivist theories of emotion. The solution, on their view, is for the evaluative content to be *non-conceptual* (Tappolet, 2016; Carruthers, 2018). But is one affectively engaged just because one does not conceptualize what one experiences? Let us return to the rally driver. If he evaluates the situation, it is on the basis on early appraisal and it is doubtful that it is conceptual. And yet he does not feel afraid or anything. There seems to be a gap between non-conceptual evaluative content and affective phenomenal character (Mendelovici, 2013). This is not to deny that emotions and affectively charged visual experiences are non-conceptual. It is only to question whether it suffices to make them hot.

What matters indeed is how the rally driver relates to his non-conceptual badness content. Merely seeing things as being bad is not the same as seeing them negatively. There is a subtle difference that we are trying to catch with the notion of affective engagement. To account for this affective phenomenal property has been one of the main challenges for pure representationalism. In particular, it has been argued that unlike sensory experiences that appear to many as diaphanous, affective states are weakly opaque (Mitchell, forthcoming). When Tidy introspectively attends to her visual experience, she is mostly aware of the properties of the rotten food that she sees but she is also aware of something beyond that, of

the special way she relates to the rotten food. If so, the phenomenal character of her visual experience is not exhausted by its content and there must be something else that determines it. One might say that evaluative content is like a fire drill: it remains empty of affective and motivational force on its own. Less metaphorically, affective engagement goes beyond the representation of the properties of the object that elicits the emotion, it involves how one represents these properties. It then seems that the notion of evaluative perception leaves out the engagement in or with the world, and thus what seems to be the signature of affective phenomenology. If so, evaluative perception is not the same as affective perception. There is a missing ingredient.

A proponent of the content-based approach may then posit a notion of affective properties conceived as *primitive phenomenal properties* that precisely match the way emotions feel: “Affective properties are exactly those familiar qualities we experience when we are angry, sad, etc” (Mendelovici, 2013). On this view, the same circumstances reliably cause more or less the same affective experiences but objects and events do not have affective properties. We thus live in an Edenic world (Chalmers, 2006), in which our affective states systematically misrepresent the objective world. In the specific case of affective perception, the Edenic view implies that we can never experience biological values but only their phenomenal proxy, so to speak. However, one can wonder about the cost of divorcing the phenomenal content from veridicality. One consequence of this view is that Tidy’s and Scruffy’s affective responses (or lack of it) equally misrepresent what they see. Nonetheless, Mendelovici (2013) is willing to make the extra claim that one is “appropriate”, and not the other. But why then posit affective properties that no object has instead of simply claiming that the content does not represent affective property? These *sui generis* phenomenal properties indeed play no role in explaining what really matters, namely that Tidy’s experience is appropriate, but not Scruffy’s.

A possibly more promising path for the content-based approach is to propose that goodness and badness are represented under a special *mode of presentation* (Döring and Lutz, 2014). On this view, thin evaluative properties can be presented under different modes, and only one of them guarantees affective phenomenal character. The challenge then is to spell out this special mode without falling into a vicious circle. To claim that there is an affective mode of presentation would simply beg the question. What is needed is an account that does not appeal to the notion of affectivity that we try to explain. But which one? One possibility is to focus on the bodily dimension of affective experiences. No matter whether one defends an embodied theory of emotion or not, there is a general agreement that affective phenomenology typically involves bodily responses. By bringing the body back into visual experiences, one may bring affect. One may then propose that one has an affectively-charged visual experience when one represents things as being good or bad under a *bodily mode* (B-good or B-bad).¹⁰ One might further note that it is precisely why one appeals to the bodily metaphor of being cold-blooded when describing the rally driver: he simply does not represent the situation as B-bad. But can one offer a satisfying characterization of the bodily mode of thin evaluative properties? Let us assume that under the bodily mode, one conceives of an event as being good or bad an event that triggers specific patterns of bodily changes. But what patterns? The difficulty is that threat can present itself in many ways and there is no unique bodily response to it. Typically, one's body does not react in the same way whether one sees rotten food or an angry face. It is already controversial whether there are specific bodily signatures for each type of emotions, but it is even more unlikely for the broader category of thin evaluative properties and one cannot leave unspecified the bodily changes if one does not want the definition to apply beyond valence.

¹⁰ For instance, see Bain (2003) for an account of pain in terms of somatosensory mode of presentation (before he started defending an evaluativist account).

There might be other ways to spell out the specificity of the mode of presentation that characterizes affective phenomenology but there is always the risk for these special modes of presentation to be mere post-hoc explanations. Some modes of presentation seem to be specifically created and designed for the sole purpose of rescuing pure representationalism but to avoid infinite multiplication, they need to have their own legitimacy and it remains to be shown that such modes exist for affective engagement.

5. Positive and negative attitudes

I have argued that if one wants to assume a ‘hot’ version of affective perception, it is difficult to see how evaluative content could suffice. Instead, I argue that the affective phenomenal character is determined by the properties of the vehicle of the representation, that is, by mental paint. Without affective mental paint, the affective phenomenology of our visual experiences remains quite bland. But what colour is affective mental paint?

One possible source of inspiration for an attitude-based approach to affective perception can be found in Millikan’s (1995) notion of primitive *pushmi-pullyu representations* (hereafter PPRs). PPRs embed two distinct attitudes into a single mental representation. Like any standard sensory representation, PPRs have a ‘mind-to-world direction of fit’. Their content is descriptive: it must fit with what is in the world in order to be true. But unlike standard sensory representation, PPRs also have a ‘world-to-mind direction of fit’. Like desires, their content is directive and it can be successful or not. PPRs blends information and motivation together inextricably: “The very same representation tells in one breath what is the case and what to do about it.” (Millikan, 2004: 18). The notion of PPRs has been used to account not only for affordances (Millikan, 1995), but also for emotions (Scarantino, 2010, 2014), pain (Martinez, 2011; Barlassina and Hayward, 2019), and smell (Martinez, 2015). For instance, according to Martinez (2011), pain consists in a specific type of PPRs, whose

indicative content represents disturbances in bodily states (sensory dimension) and the imperative content prescribes “Don’t have this bodily disturbance!” (affective dimension). He later applies the same strategy to account for olfactory experiences of disgusting smells: their imperative content then prescribes “Stay away from that source of pathogens!” (Martinez, 2015). One may then be tempted to generalize the use of PPRs to other types of affectively charged sensory experiences.¹¹

Pushmi-pullyu model: A subject’s affectively-charged visual experience consists in her having a representation that both describes how things are and prescribes either to interact with them or to stay away from them.

On this model, Tidy has a visual representation that both describes the rotten sandwich and prescribes her to stay away from it. The content is true if it fits with what is in the fridge and it is successful if Tidy does not get in contact with it. One may wonder, however, whether the opposition ‘interact’ versus ‘stay away’ neatly matches with the opposition positive versus negative. Consider the way animals react to predators. It is classically characterized in terms of the 3 Fs: fight, flight, or freeze. Out of the three, only flight corresponds to the prescription for negative valence. Even more problematic is the fact that to fight is a way of interacting with the threat, which should normally correspond to positive valence. One might reply that one sees things positively only when one’s visual experience prescribes to interact ‘positively’ with these things, which is clearly not the case when one fights. But now consider seeing your baby crying. You then visually relate to this event negatively and yet, you will try to comfort your baby, thus showing positive interaction. Consequently, the same prescription, such as “to interact”, can correspond either to positive or to negative valence, and thus cannot account for the affective phenomenal character of visual experiences. The same problem is

¹¹ One may question whether this kind of accounts falls under the attitude-based approach. One may indeed claim that it is a specific type of *content*, an imperative content, that plays the explanatory role. This content, however, is embedded in a certain attitude, a command, and Klein (2015), for instance, opposes the imperative view to pure representationalism.

encountered by the reflexive version of the pushmi-pullyu model defended by Barlassina and Hayward (2019). On their view, the directive component of PPRs targets the experience itself, instead of the perceived object. One may then claim that there is affective perception if the visual experience describes how things are and prescribes “more of seeing” for things that are good or “less of seeing” for things that are bad. This reflexive version might possibly account for the crying baby (though it may seem that you rather want your son to stop crying than your visual experience to end), but it does not work for threat. Although danger is bad, sometimes you cannot help but look at it, as when the serial killer is chasing you and you keep turning back to check on him. One may even say that one should not avoid looking at what is bad. The ostrich with its head in the sand has indeed bad reputation. Hence, it seems that once again, to the same prescription can correspond positively or negatively charged visual experiences.

Instead of an imperative attitude, one may propose to characterize the affective load of sensory perception in evaluative terms.¹² On this view, affectively-loaded visual experiences involve a feeling of favour or disfavour toward the perceived object that has been appraised as possibly harmful or beneficial, what Mitchell (forthcoming) calls “a primitive ‘yes’ or ‘no’”. The evaluative attitude is directed at the state of affairs as it is represented by the visual experience (such as the rotten sandwich), and not at the visual experience itself. Importantly, the same evaluative feeling can apply indistinctively whether one has a negatively charged visual experience of the rotten food, of a child crying or of the angry face of a man. Hence, though there are many reasons for goodness and badness, the hypothesis is that there are only two types of attitudes, positive or negative. The affective attitude then comes “in addition” to the sensory experience (Jacobson, 2019, p. 123). We may then summarize the following new account:

¹² A further alternative is in conative terms (Jacobson, 2019), but I shall not consider this option here.

Layered-attitude model: A subject's affectively-charged visual experience consists in (i) a visual experience that represents the world, and (ii) an evaluative attitude directed towards what is seen that is bound with the visual experience.

A possibly helpful model for understanding how the two layers are related is *multimodal perception*. Stretch your fingers. You can see and feel their movements and though you have two distinct modes of access, you integrate the experiences together because they are about the same object, namely, your fingers. I propose that a similar binding process operates in affective perception, integrating together the evaluative feeling and the visual experience in virtue of being about the same seen object. On this view, the evaluative attitude is conceived as a *sui generis* modality. At the computational level, it results from early appraisal, which can operate before the visual information is fully processed and reaches consciousness. This non-sensory modality can be integrated with sensory modalities. The multisensory model can then account for a number of situations. Imagine, for instance, seeing and hearing approaching danger. You then have a single affective attitude, a single primitive 'no', attached to both sensory experiences. Imagine now that you are famished and that you see an apple with a snake next to it. You then experience both a 'yes' and a 'no' towards what you see. To avoid any evaluative conflict, the seen apple needs to be integrated with the primitive 'yes', and the seen snake with the primitive 'no'. Solving such a complex case, however, is a common routine for multisensory binding, though it can sometimes go astray, leading to multisensory illusions. However, the perceptual system has the computational resources to adequately bind together distinct experiences even in complex perceptual scenes.

To summarize, I propose here that an integrative process operates for affective perception, binding together visual experiences with evaluative attitudes into a unified affectively-loaded visual experience. Of course, much more work needs to be done before any such account of affective perception is ultimately plausible but for now I prefer to return to our original

question: on this account, in what sense is affective perception different from any emotion that is a response to a given sensory content? Disgust for what one sees, for example, can also be described as a first-order affective attitude that comes in addition to the visual experience of the object. In this case the visual experience plays the role of the cognitive base of the emotion. If Tidy feels disgusted by the rotten sandwich, her visual experience of the mold provides the object of her disgust and she can reply that she feels disgust because there is mold. It may thus seem that we could describe affective perception in the same way and claim that the visual experience is the cognitive base of the evaluative feeling. In what sense, then, does it differ from perceptual emotion? We have just seen that at the functional level the evaluative feeling is bound to the visual experience, and one may argue that there is no such functional binding between the emotion and its cognitive base. Still, one may ask, does it make a phenomenological difference? Or are affective perception and perceptual emotion phenomenally indistinguishable?

I will now highlight two phenomenological differences. The first difference concerns the affective attitude itself. Arguably, the feeling of disfavour is less rich and fine-grained than the attitude of disgust. It is the most rudimentary affective attitude, even more basic than the so-called 'basic emotions'. It is completely neutral as for the specific reasons the seen object is disfavoured. It simply consists in a brutal "No!". It corresponds to the most minimal degree of affect, generated by rapid evaluative processing based on limited and rudimentary visual information. Hence, unless one is willing to claim that this super-basic attitude is an emotion, there is a phenomenological difference between affective perception and perceptual emotion because of the nature of the affective layer.

There is a further difference because of its relation with the sensory layer. The cognitive base stands in an epistemological relation with the emotion (Deonna and Teroni, 2012). It can explain why an emotion is justified or not. From a phenomenological perspective, however,

the cognitive base and the emotion remain independent. It does not directly change anything to the cognitive base whether one experiences the emotion or not. For instance, just thinking that there is something rotten in your fridge can make you feel disgusted. In that scenario, it seems that the belief that constitutes the cognitive base of your emotion remains the same whether you feel disgust or not. One can expect the same to be true when the cognitive base is perceptual instead of cognitive. It is not phenomenological bound with the emotion. By contrast, one experiences a phenomenal unity between the visual experience and the evaluative feeling in affective perception. I do not mean here that the visual experience and the evaluative feelings are unified in virtue of merely being co-conscious. I mean that one can no longer retrieve the 'pure' visual experience untainted of affective colouring once it is bound with the evaluative feeling. Once you have put affective mental paint, you cannot go back to the image in black and white, so to speak. There is then a sense in which the coloured picture is more than the combination of the black and white drawing and the coloring. Likewise, seeing negatively is more than seeing and disfavouring. The phenomenal character of the affectively-charged visual experience is not exhausted by the combination between the visual phenomenal properties combined with the affective phenomenal properties. They are fused together into a unified compound that one experiences. Still it might seem that at the end I have not shown that valence is really *in* perception but can only be fused with it. This is true but I believe that one cannot legitimately argue for more than that in the specific case of visual experiences. One should not expect the relationship between affective and sensory dimensions to be the same for affectively-charged visual experiences and for painful experiences.

Conclusion

The objective of this paper was primarily to draw the map of the various theoretical paths that a representationalist theory of affective perception could take with the aim of hopefully shaping further debates. I organized it around two main issues:

- (i) Where is valence in perception? In content or in attitude?
- (ii) What is the relationship between the sensory and the affective components?

We saw that each path is filled with obstacles. It may be possible to overcome some of them, but others might prove to be more fatal. I believe that the layered-attitude model is the most promising. First, the subject does not stay cold; she feels positively or negatively engaged towards what she sees. Secondly, the layered-attitude model is compatible with a conservative conception of perception. Thirdly, it offers one tentative story on how affective mental pain can be bound to sensory experiences and shared across sensory modalities inspired by the literature on multimodal integration. It remains to be seen whether this tentative story is actually the right one. This will, unfortunately, have to be a story for another occasion.

References

- Audi, R. 2013 *Moral Perception*, Princeton: Princeton University Press.
- Aydede, M. and Fulkerson, M. 2013. Affect: representationalists' headache, *Philosophical Studies* 170/2: 175-98.
- Bain, D. 2003. Intentionalism and pain, *The Philosophical Quarterly*, 53: 213, 502-523.
- Bain, D. 2013. What makes pains unpleasant?, *Philosophical Studies*, 166(1), 69-89.
- Barbot, A., & Carrasco, M. 2018. Emotion and anxiety potentiate the way attention alters visual appearance, *Scientific reports*, 8(1), 1-10.
- Barlassina, L. and Hayward, M. 2019. More of me! Less of me!: Reflexive Imperativism about Affective Phenomenal Character, *Mind*, 128: 512, 1013–1044.

- Bayne, T. 2009. Perception and the reach of phenomenal content, *The Philosophical Quarterly*, 59(236), 385-404.
- Block, N. 1996. Mental Paint and Mental Latex, *Philosophical Issues*, 7: 19–49.
- Carruthers, P. 2018. Valence and Value, *Philosophy and Phenomenological Research*. 97, 3, 658-680.
- Chalmers, D. J. 2004. The representational character of experience, in *The Future for Philosophy*, ed. B. Leiter, Oxford: Oxford University Press, 153–181.
- Chalmers, D. J. 2006. Perception and the fall from Eden, in *Perceptual Experience* ed T.S. Gendler and J. Hawthorne, New York: Oxford University Press, 49–125.
- Crane, T. 2003. The intentional structure of consciousness. In *Consciousness: New Philosophical Essays* ed Q. Smith & A. Jokic, Oxford: Oxford University Press, 33-56.
- Deonna, J. and Teroni, F. 2012. *The Emotions. A Philosophical Introduction*, London: Routledge.
- Dokic, J. and Lemaire, S. 2013. Are Emotions Perceptions of Value?, *Canadian Journal of Philosophy* 43(2), 227-247.
- Dokic, J. and Lemaire, S. 2015. Are Emotions Evaluative Modes?, *Dialectica* 69(3), 271-292.
- Döring, S. A. 2007. Seeing What to Do: Affective Perception and Rational Motivation, *Dialectica*, 61(3): 363–394.
- Döring, S. A. and Lutz, A. 2014. Book Review: Julien A. Deonna and Fabrice Teroni, *The Emotions. A Philosophical Introduction*, *Dialectica*, 68: 459–463.
- Fulkerson, M. 2020. Emotional Perception, *Australasian Journal of Philosophy*. 98(1), 16-30.
- Goldie, P. 2000. *The Emotions: A philosophical exploration*, Oxford: Oxford University Press.
- Harman, G. 1990. The intrinsic quality of experience, *Philosophical Perspectives*, 4, 31–52.

- Jacobson, H. 2019. Pain and mere tastes, in *Philosophy of Suffering: Metaphysics, Value, and Normativity*, ed D. Bain, M. Brady and J.Corns, London: Routledge, 123-144.
- Johnston, M. 2001. The authority of affect, *Philosophy and Phenomenological Research*, 63(1), 181-214.
- Klein, C. 2015. *What the body commands: the imperative theory of pain*, Cambridge: The MIT Press.
- Logue, H. 2018. Can we visually experience aesthetic properties?, in *Evaluative perception*, ed A. Bergqvist and R. Cowan, Oxford: Oxford University Press, 42-57.
- Lyons, J. 2018. Perception and intuition of evaluative properties. In *Evaluative perception*, ed A. Bergqvist and R. Cowan, Oxford: Oxford University Press, 183-199.
- Maratos, F. A., & Pessoa, L. 2019. What drives prioritized visual processing? A motivational relevance account, *Progress in brain research*, 247, 111-148.
- Martínez, M. 2011. Imperative Content and the Painfulness of Pain, *Phenomenology and the Cognitive Sciences* 10(1): 67–90.
- Martinez, M. 2015. Disgusting Smells and Imperativism, *Journal of consciousness studies*. 22 (5-6):191-200.
- Mendelovici, A. 2013. Pure Intentionalism About Moods and Emotions, in *Current Controversies in Philosophy of Mind*, ed Uriah Kriegel, New York: Routledge, 135-157.
- Millikan, R. 1995. Pushmi-pullyu representations, *Philosophical perspectives*, 9, 185-200.
- Millikan, R. 2004. On Reading Signs: Some Differences Between Us and the Others, in *Evolution of Communication Systems: A Comparative Approach*, ed D. K. Oller and U. Griebel, Cambridge: The MIT Press, 15-30.
- Mitchell, J. forthcoming. *Emotion as Feeling Towards Value*, Oxford: Oxford University Press

- Noordhoff, P. 2018. Evaluative perception as response-dependent representation, in *Evaluative perception*, ed A. Bergqvist and R. Cowan, Oxford: Oxford University Press, 80-108.
- Pourtois, G., Schettino, A., and Vuilleumier, P. 2013. Brain Mechanisms for Emotional Influences on Perception and Attention: What is magic and what is not, *Biological Psychology*, 92/3: 492–512.
- Prinz, J. 2004. *Gut Reactions: A Perceptual Theory of Emotions*, Oxford: Oxford University Press.
- Scarantino, A. 2010. Insights and Blindspots of the Cognitivist Theory of Emotions, *British Journal for the Philosophy of Science*, 60, 3, 729-768.
- Scarantino, A. 2014. The Motivational Theory of Emotions, in *Moral Psychology and Human Agency: Philosophical Essays on the Science of Ethics*, ed Justin D'Arms and Daniel Jacobson, Oxford: Oxford University Press, 156-185.
- Siegel, S. 2010. *The Contents of Visual Experience*, New York: Oxford University Press.
- Stokes, D. 2018. Rich perceptual content and aesthetic properties, in *Evaluative perception*, ed A. Bergqvist and R. Cowan. Oxford: Oxford University Press, 19-41.
- Tappolet, C. 2016. *Emotions, Values and Agency*, Oxford: Oxford University Press.
- Tye, M. 2008. The experience of emotion: An intentionalist theory, *Revue Internationale de Philosophie*, 62: 25–50.
- Vuilleumier, P. 2015. Affective and motivational control of vision, *Current opinion in neurology*, 28(1), 29-35.
- Werner, P.J. 2016. Moral Perception and the Contents of Experience, *Journal of moral philosophy*, 13, 294-317.