

Laura Silva¹

Towards an Affective Quality Space

Abstract: *In this paper I lay the foundations for the construction of an affective quality space. I begin by outlining what quality spaces are, and how they have been constructed for sensory qualities across different perceptual modalities. I then turn to tackle four obstacles that an affective quality space might face that would make an affective quality space unfeasible. After showing these obstacles to be surmountable, I propose a number of conditions and methodological constraints that should be satisfied in attempts to construct an affective quality space. Before concluding, I detail the high explanatory pay-off such a project promises.*

Keywords: quality space; emotion; perception; phenomenology; qualia.

1. Introduction

Experiences have particular qualities or characters often called their ‘mental’ or ‘phenomenal’ ‘quality’ (e.g. Rosenthal, 2015). That is, there is a particular way it is like to undergo a given experience. For example, smelling lavender, seeing red, and feeling pain each have their own characteristic phenomenal quality. Of our experiences, emotions are perhaps those whose phenomenal qualities are most salient. Despite this, their phenomenal qualities remain relatively under-theorized.

Correspondence:
Email: laura.silva.13@ucl.ac.uk

¹ Department of political science, Université Laval, Quebec, Canada.

In the philosophy and cognitive science of perception, ‘quality spaces’ have been constructed to map the phenomenal qualities of the senses (Clark, 1993; Gauker, 2017). These spaces map phenomenal qualities, within different sensory modalities, along specific dimensions. Colour quality space, for example, is three-dimensional as every colour that can be discriminated by humans under normal conditions varies along the axes of hue, brightness, and saturation (e.g. Keller, 2016). Experimental work has been used to map quality spaces for senses other than sight (e.g. Young, Keller and Rosenthal, 2014), and used to inform work in the philosophy of mind and perception (Berger, 2018; Clark, 1993; Rosenthal, 2015). No parallel research programme is, to my knowledge, underway in the philosophy of emotion.

My aim in this article is to do some foundational work to support the construction of an affective quality space. In other words, my aim is to clear the theoretical terrain such that a quality space of this type becomes feasible, and to provide some principles to help guide its construction. I will start by outlining in more detail what a quality space amounts to, drawing on work in the philosophy and cognitive science of perception. Then I will consider a set of potential obstacles that we might think arise for the project of constructing a quality space for the affective realm. These obstacles all concern, in some way or another, disanalogies between emotion and perception. I will argue that these obstacles are surmountable. After doing so I will turn to the provision of a few principles and methodological constraints to help guide the construction of the affective quality space. I end by outlining just how significant a well-mapped affective quality space could be for work in both philosophy and psychology, thereby stressing the high explanatory pay-off of such a project.

2. Quality Spaces

Philosophers have traditionally relied on introspection to characterize the phenomenal quality of experiences, some even claiming that these qualities are irreducibly subjective and distinct from any natural or physical properties in the world (Chalmers, 1996; Jackson, 1994). Others have argued that phenomenal qualities can fit within a physicalist picture of the mind and are not necessarily tied to conscious experience (Keller, 2016; Rosenthal, 2015). Although the notion of quality spaces has a long history (Quine, 1969; Rosenthal, 2016; Sellars, Brandom and Rorty, 1997), until the early 1990s the

literature largely ignored empirical findings in psychology and neuroscience. Austen Clark's (1993) *Sensory Qualities* notably broke this trend by bringing empirical insights into philosophy, highlighting how psychologists and neuroscientists approached the phenomenal qualities of sensory experiences in non-qualitative terms. It was arguably Clark's (1993) work, as well as Rosenthal's (1991), that popularized the notion of the 'quality space' in contemporary cognitive science, marking a paradigm shift in how to think of phenomenal qualities that now represents the dominant approach.

On this approach, called quality space theory (QST), quality spaces map capacities for perceptual discriminations for a given perceptual modality. For example, since red is experienced as more similar to orange than to blue, it will be mapped in the colour quality space as closer to orange than to blue. Similarly, as vinegar smells more similar to lemon juice than to coconut water, the smell of vinegar will be placed closer to the smell of lemon juice in the olfactory quality space than to the smell of coconut water. These similarities and differences in phenomenal qualities are ones that individuals have the capacity to subjectively discriminate, but, crucially, quality spaces are not constructed by reliance on introspection but rather on discrimination tasks. That is, phenomenal qualities are individuated by what perceptual qualities or properties we can discriminate in the world. Unconscious phenomenal qualities exist on such views, because our capacities for perceptual discriminations do not rely on conscious experiences (for example, subjects can be unconsciously primed with specific colours) (Keller, 2016; Rosenthal, 2010).

Experiments in psychology, primarily in the field of psychophysics, have gathered data on 'just noticeable differences' between stimuli, so as to construct sensory quality spaces. In brief, experiments task subjects with identifying when two stimuli are different along a given dimension. There is a threshold for the perception of differences between stimuli that plausibly track our capacities for perceptual discriminations. These discriminable properties can be mapped in a space that represents those things in the world that we can discriminate perceptually. As these perceptual discriminations depend on our capacities for perception, and perception has qualitative properties, a qualitative space can be mapped that corresponds to the discriminable phenomenal quality space. Sensory qualities mapped in quality space can represent relations between different phenomenal qualities. Importantly, quality spaces are typically multidimensional. That is, phenomenal qualities do not vary along one dimension alone. As

mentioned, the colour quality space is typically thought to be 3D (e.g. Hardin, 1993), as is the auditory quality space (pitch, timbre, and loudness), while the olfactory quality space plausibly has many more dimensions (e.g. Young, Keller and Rosenthal, 2014).

The quality space paradigm has become a popular and attractive approach to sensory qualities and has occasioned innovative contributions to the philosophy of mind and perception. For example: sensory quality spaces engender important insights about the nature of perceptual experiences that might not be attainable from introspection alone (Rosenthal, 1991); they provide a naturalistic account of perceptual experiences by correlating phenomenal qualities to empirically observable variables (Gauker, 2017); they allow us to map relations between experiences and can inform functional explanations of the phenomenal character of these experiences (Keller, 2016); philosophers have also proposed that quality spaces might map the representational content of perceptual states (Berger, 2018; Gauker, 2017), contributing to our understanding of the semantics or meaning of such states by mapping their content along a number of fine-grained dimensions.

Given the proliferation, popularity, and explanatory potential of quality spaces, one might expect philosophers to have mapped such quality spaces for arguably our most salient phenomenal experiences: emotions. Unfortunately, this is not the case. Despite some philosophers explicitly mentioning that they believe quality spaces can be constructed for the affective domain (Rosenthal, 2015; Tappolet, 2020), very little philosophical attention has been devoted to this task. Despite philosophical work on the affective quality space being scarce, there is ample work in philosophy that will be relevant to the construction of one. Work in philosophy of perception, canvassed above, will be key to constructing an affective quality space by (relative) analogy, while work in the philosophy of emotion, on the nature, phenomenology, and function of different emotion types as well as on the nature of emotion in general, will be central to the project (Deonna and Teroni, 2012; Tappolet, 2016).

In psychology, the idea of an affective quality space is far more widespread (Roseman, 1991; Russell, 2003; Scherer, Dan and Flykt, 2006). Two-dimensional models, where emotional experiences are mapped according to their valence (positive vs. negative feel) and level of arousal (or intensity) remain the most widespread (Yik, Russell and Barrett, 1999; Russell and Lemay, 2000; Kuppens *et al.*, 2013). Despite this, they have been critiqued for failing to distinguish

emotion types that we are committed to construing as distinct (Fontaine *et al.*, 2007; Trnka *et al.*, 2016). For example, on valence vs. arousal models, fear and anger would occupy the same point in affective space, as they are both negatively valenced and can be high in intensity or arousal. To address this problem, psychologists have suggested that we add further dimensions to increase the space's capacity to differentiate emotion types. A number of models propose three- (Osgood, 1966; Paramei *et al.*, 1994) or four-dimensional models (Fontaine and Scherer, 2013; Fontaine *et al.*, 2007; Sokolov and Boucsein, 2000), where emotions are typically mapped along the dimensions of valence, arousal, power and, for most 4D models, novelty. Although such models are widespread, there is little consensus on the number or nature of the dimensions of the affective quality space. Constructing an affective quality space, in philosophy, is likely to yield innovative contributions not only for debates in the philosophy of mind and emotion, but for psychology itself, as this work is likely to speak in favour of some existing empirical models over others as well as result in empirical predictions ripe for experimental exploration.

An important question that now arises concerns how the affective spaces proposed in the psychology of emotion literature relate to the quality spaces of QST. That is, do psychologists mean the same thing by 'quality space' as what is meant by philosophers that adhere to QST? There has not been sufficient dialogue between these two literatures for a clear answer to be given. On the one hand, psychologists typically hold both that emotions can occur unconsciously, and that they have what philosophers call phenomenal properties, suggesting that they may very much be in line with QST. On the other hand, affective quality spaces, in psychology, focus much more on first-personal experience than the sensory quality spaces of QST that want explicitly to move away from introspection towards third-personal measures. That is, affective quality spaces in psychology often seem to map differences and similarities between emotional experiences themselves in terms of their subjectively experienced phenomenal properties, rather than inferring phenomenal properties from their role in independently manipulable tasks, such as discrimination. This seems like a departure from QST where the main force of the theory is exactly that phenomenal properties are mapped without reference to subjective experience, based on discriminatory capacities. In what follows I aim to lay the foundations for the construction of an affective quality space in line with QST. We will see that some

departures and disanalogies from classic QST will arise, but they will not, I think, prove fatal to the project.

It is actually surprising, given the current state of debates in the philosophy of emotion, that the affective quality space is not a central topic of research. This is because one of the most popular contemporary theories of emotion is the perceptual theory, which takes emotions to be in some sense ‘perceptions of value’ (Döring, 2007; Prinz, 2006; Tappolet, 2016). It is particularly surprising, then, that perceptualists have not attempted to construct affective quality spaces given that they rely on analogies between perception and emotion² and such quality spaces have been popular in modelling perceptual experiences. Perhaps this project has not been undertaken because it faces insurmountable theoretical and practical obstacles that block the endeavour from the get go. I will consider four such putative obstacles below and show them to be anything but insurmountable.

3. Putative Obstacles

Potential obstacles to the construction of an affective quality space will likely concern differences between emotion and perception. These differences may underpin claims that perceptual quality spaces can be mapped while affective quality spaces cannot. Foreseeable obstacles include the following.

3.1. Obstacle 1: Emotions and perceptions involve different types of discriminable properties

A first obstacle that might arise for the construction of an affective quality space is the thought that while perceptual objects, in some sense, ‘have’ the properties our phenomenal experiences represent them as having, the objects of emotions surely do not. We are more inclined to grant that a fire truck is really red than we are to grant that there are values or evaluative properties out there in the world. The reason this would pose a problem for the construction of an affective quality space is, presumably, that it is harder (or impossible) to

² Perceptual theorists take analogies between emotions and perceptions to be key to understanding the emotions. These analogies include: phenomenological claims (both emotions and perceptions have distinctive subjective qualities), epistemic claims (both emotions and perceptions play important epistemic roles), and representational claims (both emotions and perceptions have non-conceptual representational content) (Cowan, 2016; Döring, 2009; Tappolet, 2016).

manipulate physical properties of objects and systematically correlate them with phenomenal qualities in the case of emotions, as there are no corresponding physical properties to be found.

There are a number of reasons why we shouldn't be too concerned by obstacle 1. First, although the objects of perception are thought to exist out there in the world, whether all their properties do is a subject of much debate. A popular view in neuroscience, for example, is that colours do not exist out there in the world, but exist only as phenomenal properties in the minds of agents (e.g. Maund, 2006). Colour realism, on the other hand, is most popular amongst philosophers (*ibid.*). Realism about evaluative properties is also popular, indeed dominant, amongst philosophers, however,³ such that differences in the nature of properties in the perceptual versus emotional realm are not so clear cut (Deonna and Teroni, 2012; Scarantino and de Sousa, 2018). That is, the metaphysics of perceptual and evaluative properties need not speak clearly in favour of a drastic difference between the two. Second, independently to considerations of metaphysics, there is considerable consensus that, much like colour correlates with observable and manipulatable physical properties such as surface reflectance and lighting, evaluative properties correlate with similarly observable and manipulatable physical properties. For example: danger plausibly correlates with size, speed, shape, and height (a large animal is typically more dangerous than a small one, a creature moving quickly towards you is more dangerous than a static creature, a sharp knife is more dangerous than a blunt one, a high cliff is more dangerous than small one, etc.). Other emotions might not bear such clear correlations to physical properties but they nonetheless correlate with features of events or states of affair. For example, anger correlates with intentional harms (if a harm is done accidentally, it triggers less anger than if done intentionally), sadness with losses (a loss peripheral to one's sphere of core concerns elicits less sadness than a loss at the core of one's concerns). These features can and are manipulated in experimental studies to elicit target emotions and study them.

³ Generally, these can take either response-dependent or response-independent forms. Fitting-attitude analyses, where objects have the evaluative properties they have by virtue of a particular emotional response being fitting, are the most popular type of the first form of realism. Response-independent forms of value realism typically employ a notion of supervenience whereby evaluative properties are thought to supervene on non-evaluative natural properties (see Deonna and Teroni, 2012).

To sum up, the first obstacle suggests that a quality space cannot be constructed for the affective realm because emotions and perceptual states concern different types of properties (perceptual and evaluative respectively), where the former are deemed to be in some sense more ‘real’ than the latter. It is important to note that in so far as we are concerned with mapping *phenomenal* properties and both perceptions and emotions seem to have these, we are concerned with mapping the same general type of property. When it comes to the question of whether and how these phenomenal properties correspond to physical or external properties in the world, we have seen that the metaphysics of perceptual and emotional objects need not provide us with a sharp difference between the two. Perhaps most importantly, we do not seem to require the exact same metaphysical verdict on this to be able to construct a quality space in the affective realm, what we need is the existence of certain external features in the world that correlate with specific affective phenomenal properties or qualities. This, few people deny.

3.2. Obstacle 2: Emotional experiences are more variable than perceptual experiences

Related to obstacle 1, we might think that even if emotional experiences correlate with features of physical objects, states of affairs, or events, this correlation is far more subjectable to cultural and individual differences than is the correlation between perceptual experiences and relevant physical features of objects. Again, there are a number of ways of responding to this concern. We might argue that, despite cultural differences in emotional expression and behaviour, there is considerable consistency in the types of emotions that exist across cultures and their correlation to specific evaluative properties (Ekman and Friesen, 1971; Izard, 1994). For example, while some cultures, such as the Japanese, might be far less expressive of some negative emotions than North American culture, these emotions are arguably still felt in response to intentional harms and offences in these cultures (Ekman and Friesen, 1971). What will differ, in addition to expression, is what sorts of things count as offences across different cultures (spitting in public, for example, is not seen as offensive in some cultures). There will then be considerable flexibility amongst what specific things fall under the purview of certain evaluative properties, such as ‘offensive’ or ‘harmful’, as well as how agents express their emotions, without these differences threatening

the connection between specific emotion types and their related evaluative properties across people and cultures (Mallon and Stich, 2000). Anger still correlates with offences and intentional harms, no matter whether or how anger is expressed, nor what counts as offensive or harmful in a particular culture, for example. What does this mean for our affective quality space? That perhaps experimenters working in different cultures will be inclined to use different vignettes and cases in their experiments and be attuned to specific expressions of emotion types in their culture. This shouldn't, however, in itself, block the construction of an affective quality space, as we continue to be able to correlate the phenomenal properties of specific emotions to specific evaluative properties.

What if the phenomenal qualities of emotions themselves, and not just its triggers and expressions, differ from culture to culture? I am sympathetic to this view. The fear of a medieval monk might be qualitatively different from my own, and likewise the anger of an oppressed group today may be qualitatively different from that of the group's allies. Accepting such differences in phenomenal quality need not threaten an affective quality space. Indeed, quality spaces might be helpful for distinguishing subtle differences between different emotional experiences that often fall under the same label. So long as there are general patterns between triggering events, contextual features (perhaps including one's social position, identity, and culture), and phenomenal qualities, mapping affective qualities seems possible. Lastly, it should be noted that in so far as quality spaces map *capacities* for affective experiences, interpersonal and intercultural differences need not pose a problem to mapping affective space. This is because capacities are distinct from occurrences; that is, everyone need not respond to an object that you find dangerous with fear for them to be capable of having a similar emotion.

3.3. Obstacle 3: Emotional experiences do not admit of similarity comparisons necessary for the construction of a quality space

One interpretation of this problem would say that, even though we might experience different emotion types as having different phenomenal qualities (most people would agree anger feels different to sadness, and that both feel different to joy or envy), we lack the capacity to make fine-grained distinctions regarding emotion phenomenology that would permit the construction of an affective quality space (Pendoley, in progress). Again, it seems the perceptual realm is our

target comparison, and that the thought is that similarity judgments are clearer or easier to make in the perceptual than the emotional realm. Note that this comparative claim is, however, entirely compatible with the construction of an affective quality space. We do not need judgments of similarity in the emotional realm to be as easy to make as they are in the perceptual realm for an affective quality space to be constructed, we just need similarity judgments in the emotional realm to be sufficiently feasible for this purpose. Do we have reason to think they are not, as the obstacle suggests? I do not think so.

Before outlining why, I would just like to note that discriminations in the perceptual realm can be deceptively difficult. The case of vision may be paradigmatic because it is particularly straightforward, but when it comes to other senses, such as olfaction and audition, it may become significantly more difficult. That this be the case, however, merely supports the construction of quality spaces for audition and olfaction that may be more coarse-grained than that of vision, rather than speaking against the possibility of olfactory or auditory phenomenal spaces. If our discriminatory capacities are less fine-grained in a given sense modality, then coarser grained mappings will presumably be sufficient to individuate sensory qualities.

Further, it is important to note that perceptual quality spaces have typically been constructed based on discrimination tasks rather than explicit judgments of similarity. That is, quality spaces are not typically based on experimental paradigms that ask subjects to compare how similar two stimuli are, but rather based on discrimination tasks where subjects are asked whether they can discriminate between two different stimuli (e.g. Keller, 2016). The reason for this is that discrimination tasks straightforwardly measure performance (whether the stimuli are in fact different) while similarity judgments are considered to rely on reports of mental content rather than performance, and there are arguably many ways of judging similarity that might be confounded in the same study (*ibid.*). This is just to say that, although quality spaces map phenomenal qualities in a manner where the relative similarities between them, along different dimensions, is represented, the experimental work that is used to construct such spaces does not typically involve explicit similarity judgments.

So, it seems that the pertinent question is whether we are capable of making sufficiently fine-grained discriminations in the emotional realm to construct a quality space. I think that literary works provide ample evidence for our capacities to describe and discriminate the emotions at a considerable level of detail. That is, even though we

might lack an extensive affective vocabulary that would allow straightforward discriminations and comparisons (I return to this below), we do not lack the capacity to appreciate the complexities of emotional experience and many have endeavoured to vividly describe them. Here are just a few illustrative examples:

On jealousy:

His jealousy, like an octopus which throws out a first, then a second, and finally a third tentacle, fastened itself irremovably first to that moment, five o'clock in the afternoon, then to another, then to another again. (Proust, 2004)

One's jealousy, ransacking the past in search of a clue, can find nothing; always retrospective, it is like the historian who has to write the history of a period from which he has no documents; always belated, it dashes like an enraged bull to the spot where it will not find the dazzling, arrogant creature who is tormenting it and whom the crown admires for his splendor and cunning. Jealousy thrashes around in the void. (*ibid.*)

On grief:

No one ever told me that grief felt so like fear. I am not afraid, but the sensation is like being afraid. The same fluttering in the stomach, the same restlessness, the yawning. I keep on swallowing. (Lewis, 2021)

Her absence is like the sky, spread over everything. (*ibid.*)

The death of a beloved is an amputation. (*ibid.*)

On anger and frustration:

My rage intensifies because I am not a victim. It burns in my psyche with an intensity that creates clarity. It is a constructive healing rage. (hooks, 1995)

I can remember the frustration of not being able to talk. I knew what I wanted to say, but I could not get the words out, so I would just scream. (Grandin, 2018)

Although often metaphorical, and rarely concerned with explicitly comparing distinct emotional experiences, the undeniable detail and differences in literary descriptions of emotional experience provide evidence that we have the capacity to discriminate emotional experience at a considerable level of detail.

Granted, it would be hard to construct a quality space based only on literary descriptions of emotional experience, but this is not what I am proposing, not least because this is not how quality spaces are constructed on QST, as we saw. Literary examples might serve as a

starting point, or data to return to so as to centre phenomenological insight in the construction of such a quality space. For example, these brief passages suggest that emotional experience involves discriminable profiles of attention, bodily feeling, and action tendencies. In grief, attention to the absence ‘spreads over everything’ while in jealousy it ‘fastened itself irremovably’ to specific moments. Grief and fear are claimed to involve similar bodily manifestations, ‘the same fluttering in the stomach, the same restlessness, the yawning’, while frustration can involve an urge to scream. Profiles of attention, bodily feeling, and action tendency plausibly contribute to qualitative differences in emotional experience. Importantly, they are not only observable from a first-person perspective, or described in works of literature; much work in psychology has investigated these as components of emotion (e.g. Carretié, 2014; Faucher and Tappolet, 2002; Frijda, 2005). For example, one phenomenal dimension along which emotional experience is thought to vary is whether the emotion is self-focused or world-focused (Lambie and Marcel, 2002). Pride is paradigmatically self-focused while sadness may typically be world-focused, for example. We can imagine that differences in focus will occur within emotion types as well — pride at having volunteered at a charity may be more world-focused than pride in one’s cleverness.

Work in psychology, canvassed above, also supports our ability to discriminate emotional experiences. The two most popular dimensions are: intensity (or arousal) and valence (e.g. Yik, Russell and Barrett, 1999). That is, we are capable of discriminating between positive and negative emotional experiences, and we are capable of discriminating between stronger and weaker emotional experiences. While anger, sadness, and fear all share a negative valence, we can have stronger or weaker experiences of all of them.

Approaches in psychology, mostly dimensional and appraisal theories, have suggested a number of other axes along which emotional experience can be mapped (Fontaine *et al.*, 2007; Russell, 2003; Scherer, Dan and Flykt, 2006). The novelty of the situation that triggers the emotion is thought to make a difference to emotional experience (e.g. Fontaine *et al.*, 2007). For example, becoming angered by an expected and foreseeable obstacle has a different qualitative feel than becoming angry at a surprising and unexpected obstacle. Similarly, anger of the first kind will differ qualitatively from a case of fear or joy caused by an unexpected stimulus. In addition to novelty or familiarity, another popular dimension is that of coping potential or control, which refers to the agent’s ability to cope

with, change, or control the eliciting stimulus (e.g. Scherer, Dan and Flykt, 2006). In sadness, for example, one often, or perhaps typically, cannot do much to change the eliciting situation (such as a death), which means there is a low coping potential or sense of control over the situation. This plausibly translates at the phenomenological level to senses of deflation, smallness, and helplessness. In many experiences of anger, on the other hand, despite valence also being negative, agents typically feel more in control or capable of addressing the eliciting situation (for example when feeling emboldened to speak up against an injustice or confront a perpetrator). This translates to feelings of power, preparedness, and possibility at the phenomenological level.

In addition to literature and psychology, philosophical work on the emotions also provides much evidence of our ability to pinpoint phenomenological characteristics of emotions and differences between them.

On shame, for example:

The core of the originating experience is an experience of damage: deep damage, or damage to the sense of self... we damaged ourselves by bringing down upon ourselves the hostile judgment. But, since the judgment is not necessarily one with which we concur, the judgment is experienced primarily as the brute impact of the world upon us. (Wollheim, 1999)

Pure shame is not a feeling of being this or that reprehensible object but, in general, of being an object, i.e., of recognizing myself in that degraded, dependent, and frozen being which I am for the Other. (Sartre, 1993)

Philosophers have been preoccupied with distinguishing shame from the closely related emotion of guilt. A central difference that is cited is that shame involves a relatively more global, or more severe, negative evaluation than guilt. Deonna, Rodogno and Teroni (2011) cash out this severity as follows: in shame, one evaluates oneself as lacking the capacity to uphold a value that one holds, while in guilt one merely evaluates one's behaviour as having transgressed a norm. One's capacity to uphold such norms is not (as severely) threatened in guilt. The phenomenology of shame is often characterized by feelings of deflation, smallness, and incapacity (Gilbert, Pehl and Allan, 1994). This fits with a more severe, and (relatively) more global, negative evaluation of oneself or one's capacities. There is a sense of helplessness in shame; while in guilt there is a stronger felt division between the self that evaluates and the evaluated behaviour/omission, one does

not wholly identify with one's actions, and plausibly positions oneself as both judge and wrongdoer. This matches differences in phenomenology between the two emotional phenomena: guilt is characterized by feelings of sinking and depletion, but is also animating, involving action readiness to seek amends (*ibid.*; Silva, 2022a). These phenomenological and conceptual distinctions between shame and guilt fit nicely with experimental work in psychology that highlights guilt as having higher coping potentials than shame (Tracy and Robins, 2006). Behaviour (in guilt) is typically easier to address, change, or excuse, than one's self, or broader capacities (in shame) (Silva, 2022a).

Work in psychology and philosophy, then, corroborates what we might think literary works suggest, that we are capable of making fine-grained phenomenal discriminations in the affective realm.

It is important to note that work in psychology and philosophy centres subjective experience far more than classic QST does for the perceptual realm; that is, it seems that experiences themselves are the object of comparisons and discriminations, rather than the objects in the world. Indeed, early psychological work was often based on introspection alone (Wundt, 1912/1924). Over the years, however, and with the development of new techniques, psychologists have been pre-occupied with correlating emotion types with subtle features of stimuli, to better understand the nature of emotions (more on this below). Many of the appraisals proposed by appraisal theory are not thought to be consciously accessible themselves, but rather thought to contribute to the overall phenomenal quality of emotions. And, in line with this, most psychologists think emotions occur not only consciously, but also unconsciously, without thereby lacking specific phenomenal qualities that their work aims to uncover.

The greater focus on phenomenal experience, then, in psychology and philosophy of emotion, need not be taken to be in tension with QST; that is, just because much attention is paid to phenomenology does not mean the affective quality space is to be mapped based on phenomenal experience alone. Although QST aims to make introspection orthogonal to the construction of sensory quality spaces, one might question whether complete disregard for phenomenal experience is really at play in this literature. Most importantly for our purposes, however, it does not seem like a decisive factor against the possible construction of an affective quality space that much attention is paid to the phenomenology of emotion. Rosenthal (2010) himself does not deny that consciousness gives us access to mental qualities, but rather holds that it should not have the last word on their nature.

With this, existing work in psychology, and my proposed quality space in philosophy, agrees.

3.4. Obstacle 4: There is no empirical programme analogous to psychophysics, which has been used to map perceptual quality spaces, that could be relevant to the construction of an affective quality space

As should be clear from the previous sections, there are a number of empirical programmes that will, I think, be relevant to the construction of an affective quality space. The most clearly relevant empirical programmes will be appraisal theories and dimensional theories. Dimensional theories of emotion attempt to describe emotions in terms of their subjective experience, ‘or what philosophers call qualia’ (Ellsworth and Scherer, 2003). Appraisal theorists attempt to go further and not only describe emotional experience but also explain it by appeal to the relation between emotions and the situations in which they occur. While some appraisal theorists posit appraisals as occurring prior to and causing emotional experiences, others take appraisals to be components of emotional experience such that ‘the subjective experience of fear, for example, is the feeling of high attention, negative valence, high uncertainty about what is happening or one’s ability to cope with it and so on’ (*ibid.*). As mentioned in the introduction, these types of theories are popular but the specific number and identity of dimensions for affective space are debated. The thought is that fine-grained appraisal dimensions underlie our experiences of thick evaluative properties. For example, negative valence, high arousal, high levels of control, and ‘other focus’ are thought to underlie our anger experiences as reactions to offence (e.g. Moors and Scherer, 2005). More will need to be said to adequately characterize the relation between fine-grained appraisals and experiences of thick evaluative properties in a full account (Lambie and Marcel, 2002; Teroni, 2021), but this is the underlying idea. Subjects can reliably feed back ratings of evaluatively significant scenarios which correlate with different appraisal dimensions, and manipulations of these appraisal dimensions by experimenters can induce different emotions in participants.

While many, although not all, appraisal theorists are concerned with investigating which appraisal dimensions contribute to emotional experience, making this empirical programme particularly relevant for the construction of an affective quality space, it will not be the only

relevant programme. Work on the attentional dynamics of emotion is likely to be relevant, as attentional focus is likely to contribute to phenomenal quality.⁴ Similarly, more wide-ranging studies in social psychology may also be relevant to our project, as many qualitative studies correlate experience to context. The terrain is undoubtedly vaster and more complex in the emotional realm than the perceptual realm, where psychophysics alone seems to be the gold standard experimental paradigm. This does not mean, however, that we lack pertinent experimental data or approaches for our task.

Another point to make here is that someone sympathetic to objection 4 might worry that we lack relevant affective vocabulary to discriminate between different experiences successfully during these various experiments. Here it is important to note that the experiments used to construct programmatic affective quality spaces in the affective sciences, as well as those used to construct sensory quality spaces, typically do not rely on capacities to explicitly conceptualize and name features of one's qualitative experience (Keller, 2016). Discrimination tasks in psychophysics just ask subjects to feed back whether two stimuli are the same or different along a relevant domain (darkness–brightness, for example), but subjects do not need to have concepts for all the varieties of shades they will encounter nor make explicit judgments of similarity. Similarly, in appraisal theory, vignettes are systematically manipulated while recording subjects' responses, typically through a questionnaire, where answers are given on rating scales. Questions such as 'How responsible do you feel x is for what happened?' and 'How strongly do you feel that this is wrong?' are asked, such that participants do not need to grasp any further concepts beyond these. Additionally, these questionnaires are sometimes coupled with non-verbal measures such as skin conductance and facial expression measures which provide additional measures that can be correlated to evaluative scenarios. Of course, questions can be raised here about methodology, and whether verbal or non-verbal responses actually correspond to differences in phenomenal quality, which is why critical engagement with experimental work will be key to the construction of any quality space.

⁴ Note that attention is thought to be necessary but not sufficient for conscious experience. That is, unconscious phenomenal qualities may involve attentional profiles (Noah and Mangun, 2020).

Despite the above, we might still worry about the concrete logistics of constructing an affective quality space based on existing empirical work. On QST, to construct the colour quality space, subjects are given stimuli of different shades, often only two stimuli at a time, in forced-choice tasks where they need to judge quickly if the stimuli are the same or different. Related tasks include organizing colour chips in terms of similarities and differences. There are two important things to underline here: there is a clear task that subjects are asked to perform (perceptual discrimination) and there are mind-independent stimuli that can be discriminated.

In the emotional realm, we saw in response to objection 1 that most philosophers are realists about evaluative properties, taking there to be features of objects and events that correlate with evaluative properties or on which such properties could supervene (Deonna and Teroni, 2012). What sort of task could one ask subjects to perform, however, that would result in something akin to discrimination tasks used in perceptual psychophysics? When presented with an image of a ferocious bear, for example, and another of an injustice, how do subjects begin to compare these stimuli and organize them or discriminate between them? What are they to discriminate in relation to? These are complex stimuli. It is certainly far less intuitive than discriminating between shades of colour. We therefore seem to lack a specific task that can be used to construct the affective quality space. Lacking such a task means that we seem to lack an observable capacity by which to map emotional experiences independently of subjective experience itself.

A few points on this. First, it is important to note that the paradigmatic quality spaces of QST are visual and the colour space is the one most often discussed. Now, colour is only one part of a visual quality space, as we are not only capable of discriminating colour but various other visual properties such as shape and size. Perhaps each emotion type should be seen as analogous to the colour quality space, in that it is only one important component of what can be felt. The thought is that, when given a range of stimuli designed to induce for example anger, that represent offences and injustices of different sorts, it becomes much easier to order them or compare them (which constitutes the worst harm? How intentional was the harm committed? How easy is it to attain rectification?). This might speak in favour of constructing the affective quality space in a piecemeal fashion, one emotion at a time, as discriminations are far more straightforward to make within an evaluative property type than between types.

It is the answers to the questions mentioned above, often called appraisal judgments, which can be seen as analogous to perceptual discrimination tasks in the emotional realm. They are gathered experimentally through rankings of emotionally relevant stimuli along various probed dimensions and have led psychologists to posit fine-grained dimensions that contribute to the phenomenal quality of emotions, such as coping potential, novelty, arousal, and valence. Our capacity to make such appraisal judgments, and the fact that they vary systematically across emotion types, suggests that they are a robust and relevant capacity for discriminating emotionally relevant stimuli, and hence for mapping the affective quality space.⁵

Note that most experimental paradigms in appraisal theory do not rely on inducing emotional experiences themselves, they present subjects with emotionally relevant stimuli (vignettes that describe situations that can be thought of as instantiating relevant evaluate properties), and ask subjects to use gradable scales to rate the situation/event in response to various questions. This seems to satisfy the main requirements of QST: there is a task performed by subjects that pertains to the stimuli presented, as opposed to emotional experiences themselves, which can inform us about the nature and structure of the phenomenal properties of emotions. As mentioned above, appraisal theorists have moved towards methodological pluralism (somatic and behavioural measures) to counter the problems of using only self-reported data (which include lack of reliability, and being easily influenceable by experimental set-up and prior assumptions of subjects) (Moors, 2017). I see no reason an affective quality

⁵ It might be contested here that, while perceptual dimensions are uncovered via perceptual discrimination tasks, appraisal dimensions seem to be known/given *a priori* in the very set-up of the experimental paradigm. That is, subjects are given the dimensions along which to rate the stimuli. I thank an anonymous referee for raising this issue. I see three options for a response which I will only be able to briefly outline here. First, we might doubt how independent the dimensions of perceptual quality spaces were from the experimenters' intuitions and preconceptions, especially once we move beyond the visual modality which is typically the main focus. Second, we should note that many questionnaires used in psychological research do not mention the appraisal dimensions themselves, but rather ask the subject in everyday language about different aspects of the vignette and then infer specific dimensions from their answers. Finally, we might accept this disanalogy between perceptual quality spaces and an affective quality space, and hold that it does not threaten the viability of the latter. We are after all dealing with very different mental phenomena, where one is far more cognitively penetrable than the other, such that it might not be surprising to expect that the relevant task required for the construction of the affective quality space be one that uses language and concepts which reflect our phenomenological intuitions about which dimensions are at play.

space in philosophy should not benefit from such methodological plurality. Note that if existing experimental work does not satisfy QST standards, we can design experimental paradigms aimed at greater similarity to those used in classic QST (complete analogy with the perceptual realm is unlikely to be possible but we can work towards structural similarity).

A fruitful direction that builds on appraisal theory would be to look to behavioural tendencies. Many theories take emotion types to be tied to action tendencies (Deonna and Teroni, 2015; Frijda, 2005) and there is mounting evidence that the properties of objects and the environment in which they occur correlate with fine-grained action patterns. For example, in fear of a predator, properties of threats (size, speed, distance to safety) determine the specific actions animals take (Evans *et al.*, 2019). It is not entirely implausible that action types correlate with differences in appraisals that, on appraisal theories, constitute emotion types and tokens within types: for example, previous encounters with predators play a role in animal actions in response to threats, which can plausibly be captured by the novelty dimension proposed by appraisal theory. Distance to safety, on the other hand, may be captured by coping potential.

Actions and behaviour, then, might be able to play the role of perceptual discriminations in the emotional realm, allowing a measure independent of emotional experience to reveal the dimensions of affective space. Fine-grained action tendencies would on such a view be the behavioural correlates of the appraisals that underly emotions, which in turn give us relevant dimensions along which to map phenomenal properties in affective space. In analogy with QST, behaviour would count as a way of discriminating emotionally relevant stimuli, and individuate affective phenomenal properties. This is not implausible given the important functional role emotions play in motivating specific actions, and the fact that emotionally relevant stimuli might be described as stimuli to be acted on in some way (Silva, 2022b). This behavioural move would, however, face more challenges than the first, appraisal judgment based, move in accounting for emotions that do not have clear behavioural correlates.

On either of these moves, which I grant would need to be far more fleshed out to gain traction, the measurable capacity does not, as it

stands, involve direct comparisons between evaluative stimuli.⁶ In both cases, the affective quality space would be mapped in a manner that is more indirect than that proposed on classic QST. Data would be gathered on correlations between stimuli and actions/motivations, or simply between stimuli and appraisal judgments, and the results would then be used to map the affective quality space. That is, the affective quality space is to be constructed using experimental data that do not rely on direct comparisons between evaluative stimuli, but rather ask subjects to make appraisal judgments, or action tendency judgments, about one situation at a time.

A few moves are available for a philosopher interested in constructing a classic QST quality space for emotions. One would be to help design experiments that incorporate direct comparisons between evaluative scenarios, including questions explicitly geared towards a comparison. It is, however, not clear that this would be necessary if sufficiently robust data are gathered for each scenario independently, and relative comparisons can be inferred, but it certainly would be a good way to corroborate these inferences. Again, I presume these comparative tasks would be more easily completed by subjects when the evaluative scenarios differ in detail but not overall evaluative property, suggesting a two-step procedure for the construction of affective space: begin with the construction of affective space for specific emotion types, using experimental evidence that explicitly probes this, and then either systematically infer how these emotion-specific spaces might be put together into an encompassing affective space, and/or use this first batch of data to design subsequent experiments that specifically probe cross-evaluative comparisons. This might be done by keeping as many features of a stimulus scenario constant as possible (proximity to incident and novelty of situation, for example), while varying the evaluative import alone (whether it's a threat, a loss, or an injustice, etc.). Doing so would allow

⁶ Above, I mentioned that explicit judgments of similarity between stimuli are not actually the primary paradigm used for the construction of sensory spaces, but rather discriminatory tasks are (p. 173). These discriminatory tasks are, however, comparative in nature, as subjects are asked to discriminate which stimuli is darker than the other(s), for example. What I am proposing for the affective realm respects this, as subjects are asked to judge which of two (or more) distinct evaluative stimuli ranks higher (or lower) on a number of sliding scales from which dimensions for affective space can be inferred.

experimental insight into how differences in evaluative properties impact the previous dimensions identified for each emotion type.

4. Guidelines for the Construction of an Affective Quality Space

4.1. Identifying dimensions for a quality space

To go about identifying dimensions for a systematic quality space, we should, in the first instance, think carefully about what conditions candidate dimensions should satisfy. I am heavily informed by Cochrane (2009) on this. Cochrane's article is, to my knowledge, the only one in the literature devoted to the topic of identifying dimensions for a potential conceptual space for the emotions. Crucially, however, his project concerns mapping a dimensional model for emotion *concepts* rather than for the *phenomenal qualities* of emotion. As these projects are undoubtedly related, however, in this section I attempt to build on his criteria and adapt them for our specific purposes. I've adopted (although sometimes renamed) three of his nine conditions (those I think are most relevant for our project) and I've importantly modified a further two more for our purposes.

4.1.1. True-scale condition

Dimensions should be such that a given emotion cannot simultaneously occupy more than one position on it. If more than one position is occupied on a specific dimension by a given emotional experience then it may be evidence that one is in a mixed emotional state.⁷

4.1.2. Continuous condition

Emotions should be able to occupy a range of continuous locations along a given dimension (as opposed to dimensions only admitting binary or a small set of predetermined possible locations for emotions to lie on). That is, where a given emotion lies with respect to a given phenomenal dimension should be a matter of degree.

⁷ That being said, see section on the empirical constraint page 186.

4.1.3. *Applicability condition*

Every dimension identified should in principle be applicable to each emotion; even if specific emotions end up occupying a zero value on the relevant dimension, the dimensions should be variables along which the emotion could conceivably vary.

4.1.4. *Stability condition*

Cochrane (*ibid.*) believes emotions should remain relatively static on at least some dimensions for the dimensions to be adequately mapping emotions of particular types. If a putative emotion demonstrates no stability whatsoever, across any of the various dimensions, it is doubtful that a specific emotion is being mapped at all. As emotional experiences evolve over time, however, I think it best that we allow stability to occur not only regarding points on a given number of dimensions at a given time, but of particular patterns of positions along various dimensions, over time. This would better reflect the dynamic profiles of emotions, and even allow us to map these profiles, while remaining committed to the importance of stability in affective space.

4.1.5. *Relative independence condition*

Cochrane subscribes to an ‘independence condition’ whereby ‘If occupying a point on one dimension limits the range of points that an emotion can occupy on another dimension’ (*ibid.*), then one of the dimensions should be discarded. I think this is too strong if what we are concerned with is mapping phenomenal qualities. Sometimes objectively discriminable properties will interact with each other at the level of experience without this meaning that one of the phenomenal dimensions should be abandoned. Take colour quality space, for example, structured by the dimensions of hue, saturation, and brightness. Colours high in saturation are often experienced as being high in brightness — this has been called the Helmholtz-Kohlrausch effect and violates Cochrane’s strict independence condition (Nayatani, 1997). In the emotional realm too, then, we should not be too quick to discard dimensions merely because there are strong relationships or interactions between them. If a dimension satisfies other conditions

and constraints, merely admitting of interactions with other dimensions should not count as a reason for disqualification.⁸

4.2. *The constraints of an affective quality space*

The construction of affective quality space should be conducted within the following methodological constraints.

4.2.1. *Empirical constraint*

This constraint involves, firstly, the identification of relevant empirical work for our project. I have already suggested a number of fruitful places to start. Secondly, it involves critically engaging with this work rather than simply importing purported conclusions. Often this will involve making more fine-grained conceptual distinctions than might currently exist in the empirical literature. For example, the dimension of valence in particular has already come under philosophical scrutiny (Carruthers, 2018; Colombetti, 2005; Solomon, 2006). There are many different ways in which an emotion can be said to be ‘good’ or ‘bad’ (Solomon, 2006, highlights 17 different ways, including pleasure, virtue, social status, and health). In so far as we are interested in valence as a specifically phenomenal quality though we might think pleasure more aptly tracks this, there still might be important distinctions to be made between whether the emotion as a whole is positive/negative in this sense and whether its components or aspects are (Colombetti, 2005). For example, there might be a sense of pleasure in anger despite this being a paradigmatic negative emotion. This relates to a second theoretical concern, which is whether valence should be seen as consisting of two mutually exclusive poles or not (answering in the affirmative might violate the complexity of emotion experience) (*ibid.*). One option would be to propose a number of finer-

⁸ Note that Cochrane (2009) takes independent dimensions to be preferable to non-independent ones, other things being equal, such that he is likely to agree with my relative independence condition if further considerations speak strongly in favour of non-independent dimensions. Nonetheless, Cochrane is right to hold that independent dimensions are explanatorily preferable for a number of reasons (for example, independent dimensions respect our intuitive and mathematical conception of dimensions, and they are plausibly clues to underlying components of emotions in a way that non-independent dimensions are not). It remains, however, a central question whether an affective space made up of entirely independent dimensions is feasible (see *ibid.*, p. 384), and my ‘relative independence condition’ aims to highlight that we shouldn’t be discouraged even if it is not.

grained concepts that might fall under ‘valence’, conceived as an umbrella term. In any case, this is the sort of work that needs to be done so as to comply with the empirical constraint. The affective quality space must be empirically informed but only contributions from empirical work that survive, or are modified by, careful conceptual work will be deemed acceptable.

Finally, a key component of the empirical constraint is to verify that there are robust correlations between certain phenomenal qualities and observable properties (broadly construed). Attention must be paid not only to the results of experiments but also their methods. Which experimental paradigm is best will likely depend on the candidate dimension under consideration. For example, non-verbal measures such as eye-tracking might be key to studying the attentional profiles of emotion and exploring whether they correlate with differences in appraisal judgments, while manipulating emotionally salient components of vignettes or computer tasks will plausibly be better placed to study coping potential and valence dimensions. In general, we can consider that for a correlation between phenomenal and external features to be considered robust, it would be ideal to observe this correlation across a range of distinct experimental paradigms. Systematic interdisciplinary work of this kind will then be necessary to comply with the empirical constraint and construct a viable affective quality space.

4.2.2. Phenomenological constraint

This constraint involves centring subjective emotion experience. The above discussion should have already suggested some promising directions. In addition to heeding empirical work that centres subjective experience it will be important to come back to literary and philosophical work that probes emotional experience. As mentioned above, this is a key difference between the affective quality space and classic quality spaces on QST. Again, I do not take this difference to be fatal to the project for a number of reasons. First, we may doubt classic QST quality spaces are immune to phenomenologically guided intuitions and experimental directions. Second, I agree with QST that subjective experience should not be the main guide in the construction of an affective quality space, and the direction recent psychological work has taken reflects the same commitment. Nonetheless, results that starkly violate our phenomenological data should be submitted to intense scrutiny. Lastly, we should perhaps expect phenomenology to be more relevant to the construction of an affective quality space than

a sensory space because emotions are phenomenologically far more complex than perceptions.

The relationship between the empirical and phenomenological constraint is not linear or one-directional. We must seek a reflective equilibrium of sorts between the two, where each informs the other in a continual feedback loop. What will this mean concretely? Well, for one, our robust phenomenological commitments should constrain the affective quality space. For example, if two emotions that we take to be different are appearing in the same location in the quality space (or far too close together to match our phenomenological evidence), then this should give us reason to revise the quality space and our experimental methods. This can be seen as an example of the phenomenological constraint taking primacy over the empirical constraint. On the other hand, we should be prepared to import new vocabulary from sound empirical work to better characterize our phenomenal experience. This would be an example of the empirical work taking primacy over, or fundamentally influencing, our phenomenological evidence.

Although some of the terms used in appraisal theory may seem contrived and unintuitive to apply to our everyday emotional experiences (for example: coping potential), this in itself is not reason enough to reject that these are real dimensions along which our phenomenal experience varies. We have only to look at the sensory quality spaces to find evidence of conceptual innovation. It is uncontroversial that experiences, sensory and certainly emotional, are often ineffable or hard to put into words. This is because experiences, many think, are more fine-grained than words or concepts. It is therefore no surprise that we might find our existing vocabulary lacking when it comes to the task of describing our experiences in detail. Indeed, many of the terms that we use for sensory qualities today are owed to the conceptual innovations of experts and scientists. In the case of vision, the word 'hue' was an Old English term revived by scientists in the nineteenth century to designate something specific and separate from 'colour'. Similarly, the word 'pitch' was first used in the sixteenth century by experts to designate the 'height of the tone' as it appears in musical annotation (Keller, 2016). As the scientific study of emotion is far younger than that of perception, it is not implausible to expect similar linguistic innovations to spread from the academy into popular vernacular. Indeed, if we are in the midst of this process, given the relatively recent explosion in theoretical and empirical work on emotion, it would be important and exciting for researchers trained in philosophy to contribute to it.

The thought here then is that, in so far as dimensions proposed by empirical work (such as appraisal dimensions) seem to add to our experiential vocabulary and help make sense of our experience from a first-person perspective, they can, and perhaps should, be adopted as genuine phenomenal dimensions.

5. Potential Applications

Why should we put in the work to construct an affective quality space? Many reasons. In the philosophy of emotion, the quality space may help distinguish emotions from other types of affective state.⁹ It is likely, for example, that moods can be plotted along a small subset of those dimensions needed to plot occurrent emotions. Additionally, the quality space may allow us to plot the dynamic structure of conscious emotions over time in an unprecedented manner, as experiences will leave a trajectory in the affective space as experience waxes, wanes, or evolves (Cochrane, 2009; Silva, 2022b). Crucially, the affective quality space may be able inform questions on how and what emotions represent. The quality space has an analogue structure characteristic of non-conceptual representational formats. This is because it is an informationally rich, fine-grained system of representation where representations vary continuously along dimensions (Cochrane, 2009; Maley, 2011; Silva, 2022b). This type of representation contrasts with conceptual representations typical of beliefs, which are coarse-grained and discrete. In occupying specific points in the quality space, emotions may represent evaluative properties non-conceptually.¹⁰

The thought is that emotions represent evaluative properties in a manner analogous to how perceptual experiences are thought to represent their objects; that is, without making use of concepts, or non-conceptually. Although the view that emotions represent evaluative properties non-conceptually is quite popular in the philosophy of emotion literature (Deonna and Teroni, 2012; Tappolet, 2020), specific accounts of this are largely lacking.¹¹ Most existing efforts

⁹ Cochrane (2009) makes a similar claim (see p. 397 and p. 407).

¹⁰ See Berger (2021; 2018) for functionalist and holistic versions of this claim, respectively, for perception.

¹¹ Note that many believe that, while certain emotion types represent non-conceptually (typically the more basic emotion types that we share with young children and non-human animals), other 'higher-order' emotions, such as guilt, involve conceptual

take a ‘negative’ form, in that they seek to establish that emotions do not make use of concepts in representing their objects (Döring, 2009; Tappolet, 2016). For example, emotions can be felt when relevant concepts are lacking, and they do not bear the types of inferential relations to each other nor to other mental states characteristic of beliefs, which are conceptually structured. One cannot infer sadness from nostalgia, nor from a belief about a loss, for example. Positive arguments that seek to establish how emotions might actually represent evaluative properties non-conceptually are hard to come by (see Tappolet, 2020, for an exception). The affective quality space could provide just the positive argument to this effect, modelling representations in terms of phenomenal qualities that vary along fine-grained dimensions, and occupying points in an affective space that resists easy conceptualization.

The affective quality space may also contribute to our understanding of the epistemology of emotions. First, by establishing that emotions have evaluative content many think necessary for their ability to provide justification for evaluative beliefs (Tappolet, 2016), including beliefs in political contexts where emotions bypass explicit conceptual reasoning (Jaggar, 1989; Silva, 2021). Additionally, the quality space may help map justificatory relations as we might think that clusters formed in the quality space are experiences that are all apt to justify beliefs that employ a relevant evaluative concept (the anger cluster of experiences is apt to justify beliefs about offence, for example). Other geometric properties of affective space may similarly prove to be of extreme explanatory power. For example, distances between different clusters might map logical or normative relations between affective states. That guilt and shame appear close together in affective space might indicate that objects that merit shame often merit also guilt.

In psychology, beyond contributing to the fine-tuning of models of affective space, the construction of an affective quality space is likely to open up a number of novel questions ripe for empirical investiga-

representations (Deonna and Teroni, 2012; Griffiths, 1997). For those that subscribe to such a view, the above statements would only apply to those emotions that do not have conceptual content. Although I do not have time to develop this further here, I find it a plausible view that while some emotions might presuppose concepts, and even incorporate conceptual representations, the manner in which evaluative properties *themselves* are represented, in emotions, is always non-conceptual, via their phenomenal qualities. Developing such a view is a topic for future work and adjudicating between it and alternative views would of course depend on cashing out what exactly is meant by ‘conceptual’, which varies in the work of different authors.

tion. First, experiments aimed at validating proposed dimensions should be devised. Second, experiments aimed at testing conclusions drawn for the proposed affective quality space will also be called for. For example, psychologists can empirically investigate whether emotion types that appear closer together in affective space co-occur more frequently. Crucially, psychologists can probe whether proximity along some dimensions is more predictive of co-occurrence than proximity along other dimensions (for example, we might expect that the control/coping dimension to predict co-occurrence more than valence). Empirical studies of this sort might be apt to inform arguments for the normative implications of the affective quality space.

6. Conclusion

In this paper, I have taken the initial steps towards the construction of an affective quality space. This project is surprisingly peripheral to contemporary emotion theorists' concerns. It is surprisingly peripheral given both the high explanatory pay-off such an affective quality space might yield, and the current state of the emotion literature, where perceptual theories, which push analogies between perceptions and emotions, are popular. The success of perceptual quality spaces should, I think, motivate us to construct an affective quality space, whether or not we are committed perceptualists regarding emotion. In fact, I do not think the affective quality space is the sole purview of the perceptualist, I merely think it would be an interesting and expected avenue for them to pursue given their commitments.

I have attempted to start to clear the theoretical terrain such that an affective quality space can begin to be constructed. I did so by addressing and dismissing four putative obstacles that such a project might face. I then outlined a number of conditions and methodological constraints that should be satisfied in the construction of the affective quality space. Most notably, the project emerges as a deeply interdisciplinary one, with vast potential applications in both philosophy and psychology. I have nowhere claimed that such a project will be easy, only that it is theoretically and practically feasible, and that it holds great explanatory promise.

References

- Berger, J. (2018) A defense of holistic representationalism, *Mind & Language*, **33** (2), pp. 161–176. doi: 10.1111/mila.12163

- Berger, J. (2021) Quality-space functionalism about color, *Journal of Philosophy*, **118** (3), pp. 138–164. doi: 10.5840/jphil2021118311
- Carretié, L. (2014) Exogenous (automatic) attention to emotional stimuli: A review, *Cognitive, Affective, & Behavioral Neuroscience*, **14** (4), pp. 1228–1258. doi: 10.3758/s13415-014-0270-2
- Carruthers, P. (2018) Valence and value, *Philosophy and Phenomenological Research*, **97** (3), pp. 658–680. doi: 10.1111/phpr.12395
- Chalmers, D.J. (1996) *The Conscious Mind: In Search of a Fundamental Theory*, New York: Oxford University Press.
- Clark, A. (1993) *Sensory Qualities*, Oxford: Oxford University Press.
- Cochrane, T. (2009) Eight dimensions for the emotions, *Social Science Information*, **48** (3), pp. 379–420. doi: 10.1177/0539018409106198
- Colombetti, G. (2005) Appraising valence, *Journal of Consciousness Studies*, **12** (8–10), pp. 8–10.
- Cowan, R. (2016) Epistemic perceptualism and neo-sentimentalist objections, *Canadian Journal of Philosophy*, **46** (1), pp. 59–81. doi: 10.1080/00455091.2015.1123037
- Deonna, J., Rodogno, R. & Teroni, F. (2011) *In Defense of Shame: The Faces of an Emotion*, Oxford: Oxford University Press.
- Deonna, J.A. & Teroni, F. (2012) *The Emotions: A Philosophical Introduction*, London: Routledge.
- Deonna, J.A. & Teroni, F. (2015) Emotions as attitudes, *Dialectica*, **69** (3), pp. 293–311. doi: 10.1111/1746-8361.12116
- Döring, S.A. (2007) Seeing what to do: Affective perception and rational motivation, *Dialectica*, **61** (3), pp. 363–394. doi: 10.1111/j.1746-8361.2007.01105.x
- Döring, S.A. (2009) The logic of emotional experience: Noninferentiality and the problem of conflict without contradiction, *Emotion Review*, **1** (3), pp. 240–247. doi: 10.1177/1754073909103592
- Ekman, P. & Friesen, W.V. (1971) Constants across cultures in the face and emotion, *Journal of Personality and Social Psychology*, **17** (2), pp. 124–129. doi: 10.1037/h0030377
- Ellsworth, P.C. & Scherer, K.R. (2003) Appraisal processes in emotion, in Davidson, R.J., Scherer, K.R. & Goldsmith, H.H. (eds.) *Handbook of Affective Sciences*, pp. 572–595, Oxford: Oxford University Press.
- Evans, D.A., Stempel, A.V., Vale, R. & Branco, T. (2019) Cognitive control of escape behaviour, *Trends in Cognitive Sciences*, **23** (4), pp. 334–348. doi: 10.1016/j.tics.2019.01.012
- Faucher, L. & Tappolet, C. (2002) Fear and the focus of attention, *Consciousness & Emotion*, **3** (2), pp. 105–144. doi: 10.1075/ce.3.2.02fa
- Fontaine, J.R.J., Scherer, K.R., Roesch, E.B. & Ellsworth, P.C. (2007) The world of emotions is not two-dimensional, *Psychological Science*, **18** (12), pp. 1050–1057. doi: 10.1111/j.1467-9280.2007.02024.x
- Fontaine, J.R.J. & Scherer, K.R. (2013) The global meaning structure of the emotion domain: Investigating the complementarity of multiple perspectives on meaning, in Fontaine, J.R.J., Scherer, K.R. & Soriano, C. (eds.) *Components of Emotional Meaning: A Sourcebook*, pp. 106–125, Oxford: Oxford University Press. doi: 10.1093/acprof:oso/9780199592746.003.0008
- Frijda, N. (2005) *Emotion Experience: Cognition and Emotion*, vol. 19, no. 4, Oxford: Taylor and Francis. doi: 10.1080/02699930441000346

- Gauker, C. (2017) Three kinds of nonconceptual seeing-as, *Review of Philosophy and Psychology*, **8** (4), pp. 763–779. doi: 10.1007/s13164-017-0339-2
- Gilbert, P., Pehl, J. & Allan, S. (1994) The phenomenology of shame and guilt: An empirical investigation, *The British Journal of Medical Psychology*, **67** (Pt 1), pp. 23–36. doi: 10.1111/j.2044-8341.1994.tb01768.x
- Grandin, T. (2018) *The Stories I Tell My Friends*, Arlington, TX: Future Horizons.
- Griffiths, P. (1997) *What Emotions Really Are: The Problem of Psychological Categories*, Chicago, IL: University of Chicago Press.
- Hardin, C.L. (1993) *Color for Philosophers*, Cambridge: Hackett.
- hooks, b. (1995) *Killing Rage: Ending Racism*, New York: H. Holt and Co.
- Izard, C.E. (1994) Innate and universal facial expressions: Evidence from developmental and cross-cultural research, *Psychological Bulletin*, **115** (2), pp. 288–299. doi: 10.1037/0033-2909.115.2.288
- Jackson, F. (1994) Armchair metaphysics, in Michael, M. & O’Leary-Hawthorne, J. (eds.) *Philosophy in Mind: The Place of Philosophy in the Study of Mind*, pp. 23–42, Dordrecht: Springer. doi: 10.1007/978-94-011-1008-2_3
- Jaggar, A. (1989) Love and knowledge: Emotion in feminist epistemology, *Inquiry*, **32** (2), pp. 151–176. doi: 10.1080/00201748908602185
- Keller, A. (2016) Perceptual quality space, in *Philosophy of Olfactory Perception*, pp. 7–38, Cham: Palgrave Macmillan. doi: 10.1007/978-3-319-33645-9_1
- Kuppens, P., Tuerlinckx, F., Russell, J.A. & Barrett, L.F. (2013) The relation between valence and arousal in subjective experience, *Psychological Bulletin*, **139**, pp. 917–940. doi: 10.1037/a0030811
- Lambie, J.A. & Marcel, A.J. (2002) Consciousness and the varieties of emotion experience: A theoretical framework, *Psychological Review*, **109** (2), pp. 219–259.
- Lewis, C.S. (2021) *A Grief Observed*, Hawthorne, CA: ValdeBooks.
- Maley, C.J. (2011) Analog and digital, continuous and discrete, *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, **155** (1), pp. 117–131.
- Mallon, R. & Stich, S.P. (2000) The odd couple: The compatibility of social construction and evolutionary psychology, *Philosophy of Science*, **67** (1), pp. 133–154. doi: 10.1086/392765
- Maund, B. (2006) The illusory theory of colours: An anti-realist theory, *Dialectica*, **60** (3), pp. 245–268.
- Moors, A. (2017) Appraisal theory of emotion, in Zeigler-Hill, V. & Shackelford, T. (eds.) *Encyclopedia of Personality and Individual Differences*, Cham: Springer. doi: 10.1007/978-3-319-28099-8_493-1
- Moors, A. & Scherer, K.R. (2013) The role of appraisal in emotion, in Robinson, M.D., Watkins, E. & Harmon-Jones, E. (eds.) *Handbook of Cognition and Emotion*, pp. 135–155, New York: The Guilford Press.
- Nayatani, Y. (1997) Simple estimation methods for the Helmholtz-Kohlrusch effect, *Color Research & Application*, **22** (6), pp. 385–401. doi: 10.1002/(SICI)1520-6378(199712)22:6<385::AID-COL6>3.0.CO;2-R
- Noah, S. & Mangun, G.R. (2020) Recent evidence that attention is necessary, but not sufficient, for conscious perception, *Annals of the New York Academy of Sciences*, **1464** (1), pp. 52–63. doi: 10.1111/nyas.14030
- Osgood, C.E. (1966) Dimensionality of the semantic space for communication via facial expressions, *Scandinavian Journal of Psychology*, **7** (1), pp. 1–30. doi: 10.1111/j.1467-9450.1966.tb01334.x

- Paramei, G.V., Schneider, K., Josephs, I. & Slusarek, M. (1994) *Identification of Emotional Meaning in Line Drawings and Faces*, unpublished report of the University of Bochum.
- Pendoley, K. (in progress) How to think about emotional feelings.
- Prinz, J.J. (2006). *Gut Reactions: A Perceptual Theory of Emotion*, New York: Oxford University Press.
- Proust, M. (2004) *Swann's Way: In Search of Lost Time, Volume 1*, Prendergast, C. (ed.), Davis, L. (trans.), revised ed., London: Penguin Classics.
- Quine, W.V.O. (1969) Natural kinds, in Rescher, N. (ed.) *Essays in Honor of Carl G. Hempel*, Dordrecht: D. Reidel.
- Roseman, I.J. (1991) Appraisal determinants of discrete emotions, *Cognition and Emotion*, **5** (3), pp. 161–200. doi: 10.1080/02699939108411034
- Rosenthal, D. (1991) The independence of consciousness and sensory quality, *Philosophical Issues*, **1**, pp. 15–36. doi: 10.2307/1522921
- Rosenthal, D. (2010) How to think about mental qualities, *Philosophical Issues*, **20** (1), pp. 368–393. doi: 10.1111/j.1533-6077.2010.00190.x
- Rosenthal, D. (2015) Quality Spaces and sensory modalities, in Coates, P. & Coleman, S. (eds.) *Phenomenal Qualities: Sense, Perception, and Consciousness*, Oxford: Oxford University Press. doi: 10.1093/acprof:oso/9780198712718.003.0002
- Rosenthal, D. (2016) Quality spaces, relocation, and grain, in O'Shea, J.R. (ed.) *Sellars and his Legacy*, Oxford: Oxford University Press. doi: 10.1093/acprof:oso/9780198766872.003.0009
- Russell, J.A. (2003) Core affect and the psychological construction of emotion, *Psychological Review*, **110** (1), pp. 145–172. doi: 10.1037/0033-295x.110.1.145
- Russell, J.A. & Lemay, G. (2000) Emotion concepts, in Lewis, M. & Haviland-Jones, J.M. (eds.) *Handbook of Emotions*, pp. 491–503, New York: Guilford Press.
- Sartre, J.-P. (1993) *Being and Nothingness*, reprint ed., New York: Washington Square Press.
- Scarantino, A. & de Sousa, R. (2018) Emotion, in Zalta, E.N. (ed.) *The Stanford Encyclopedia of Philosophy (Winter 2018 Edition)*, [Online], <https://plato.stanford.edu/archives/win2018/entries/emotion/>.
- Scherer, K., Dan, E. & Flykt, A. (2006) What determines a feeling's position in affective space? A case for appraisal, *Cognition and Emotion*, **20** (1), pp. 92–113. doi: 10.1080/02699930500305016
- Sellars, W., Brandom, R.B. & Rorty, R. (1997) *Empiricism and the Philosophy of Mind*, Cambridge, MA: Harvard University Press.
- Silva, L. (2021) The epistemic role of outlaw emotions, *Ergo*, **8**, art 23. doi: 10.3998/ergo.1160
- Silva, L. (2022a) On being angry at oneself, *Ratio*, **35** (3), pp. 236–244. doi: 10.1111/rati.12340
- Silva, L. (2022b) The ineffable as radical, in Deonna, J., Tappolet, C. & Teroni, F. (eds.) *A Tribute to Ronald de Sousa*, [Online], <https://www.unige.ch/cisa/related-sites/ronald-de-sousa/>.
- Sokolov, E.N. & Boucsein, W. (2000) A psychophysiological model of emotion space, *Integrative Physiological and Behavioral Science*, **35** (2), pp. 81–119. doi: 10.1007/BF02688770
- Solomon, R.C. (2006) *True to Our Feelings: What Our Emotions are Really Telling Us*, Oxford: Oxford University Press.

- Tappolet, C. (2016) *Emotions, Values, and Agency*, Oxford: Oxford University Press.
- Tappolet, C. (2020) Emotions inside out: The content of emotions, in Demmerling, C. & Schröder, D. (eds.) *Concepts in Thought, Action, and Emotion: New Essays*, New York: Routledge, [Online], <https://philpapers.org/rec/TAPEIO>.
- Teroni, F. (2021) Evaluative theories in psychology and philosophy of emotion, *Mind and Language*, **38** (1), pp. 81–97.
- Tracy, J.L. & Robins, R.W. (2006) Appraisal antecedents of shame and guilt: Support for a theoretical model, *Personality & Social Psychology Bulletin*, **32** (10), pp. 1339–1351. doi: 10.1177/0146167206290212
- Trnka, R., Lačev, A., Balcar, K., Kuška, M. & Tavel, P. (2016) Modeling semantic emotion space using a 3D hypercube-projection: An innovative analytical approach for the psychology of emotions, *Frontiers in Psychology*, **7**, [Online], <https://www.frontiersin.org/articles/10.3389/fpsyg.2016.00522>.
- Wollheim, R. (1999) *On the Emotions*, Newhaven, CT: Yale University Press.
- Wundt, W. (1912/1924) *An Introduction to Psychology*, Pintner, R. (trans.), London: Allen & Unwin.
- Yik, M.S.M., Russell, J.A. & Barrett, L.F. (1999) Structure of self-reported current affect: Integration and beyond, *Journal of Personality and Social Psychology*, **77** (3), pp. 600–619. doi: 10.1037/0022-3514.77.3.600
- Young, B., Keller, A. & Rosenthal, D. (2014) Quality-space theory in olfaction, *Frontiers in Psychology*, **5**, [Online], <https://www.frontiersin.org/articles/10.3389/fpsyg.2014.00001>.