

Emotionless animals?

Constructionist theories of emotion beyond the human case

Jonathan Birch

Centre for Philosophy of Natural and Social Science,
London School of Economics and Political Science.

[@birchlse](mailto:jbirchlse) | j.birch2@lse.ac.uk | <http://personal.lse.ac.uk/birchj1>

Abstract

Could emotions be a uniquely human phenomenon? One prominent theory in emotion science, Lisa Feldman Barrett's "Theory of Constructed Emotion" (TCE), suggests they might be. The source of the sceptical challenge is that TCE links emotions to abstract concepts tracking socio-normative expectations, and other animals are unlikely to have such concepts. Barrett's own response to the sceptical challenge is to relativize emotion to the perspective of an interpreter, but this is unpromising. A more promising response may be to amend the theory, dropping the commitment to the abstract nature of emotion concepts and allowing that, like olfactory concepts, they have disjunctive sensory groundings. Even if other animals were emotionless, this would not imply they lack morally significant interests. Unconceptualized valenced experiences are a sufficient basis for morally significant interests, and such experiences may occur even in the absence of discrete, constructed emotions.

Could emotions be a uniquely human phenomenon? On the face of it, the suggestion is wildly implausible. But we cannot leave things there. In the neuroscience of emotion, an important class of theories— *constructionist* theories—has risen to prominence. Arguably the highest profile theory in this family is Lisa Feldman Barrett’s “theory of constructed emotion” (TCE). And it is a theory that raises a serious possibility that emotions are indeed, despite appearances to the contrary, uniquely human.

My goals here are, firstly, to examine the complicated relationship between Barrett’s theory and emotions in other animals and, secondly, to reflect on the ethical implications of the debate. The discussion will involve elements of philosophy of science, philosophy of mind, and ethics, and it will be awash with “ifs”. My concern is with the consequences *if* Barrett’s theory is true or close to the truth. The theory, in my view, does not yet have compelling evidence in its favour, but discussion of its consequences for animals need not await a scientific consensus either for or against it.

To set up the debate, it will be necessary to locate constructionist theories in relation to their main opponent: the “basic emotions” family of theories. The clash between the two pictures is persistent and hard for any amount of empirical evidence to resolve. From here, I will zoom in on Barrett’s theory. I will aim to tease out (i) why the theory, on the face of it, makes the human uniqueness of emotions likely; (ii) how Barrett has attempted to resist this result, at least partially, by relativizing emotions to the perspective of an interpreter; and (iii) why this attempt is unconvincing.

I will then ask: are there small modifications to Barrett’s theory that would avoid the human uniqueness of emotion? I will contrast *abstract* and *grounded* versions of the theory. Barrett’s own version falls on the abstract side of this distinction, but the nearby grounded version, I will suggest, retains its main attractions while achieving compatibility with the existence of emotions in many other animals. On the grounded version of constructionism, there are strong analogies between emotion and olfaction: both processes build from real patterns of similarity and difference in the flow of sensory signals, despite a major role for cognitive interpretation in the building process. It is a view on which many other animals plausibly experience emotions, just as many other animals plausibly experience odours.

Finally, I will turn to another “if”: if the abstract version of the theory is true, what ethical implications would result? Are we staring down the barrel of a collapse of the field of animal ethics? I argue that, thankfully, we are not. The denial of emotions in other animals is compatible with a robust defence of the moral significance of their interests, provided they still have valenced experiences. A constructionist, crucially, can allow that valenced experiences are much more widespread in the animal kingdom than emotions.

1. Emotion and the brain: Two pictures

1.1 Basic emotions

The “basic emotions” picture (defenders of which have included Paul Ekman, Jaak Panksepp, Carroll Izard, and recently Mark Solms¹) is broadly supportive of ascribing human-like emotions to other animals, especially other mammals. The central idea is that basic emotions are *neurobiological natural kinds*: “joints in nature”, as philosophers like to say. There is no agreement among basic emotions theorists on a precise list of these kinds. Here I will take Jaak Panksepp’s version of the view as my focal example. Panksepp posited seven basic emotions that he called SEEKING, FEAR, RAGE, CARE, LUST, PANIC and PLAY. These line up quite well the extensions of the same words in ordinary English, but the capitals are intended to convey that they need not line up perfectly, since we are talking about evolutionarily ancient neurobiological mechanisms in the brain, not folk-psychological categories.

Panksepp’s view was that, for each of these basic emotions, there are distinctive brain mechanisms that generate them. He took these mechanisms to be located in the midbrain, an area of the brain at the top of the brainstem. He saw one particular region, the periaqueductal gray (PAG), as having special importance, calling it “the major epicenter for the generation of emotional feelings” (Panksepp 1998, p. 572). Moreover, each basic emotion was posited to have a characteristic behavioral expression that could be elicited by electrically stimulating the relevant brain mechanisms.

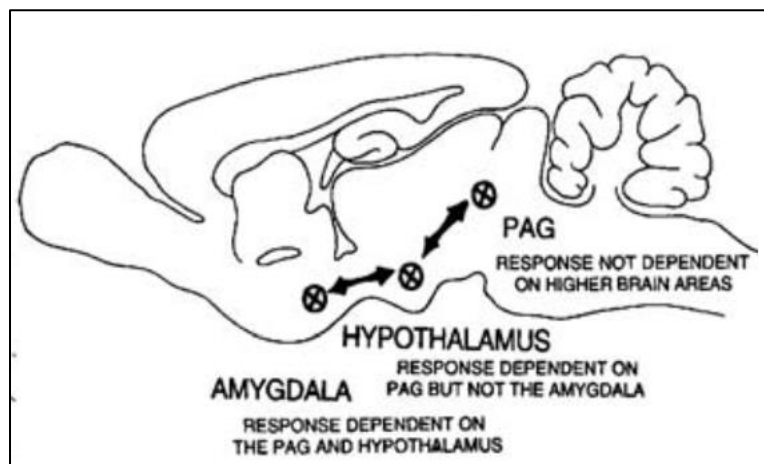


Figure 1. A drawing from Panksepp of the brain regions (in a generic, rat-like mammalian brain) primarily involved in generating ANGER. The image captures the core idea that, while other subcortical and cortical

¹ Ekman and Cordaro (2011); Panksepp (1998, 2005, 2007, 2011); Izard (2007); Solms (2021). Disagreements within this group are reviewed by Tracy and Randles (2011).

regions are influenced by PAG activity, the PAG itself is an initiator of emotion not significantly influenced by downstream brain mechanisms—hence Solms’s metaphor of the “hidden spring”.

Panksepp regarded the PAG as serving this emotion-generating function in all mammals, and, moreover, held that his seven basic emotion circuits were shared by all mammals, so that, for example, the PLAY circuits that drive rough-and-tumble play in rats are homologous with those driving the same behaviour in human children. He was heavily invested in the value of inferences across mammalian species: much of his experimental research consisted of stimulating brainstem regions in rats to produce behavioural responses, and he hoped this work would lead to the development of psychiatric therapies for humans (Panksepp 2016b).

Panksepp was clear that, in his view, these basic emotions “require no readout by a higher cognitive apparatus” in order to be experienced (2005, p. 64). He took them to be “re-represented” in the cerebral cortex, and thought this re-representation was critical in allowing our emotional experiences to integrate with learning and memory, but did not see this re-representation as necessary for basic emotional experiences (Panksepp 2011, p. 1795).

1.2 Constructionism

“Constructionist” theories of emotion take the “basic emotions” picture as their critical target. Here too, we are talking about a family of different theories, not a single theory. Prominent proponents of constructionism include Joseph E. LeDoux, James A. Russell, and Lisa Feldman Barrett, all of whom develop in different ways the idea that even apparently “basic” emotions are constructed by cognitive processing in the cerebral cortex rather than simply “upwelling” ready-made from a subcortical “spring” or “epicentre”.² They replace the image of a river flowing from the brainstem with that of a building site where interoception—monitoring of the internal milieu—supplies raw materials that must be sculpted and assembled for even the most basic emotional experiences. And they replace the metaphor of “joints in nature” with the metaphor of the “cookie cutter”: the raw materials can be assembled into emotions in many ways.

I will be focusing here on Lisa Feldman Barrett’s version of the view: the “theory of constructed emotion” (TCE) (Barrett 2017a, b). This is descended from an earlier version she called the “conceptual act theory” (Barrett 2014), and the old name still captures the theory’s most distinctive feature: the thesis that emotional experiences are constituted by the application of concepts to interoceptive signals. The thought is that emotions arise because the brain needs to predict, categorize and interpret the signals flowing in from the body and world and respond appropriately. These signals contain information about allostatic needs—predicted or realized

² Barrett (2017a, b); Barrett and Russell (2015); LeDoux and Pine (2016); LeDoux and Brown (2018); LeDoux (2019); Russell (2003, 2009).

departures from homeostasis requiring corrective action—together with information about their causes.

For example, an experience of fear involves applying a concept, “fear”, to make sense of the allostatic needs arising from your current internal milieu and the external context. The external context might include (in the cliché example) a bear running at you. Your internal milieu might include arousal of the autonomic nervous system, leading to increased heart rate and a cold sweat. For a constructionist, neither the external context nor the interoceptive signalling nor their combination suffices for an experience of fear. That experience comes only when you *conceptualize* this combination of internal milieu and external context as falling under the concept “fear”. An implication of this view is that if there is no concept, there can be no corresponding emotion. No concept of fear, no fear.

For Barrett, the processes that do this conceptualization occur primarily in the cerebral cortex. Moreover, the concepts and the corresponding emotional experiences are highly susceptible to cultural and linguistic influence. This is manifested, superficially, in the way emotion words vary so much across cultures. People also vary in their level of classificatory skill: what Barrett calls their emotional *granularity*, and this is posited to influence the character of their emotional experiences. Some people struggle to categorize their internal states at a fine grain. They can tell you if they're feeling *bad* or feeling *good*, and they can tell you if they're feeling *tired* or *energized*, but they struggle to make any more subtle discriminations. If you ask: “In what way are you feeling bad? Are you sad, anxious, depressed, angry, afraid?”, they will find it difficult to answer. Barrett contrasts this with people at the other extreme: people who are, in her memorable phrase, “sommeliers of emotion” with fine-grained classification skills.³ Such people will be able to tell with ease the difference between feeling anxious, depressed, weary or jaded, or the difference between feeling contented, thrilled, elated or ecstatic. It is an important part of this view that this alters the phenomenology, just as it would for a skilled perfumer or sommelier. A sommelier does not just have a purely cognitive skill that makes no experiential difference. They experience the wine differently, because they can attend sequentially to various notes that the novice will entirely miss (Barwich 2020). And so it is with emotional granularity, for Barrett. People with a high level of emotional granularity have different emotional experiences to those who can only make very crude discriminations.

It should be clear, given this, why Barrett is so attracted to the cookie cutter metaphor. The raw materials of internal milieu plus external context can be carved up in many ways, depending on our stock of emotion concepts, their granularity, and our skill at applying them. Your depression could be my anxiety—or my “feeling bad, not sure why”. And the emotions we feel, not just our reports of them, depend on how we conceptualize the raw dough.

³ Barrett (2017a, p. 108). The analogies between emotion and olfaction may be even richer than Barrett suggests—this theme is picked up in Section 5.

2. The current deadlock

The clash between the two pictures has proven very persistent. Studies of cultural variation in emotion concepts have revealed complicated patterns of similarity and difference: enough similarity to allow translation of concepts like “anger” and “fear” across more than two thousand languages, yet enough differences to create a strong suspicion that the conceptualization and phenomenology of these emotions varies cross-culturally (Jackson et al. 2019). This has failed to resolve the debate.

It may seem, on the face of it, that the existence of some emotion concepts with no direct translation favours constructionism. One often reads about words in other languages for emotions that have no English equivalent, such as *hygge* (Danish), the feeling of comfort and cosiness experienced when enjoying food or drink with friends or family, or *koi no yokan* (Japanese), ‘the feeling upon first meeting someone that you will inevitably fall in love with them’.⁴ However, the basic emotions picture can accommodate this variation, for it is only a theory of *basic* emotions. It has always been part of the picture that there are also *complex* emotions, formed from combinations of basic ingredients, just as a complex fragrance may be formed from a combination of more basic elements. For example, Jesse Prinz (2007, p. 67) has proposed that contempt is a blend of anger and disgust and that exhilaration is a blend of joy and fear. These combinations may vary across cultures. Moreover, display rules, which regulate the way emotions manifest behaviourally, can also vary cross-culturally. A culture may demand strong inhibition of fear (for example) in specific contexts, while encouraging its expression in others.

Some of the most persuasive evidence for constructionism comes from the phenomenon of emotional granularity, described earlier. But the basic emotions picture can accommodate this evidence too, because it can allow that people in vary in the skill with which they are able to parse complex blends of basic emotions. Just as people vary in their ability to pick out the ingredients of a perfume or a meal, they may vary in their ability to resolve the basic notes from which a complex emotion state is assembled.

Meanwhile, one might assume that any evidence of cross-cultural universals at the level of basic emotions would favour the basic emotions picture. However, constructionism can accommodate some degree of universality. For the constructionist can allow that people in different cultures often faces the same environmental challenges, leading to the same allostatic needs. This can lead to powerful attractors in the construction process: concepts on which many different cultures reliably converge (Lindquist et al. 2022). We should not be surprised, for example, to find that it is useful in many cultures to have a concept for a negatively valenced, high arousal

⁴<https://www.bbc.com/culture/article/20180103-the-untranslatable-japanese-phrase-that-predicts-love>

state elicited by threat and normally motivating heightened vigilance and/or escape—a concept that will be translated as “fear”.

Perhaps the strongest evidence for basic emotions comes from experimental investigations of subcortical circuits in other animals, mainly rats, identifying distinctive mechanisms credibly linked to basic emotions. Critics of the basic emotions picture, such as Barrett and LeDoux, have responded to this evidence in two main ways. In some cases, they question whether a distinctive circuit really exists: see Barrett et al. (2007, pp. 310-11) on PLAY and PANIC. In other cases, they accept the existence of the relevant circuits but regard them as insufficient for a conscious experience of emotion. For example, in rats, circuits controlling freeze responses (Panksepp’s FEAR), addiction and reward-seeking (SEEKING), and maternal (CARE) and sexual (LUST) behaviours have all been well documented, as have the circuits controlling attack responses in cats (RAGE). Indeed, in the case of the freeze response, LeDoux’s own lab has done much of the relevant work. As LeDoux exemplifies, it is possible to believe in the existence of these circuits without believing they correspond to basic emotions (LeDoux and Pine 2016; LeDoux and Brown 2017). Since other animals cannot direct report their emotions, it is always possible to doubt whether they have the same phenomenology we do. While Panksepp theorized that no “readout by higher cortical apparatus” was required to feel the emotion, constructionists theorize that cortical re-representation of the subcortical activity is needed. The clash reflects our deep ignorance about the neural basis of conscious experience more generally, and it is hard to imagine any decisive resolution while that wider question remains so mysterious (He 2023).

Given the current state of deadlock, we should reflect on the implications of *both* families of theories for other animals. If there were points of consensus, that would be reassuring. Unfortunately, the two pictures appear to have radically different implications beyond the human case. As already noted, the basic emotions picture is deeply committed to the universality of basic emotions across, at least, all mammals. The question concerning non-mammals (birds, reptiles, fish, invertebrates) becomes one not simply of whether they possess the PAG (all vertebrates do), but one of whether the PAG contains functionally similar circuitry. For Panksepp, this question could only be addressed through invasive neuroscientific research—research he was not able to complete himself, and that no one has taken up since—and so he was led to an agnostic attitude (Panksepp 2016a). By contrast, TCE is not committed to any human emotion being shared with any other animal. Indeed, for reasons discussed in the next section, the theory makes it quite unlikely that any other animal has emotions.

3. TCE and other animals: The sceptical challenge

The source of the sceptical challenge to emotions in other animals is *not* this: TCE makes emotions necessarily language-dependent, other animals lack language, so they lack emotions. It is not that simple. Barrett has been clear in recent work that there is no *necessary* link, in her view, between emotion and language. The challenge, rather, begins with this: emotion concepts

are, according to TCE, *abstract concepts* in the sense that there is *no common sensory grounding* that unites all of their instances.

Some concepts do have a concrete sensory grounding. Basic colour concepts, such as “red”, work this way. There are specific chromatic receptors in the human eye that are associated with sensations of red. The relationship is not straightforward, but, nonetheless, triggering these receptors in a normal human eye under normal conditions is important to an object’s being classified as red. And we can imagine a fictional creature whose emotion concepts worked in a similar way. They might have a “fear receptor”, an “anger receptor”, and so on, akin to taste buds, and readily form concepts for the classes of stimuli that trigger these receptors under normal conditions.

At the core of TCE, though, is the idea that emotion concepts are fundamentally *not* like this—they are *not* concepts that refer to the causes of a certain characteristic kind of stimulus. Instead, they are concepts that refer to diverse classes of combinations of internal milieu and external context. For example, ‘fear’ picks out a heterogeneous class of states with a very wide range of sensory triggers and physiological/behavioural manifestations: fleeing when one is being chased by a bear, hesitating when one is about to step out on to a comedy stage, freezing when reading an email with subject line “Your job application...”, refreshing the webpage containing the latest reports from some ongoing disaster on hearing a friend may be involved, and countless other responses.

The members of the ‘fear’ class may have nothing in common on the sensory side and very little in common on the behavioural side: not much, after all, unites seeing a bear, receiving a cue from a stagehand, reading an email, hearing bad news, and so on, in terms of the sensory processing involved. But they are united at a more abstract level by having a certain kind of significance in our social lives and our shared social world. We find it useful, when making sense of our own and each other’s behaviour, to group these states under a single concept, because of the abstract similarities these situations share: they are all cases of perceived threat leading to high physiological arousal and negative valence, normally motivating heightened vigilance and/or escape, where “normally” alludes to the existence of a wider community and a social norm regulating appropriate response. On TCE, the application of a common concept to these states produces common phenomenology—the phenomenology of fear.

The problem regarding other animals is now coming into view. On TCE, to acquire an emotion concept is to grasp an *abstract pattern of similarity* shared by a variety of scenarios that vary massively at the level of immediate sensation and behaviour. The pattern is to be found at the level of normative expectations regarding appropriate behavioural response. That is a hard concept to acquire. It requires sensitivity to subtle, abstract similarities between situations, plus membership of a social group with a norm-governed social life.

Here language becomes *contingently* relevant, because linguistic beings like us can scaffold the acquisition of such concepts through language: adults can help children acquire a concept of “fear” by pointing to various apt, situated responses and calling them all “fear”, just as they can point to a plastic card, a paper note, a metal coin and a number on a screen and call them all “money”. It is very difficult to imagine how humans might acquire an abstract concept like “money” without linguistic scaffolding, and the same goes, Barrett argues, for emotion concepts.

This is where Barrett is coming from when she writes “some animal brains, such as those of chimps, can categorize abstractly. But, to the best of our knowledge, only we have the wiring to compute abstractions of this magnitude” (Barrett 2022b; see also Barrett 2017a, chapter 12; 2022a). On the face of it, the “no concept, no emotion” thesis, when conjoined with the further thesis that the abstract and socio-normative nature of the relevant concepts puts them beyond the reach of any other animal, entails the human uniqueness of emotion.

4. The most direct way out: other animals as abstract, socio-normative thinkers

A way to confront the sceptical challenge head-on is to argue for a more optimistic picture of the abstractive and socio-normative abilities of other animals. After all, evidence one way or the other regarding these abilities is thin, leaving room for both optimistic and sceptical outlooks. In this vein, constructionist Eliza Bliss-Moreau (2017) has discussed the abstractive abilities of macaque monkeys and corvids as cases worthy of further investigation. We might even add bees to the list, given evidence that bees can succeed in tasks that, in humans, require recognition of abstract relations such as sameness and difference (Giurfa et al. 2001), probably using surprisingly simple neural mechanisms (Cope et al. 2018). Yet, as Bliss-Moreau notes, abstraction alone is not enough. An emotion concept, on TCE, involves recognition of a *special kind* of abstract pattern, namely a socially significant pattern of appropriate response. For emotions in other animals, then, we need not just abstract pattern recognition abilities, but also a capacity for acquiring social norms and judging some responses as appropriate, others not.

I see this head-on response as unpromising for three reasons: (i) on my own reading of the evidence, it is unlikely that even our fellow great apes have socio-normative abstractive abilities of the relevant kind (I think Barrett’s assessment that “only we have the wiring to compute abstractions of this magnitude” is more likely than not to be correct);⁵ (ii) I find it implausible at face value that the existence, or not, of emotions in other animals could turn on these questions—the cognitive requirements for emotion seem to be being ramped up to improbably high levels; and (iii) I think this route could only ever vindicate the idea of emotions in a small minority of cognitively sophisticated animals (even if this select group does turn out to include bees) so we have not escaped friction between TCE and our ordinary attributions of emotion to animals.

⁵ Birch (2021a, b). This is a point of agreement with Tomasello (2016, 2021) and Sterelny (2021a,b), but cf. Andrews (2020); Westra and Andrews (2021, 2022).

What other responses to the sceptical challenge are available? The next section considers Barrett's own response and finds it unsatisfactory. The subsequent section turns to a way of modifying TCE to weaken its commitment to the abstract nature of the relevant concepts.

5. Barrett's way out: The relativization of emotion

Barrett herself defends a version of TCE on which emotions are relativized to the perspective of an interpreter, writing that "a non-human animal's fearful state is real for human observers, but not necessarily for the creature itself" (Barrett 2022b; see also Barrett 2012, p. 422; 2017a, chapter 12). We apply emotion concepts such as fear to our own states, and we also apply those same concepts to states of others, including states of other animals. Those applications to other animals can be correct applications if they conform to (human) social norms of concept use. We are not wrong, on this view, to classify mice freezing in response to a threat as "afraid", as long as we understand them to be, strictly speaking, *afraid-from-our-human-perspective*, not *afraid-from-their-own-first-person-perspective*.

A serious problem for this relativist picture arises from the fact that emotions are by their nature *experienced by the subject who has them*: that is, they constitutively involve first-person phenomenology. This generates a dilemma: Is the mouse's being afraid-from-our-human-perspective something that involves fear phenomenology from the mouse's point of view—or not? If not, then the mouse simply cannot be instantiating fear. At most, it is experiencing a functionally fear-like state that lacks an essential feature of fear, namely its first-person phenomenology.

Yet if we say instead that our interpretation of the mouse as afraid *does* somehow induce fear phenomenology in the mouse, the picture is stranger still. For we know, in our own case, that emotion phenomenology does not work like this. If you interpret me as being afraid, there may be causal feedback from your interpretation to my experience (your interpretation may frighten me!), but my fear phenomenology does not *constitutively* depend on you interpreting me as being afraid. The neural basis of my fear phenomenology is in my brain, not yours. So, if we say that the mouse's fear phenomenology *does* constitutively depend on the interpretations of human observers, we are conjecturing that something manifestly false in our case is nonetheless true of other animals. We are thereby introducing a mysterious metaphysical chasm between the bases of human and non-human emotional experience. That is a very unpromising route to take.

So, I conclude that the proposed relativization of emotion facts to the perspective of an observer is untenable. The mouse either experiences fear or does not, and the mouse's own first-person perspective is what determines the matter. A constructionist should simply hold that, as in the human case, the facts depend on the mouse's *own* conceptualizations of its internal milieu and external context, not on those of outside observers. What human observers think may be a source

of causal feedback (how we feel may influence, causally, how the mouse feels) but there is no constitutive relationship here.

Unfortunately, this brings us back to square one: in so far as we accept both TCE and the claim that a mouse (or any other non-human animal) cannot form emotion concepts of its own, we should also accept that mice and other non-human animals lack emotions. The sceptical challenge remains unmet. Can we find another way out?

6. A better way out? From abstract to grounded construction

We can distinguish *abstract* and *grounded* versions of constructionism. The abstract version is Barrett's, and it involves the claim that the instances of an emotion concept like "fear" lack a common sensory grounding, their similarity becoming apparent only to creatures who can conceptualize abstract patterns. For fear as for money, there is only the cultural-linguistic fact that all the instances fall under the word "fear" and that we find it useful to classify them so, finding it worthwhile to group together situations in which we aptly respond to perceived threat with forms of heightened vigilance and escape. Fear thereby becomes part of "social reality" (Barrett 2012).

The grounded alternative retains the core idea—that emotion experiences involve the application of emotion concepts to make sense of signals from the internal milieu and external situation—but drops the commitment to the abstract, social nature of the patterns of similarity. For a grounded constructionist, the process of construction begins with recognizing, from the inside, subtle objective similarities in the neurobiological circuits being activated. It's just that, in a departure from the basic emotions picture, there are many patterns of similarity from which such processes could build.

If the basic emotions picture invites the metaphor of "carving nature at the joints" and the abstract constructionist invites the metaphor of the "cookie cutter", the grounded constructionist view can be seen as combining the two ideas: we carve at real neurobiological joints, but there are many cross-cutting joints we might follow, depending on the emotion concepts we have, and so the emotions we construct are still sensitive to those concepts.⁶

By way of comparison, I take grounded constructionism to be very plausible in relation to *olfaction*. Our odour/flavour concepts influence the phenomenology of olfactory experience: a skilled sommelier or perfumer has different phenomenology to a novice, their fine-grained concepts generating specific predictions that can be compared to the incoming stimulus in a very

⁶ The view resembles Dupré's (1993) "promiscuous realism" in the philosophy of science, which tries to find a middle path between "carving at the joints" and "cookie cutter" constructivism in a similar way. But we are thinking here of the process of psychological construction in an individual brain rather than the construction of taxonomies in science.

different way to the novice’s much coarser-grained predictions (Barwich 2020). Although the combination of odours into more complex ones undoubtedly occurs, a “basic odours” theory that tried to explain *all* olfactory experience as formed through the combination of 5-7 core odours would be rightly considered outlandish. At the same time, abstract constructionism about odour is also clearly false. It would be plainly wrong to suggest instances of an odour concept, like “lavender”, “vanilla” or “citrus”, have nothing in common *at all* at the level of the olfactory stimulus. They do (and perfumery is only possible because they do). It is just that the similarities from which the process builds are poorly understood, hard to detect from the outside, do not map on to obvious stimulus properties (since very similar molecular structures can produce very different odours, while very different structures can produce very similar odours) and are disjunctive (a matter of having this or this or this blend of molecules...), allowing different routes to qualitatively similar sensations.

The grounded constructionist about emotion proposes that, with emotions too, there are objective neurobiological similarities, and we latch on to them with concepts such as “fear”, “anger”, “panic” and “lust”. There are similar subcortical circuits being activated, as Panksepp proposed. This is still far from the “basic emotions” picture for five main reasons: (i) the application of emotion concepts still shapes the phenomenology—the felt experience of emotion; (ii) there are many real patterns of similarity, and we have concepts for only a small subset of them; (iii) we concede the point that we have not so far been able to identify these patterns *from the outside*, e.g. through neuroimaging, and should not expect to be able to do so easily; because (iv) the patterns may be quite disjunctive (this circuit or this or this...);⁷ and (v) social considerations may still play a role in determining which neurobiological similarities we find most salient and worth forming a concept of.

If there are real patterns of neurobiological similarity that unite the instances of emotion concepts, it would be no surprise to find that many other animals are able to acquire the relevant interoceptive categorization skills, just as they do with colours and odours. Some may even do it better than us, as we find with odour. Note, though, that grounded constructionism still calls for humility and caution concerning attributions of *human-like* emotions to other animals. For we remain largely in the dark as to what patterns of interoceptive signals other animals have concepts of, and so there is still a sceptical challenge to any suggestion that we can simply read off human-like emotions in other animals on the basis of their outward behaviour.

Other animals *may* have converged on similar emotion concepts (such as “fear” and “anger”), using them to pick out the same patterns of similarity we pick out. But there is no strong reason

⁷ Scarantino (2015) has argued that “basic emotions” can also be disjunctive at the level of neural mechanisms. The size of the gap between the traditional basic emotions picture and TCE can be reduced from both directions: Scarantino is closing the gap from the basic emotions side, whereas here I am reducing it from the constructionist side.

to expect this, and we should also expect a large amount of variation. Other animals will have whatever classification schemes are most useful for them. Again, the odour analogy pushes our intuitions in a good direction. Do dogs have concepts of “vanilla” or “lavender”? They may well do, but only if these concepts capture patterns of similarity that are significant for them. We cannot take this for granted. They may have found other ways more useful to them of carving up the space of odour patterns—and so it is with “fear” or “anger”.

Grounded constructionism is not Barrett’s TCE: it modifies the theory, dropping the commitment to the abstract nature of emotion concepts. So, I have not successfully reconciled TCE with emotions in other animals. But, in my view, the proposed modification is an improvement. The current empirical evidence does not force us all the way to the idea that emotion concepts are like “money”, referring to abstract features of social reality. We can hold on to the attractive features of constructionism while allowing that emotion concepts may resemble olfactory concepts like “citrus” more than they resemble “money”, latching on to real patterns in the relevant stimulus (in this case, interoceptive rather than olfactory) that we find significant. The limited abstractive abilities of other animals are then no longer an obstacle to the suggestion that they have rich emotional lives.

7. Implications for animal ethics

I do not rule out TCE in its abstract version. This could be true, and the current evidence is indecisive for the reasons reviewed in Section 2. Our intuitive attachment to the idea that some animals, especially our pets, have human-like emotions is not particularly strong evidence in favour of the reality of these emotions. We know that our mindreading abilities can easily misfire when used outside their core domain of application (Mameli and Bortolotti 2006). Even on the grounded version of constructionism I prefer, the emotions of other animals will not be very *human-like* unless they have converged on a similar set of emotion concepts.

We should ask, then: what would the ethical implications be (a) if emotions in general turn out to be uniquely human, or (b) if specific emotional kinds like fear and anger turn out to be uniquely human? Would this threaten the idea that other animals have morally significant interests?

7.1 Valenced experience, not emotion, is the source of morally significant interests

It is important to distinguish emotion from a broader category of *valenced experiences*: experiences that feel bad or feel good to the subject who has them.⁸ For most positions in

⁸ I set aside here a different way of thinking about “valence” that detaches it from felt hedonic quality, instead defining it as the nonconceptual representation of value, whether conscious or unconscious (Carruthers 2017). These senses of “valence” can coexist peacefully, in my opinion. Viola (2017) has expressed doubt about whether “valence” is any less constructed, or any more projectible to other animals, than emotion—but that is a topic for another occasion. Here my point is simply that TCE does not rule out valenced experience being widespread.

contemporary animal ethics, this broader category is the source of morally significant interests, not emotions as such.

This idea that valenced experience has a strong and special ethical significance can be reached from a variety of ethical starting points. One is classical (hedonic) utilitarianism, on which goodness is happiness, and happiness is understood as a state of pleasure and the absence of pain, where both ‘pleasure’ and ‘pain’ are interpreted broadly, capturing what many would now call positively and negatively valenced experiences, respectively (Driver 2022). The valence of an experience is crucially important, for it determines whether that experience contributes positively or negatively to the calculation that determines the right action.

A second route runs through Christine Korsgaard’s neo-Kantian ethics, on which we have a fundamental obligation to respect all sentient beings as “ends-in-themselves” and to avoid using them as “mere means” to our own ends (Korsgaard 2018). The terminology comes from Kant, but the picture is vastly more inclusive of other animals than Kant’s own ethics. Respecting animals as ends-in-themselves requires treating as valuable those goals that animals themselves value. For Korsgaard, it is through having valenced experiences that animals are able to value and disvalue ends in the sense that matters ethically.

A third route, different again, runs through Martha Nussbaum’s (2023) theory of justice. For Nussbaum, the fundamental imperative of justice is to help all beings capable of “significant striving” to achieve a kind of flourishing appropriate to their form of life, without being “wrongfully thwarted” by others. But which beings can be “wrongfully thwarted”? The view initially sounds like it will be radically inclusive of all living things, since one can talk even of plants and bacteria striving. In fact, Nussbaum sees only beings with “a felt orientation towards what is seen as good and a felt aversion to what is seen as bad” (2023, p. 119) as capable of being treated justly or unjustly. A plant can be thwarted but not wrongfully thwarted, because the frustration of its goal-directed behaviours will not register unpleasantly in its subjective point of view on the world. Here too, valenced experiences are what matters.

While far from exhaustive, this brief survey is enough to show that valenced experience is widely seen as implying morally significant interests. Crucially, none of the three lines of argument just considered requires that the animal experiences *emotions*, let alone human-like emotions. *Any* valenced experience is enough.

TCE, even in its abstract version, allows that there can be valenced experiences in beings who lack emotion concepts. These experiences will not be *emotions*, but they will still have hedonic quality: they will feel bad or feel good. It is very hard for us to imagine such experiences, since we are constantly engaged in conceptualization, and this (according to TCE) transforms the phenomenology. We try to imagine “feeling bad without applying any emotion concept” and we

just end up imagining ourselves applying the concept of “feeling bad”. But this imaginative limitation is no reason to think that such experiences of raw, uninterpreted valence cannot happen in other animals, any more than our inability to imagine what it’s like to echolocate is a reason to think there is nothing it’s like to be a bat.⁹

Admittedly, TCE might be combined with a cognitively demanding theory of conscious experience, such as a higher-order theory, and the combined package might end up ruling out raw valenced experiences in a wide range of animals, depending on the details.¹⁰ But it would not be TCE driving the result. TCE is compatible with the presence of raw valenced experience in wide range of animals and will imply this if combined with a relatively inclusive picture of the basic system requirements for conscious experience and valence.

7.2 Emotion is of pragmatic but not (fundamental) ethical importance

Residual unease may remain. Might raw valenced experience confer a kind of second-tier moral significance, lesser than that conferred by full-blown emotions? I see no reason to think this.

Admittedly, it would be *useful* if human emotion categories turned out to be widely applicable to other animals, because it would help us design appropriate interventions to improve welfare. This is because appropriate interventions depend on the finer details of a valenced state, not just its valence. Good interventions to alleviate fear are not the same as interventions to alleviate stress, boredom, discomfort or pain, even though all of these states have negative valence. Agnosticism about whether our human emotion concepts are applicable threatens indecision about what type of welfare interventions are appropriate.¹¹

Yet the *usefulness* of being able to apply human emotion categories to animals is no reason to think the *moral significance* of their interests is elevated if these categories apply. On all three of the theoretical standpoints mentioned earlier, valenced experience alone grounds interests, and the moral weight of the interest is not amplified by virtue of falling under a human emotion concept—or any emotion concept.

⁹ Nagel (1974). Schwitzgebel (2023) has recently made a similar point in relation to borderline cases of conscious experience.

¹⁰ LeDoux has combined a constructionist view of emotion with a higher-order theory of consciousness (LeDoux and Pine 2016; LeDoux and Brown 2017). Yet his picture still allows for the possibility of unconceptualized, “anoetic” consciousness in many other animals, probably including all mammals and perhaps other vertebrates too (LeDoux et al. 2023). The view is much less hostile to valenced experience in other animals than may at first seem to be the case.

¹¹ Perhaps even the pragmatic significance of emotion is overstated. One mainstream position in animal welfare science is that interventions should be designed based on an animal’s health needs and preferences, not on attempts to infer the underlying emotion from those health needs and preferences (Dawkins 2019).

7.3 Unconceptualized experiences need not be less “rich”

I have argued elsewhere that, although it is a mistake to think we could ever classify species of animal as “more or less conscious/sentient” on a single scale, it may be possible to make comparisons along the dimension of *evaluative richness*: the richness of an animal’s valenced experiences (Birch et al. 2020). Against this background, a concern arises that, if we can attribute only unconceptualized valenced experience to other animals, we seem to be leaving them with an impoverished form of sentience, much less rich on the evaluative side than the human form. That might result in them counting for less on a framework that ties moral weight to the richness of valenced experience.

There are several components to the “richness” of valenced experience, as Birch et al. (2020) noted, and they have the potential to come apart. One thing we may mean by richness is *categorization power*: the fineness of grain with which an animal categorizes stimuli in a particular modality, as reflected at the level of experience. This is the sense of “richness” at issue when we talk about the granularity of olfactory or emotional experience and say that the sommelier has “richer” experiences than the novice. That is a legitimate sense of the word, and an animal with no classificatory ability at all in the relevant domain will be at zero for this component of richness.

However, three other aspects of richness have no necessary relationship to categorization power. One is *acuity*: the fineness of the just-noticeable differences to which we are sensitive in a particular modality. We may well surpass falcons for categorization power in the visual domain, but they surpass us in acuity. In the olfactory domain too, acuity and categorization power seem to have little to do with each other. A person can notice tiny differences in an odour (“that’s changed”), and yet struggle to categorize the odour or its elements (Barwich 2020). An analogous point applies regarding emotions. Even supposing other animals are poor categorizers of interoceptive signals and their external contexts, they might nonetheless notice very fine differences (i.e. this feeling just got slightly worse/better).

A third aspect of richness is *bandwidth*. Bandwidth refers to the animal’s ability to experience multiple valenced contents simultaneously. We humans, for example, can simultaneously experience the pain of exercising hard and the pleasure of enduring the pain, leading to an experience Colin Klein (2014) has called a feeling of penumbrality: pushing oneself to the edge. Whether one is a basic emotions theorist (who sees complex experiences resulting from blending of basic emotions) or a constructionist (who sees them resulting from the skilful interpretation of an internal state in an external context), it can be a point of agreement that complex emotions rest on our ability to process multiple valenced states at once, sometimes including states with opposite valence. Other animals may lack our conceptual apparatus and yet still possess significant processing bandwidth. This is likely to vary a great deal across species depending on their general information processing capabilities.

A fourth aspect is *maximum intensity*. Whereas acuity concerns the fineness of discriminations being made, this dimension aims to capture the length of the ruler. It is conceivable that some other animals experience far greater intensities of negative and positive valence than we do. The reverse is also possible: perhaps some animals experience the most serious injuries as mild irritations, feeling them with an intensity akin to a heavily sedated human rather than a fully awake one. This is an area of massive unknowns, with no consensus about how to probe intensity experimentally in other animals (Browning 2023). Importantly, though, intensity need not be tied to categorization power. Raw, uninterpreted valenced experiences could be more intense than constructed emotions, less intense, about the same, or there might be no systematic relationship at all. All options are conceivable.

In sum, the valenced experiences of a being without constructed emotions will be less rich only if we privilege one specific meaning of “rich”—categorization power—while neglecting at least three other meanings. We should not neglect these other meanings. Their relative ethical significance could be a useful topic of further work in animal ethics. On classical utilitarianism, at least, maximum intensity—the most inscrutable of the four—is arguably the one that matters most.

Acknowledgements

My thanks to audiences at the inaugural meeting of the Philosophy of Animal Minds and Behaviour Association (PAMBA) in Madrid (April 2023), the Philosophy in Biology and Medicine meeting in Cambridge (May 2023), and the University of Minnesota (October 2023) for their valuable comments, questions, and encouragement. Special thanks to: Kristin Andrews, Colin Allen, Lisa Feldman Barrett, Aliya Dewey, Joseph LeDoux, Camilo Martinez, Susana Monsó, Victor Carranza Pinedo, Adina Roskies and Antonella Tramacere.

References

- Andrews, Kristin 2020: ‘Naïve Normativity: The Social Foundation of Moral Cognition’. *Journal of the American Philosophical Association*, 6(1), pp. 36-56.
- Birch, Jonathan 2021a: ‘Toolmaking and the Evolution of Normative Cognition’. *Biology & Philosophy*, 35, 4.
- Birch, Jonathan 2021b: ‘The Skilful Origins of Human Normative Cognition’. *Analyse & Kritik*, 43(1), pp. 191-201.
- Birch, Jonathan, Alexandra K. Schnell, and Nicola S. Clayton 2020: ‘Dimensions of Animal Consciousness’. *Trends in Cognitive Sciences*, 24(10), pp. 789-801.
- Dawkins, Marian Stamp 2019: *The Science of Animal Welfare: Understanding What Animals Want*. Oxford: Oxford University Press.
- Driver, Julia 2022: ‘The History of Utilitarianism’. In Edward N. Zalta and Uri Nodelman (eds.), *The Stanford Encyclopedia of Philosophy* (Winter 2022 Edition).

- Ekman, Paul 1992: 'An Argument for Basic Emotions.' *Cognition and Emotion*, 6, pp. 169-200.
- Barrett, Lisa Feldman 2012: 'Emotions Are Real'. *Emotion*, 12(3), pp. 145-172.
- Barrett, Lisa Feldman 2014: 'The Conceptual Act Theory: A Précis'. *Emotion Review*, 6(4), pp. 292-297.
- Barrett, Lisa Feldman 2017a: *How Emotions Are Made: The Secret Life of the Brain*. London: Macmillan.
- Barrett, Lisa Feldman 2017b: 'The Theory of Constructed Emotion: An Active Inference Account of Interoception and Categorization'. *Social Cognitive and Affective Neuroscience*, 12(1), pp. 1-13.
- Barrett, Lisa Feldman 2022a: 'Context Reconsidered: Complex Signal Ensembles, Relational Meaning, and Population Thinking in Psychological Science'. *American Psychologist* 77(8), pp. 894–920.
- Barrett, Lisa Feldman 2022b: 'The Big Idea: Do Animals Have Emotions?' *The Guardian*, 29 August 2022.
- Barrett, Lisa Feldman, Kristen A. Lindquist, Eliza Bliss-Moreau, Seth Duncan, Maria Gendron, Jennifer Mize and Lauren Brennan 2007: 'Of Mice and Men: Natural Kinds of Emotions in the Mammalian Brain? A Response to Panksepp and Izard'. *Perspectives on Psychological Science*, 2(3), pp. 297-312.
- Barrett, Lisa Feldman, and James A. Russell (eds.) 2015: *The Psychological Construction of Emotion*. New York: Guilford Press.
- Barwich, A. S. 2020: *Smellosophy: What the Nose Tells the Mind*. Cambridge, MA: Harvard University Press.
- Bliss-Moreau, Eliza 2017: 'Constructing Nonhuman Animal Emotion'. *Current Opinion in Psychology*, 17, pp. 184-188.
- Browning, Heather 2023: 'Welfare Comparisons Within and Across Species'. *Philosophical Studies*, 180, 529-551.
- Carruthers, Peter 2017: 'Valence and Value.' *Philosophy and Phenomenological Research*, 97(3), pp. 658-680.
- Cope, Alex J., Eleni Vasilaki, Dorian Minors, Chelsea Sabo, James A. R. Marshall, and Andrew Barron 2018: 'Abstract Concept Learning in A Simple Neural Network Inspired by the Insect Brain'. *PLoS Computational Biology*, 14(9), e1006435.
- Dupré, John 1993: *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Cambridge, MA: Harvard University Press.
- Ekman, Paul and Daniel Cordaro 2011: 'What is Meant by Calling Emotions Basic?' *Emotion Review*, 3(4), pp. 364-370.
- Giurfa, Martin, Shaowu Zhang, Arnim Jenett, Randolph Menzel, and Mandyam Srinivasan 2001: 'The Concepts of "Sameness" and "Difference" in an Insect'. *Nature*, 410(6831), pp. 930-933.
- He, Biyu J. 2023: 'Next Frontiers in Consciousness Research'. *Neuron*, 111(10), pp. 3150-3153.

- Izard, Carroll 2007: 'Basic Emotions, Natural Kinds, Emotion Schemas, and A New Paradigm.' *Perspectives in Psychological Science*, 2, 260-280.
- Jackson, Joshua Conrad, Joseph Watts, Teague R. Henry, Johann-Mattis List, Robert Forkel, Peter J. Mucha, Simon J. Greenhill, Russell D. Gray, and Kristen A. Lindquist 2019: 'Emotion Semantics Show Both Cultural Variation and Universal Structure.' *Science* 366(6472), pp. 1517-1522.
- Klein, Colin 2014: 'The Penumbra Theory of Masochistic Pleasure'. *Review of Philosophy and Psychology*, 5(1), pp. 41-55.
- Korsgaard, Christine M. 2018: *Fellow Creatures: Our Obligations to the Other Animals*. Oxford: Oxford University Press.
- LeDoux, Joseph E. 2019: *The Deep History of Ourselves: The Four-Billion-Year Story of How We Got Conscious Brains*. New York: Viking Press.
- LeDoux, Joseph E., Jonathan Birch, Kristin Andrews, et al. 2023: 'Consciousness Beyond the Human Case'. *Current Biology*, 33, R829-R854.
- LeDoux, Joseph E. and Richard Brown 2017: 'A Higher-Order Theory of Emotional Consciousness'. *Proceedings of the National Academy of Sciences of the United States of America*, 114(10), E2016-E2025.
- LeDoux, Joseph E. and Daniel S. Pine 2016: 'Using Neuroscience to Help Understand Fear and Anxiety: A Two-System Framework'. *American Journal of Psychiatry*, 173(11), 1083-1093.
- Lindquist, Kristen A., Joshua Conrad Jackson, Joseph Leshin, Ajay B. Satpute, and Maria Gendron 2022: 'The Cultural Evolution of Emotion'. *Nature Reviews Psychology*, 1, 669-681.
- Mameli, Matteo and Lisa Bortolotti 2006: 'Animal Rights, Animal Minds, and Human Mindreading'. *Journal of Medical Ethics*, 32(2), pp. 84-89.
- Nagel, Thomas 1974: 'What Is It Like to Be a Bat?' *Philosophical Review*, 83(4), pp. 435-450.
- Nussbaum, Martha C. 2023: *Justice for Animals: Our Collective Responsibility*. New York: Simon & Schuster.
- Panksepp, Jaak 1998: *Affective Neuroscience: The Foundations of Human and Animal Emotions*. New York: Oxford University Press.
- Panksepp, Jaak 2005: 'Affective Consciousness: Core Emotional Feelings in Animals and Humans'. *Consciousness and Cognition*, 14, pp. 30-80.
- Panksepp, Jaak 2007: 'Neurologizing the psychology of affects: How appraisal-based constructivism and basic emotion theory can coexist'. *Perspectives in Psychological Science*, 2, pp. 281-296.
- Panksepp, Jaak 2011: 'The Basic Emotional Circuits of Mammalian Brains: Do Animals Have Affective Lives?' *Neuroscience and Biobehavioral Reviews*, 35, pp. 1791-1804.
- Panksepp, Jaak 2016a: 'Brain Processes for "Good" and "Bad" Feelings: How Far Back in Evolution?'. *Animal Sentience* 1(3), 24.

- Panksepp, Jaak 2016b. 'The Cross-Mammalian Neurophenomenology of Primal Emotional Affects: From Animal Feelings to Human Therapeutics.' *Journal of Comparative Neurology*, 524(8), pp. 1624-35.
- Prinz, Jesse J. 2007. *The Emotional Construction of Morals*. Oxford: Oxford University Press.
- Russell, James A. 2003: 'Core Affect and the Psychological Construction of Emotion'. *Psychological Review*, 110(1), 145-172.
- Russell, James A. 2009: 'Emotion, Core affect, and Psychological Construction'. *Cognition and Emotion*, 23(7), 1259-128.
- Scarantino, Andrea 2015: 'Basic Emotions, Psychological Construction, and the Problem of Variability.' In Lisa Feldman Barrett and James A. Russell (eds.), *The Psychological Construction of Emotion*. New York: Guilford Press, pp. 334–376.
- Schwitzgebel, Eric 2023: 'Borderline Consciousness, When It's Neither Determinately True nor Determinately False that Experience is Present'. *Philosophical Studies*. Advance online publication.
- Solms, Mark 2021: *The Hidden Spring: A Journey to the Source of Consciousness*. London: Profile Books.
- Sterelny, Kim 2021a: *The Pleistocene Social Contract: Culture and Cooperation in Human Evolution*. New York: Oxford University Press.
- Sterelny, Kim 2021b: 'The Skill Hypothesis: A Variant'. *Analyse & Kritik*, 43(1), pp. 225–234.
- Tomasello, Michael 2018: *A Natural History of Human Morality*. Cambridge, MA: Harvard University Press.
- Tomasello, Michael 2021: 'Norms Require Not Just Technical Skill and Social Learning, but Real Cooperation.' *Analyse & Kritik*, 43(1), pp. 219–223.
- Tracy, Jessica L., and Daniel Randles. 2011: 'Four Models of Basic Emotions: A Review of Ekman and Cordaro, Izard, Levenson, and Panksepp and Watt.' *Emotion Review*, 3(4), pp. 397–405.
- Viola, Mark 2017: 'Commentary: Constructing Nonhuman Animal Emotion', *Frontiers in Psychology*, 8, 2070.
- Westra, Evan, and Kristin Andrews 2021: 'If Skill is Normative, Then Norms are Everywhere.' *Analyse & Kritik*, 43(1), 203-218.
- Westra, Evan, and Kristin Andrews 2022: 'A Pluralistic Framework for the Psychology of Norms.' *Biology & Philosophy*, 37, 40.