

Understanding the multidimensionality of sentience in interspecies welfare comparisons

Víctor Carranza-Pinedo
(draft – May 2024)

Abstract

Are some organisms more sentient than others? Recent attention within animal welfare research centres around which and how much evidence is sufficient to ascertain whether a species' members are sentient. However, as more species are recognised as potentially sentient, a pressing issue arises in policymaking: should all sentient species be regarded as sentient to the same extent? While a degreed notion of sentience has been criticised as conceptually implausible or ethically problematic, this paper argues that these objections are flawed. By employing formal semantic tools, this paper proposes a delineation of the multidimensional structure of sentience that can serve as the basis for a framework for responsibly comparing degrees of sentience across species. The framework proposed underscores that the current debate regarding cross-species comparisons will only progress through an overall understanding of the different commitments that achieving welfare comparisons involves within the science-policy interface.

Keywords. Sentience – Consciousness – Animal Welfare – Animal Emotion

1. Introduction

In animal welfare and ethics research, sentience is standardly considered to ground the moral status of species. In a few terms, organisms with an interest—an interest in avoiding suffering—are deemed to deserve moral consideration (Singer, 1977). What is, however, sentience? In a “broad” sense, it denotes the capacity to undergo subjective conscious states, also called “phenomenal consciousness”. In this sense, an organism has conscious experiences when “there is something it is like to be that organism” (Nagel, 1974). In another, “strict” sense, sentience denotes the capacity to undergo valenced or “hedonic” states, such as pleasure or pain. This latter notion of sentience, also called “affective sentience” (Powell & Mikhalevich, 2021), holds particular ethical relevance (e.g., Duncan, 2002; Mellor, 2016; Browning, 2023) and hence will be the focus of this paper.

Contemporary discussions in animal welfare research revolve around which and how much evidence is sufficient to determine whether members of a particular species are sentient (Browning & Birch, 2022). However, our actions and policies, including the increasing acknowledgement of species as sentient, may negatively affect other sentient species in expectation (Fischer & Sebo, 2023). For instance, in biomedical research, marine safety tests employ decapod crustaceans instead of fish based on extant evidence that fish are sentient (reviewed in Brown, 2015). Nevertheless, new findings suggesting that decapod crustaceans are also potentially sentient (Birch et al., 2021) may introduce a “perverse incentive” to employ either crustaceans or fish in marine safety tests. As Birch (2017) suggests, one way to prevent this situation is to argue that fish have a more “complex” form of sentience than decapod crustaceans and should, therefore, be allocated more protection.

The foundation of these comparative assessments lies in the presumption that sentience comes in varying degrees. By making such assessments, one assumes it is possible to locate members

of distinct species on a unified or “overall” scale based on their capacity to experience a broader or narrower range of hedonic states relative to others (Browning, 2023; Schukraft, 2020; Višak, 2017). Given that there is strong evidence of sentience in mammals and birds, as well increasing evidence of the realistic possibility that invertebrates, including insects, are sentient (cf. Barron & Klein 2016, and the New York Declaration on Animal Consciousness for an overview of the current state of the research), it is a priority to work on the theoretical foundations of responsible cross-species sentience comparisons and the manifold challenges these give rise to.

There are at least three domains regarding cross-species comparisons that are currently being explored. One of them is the distribution of sentience across living organisms. Višak (2022) argues from an evolutionary perspective that all animals have an equal hedonic capacity, while Veit (2023) argues that differences in life history can track phenomenological differences. Another topic is the empirical grounds for making sentience comparisons in a principled way. Browning (2023) argues that key similarity assumptions (e.g., similarity in the physiological responses) can form the basis for sentience comparisons, whereas Gaffney et al. (2023) emphasise the use of a variety of empirical proxies to measure degrees of sentience. Lastly, a question that has also received attention is how to allocate resources given the uncertainty regarding sentience involved. Sebo (2018), for instance, proposes using principles of risk to make welfare estimates under uncertainty, while Budolfson & Spears (2019) adapt formal tools from economic policy analysis.

However, a relatively neglected issue in these discussions concerns sentience’s multidimensional character. Sentience, it is said, can vary along different dimensions, such as duration, intensity, or salience. Accordingly, comparing the sentience of different species requires specifying the specific dimension on which the comparison is based. Hypothetically, fish may exhibit greater sentience than crustaceans regarding intensity, while crustaceans may exhibit more sentience than fish regarding salience. Therefore, follows the argument, it cannot be claimed that all-things-considered fish are more sentient than crustaceans or vice versa (Bayne et al., 2016; Birch et al., 2020; Carruthers, 2019). A graded concept of sentience would, at best, lack conceptual soundness and, at worst, provide an unprincipled basis for policy-making decisions.

The problem of multidimensionality, hence, is two-fold: on the one hand, it is unclear whether sentience’s multidimensional structure allows for overall comparisons across species; on the other, even assuming it does, it remains unclear whether such comparisons can be established non-arbitrarily. In this article, I have two main objectives. The first is to analyse the concept of sentience, formally distinguishing sentience’s dimensionality from other related features, such as degrees and thresholds of sentience. Drawing upon D’Ambrosio & Hedden’s (2023) semantics of multidimensional expressions, I argue that sentience’s multidimensionality does allow for overall comparisons.

The second objective is to employ the proposed formal analysis to outline the milestones to achieve responsible comparisons of sentience across species. This roadmap includes the following steps: (i) determining sentience’s dimensions, (ii) calculating partial orderings based on these dimensions, (iii) merging those partial orderings into an overall ordering, and (iv) developing and applying a decision-making formula that takes the overall ordering as input. By specifying the more specific challenges and commitments each step involves, I show how formally clarifying the concept of sentience can enhance progress in the animal welfare science-policy interface.

The paper is organised as follows: Section 2 analyses in more detail what is meant by degrees and dimensions of sentience and the worries regarding overall comparisons. Section 3 addresses the formal issues by characterising sentience's conceptual structure. Section 4 addresses empirical issues by outlining the main steps required for making informed decisions regarding sentience and highlights the primary sources of uncertainty in this process. Section 5 briefly reflects on potential objections to the proposal and alternative perspectives regarding cross-species sentience comparisons.

2. Background: trade-offs, gradability, and multidimensionality

Several indicators of sentience have been proposed, prompting debates on their implications for animal protection legislation (Birch, 2017). For instance, motivational trade-off behaviour, whereby an animal behaves as if weighing its preference to avoid a noxious stimulus against other preferences, has been deemed a reliable indicator of sentience (Sneddon et al., 2014). The ensuing discussion revolves around whether the presence of such behaviour is sufficient to count members of a species as sentient or if additional proof is required. Hence, establishing a threshold of evidence becomes crucial for confidently attributing the capacity for sentience to a particular animal species and, thereby, including it in our moral circle.

However, establishing an evidence threshold is closely connected to another significant concern. Namely, to the implications of increasingly bringing different species under the 'sentient organism' category. Given that resources (e.g., time, attention, money, etc.) are limited, distributing them in a way that provides the most significant welfare increase requires trade-offs among sentient species. These decisions arise in various contexts, including local management decisions such as prioritising the welfare of one sentient individual over others within a zoo (Browning, 2023), analysing optimal philanthropic investment (Gaffney et al., 2023), and ultimately shaping policy design (Budolfson & Spears, 2019).

As mentioned above, tests aimed at detecting harmful chemicals in marine life may opt to employ crustaceans instead of fish, based on Russell & Burch's (1959) principle of replacing sentient with non-sentient animals in conducting experiments. However, given the increasing evidence that crustaceans are also sentience, how can these and other types of trade-offs be determined in a non-arbitrary, principled way? One proposed solution is to acknowledge that the capacity for sentience is realised differently across species, meaning that species' capacity to undergo (un)pleasant states comes in degrees. Considering extant evidence, we may grant that crustaceans and fish are sentient while acknowledging that their respective capacities for sentience differ in complexity. In this view, differences in legislation and treatments would be adjusted to differences in sentience.

§

What does a graded notion of sentience amount to? On a standard interpretation, this notion is spelt out as indicating that if we quantify the welfare—that is, according to a subjectivist view, the sentience—of different species in a ratio scale, these will potentially have different maximum and minimum levels (Browning, 2023; Dung, 2023; Fischer & Sebo, 2024; Gaffney et al., 2023; Schukraft, 2020). In this interpretation, we say that a creature's capacity for sentience is more complex than another's if the distance between its maximum and minimum well-being levels is

larger. The rationale behind is this: whereas an animal's welfare level refers to how well or poorly an individual is faring at a time (Broom, 1986), an animal's welfare *capacity* refers to the difference between how well or poorly an animal can fare at a time (Schuhkraft 2020).

Upon closer examination, however, this metric interpretation presents some shortcomings. The first issue is empirical. According to the standard view, organisms' capacity for positive and negative states can vary independently. However, it is unclear what it would mean for an individual to possess asymmetric capacities for positive and negative states. A single organism may *display* positive and negative states to different degrees (Nielsen et al., 2005) or may possess different capacities for coping with adverse events compared to savouring positive ones. However, this does not imply that such an organism possesses different capacities to *undergo* positive and negative experiences.

The second issue is ethical. We can assume that members of one species, e.g., crustaceans, exhibit a heightened capacity for negative experiences but a limited capacity for positive ones and that members of another species, fish, possess a moderate capacity for both positive and negative experiences. If the range between crustaceans and fish's minimum and maximum levels of welfare is equal, they would be predicted to have the same capacity for sentience according to this view. Nonetheless, it would seem that crustaceans' interests should be given greater consideration than fish', as crustaceans' lives would be worse off, other things being equal.

Lastly, another issue concerns the focus of attention. In the standard interpretation, the focus is on how well or bad things can go for an organism, namely, how intense an organism's positive and negative states can be at a given time. Hence, this view appeals to affective states' dimensions: valence (positive or negative) and arousal (high or low). However, determining cross-species comparisons requires focusing on sentience's dimensions rather than on affective states' dimensions. How can this notion of multidimensionality be understood?

§

Sentience concerns organisms' biological capacity to undergo affective states. In this sense, sentience's dimensions ought to be defined as clusters of specialised cognitive, physiological, or behavioural dispositions that provide an organism with such a specific ability—the ability to undergo affective experiences. Hence, due to its ethical relevance, I propose departing from the assumption that cross-species comparisons need to rely on the dimensions defining affective states (valence and arousal) and focus on sentience's biological bases instead.

To provide a first insight into this notion of sentience's multidimensionality, we might briefly draw an analogy with Uexküll's (1934/2010) concept of *Umwelt*. This concept denotes organisms' subjectively perceived surroundings as available through their senses. For example, as Figure 1 illustrates, bees have evolved a distinct sensory apparatus that enables them to perceive some stimuli with different bandwidth, intensity and clarity compared to humans.

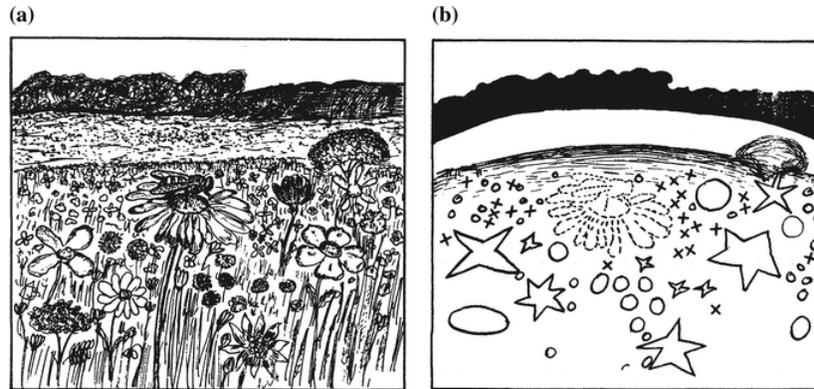


Figure 1 (a) represents the environment of a bee, a blooming field, as perceived by a human observer. Figure (b) represents, in contrast, the same environment from the point of view of the bees' *Umwelt*, where blossoms are perceived as stars and buds as circles according to their shape (Uexküll, 1934/1957).

Similarly, we can assume that affective experiences and their underlying physiological bases are also products of selective processes (Godfrey-Smith, 2017). Under such an assumption, just as animals have evolved different sensory systems, they have also evolved various ways of affectively reacting to the world. As a result of these selection pressures, some species may experience affective states with respect to more or less events, with more or less clarity, or with more or less salience than others. Different species inhabit, so to say, different 'affective Umwelts' or "affective niches".

§

Yet, a problem arises: each affective *Umwelt* would be incommensurable with each other. Korsgaard (2018), for instance, argues that members of each species embody a distinct form of life, making assessment feasible only within the standards established by that particular form. While it might be reasonable to evaluate whether one elephant's life is better or worse than another's, such a comparative judgment becomes pointless when applied to members of different species, like an elephant and an ant, as there is no common standard to apply to both. Hence, overall cross-species welfare comparisons are implausible.

Similarly, notwithstanding the potential value of sentience comparisons for determining principled trade-offs, various authors have argued that the multidimensional character of sentience blocks overall comparisons (Bayne et al., 2016; Birch et al., 2020; Carruthers, 2019). To decide whether the property picked up by 'sentient' in a given context applies to a particular organism, several incommensurable dimensions need to be considered. When assessing whether fish are more sentient than crustaceans, we may ask: are fish more sentient in terms of bandwidth, salience, or granularity? The problem of the incommensurability of sentience's dimensions gives rise to two interconnected challenges.

On the one hand, multidimensionality raises concerns regarding whether sentience's conceptual structure allows for overall orderings. Since sentience is multidimensional, each dimension establishes a different but equally valid ranking. For instance, one species, fish, may surpass elephants regarding the potential bandwidth of their affective experiences, whereas elephants may be more sentient than fish regarding the granularity of such experiences. Consequently,

placing elephants and fish on a unique ranking in terms of overall sentience-ness would be conceptually untenable.¹

On the other hand, multidimensionality also raises empirical concerns. According to this worry, the multidimensionality of sentience renders any attempt at achieving overall comparisons scientifically inaccurate. If a unified scale for comparing different species regarding sentience was established, such comparisons would remain arbitrary. Even if we were to devise a method for converting welfare units across elephants and fish, this conversion method would inevitably overlook some dimensions of their capacity to undergo affective experiences. Consequently, sentience comparisons would provide a deficient guide for determining trade-offs, designing policies, allocating resources, etc.

In response to this problem, researchers argue that it is preferable to contrast each species' individual sentience “profiles” (Birch et al., 2020; Bayne et al., 2016; Dung & Newen, 2023; Veit, 2023). Using Korsgaard’s (2018) terms, this proposal implies focusing on species’ specific standards for how good or bad their lives can go. However, when trade-offs are required, comparisons in terms of profiles are not enough. Hence, in the next section, I propose a multidimensional approach which allows for degrees of sentience, thus showing that these concerns regarding cross-species sentience comparisons lack substantial support. In my view, establishing a robust analysis of the structural aspects of sentience is the initial stride towards a clear understanding of the precise challenges that comparisons regarding sentience impose.

3. The logic of multidimensional comparisons

A property F is multidimensional just in case whether and to what extent something is F depends on how it stands along multiple underlying dimensions, or respects, of F ness. Properties like health, intelligence, or fitness are paradigmatic examples: one can be healthy in various respects, such as blood pressure, cholesterol, or blood sugar level. By contrast, unidimensional properties such as tallness or temperature are associated with a single dimension—height and hotness, respectively.

Following D’Ambrosio & Hedden’s (2023), we can characterise multidimensional properties using a function DIM . This function takes the dimensions of a property F as input and outputs how objects rank along those dimensions. More specifically, $DIM(F, c, w)$ takes a property F , a context c , and a world w as inputs, yielding a profile of orderings $\langle \succsim_{F_1}, \dots, \succsim_{F_n} \rangle$ for the objects in the domain O within that context. Each ordering \succsim_{F_i} in this profile represents how objects rank on underlying dimension i of F . Moreover, each ordering \succsim_{F_i} can be represented by a dimensional value function $V_{F_i} : O \rightarrow \mathbb{R}$ from objects in the domain O to real numbers, such that $V_{F_i}(x) \succsim V_{F_i}(y)$ iff $x \succsim_{F_i} y$. Thus, given a value function V_{F_i} that represents the dimensional ordering \succsim_{F_i} , we can treat $V_{F_i}(x)$ as the degree to which an object x is F along dimension i .

¹ A related criticism concerns sentience’s ‘sharpness’: an organism, it is said, is either sentient or not (however, see Birch, 2020). That is, it is either capable or incapable of having positive or negative experiences. There cannot be borderline cases of sentience—cases where there is no fact about whether something is sentient (e.g., Carruthers, 2019). For this reason, it is argued that sentience cannot be sharp and graded. I will address this issue in Section 3.

Sometimes, we may only discuss whether and to what extent something is F along a specific dimension i . But often, we need to make overall judgments. For instance, we may want to determine which individuals are more athletic overall, which are healthier overall, and so on. According to D'Ambrosio & Hedden (2023), one way to establish overall comparisons consists of aggregating dimensional value functions. What they call a “dimensional aggregation function” $a : V^n \rightarrow \wp(O^2)$ takes a profile of n dimensional value functions $\vec{v} = \langle V_{F_1}, \dots, V_{F_n} \rangle$ as input and returns an ‘overall’ or ‘all-things-considered’ ordering $\succsim_{\vec{v}}^a$ of the objects in the domain O . Hence $a(\vec{v}) = a(\langle V_{F_1}, \dots, V_{F_n} \rangle) = \succsim_{\vec{v}}^a$. Moreover, we can also specify a designated object d as the ‘standard’ for a property F , such that ‘ X is F ’ is true (relative to a and \vec{v}) if and only if $x \succsim_{\vec{v}}^a d$.

For instance, assuming that intelligence (I) differs with regard to attention control (I_1), working memory (I_2), and inference capacities (I_3), DIM 's output would be a profile of orderings $\langle \succsim_{I_1}, \succsim_{I_2}, \succsim_{I_3} \rangle$, where each ordering is represented by a dimensional value function V_{I_i} . Then, even though some individuals may be more intelligent than others regarding attention control, but not their working memory or inference capacities, overall comparisons are obtained by postulating an aggregation function a which yields an overall intelligence ordering $\succsim_{\vec{v}}^a$ as output. Notably, the admissibility of any aggregation function a is dependent on the specific context. For instance, on the set of individuals to which the property is applied. Although working memory may be considered less significant than inference capacities in evaluating the intelligence of some species, the reverse may be true for others.

§

The semantic framework sketched above assists us in demonstrating that, in principle, sentience's multidimensionality does not block overall comparisons. Assuming that the dimensions of sentience (S) relevant in a context c are bandwidth (S_1), granularity (S_2), and salience (S_3), $DIM(S, c, w)$ provides a profile of dimensional orderings $\langle \succsim_{S_1}, \succsim_{S_2}, \succsim_{S_3} \rangle$. This profile, in turn, is represented by the profile of value functions $\vec{v} = \langle V_{S_1}, V_{S_2}, V_{S_3} \rangle$. The conceptual criticism of multidimensionality can be thus translated as follows: sentience is associated with multiple value functions V_{S_i} , each leading to diverse yet equally legitimate rankings of species. Therefore, numerous legitimate ways to rank species prevent the creation of overall comparisons.

However, determining an aggregation function is a crucial additional step for making such comparisons. As with the case of intelligence, an aggregation function a for sentience takes \vec{v} as input and returns an overall ordering $\succsim_{\vec{v}}^a$ as output. For instance, in a context where one needs to compare fish and crustaceans, this would provide us with a precise way to locate them on a single sentience scale. Within this ordering, fish would be at least as sentient as crustaceans relative to a and \vec{v} if and only if $fish \succsim_{\vec{v}}^a crustaceans$. The particularly thorny problem of how any such function a might, in practice, be responsibly determined will be examined in the next section.

This analysis, moreover, shows that sentience can be graded and “sharp”. As mentioned in footnote 1, it is standardly assumed that an organism is either sentient or incapable of experiencing positive or negative hedonic states altogether. Therefore, it is argued, sentience

cannot exhibit gradation. However, a multidimensional framework must only specify a threshold d as the standard for a property F to accommodate sharpness. A judgment ‘X is sentient’ is true (relative to \mathbf{a} and \vec{v}) iff X’s degree of sentience equals or exceeds the standard d and false otherwise. Arguably, as Lee (2023) points out, sentience’s threshold ought to be set to the *lowest* point: any capacity for experiencing pain or joy, even to the dimmest, will qualify an organism as sentient.² Therefore, multidimensionality does not entail indeterminacy.

§

Before moving on, let us address a potential objection to this formal characterisation. The concern revolves around another structural conflict between the sharpness and multidimensionality of sentience. As observed above, a dimensional aggregation function \mathbf{a} can be employed to derive overall sentience comparisons between species in the domain O . However, it may be claimed that for any context c there will always be various competing aggregation functions \mathbf{a} that are deemed admissible in that context. Therefore, since there may be not only one but multiple aggregations functions \mathbf{a} that can be used, there will be multiple, equally valid, overall orderings. As a result, members of one species X may meet the standard d in one overall ordering but not another, so it will be vague whether members of X are sentient or not.

The multidimensional framework, however, remains neutral regarding the number of admissible aggregation functions for a property F . Compare, for instance, the multidimensional properties ‘athletic’ and ‘size’ (understood as ‘volume’). Whereas ‘athletic’ encompasses strength, speed, and endurance, ‘size’ relies on an object’s length, width, and depth. However, ‘athletic’ allows for various aggregation functions contingent upon the relative weight assigned to each dimension, leading to indeterminacy. In some context c , endurance may be prioritised over strength, and vice versa. In contrast, there is a unique admissible aggregation function \mathbf{a} for establishing comparisons of three-dimensional objects based on size. Therefore, introducing aggregation functions does not necessarily imply the vagueness of a property F .

4. A roadmap for navigating the science-policy interface

The last section delineated sentience’s multidimensionality, degreed, and sharp character, showing their simultaneous plausibility. However, the crucial question remains: can cross-species comparisons be conducted responsibly? More precisely, can the values assigned to the aggregation function \mathbf{a} be non-arbitrary? In this section, I address this and related concerns by outlining the commitments that each step involved in making these comparisons entails. While not providing an exhaustive picture, the formalism will assist in organising the decision process, highlighting areas where authors have addressed specific issues and areas where further work is still needed (see Figure 3).

² It is important here to distinguish evidence thresholds from thresholds for an animal to be sentient, here represented by d . The evidence threshold for accepting that a species is sentient can be set to the highest level: we may, for instance, require the empirical observation of various indicators of sentience before acknowledging a species as sentient. Yet, the threshold d for an animal to be considered sentient can be simultaneously set to the lowest point.

Steps	Proposals	Potential issues
I. Which dimensions? (i)	Bandwidth	Some dimensions may be unknown to us.
	Accuity	
	Intensity (?)	Not all dimensions may be equally measurable.
II. Which dimensional value function? ($\forall s_i : O \rightarrow \mathbb{R}$)	Passive value function	Passive view is more tractable but less accurate.
	Active value function	
III. Which aggregation functions? ($a : V^n \rightarrow \wp(O^2)$)	Assume similar weights	Assuming different weights may require appealing to, e.g., social choice theory
	Assume different weights	
IV. Which decision-making formula?	Consider species' actual (and not potential) welfare levels, their population size, their average lifespan, etc.	This involves developing an interface in which welfare scientists, policymakers and stakeholders participate

Fig. 3 This table summarises the main steps involved in cross-species comparisons, from deciding which dimensions compose sentience to formulating a decision-making formula that takes degrees of sentience as input. Note that this decision process characterises a policy-making interface. Whereas steps I and II will be taken care of by welfare scientists, steps III and IV will also involve policymakers and stakeholders as well, without assuming sharp boundaries.

I. Determining sentience's dimensions: The initial query revolves around which dimensions sentience should be associated with. Given that affective sentience is standardly considered to presuppose phenomenal consciousness, some dimensions associated with phenomenal consciousness are also considered integral to the capacity for affective sentience. For instance, Birch et al. (2020) argue that sentience (referred to as 'evaluative-richness' in their paper) involves i) 'evaluative bandwidth', that is, how many inputs the affective state can consider at once, and ii) 'evaluative acuity', that is, how sensitive to slight differences in those inputs an organism is. Dung & Newen (2023) propose that, in addition to these two, sentience's dimensions should also include iii) 'evaluative intensity', that is, how strongly the organism can feel a positive or negative experience.

Importantly, each of these dimensions requires establishing proxies for effective measurement. For instance, one might measure an organism's evaluative bandwidth by observing its reaction times to varying types of situations known to elicit valenced responses (e.g., rewarding vs. punishing scenarios). Determining measurements for a species' intensity, however, may pose more complex challenges. As mentioned in Section 2, pleasure or pain's intensity may be *expressed* differently by different species without entailing that they are experienced differently. To overcome this problem, Browning (2023) argues that, in cases where the species to be compared are phylogenetically close, one can assume that similar behavioural responses map similar degrees of affective intensity. Alternatively, in cases of less closeness among the species, one may proceed by assuming that species possess similar capacity for intensity but manifest it differently in their behaviours.

Some general remarks are in order. First, note that some indicators may inform us that a species is sentient without informing us about its degree of sentience. For example, observing trade-off

behaviour, whereby an animal behaves as if weighing its preference to avoid a noxious stimulus against other preferences, may indicate sentience but not its degree of complexity. Second, once the domain of species in the set O for comparison is determined, the exact dimensions should be applied to all of them to proceed with the subsequent steps. This entails that if one dimension is still unknown to us or can, at the present state of methodology, only be studied in one species but no others (e.g., intensity), then cross-species comparisons should only consider the remaining dimensions (e.g., bandwidth and acuity). Third, in making cross-species comparisons, it is also essential to consider whether all dimensions carry the same weight or whether one may be more relevant than others, a matter that will be addressed in step III.

II. Determining a dimensional value function: The next step involves determining orderings based on sentience’s dimensions. Each of these orderings, represented by a dimensional value function \succsim_{s_i} , provides a ranking describing how sentience varies across species relative to a particular dimension i . As discussed in Section 2, the standard approach involves measuring the distance between each species’ maximum and minimum sentience levels and then comparing them. Yet, this bears some conceptual shortcomings, especially given that it is uncertain that an organism may have divergent capacities for undergoing positive and negative states.

Hence, I propose considering a species’ maximum attainable degree of sentience with respect to each dimension, such as bandwidth, acuity, and intensity, without specifying whether such maximum degree corresponds to positive or negative states. In particular, for each dimension, we can employ a zero-one method, where 0 corresponds to being non-sentient, and 1 corresponds to the maximum possible level of sentience attainable. As discussed in Section 2, the standard d to qualify as sentient ought to be set at the lowest possible level (‘ x is sentient’ is true iff $x \succsim_{\frac{a}{v}} d$); this captures the idea that even the slightest degree of capacity to undergo valenced experiences qualifies an organism as sentient.

For instance, when comparing crabs and fish, fish may (hypothetically) qualify higher in evaluative bandwidth and acuity, but nonetheless qualify lower in evaluative intensity, as illustrated in Figure 4:

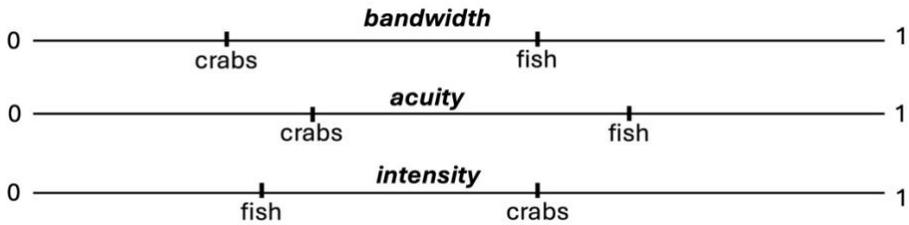


Fig. 4 Hypothetical comparison of fish and crabs’ degree of sentience relative to bandwidth, acuity, and intensity

How does the dimensional value function $\forall s_i$ operate to yield the results in Figure 4? Here, it is assumed that more ‘complex’ organisms possess a heightened capacity for sentience (Nussbaum 2004, p. 309). For example, species that are able to process a more significant number of inputs simultaneously are expected to experience heightened (un)pleasantness compared to those with lesser-developed perceptual capabilities. This entails a particular interpretation of the dimensional value function \succsim_{s_i} , which we may call the ‘passive view’. In this view, an organism’s

degree of sentience is “amplified” in a directly proportional way to its cognitive capacities (e.g., Vallentyne, 2007, p. 213). In other terms, the dimensional function VF_i would map O to \mathbb{R} according to whether a species X belonging to O displays a more cognitively sophisticated skill in a given dimension i .

It should be noted, however, that the passive view raises some concerns (e.g., Akhtar, 2011; Broom, 2007; Wong, 2016; Schuhkraft, 2020). While compelling, it remains unclear whether having more advanced cognitive capacities, such as a greater degree of evaluative bandwidth, implies having a greater capacity to experience pleasure and pain. More complex organisms may have developed additional resources to deal with adverse events, potentially resulting in experiencing lower degrees of pleasure or pain. As a result, this notion of “complexity” and its associated value function might fall short of accurately capturing varying degrees of sentience across species.

As an alternative, I propose that efforts should be made to shift to what may be called an ‘active view’. In this approach, the emphasis would be on how affective experiences influence cognitive capacities. Whereas the passive view sees cognitive bases as enhancing the capacity for sentience, the active view highlights how affective states influence such cognitive bases in return. This perspective involves measuring how organisms’ emotions or moods impact, for example, their evaluative bandwidth, independently of whether such bandwidth is complex or fully developed (perhaps by employing a cross-species judgment bias task paradigm, cf. Hintze, 2016). Here, the dimensional value function VF_i would map O to \mathbb{R} according to the degree to which species’ positive or negative states affect their evaluative bandwidth, acuity, or intensity.

III. Determining an aggregation function: Once orderings of species with respect to each dimension i have been established, the following task is to determine a dimensional aggregation function a that merges them into an overall ranking. This aggregation function leads to the overall computation of degrees of sentience, namely, a ranking of species based on their all-things-considered capacity for (un)pleasant states.

Is an aggregation function a invariably arbitrary? Let’s consider the example depicted in Figure 4. There, we identified the maximum levels of sentience for fish and crabs with respect to each dimension i . To merge these dimensions, we can assume that each dimension has the same weight, that is, that each is equally important in determining an organism’s overall degree of sentience. Hence, such aggregation function a outputs a general rank that averages species’ maximum degree of sentience for each dimension, revealing that fish are more sentient than crabs (as illustrated in Figure 5). In this view, comparing species regarding sentience is analogous to comparing three-dimensional objects based on their volume. In both scenarios, every dimension is considered to bear the same weight, thus rejecting the idea that an aggregation function a would invariably neglect essential aspects of sentience.

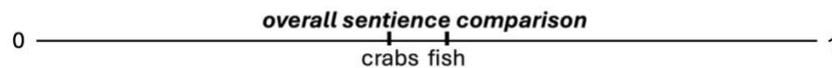


Fig. 5 Hypothetical comparison of fish and crabs’ overall degree of sentience

Yet, it may be claimed that not all dimensions weigh the same, thus undermining the comparison illustrated in Fig. 5. For instance, Dung & Newen (2023) argue that evaluative intensity holds

more ethical relevance than richness and bandwidth, as ethical considerations regarding animal well-being ought to derive from the intensity of the pain experienced. As a consequence, by merging the dimensional rankings in Figure 4, we would obtain that crabs are, in fact, more sentient than fish, as their capacity for evaluative intensity is higher. However, notwithstanding the ethical relevance of intensity, proxies of sentience's intensity may be less reliable than those of evaluative bandwidth and acuity (Stasiak et al., 2003). Hence, to compensate for such unreliability, one may maintain the assignment of equal weight to intensity, bandwidth, and acuity in the aggregation function.

Alternatively, if the need to assign different weights sentience's dimensions persists, we can employ formal tools from welfare economy. As D'Ambrosio & Hedden (2023) argue, the problem of dimensional aggregation is analogous to the issue of 'preference' aggregation in social choice theory. Broadly speaking, social choice theory is concerned with whether and how it is possible to aggregate individual preferences into an overall social ranking. In this context, the key is to think of each underlying dimension of a multidimensional concept as akin to an individual whose preferences correspond to that dimension's ranking of alternatives. Although spelling out the details of how social choice theory can be applied to determine sentience's aggregation function goes beyond the scope of this paper, it is worth noting that there exist formal tools to discuss and ultimately agree on the weights sentience's dimensions should possess (see also, e.g., Hsieh & Andersson, 2021 on comparisons of allegedly incommensurable goods). Hence, the problem of how responsible cross-species comparisons can be determined is, if not totally mitigated, at least responsibly tractable.

IV. Determining a decision-making formula: An organism's degree of sentience is only one among many other inputs required for making informed decisions about resource allocation. That is, for deciding whether to help some organisms rather than others. To take such decisions, these cross-species comparisons must be integrated into a more comprehensive calculation that factors in, for instance, the actual (and not potential) welfare levels of the members of the species compared, the species' population size, their average lifespan, etc.

For instance, Budolfson & Spears (2019) argue that deciding whether to allocate resources to a species not only depends on the degree of sentience of members of *s* but also the average duration of a life of a member of *s*, and a "quality of life adjustment term" that estimates how well members of the species are actually flourishing relative to their species capacity. The rationale behind this is that, although we start by focusing on comparing organisms with respect to their capacity for sentience at a given time, we also need to consider differential lifespans in order to estimate an organism's potential for well-being over their lifetime (Schuhkraft, 2020).

Moreover, Budolfson & Spears (2019) argue that the challenge of making cross-species comparisons is analogous to the challenge of making interpersonal comparisons in economy. The analogy they propose lies in the idea that, just as "consumption" can serve as a proxy for estimating well-being across humans, policymakers and animal welfare scientists could employ a measurable proxy to estimate well-being potentials across species, such as "number of neurons". In technical terms, their proposal consists of developing a formula which involves a concave transformation of the average number of neurons of a species to their well-being in the same way that similar economic formulae transform individuals' consumption into utility. Despite its promising character, however, this method is designed to involve a unidimensional proxy (namely, numbers of neurons), rather than a multidimensional proxy like sentience.

Before moving to the final part of the paper, it is important to note that achieving responsible cross-sentience comparisons is a multidisciplinary task that requires working at the interface between science and policymaking. From step III, it becomes clearer that decisions not only rely on empirical concerns but also involve policymakers' values (e.g., in determining how much each dimension of sentience should weigh) and stakeholders' interests (e.g., in deciding which and how resources should be ultimately allocated). Although the result of these comparisons will always be imperfect and may not be followed to the letter, having clearly specified our assumptions will help spot those places where refinement is needed. Ultimately, working towards a clear decision path is more responsible than deciding based on our intuitions or dubious proxies such as a number of neurons. As Gaffney et al. (2023) propose, we could even use different decision processes and then converge the results; that is, we can aggregate proposals like this one and his to arrive at the outcome that best maximises welfare.

5. Addressing some objections

Each of the steps outlined above entails commitments that have to be clearly stated. However, concerns about the notion of welfare we have assumed or the proposed methodology itself also arise. In this section, I address some of these issues.

Choosing a welfare proxy: There is wide disagreement about the basis for welfare. In this paper, I assumed that welfare is a matter of the capacity to experience valenced states like pleasure and pain. However, this assumption might make inter-species welfare comparisons harder since it implies that the determinants of welfare are not directly observable. Other perspectives, in contrast, might employ observable determinants of welfare, thereby making cross-species comparisons more tractable in contrast. Some researchers, for example, have proposed to think of welfare in terms of an objective list of goods (e.g., sociality, intelligence, love, to name a few) which might vary within and across species depending on the form of life that particular welfare subjects have (Moore, 2000).

However, it is worth noting that this alternative account, in terms of a list of objective goods, faces a problem of aggregation similar to that which we have treated in this paper. Indeed, this approach not only requires determining the different parameters by which a life could be said to have flourished but also integrating them into an overall measure of the individual's degree of flourishing. Moreover, because the proposed determinants of welfare will probably have different weights for different species, the objective list theorists will also have to recur to an aggregation function. For instance, even though octopuses are incredibly intelligent, they are also deeply asocial. Ants, in contrast, are plausibly much less intelligent, but they tend to live in densely populated mounds containing millions of individual ants. Hence, a trade-off involving these two species will have to assign a species-specific weight to each factor and aggregate them for determining trade-offs.

Gathering evidence: As we saw in Section 4, comparing two species with respect to their capacity for sentience requires using the exact dimensions for both species. However, in measuring an organism's capacity for sentience with respect to each dimension, specific experimental paradigms tailored to the requirements of each species are often necessary. For

example, when studying elephants' evaluative bandwidth, measurements ought to be adjusted to the perceptual abilities of that particular species. Consequently, a potential concern arises: the results of a study on evaluative bandwidth in one species (e.g., elephants) might not be directly comparable to the findings of a similar study in another species (e.g., ants), undermining its effectiveness as a guide for cross-species sentience comparisons.

However, a comparable issue arises in inter-individual comparisons, such as when comparing the cognitive skills of humans from different cultures. The challenge here is analogous: experimental paradigms must be tailored to the characteristics of the study's subjects and their culture. For example, when examining literacy skills in children from diverse cultures, experiments need to account for the preferred forms of information transmission in those cultures (e.g., print vs. digital). The need to adjust paradigms for different subjects does not automatically invalidate comparisons across the studied subjects of the same species. Hence, despite the uncertainty involved, we can expect cross-species comparisons to be still preferable to not attempting them at all, as emphasised by Fischer & Sebo (2024).

The relevance of sentience: Lastly, it may be argued that even if sentience varies in degrees, the decision-making process required for resource allocation does not need to hinge on degrees of sentience. This critique emerges from viewpoints that advocate separating research on animal sentience from policymaking. Dawkins (2022), for instance, argues that these domains follow distinct rules and should thus be kept separate. However, this remains a minority viewpoint, as there is a growing consensus that policymaking should, whenever possible, incorporate information from animal welfare science.

A way to observe the relevance of the science-policymaking interface is the following: policymakers' deliberations regarding the prospective allocation of resources to one species rather than another can, in turn, have an impact on scientists' decisions regarding the preferred evidence threshold for sentience. Prospective scenarios requiring more decisions about which species to prioritise (e.g., fish vs crustaceans, crustaceans vs insects, and so on) may lead stakeholders to lobby for a stricter evidence threshold (e.g., one threshold where insects, or even crustaceans, would not count as sentient). Yet, since this would put many potentially sentient species at risk, understanding how such sentience comparisons can be responsibly determined becomes crucial.

6. Conclusion

This paper has pursued two interconnected objectives. Firstly, it has argued that the structure of sentience possesses sharp and multidimensional features while still allowing for overall comparisons without resulting in contradictions. Secondly, it has utilised this formal characterisation to outline the decision-making process that ought to form the basis for establishing welfare comparisons across species. Through these proposals, my goal was to show that the idea that sentience comes in degrees possesses not only conceptual validity but also practical utility in real-world scenarios.

Moreover, the methodology used in this paper can serve to shed light on the complex issues arising at the crossroads of welfare sciences, policymaking, and ethics. Its utility, however, will be largely dependent on the progress of animal welfare scientist's work. Species which are not studied enough will not be able to be considered in deliberations regarding sentience

comparisons, potentially causing unnecessary harm. Still, this paper can be seen as contributing to the idea that animal sentience is not only an essential field of research per se but also can affront and detect sources of uncertainty, thus fostering consensus about the distribution and quality of species' sentience when shaping policymaking is required.

Lastly, although there are alternative approaches to animal welfare that merit discussion, such as considering welfare as an objective list of goods, this paper has emphasised that many of the issues raised in this paper will also arise in those proposals. Thus, this paper indirectly but meaningfully contributes to such discussions. Similarly, when the time comes to compare the average capacity for welfare not only of different species but also of organisms of different types of substrates (animal vs AI), the problems raised in this paper will remain relevant.

References

Akhtar, S. (2011). Animal pain and welfare: Can pain sometimes be worse for them than for us?.

Barron, A. B., & Klein, C. (2016). What insects can tell us about the origins of consciousness. *Proceedings of the National Academy of Sciences*, 113(18), 4900-4908. <https://doi.org/10.1073/pnas.1520084113>

Bayne, T., Hohwy, J. &, & Owen, A. M. (2016). Are There Levels of Consciousness? *Trends in Cognitive Sciences*, 20(6), 405–413.

Birch, J. (2017). Animal sentience and the precautionary principle. *Animal sentience*, 2(16):1.

Birch, J. (2020). Global Workspace Theory and Animal Consciousness. *Philosophical Topics*.

Birch, J., Schnell, A. K., and Clayton, N. S. (2020). Dimensions of animal consciousness. *Trends in cognitive sciences*, 24(10):789–801.

Budolfson, M. and Spears, D. (2019). Quantifying animal well-being and overcoming the challenge of interspecies comparisons. In *The Routledge Handbook of Animal Ethics*, pages 92–101. Routledge.

Broom, D. M. (2007). Cognitive ability and sentience: which aquatic animals should be protected?. *Diseases of Aquatic Organisms*, 75(2), 99-108.

Brown, C. (2015). Fish intelligence, sentience and ethics. *Animal cognition*, 18(1), 1–17.

Browning, H. (2023). Welfare comparisons within and across species. *Philosophical Studies*, 180(2), 529–551.

Carruthers, P. (2019). *Human and Animal Minds: The Consciousness Questions Laid to Rest*. Oxford University Press.

Dawkins, M. S. (2022). The science of animal sentience and the politics of animal welfare should be kept separate. *Animal Sentience*, 6(31):8.

Dung, L. (2023). Dimensions of animal wellbeing. *Philosophy and the Mind Sciences*, p. 4.

Dung, L. and Newen, A. (2023). Profiles of animal consciousness: A species-sensitive, two-tier account to quality and distribution. *Cognition*, 235:105409.

D'Ambrosio, J. and Hedden, B. (2023). Multidimensional adjectives. *Australasian Journal of Philosophy*, pages 1–25.

Fischer, B., & Sebo, J. (2024). Intersubstrate Welfare Comparisons: Important, Difficult, and Potentially Tractable. *Utilitas*, 36(1), 50-63.

Gaffney, L. P., Lavery, J. M., Schiestl, M., Trevarthen, A., Schukraft, J., Miller, R., Schnell, A. K., & Fischer, B. (2023). A theoretical approach to improving interspecies welfare

comparisons. *Frontiers in Animal Science*, 3, 1062458. <https://doi.org/10.3389/fanim.2022.1062458>

Godfrey-Smith, P. (2017). The evolution of consciousness in phylogenetic context. In *The Routledge Handbook of Philosophy of Animal Minds* (pp. 216–226). Routledge.

Hintze, S., Melotti, L., Colosio, S., Bailoo, J. D., Boada-Sañna, M., Würbel, H., and Murphy, E. (2018). A cross-species judgement bias task: integrating active trial initiation into a spatial go/no-go task. *Scientific reports*, 8(1):5104.

Hsieh, Nien-hê and Henrik Andersson, "Incommensurable Values", *The Stanford Encyclopedia of Philosophy* (Fall 2021 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/fall2021/entries/value-incommensurable/>>.

Korsgaard, C. (2018). *Fellow Creatures: Our Obligations to the Other Animals* (New York: Oxford University Press), 3–15.

Lagisz, M., Zidar, J., Nakagawa, S., Neville, V., Sorato, E., Paul, E. S., Bateson, M., Mendl, M., and Løvlie, H. (2020). Optimism, pessimism and judgement bias in animals: a systematic review and meta-analysis. *Neuroscience & Biobehavioral Reviews*, 118:3–17.

Lee, A. Y. (2023). Degrees of consciousness. *Noûs*, 57(3), 553–575.

Mendl, M., Burman, O. H., and Paul, E. S. (2010). An integrative and functional framework for the study of animal emotion and mood. *Proceedings of the Royal Society B: Biological Sciences*, 277(1696):2895–2904.

Mill, J. S. (2016). Utilitarianism. In *Seven Masterpieces of Philosophy*, pages 329–375. Routledge.

Moore, A., 2000, “Objective Human Goods”, in B. Hooker and R. Crisp (eds.), *Well-being and Morality: Essays in Honour of James Griffin*, Oxford: Clarendon Press, 75–89.

Nagel, T. (1974). What is it like to be a bat? *Philosophical Review*, 83(4), 435–450. <https://doi.org/10.2307/2183914>

Nussbaum, M. C. (2004). Beyond “Compassion and Humanity:” Justice for Nonhuman Animals. in C. R. Sunstein and M. Nussbaum (eds) *Animal Rights: Current Debates and New Directions*. Oxford: Oxford University Press, pp. 299–320

Powell, R., & Mikhalevich, I. (2021). Affective sentience and moral protection. *Animal Sentience*, 29. <https://www.doi.org/10.51291/2377-7478.1668>

Russell, W. M. S., & Burch, R. L. (1959). *The principles of humane experimental technique*. Methuen.

Schukraft, J. (2020). Comparisons of capacity for welfare and moral status across species. *Rethink Priorities*. <https://forum.effectivealtruism.org/s/Yw5JeJHyhYQtYw8Z6/p/4Xg3dC6rrW4WFnSne#fnref-auGCKNd6gDdA86tLa-6>

Sebo, J. (2018). The Moral Problem of Other Minds. *The Harvard Review of Philosophy*, 25,51–70 <<https://doi.org/10.5840/harvardreview20185913>>.

Singer, P. (1977). *Animal liberation*. Avon Books.

Stasiak, K. L., Maul, D. O. N., French, E., Hellyer, P. W., & Vandewoude, S. (2003). Species-specific assessment of pain in laboratory animals. *Journal of the American Association for Laboratory Animal Science*, 42(4), 13–20.

Vallentyne, P. (2007). Of mice and men: Equality and animals. in N. Holtug and K. Lippert-Rasmussen (eds.) *Egalitarianism: New Essays on the Nature and Value of Equality*, Oxford University Press, pp. 211–238.

Veit, W. (2023). *A Philosophy for the Science of Animal Consciousness*. New York: Routledge.

Višak, T. (2017). Cross-species comparisons of welfare. In A. Woodhall & G. G. da Trindade (Eds.), *Ethical and political approaches to nonhuman animal issues* (pp. 347–363). Springer International Publishing. https://doi.org/10.1007/978-3-319-54549-3_15

Višak, T. (2022). *Capacity for welfare across species*. Oxford University Press.

von Uexküll [1934] 1957, "A Stroll Through the Worlds of Animals and Men: A Picture Book of Invisible Worlds." In *Instinctive Behavior: The Development of a Modern Concept*, edited and translated by Claire H. Schiller, New York: International Universities Press, pp. 5–80.

von Uexküll, [1934] 2010. *A Foray Into the Worlds of Animals and Humans with a Theory of Meaning*. Minneapolis: University of Minnesota Press

Wong, K. (2016). *Counting Animals: On Effective Altruism and the Prospect of Interspecies Commensurability*, B.A. thesis, Princeton University, Princeton, NJ.