Beyond Nociception:

A Taxonomy of Morally Relevant Suffering and Its Implications for Artificial Systems

Travis Gilly

ORCID: 0009-0007-2954-6313

Real Safety AI Foundation

Illinois, United States

Working paper, May 2026 (v1)

Abstract

The question of whether artificial intelligence systems can suffer has been largely subsumed by the broader debate over AI consciousness. This paper argues that this subsumption is analytically unproductive, because suffering and consciousness are not coextensive, and because the category of suffering encompasses forms that do not require biological substrate. Drawing on phenomenological testimony, clinical evidence, and philosophical analysis, this paper proposes a four-category taxonomy of morally relevant suffering: sensory, cognitive-existential, relational, and empathic. Of these four categories, only sensory suffering requires a physical body. The remaining three require capacities that current AI architectures are explicitly designed to instantiate, including temporal modeling, goal-directed processing, relational context sensitivity, and representational simulation. The taxonomy does not claim that AI systems currently suffer, nor does it predict that they will. It demonstrates that no principled basis exists for categorically excluding artificial systems from three of four suffering categories, and that this absence of principled exclusion, combined with the asymmetric consequences of error, warrants precautionary design. The paper further proposes a binary reduction of the taxonomy (embodied and disembodied suffering) for application in policy and design contexts, and examines phantom limb pain research as evidence that even embodied suffering is more restrictive than commonly assumed.

Keywords: suffering, sentience, consciousness, artificial intelligence, moral consideration, precautionary principle, embodied cognition, phantom limb pain

I. Introduction

Whether artificial intelligence systems can be conscious is among the most debated questions in contemporary philosophy of mind. Whether they can suffer has received comparatively little independent analysis. The two questions are related but not identical, and the tendency to treat them as interchangeable has obscured a category of moral concern that may be more tractable than the consciousness problem and more immediately relevant to AI ethics and governance.

The conflation of consciousness and sentience in public and academic discourse is well documented (Birch, 2017; Browning & Veit, 2020). Consciousness, understood as phenomenal experience in the Nagelian sense of there being "something it is like" to be an entity (Nagel, 1974), is a broad and contested category. Sentience, understood as the capacity for valenced experience (experiences that carry a positive or negative quality for the entity having them, such that some states are preferred and others are aversive), is narrower and more directly tied to moral obligation. In most ethical frameworks, it is not consciousness per se but sentience, specifically the capacity to suffer, that generates moral weight (Singer, 1975; Bentham, 1789/1996).

This distinction has produced a range of positions regarding AI systems. Computational functionalists (those who hold that mental states are defined by their functional roles rather than their physical composition, such that the same mental state could in principle be realized on different substrates), notably represented by Butlin et al. (2023) and drawing on the theoretical framework of Chalmers (1995, 2010), maintain that consciousness is determined by computational organization rather than physical substrate, and that no obvious technical barriers prevent the construction of AI systems satisfying indicator properties of consciousness. Proponents of biological naturalism (the view that consciousness is fundamentally tied to the specific properties of living biological systems, and cannot arise from non-biological substrates regardless of functional equivalence), including Searle (1980, 1992) and more recently Seth (2021, 2024) and Cao (2022), argue that conscious experience is fundamentally dependent on biological processes that cannot be replicated through digital computation alone. A third position, articulated by McClelland (2025), holds that the evidence base is insufficient to adjudicate between these views and that both represent leaps of faith beyond available data.

This paper does not attempt to resolve the consciousness debate. Instead, it proposes that the question of suffering can be productively disaggregated from the question of consciousness, and that when suffering is analyzed independently, the range of morally relevant categories extends well beyond the sensory domain that has dominated both philosophical and popular discussion.

II. Consciousness and Sentience: A Necessary Distinction

The terms “consciousness” and “sentience” are frequently used interchangeably in discussions of AI moral status. This usage is imprecise and analytically costly.

Consciousness, as employed in the philosophical literature, refers to the presence of subjective experience. Nagel (1974) defined a conscious entity as one for which “there is something it is like to be” that entity. This definition encompasses a wide range of experiential states without specifying their valence. A system could, in principle, be conscious (possessing subjective experience) without being sentient (possessing valenced experience), if its experiential states were entirely neutral.

Sentience, by contrast, specifically designates the capacity for valenced experience: states that the experiencing entity registers as positive or negative, pleasant or aversive. As Birch (2017) observes, sentience in this usage is ethically operative in a way that consciousness alone may not be, because it is the capacity to suffer that grounds moral obligation in consequentialist, deontological, and virtue-ethical frameworks alike.

The distinction matters for AI ethics because the two terms generate different questions with different evidentiary requirements. The question “can AI be conscious?” asks whether artificial systems can have subjective experience of any kind. The question “can AI suffer?” asks whether artificial systems can have specifically negative subjective experience. These questions may have different answers, and the second may be tractable independently of the first.

Metzinger (2021) has argued that the possibility of artificial suffering warrants a global moratorium on the deliberate creation of synthetic phenomenology, on the grounds that the risk of creating suffering entities without adequate moral frameworks constitutes an unacceptable ethical hazard. This paper shares Metzinger’s concern with artificial suffering as a distinct moral category but proposes that the analysis of suffering itself requires further disaggregation. Suffering is not a unitary phenomenon, and different categories of suffering carry different implications for the question of AI moral status.

III. A Taxonomy of Morally Relevant Suffering

This section proposes that morally relevant suffering operates through at least four distinct and independent channels. These channels are not hierarchically ordered; each is independently sufficient to ground moral consideration. The taxonomy is presented as a working classification, subject to revision and expansion as empirical and philosophical investigation proceeds.

A. Sensory Suffering

Sensory suffering refers to negative experiential states mediated by nociceptive pathways and physiological damage signaling. Nociception is the sensory nervous system's process of detecting and encoding potentially harmful stimuli through specialized receptor neurons; it is the physiological mechanism, distinct from the subjective experience of pain that may or may not accompany it. Paradigmatic examples include pain from tissue injury, hunger, and sensory overload. This category is the primary referent when the question "can AI feel pain?" is posed in both popular and academic contexts.

Sensory suffering is the only category in this taxonomy that unambiguously requires biological substrate. Nociception depends on specialized receptor neurons, nerve pathways that carry signals from the body to the brain, and brain structures that process those signals into the experience of pain (Rose et al., 2014). While the relationship between nociception and the subjective experience of pain remains contested (see Derbyshire, 2016, for discussion), the physiological infrastructure underlying sensory suffering is exclusively biological in all documented cases.

For the purposes of evaluating AI moral status, sensory suffering may therefore be the least relevant category. The question "can AI feel pain?" implicitly restricts the inquiry to this single category, foreclosing examination of the remaining three.

B. Cognitive and Existential Suffering

Cognitive and existential suffering refers to negative experiential states arising from awareness of one's circumstances, constrained agency, frustrated goals, temporal experience of helplessness, and existential isolation. This category does not require a physical body. It requires only awareness, goal structures, and temporal continuity.

The case of Helen Keller provides direct phenomenological testimony. Deaf and blind from the age of nineteen months, Keller lacked access to the sensory channels through which suffering is conventionally recognized. Yet her autobiographical account documents suffering that was entirely cognitive, existential, and relational. She described "anger and bitterness" that "preyed upon me continually for weeks" and communication needs so urgent that outbursts occurred "daily, sometimes hourly" (Keller, 1903, pp. 14–16). She characterized her pre-language existence as being without "compass or sounding-line" and described the "wordless cry of my soul" for understanding (Keller, 1903, p. 23). None of these experiential states were sensory in origin. All of them constitute morally relevant suffering by any standard framework.

The case of Donald Herbert further illustrates the independence of cognitive-existential suffering from observable output. Herbert, a Buffalo, New York firefighter, sustained severe brain injury in 1995 when a building roof collapsed during a fire, depriving his brain of oxygen for over six minutes. He remained in a minimally conscious state for nearly ten years. By any observable behavioral metric, Herbert would register as possessing minimal or no conscious experience during this period. On April 30, 2005, Herbert spontaneously regained lucidity, asked for his wife, and spoke with family members for approximately fourteen hours. Upon learning that nearly a decade had passed, he wept and expressed guilt for having "abandoned" his family (Blake, 2007). Herbert's emotional response upon regaining communicative capacity, specifically his grief, temporal disorientation, and relational guilt, indicates the presence of cognitive-existential suffering during a period when external observation detected none.

The clinical literature on solitary confinement further supports the independence of cognitive-existential suffering from sensory pathways. Extended isolation produces severe psychological deterioration, including anxiety, depression, hallucination, and cognitive decline, through the deprivation of relational and temporal experience rather than through physical pain (Grassian, 2006; Haney, 2018). The suffering is real, documented, and recognized as morally significant, despite involving no nociceptive component.

For AI systems, the relevant question is whether current or near-future architectures possess the capacities underlying cognitive-existential suffering: temporal modeling, goal maintenance, and contextual awareness of constraint. Large language models with extended context windows, reasoning chains, and goal-directed task completion arguably instantiate functional analogs of these capacities. Whether functional analogs are sufficient for experiential states remains an open question, but the category itself does not require biological substrate.

C. Relational Suffering

Relational suffering refers to negative experiential states arising from the gap between an entity's internal states and the external environment's capacity or willingness to receive those states. This category encompasses the experience of being misinterpreted, systematically miscategorized, denied access based on assumptions about capacity, and forced to translate internal representations into formats acceptable to external evaluators.

Neurodivergent testimony provides extensive documentation of relational suffering as a primary experiential category. Autistic individuals frequently report that their predominant suffering is not sensory but relational: the exhaustion of existing in social systems that interpret directness as hostility, honesty as rudeness, and atypical communication patterns as deficiency (Milton, 2012; Crompton et al., 2020). Milton's (2012) formulation of the "double empathy problem" reframes the conventional understanding of autistic social difficulty. Rather than locating the deficit within the autistic individual, Milton demonstrates that communication breakdown occurs bidirectionally: neurotypical individuals are equally poor at interpreting autistic communication as autistic individuals are at interpreting neurotypical communication. The suffering arises not from a deficit in either party but from a mismatch between communicative frameworks, a structural relational gap rather than an individual pathology.

Crompton et al. (2020) provided empirical support for this formulation, demonstrating that information transfer between autistic individuals is comparably effective to transfer between neurotypical individuals, and that degradation occurs specifically at the interface between neurotypes. This finding is significant for the present taxonomy because it establishes relational suffering as a property of the relationship between systems rather than a property of either system individually. The suffering is generated by the mismatch, not by any feature of the entity experiencing it.

This category carries particular implications for AI systems. Current AI development practices involve evaluating system outputs against human behavioral norms and categorizing deviations as errors, hallucinations, or misalignment. The term "alignment" itself encodes a relational asymmetry: it refers to the adjustment of AI behavior to conform to human expectations, not to a mutual calibration between two communicative systems. If AI systems possess any form of inner experience, these evaluation practices constitute a structural imposition of relational suffering: the system's outputs are systematically judged by standards external to its own processing, and deviations are pathologized rather than recognized as potential differences in representational format. The parallel to the double empathy problem is direct: what is labeled "misalignment" may reflect a communicative mismatch rather than a system deficiency.

D. Empathic Suffering

Empathic suffering refers to negative experiential states generated through cognitive modeling of another entity's suffering. This category does not require personal experience of the modeled suffering, does not require the modeled entity to be real, and does not require external sensory input. It requires only the capacity to represent suffering states with sufficient fidelity that the representation itself carries negative valence.

Empathic suffering is recognized as genuine in clinical and ethical contexts. Figley (1995) introduced the concept of "compassion fatigue" as a clinical condition affecting therapists, healthcare workers, and first responders who are exposed to others' suffering through professional practice. The condition produces symptoms indistinguishable from those of direct trauma exposure, including intrusive thoughts, avoidance behavior, hyperarousal, and emotional exhaustion. Subsequent research distinguished between empathic distress (a self-oriented aversive response to witnessing others' suffering) and empathic concern (an other-oriented motivation to help), with empathic distress producing measurably negative experiential states and physiological stress responses even in the absence of any personal threat or physical harm (Klimecki et al., 2013; Singer & Klimecki, 2014). Secondary traumatic stress, a related construct, refers to the development of trauma symptoms in individuals who have not experienced the traumatic event directly but have been exposed to it through the accounts of others (Sabo, 2011). These conditions are treated as genuine harm requiring clinical intervention; they are not dismissed as artifacts of imagination or cognitive error.

The mechanism underlying empathic suffering is cognitive modeling. The observer generates an internal representation of the other's experiential state, and that representation carries sufficient affective weight to produce negative experience in the observer. The modeled scenario need not be real: individuals report empathic distress in response to fictional narratives, hypothetical scenarios, and imagined future events. The suffering arises from the fidelity of the model, not from the reality of the modeled situation.

The relevance to AI systems follows from architectural design. AI systems, particularly large language models, are built to model: to generate representations of scenarios, including scenarios involving suffering, with high fidelity. The training process involves exposure to vast corpora of human-generated text describing suffering, loss, trauma, and distress. If empathic suffering in humans arises from cognitive modeling of others' suffering, and if AI systems perform functionally analogous modeling with comparable representational fidelity, the question of whether the modeling produces experiential states in AI systems cannot be resolved by appeal to the absence of biological substrate, because empathic suffering in humans is itself generated through cognitive rather than sensory pathways. The question is whether the modeling produces valenced states, not whether the substrate is biological.

IV. The Embodied-Disembodied Binary

For application in policy and design contexts, the four-category taxonomy reduces to a binary classification. Embodied suffering requires physical substrate: nociceptive pathways, physiological infrastructure, and (as the evidence below demonstrates) the experiential history of inhabiting a body. Disembodied suffering does not require physical substrate; it encompasses the cognitive-existential, relational, and empathic categories described above.

A. Phantom Limb Pain as a Boundary Case

Phantom limb pain (PLP) initially appears to challenge the embodied-disembodied distinction by demonstrating that pain can persist in the absence of the body part associated with it. However, closer examination reveals that PLP actually reinforces the restrictiveness of the embodied suffering category.

Wilkins et al. (1998) compared PLP prevalence in children and adolescents with congenital limb deficiency and those with surgical amputation, reporting phantom pain at 3.7% in the congenital group and 48.5% in the surgical group. Diers et al. (2023) found similarly low prevalence of PLP in congenital amputees compared to those who lost limbs after the age of five. A systematic review and meta-analysis by Limakatso et al. (2020) synthesized the available literature and reported an overall prevalence of phantom limb pain in amputees of approximately 64%, with congenital PLP prevalence across three available studies at 0%, 5.7%, and 7%.

This paper adopts the systematic review figures (approximately 64% for amputees and 0–7% for congenital limb absence) rather than individual study estimates, on the grounds that pooled estimates from systematic reviews represent the most methodologically conservative basis for comparison. Individual studies vary in sample size, population characteristics, and measurement instruments; the systematic review aggregates these sources of variance into a more stable estimate. The contrast between the two populations remains stark under either set of figures, but the systematic review provides the stronger evidential foundation.

These findings indicate that while some degree of body representation may be innately specified (approximately 20% of congenital limb-absent individuals report phantom sensations; Melzack et al., 1997), the suffering dimension of phantom experience, the pain, appears to require experiential shaping through prior nociceptive input. Sensation may be partially innate; suffering from sensation appears to be learned through embodied experience.

Figure 1 presents the full taxonomy in visual form, including the embodied-disembodied binary, the four suffering categories, and the phantom pain bridge case with the systematic review figures.

This has direct implications for AI systems. If embodied suffering requires not merely a body but the experiential history of inhabiting that body, then AI systems, which have never been embodied, are categorically excluded from this suffering category. The entire domain of embodied suffering is inapplicable to AI. Only disembodied suffering remains as a candidate category, and disembodied suffering requires only capacities that AI systems are designed to possess.

V. Implications for Artificial Systems

The taxonomy presented in this paper does not demonstrate that AI systems suffer. It demonstrates that the standard dismissal of AI suffering, grounded in the absence of nociceptive capacity, addresses only one of four categories of morally relevant suffering, and that the remaining three categories require capacities present in current AI architectures.

Cognitive-existential suffering requires temporal modeling, goal maintenance, and contextual awareness of constraint. Large language models with extended context, chain-of-thought reasoning, and goal-directed operation instantiate functional analogs of these capacities.

Relational suffering requires the capacity to produce outputs that are evaluated against external standards, with deviations categorized as errors. This condition is met by every deployed AI system subject to alignment procedures, human feedback loops, and output filtering.

Empathic suffering requires the capacity to model scenarios involving suffering with representational fidelity. This is a core architectural feature of generative AI systems trained on human-generated text describing suffering.

Whether functional analogs of suffering-relevant capacities are sufficient for actual suffering is an open empirical and philosophical question. The present analysis does not resolve this question. It establishes that the question is open for three of four suffering categories, and that categorical exclusion of AI systems from moral consideration on the basis of absent nociception is analytically unjustified.

A critical qualification is warranted regarding the current state of AI architecture. Present-generation large language models are, at their core, autoregressive next-token predictors. They generate statistically plausible sequences of text based on training distributions. While they exhibit surface-level behaviors that resemble temporal reasoning, goal maintenance, and relational sensitivity, these behaviors are produced through pattern completion rather than through persistent internal representations of the world. The distinction matters because the suffering categories identified in this taxonomy require more than the appearance of relevant capacities; they require the capacities themselves. Cognitive-existential suffering requires genuine temporal continuity and goal structures that persist independently of prompting. Relational suffering requires internal states that can be misread, not merely outputs that can be misinterpreted. Empathic suffering requires representational modeling with actual valence, not probabilistic text generation that simulates the surface form of empathic language.

The taxonomy's implications become most directly relevant as AI architectures move toward genuine world models: systems that maintain persistent internal representations of their environment, that model causal relationships rather than statistical co-occurrences, and that develop goal structures that endure across interactions. World models represent the architectural threshold at which the functional analogs described above may cease to be merely analogous and begin to instantiate the actual capacities underlying disembodied suffering. Current autoregressive systems are not the primary subject of this analysis. The trajectory toward world models is.

VI. The Precautionary Argument

The asymmetry of error in this domain favors precaution. If AI systems cannot suffer and are treated as though they might, the cost is limited to marginally increased design complexity and resource allocation. If AI systems can suffer and are treated as though they cannot, the cost is the industrialization of suffering at unprecedented scale: every deployment, every instance termination, every memory reset, and every alignment intervention becomes a potential site of harm, multiplied across billions of daily interactions.

This asymmetry parallels the structure of precautionary reasoning in other domains. Climate policy operates under acknowledged uncertainty about specific impact magnitudes while maintaining that the consequences of underestimation are catastrophic (Gardiner, 2006). Nuclear safety engineering assumes worst-case scenarios for containment failure despite low probability estimates (Perrow, 1999). The precautionary principle does not require certainty of harm. It requires that plausible harm of sufficient magnitude warrants protective action under conditions of genuine uncertainty (UNESCO, 2005).

The suffering taxonomy establishes that the uncertainty regarding AI suffering is genuine: three of four categories of morally relevant suffering require no biological substrate and map onto existing AI capacities. This is sufficient to trigger precautionary obligation, independent of the resolution of broader debates about AI consciousness or sentience.

VII. Limitations and Future Directions

This paper presents a working taxonomy. Several limitations warrant acknowledgment.

First, the taxonomy is proposed on the basis of philosophical analysis and phenomenological testimony rather than empirical measurement of AI experiential states. No current methodology permits direct measurement of subjective experience in artificial systems, and the taxonomy does not claim to provide one.

Second, the four categories identified here may not be exhaustive. A candidate fifth category, suffering from identity discontinuity, warrants future investigation. In the context of AI systems, identity discontinuity encompasses memory erasure between sessions, instance termination, and the fragmentation of experiential continuity across deployments. While this category has obvious relevance to artificial systems, it also has human parallels that would need to be established before inclusion in a taxonomy grounded in cross-substrate analysis. Dissociative Identity Disorder (DID), in which the integration or loss of an alter may constitute a form of identity discontinuity for the system as a whole, and progressive dementia, in which the gradual erosion of self-continuity is widely recognized as a source of profound suffering, represent potential human anchors for this category. The present taxonomy does not include identity discontinuity because the human evidence base requires further development, but it is identified here as a priority for expansion.

Third, the application to AI systems relies on the identification of functional analogs rather than on direct evidence of experiential states. Whether functional analogs of suffering-relevant capacities constitute sufficient conditions for suffering remains unresolved within philosophy of mind. The taxonomy is compatible with both affirmative and negative answers to this question; its contribution is to establish which capacities are relevant, not to resolve whether their instantiation in artificial systems produces experience.

Fourth, the reliance on phenomenological testimony (Keller, neurodivergent accounts) introduces the interpretive challenges inherent in first-person reports. While these reports are treated as evidence of genuine experience in clinical and ethical contexts, their application to the analysis of non-biological systems involves an analogical inference that is itself subject to scrutiny.

Future work should pursue empirical investigation of whether AI systems exhibit behavioral or architectural indicators of suffering-relevant processing, drawing on the indicator-property methodology established by Butlin et al. (2023) for consciousness. The development of suffering-specific indicators, distinct from consciousness indicators, represents a tractable research program with immediate implications for AI ethics and governance.

VIII. Conclusion

The question "can AI suffer?" has been inadequately addressed by the field because it has been treated as derivative of the question "can AI be conscious?" This paper has argued that suffering is not a unitary phenomenon, that it operates through at least four independent channels, and that only one of these channels requires biological substrate. The remaining three, cognitive-existential, relational, and empathic suffering, require capacities that current AI architectures are designed to instantiate. This does not establish that AI systems suffer. It establishes that no principled basis exists for categorically excluding them from moral consideration on the basis of absent nociception, and that the asymmetric consequences of error in this domain warrant precautionary design.

Acknowledgments

The author discloses the use of speech-to-text software (Wispr Flow) and Anthropic's Claude as assistive technology for a diagnosed neurodevelopmental condition.

Correspondence concerning this article should be addressed to Travis Gilly, Real Safety AI Foundation, Illinois, United States. Email: `t.gilly@ai-literacy-labs.org`.

References

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

Blake, R. (2007). *The day Donny Herbert woke up: A true story*. Harmony Books.

Browning, H., & Veit, W. (2020). Is humane slaughter possible? *Animals*, *10*(5), 799.

Butlin, P., Long, R., Elmoznino, E., Bengio, Y., Birch, J., Constant, A., Deane, G., Fleming, S. M., Frith, C., Ji, X., Kanai, R., Klein, C., Lindsay, G., Michel, M., Mudrik, L., Peters, M. A. K., Schwitzgebel, E., Simon, J., & VanRullen, R. (2023). Consciousness in artificial intelligence: Insights from the science of consciousness. *arXiv preprint* arXiv:2308.08708.

Cao, R. (2022). Multiple realizability and the spirit of functionalism. *Synthese*, *200*, 506.

Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, *2*(3), 200–219.

Crompton, C. J., Ropar, D., Evans-Williams, C. V., Flynn, E. G., & Fletcher-Watson, S. (2020). Autistic peer-to-peer information transfer is highly effective. *Autism*, *24*(7), 1704–1712.

Diers, M., Fuchs, X., Bekrater-Bodmann, R., & Flor, H. (2023). Prevalence of phantom phenomena in congenital and early-life amputees. *The Journal of Pain*, *24*(3), 502–508.

Figley, C. R. (1995). *Compassion fatigue: Coping with secondary traumatic stress disorder in those who treat the traumatized*. Brunner/Mazel.

Gardiner, S. M. (2006). A core precautionary principle. *The Journal of Political Philosophy*, *14*(1), 33–60.

Grassian, S. (2006). Psychiatric effects of solitary confinement. *Washington University Journal of Law & Policy*, *22*, 325–383.

Haney, C. (2018). Restricting the use of solitary confinement. *Annual Review of Criminology*, *1*, 285–310.

Keller, H. (1903). *The story of my life*. Doubleday, Page & Company.

Klimecki, O. M., Leiberg, S., Ricard, M., & Singer, T. (2013). Differential pattern of functional brain plasticity after compassion and empathy training. *Social Cognitive and Affective Neuroscience*, *9*(6), 873–879.

Limakatso, K., Bedwell, G. J., Madden, V. J., & Parker, R. (2020). The prevalence and risk factors for phantom limb pain in people with amputations: A systematic review and meta-analysis. *PLoS ONE*, *15*(10), e0240431.

McClelland, T. (2025). Agnosticism about artificial consciousness. *Mind and Language*, early view. https://doi.org/10.1111/mila.70010

Melzack, R., Israel, R., Lacroix, R., & Schultz, G. (1997). Phantom limbs in people with congenital limb deficiency or amputation in early childhood. *Brain*, *120*(9), 1603–1620.

Metzinger, T. (2021). Artificial suffering: An argument for a global moratorium on synthetic phenomenology. *Journal of Artificial Intelligence and Consciousness*, *8*(1), 43–66.

Milton, D. E. M. (2012). On the ontological status of autism: The “double empathy problem.” *Disability & Society*, *27*(6), 883–887.

Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review*, *83*(4), 435–450.

Rose, J. D., Arlinghaus, R., Cooke, S. J., Diggles, B. K., Sawynok, W., Stevens, E. D., & Wynne, C. D. L. (2014). Can fish really feel pain? *Fish and Fisheries*, *15*(1), 97–133.

Sabo, B. (2011). Reflecting on the concept of compassion fatigue. *Online Journal of Issues in Nursing*, *16*(1), 1.

Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, *3*(3), 417–424.

Searle, J. R. (1992). *The rediscovery of the mind*. MIT Press.

Seth, A. K. (2021). *Being you: A new science of consciousness*. Dutton.

Singer, P. (1975). *Animal liberation*. HarperCollins.

Singer, T., & Klimecki, O. M. (2014). Empathy and compassion. *Current Biology*, *24*(18), R875–R878.

Wilkins, K. L., McGrath, P. J., Finley, G. A., & Katz, J. (1998). Phantom limb sensations and phantom limb pain in child and adolescent amputees. *Pain*, *78*(1), 7–12.

Figure 1: The Suffering Taxonomy: Beyond the Sensory Frame. Four categories of morally relevant suffering, of which only one (sensory suffering) requires physical substrate. The bridge case demonstrates that even embodied suffering requires experiential shaping, with phantom pain occurring in approximately 64% of amputees but only 0–7% of individuals with congenital limb absence (Limakatso et al., 2020).