AI, Brain Death Detection, and Islamic Law

Muhammad Aurangzeb Ahmad [1]

Abstract

The deployment of machine learning systems capable of detecting covert consciousness in neurologically injured patients creates a profound challenge at the intersection of clinical medicine, AI ethics, and Islamic jurisprudence. We argue that the shift from binary clinical verdicts to probabilistic, temporally granular neural-state estimates should be addressed through three foundational constructs in Islamic legal epistemology: *bayyina* (clear evidentiary proof), *yaqīn* (epistemic certainty), and the theologically mandated agnosticism about the *rūḥ* (soul). We survey the current technical literature on AI-based consciousness detection, map it onto the landscape of Islamic brain death scholarship, and identifykey challenges. We also discuss its implicationsfor AI surrogate decision systems.

1. Introduction

In 2006, Owen et al. (2006) demonstrated that a patient in a behaviourally unresponsive vegetative state could follow commands by wilfully modulating her brain's metabolic activity as measured by functional MRI. This was a patient who, by every prior clinical criterion, had no detectable inner life. This landmark finding opened a field that has since revealed covert consciousness in up to 15–20% of patients formally diagnosed as unresponsive (Monti et al., 2010; Giacino et al., 2014; Schnakers et al., 2009). The domain has since been transformed by machine learning and AI. Deep neural networks trained on hundreds of thousands of EEG, electrocorticography (ECoG), and local field potential (LFP) samples now achieve high accuracy in predicting consciousness levels across diverse neurological populations, with results correlating significantly with gold-standard behavioural measures in held-out validation data (Toker et al., 2024; Bonanno et al., 2025). Multimodal approaches, fusing resting-state fMRI functional connectivity, quantitative EEG features, diffusion tractography, and PET metabolic imaging through classification models are now being evaluated in multi-centre European clinical trials (PerBrain Consortium, 2024).

This technical progress carries profound and unexamined consequences for ethical issues in end of life medical care, especially around brain death and organ transplantation. The AI community has discussed algorithmic fairness, data bias, and clinical-deployment ethics extensively, but has not engaged the world's second-largest religious tradition, Islam, on such questions. This paper makes the following contributions:

1. We survey the current technical literature on AI-based DoC detection for ML readers unfamiliar with its clinical stakes.
2. We map these developments onto the Islamic jurisprudential landscape on brain death.
3. We propose a cross-disciplinary research agenda for ML researchers and Islamic bioethicists.

A note on scope: this paper addresses Sunni jurisprudential frameworks. Shia maraji positions on brain death differ in important respects, most notably in the greater weight accorded to the heartbeat criterion and in the specific rulings of Grand Ayatollah Sīstani and others within the Twelver tradition. This would require separate treatment.

2. The Technical Landscape: What AI Now Detects

2.1. Disorders of Consciousness: Taxonomy and Stakes

Disorders of consciousness (DoC) refer to neurological conditions, typically resulting from severe acquired brain injury, in which the relationship between neural activity and behavioural responsiveness is disrupted. The principal diagnostic categories are as follows: (i) *coma*: absence of wakefulness and awareness; (ii) *unresponsive wakefulness syndrome* (UWS), formerly vegetative state: eyes open, sleep-wake cycles preserved, no behavioural evidence of awareness; (iii) *minimally conscious state* (MCS): reproducible but inconsistent behavioural signs of awareness;

[1]Department of Computer Science, University of Washington Bothell, Bothell, USA. Correspondence to: Muhammad Aurangzeb Ahmad <maahmad@uw.edu>.

Proceedings of the 43^{rd} International Conference on Machine Learning, Seoul, South Korea. PMLR 267, 2026.

and (iv) *cognitive motor dissociation* (CMD): intact neural command-following with no behavioural output.

Given that we are dealing with a problem where the stake are very high, misdiagnosis has profound consequences for withdrawal-of-treatment decisions, family communication, pain management, and organ-donation eligibility. Diagnostic error rates between UWS and MCS have historically been estimated at approximately 40% using bedside behavioural assessment alone (Andrews et al., 1996; Schnakers et al., 2009).

2.2. AI-Powered Detection: State of the Art

A 2025 scoping review of the DoC-ML literature identified 49,417 candidate articles across PubMed, Embase, Scopus, and Cochrane Library (Bonanno et al., 2025). High-quality studies span supervised ML (SVM, XGBoost, random forests), deep learning (CNNs, LSTMs, graph neural networks), and hybrid IoT-ML approaches for continuous bedside monitoring. The following is a high level overview of some of the apprahces that are used in this domain. It is meant to be non-exhasustive and we only discuss the represenattive example given the limitations of space.

EEG-based approaches: Quantitative EEG (qEEG) features, including microstate analysis, P300 signal detection, fractal-dimension analysis, and functional-connectivity metrics, combined with ML classifiers achieve high accuracy in DoC classification (Bonanno et al., 2025). Di Gregorio et al. (2022) showed that combining multiple EEG biomarkers with ML enhances predictive accuracy beyond any single biomarker.

Deep learning for consciousness detection: Toker et al. (2024) describe a generative-discriminative architecture in which deep convolutional neural networks (DCNNs), trained on over 680,000 EEG, ECoG, and LFP samples from humans and animals, are pitted against biophysically grounded dynamical brain models in an adversarial loop. The DCNN trained on acute TBI coma patients achieved statistically significant correlation with GCS scores in held-out validation ($p < 0.0001$), and with CRS-R scores in chronic DoC patients ($p < 0.0001$). It should be noted that the system produces *probabilistic continuous outputs* and not necessarily binary classifications.

Multimodal fusion: The EU-funded PerBrain Consortium is conducting multi-centre trials combining EEG, resting-state fMRI (rs-fMRI), diffusion MRI, anatomical MRI, and FDG-PET, analysed through classification models across three European sites (PerBrain Consortium, 2024). Chen et al. (2025) used rs-fMRI to measure inter-regional connectivity, identifying subtle neural patterns indicative of awareness in patients with DoC.

Brain-muscle network analysis: Fló et al. (2025) demonstrated that combining EEG with electromyography (EMG) and cardiac recordings through a network-analysis approach detects covert command-following, with heart activity and cortical power jointly predicting mental movement rehearsal. This is a promising direction for ICU bedside deployment.

TMS-EEG perturbational complexity. Transcranial magnetic stimulation combined with EEG can assess the capacity for consciousness via the perturbational complexity index (PCI) without requiring task performance. This provides a measurement independent of sensory processing or motor output (Casali et al., 2013).

2.3. Covert Consciousness

An important finding in recent years is the detection of *cognitive motor dissociation* (CMD). This refers to covert cortical command-following in patients whose bedside behavioural assessment indicates no awareness. Task-based fMRI and EEG can reveal preserved covert consciousness in up to 14–25% of UWS patients (Lo et al., 2024) (Bodien et al., 2024).

A positive result is strong evidence of awareness. However, a negative result cannot rule out awareness(Schnakers et al., 2020; Edlow & Fins, 2018). Researchers therefore proposed the term "covert cortical processing" rather than "covert consciousness" for patients with intact cortical responses to passive stimuli but no discernible active-task responses. This framing acknowledges an epistemic gap between neural measurement and subjective experience that no technology can close (Young et al., 2026).

The gap that exists here is the permanent underdetermination of subjective experience by neural measurement. This is a fundamental feature of the hard problem of consciousness (Chalmers, 1995) that neither deeper networks nor richer imaging modalities will dissolve. We revisit this issue in Section 4.

3. The Jurisprudential Landscape: Brain Death in Islamic Law

3.1. Classical Definition of Death

Classical Islamic jurisprudence defined death through observable cessation of breath (*nafas*) and heartbeat. These are the signs that are accessible to the senses, requiring no instrumentation. The *rūḥ* (spirit/soul) understood as the animating principle breathed by God into the human being (*wa nafakhtu fīhi min rūḥī*: "and I breathed into him of My spirit," Qur'an 15:29), departs at death and returns to God. Crucially, the Qur'an explicitly forecloses human knowledge of the *rūḥ*'s nature: "They ask you about the spirit. Say: the spirit is from the command of my Lord, and you have been given of knowledge only a little" (Qur'an

17:85). For Muslims, this also meant a type of agnosticism about the nature of the soul and a theological instruction and also the limits of human reason The eleventh-century scholar al-Ghazālī formulated the canonical position in Islamic thought: separation of the soul from the body is the end of its dominance over the body, and this event, not cardiac or respiratory function per se, constitutes death (Albar, 2012). Some Muslim scholars even argued that the soul's relationship to its substrate is therefore theologically prior to, and cannot be reduced to, any biological criterion, however sophisticated. Death in classical *fiqh* is thus *defined* by empirical signs while its *meaning*, the departure of the *rūḥ*, transcends empirical determination. This structural duality, latent in every pre-modern discussion of death, is important when we discuss the use of instrumentation in determination of life and death.

3.2. The Brain Death Controversy

When the Harvard Ad Hoc Committee formalised neurological criteria for death in 1968 (Beecher et al., 2007), Islamic scholars faced a new problem to address. The new criteria described a state in which the heart continued to beat, the body remained warm, and the organism appeared superficially alive; yet apnoea, absent brainstem reflexes, and (in some protocols) electrocerebral silence were taken to indicate the irreversible loss of the person. As a response to this challenge, three positions have emerged.

Position 1: Brain death is true death: The irreversible destruction of the brain's integrative capacity constitutes the loss of the human person as a whole. Albar (2012) and Pasha & Albar (2017) argue that the absence of *nafs* (personhood) and *nafas* (breath) in apnoeic coma constitutes departure of the *rūḥ*, and that the ventilated body is maintained artificially rather than alive in a meaningful sense. The Islamic Organisation of Medical Sciences (IOMS, 1986) and the Islamic Fiqh Academy of the Organisation of Islamic Cooperation (OIC, 1987) issued resolutions broadly aligned with this view, enabling transplantation *fatwās* across Gulf states and much of the Arab world. Notably, Saudi Arabia's High Committee on Brain Death requires EEG confirmation before establishing the diagnosis. This is an acknowledgement that the standard clinical examination is insufficient without additional neurophysiological corroboration (Albar, 2012).

Position 2: Brain death is not death: The heartbeat is the sign of life; a perfused body is a living body. Death occurs when the heart irreversibly stops. Rady & Verheijde (2018) argue that the secular concept of neurological death was constructed specifically to enable organ procurement and represents a novel imposition on Islamic communities that is at odds with the Qur'anic definition of death as biological disintegration. On this view, withdrawing ventilation from a brain-dead patient constitutes an act of killing, and many South Asian Deobandi and Barelvi scholars, as well as several Shia *marāji*^{c} including Grand Ayatollah Sīstānī, align with this position (Miller et al., 2014). Jurists in this camp further dispute whether the brain is the "seat" of the soul at all, a prerequisite for the whole-brain formulation of death (Miller et al., 2014). A full analysis of Shia positions, including the institutional *fatwa* landscape is beyond the scope of this paper.

Position 3: Deliberate agnosticism: A third position holds that the theological injunction to humility about the *rūḥ* (Qur'an 17:85) demands that Islamic jurisprudence not issue a binding ruling. Uncertainty is itself theologically appropriate; the diversity of *fatwā* positions should be preserved as principled epistemic humility rather than resolved by scholarly majority vote. This view underpins the refusal of some *fuqahā'* to issue any ruling on organ procurement from brain-dead patients, not because they reject it but because they hold the evidentiary threshold for such a determination to be inherently unattainable.

Current landscape. The problem of proof of consciousness and death has implications for things like organ donation. A systematic review of *fatwās* around this domain found that organ donation is broadly permitted within Islamic law, though conditions and scope vary significantly across schools and regions (Pasha & Albar, 2017). The Fiqh Council of North America (FCNA) issued a detailed ruling in 2018 concluding that brain death satisfies Islamic criteria provided specific clinical safeguards are met (Padela & Auda). The Islamic juridical deliberations around brain death largely took place over twenty-five years ago; the debates within Muslim bioethics require both updating and deepening with regard to these early rulings (Miller et al., 2014).

3.3. The Epistemological Framework

Islamic legal epistemology provides a graded taxonomy of certainty that directly governs what kinds of evidence can support what kinds of legal conclusions. The relevant grades are: *yaqīn* (certainty, knowledge admitting no doubt); *ẓann ghālib* (dominant probability, sufficient for most legal purposes under normal circumstances); *shakk* (doubt, which suspends or blocks action); and *wahm* (mere conjecture, which is legally ineffective). It should be noted that given the complexity and depth of Islamic jurisprudence, it is not possible to cover this topic given the limited space. We however give a high level overview of some of the concepts in Islamic *fiqh* relevant to our current discussion.

Bayyina iterally "that which makes clear", is the evidentiary standard for facts carrying grave legal consequences, including death certification, criminal conviction, and the dissolution of marriage. Classical *fiqh* construed *bayyina* narrowly as personal testimony by qualified witnesses; con-

temporary scholars have debated whether the concept can accommodate scientific evidence not based on human testimony, as discussed in Section 3.4. The principle of *iḥtiyāṭ* (precaution) requires that where a doubt concerning a life exists and cannot be resolved, the more protective conclusion must prevail. This operates as a structural constraint. Thus, t is not sufficient to show that brain death is probably present, one must show that the possibility of remaining life has been positively excluded.

Two further maxims are relevant. *Al-yaqīn lā yazūlu bil-shakk* (certainty is not removed by doubt) means that the established presumption of life persists until overturned by evidence of equivalent certainty. *Ḍarūra* (necessity) can relax evidentiary standards under acute exigency, and has been invoked to justify organ procurement but *ḍarūra* doctrines are bounded (*tuqaddar bi-qadrihā* necessity is assessed proportionately), and cannot be extended to routinise what began as an exceptional permission.

The question this paper raises is: *what evidentiary grade does an AI probabilistic consciousness score achieve?* As we argue in Section 4, the answer is that it achieves neither *yaqīn* nor *bayyina*. However, it may do more than merely generate *shakk*, because it provides a quantified, actionable estimate of residual awareness that the precautionary framework must take into account.

3.4. Precedents: When Islamic Law Has Engaged New Diagnostic Technology

The problem of new diagnostic technology disrupting established evidentiary categories is not novel to Islamic jurisprudence. Three precedents illuminate the structural challenge AI monitoring now poses, and suggest both the resources available to Islamic law and the limits of those resources.

Blood transfusion and the *darūra* model: Blood transfusion presented one of the first major intersections of modern medical technology and Islamic law. The Qur'an prohibits the consumption of blood (Qur'an 5:3). Scholars thus debated whether transfusion, blood introduced into the body by a different route, fell under this prohibition. Beginning in 1959, the Grand Mufti of Egypt and the Grand Mufti of Tunisia both issued *fatwās* permitting transfusion under *ḍarūra*: necessity permits what is otherwise prohibited when life is at stake (Al-Bar & Chamsi-Pasha, 2015). This set the template for Islamic bioethical engagement with new medical technology: (i) identify the relevant Qur'anic or *Sunna* prohibition; (ii) assess whether the new technology falls within its scope; (iii) if so, ask whether *ḍarūra* applies. The blood transfusion case resolved relatively cleanly because the question was binary i.e., permitted or not, and the life-saving benefit was direct and unconditional. AI consciousness monitoring does not resolve cleanly on this model: the technology does not save lives by its use. It may, under Islamic precautionary logic, *prevent* deaths by blocking premature certification, but this protective function conflicts with the *maṣlaḥa* of organ availability.

DNA evidence and the reconstruction of *bayyina*: The most instructive precedent is the long and contested debate over DNA evidence in Islamic family law. Classical Islamic law establishes paternity (*nasab*) through the marital bed (*farāsh*), voluntary acknowledgment, and in contested cases *bayyina* (witness testimony). The Qur'anic procedure of *liʿān* (mutual oath-swearing) provides a mechanism for disputed paternity that deliberately avoids biological certainty. DNA testing, which can establish biological paternity with near-certainty, raised the question of whether it could constitute *bayyina* or override *liʿān*.

The Islamic Fiqh Council of the Muslim World League (Mecca, 2002) issued a resolution acknowledging DNA testing as an effective scientific method yielding "certain or near-certain results," but specified that it should support *sharīʿa*-based methods rather than replace them (Shabana, 2012). The Malaysian National Fatwa Council (2012) similarly restricted DNA's role, favouring traditional methods. Several contemporary *fuqahā'* have argued, following Ibn Qayyim al-Jawziyya's expansive reading of *bayyina* as "anything that reveals the truth," that DNA should be admitted as an independent evidentiary basis (Al-Bar & Chamsi-Pasha, 2015). Others insist that DNA can function only as *qarīna* (circumstantial evidence) rather than *bayyina* proper, and cannot override established presumptions.

This debate is directly analogous to the AI consciousness monitoring question: A new scientific method produces near-certain biological information that disrupts an established legal framework designed around a different evidentiary basis. The DNA case shows that Islamic jurisprudence is capable of accommodating new evidence types but only after a sustained, contentious, multi-decade deliberative process, and with significant residual disagreement. It also shows the characteristic Islamic response: admit the new evidence as supplementary, not constitutive. Preserve the existing framework's authority over ultimate determinations and impose procedural conditions (government-authorised labs, multiple independent analyses, restrictions on who may request testing) that manage the technology's disruptive potential (Shariff et al., 2019). A doctrine of AI-based *bayyina* will need to develop analogous constraints if it is to be accepted.

EEG as confirmatory test: a partial precedent within brain death. There is a smaller-scale precedent internal to the brain death debate itself. Saudi Arabia's High Committee on Brain Death already insists on EEG as a confirmatory test before establishing the diagnosis (Albar, 2012). This represents an implicit jurisprudential judgment: the standard clinical examination alone does not produce *bayyina*

sufficient for death certification in an Islamic context; an additional neurophysiological test is required. This is noteworthy precisely because it shows that Islamic scholarly bodies have already, in effect, elevated the evidentiary threshold for brain death beyond what secular protocols require, not on clinical grounds but on epistemological ones. AI continuous monitoring can be understood as the logical extension of this precedent: if a flat-line EEG is required to supplement clinical signs, what is the status of an AI system that detects transient neural activity the EEG would miss? The Saudi requirement, far from resolving the AI question, sharpens it: the same precautionary logic that motivated the EEG requirement now generate a demand for AI monitoring that the EEG precedent cannot itself satisfy.

4. Structural Problems in AI Based Death Detection

4.1. The Probabilistic Output Problem

In section 2 we addressed disorders of consciousness for the patients who are not brain-dead. Brain death, by contrast, is a distinct clinical and legal determination: the irreversible cessation of all brain function, including the brainstem. While these are not the same population, the jurisprudential challenge this paper addresses arises at their intersection: AI monitoring systems developed for DoC classification are increasingly being evaluated as confirmatory tools in brain death protocols, and probabilistic outputs generated in that context attach to death certification rather than to treatment planning. The example that follows should be read in this light i.e., a residual probability of neural activity flagged during a brain death evaluation, not a DoC classification in a patient already known to be in MCS. Thus consider, a system might output: 87% probability of irreversible loss of integrative brain function; 13% probability of residual activity consistent with minimally conscious state. This is genuinely more informative and may be more accurate.

AI consciousness scores occupy an intermediate evidentiary position, above shakk but below bayyina. This is so because *Yaqin* (certainity) is unattainable for three independent reasons: first, the system produces a probability distribution, not a verdict, and no threshold transformation of a probability into a binary output preserves the certainty the concept requires. Second, the hard problem of consciousness (Section 4.3) means that even a perfect neural measurement would not establish the presence or absence of the *ruh* (soul). Third, current systems are trained predominantly on European and North American cohorts, introducing distributional uncertainty that compounds the first two. Yet the score exceeds mere *shakk* (doubt) in ways that matter legally. It is reproducible across independent evaluations, externally validated against gold-standard behavioural measures. Islamic legal history accommodates intermediate evidentiary categories: the DNA precedent shows that a new evidence type can be admitted as *qarina* (circumstantial corroboration) without achieving the status of bayyina proper, and without being dismissed as legally ineffective conjecture. AI consciousness scores may be candidates for an analogous category, subject to the procedural conditions discussed in Section 5.

The 13% residual probability needs to be thought through within an Islamic precautionary framework. The *iḥtiyāṭ* principle demands that doubt be resolved in favour of the more protective conclusion. The AI makes uncertainty *explicit and quantified*, which paradoxically may make precautionary withdrawal *less*, not more, permissible under *fiqh*. Epistemic uncertainty quantification, a feature widely advocated in trustworthy clinical AI precisely because it prevents false confidence and supports informed human oversight, here produces the opposite of its intended effect. It does so by making the residual probability of preserved consciousness legible as a number rather than leaving it as a vague clinical impression, the system transforms what was previously a matter of professional judgment (and therefore *iḥtiyāṭ*'s domain) into a documented, actionable quantity that the precautionary framework *cannot* set aside. In Islamic legal terms, an uncertain clinical impression might be navigated through the physician's discretion and the family's *walī* authority; a logged 13% posterior probability of awareness, attached to a certified death determination, is a different kind of object. It may generates an obligation to justify why the precautionary conclusion was not taken, and for which no existing doctrine of computational *bayyina* yet supplies that justification.

4.2. The Temporal Granularity Problem

Death in Islamic jurisprudence is a *moment*, an event with a before and an after. The *rūḥ* departs; the person is gone. Legal consequences are attach to this moment: inheritance distributes, marriage dissolves, organ retrieval becomes permissible. AI continuous monitoring introduces *temporal trajectories of neural state*. A monitoring system might show that what appeared as flat neural activity at 14:00 was preceded by transient fluctuations at 11:00 and followed by a possible perturbational response at 17:00. Death becomes not a discrete event but a process with uncertain boundaries and probabilistically detected waypoints.

4.3. The Hard Problem Cannot Be Dissolved

A central claim of Islamic metaphysics is that the *rūḥ* is breathed into the human being by God. The classical Islamic position on whether it has a biological substrate or not is left ambiguous. The ML community has implicitly adopted the working assumption of neural correlates of consciousness (NCC) that consciousness supervenes on, or is identical to,

certain neural processes. This is a *philosophical* assumption, not an established fact, and it is precisely the assumption that Islamic theology is agnostic about. AI consciousness scores cannot be interpreted as evidence about the *rūḥ*'s presence or absence.

4.4. The *Ijtihād* Speed Mismatch

Islamic jurisprudence develops through *ijtihād*, independent scholarly reasoning, and the issuance of *fatwās*. This process operates on timescales of months to years. The 2018 FCNA *fatwā* on organ donation resulted from two years of multidisciplinary deliberation (Padela & Auda). AI clinical capabilities are advancing on timescales of weeks to months. The PerBrain Consortium's multimodal results appeared as a preprint in November 2024; the 680,000-sample deep-learning consciousness-detection system in October 2024. Each advance potentially shifts the jurisprudential ground under existing *fatwās*. There is currently no institutional mechanism in any Islamic scholarly body for ongoing iterative engagement with rapidly evolving AI capabilities. The AI ethics community has begun to discuss "pacing" as a governance challenge (Dafoe, 2018), but exclusively within secular frameworks.

4.5. The AI Surrogate Problem

A related trajectory compounds the issues above. AI-based *patient preference predictors* (PPPs) i.e., systems that fine-tune large language models on a patient's prior decisions, values, and behavioural data to infer what they would have wanted when incapacitated, are now technically feasible and actively debated (Earp et al., 2024; Starke et al., 2025). Proponents argue they outperform human surrogates, whose accuracy at identifying p atient preferences is near chance (Refolo et al., 2025). Critics counter that clinical teams may treat probabilistic outputs as determinative in a domain that is "relational, existential, and culturally diverse" (Ahmad, 2025; Blumenthal-Barby et al., 2024).

PPPs are designed within the Western bioethical tradition of *individual autonomy*: their goal is to extend, computationally, the self-determination of an incapacitated patient. Islamic end-of-life authority is not located in the patient's preferences but in a tripartite structure of obligations: the *walī* (guardian), whose duty is to protect the patient's *maṣlaḥa* within Islamic law rather than to substitute the patient's judgment; the *tabīb* (physician), whose *amāna* (trusteeship) is independent of family or patient preferences; and the *faqīh* (jurist), whose consultation is required in complex cases (Larijani et al., 2024; Mohamad et al., 2025). The *walī* asks not "what would the patient have wanted?" but "what does Islamic law require, and what serves this patient's *maṣlaḥa*?" These are different questions that no preference-trained model can answer. The accountability gap further matters: Islamic ethics requires every morally significant act to be traceable to a responsible agent (*mukāllaf*) who bears *taklīf* before God. This is a chain that is obscured, not replaced, by algorithmic recommendations (Ahmad, 2025).

5. Implications and Research Agenda

5.1. For Machine Learning Researchers

Transparency about population-level epistemics. DoC papers should report population-level uncertainty distributions alongside aggregate performance metrics, particularly for patients near diagnostic boundaries. **Cross-cultural validation.** Large-scale training datasets are predominantly drawn from European and North American clinical populations (PerBrain Consortium, 2024). Muslim-majority populations are systematically underrepresented in the training data of systems that will be deployed in their healthcare systems, a specific instance of the distribution-shift and representational fairness problems documented in medical AI (Ahmad, 2025). **Interpretability in high-stakes settings.** Deep learning approaches offer "superior predictive power but with higher cost of interpretability" (Bonanno et al., 2025). A decision that cannot be explained cannot be adjudicated within Islamic legal process. Investment in interpretable architectures is therefore a jurisprudential as well as a technical priority. **Acknowledgement of the hard problem.** ML consciousness-detection papers should explicitly state that their systems measure neural correlates of consciousness, not consciousness itself, and that the relationship between these is a philosophical working assumption, not an established fact.

5.2. For Islamic Bioethicists

Direct engagement with the technical literature. Most Islamic bioethics papers on brain death cite clinical criteria that has been superceded (Pasha & Albar, 2017). To the best of our knowledge, the CMD literature, the PerBrain trial results, and the deep-learning systems described above have not entered Islamic bioethics discourse. **A doctrine of computational *bayyina*.** A developed doctrine governing probabilistic machine outputs, with explicit conditions for sufficiency, weighting, and mandatory precautionary adjustments, may be needed to address problems like the one discussed here. **Institutional mechanisms for rapid *ijtihād*.** Existing Islamic scholarly bodies should establish standing technical advisory committees with ongoing access to ML researchers and clinical neurologists.

6. Broader Significance

The ML community's engagement with value-aligned and culturally situated ethics (Busch et al., 2023; Ersoy, 2025) has remained predominantly within Western philosophical

frameworks: principlism, contractarianism, utilitarian welfare metrics. The Islamic ethical tradition represents not merely a different set of cultural values to be accommodated, but a different epistemological framework that challenges some working assumptions of the ML ethics enterprise itself. The Qur'an's instruction that the *rūḥ* is beyond human determination is a philosophically serious position about the underdetermination of soul by substrate and the category error involved in treating ultimate questions about personhood as tractable by measurement. The hard problem of consciousness which neuroscience cannot solve (Chalmers, 1995) represents, in secular philosophical terms, the same epistemic gap that Islamic theology names with the concept of *ghayb* (the unseen).

7. Conclusion

AI-powered disorders-of-consciousness detection is a clinical diagnostic advance. In the context of brain death determination, it is a jurisprudential event as well. It complicates existing Islamic legal rulings, challenges traditional Islamic epistemic standards for death evidence. The AI community has a responsibility here that is not covered by standard algorithmic fairness analysis. Researchers must engage the jurisprudential literature directly and acknowledge that their systems are being deployed into a world where the most consequential interpretations of their outputs will be made within frameworks their models cannot address.

References

Ahmad, M. A. Algorithmic fairness in ai surrogates for end-of-life decision-making. *arXiv preprint arXiv:2510.16056*, 2025.

Al-Bar, M. A. and Chamsi-Pasha, H. *Contemporary Bioethics: Islamic Perspective*. Springer International Publishing, Cham, 2015.

Albar, M. Organ transplantation: a sunni islamic perspective. *Saudi Journal of Kidney Diseases and Transplantation*, 23(4):817–822, 2012.

Andrews, K., Murphy, L., Munday, R., and Littlewood, C. Misdiagnosis of the vegetative state: retrospective study in a rehabilitation unit. *Bmj*, 313(7048):13–16, 1996.

Beecher, H. K. et al. A definition of irreversible coma. *International Anesthesiology Clinics*, 45(4):113–119, 2007.

Blumenthal-Barby, J., Fletcher, F. E., Taylor, L., Nelson, R. H., Moore, B., Saloner, B., and Ubel, P. A. Ethical complexities in utilizing artificial intelligence for surrogate decision making, 2024.

Bodien, Y. G., Allanson, J., Cardone, P., Bonhomme, A., Carmona, J., Chatelle, C., Chennu, S., Conte, M., Dehaene, S., Finoia, P., et al. Cognitive motor dissociation in disorders of consciousness. 391(7):598–608, 2024.

Bonanno, M., Cardile, D., Liuzzi, P., Celesti, A., Micali, G., Corallo, F., Quartarone, A., Tomaiuolo, F., and Calabrò, R. S. Can artificial intelligence improve the diagnosis and prognosis of disorders of consciousness? a scoping review. *Frontiers in Artificial Intelligence*, 8:1608778, 2025.

Busch, F., Keller, S., Rueger, C., Kader, A., Ziegeler, K., Bressem, K. K., and Adams, L. C. Mapping gender and geographic diversity in artificial intelligence research: editor representation in leading computer science journals. *Acta Radiologica Open*, 12(10):20584601231213740, 2023.

Casali, A. G., Gosseries, O., Rosanova, M., Boly, M., Sarasso, S., Casali, K. R., Casarotto, S., Bruno, M.-A., Laureys, S., Tononi, G., et al. A theoretically based index of consciousness independent of sensory processing and behavior. *Science translational medicine*, 5(198): 198ra105–198ra105, 2013.

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

Chen, S., Wang, L., Wang, Y., and Yang, Z. From scales to circuits: integrating behavioral diagnosis and neural biomarkers for improved classification in disorders of consciousness. *Frontiers in Neuroscience*, 19:1725420, 2025.

Dafoe, A. AI governance: A research agenda. Technical report, Future of Humanity Institute, University of Oxford, 2018.

Di Gregorio, F., La Porta, F., Petrone, V., Battaglia, S., Orlandi, S., Ippolito, G., Romei, V., Piperno, R., and Lullini, G. Accuracy of eeg biomarkers in the detection of clinical outcome in disorders of consciousness after severe acquired brain injury: preliminary results of a pilot study using a machine learning approach. *Biomedicines*, 10(8):1897, 2022.

Earp, B. D., Porsdam Mann, S., Allen, J., Salloch, S., Suren, V., Jongsma, K., Braun, M., Wilkinson, D., Sinnott-Armstrong, W., Rid, A., et al. A personalized patient preference predictor for substituted judgments in healthcare: technically feasible and ethically desirable. *The American Journal of Bioethics*, 24(7):13–26, 2024.

Edlow, B. L. and Fins, J. J. Assessment of covert consciousness in the intensive care unit: clinical and ethical considerations. *The Journal of head trauma rehabilitation*, 33(6):424–434, 2018.

Ersoy, C. Inclusive transformation with artificial intelligence. *Advances in Computational Intelligence and Robotics*, pp. 275–296, 2025.

Fló, E., Fraiman, D., and Sitt, J. D. Assessing brain-muscle networks during motor imagery to detect covert command-following. *BMC medicine*, 23(1):68, 2025.

Giacino, J. T., Fins, J. J., Laureys, S., and Schiff, N. D. Disorders of consciousness after acquired brain injury: the state of the science. *Nature Reviews Neurology*, 10 (2):99–114, 2014.

Larijani, B., Mobasher, M., Zahedi, F., and Tahmasebi, M. Ethical decision-making regarding life sustaining treatment in end-of-life care: a scoping review of the similarities and differences between two viewpoints. *Archives of Academic Emergency Medicine*, 13(1):e17, 2024.

Lindenmeyer, A., Blattmann, M., Franke, S., Neumuth, T., and Schneider, D. Towards trustworthy ai in healthcare: epistemic uncertainty estimation for clinical decision support. *Journal of Personalized Medicine*, 15(2):58, 2025.

Lo, C. C. H., Woo, P. Y. M., and Cheung, V. C. Task-based eeg and fmri paradigms in a multimodal clinical diagnostic framework for disorders of consciousness. *Reviews in the Neurosciences*, 35(7):775–787, 2024.

Miller, A. C., Ziad-Miller, A., and Elamin, E. M. Brain death and islam: the interface of religion, culture, history, law, and modern medicine. *Chest*, 146(4):1092–1101, 2014.

Mohamad, M. I. B. K., Hamdan, M. N., and Yusof, A. N. M. Third-party refusal of medical treatment–a critical analysis of case report from islamic ethical perspectives. *BMC Medical Ethics*, 26(1):48, 2025.

Monti, M. M., Vanhaudenhuyse, A., Coleman, M. R., Boly, M., Pickard, J. D., Tshibanda, L., Owen, A. M., and Laureys, S. Willful modulation of brain activity in disorders of consciousness. *New England journal of medicine*, 362 (7):579–589, 2010.

Nayak, B. S. and Walton, N. Artificial intelligence (ai) and ideology. *Information Technology & People*, pp. 1–19, 2026.

Owen, A. M., Coleman, M. R., Boly, M., Davis, M. H., Laureys, S., and Pickard, J. D. Detecting awareness in the vegetative state. *science*, 313(5792):1402–1402, 2006.

Padela, A. and Auda, J. The moral status of organ donation and transplantation within islamic law. *the Fiqh council of North America's position*, 2020:6.

Pasha, H. C. and Albar, M. A. Do not resuscitate, brain death, and organ transplantation: Islamic perspective. *Avicenna journal of medicine*, 7(02):35–45, 2017.

PerBrain Consortium. Integrative electrophysiology and neuroimaging approach in assessing disorders of consciousness: A multimodal multicentric machine learning study. medRxiv, 2024.

Rady, M. Y. and Verheijde, J. L. Legislative enforcement of nonconsensual determination of neurological (brain) death in muslim patients: A violation of religious rights. *Journal of religion and health*, 57(2):649–661, 2018.

Refolo, P., Sacchini, D., Raimondi, C., Masilla, S. S., Corsano, B., Mercuri, G., Oliva, A., and Spagnolo, A. G. Should artificial intelligence-based patient preference predictors be used for incapacitated patients? a scoping review of reasons to facilitate medico-legal considerations. In *Healthcare*, volume 13, pp. 590. MDPI, 2025.

Schnakers, C., Vanhaudenhuyse, A., Giacino, J., Ventura, M., Boly, M., Majerus, S., Moonen, G., and Laureys, S. Diagnostic accuracy of the vegetative and minimally conscious state: clinical consensus versus standardized neurobehavioral assessment. *BMC neurology*, 9(1):35, 2009.

Schnakers, C., Hirsch, M., Noé, E., Llorens, R., Lejeune, N., Veeramuthu, V., De Marco, S., Demertzi, A., Duclos, C., Morrissey, A.-M., et al. Covert cognition in disorders of consciousness: a meta-analysis. *Brain Sciences*, 10 (12):930, 2020.

Shabana, A. Paternity between law and biology: The reconstruction of the islamic law of paternity in the wake of dna testing. *Zygon®*, 47(1):214–239, 2012.

Shariff, A. A. M., Azhar, A., Rajamanickam, R., Manap, N. A., Hussein, S. M., Said, M. H. M., Ab Halim, A. H., and Markom, R. Analysis on admissibility of dna evidence in malaysian syariah courts. *Academic Journal of Interdisciplinary Studies*, 8(4):159–159, 2019.

Starke, G., Schopp, L., Meier, C., Baffou, J., Thanou, D., Maurer, J., and Jox, R. J. Machine learning–based patient preference prediction: a proof of concept. *NEJM AI*, 2 (10):AIoa2500265, 2025.

Toker, D., Zheng, Z. S., Thum, J. A., Guang, J., Annen, J., Miyamoto, H., Yamakawa, K., Vespa, P. M., Laureys, S., Schnakers, C., et al. Ai discovery of mechanisms of consciousness, its disorders, and their treatment. *bioRxiv*, pp. 2024–10, 2024.

Young, M. J., Nagel, T., Mudrik, L., Roskies, A., Diringer, M., and Fischer, D. Covert consciousness: epistemic limits of its definition and detection. *Neurocritical Care*, 44(2):506–517, 2026.