




SvedbergOpen
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: <https://www.svedbergopen.com/>



Research Paper

Open Access

Algorithmic Conceptions: Regulating Artificial Intelligence in Assisted Reproductive Technologies through an Ethical and Legal Lens

Dr. Saif Rasul Khan^{1*}, Dr. Jafreen Haque², Ms. Shruti Hazarika Borpujari³

¹Assistant Professor, Royal School of Law and Administration (RSLA), The Assam Royal Global University, Guwahati, Assam, India.
E-mail: saifrasul khan@gmail.com

²Assistant Professor, Centre for Legal Studies, Girijananda Chowdhury University, Guwahati, Assam, India.
E-mail: jafreenhaque24@gmail.com

³Assistant Professor, Royal School of Law and Administration (RSLA), The Assam Royal Global University, Guwahati, Assam, India.
E-mail: shrutihazarika18@gmail.com

ABSTRACT

Artificial intelligence is increasingly integrated into assisted reproductive technology (ART), with regulator-approved systems now assessing embryo viability through time-lapse imaging and multimodal clinical data. This review investigates whether current legal frameworks sufficiently regulate algorithmic decision-making in reproductive practices. Employing a comparative doctrinal methodology, it examines the European Union's Artificial Intelligence Act (Regulation (EU) 2024/1689), the General Data Protection Regulation, the Medical Device Regulation, and India's Assisted Reproductive Technology (Regulation) Act, 2021, alongside relevant technical literature on operational embryo-assessment systems. The analysis reveals a structural deficiency: each regulatory framework treats algorithmic embryo selection as incidental to other primary regulatory objectives, namely, product safety, data processing, or clinic licensing, and thus fails to address the technology's reproductive and intergenerational implications directly. The review advocates a rights-based governance model grounded in algorithmic transparency, compulsory human oversight, bias auditing, sui generis reproductive-data protection, and a legally enforceable right to an explanation.

Keywords: Artificial Intelligence; Assisted Reproductive Technology; AI Act; Algorithmic Accountability; Bioethics; Data Protection; Embryo Selection; Reproductive Autonomy; Right to Explanation; Comparative Law

1. Introduction

Artificial intelligence (AI) has moved from the margins of assisted reproductive technology (ART) into the clinical decisions that precede conception. Fertility clinics now use regulator-cleared algorithms to assess embryo viability from time-lapse images, predict in vitro fertilisation (IVF) outcomes, and optimise hormonal stimulation protocols. In September 2025, the United States Food and Drug Administration granted its first 510(k) clearance for a machine-learning embryo-assessment tool, Alife Health's Embryo Predict, based on a prospective, randomised, multicentre trial involving 440 patients at seven clinics (Fertility Bridge, 2026). Fairtility's CHLOE Blast received comparable clearance that month. At the same time, time-lapse systems such as Vitrolife's iDAScore and AIVF's EMA are already widely integrated into incubator workflows (see review in Nicolielo et al., 2025). By reducing diagnostic uncertainty and clinician subjectivity to a single embryo ranking, these systems influence which potential life is transferred and which is not.

Concentrating such consequential decisions in a single algorithmic score exposes a gap that existing regulatory frameworks were not designed to address. Medical-device law focuses on product safety, data-protection law on lawful data processing, and ART legislation on clinic licensing. Considered separately, none fully addresses the reproductive and intergenerational implications of ranking embryos through probabilities generated by opaque models. This review examines that gap through comparative legal and ethical analysis. It has three objectives: to identify the ethical challenges posed by algorithmic decision-making in assisted reproduction; to assess how effectively the European Union's Artificial Intelligence Act, India's Assisted Reproductive Technology (Regulation) Act, 2021, and the General Data Protection Regulation (GDPR) respond to those challenges; and to propose a governance model that addresses the shortcomings revealed by the comparison.

2. Methodology

This review adopts a doctrinal and comparative-legal methodology rather than an empirical or quantitative design; no primary clinical, patient-level, or experimental data were collected, and the review accordingly did not require institutional ethics or animal-subject approval. Primary legal sources were read directly from the promulgating instruments: Regulation (EU) 2024/1689 (the Artificial Intelligence Act) and its Annexes I and

III; Regulation (EU) 2016/679 (the General Data Protection Regulation); Regulation (EU) 2017/745 (the Medical Device Regulation); and the Assisted Reproductive Technology (Regulation) Act, 2021 (Act No. 42 of 2021, India). These instruments were analysed against technical and clinical literature describing currently deployed or regulator-cleared AI systems in embryo assessment, drawn from peer-reviewed sources and publicly reported regulatory-clearance records for the period 2023–2026. Where a claim originates in vendor-published or promotional material rather than independently peer-reviewed evidence, it is identified as such in the text and treated as illustrative of industry practice rather than as validated clinical evidence, a distinction that itself bears on the regulatory analysis in Section 5.

3. AI-Enabled Systems In Assisted Reproduction: A Technical Landscape

AI-enabled embryo assessment rests on two main technical approaches. The first, exemplified by Vitrolife's iDAScore and Fairtility's CHLOE platform, applies convolutional neural networks to morphokinetic time-lapse video captured inside the incubator, learning to associate frame-level developmental patterns with implantation or live-birth outcomes and outputting a continuous viability score (see review in Nicolielo et al., 2025). The second, exemplified by AIVF's EMA system and by the IVFormer architecture described by Wang et al. (2024) in the journal *Patterns*, extends this further through multimodal contrastive learning: a transformer-based backbone (IVFormer), trained via a self-supervised visual-temporal contrastive framework (VTCLR) on large unlabelled datasets, integrates embryo imaging with structured clinical variables across the full IVF cycle rather than a single time point. Wang et al. (2024) report that their pre-trained model outperformed physician assessment across evaluated categories for euploidy ranking, though this claim, like most in the field, rests substantially on retrospective validation rather than the prospective randomised design that regulators increasingly require.

That distinction is not academic. Alife Health's Embryo Predict is, to date, the only embryo-assessment tool cleared on the basis of a prospective multicentre randomised trial (Fertility Bridge, 2026); most competing claims of improved implantation or pregnancy rates circulate instead through vendor websites and single-clinic physician testimonials that have not been independently peer-reviewed. A commonly cited example, a fertility clinic's reported correlation between a proprietary AI score threshold and a specific implantation rate, appears on the vendor's own marketing material rather than in a peer-reviewed journal (AIVF, 2026). A comparatively rare example of prospective, peer-reviewed validation outside the regulatory-clearance route is Brazil's MAIA platform: trained on 1,015 embryo images and then evaluated prospectively across three fertility centres, it achieved 66.5% overall accuracy in predicting clinical pregnancy from 200 single-embryo transfers, rising to 70.1% in elective-transfer cases where more than one embryo was available for selection (Nicolielo et al., 2025). The gap between that disclosed, peer-reviewed accuracy figure and an undated marketing claim is itself a governance problem for a technology that ranks which embryo is transferred, addressed further in Section 5.

4. Ethical And Legal Challenges

4.1 Reproductive Autonomy and Informed Consent

In clinical medicine, informed consent is not satisfied by a signed form or a generic description of treatment. It requires a dialogue in which the clinician takes reasonable care to ensure that the patient understands material risks and reasonable alternatives. Under *Montgomery v Lanarkshire Health Board* [2015] UKSC 11, a risk is material when a reasonable person in the patient's position would likely attach significance to it, or when the clinician knows or ought reasonably to know that the particular patient would do so. Professional guidance similarly treats consent as an ongoing process of shared decision-making and requires patients to receive information they can understand, together with time and support to consider it (General Medical Council, 2020). In IVF, this duty should extend beyond consent to ovarian stimulation, retrieval, fertilisation, and transfer. Where an algorithm materially influences which embryo is transferred, the basis and limitations of that recommendation form part of the decision the patient is being asked to authorise.

Algorithmic embryo ranking creates a specific disclosure problem because the score may appear precise while concealing substantial uncertainty. Two embryos with similar morphological grades may receive different rankings because the model has detected a complex temporal pattern in time-lapse images or combined imaging with clinical variables. Yet a clinician may be unable to translate that pattern into a stable, patient-comprehensible reason for preferring one embryo. Research on explainability in AI-assisted embryo selection identifies this opacity as a barrier to clinical acceptance. It stresses that explanations must be useful to both clinicians and patients, rather than merely technically available to developers (Urcelay et al., 2023). The difficulty intensifies when the model has been validated retrospectively, trained on a population unlike the patient, or designed to predict a proxy such as implantation rather than the outcome the patient is likely to value most, such as live birth. In those circumstances, disclosing a score without disclosing its target outcome, evidential limits, and uncertainty may create an impression of objectivity that the underlying evidence does not support.

The relevant question is therefore not whether every parameter of a deep-learning model must be explained. That standard would often be technically impossible and clinically unhelpful. The legally material information

is functional: that AI was used; what outcome the model predicts; what data it considered; whether it has been validated in comparable patients and clinical settings; how accurate and uncertain its prediction is; whether an embryologist reviewed or departed from the ranking; and what reasonable alternatives remain, including selection by conventional morphological assessment or additional human review. Patients may also attach significance to whether the system is advisory or effectively determinative, whether commercial interests shape its use, and whether declining algorithmic assessment would alter access, cost, or treatment timing. The World Health Organization's guidance on AI for health supports this approach by linking autonomy protection to transparency, explainability, intelligibility, and accountability across the AI lifecycle (World Health Organization, 2021).

Consent may therefore be valid for IVF and embryo transfer in general while remaining deficient in relation to the particular embryo-selection process. Framed in this way, the problem is not a wholly new category of consent but a disclosure gap within established materiality doctrine. A proportionate remedy would require a plain-language explanation to accompany the AI output and be discussed before transfer. That explanation should identify the model's intended use; the principal factors and clinical inputs informing the score, where these can be meaningfully described; the limits of validation; the role of human judgment; and the consequences of relying on or rejecting the recommendation. The explanation should also be documented in the clinical record, together with the model version used and any clinician override, so that consent and accountability can later be assessed. As proposed in Section 6, such a requirement would adapt existing informed-consent principles to AI-mediated reproductive decisions without demanding disclosure of source code or creating an entirely separate consent regime.

4.2 Data Privacy and Genetic Sensitivity

Data generated through artificial intelligence is not simply categorised as 'health data' under most data protection frameworks. Embryo imaging, genetic screening results, and gamete-donor records have implications not only for the individuals directly involved but also for the resulting child and their descendants, all of whom were never in a position to consent to their collection. The General Data Protection Regulation (GDPR) designates genetic and health data as 'special category' data pursuant to Article 9, thereby imposing more stringent processing requirements. Furthermore, the regulation considers algorithmic embryo ranking as a form of automated decision-making with legal or similarly significant effects under Article 22, which grants the data subject the right to human intervention and an explanation of the underlying logic. However, both provisions define the data subject as the patient undergoing treatment; they do not account for the later interests of the resulting child in the data used to select the embryo from which they developed. Currently, this intergenerational interest lacks a clear rights holder under Articles 9 and 22. This review addresses this specific, well-defined gap by proposing reproductive data as a sui generis protected category in Section 6, rather than treating it as an ordinary subset of health data.

4.3 Algorithmic Bias and Equity

Embryo-assessment models are trained predominantly on records generated by patients who can access advanced fertility care in the health systems where the tools are developed and validated. Those datasets are often concentrated in a limited number of high-volume clinics in North America, Europe, and East Asia. They may encode local referral patterns, laboratory protocols, incubator conditions, imaging equipment, embryo-grading practices, and patient eligibility criteria. A viability score is therefore not a universal measure of embryo quality; it is a probability learned from a particular clinical environment. When the model is deployed in a clinic serving patients or using workflows that differ from that environment, distribution shift may reduce its reliability. The resulting error need not appear as an obviously discriminatory rule. It may instead emerge through lower calibration or ranking accuracy for patients of different ethnic backgrounds, maternal ages, ovarian-response profiles, infertility diagnoses, treatment histories, or gamete sources. Because embryo selection is comparative, even a modest subgroup-specific error can have cumulative consequences: embryos from one group may be consistently ranked lower, apparently precise scores may falsely reassure clinicians, and patients may undergo additional transfer cycles before selecting a viable embryo.

Unlike the autonomy and privacy concerns discussed above, this risk is empirically testable and should be governed through measurable performance obligations. Premarket evaluation should disclose the composition and provenance of training, tuning, and validation datasets; report missingness and exclusion criteria; and assess discrimination, calibration, sensitivity, specificity, false-positive and false-negative rates, and ranking performance across prespecified demographic and clinical subgroups. Validation should also be external and multicentre, using data from clinics, equipment, and patient populations not represented in model development, with uncertainty intervals reported so that small subgroup samples are not mistaken for evidence of equivalence. Where sample sizes are insufficient, the limitation should trigger restricted indications, additional data collection, or heightened human review rather than an unsupported claim of fairness. Clearance should further require a postmarket monitoring plan capable of detecting performance drift after changes in laboratory practice, patient mix, imaging systems, or model updates. Audit results should be available to regulators and, in an intelligible form, to clinics and patients. Such requirements convert fairness from a broad ethical aspiration into a verifiable condition of clinical use.

4.4 Commodification, Eugenics, and Accountability

Ranking embryos by a numerical probability of success resembles, but is not equivalent to, the selection logic historically associated with eugenics. The relevant distinction lies in the purpose and evidentiary basis of the ranking. A model limited to predicting implantation or live birth seeks to support a defined therapeutic objective within an existing IVF cycle; a model that ranks embryos according to traits unrelated to clinical viability instead moves toward selecting the characteristics of a future child. International bioethics instruments already recognise that this boundary has human-rights significance. UNESCO's Universal Declaration on the Human Genome and Human Rights places developments concerning the human genome within a framework of dignity, equality, and responsibility to future generations (UNESCO, 1997). More specifically, Article 14 of the Council of Europe's Convention on Human Rights and Biomedicine prohibits the use of medically assisted procreation to choose a future child's sex, except to avoid a serious hereditary sex-related disease (Council of Europe, 1997). Although these instruments do not directly govern AI-based embryo scoring, they show that reproductive selection is not ethically neutral merely because it is performed through prediction rather than direct genetic intervention.

Current commercial embryo-assessment systems are generally presented as tools for predicting implantation, euploidy, clinical pregnancy, or live birth. Yet the technical architecture does not itself enforce that clinical limit. The same imaging and multimodal pipelines can be retrained to predict any label sufficiently represented in the data. At the same time, genomic inputs may enable inference of characteristics correlated with, but not necessary for, reproductive success. Scope creep could occur incrementally: a clinic might first add a secondary score marketed as improving personalised treatment, then allow that score to influence embryo ranking, and eventually normalise selection for traits whose clinical relevance is weak or contested. The concern is therefore not that existing viability tools are already eugenic, but that a system designed to optimise one outcome can be repurposed without a corresponding change in its visible interface, institutional setting, or regulatory classification. The WHO's governance recommendations for human genome editing emphasise robust oversight, registries, international collaboration, public engagement, and mechanisms for identifying unsafe or unethical activity; those principles are instructive here because embryo-selection algorithms also create consequences that may extend beyond the immediate patient and across generations (WHO Expert Advisory Committee, 2021).

This possibility makes accountability inseparable from explainability. Deep-learning scores often cannot be reduced to a stable rule that a clinician, regulator, or ombudsperson can reconstruct after a disputed transfer decision. Research on AI-assisted embryo selection accordingly identifies opacity as a barrier to clinical acceptance and calls for interpretability tailored to clinicians' and patients' needs (Urcelay et al., 2023). A general statement that the system analysed embryo morphology is insufficient: meaningful accountability requires a record of the model version used, the outcome it was trained to predict, the variables and data sources considered, the relative weight assigned to AI and embryologist judgment, and any manual departure from the ranking. It also requires periodic inspection for changes in labels, input features, intended use, and commercial claims. Without these safeguards, a provider could expand a model from viability assessment to trait-oriented prediction while still presenting it as ordinary clinical decision support.

An enforceable right to explanation should therefore operate alongside substantive limits on permissible objectives. Regulators should require developers to specify the target outcome in advance, prohibit undisclosed secondary optimisation, and treat adding a new predicted trait or genomic input as a material change requiring renewed review. Clinics should disclose to patients which outcomes the algorithm predicts, whether the score is advisory or determinative, and how non-AI clinical factors affected the final choice. Audit logs should be retained so that an independent reviewer can determine whether the system remained within its authorised purpose. These measures would not eliminate every normative dispute about embryo selection. Still, they would make the shift from viability ranking to trait selection visible, contestable, and legally attributable, rather than allowing it to occur through an opaque update to an otherwise familiar clinical tool.

5. COMPARATIVE REGULATORY ANALYSIS

The three instruments examined here do not compete to regulate a single legal object. Each approaches algorithmic embryo selection through a different jurisdictional trigger: the EU Artificial Intelligence Act through the placing on the market or use of specified AI systems, the GDPR through the processing of personal data, and India's Assisted Reproductive Technology (Regulation) Act, 2021 through the registration and supervision of ART clinics and banks. Their provisions may apply concurrently to the same clinical episode, but they allocate duties to different actors, pursue different regulatory purposes, and activate at different stages of the technology's lifecycle. The governance gap therefore arises less from complete silence than from fragmentation: no instrument treats the embryo-ranking decision itself—its evidentiary basis, permissible objective, effect on embryo choice, and implications for a future child—as the primary object of regulation.

Under the EU AI Act (Regulation (EU) 2024/1689), the most plausible route to high-risk classification is derivative. Article 6(1), read with Annex I, applies where the AI system is itself a product, or a safety component of a product, covered by listed European Union harmonisation legislation and that product is required to undergo third-party conformity assessment before being placed on the market or put into service (European Parliament and Council of the European Union, 2024). The Medical Device Regulation (Regulation (EU)

2017/745) and the In Vitro Diagnostic Medical Devices Regulation (Regulation (EU) 2017/746) fall under that route. Embryo-assessment software would therefore attract the AI Act's high-risk requirements only if its intended purpose and risk classification first bring it within the MDR or IVDR and trigger notified-body assessment.

That antecedent inquiry is consequential because European medical-device law classifies software by its manufacturer-defined intended purpose and by the significance of the information it provides for diagnosis or therapeutic decisions. The Medical Device Coordination Group's software guidance explains that qualification, classification, modularity, and changes to medical-device software must be assessed under the MDR and IVDR, with Rule 11 of the MDR addressing software that supplies information used for diagnostic or therapeutic decisions (Medical Device Coordination Group, 2025). A viability-ranking tool expressly intended to guide embryo transfer is therefore likely to be treated as medical-device software, but classification remains fact-sensitive; it cannot be assumed solely from the presence of machine learning. Conversely, software presented as workflow support, research functionality, or non-medical information may be argued to fall outside the medical-device route, even when clinicians rely on its output in practice. The regulatory object remains the marketed product and its intended purpose, not every consequential use of a score at the bedside.

Even when the high-risk regime applies, its core duties—risk management, data governance, technical documentation, logging, transparency to deployers, human oversight, accuracy, robustness, cybersecurity, conformity assessment, and post-market monitoring—are horizontal controls on the AI system. They can improve the reliability and traceability of embryo-ranking software. Still, they do not determine which reproductive outcomes may lawfully be optimised, what explanation a patient choosing between similarly graded embryos is owed, or whether intergenerational interests require distinct protection. The AI Act thus supplies an important compliance architecture without becoming a complete law of algorithmic reproductive decision-making.

The GDPR regulates a different object: the collection, inference, use, disclosure, and retention of personal data. Embryo images linked to a patient, treatment records, donor information, and genetic test results may engage the GDPR's general principles and its heightened protection for health and genetic data under Article 9. The framework can require a lawful basis, purpose limitation, data minimisation, security, transparency, data-protection impact assessment, and enforceable data-subject rights. These obligations reach important parts of the pipeline—from training-data acquisition to clinical deployment. Still, they evaluate whether data processing is lawful, fair, and transparent rather than whether the embryo-ranking criterion is clinically or ethically permissible.

Article 22 requires particular care. Its protection concerns decisions based solely on automated processing that produce legal effects or similarly significant effects for the data subject; where an exception applies, safeguards include the ability to obtain human intervention, express a view, and contest the decision. Official guidance emphasises that meaningful human involvement must be genuine, not tokenistic (Article 29 Data Protection Working Party, 2018). In embryo selection, Article 22 may apply if the system effectively determines a patient's treatment without substantive clinical review. It may not apply where an embryologist independently evaluates the embryos and retains real authority to accept or reject the ranking. The provision is therefore not a general right to an explanation whenever AI contributes to care, and its applicability turns on the degree of automation, the significance of the effect on the patient, and the reality of human oversight.

A further limitation is temporal and relational. The treated patient, donor, or other identifiable living person may be a data subject. Still, the GDPR does not readily represent a child's later interest in the data and decision-making process through which a particular embryo was selected before that child existed as an identifiable rights holder. Nor does it assign rights to an embryo as such. The patient's rights may indirectly protect aspects of that future interest, and once born, the child may acquire rights in personal data relating to them. Still, the framework does not create a distinct intergenerational claim to the provenance, logic, or long-term use of the embryo-selection record. That gap cannot be closed simply by stretching Article 22 beyond its text; it requires an express legislative or sector-specific mechanism.

India's Assisted Reproductive Technology (Regulation) Act, 2021 (Act No. 42 of 2021) regulates a third object: the institutional provision of ART services. Its stated architecture centres on ART clinics and banks, national and state supervisory bodies, a National Registry, registration, inspection, record-keeping, consent, handling of gametes and embryos, and specified prohibited practices (Government of India, 2021). Sections 9 to 11 establish and define the National Registry, while sections 15 to 20 govern registration, renewal, suspension, cancellation, appeal, and inspection. Sections 21 to 24 impose duties concerning treatment, written consent, records, and the handling of human gametes and embryos. These provisions substantially control who may provide ART services and how clinics must document and conduct them.

The statute nevertheless does not expressly identify algorithmic clinical decision-support, software validation, model updates, training-data quality, subgroup performance, or explainability as regulatory subjects. Existing powers could reach an algorithm indirectly if its use caused a clinic to breach duties of consent, safe practice, record-keeping, confidentiality, non-discrimination, or embryo handling. Inspection and registration conditions might also provide a basis for future subordinate rules or standards. But indirect enforceability is not equivalent to a dedicated software-assurance regime. The Act does not specify when an embryo-ranking

model must be independently validated, what performance evidence a clinic must obtain before deployment, how changes to the model should be reviewed, or whether clinics must report adverse algorithmic incidents. Accordingly, the strongest claim is not that a clinic could deploy any unvalidated score without legal consequence. General duties and professional standards may still be engaged, particularly where reliance on the tool compromises consent or safe treatment. The narrower structural point is that compliance is assessed primarily through the clinic's registration, personnel, procedures, records, and observance of statutory duties, rather than through an ex ante regulatory dossier for the algorithm itself. Unless another medical-device or professional-regulation route applies, the ART Act supplies no express trigger for independent review merely because software materially influences which embryo is transferred.

The resulting architecture is therefore best described as partially overlapping but functionally incomplete. The AI Act can regulate a qualifying high-risk system as a product; the GDPR can regulate personal-data processing and, in narrower circumstances, solely automated decisions; and India's ART Act can regulate the clinic that uses the tool. The same deployment may fall within all three, yet each regime asks a different compliance question. None directly supplies a unified answer to whether the model's target outcome is permissible, whether its comparative ranking is sufficiently validated for the patient population, what explanation must precede embryo transfer, how model drift should affect continued use, or who may assert interests connected to the resulting child.

This mismatch also creates enforcement discontinuities. A notified body may assess conformity without supervising the clinic's day-to-day reliance on the score; a data-protection authority may examine lawful processing without determining clinical validity; and an ART inspector may review consent forms and records without access to the developer's training data, technical file, or post-market performance evidence. Responsibility can consequently be dispersed among the developer, device manufacturer, clinic, clinician, cloud provider, and data controller, with no single actor obliged to demonstrate the integrity of the complete decision pathway from dataset to embryo transfer.

Algorithmic embryo selection thus falls not into a wholly unregulated void, but into the seams between product safety, data protection, and professional or institutional oversight. That distinction matters for reform. The objective should not be to duplicate existing controls, but to connect them through reproductive-specific requirements: a clear regulatory trigger based on material influence over embryo choice; independent clinical and subgroup validation; versioned audit logs and post-market drift monitoring; patient-facing explanation; meaningful human review; and rules governing the long-term stewardship of reproductive data. Section 6 develops that connective layer.

6. TOWARD A RIGHTS-BASED ADAPTIVE GOVERNANCE FRAMEWORK

Closing the structural gap requires a connective regulatory layer directed at both the algorithm and the embryo-selection decision it materially shapes, rather than another instrument confined to the device, the data-processing operation, or the clinic. The proposed framework should apply whenever software has a material influence on embryo ranking or treatment choice, irrespective of how the developer markets the product or whether it independently satisfies a medical-device classification test. Its obligations should follow the system across development, procurement, deployment, model updating, and post-market monitoring, and should be allocated among developers, clinics, and clinicians according to their control over each stage. Five mutually reinforcing elements respond to the specific deficiencies identified above. They draw on, but deliberately extend beyond, the AI Act's requirements for data governance, transparency, human oversight, accuracy, logging, and post-market monitoring, and the World Health Organization's principles of autonomy, explainability, accountability, inclusiveness, and equity (European Parliament and Council of the European Union, 2024; World Health Organization, 2021).

1. Reproductive-specific transparency and explainability. Providers should submit to regulators a standardised technical dossier identifying the model's intended clinical purpose, target outcome, development and validation datasets, subgroup composition, exclusion criteria, missingness, external-validation design, performance metrics, uncertainty, known limitations, model version, and rules governing updates. This proposal builds on Articles 10 and 13 of the AI Act, which require data governance and sufficiently transparent information for deployers of high-risk systems, including information about accuracy, limitations, relevant datasets, group-specific performance where appropriate, and capabilities for explaining output (European Parliament and Council of the European Union, 2024). Embryo selection, however, also requires a patient-facing layer. Before transfer, the clinic should explain in plain language that AI was used, what outcome the score predicts, whether the model was validated in comparable patients and settings, how much weight the clinician gave the score, and what reasonable alternatives exist.

Transparency should therefore be audience-specific rather than equated with indiscriminate disclosure of source code. Regulators need enough information to reproduce and audit performance; clinics need instructions enabling safe interpretation; patients need information material to consent; and independent researchers may require controlled access to de-identified evidence. Research on explainability in AI-assisted embryo selection likewise indicates that explanations must be designed around the distinct needs of clinicians and patients (Urcelay et al., 2023). These duties would close the disclosure gap identified in Section 4.1 even

where software falls outside the MDR-derived high-risk route.

2. Meaningful human oversight and retained clinical authority. The embryologist or responsible clinician should retain final authority over embryo and treatment decisions, with the capacity to disregard, override, or suspend the model's output. Article 14 of the AI Act offers a useful baseline: oversight must enable a natural person to understand the system's capabilities and limitations, monitor for anomalies, recognise automation bias, interpret outputs, and decide not to use or to reverse them (European Parliament and Council of the European Union, 2024). In reproductive care, those safeguards should be translated into professional duties rather than satisfied by placing a nominal reviewer at the end of an automated workflow.

Meaningful oversight requires adequate training, sufficient time to review competing embryos, access to morphology and clinical context independent of the AI score, and freedom from contractual or throughput pressures that make adherence to the ranking the default. Clinics should document the score shown, the human assessment, the final decision, and any override with a brief clinical reason. Recurrent divergence between human judgment and model output should trigger review rather than be treated as individual error. These measures answer the accountability erosion described in Section 4.4 by making the AI score one evidential input and preserving an identifiable professional decision-maker.

3. Mandatory stratified performance and bias auditing. Regulatory clearance and continued clinical use should depend on disclosed premarket and post-market audits rather than a general assurance of fairness. Audits should report discrimination, calibration, sensitivity, specificity, false-positive and false-negative rates, ranking stability, and clinically relevant outcomes across prespecified demographic and clinical subgroups. They should also identify subgroup sample sizes and uncertainty intervals, because an apparent absence of disparity in a small sample is not evidence of equivalent performance. The AI Act's data-governance and accuracy requirements provide a general foundation, while the NIST AI Risk Management Framework identifies validity, reliability, accountability, transparency, explainability, privacy, and the management of harmful bias as components of trustworthy AI (National Institute of Standards and Technology, 2023).

The audit should be external and multicentre where feasible, cover clinics and equipment not represented in development, and continue after deployment to detect drift caused by changing patient populations, laboratory protocols, incubators, imaging systems, or software updates. Material underperformance should lead to a corrective-action plan, restricted indications, heightened human review, suspension, or withdrawal. A public summary should be intelligible to clinics and patients, while regulators receive the complete results. This converts the empirical concern identified in Section 4.3 into a measurable and enforceable condition of use.

4. Sui generis protection for reproductive and embryonic data. Data-protection law should recognise a protected category encompassing embryo images, morphokinetic recordings, gamete and donor information, genetic or genomic results, inferred viability traits, model-generated scores, selection rationales, and linked treatment outcomes. The protection should be justified not by assigning legal personhood to an embryo, but by the relational and longitudinal character of the data: the same record may concern the treated patient, donors, genetic relatives, and, after birth, the resulting child. Existing instruments such as the GDPR and Convention 108+ protect identified or identifiable individuals and provide a foundation of purpose limitation, security, transparency, and accountability, but they do not establish a distinct intergenerational claim to the provenance and later use of an embryo-selection record (European Parliament and Council of the European Union, 2016; Council of Europe, 2018).

Legislation should therefore specify strict purpose limitation, enhanced security, controlled reuse for model development, retention periods calibrated to clinical follow-up and the later interests of a resulting child, and restrictions on commercial transfer or trait prediction unrelated to treatment. Once the child is born and capable of exercising rights, the framework should provide age-appropriate access to information about the existence, provenance, and authorised uses of records relating to the embryo-selection process, subject to safeguards for the privacy of patients, donors, and relatives. An independent trustee or designated authority could protect these future interests before the child can act. This approach supplies the legislative bridge identified in Sections 4.2 and 5 without assuming that the GDPR's existing definition of a data subject can simply be extended retrospectively.

5. An enforceable right to explanation and contestation. Patients should have a sector-specific right to receive an explanation whenever AI materially influences embryo ranking, whether or not the decision is "solely automated" within Article 22 GDPR. The right should arise before transfer, when explanation can still support choice, and again after the decision if a patient seeks review. It should cover the model's intended outcome; the principal categories of data used; relevant validation limits and uncertainty; the reasons the recommended embryo ranked above clinically comparable alternatives, to the extent these can be meaningfully stated; the role of human judgment; and the available alternatives. It should also include access to a qualified human reviewer and a practical route to challenge the use or interpretation of the score. Existing GDPR guidance on automated decision-making remains an important baseline, but its trigger and safeguards do not create a general explanation right for every AI-assisted clinical decision (Article 29 Data Protection Working Party, 2018).

The right should be enforceable against the clinic as the patient-facing decision-maker, while developers remain obliged to provide clinics and regulators with the information necessary to discharge it. Failure to

disclose material limitations should support regulatory sanctions and ordinary remedies where causation and loss are established. Providers should also retain versioned logs showing the model used, inputs considered, output generated, human review performed, and final embryo selected. This remedy is not a demand for a mathematically complete account of a neural network; it is a right to a clinically meaningful, timely, and contestable explanation calibrated to the reproductive and intergenerational stakes of the decision.

Taken together, the five elements create an adaptive governance cycle rather than a one-time approval event: transparency makes the system inspectable; human oversight preserves professional responsibility; audits test performance and equity; data rules protect interests across time; and explanation enables patient agency and legal contestation. Their effectiveness depends on coordination. A reproductive-technology regulator or designated lead authority should maintain a registry of approved systems and versions, receive serious-incident and drift reports, coordinate with medical-device and data-protection authorities, and require renewed review when the model's intended purpose, inputs, target outcome, or performance materially changes.

7. Discussion And Implications

For AI/ML practitioners developing embryo-assessment systems, the preceding analysis translates into concrete design and validation requirements rather than abstract ethical aspirations. Claims intended to support regulatory approval or clinical adoption should, wherever feasible, be grounded in prospective, multicentre, and randomised evaluation, as in the study design reported for Embryo Predict, rather than in retrospective accuracy alone. Model development should incorporate prespecified subgroup analyses, external validation, calibration assessment, uncertainty reporting, and post-deployment drift monitoring from the outset. These measures allow developers to identify where performance deteriorates across patient populations, clinics, imaging systems, or laboratory protocols before those differences affect treatment decisions. For complex multimodal architectures such as IVFormer, explainability must likewise be treated as a design constraint: the system should retain traceable inputs, versioned outputs, clinically intelligible reasons for rankings, and mechanisms for human review. Retrofitting explanations onto a deployed black-box model is both technically weaker and less useful for consent, audit, and accountability than building those capabilities into the development lifecycle.

For policymakers, the comparison in Section 5 shows why isolated amendments are unlikely to be sufficient. Extending Annex III of the AI Act, adding embryo-related examples to GDPR guidance, or inserting the word "software" into India's ART Act would improve coverage at particular points but preserve the same fragmentation, because each regime continues to organise responsibility around a different regulatory object. A more durable response would recognise AI-mediated embryo selection as a distinct regulatory subject and create a coordinating layer across product safety, data protection, professional responsibility, and clinic licensing. That layer should allocate duties across developers, manufacturers, clinics, and clinicians; establish a trigger based on material influence over embryo choice; and connect premarket validation with patient-facing explanation, audit records, adverse-event reporting, and post-market monitoring. The practical implication is not necessarily a wholly separate regulator, but a clear lead authority and interoperable obligations that prevent responsibility from disappearing at institutional boundaries.

The clinical promise of these systems is real and should not be understated. Earlier, non-invasive identification of embryos with a lower probability of implantation may reduce avoidable transfer cycles and the associated physical, emotional, temporal, and financial burdens of IVF. Better decision support may also improve consistency between embryologists and help clinics use limited time and laboratory resources more effectively. Those benefits, however, depend on evidence that the model performs reliably in the population and setting in which it is used, that clinicians retain meaningful authority, and that patients understand the basis and limits of the recommendation. The argument of this review is therefore not that AI has no legitimate place in embryo selection, but that its place must be defined with the legal and ethical specificity warranted by a decision that is simultaneously clinical, reproductive, data-intensive, and potentially intergenerational. Governance should enable responsible innovation by making trustworthy systems easier to validate and adopt while preventing opaque, weakly evidenced, or repurposed systems from becoming embedded in care. The central policy task is to ensure that technological capability does not mature faster than the safeguards required to make its use clinically defensible and publicly legitimate.

8. Conclusion

This review's central finding is that the regulation of AI-assisted reproductive technologies is not absent; rather, it is dispersed across various legal frameworks that do not explicitly address the pivotal act: employing an algorithm to prioritise one embryo over another. The EU AI Act and medical-device law categorise qualifying systems as products; the GDPR governs personal data processing and, under specific circumstances, automated decision-making; and India's ART Act oversees clinics and reproductive services. Each framework provides essential safeguards, yet none singularly focuses on the embryo-selection process, the core outcome, its

evidentiary foundation, implications for patient autonomy, and potential consequences for the resulting child. Consequently, the primary issue lies in structural fragmentation rather than legislative neglect.

The appropriate response is not to displace these regimes or duplicate their controls, but to connect them through a reproductive-specific governance layer triggered whenever AI materially influences embryo ranking or treatment choice. That layer should require prospective and external validation, disclosed subgroup performance and continuing drift monitoring, meaningful human authority, patient-facing explanation and contestation, versioned audit records, and enhanced stewardship of reproductive data over time. Duties should be allocated according to control: developers must substantiate performance and disclose limitations; clinics must procure, monitor, and document systems responsibly; clinicians must retain and exercise independent judgment; and a designated lead authority must coordinate medical-device, data-protection, and ART oversight. This framework would convert broad commitments to transparency, fairness, and human oversight into measurable conditions of lawful clinical use.

Such regulation need not hinder innovation. Conversely, well-defined, interoperable standards will give responsible developers a credible pathway to clinical adoption. These standards will help clinics distinguish validated systems from promotional claims and enable patients to benefit from decision support without compromising informed consent. Although the technology can reduce unnecessary transfer cycles and ease IVF-related burdens, these benefits are justifiable only when the score is reliable for the patient before clinical use, its limitations are clearly understood, and accountability remains traceable from dataset to transfer decision. The fundamental principle should therefore be straightforward: the greater the algorithm's influence on reproductive choice, the stronger the evidence, explanation, oversight, and long-term accountability must be. Establishing such a framework during the ongoing deployment phase is both more feasible and more legitimate than attempting to reconstruct accountability after opaque systems have become entrenched in routine care.

Conflicts of Interest

The author(s) declare no conflict of interest.

Acknowledgments

Not applicable.

Declarations

Use of Generative AI

During manuscript preparation, the author(s) used Claude (Anthropic) for literature review and editing.

References

1. AIVF. (2026). AIVF: Revolutionizing IVF with AI technology. Retrieved from <https://aivf.co/>
2. European Parliament and Council of the European Union. (2016). Regulation (EU) 2016/679 (General Data Protection Regulation). Official Journal of the European Union, L119.
3. Council of Europe. (1997). Convention for the Protection of Human Rights and Dignity of the Human Being with regard to the Application of Biology and Medicine: Convention on Human Rights and Biomedicine (Oviedo Convention), ETS No. 164.
4. Council of Europe. (2018). Protocol amending the Convention for the Protection of Individuals with regard to Automatic Processing of Personal Data (Convention 108+), CETS No. 223, and Explanatory Report.
5. Article 29 Data Protection Working Party. (2018). Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679 (WP251rev.01), endorsed by the European Data Protection Board.
6. European Parliament and Council of the European Union. (2017). Regulation (EU) 2017/745 on medical devices. Official Journal of the European Union, L117.
7. European Parliament and Council of the European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). OJ L, 2024/1689.
8. Fertility Bridge. (2026). FDA clears first AI-powered embryo selection tool in IVF labs. Retrieved from <https://www.fertilitybridge.com/news-articles/fda-clears-ai-embryo-selection-alife>
9. Government of India. (2021). The Assisted Reproductive Technology (Regulation) Act, 2021 (Act No. 42 of 2021).
10. General Medical Council. (2020). Decision making and consent. London: General Medical Council.
11. Montgomery v Lanarkshire Health Board [2015] UKSC 11.
12. Medical Device Coordination Group. (2025). MDCG 2019-11 Rev. 1: Guidance on Qualification and Classification of Software in Regulation (EU) 2017/745—MDR and Regulation (EU) 2017/746—IVDR. Brussels: European Commission.
13. National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1). Gaithersburg, MD: U.S. Department of Commerce.

- <https://doi.org/10.6028/NIST.AI.100-1>
14. Nicolielo, M., Jacobs, C. K., Lourenço, B., Maffei, M. C., Chéles, D. S., Duarte, M. B., Mendes, B. A., Moraes, V. C., Chehin, M. B., Alegretti, J. R., da Motta, E. L. A., Lorenzon, A. R., Nogueira, M. F. G., and Rocha, J. C. (2025). MAIA platform for routine clinical testing: An artificial intelligence embryo selection tool developed to assist embryologists. *Scientific Reports*, 15, Article 32273. <https://doi.org/10.1038/s41598-025-17755-y>
 15. PR Newswire / Alife Health. (2025). Alife Health receives CE Mark under European Medical Device Regulation for Embryo Predict™. Retrieved from <https://www.prnewswire.com/news-releases/alife-health-receives-ce-mark-under-european-medical-device-regulation-for-embryo-predict-its-ai-powered-embryo-selection-tool-302592126.html>
 16. UNESCO. (1997). Universal Declaration on the Human Genome and Human Rights. Paris: United Nations Educational, Scientific and Cultural Organization.
 17. Urcelay, L., Hinos, D., Martin-Torres, P., González, M., Méndez, M., Cívico, S., Álvarez-Napagao, S., and Garcia-Gasulla, D. (2023). Exploring the role of explainability in AI-assisted embryo selection. *Frontiers in Artificial Intelligence and Applications*, 375, 153–162. <https://doi.org/10.3233/FAIA230678>
 18. Wang, G., Wang, K., Gao, Y., Chen, L., Gao, T., Ma, Y., Jiang, Z., Yang, G., Feng, F., Zhang, S., Gu, Y., Liu, G., Chen, L., Ma, L. S., Sang, Y., Xu, Y., Lin, G., and Liu, X. (2024). A generalized AI system for human embryo selection covering the entire IVF cycle via multi-modal contrastive learning. *Patterns*, 5(7), 100985. <https://doi.org/10.1016/j.patter.2024.100985>
 19. WHO Expert Advisory Committee on Developing Global Standards for Governance and Oversight of Human Genome Editing. (2021). Human genome editing: Recommendations. Geneva: World Health Organization. ISBN 978-92-4-003038-1.
 20. World Health Organization. (2021). Ethics and governance of artificial intelligence for health: WHO guidance. Geneva: World Health Organization. ISBN 978-92-4-002920-0.