

Ethical Implementation of Artificial Intelligence to Select Embryos in *In Vitro Fertilization*

Michael Anis Mihdi Afnan

Department of Medicine
Imperial College London
London, UK
michaelafnan@icloud.com

Cynthia Rudin

Departments of Computer Science,
Electrical Engineering and
Statistical Science
Duke University
Durham, North Carolina, USA
cynthia@cs.duke.edu

Vincent Conitzer

Departments of Computer Science,
Economics and Philosophy & Institute
for Ethics in AI and Departments of Computer
Science and Philosophy
Duke University & Oxford University
Durham, North Carolina, USA
conitzer@cs.duke.edu

Julian Savulescu

Uehiro Centre for Practical Ethics & Wellcome
Centre for Ethics and Humanities & Murdoch Children's
Research Institute
Oxford University & Oxford University & Royal Children's
Hospital, Oxford, UK
julian.savulescu@philosophy.ox.ac.uk

Abhishek Mishra

Uehiro Centre for Practical Ethics
Oxford University
Oxford, UK
abhishek.mishra@philosophy.ox.ac.uk

Yanhe Liu

Monash IVF Group & School of Human Sciences
& School of Medical and Health Sciences
Monash IVF Group & University of Western Australia
& Edith Cowan University, Southport, Australia
gift0409@yahoo.com.au

Masoud Afnan

Department of Obstetrics and Gynaecology
Qingdao United Family Hospital
Qingdao, China
masoudafnan@me.com

ABSTRACT

AI has the potential to revolutionize many areas of healthcare. Radiology, dermatology, and ophthalmology are some of the areas most likely to be impacted in the near future, and they have received significant attention from the broader research community. But AI techniques are now also starting to be used in *in vitro fertilization (IVF)*, in particular for selecting which embryos to transfer to the woman. The contribution of AI to IVF is potentially significant, but must be done carefully and transparently, as the ethical issues are significant, in part because this field involves creating new people.

We first give a brief introduction to IVF and review the use of AI for embryo selection. We discuss concerns with the interpretation of the reported results from scientific and practical perspectives. We then consider the broader ethical issues involved. We discuss in detail the problems that result from the use of black-box methods in this context and advocate strongly for

the use of *interpretable* models. Importantly, there have been no published trials of clinical effectiveness, a problem in both the AI and IVF communities, and we therefore argue that clinical implementation at this point would be premature. Finally, we discuss ways for the broader AI community to become involved to ensure scientifically sound and ethically responsible development of AI in IVF.

CCS CONCEPTS

• Computing methodologies-Machine learning algorithms; Computer vision; Artificial intelligence; Machine learning; Machine learning approaches • Applied computing-Life and medical sciences

KEYWORDS

IVF, *In Vitro Fertilization*, Embryo Selection, Artificial Intelligence, AI, Machine Learning, Interpretable, Black-Box, Randomised Controlled Trials, RCT, Ethics

ACM Reference format:

Michael Anis Mihdi Afnan, Cynthia Rudin, Vincent Conitzer, Julian Savulescu, Abhishek Mishra, Yanhe Liu and Masoud Afnan. 2021. Ethical Implementation of Artificial Intelligence to Select Embryos in *In Vitro Fertilization*. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AIES '21)*, May 19-21, 2021, Virtual Event, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3461702.3462589>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org. AIES '21, May 19-21, 2021, Virtual Event, USA.

© 2021 Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-1-4503-8473-5/21/05...\$15.00.

DOI: <https://doi.org/10.1145/3461702.3462589>

1 Introduction

In vitro fertilization (IVF) is a clinical technique which has revolutionized the treatment of infertility. The process involves fertilizing the egg in a laboratory and replacing the resultant embryo into the uterus. Natural fertilization and conception is an inefficient process, with low chances of a live birth for any particular embryo. The solution both in nature and with medical treatment is to create multiple embryos, so that ultimately one will probably implant. In nature, the cost is time to pregnancy or, in the event of no embryos implanting, the pain of childlessness. In clinical practice, the cost is additionally measured in dollars. To increase the efficiency of clinical practice, much attention has been given to selecting the embryo that is most likely to implant. A recent innovation in the laboratory is time-lapse imaging of the embryo in culture over a number of days. This gives rise to thousands of visual data points, and with it the promise of augmenting the embryo selection process with artificial intelligence (AI)-based models. In this paper, we provide an overview of the IVF process, review current approaches to using AI in embryo selection, discuss ethical issues of using AI in this specific field, and make proposals for the ethical implementation of this new technology. We finish with encouragement for AI researchers to collaborate with fertility clinicians to take this research forward in a meaningful and ethical way.

2 The In Vitro Fertilization (IVF) Process

Each year, millions of couples who suffer from infertility pin their hopes of starting or growing their family on IVF [1]. Heavily criticized by many at first as an unethical human experiment [2], the technique has become one of the most successful therapeutic innovations of the past half-century, leading to over 9 million babies born since the first IVF birth in 1978 [1]. The limit of IVF's success, however, is reflected in the millions more whose hopes have not been fulfilled. Particularly for those with advancing age and comorbidities, but also for every couple who tries, success is not guaranteed. On average, across all age groups, the live birth rate per treatment cycle is 26.1% in the UK [3].

To maximize the chance of retrieving a good quality egg and subsequent embryo, women are given hormone treatment to stimulate development of multiple eggs, which are then harvested, inseminated and the resultant embryos cultured in the laboratory for 2-6 days. Typically 2-4 would reach the blastocyst stage around day 5 or 6. The embryologist would then select 1 blastocyst for transfer to the uterus. Any unused embryos thought to be viable are then frozen for use later in case the initial transfer following egg collection is unsuccessful, or if successful, for a future sibling [4]. It is important to note that competitive embryo selection is unique to IVF, and does not occur in nature. IVF clinics are keen to improve on current embryo selection strategies in order to maximise the chance of pregnancy at an early stage of the couple's treatment.

Early embryo development at the preimplantation stage is a very dynamic process. Hours after fertilisation, two pronuclei are formed carrying DNA material contributed by the sperm and the egg. The pronuclear membrane breaks down shortly before the

first cell division, leading to a 2-cell embryo. As cells continue to divide, they become more compact with increased cell-to-cell interaction from 3 days post-fertilization. On day 4, the embryo reaches the "morula" stage, where borders between cells become invisible. During the next 1 to 2 days, cells separate into 2 layers, with a growing cavity formed between them; at this point, the embryo is called a "blastocyst". The outer layer of cells, also known as the trophectoderm, will become part of the placenta following implantation, while the inner layer (inner cell mass) will become the fetus. Both layers hatch out of the shell around the embryo (zona pellucida) before implantation into the uterus.

Traditional embryo selection is based on several snapshot observations of an embryo under a microscope, at specific time points during culture (Figure 1). Considering the dynamic nature of embryo development, the static nature of the information collected in this method limits the accuracy of embryo selection [5].

Examples of other developments to select embryos more likely to implant include (1) allowing embryos to self-deselect via extended culture in the laboratory [6], (2) metabolomic profiling of spent culture media [7] or (3) adding extra genetic testing such as pre-implantation genetic testing for aneuploidy (PGT-A) which is controversial because of its invasive nature and diagnostic imperfection [8].




Figure 1: Embryo development from a few hours after fertilisation (top-left) with 2 pronuclei to the "blastocyst" with trophectoderm and inner cell mass (bottom-right).

Evaluation of embryo quality by the embryologist is limited by considerable inter-operator variability, due to the current mix of objective and subjective measures in assessment, and human factors, such as being influenced by confounders [9]. However, the availability of numerous data points made available by recent advances, such as with time-lapse videography [10], has enabled novel parameters to be measured for embryo selection [11, 12, 13]. Debate is still ongoing regarding the best approach of using such time-lapse images for embryo selection [14].

3 AI as an Embryo Selection Tool to Improve the Success Rate Per Transfer

The application of AI in IVF has the potential to provide more objective, more rapid, and potentially more accurate evaluation of key steps in the IVF process, to make it more reproducible and repeatable when compared with a purely human approach [15]. In particular AI for embryo selection has attracted much interest, and potentially holds much promise [16].

The type of AI that can help embryo selection is Machine Learning (ML) – models that can automatically learn and adapt as they are exposed to more data (whether images or other data). This is particularly useful when there is access to lots of data, but we do not immediately know how to leverage it to make better predictions, or when we cannot manually process it all to generate meaningful knowledge. Potential variables include morphological features, such as cleavage of the embryo cells (blastomeres), fragmentation, morphokinetic features (including time intervals between certain features), and clinical factors such as age of the woman or cause of infertility [17]. Computer Vision (CV) allows large amounts of image data to be automatically analyzed by algorithms, and rapid recent advances in this field offer great promise to improve embryo selection.

4 Current State of Research on the Use of AI to Select Embryos

We searched MEDLINE, Embase and Google Scholar for full-text studies evaluating AI to select embryos using the strategy included in the appendix. We checked the citations of papers we identified in the search for any publications we might have missed.

Studies evaluating AI for embryo selection make impressive accuracy claims for their ML models [18, 19]. One commonly reported performance measure is the receiver operating characteristic (ROC) curve which shows how a test's sensitivity and specificity correlate at different thresholds. The area under this curve (AUC) indicates the test's performance. An AUC >0.9 usually indicates outstanding performance, and the ML models from the studies cited above surpass this benchmark.

Studies that evaluate the efficacy of AI models for embryo selection do so for 2 types of outcomes; a) outcomes meaningful to the patient, such as a live birth or a fetal heartbeat positive pregnancy, or b) agreement with the existing standard, which in this case would be assessment by embryologists. One of the challenges of using live birth as the meaningful outcome (ground truth) is that a potentially viable embryo can result in either a live birth, or no live birth, depending on other, non-embryo factors, such as the health of the mother.

Tran, Cooke and Illingworth's [18] study belongs to the former category. They evaluated a model called IVY which rates how likely an embryo is to lead to a fetal heartbeat (FH) pregnancy on a confidence scale of 0 (definitely will not implant) to 1 (definitely will implant). Their ROC curve's AUC was 0.93. However, as Kan-Tor, Ben-Meir and Buxboim [20] point out, the majority of the embryos on which the algorithm had been trained and tested were of such poor quality that they would have been discarded in any event, thereby artificially inflating the AUC. As Kan-Tor, Ben-Meir

and Buxboim explain, the clinical need is to identify the embryo with the highest chance of success among a set of embryos that appear to be potentially viable, and not from embryos which embryologists readily discard.

Khosravi et al.'s study [19], on the other hand, belongs to the latter category. They categorized embryos into 3 groups – good-, fair-, and poor-quality embryos according to a consensus of multiple embryologists. They then evaluated their AI algorithm's ability to identify the good- and the poor-quality embryos (but not the fair-quality embryos); for this task the algorithm achieved 96.94% accuracy. This was better than the performance of individual embryologists. However, broad categorizations into "good" or "poor" quality are of limited benefit when trying to find the best embryo in a group of similar-quality embryos.

The above analyses of Khosravi et al. [19] and Tran, Cooke and Illingworth's [18] studies demonstrate the importance of understanding exactly how researchers test their algorithms before drawing conclusions from headline statistics. These studies are important steps to investigate efficacy (the ability to produce a specified outcome in experimental circumstances), to develop the tool, and establish proof of principle. However, they are only a prelude to testing in the clinic. When Curchoe et al. [21] reviewed how the results of AI studies in reproductive medicine relate to real-life clinical practice, they highlighted four pitfalls that are common throughout the literature: small sample sizes, imbalanced datasets, non-generalizable settings and limited performance metrics.

Furthermore, to date, no AI studies for embryo selection using a Randomized Controlled Trial (RCT) have been published, though 1 is registered [22]. The lack of RCTs appears to be typical of much of AI in medicine [23]. The problem of lack of evidence before implementation is compounded by the IVF industry which is notorious for aggressively marketing unproven clinical and laboratory "add-ons" [24, 25]. The problem is compounded because clinicians who do not have an adequate understanding of AI will find it difficult to critically navigate the literature which contains unfamiliar concepts and terminology.

Many AI studies for embryo selection use uninterpretable ("black-box") machine learning models. These models are either too complicated for any human to understand, or they are proprietary – in which case, comprehension of such a model is not possible for outsiders [26]. Specifically, many studies in this field use neural networks that are not interpretable, and not designed to be interpretable (e.g., Chen et al. [27]). Other approaches use interpretable features (whether they are labeled manually by doctors or labeled by neural networks whose output can be manually verified) but combine them in uninterpretable ways, such as using principal component analysis (PCA) pre-processing (which forces a dependence on all variables) followed by a machine learning method such as a neural network or random forest [28, 29]. The work of Leahy et al. [30] is interesting because the model is decomposable into separate neural network models that each extract different information (e.g., measurements of an embryo) that can be directly checked by an embryologist. Combining these separate models into an interpretable "combined" model to form the final prediction would be something we could

potentially recommend; if their final combined model was interpretable (such as a scoring system or sparse decision tree), then each piece of the whole system would be directly checkable by an embryologist for correctness, and thus interpretable. A third category of studies use fully interpretable features (e.g., measurements of the embryo taken by embryologists), but use older techniques that are not particularly accurate and do not explicitly optimize for interpretability (for instance, the models are not sparse). These works generally do not apply any computer vision techniques, relying instead on humans to estimate measurements from the embryo images. Examples include the works of Raef, Maleki and Ferdousi [31] and Morales et al. [32], who created interpretable hand-calculated features and applied a variety of classical machine learning algorithms to them.

The opaqueness or ‘black-box’ nature of AI models is problematic for two main types of reasons: ethical, and epistemic, which we will describe next.

5 Ethical Concerns with Black-Box AI Models

5.1 Failure to Perform Randomized Controlled Trials

The most important ethical issue facing the adoption of AI assisted IVF is the need for careful RCTs against best current approaches. Whilst 1 RCT has been registered [22], it is premature to implement a technology in the clinical setting before the trial results are made available. No matter how promising a new intervention appears to be, the gold standard for evaluation is the RCT. Failing to do such trials risks harming patients, as does failing to perform systematic reviews of existing evidence and failing to publish negative results [33].

Both black-box and interpretable models must be systematically studied using RCTs. Even highly promising interventions can do more harm than good [33].

5.2 Equipose

The ethical justification for RCTs is that *equipose* exists between the proposed new intervention and existing care. Equipose exists when the new intervention is not known or reasonably believed to be significantly better or worse than the existing intervention. For embryo selection, equipose would therefore exist if the use of AI for selection is neither known to be better or worse in achieving successful pregnancies than existing approaches. From the literature, it is very difficult to determine whether equipose exists for AI models, and black-box models in particular: one would need to determine exactly when AI would be used (e.g., for obvious cases? For non-obvious cases only?), how it would be used in the context of the clinic (e.g., does the clinician always follow the AI prediction?), and for what populations and what settings the reported AUCs would be representative. As we discuss below, black-box models have serious problems with robustness, and as discussed above, AUC values reported from published studies cannot necessarily be trusted. Also, since interpretable AI methods have not been heavily developed, we cannot assess the value of interpretability in the decision-making process. In other words, we

currently cannot assess whether a human/AI “centaur” could be better than either one alone. This information would be essential to the question of equipose and whether to conduct an RCT.

Validation studies would help to assess these questions; typically these are conducted on a new sample from a different clinic than the AI method was trained on. However, we are not aware of any such validation studies.

5.3 Compromised Shared Decision-Making

An important concern centers on the use of opaque AI models potentially compromising shared decision-making and patient-centered medicine more broadly. Over the past few decades, the accepted model of clinical practice has shifted from a paternalistic one, where the clinician’s opinion and recommendation are simply accepted by the patient, to one of shared decision-making where this power and informational asymmetry is reduced to the benefit of patient autonomy [34]. Clinical AI models that are opaque (so that medical explanations for a model’s recommendation are inaccessible) compromise this shared decision-making due to the inability of the clinician and the patient to understand the model’s decision [35]. While there have been some counterarguments raised as to whether shared decision-making is truly compromised by opaque AI models [36], application of AI in embryo selection should be guided by an awareness of such potential dangers. It will be important to fully explain what is known about how the AI model comes to a “decision” (nature and size of dataset, reasons for confidence in prediction, possible alternative lines of justification, etc.), and further examine how interactions between clinicians and patients may change, both at the point of embryo selection as well as at the point of implantation failure. Clinicians should explain the basis of how embryos are selected for transfer, whether it is clinical or AI-assisted. If information that is traditionally conveyed to the patient as to why a particular embryo is selected – for example the number and symmetry of the cells, or if the cells are fragmented, and therefore what the chances of implantation are, and why implantation may fail – are no longer accessible, shared decision-making might indeed be compromised. Existing measures of shared decision-making and decision quality, such as the Decision Conflict Scale [37], the OPTION Scale [38], and the SURE Test [39] (among other patient-reported measures) can be used to guide such an evaluation.

It is important, however, not to overstate this concern. Firstly, AI-assisted decision-making should be compared to the status quo. Current expert judgment is based on biologically meaningful measures, which, although more broadly communicable than decisions of opaque models, are not very accurate for predicting a live birth. AI-assisted decision-making may not be worse (but may also not be better). More importantly, autonomy requires understanding information relevant and meaningful to one’s values. Knowing the basis of a prediction (cleavage rate, symmetry, etc) is not relevant: what is relevant are the risks, side-effects and benefits, and the confidence attached to these assessments.

However, black-box AI has the potential to significantly undermine shared decision-making in a way that interpretable AI does not. Although they are marketed and approved as decision

aids to the clinician, where the decision finally rests with the clinicians, black-box AI will in practice have the potential to introduce a new form of paternalism: “machine paternalism.” It is known that people tend to be complacent about use of automation [40] and tend to be accepting of AI in a role when they are familiar with AI in that role or when they believe it performs well in that role [41]. In practice it is hard to see how clinicians will challenge the deliverances of black-box AI. Indeed, doing so without good reason might open them up to legal liability. So in practice, black-box AI risks instrumentalizing the clinicians in a way that interpretable AI does not. While interpretable AI is an enhancement of human decision-making, black-box AI is a replacement for it.

5.4 Misrepresentation of Patient Values

Another ethical issue concerns potential harms from a misrepresentation of patient-values in the decision process. For example, there are reported differences between early morphokinetic profiles between male and female embryos [42, 43, 44, 45] (and other traits might be similarly differentially represented at this early stage). Models for embryo selection run the risk of systematically selecting for these traits if they are perceived by the model to be correlated with implantation success. For example, if a patient prefers that sex be randomly selected, this model may run counter to those values. If such models are opaque, this systematic favoring of particular traits may not be detected at the time of decision-making and so cannot be corrected for in a way that it could be with interpretable models. If we know exactly which (interpretable) characteristics of the embryo are correlated with various traits, it would be much easier to detect them at the time of decision-making. Otherwise, it may take considerable time for such a bias to be detected as a result of the use of AI because it will require large scale, statistically significant changes to manifest in evaluation, and a causal analysis that determines that the cause of bias was indeed the use of AI, and not another source.

If some of these traits are ethically salient ones for the patient, then this creates a scenario where the patient’s values may not be sufficiently represented to guide the decision-making process for embryo selection. Such concerns have also been raised for other clinical models [46], calling for the design of such systems to be ‘value-flexible’ so that in clinical settings, both clinicians and patients are (1) aware of what metrics are driving a model’s recommendations (either directly or as a proxy for some other medical fact/trait), and (2) able to appropriately reflect the patient’s values in the decision process either directly through the model, or in subsequently adjusting the recommendation.

Again, it is important not to overstate this concern. The patient’s own values could be inserted into AI algorithms (e.g., preference for sex and other non-disease characteristics) and AI might bring to the surface the importance of these values in decision-making. Of course, valuing and selecting non-disease traits (such as sex or intelligence) raises the debate around designer babies, but some have argued that such selection is permissible [47] or even a moral obligation when it relates to the well-being of a future child [48, 49, 50]

5.5 Health and Well-Being of Future Children

Such potential biasing of AI-selection might also have impacts on the health or well-being of future children. For example, it is possible that some disadvantageous trait (such as increased risk of cancer or mental disorder) correlates with a higher chance of implantation. However, this risk might be present unknowingly in ordinary clinical decision-making. This also underscores the importance of clinical trials not merely measuring implantation or even healthy live birth but long-term well-being of the child created by IVF through long-term (decades) follow up.

Reproduction is also unique because selection determines who will come into existence. This creates the so-called “non-identity problem” which has spawned decades of unresolved philosophical debate, sparked by [51]. Imagine Embryo A has a higher chance of implantation but unknowingly a higher chance of cancer later in life than embryo B. AI selects A. A is born but gets cancer at the age of 30. Was A harmed by the decision to select A rather than B? No, a different person (B) would have been otherwise selected. Provided that the disadvantageous trait or genes do not make A’s life so bad as to have been not worth living, then A cannot be harmed by selection. On this ground, greater risks can be taken in embryo selection than with interventions on a specific embryo (such as A) which do risk harm to a specific individual [52]. Nonetheless, some have argued that parents (and clinicians) still have a moral obligation to select the embryo with the best chance of the best life [48, 49, 50].

5.6 Impacts of Devaluing Disability

There is a general problem with embryo selection raised by disability activists: any kind of selection based on predicted health or well-being discriminates against the disabled and expresses a negative message about the value of their lives - the expressivist objection [53]. For example, screening for Down Syndrome has been said to express a negative view about the value of people with Down Syndrome [54]. This applies not only to AI selection but to clinical selection more generally, and there are numerous responses [53]. However, AI might considerably expand the scope of this objection: any trait that lowers chance of implantation might result in selection against that group, e.g., sex as previously discussed.

The best response to these concerns would be to monitor such effects and ensure social responses that reinforce the equality of all people, including people with disabilities. Thus, rather than forgoing selection, it is better to ensure there are sufficient social resources so that all existing people have a reasonable chance of a good life [50].

Interpretable AI may allow issues of bias to be identified earlier and in a more actionable way. Specifically, if we discover interpretable features that can be linked to various outcomes, it is easier to monitor for possible bias and avoid it.

5.7 Societal Impacts of AI for Embryo Selection

Successful AI models might be deployed at scale, and if such models systematically favor certain traits represented in early morphokinetic profiles, this might impact society. Even if would-

be parents might not care about the sex of their future child and might be willing to accept a higher likelihood of one sex for a higher likelihood of implantation success, this will still have societal ramifications through a skewed population ratio. The scale of these ramifications will correlate with rates of IVF use in the future; the more individuals opt for IVF, the greater the impact. While such possibilities are at this stage mostly speculative, they represent a scale of impact that is significant and should therefore be considered. Since black-box models do not aim to identify specific aspects of an embryo with specific traits, it might make these issues more difficult to detect until it is too late and there are societal-level impacts. In this regard, interpretable AI may allow earlier detection of systematic favouring of certain traits.

Further, if AI-assisted IVF works better for some races than others, this could have serious societal implications. While such differences should be monitored whether a black-box or interpretable AI method is used, interpretable AI may help to detect such issues more easily, for instance, if the AI model is known to leverage factors that differ among ethnic groups (e.g., the relationship of age to fertility).

5.8 Black-Box Models Pose a Responsibility Gap

The final ethical issue concerns a potential erosion of ethical and legal accountability through the use of opaque AI models. If it is determined that clinicians cannot be held responsible for injuries sustained by the patient due to a reliance on opaque AI models, the responsibility for this class of errors would need to be borne by another agent. In the absence of institutionalized accountability mechanisms that hold other agents, like model developers, responsible, this creates a ‘responsibility gap’ when it comes to the use of AI models.

The most straightforward case in which accountability is required would be repeated implantation failure or low success rates due to suboptimal embryo selection processes, and/or injury being sustained by the patient as a result of implementation of a model recommendation (either to the mother through surgical complications or the child when he/she is born - wrongful life or birth). Under such circumstances, if the patient seeks an account of what happened or advances a charge of negligence against the clinician, the decision-making process needs to be explicable. Traditionally, if a charge of negligence is advanced, experts assess the clinician’s decision-making process, and depending on whether they deem this to be medically reasonable, the clinician is either acquitted or held culpable. If AI models used for embryo selection reason in uninterpretable ways, it is unclear how a court might evaluate the doctor’s decision-making, and subsequently, it would be unclear how responsibility for injury may be adjudicated [55, 56].

6 Epistemic Concerns with Black-Box AI Models

There are technical challenges posed by black-box or opaque systems; it is unclear how we might assess the reliability of the model’s predictions, eliminate potentially confounding factors at the decision-making point, and assess to what extent the model’s

accuracy is representative in a given use-case. In a field where ‘add-on’ clinical offerings are already widespread despite inadequate evidence of effectiveness, this epistemic problem is especially troubling [25].

6.1 Black-Box Models Create Information Asymmetries

The use of black-box models creates an information asymmetry between the company selling the tool and the clinicians having to make daily decisions as to which embryo to transfer. Using such models would force the embryologist to abrogate decision-making to programs they themselves do not understand. It is not possible to fully evaluate whether to trust these complex models without an understanding of their reasoning processes.

6.2 Confounders are Rampant

If we do not understand what a black-box model is doing, it is entirely possible that its predictions are based on confounders that should not be used as predictors. Confounders are often difficult to detect and cause models not to generalize. When coupled with a poor choice for evaluation metric, the confounding might not be noticed [57, 58].

Let us construct a simple example where an obvious confounder and a standard (but ill-chosen) evaluation metric provide a situation where a useless model would appear to be excellent. In this example, the confounder is the mother’s age, and the metric is overall AUC (not the AUC for an individual couple). It is largely possible that the mother’s age is a major factor in predictions; what if it were the sole factor, so that a model based on an image of the embryo is predictive of mother’s age only? If the model were a mere proxy for age (and we did not know it), it would be entirely useless in discriminating between embryos from the same couple, yet it may still score highly on AUC because age alone is predictive of success in implantation, and we may be fooled into thinking that it is a good model. Here there are two problems that combine to be worse than either alone: a mismatched evaluation metric, and an inscrutable model that does not reveal the problem with either the predictions or the metric.

6.3 Real-Time Error-Checking is Harder with Black-Box Models

The two problems discussed above (information asymmetry and the possibility of confounders) lead to a third problem, namely the difficulty of error-checking the model in real-time as it makes predictions in the clinic. We would want the clinician to be able to determine whether the model is reasoning in a way that is obviously wrong and catch new problems immediately should they arise. For instance, after a change in camera setting, an algorithm might suddenly start thinking that the shape of a current embryo looks like an embryo from the training set with a completely different shape. A clinician could potentially catch that problem immediately if they knew the reasoning process of the model.

6.4 The Economics of “Buying Into” a Brittle Model Does Not Favor Clinicians or Patients

A potential consequence of the problems of information asymmetry and confounders listed above would be that black-box model performance may be brittle to changes from the system it was trained on, and thus would likely be limited to the ecosystem in which it has been shown to work. This means that a clinic using this model would need to buy into that ecosystem, ovarian stimulation regimens, use of incubators and culture medium amongst other potential variables. This gives AI companies a great deal of *economic power* over clinics, potentially increasing treatment costs.

6.5 Overall Troubleshooting is Difficult for Black-Box Models

If the model were more interpretable, it might be easier to troubleshoot broad problems in the model (beyond serious issues that might be noticed in real-time usage). This includes ethical concerns such as sex, disability or racial bias (as we discussed), as well as epistemic issues with accuracy or subtle confounding. If interpretability reveals flawed reasoning processes, the designer would be forced to alter the model to use correct reasoning, leading potentially to more robustness across ecosystems.

7 Interpretable ML as the Way Forward in Embryo Selection

An interpretable ML model is a predictive model that is constrained so that a human can better understand its reasoning process [26, 59]. Interpretable ML is a field that dates to the beginning of AI, back to the days of expert systems. The benefits of interpretable models are clear: by understanding the reasoning processes of predictive models, physicians can troubleshoot them and justify their decisions (to patients, other physicians, and during lawsuits). They would not need to place blind trust in a black-box model. Physicians and interpretable ML models can create a “centaur” that leverages both the information in a database (through ML) and a human’s system-level way of thinking about problems. Furthermore, correlating patient values (chance of disability, sex, single vs. double embryo transfer and chance of implantation) with outcomes can be more easily accommodated by interpretable models.

Focusing on increasing the use of interpretable AI models is an elegant approach to both the epistemic and ethical concerns, by dispelling the opaqueness and allowing precise explanations of model predictions. For instance, models that are not opaque have an advantage because their use preserves existing mechanisms of accountability to a greater extent by allowing clinicians to understand model decisions better and thus retain the responsibility. For models that are opaque, revisions to such mechanisms would be necessary. This argument is far from resolved, but it is a promising reason to favor interpretable models over black-box ones.

To the extent that there exists little or no trade-off between how interpretable a model is and how accurate it is for embryo

selection, interpretable models are thus a promising solution. If it is the case that there is a salient difference in performance between interpretable and non-interpretable models, alternative solutions to both of the above epistemic and ethical concerns might have to be developed, so that we may benefit from the higher predictive accuracy of non-interpretable models. For now, there is no reason to believe that a salient performance difference between interpretable and non-interpretable models would exist. Interpretable models perform just as well even for benchmark datasets in computer vision, and we will go into more detail on this important point shortly. In fact, interpretable models are easier to troubleshoot (as domains change, as unusual cases arise, as cases indicating bias need to be investigated), and thus lead to overall better performance of the model.

A major question in interpretable ML is what interpretability metric to use, as these metrics must (by definition of interpretability in that domain) be domain dependent. For computer vision for natural images, there have been major successful efforts by numerous groups of researchers to create interpretable neural networks that do not lose accuracy over their black-box counterparts. These neural models go well beyond modeling only the “attention” of the network (that is, where the network is looking within an image), and are particularly useful for computer vision problems. Such interpretable neural networks could use different types of logical reasoning processes, including:

- Case-based reasoning (variations of k-nearest neighbors): In this case, the network would point out which parts of a test image are similar to prototypical past cases. The prototypical cases are chosen by the network along with the ways in which images are similar to each other (e.g., Chen et al. [27]). One could envision an embryologist looking at a test image of an embryo, with an interpretable ML model pointing out how parts of it look similar to other prototypical known embryos whose outcome is known. Case-based reasoning models have been developed for different problems in radiology using interpretable ML [60, 61], and thus such methods would be poised for use in embryology.
- Latent space disentanglement, where all information about a single concept (such as mother’s age, or embryo size, density or color) is forced to travel through a single node of a network. Another way to say this is that each axis of a latent space (where an axis corresponds to activation of a node) represents a concept. This helps to understand information flow through the network (e.g., Chen, Bei and Rudin [62]). These types of disentangled models could potentially be useful for separating out the type of equipment, the age of the mother, and other pieces of information that might be embedded within the image of the embryo.
- Networks imbued with logical structure, e.g., probabilistic decision trees. By forcing the network to reason logically, we may be better able to understand its reasoning process (e.g., Wu and Song [63], Li, Song and Wu [64]).

There are many challenges still in designing interpretable neural networks, particularly when the domain experts themselves do not know what constitutes interpretability; in other words, there are many directions for future research.

For problems involving categorical or real data (“tabular” data, rather than image data, time sequence data, or text data), interpretable ML models can also be developed. These can potentially take the form of a medical scoring system, which means a small number of integer “point” values that sum, and translate into a risk (e.g., Ustun and Rudin [65]). For tabular data, neural networks and other forms of black-box models do not seem to provide additional accuracy, which means that optimized medical scoring systems might be as accurate as one could get (depending on the dataset) [66, 67].

An interesting direction for future research is to combine interpretable neural networks for computer vision (to handle the visual data, similar to the work of Leahy et al. [30]) with interpretable models for tabular data (rule lists or decision trees, for instance) to form a global interpretable model that handles these heterogeneous data types.

One key point that has emerged from past research is that as long as one can design the interpretability metric carefully to match the domain, interpretable models tend not to lose accuracy relative to their black-box counterparts [26, 27]. As far as we know, modern interpretable ML methods have not yet been fully applied to IVF.

In Table 1, we summarize the advantages of interpretable AI models over black-box models.

Normative Considerations	Black-Box ML System	Interpretable ML System
Clinical decision-making	Replaced (machine paternalism)	Enhanced
“Centaur” arm of RCT which combines ML and human expertise	Not yet possible – acceptance of black-box recommendation is all-or-nothing	Permits such an arm
Responsibility for treatment success	Unclear	Clearly remains with the clinician
Biases that misrepresent patient values, and have unintended consequences for future people and society	May go unnoticed until societal effects become significant	Easier to detect earlier: the model reports the parameters on which it is basing its decision
Confounders which reduce the true predictive power of the model	May go unnoticed until clinics report underwhelming success rates	Easier to detect earlier: the model reports the parameters on which it is basing its decision
Error-checking	May go unnoticed until clinics report underwhelming success rates	Clearly-erroneous decision-making could be identified by an embryologist

Economics	Clinics may need to purchase specific equipment to guarantee the success of models which require the same equipment used for evaluation	Clinicians could modulate their interpretation of a model’s suggestion under different conditions, if the model can explain the basis of its decisions
Accuracy/capability	There is no evidence to suggest that black-box models are any more accurate than interpretable models	

Table 1: Summary box of reasons why interpretability gains an advantage over black-box models

8 Recommendations

8.1 Rigorous Evaluation with RCTs

Researchers must evaluate AI models to select embryos using the gold standard of RCTs against best clinical judgement or black-box AI, if these have been deployed into practice or show promising results. Key outcomes for evaluation include time to live birth, number of embryo transfers before live birth and associated cost analysis, as well as live birth per egg collection, and health of the baby. Researchers should monitor the effects of the new technology with post-implementation surveillance. Before an RCT is performed, a validation study should be performed on data from different clinics than those used to train the model.

8.2 Interpretable AI

Programmers should aim to build interpretable machine learning models where biologically meaningful parameters guide embryo assessment, reducing the risk of hidden biases in algorithms causing unintended harms to society, permitting better troubleshooting, and better enabling clinicians to counsel their patients on the thinking underlying their treatment.

8.3 Regulatory Oversight for Interpretable AI

The importance of interpretability should be captured in mechanisms of regulatory oversight. Current regulatory approaches attempt to capture medical AI models as a type of medical device; they should further require either that AI model developers not produce black-box models if interpretable models are shown to have similar performance, or that any black-box model must come with the next-best interpretable model considered and trialed. Further, despite the fact that the field of assisted reproductive technology utilizes ‘good practice’ regulation for many advancements (such that violations are not legally punished), this would not suit the many risks of AI in embryo selection as outlined above. A ‘hard’ regulatory stance that promotes interpretable models would be a more advisable approach.

8.4 Access to Code and Data

Data and code used to create ML models should be made publicly accessible. This would enable reproducible research and the advancement of an exciting and important academic field. A high-quality public model would, at the very least, provide a performance baseline for other models.

8.5 Respect for Patient Privacy and Autonomy

Procedures should be put in place for securing patient privacy when data is shared, such as data anonymization. All patients who use AI to select embryos should give fully informed consent, including knowledge of limitations and unknowns, use of data and images, and harms and benefits as shown by RCTs. They should be informed of how a model arrives at a recommendation. Where possible, patient values should be inserted into the reasoning process of selection models.

8.6 Involving the Broader AI Community

Many young ML researchers are eager to get their hands on data to try out the latest techniques, and are passionate about using the technology to make the world a better place. Their participation should be encouraged. Currently, datasets for embryo selection are not broadly available. A naïve release of such data may do more harm than good, potentially inviting simplistic evaluations of ML techniques that fall prey to many of the criticisms we have discussed. Releasing a dataset and suggesting evaluation criteria for it which reflects actual practice, and takes ethical concerns into account, will require a broader discussion between embryologists, ethicists, and researchers in AI and statistics, and will also require addressing privacy concerns. This discussion ought to continue after the data are released. Nonetheless, allowing the broader AI community to see the data and get involved in their analysis will ensure that flawed and biased evaluations do not easily fly under the radar. It will also likely bring other important issues into the open that we have not yet recognized.

We summarize our recommendations in Table 2 below.

<ul style="list-style-type: none">• Use of replicable, interpretable machine learning tools and data• Well designed and conducted RCTs• Post implementation surveillance• Regulatory oversight requiring interpretable AI whenever possible• Funding for public institutions to transparently develop and evaluate machine learning models, and open access to code used in models• Procedures for maintaining security of patient/embryo data whilst permitting ethical data sharing• Fully informed consent to use AI• Inclusion of patient values into AI programmes where possible• Participation from the broader AI community

Table 2: Summary box of recommendations

9 Conclusion

Starting or growing a family is an immensely significant decision; technology which could help individuals who make that decision realize their goal would be invaluable. We see potential for AI in IVF to help couples have children earlier in their treatment and at a lower cost. However, researchers, companies and clinics must ensure that the technology they promote or adopt brings real, measurable benefits to patients and, most importantly, does no harm. In this article, we highlighted limitations of current ML models and the studies which evaluate them, we drew specific attention to the ethical concerns that this technology could introduce in its current form, and suggested a path forward in terms of model design and evaluation. Most importantly, we hope to see interpretable machine learning models that clinicians could understand, troubleshoot and explain to their patients, rigorously evaluated with RCTs. We believe these are essential for creating tools which are fit for use for real individuals, hoping to start or grow a family, in the clinical setting.

10 Conflict of Interest

The authors declare no conflict of interest.

ACKNOWLEDGMENTS

The authors received no funding for this review.

REFERENCES

- [1] European Society for Human Reproduction and Embryology. 2021. Resources. Retrieved January 26, 2021 from <https://www.eshre.eu/Press-Room/Resources>.
- [2] Bart C. Fauser and Robert G. Edwards. 2005. The early days of IVF. *Hum. Reprod. Update*, 11, 5 (Oct. 2005), 437–438. DOI: 10.1093/humupd/dmi026
- [3] Human Fertilisation and Embryology Authority. 2020. Fertility treatment 2018: trends and figures Quality and Methodology Report. Retrieved April 29, 2021 from <https://www.hfea.gov.uk/about-us/publications/research-and-data/fertility-treatment-2018-trends-and-figures/fertility-treatment-2018-quality-and-methodology-report/>.
- [4] US Department of Health and Human Services Centers for Disease Control. 2018. 2018 Assisted Reproductive Technology Fertility Clinic Success Rates Report. Retrieved April 29, 2021 from <https://www.cdc.gov/art/pdf/2018-report/ART-2018-Clinic-Report-Full.Pdf>.
- [5] David K. Gardner, Marcos Meseguer, Carmen Rubio, and Nathan R. Treff. 2015. Diagnosis of human preimplantation embryo viability. *Hum. Reprod. Update*, 21, 6 (Nov. 2015), 727–747. DOI: 10.1093/humupd/dmu064
- [6] David K. Gardner, Michelle Lane, John Stevens, Terry Schlenker, and William B. Schoolcraft. 2000. Blastocyst score affects implantation and pregnancy outcome: towards a single blastocyst transfer. *Fertil. Steril.*, 73, 6 (June 2000), 1155–1158. DOI: 10.1016/S0015-0282(00)00518-5
- [7] Raminta Zmuidinaite, Fady I. Shahara, Ray K. Iles. 2021. Current advancements in noninvasive profiling of the embryo culture media secretome. *Int. J. Mol. Sci.*, 22, 5 (March 2021), 1–13. DOI: 10.3390/ijms22052513
- [8] James M. Kemper, Rui Wang, Daniel L. Rolnik, and Ben W. Mol. 2020. Preimplantation genetic testing for aneuploidy: are we examining the correct outcomes? *Hum. Reprod.*, 35, 11 (Nov. 2020), 2408–2412. DOI: 10.1093/humrep/deaa224
- [9] Yanhe Liu, Katie Feenan, Vincent Chapple, and Phillip Matson. 2019. Assessing efficacy of day 3 embryo time-lapse algorithms retrospectively: impacts of dataset type and confounding factors. *Hum. Fertil. (Camb.)*, 22, 3, 182–190. DOI: 10.1080/14647273.2018.1425919
- [10] Marcos Meseguer, Javier Herrero, Alberto Tejera, Karen M. Hilligsøe, Niels B. Ramsing, and Jose Remoh. 2011. The use of morphokinetics as a predictor of embryo implantation. *Hum. Reprod.*, 26, 10 (Oct. 2011), 2658–2671. DOI: 10.1093/humrep/der256
- [11] Yanhe Liu, Vincent Chapple, Peter Roberts, and Phillip Matson. 2014. Prevalence, consequence, and significance of reverse cleavage by human embryos viewed with the use of the Embryoscope time-lapse video system. *Fertil. Steril.*, 102, 5 (Nov. 2014), 1295–1300.e2. DOI:

- 10.1016/j.fertnstert.2014.07.1235
- [12] Yanhe Liu, Vincent Chapple, Katie Feenan, Peter Roberts, and Phillip Matson. 2015. Clinical significance of intercellular contact at the four-cell stage of human embryos, and the use of abnormal cleavage patterns to identify embryos with low implantation potential: a time-lapse study. *Fertil. Steril.*, 103, 6 (June 2015), 1485–1491.e1. DOI: 10.1016/j.fertnstert.2015.03.017
- [13] Yanhe Liu, Vincent Chapple, Katie Feenan, Peter Roberts, and Phillip Matson. 2016. Time-lapse deselection model for human day 3 in vitro fertilization embryos: The combination of qualitative and quantitative measures of embryo growth. *Fertil. Steril.*, 105, 3, 656–662.e1. DOI: 10.1016/j.fertnstert.2015.11.003
- [14] Yanhe Liu, Fang Qi, Phillip Matson, Dean E. Morbeck, Ben W. Mol, Sai Zhao, and Masoud Afnan. 2020. Between-laboratory reproducibility of time-lapse embryo selection using qualitative and quantitative parameters: a systematic review and meta-analysis. *J. Assist. Reprod. Genet.*, 37, 6, 1295–1302. DOI: 10.1007/s10815-020-01789-4
- [15] Zev Rosenwaks. 2020. Artificial intelligence in reproductive medicine: a fleeting concept or the wave of the future? *Fertil. Steril.*, 114, 5, 905–907. DOI: 10.1016/j.fertnstert.2020.10.002
- [16] Eleonora I. Fernandez, André S. Ferreira, Matheus H. M. Cecílio, Dóris S. Chêles, Rebeca C. M. de Souza, Marcelo F. G. Nogueira, and José C. Rocha. Artificial intelligence in the IVF laboratory: overview through the application of different types of algorithms for the classification of reproductive data. *J. Assist. Reprod. Genet.*, 37, 10, 2359–2376, 2020.
- [17] Susanna Apter, Thomas Ebner, Thomas Freour, Yves Guns, Borut Kovacic, Nathalie Le Clef, Monica Marques, Marcos Meseguer, Debbie Montjean, Ioannis Sfountouris, Roger Sturmey, and Giovanni Coticchio. 2020. Good practice recommendations for the use of time-lapse technology. *Hum. Reprod. Open*, 2020, 2, 1–26. DOI: 10.1093/hropen/hoaa008
- [18] Dimitry Tran, Simon Cooke, Peter J. Illingworth, and David K. Gardner. 2019. Deep learning as a predictive tool for fetal heart pregnancy following time-lapse incubation and blastocyst transfer. *Hum. Reprod.*, 34, 6, 1011–1018. DOI: 10.1093/humrep/dez064
- [19] Pegah Khosravi, Ehsan Kazemi, Qiansheng Zhan, Jonas E. Malmsten, Marco Toschi, Pantelis Zisimopoulos, Alexandros Sigaras, Stuart Lavery, Lee A. D. Cooper, Cristina Hickman, Marcos Meseguer, Zev Rosenwaks, Olivier Elemento, Nikica Zaninovic, and Iman Hajirasouliha. 2019. Deep learning enables robust assessment and selection of human blastocysts after in vitro fertilization. *npj Digit. Med.*, 2, 1, 1–9. DOI: 10.1038/s41746-019-0096-y
- [20] Yoav Kan-Tor, Assaf Ben-Meir, and Amnon Buxboim. 2020. Can deep learning automatically predict fetal heart pregnancy with almost perfect accuracy? *Hum. Reprod.*, 35, 6, 1473–1473. DOI: 10.1093/humrep/deaa083
- [21] Carol L. Curchoe, Adolfo Flores-Saiffe Farias, Gerardo Mendizabal-Ruiz, and Alejandro Chavez-Badiola. 2020. Evaluating predictive models in reproductive medicine. *Fertil. Steril.*, 114, 5, 921–926. DOI: 10.1016/j.fertnstert.2020.09.159
- [22] Australia New Zealand Clinical Trials Registry. 2020. Trial Registered on ANZCTR. Retrieved April 29, 2021 from <https://www.anzctr.org.au/Trial/Registration/TrialReview.aspx?id=379161&isReview=true>.
- [23] Myura Nagendran, Yang Chen, Christopher A. Lovejoy, Anthony C. Gordon, Matthieu Komorowski, Hugh Harvey, Eric J. Topol, John P. A. Ioannidis, Gary S. Collins, and Mahiben Maruthappu. 2020. Artificial intelligence versus clinicians: systematic review of design, reporting standards, and claims of deep learning studies. *BMJ*, 368 (March 2020), m689. DOI: 10.1136/bmj.m689
- [24] Masoud A. M. Afnan, Khalid S. Khan, and Ben W. Mol. 2020. Generating translatable evidence to improve patient care: the contribution of human factors. *Reprod. Biomed. Online*, 41, 3 (Sep. 2020), 353–356. DOI: 10.1136/bmj.m689
- [25] Jack Wilkinson, Phillipa Malpas, Karin Hammarberg, Pamela Mahoney Tsigdinos, Sarah Lensen, Emily Jackson, Joyce Harper and Ben W. Mol. 2019. Do à la carte menus serve infertility patients? The ethics and regulation of in vitro fertility add-ons. *Fertil. Steril.*, 112, 6 (Dec. 2019), 973–977. DOI: 10.1016/j.fertnstert.2019.09.028
- [26] Cynthia Rudin. 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat. Mach. Intell.*, 1, 5, 206–215. DOI: 10.1038/s42256-019-0048-x
- [27] Chaofan Chen, Oscar Li, Daniel Tao, Alina Barnett, Cynthia Rudin, and Jonathan K. Su. 2019. This Looks Like That: Deep Learning for Interpretable Image Recognition. In *33rd Conference on Proceedings of Neural Information Processing Systems (NeurIPS)*, December 8 – 14, 2019, Vancouver, Canada. 8930–8941.
- [28] Alejandro Chavez-Badiola, Adolfo Flores-Saiffe Farias, Gerardo Mendizabal-Ruiz, Rodolfo Garcia-Sanchez, Andrew J. Drakeley, and Juan P. Garcia-Sandoval. 2020. Predicting pregnancy test results after embryo transfer by image feature extraction and analysis using machine learning. *Sci. Rep.*, 10, 1, 4394. DOI: 10.1038/s41598-020-61357-9
- [29] Robert Milewski, Agnieszka Kuczyńska, Bożena Stankiewicz, and Waldemar Kuczyński. 2017. How much information about embryo implantation potential is included in morphokinetic data? A prediction model based on artificial neural networks and principal component analysis. *Adv. Med. Sci.*, 62, 1, 202–206. DOI: 10.1016/j.advms.2017.02.001
- [30] Brian D. Leahy, Won-Dong Jang, Helen Y. Yang, Robbert Struyven, Donglai Wei, Zhe Sun, Kylie R. Lee and Charlotte Royston. 2020. Automated Measurements of Key Morphological Features of Human Embryos for IVF. *Med. Image Comput. Comput. Assist. Interv.*, 12265, 25–35. DOI: 10.1007/978-3-030-59722-13
- [31] Behnaz Raef, Masoud Maleki, and Reza Ferdousi. 2020. Computational prediction of implantation outcome after embryo transfer. *Health Informatics J.*, 26, 3, 1810–1826. DOI: 10.1177/1460458219892138
- [32] Dinora A. Morales, Endika Bengoetxea, Pedro Larrañaga, Miguel García, Yosu Franco, Mónica Fresnada, and Marisa Merino. 2008. Bayesian classification for the selection of in vitro human embryos using morphological and clinical data. *Comput. Methods Programs Biomed.*, 90, 2, 104–116. DOI: 10.1016/j.cmpb.2007.11.018
- [33] Julian Savulescu, Iain Chalmers, and Jennifer Blunt. 1996. Are research ethics committees behaving unethically? Some suggestions for improving performance and accountability. *BMJ*, 313, 7069 (Nov. 1996), 1390–1393. DOI: 10.1136/bmj.313.7069.1390
- [34] Cathy Charles, Amiram Gafni, and Tim Whelan. 1997. Shared decision-making in the medical encounter: What does it mean? (or it takes at least two to tango). *Soc. Sci. Med.*, 44, 5, 681–692. DOI: 10.1016/S0277-9536(96)00221-3
- [35] Jens C. Bjerring and Jacob Busch. 2020. Artificial Intelligence and Patient-Centered Decision-Making. *Philos. Technol.* DOI: 10.1007/s13347-019-00391-6
- [36] Abhishek Mishra, Julian Savulescu, and Alberto Giubilini. Forthcoming. Ethics of Medical AI. In *Oxford Handbook of Digital Ethics*.
- [37] Mirjam M. Garvelink, Laura Boland, Krystal Klein, Don Vu Nguyen, Matthew Menear, Hilary L. Bekker, Karen B. Eden, Annie LeBlanc, Annette M. O'Connor, Dawn Stacey, and France Légaré. 2019. Decisional Conflict Scale Use over 20 Years: The Anniversary Review. *Med. Decis. Mak.*, 39, 4, 301–314. DOI: 10.1177/0272989X19851345
- [38] Glyn Elwyn, Adrian Edwards, Michelle Wensing, Karenza Hood, Christine Atwell, and Richard Grol. Shared decision making: developing the OPTION scale for measuring patient involvement. *Qual. Saf. Heal. Care*, 12, 2 (April 2003), 93 LP – 99.
- [39] France Légaré, Stephen Kearing, Kate Clay, Susie Gagnon, Denis D'Amours, Michel Rousseau and Annette O'Connor. 2010. Are you SURE?: Assessing patient decisional conflict with a 4-item screening test. *Can. Fam. Physician*, 56, 8 (Aug. 2010), e308–e314.
- [40] Raja Parasuraman and Dietrich H. Manzey. 2010. Complacency and Bias in Human Use of Automation: An Attentional Integration. *Hum. Factors*, 52, 3 (June 2010), 381–410.
- [41] Max F. Kramer, Jana Schaich Borg, Vincent Conitzer, and Walter Sinnott-Armstrong. 2018. When Do People Want AI to Make Decisions?. In *AIES 2018 - Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*. ACM Inc., New York, NY, 204–209. <https://doi.org/10.1145/3278721.3278752>
- [42] Ange Wang, Jonathan Kort, Barry Behr, and Lynn M. Westphal. 2018. Euploidy in relation to blastocyst sex and morphology. *J. Assist. Reprod. Genet.*, 35, 9 (Sep. 2018), 1565–1572. DOI: 10.1007/s10815-018-1262-x
- [43] Juan J. Tarin, Miguel A. García-Pérez, Carlos Hermenegildo, and Antonio Cano. 2014. Changes in sex ratio from fertilization to birth in assisted-reproductive-treatment cycles. *Reprod. Biol. Endocrinol.*, 12, 56 (June 2014). DOI: 10.1186/1477-7827-12-56
- [44] Fernando Bronet, Maria-Carmen Nogales, Eva Martinez, Marta Ariza, Carmen Rubio, Juan-Antonio Garcia-Velasco, and Marcos Meseguer. 2015. Is there a relationship between time-lapse parameters and embryo sex?. *Fertil. Steril.*, 103, 2, 396. DOI: 10.1016/j.fertnstert.2014.10.050
- [45] Bo Huang, Xinling Ren, Lixia Zhu, Li Wu, Huiping Tan, Na Guo, Yulan Wei, Juan Hu, Qun Liu, Wen Chen, Jing Liu, Dan Li, Shujie Liao, and Lei Jin. 2019. Is differences in embryo morphokinetic development significantly associated with human embryo sex?. *Biol. Reprod.*, 100, 3 (March 2019), 618–623. DOI: 10.1093/biolre/iy0229
- [46] Rosalind J. McDougall. Computer knows best? The need for value-flexibility in medical AI. 2019. *J. Med. Ethics*, 45, 3 (March 2019), 156 LP – 160. DOI: 10.1136/medethics-2018-105118
- [47] Nicholas Agar. 2004. *Liberal Eugenics: In Defense of Human Enhancement*. Blackwell Publishing, Oxford, UK.
- [48] Julian Savulescu. 2001. Procreative beneficence: why we should select the best children. *Bioethics*, 15, 5–6, 413–426. DOI: 10.1111/1467-8519.00251
- [49] Julian Savulescu and Guy Kahane. 2009. The Moral Obligation to Create Children With the Best Chance of the Best Life. *Bioethics*, 23, 5 (June 2009), 274–290. DOI: 10.1111/j.1467-8519.2008.00687.x

- [50] Julian Savulescu and Guy Kahane. Understanding procreative beneficence. 2016. *Oxford Handbooks Online OP*. DOI: 10.1093/oxfordhb/9780199981878.013.26
- [51] D. Parfit. 1984. *Reasons and Persons*. Oxford University Press, Oxford, UK.
- [52] Julian Savulescu, Melanie Hemsley, Ainsley Newson, and Bennett Foddy. 2006. Behavioural Genetics: Why Eugenic Selection is Preferable to Enhancement. *J. Appl. Philos.*, 23, 2 (May 2006), 157–171. DOI: 10.1111/j.1468-5930.2006.00336
- [53] Allen Buchanan, Dan W. Brock, Norman Daniels, and Daniel Wikler. 2000. *From Chance to Choice: Genetics and Justice*. Cambridge University Press, Cambridge, UK.
- [54] Bjørn Hofmann. 2017. ‘You are inferior!’ Revisiting the expressivist argument. *Bioethics*, 31, 7, 505–514. DOI: 10.1111/bioe.12365
- [55] Daniel Schönberger. 2019. *Artificial intelligence in healthcare: A critical analysis of the legal and ethical implications*, 27, 2. DOI: 10.1093/ijlit/ez004
- [56] W. Nicholson Price II, Sara Gerke, and I. Glenn Cohen. 2019. Potential Liability for Physicians Using Artificial Intelligence. *JAMA*, 322, 18 (Nov. 2019), 1765–1766. DOI: 10.1001/jama.2019.15064
- [57] Matt O’Connor. 2021. Algorithm’s ‘unexpected’ weakness raises larger concerns about AI’s potential in broader populations. Retrieved April 29, 2021 from https://www.healthimaging.com/topics/artificial-intelligence/weakness-ai-broader-patient-populations?utm_source=newsletter&utm_medium=rb_news.
- [58] Andrew F. Voter, Ece Meram, John W. Garrett, and John-Paul J. Yu. 2010. Diagnostic Accuracy and Failure Mode Analysis of a Deep Learning Algorithm for the Detection of Intracranial Hemorrhage. *J. Am. Coll. Radiol.* DOI: 10.1016/j.jacr.2021.03.005
- [59] Cynthia Rudin, Chaofan Chen, Zhi Chen, Haiyang Huang, Lesia Semenova, and Chudi Zhong. 2021. Interpretable Machine Learning: Fundamental Principles and 10 Grand Challenges. arXiv: 2103.11251. Retrieved from <http://arxiv.org/abs/2103.11251>.
- [60] Alina J. Barnett, Fides R. Schwartz, Chaofan Tao, Chaofan Chen, Yinhao Ren, Joseph Y. Lo, Cynthia Rudin. 2021. IAIA-BL: A Case-based Interpretable Deep Learning Model for Classification of Mass Lesions in Digital Mammography. arXiv: 2103.12308. Retrieved from <http://arxiv.org/abs/2103.12308>.
- [61] Eunji Kim, Siwon Kim, Minji Seo, and Sungroh Yoon. 2021. XProtoNet: Diagnosis in Chest Radiography with Global and Local Explanations. arXiv: 2103.10663. Retrieved from <http://arxiv.org/abs/2103.10663>.
- [62] Zhi Chen, Yijie Bei, and Cynthia Rudin. 2020. Concept whitening for interpretable image recognition. *Nat. Mach. Intell.*, 2, 12, 772–782. DOI: 10.1038/s42256-020-00265-z
- [63] Tianfu Wu and Xi Song. 2019. Towards interpretable object detection by unfolding latent structures. In *Proceedings of the IEEE International Conference on Computer Vision*, 6032–6042. <https://doi.org/10.1109/ICCV.2019.00613>
- [64] X. Li, X. Song, and T. Wu. 2017. AOGNets: Compositional grammatical architectures for deep learning. arXiv: 1711.05847. Retrieved from: <https://arxiv.org/abs/1711.05847>.
- [65] Berk Ustun and Cynthia Rudin. 2019. Learning optimized risk scores. *J. Mach. Learn. Res.*, 20, 1–75.
- [66] Cynthia Rudin and Berk Ustun. 2018. Optimized Scoring Systems: Toward Trust in Machine Learning for Healthcare and Criminal Justice. *INFORMS J. Appl. Anal.*, 48, 5 (Oct. 2018), 449–466. DOI: 10.1287/inte.2018.0957
- [67] Yin Lou, Rich Caruana, Johannes Gehrke, and Giles Hooker. 2013. Accurate intelligible models with pairwise interactions. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. Part F1288, 623–631. <https://doi.org/10.1145/2487575.2487579>