Embryo Graphs: Predicting Human Embryo Viability from 3D Morphology

Chloe He[1,2,3], Neringa Karpavičiūtė[3,4], Rishabh Hariharan[3], Céline Jacques[5], Jérôme Chambost[5], Jonas Malmsten[6], Nikica Zaninovic[6], Koen Wouters[7], Thomas Fréour[8], Cristina Hickman[9], and Francisco Vasconcelos[1,2]

[1] Wellcome/EPSRC Centre for Interventional and Surgical Sciences, University College London, UK
chloe.he.21@ucl.ac.uk
[2] Department of Computer Science, University College London, UK
[3] Apricity, UK
[4] Department of Life Sciences, Manchester Metropolitan University, UK
[5] Apricity, France
[6] Ronald O. Perelman and Claudia Cohen Center for Reproductive Medicine, Weill Cornell Medicine, USA
[7] Brussels IVF, University Hospital Brussels, Belgium
[8] Department of Reproductive Biology, University Hospital of Nantes, France
[9] Institute of Reproductive and Developmental Biology, Imperial College London, UK

Abstract. Embryo selection is a critical step in the process of in-vitro fertilisation in which embryologists choose the most viable embryos for transfer into the uterus. In recent years, numerous works have used computer vision to perform embryo selection. However, many of these works have neglected the fact that the embryo is a 3D structure, instead opting to analyse embryo images captured at a single focal plane. In this paper we present a method for the 3D reconstruction of cleavage-stage human embryos. Through a user study, we validate that our reconstructions align with expert assessments. Furthermore, we demonstrate the utility of our approach by generating graph representations that capture biologically relevant features of the embryos. In pilot experiments, we train a graph neural network on these representations and show that it outperforms existing methods in predicting live birth from euploid embryo transfers. Our findings suggest that incorporating 3D reconstruction and graph-based analysis can improve automated embryo selection.

Keywords: 3D Reconstruction · Embryology · Microscopy · Graph Neural Networks

1 Introduction

In-vitro fertilisation (IVF) is a common treatment for infertility, a condition which globally affects 1 in 6 people at some point in their lives [36]. It involves creating and incubating embryos outside the body before transferring them back into the patient. A key task in the IVF workflow is deciding which embryos to

prioritise for transfer as not all embryos created during IVF are viable and transferring multiple embryos at a time increases the risk of pregnancy complications. Among the simplest and most widely used methods of embryo selection is morphological assessment which is conducted by specialist embryologists under a microscope [26].

In recent years, the medical imaging community has tried its hand at automating morphological embryo assessment, the tool of choice: convolutional neural networks (CNNs). Nevertheless, these approaches have neglected the fact that embryos are 3D objects. Most CNN-based methods are only applied to images captured at a single focal plane, ignoring the rich structural information present in focal planes above and below. In the few works that do consider multiple (typically up to three) focal planes, 3D structure is only implicitly captured by the learned convolutional filters [20].

Concurrently, there has been a growing body of work suggesting the spatial organisiation of cells in a pre-implantation human embryo to be predictive of its viability from as early as the 4-cell stage [2,13]. Notwithstanding, owing to the time-consuming nature of assessing the spatial organisation of embryos in a clinical setting without the use of dyes or advanced optical setups - a task typically undertaken manually by embryologists - the topic remains relatively under-explored.

In this work, we present an automated pipeline for the 3D reconstruction of clinical human embryos imaged through the Hoffman modulation contrast (HMC) microscopes commonly built into embryo incubators. Through a user study, we verify that the 3D reconstructions are mostly consistent with embryologist assessments and thus that our system shows promise in enabling future large-scale investigations into the spatial organisation of embryos. We further demonstrate how a novel graph-based representation of embryos that explicitly captures biologically-relevant information on the spatial organisation of an embryo can be derived from 3D reconstructions and be used for embryo selection. To our knowledge, our approach marks the first application of graph-based deep learning to human embryology.

2 Related Work

Computer Vision in Embryo Selection The majority of works exploring vision-based embryo selection have made use of CNNs trained in a supervised fashion on images captured through time-lapse embryo incubators. Such approaches have yielded promising results for the prediction of clinical endpoints such as blastocyst formation [23,6], blastocyst grade [19,20,22], ploidy [10], pregnancy [9,4,8] and live birth [3,16]. More recent works have attempted to build on the performance of these through techniques such as self-supervised pre-training [21], semi-supervised learning [9,27] and transformer-based approaches [22]. Nevertheless, although most incubators are capable of capturing the embryo across various focal planes (thereby providing some 3D information), most works opt to work with one or a few central focal planes. To our knowledge, only a few

works have considered explicitly modelling and working with the 3D spatial organisation of human embryos [11].

3D Reconstruction in Live-Cell Imaging Microscopic imaging of live specimens presents several challenges, especially when said live specimens are clinically relevant such as in clinical microendoscopy and embryology. Typically, 3D reconstruction is conducted using confocal and light-sheet microscopy on stained or autoflourescent samples [7]. These methods aim to mitigate the impact that light scattering has on images, namely the presence of out-of-focus light and thus 3D reconstruction of specimens captured with these modalities are relatively straightforward to reconstruct by segmentation and registration [18]. In settings such as clinical embryology, however, confocal (and, to a somewhat lesser extent, light-sheet microscopy), are infeasible due to expense, the toxicity of stains, as well as the risk of phototoxicity to the live specimens [17,7]. As a result, focal stacks captured with conventional light microscopy techniques such as phase contrast and HMC microscopy are commonly used in clinical settings, as their safety is more well-established [17]. 3D reconstruction from such modalities is complicated by the impact of light scattering. Many existing methods attempt to address this challenge by identifying in-focus regions of the image through techniques such as texture analysis [11]. Other methods have utilised forward models describing the physics of the optical setups, casting 3D reconstruction as an optimisation problem recovering the input volume that formed the image [38].

Graph Neural Networks in Histopathology Over recent years, graph neural networks (GNNs) have become increasingly adopted in computational histopathology. A key driver from this stems from the necessity of capturing long-range context when modelling tissue architecture in tasks such as whole-slide imaging [32]. To date, GNNs have been used across multiple tissue types (such as for breast [25,28,29], gastointestinal [33,34,35] and lung [1] cancers) and scales (ranging from individual cells [39,33] to whole regions of tissue [25,29]). Across most of these works, graphs are constructed with nodes representing histological points of interest (such as cells or tissues), and edges representing potential interactions between nodes. In practice, edges are typically created according to locality criteria [25,34,29]. To our knowledge, however, no prior works have applied graph-based deep learning approaches to human embryology.

3 Methods

In this section, we introduce our methodology for producing 3D reconstructions from embryo focal stacks. We then outline how these reconstructions can be used to construct biologically-grounded graph representations of embryos to predict clinical outcomes. An overview of our methods can be seen in Figure 1.

Cell Segmentation We consider an arbitrary HMC focal stack comprising 11 evenly spaced focal planes. We use a Mask-RCNN [15] with a ResNet-50-FPN

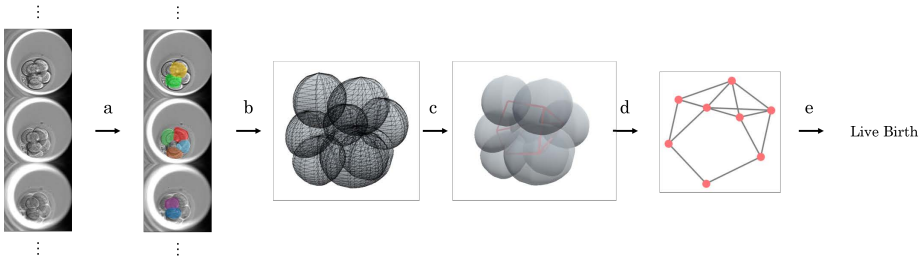

Fig. 1. Summary of our methodology. (a) We start by performing cell segmentation on a focal stack. (b) From these segmentations, we generate a 3D reconstruction. (c) From the 3D reconstruction, we compute intercellular contacts and (d) derive a graph representation. (e) We predict the embryo's fate using a GNN.

backbone to detect individual cells with the class label corresponding to the focal plane at which the cell is most in-focus. The network's input layer accepts an 11-channel 400×400 image where each channel represents a focal plane. We use the standard Mask R-CNN training setup and multitask loss defined in [15]. We introduce a simple yet, to our knowledge, novel non-maximum suppression (NMS) algorithm for focal stacks which we use to post-process detector outputs and reduce duplicate detections. The algorithm, which we shall refer to as *Stack NMS*, can be found in Algorithm 1. The key difference between our method and the classical NMS algorithm is that we introduce the possibility that two overlapping detections may indeed be valid (for example, representing cells that are directly on top of each other). In particular, we only suppress a detection if an overlapping detection of higher confidence is within some window size ϕ.

3D Reconstruction A mesh is generated for each detected cell, given several shape priors. First, we assume that cells are roughly the same dimensions along each axis. In particular, we take the radii of the cells along the depth axis to be the mean of their width and height in the most in-focus plane. Second, we assume that the cross section of the cell remains the same at each focal plane bar its scale. We can thus generate meshes for each blastomere in the following manner (a helpful diagram of the process can be found in the Supplementary Materials):

1. Convert the outline of the blastomere's equatorial plane into a polygon and place the polygon at depth d corresponding to the depth at which the equatorial cross section is most in-focus.
2. Create copies of the outline polygon and place them above and below the equatorial focal plane at regular fixed intervals within the range $[d - \frac{w+h}{4}, d + \frac{w+h}{4}]$ where w and h are the width and height of the blastomere's equatorial cross section respectively.
3. Scale each polygon such that the polygon at depth $d_{current}$ is scaled by a factor of $\sqrt{1 - 16(\frac{d - d_{current}}{w+h})^2}$.

Algorithm 1: Stack NMS

Input: Predicted bounding boxes $b_k \in B^{pred}$, predicted depths $d_k \in D^{pred}$, prediction confidences $c_k \in C^{pred}$, intersection-over-union threshold η, confidence threshold ζ and window size ϕ.

Output: A set of indices of bounding boxes to be suppressed O^{sup}.

Initialise $O^{sup} := \{k \mid c_k \leq \zeta,\ c_k \in C^{pred}\}$.
Compute sets $O_i \in O$ such that $\forall j, k \in O_i,\ IoU(b_j, b_k) > \eta$, $\forall j \neq i, k \in O_i \iff k \notin O_j$ and $\forall k \in B^{pred} \exists i,\ k \in O_i$.
for $i = 1..|O|$ **do**
  if $|O_i| = 1$ **then**
    continue
  else
    for $j \in O_i$ **do**
      $w := \{c_k \mid |d_j - d_k| \leq \phi, k \in O_i\}$
      if $c_j \neq max(w)$ **then**
        $O^{sup} := O^{sup} \cup \{j\}$
      end
    end
  end
end
return O^{sup}.

4. Generate triangles between the adjacent polygons (this is straightforward as there are an equal number of vertices in each outline mesh).

Biologically-Grounded Graph Representations We consider an undirected graph $G = \langle V, E \rangle$, representing a particular cleavage stage embryo. Each $v_i \in V$ represents a set of features associated with the ith cell in the embryo. Each edge $e_{ij} \in E$ represents intercellular contact between cells i and j.

In this work, node features v_i comprise both features pertaining to a cell (namely its width and height) as well as graph-theoretic features (in particular, the degree of the node). The presence of intercellular contact between two arbitrary cells i and j is assessed by scaling the mesh of each cell by a factor of 1.05, centered at the center of the cell. If the scaled meshes overlap, we assume there is cell contact and create edge e_{ij}.

We make predictions on the fate of an embryo from its graph representation using a graph isomorphism network (GIN) [37]. Node embeddings x_i are obtained through three rounds of GIN convolutions $x_i := h_\theta(x_i + \sum_{j \in \mathcal{N}(i)} x_j)$ where $\mathcal{N}(i)$ denotes the set of neighbors of node i and h_θ is a perceptron with a single hidden layer parameterised by θ. Global mean pooling is performed across the whole graph to aggregate node embeddings into a single vector which is then passed through a single linear layer. Further architecture details can be found in the Supplementary Materials.

4 Experiments

Experiments were performed on de-identified private embryo imaging datasets derived from three IVF clinics across three countries. The dataset consisted of HMC focal stacks consisting of 7-11 focal planes captured on Embryoscope time-lapse incubators between 2018 and 2021. Further details on the specific data used in each experiment can be found in the Supplementary Materials. Focal stacks were resized to dimensions 400×400. Contrast limited adaptive histogram equalisation was applied to each focal stack and a circular mask was applied to hide the edges of the microscope well. Super-Focus standardisation [12] was applied where appropriate.

All experiments were implemented using PyTorch (v1.13.1) and PyG (v2.2.0) on a Ubuntu 20.04 machine with an NVIDIA Titan X GPU. Mesh generation was conducted in the Unity engine (v2021.3.7f1). Further details on training and hyperparameter search for each experiment can be found in the Supplementary Materials. Code is available at `https://github.com/chlohe/embryo-graphs`.

Evaluation of Cell Segmentation We evaluate the performance of our cell segmentation model, along with our proposed NMS algorithm. We use a dataset of 498 cleavage-stage embryos annotated by a team of embryologists and researchers. The annotators were instructed to identify cells and large fragments in the stacks and to annotate them on the focal planes they were most in focus. Five instances of the segmentation model were trained for 2000 epochs ($\sim$56 hours) each using a 5-fold cross-validation setup.

We benchmark our NMS algorithm against the classic NMS algorithm [15] and Crowd NMS [24], an NMS algorithm specifically developed for embryo cell segmentation. Hyperparameters for the NMS algorithms were calibrated by averaging their performance when applied to the outputs of the different segmentation model instances on a separate dataset consisting of 92 focal stacks. The final segmentation model-NMS pipeline was tested on an independent testing set consisting of 62 focal stacks, once again making use of the different segmentation model instances. The quality of segmentation was evaluated using the Dice score, while the performance at detecting cells was evaluated using the precision,

Table 1. Performance comparison between non-maximum suppression (NMS) algorithms averaged over 5-fold cross validation. Standard deviations are in parentheses. Highest performances for each metric are highlighted in **bold**. An asterisk (*) is used to denote that a given result is significantly higher than the next best method ($p < 0.05$).

Method	Precision	Recall	F1
Classic NMS	**0.93 (0.01)**	0.81 (0.01)	0.87 (0.01)
Crowd NMS	0.87 (0.07)	0.84 (0.02)	0.86 (0.01)
Stack NMS	**0.93 (0.01)**	**0.89 (0.01)***	**0.91 (0.01)***

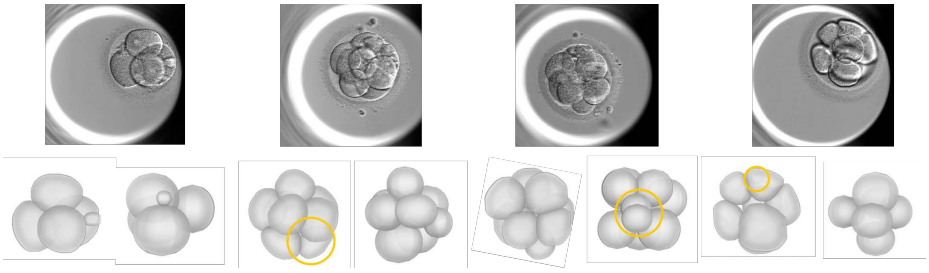

Fig. 2. Examples of 3D reconstructions from the user study. For each embryo, a central plane from the original focal stack is provided along with a top-down and side views. Issues raised by the embryologists are circled. The leftmost embryo received a mean score of 4.4 with no issues identified. The second embryo from the left received an mean score of 3.7, with embryologists noting an incorrectly shaped cell. The third embryo from the left received a mean score of 3.7 and contained a duplicated cell. The rightmost embryo received a mean score of 4.1, though embryologists noted a hallucinated cell.

recall and F1 metrics. Two-tailed unpaired t-tests at a 95% confidence level with Bonferroni correction were used to compare methods for each metric.

Our segmentation model paired with Stack NMS ($\phi = 1$) achieved a Dice score of 0.944 ± 0.001 (M±SD). Results on cell detection can be found in Table 1. Our method outperformed the others in terms of recall and F1 score, surpassing the purpose-built Crowd NMS in recall while retaining high precision.

Evaluation of 3D Reconstruction We evaluate the quality of 3D reconstructions through a user study involving 9 embryologists with a total of 131 years of clinical experience between them. 3D reconstructions of 15 cleavage-stage embryos captured between 35 and 70 hours post insemination were evaluated by individual embryologists using the following scale: 1 (very poor - unrecognisable from focal stack), 2 (poor - major issues such as missing cells), 3 (ok - minor issues such as duplicated cells), 4 (good - only very minor issues), 5 (perfect). The Sketchfab platform was used to display the 3D models and responses were collected through Google Forms. Embryologists were also asked to note down issues they observed with the reconstructions. We considered a reconstruction to have a certain issue if at least a third of the embryologists brought it up.

The reconstructions received a mean rating of 3.8 ± 0.4 (M±SD) indicating, for the most part, only minor issues with reconstructions. The most common issues raised were incorrect cell shapes (N=6), duplicate cell detections (N=5), missing large fragments (N=3) and non-duplicate false positive detections (or "hallucinations", N=3). On closer inspection of the comments, many reports of incorrect cell shapes stemmed from certain cells seeming too large. This may have been a side-effect of the visualisation platform, which used a perspective camera model that may not completely recapitulate microscope optics which may look more like an orthographic projection - a matter that should be considered in future work. An illustration of this phenomenon can be found in the Supplemen-

tary Materials. Moreover, though our NMS method outperformed the others in terms of F1 score, the results of our user study indicate that there is still much room for improvement on NMS and indeed segmentation in this challenging setting. In light of these results, we add an additional layer of human-in-the-loop validation to all subsequent experiments, removing any false positive detections manually.

Predictive Performance Comparison We apply our GNN-based approach to the task of predicting whether an euploid stage human embryo results in a live birth using information available at the 8-cell stage. This is a task of significant clinical interest, as it is not yet fully understood why embryos deemed chromosomal normal do not always lead to live births [5]. Our method was evaluated against three baseline approaches: a ResNet-50 CNN [14] (used as the basis for many other works in the embryo selection literature [9,31]), logistic regression on the KIDScore D3 [30] (an established commercially-available embryo quality scoring algorithm based on the timing of cell divisions) and logistic regression on the total number of cell contacts in the embryo [13]. The dataset used comprised data from 80 euploid embryo transfers and included focal stacks captured at the start of the 8-cell stage along with expert annotations of the cell division timings necessary to calculate the KIDScore D3.

All methods were benchmarked using 10-times-repeated 5-fold cross validation. While the GNN was trained for 5 epochs ($<$ 1 minute) from scratch, the CNN was trained for 20 epochs ($<$ 1 minute) using transfer learning with weights pre-trained on ImageNet. The Adam optimizer was used to fit both models. The total number of cell contacts in the embryo was computed from the same graph representations used by the GNN. To determine the statistical significance of performance differences, we conducted two-tailed unpaired t-tests at a 95% confidence level with Bonferroni correction, comparing the highest-performing model for each metric against the next best-performing model.

The results of the performance comparison can be seen in Table 2. Our method outperformed the baselines on accuracy, precision, recall and F1 score, and matched the total contacts baseline in terms of AUC. Our results suggest

Table 2. Performance comparison between methods averaged over 10-times repeated 5-fold cross validation. Standard deviations are given in parentheses. Highest performances for each metric are highlighted in **bold**. An asterisk (*) is used to denote that a given result is significantly higher than the next best method ($p < 0.05$).

Method	Accuracy	Precision	Recall	F1	AUC
CNN	0.52 (0.10)	0.55 (0.10)	0.71 (0.18)	0.61 (0.10)	0.50 (0.10)
KIDScore D3	0.56 (0.07)	0.58 (0.42)	0.14 (0.11)	0.21 (0.14)	0.53 (0.07)
Total Contacts	0.62 (0.12)	0.59 (0.18)	0.57 (0.18)	0.57 (0.14)	**0.62 (0.12)**
GNN (Ours)	**0.64 (0.09)**	**0.64 (0.08)**	**0.80 (0.15)***	**0.71 (0.09)***	**0.62 (0.09)**

that graph representations in embryology may prove a powerful tool, especially with smaller datasets. Moreover, our approach may also be more interpretable than CNN-based systems given the strong biological priors introduced by the graph representation.

5 Conclusions and Future Work

In this paper, we propose a technique to create 3D reconstructions of early-stage human embryos without the need for invasive procedures. We demonstrate that these reconstructions can be used to derive biologically-grounded graph representations of embryos. In initial experiments on a small dataset, a GNN trained on these representations outperforms existing techniques. Beyond its direct relevance to clinical embryology, our technique offers a pathway to gather spatial data from embryos in clinical settings. This could substantially augment the pool of data available for fundamental research into embryo development and cell fate determination making our non-invasive approach a potentially worthwhile advancement.

Disclosure of Interests. Chloe's PhD is supported by Apricity and the Wellcome/EPSRC Centre for Interventional and Surgical Sciences. Cristina is founder of Avenues, co-founder of Ovom Care, and chief clinical officer of Fairtility.

References

1. Adnan, M., Kalra, S., Tizhoosh, H.R.: Representation learning of histopathology images using graph neural networks. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). IEEE (Jun 2020)
2. Ajduk, A., Zernicka-Goetz, M.: Polarity and cell division orientation in the cleavage embryo: from worm to human. Molecular Human Reproduction **22**(10), 691–703 (Dec 2015)
3. Bodri, D., et al.: Predicting live birth by combining cleavage and blastocyst-stage time-lapse variables using a hierarchical and a data mining-based statistical model. Reproductive Biology **18**, 355–360 (12 2018)
4. Bormann, C.L., et al.: Performance of a deep learning based neural network in the selection of human blastocysts for implantation. eLife **9**, 1–14 (9 2020)
5. Cimadomo, D., et al.: Opening the black box: why do euploid blastocysts fail to implant? a systematic review and meta-analysis. Human Reproduction Update **29**(5), 570–633 (May 2023)
6. d'Estaing, S.G., et al.: A machine learning system with reinforcement capacity for predicting the fate of an art embryo. Systems biology in reproductive medicine **67**, 64–78 (2 2021)
7. Domingo-Muelas, A., et al.: Human embryo live imaging reveals nuclear dna shedding during blastocyst expansion and biopsy. Cell **186**(15), 3166–3181.e18 (Jul 2023)
8. Enatsu, N., et al.: A novel system based on artificial intelligence for predicting blastocyst viability and visualizing the explanation. Reproductive Medicine and Biology **21** (1 2022)

9. Erlich, I., et al.: Pseudo contrastive labeling for predicting ivf embryo developmental potential. Scientific Reports **12** (12 2022)
10. Gheselle, S.D., et al.: Machine learning for prediction of euploidy in human embryos: in search of the best-performing model and predictive features. Fertility and Sterility **117**, 738–746 (4 2022)
11. Giusti, A., et al.: Blastomere segmentation and 3d morphology measurements of early embryos from hoffman modulation contrast image stacks. In: 2010 IEEE International Symposium on Biomedical Imaging: From Nano to Macro. IEEE (2010)
12. He, C., et al.: Super-Focus: Domain Adaptation for Embryo Imaging via Self-supervised Focal Plane Regression, p. 732–742. Springer Nature Switzerland (2022)
13. He, C., et al.: Seeking arrangements: cell contact as a cleavage-stage biomarker. Reproductive BioMedicine Online p. 103654 (2023)
14. He, K., et al.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2016)
15. He, K., et al.: Mask r-cnn. In: 2017 IEEE International Conference on Computer Vision (ICCV). pp. 2980–2988 (2017). https://doi.org/10.1109/ICCV.2017.322
16. Huang, B., et al.: Using deep learning to predict the outcome of live birth from more than 10,000 embryo data. BMC pregnancy and childbirth **22**, 36 (1 2022)
17. Iyer, S., Mukherjee, S., Kumar, M.: Watching the embryo: Evolution of the microscope for the study of embryogenesis. BioEssays **43**(6) (Apr 2021)
18. Kar, A., et al.: Benchmarking of deep learning algorithms for 3d instance segmentation of confocal image datasets. PLOS Computational Biology **18**(4), e1009879 (Apr 2022)
19. Khosravi, P., et al.: Deep learning enables robust assessment and selection of human blastocysts after in vitro fertilization. npj Digital Medicine **2**(1) (Apr 2019)
20. Kragh, M.F., et al.: Automatic grading of human blastocysts from time-lapse imaging. Computers in Biology and Medicine **115**, 103494 (Dec 2019)
21. Kragh, M.F., et al.: Predicting embryo viability based on self-supervised alignment of time-lapse videos. IEEE Transactions on Medical Imaging **41**(2), 465–475 (Feb 2022)
22. Kromp, F., et al.: An annotated human blastocyst dataset to benchmark deep learning architectures for in vitro fertilization. Scientific Data **10**(1) (May 2023)
23. Liao, Q., et al.: Development of deep learning algorithms for predicting blastocyst formation and quality by time-lapse monitoring. Communications Biology **4** (12 2021)
24. Liao, Z., et al.: A clinical consensus-compliant deep learning approach to quantitatively evaluate human in vitro fertilization early embryonic development with optical microscope images. Artificial Intelligence in Medicine **149**, 102773 (2024)
25. Lu, W., et al.: Slidegraph: Whole slide image level graphs to predict her2 status in breast cancer. Medical Image Analysis **80**, 102486 (Aug 2022)
26. Montag, M., Toth, B., Strowitzki, T.: New approaches to embryo selection. Reproductive BioMedicine Online **27**(5), 539–546 (Nov 2013)
27. Nagaya, M., Ukita, N.: Embryo grading with unreliable labels due to chromosome abnormalities by regularized pu learning with ranking. IEEE transactions on medical imaging **41**, 320–331 (2022)
28. Patel, V., Chaurasia, V., Mahadeva, R., Patole, S.P.: Garl-net: Graph based adaptive regularized learning deep network for breast cancer classification. IEEE Access **11**, 9095–9112 (2023)

29. Pati, P., et al.: Hierarchical graph representations in digital pathology. Medical Image Analysis **75**, 102264 (Jan 2022)
30. Petersen, B.M., et al.: Development of a generally applicable morphokinetic algorithm capable of predicting the implantation potential of embryos transferred on day 3. Human Reproduction **31**(10), 2231–2244 (Sep 2016)
31. Salih, M., et al.: Embryo selection through artificial intelligence versus embryologists: a systematic review. Human Reproduction Open **2023**(3) (Jan 2023)
32. Song, A.H., et al.: Artificial intelligence for digital and computational pathology. Nature Reviews Bioengineering **1**(12), 930–949 (Oct 2023)
33. Studer, L., et al.: Classification of intestinal gland cell-graphs using graph neural networks. In: 2020 25th International Conference on Pattern Recognition (ICPR). IEEE (Jan 2021)
34. Su, Y., et al.: Hat-net: A hierarchical transformer graph neural network for grading of colorectal cancer histology images. In: The British Machine Vision Conference (BMVC) (2021)
35. Wang, Y., et al.: Cell graph neural networks enable the precise prediction of patient survival in gastric cancer. npj Precision Oncology **6**(1) (Jun 2022)
36. World Health Organization: Infertility prevalence estimates, 1990-2021 (2023)
37. Xu, K., et al.: How powerful are graph neural networks? In: International Conference on Learning Representations (2019)
38. Yamaguchi, T., et al.: 3D Image Reconstruction from Multi-focus Microscopic Images, p. 73–85. Springer International Publishing (2020)
39. Zhou, Y., et al.: Cgc-net: Cell graph convolutional network for grading of colorectal cancer histology images. In: 2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW). IEEE (Oct 2019)