




OPEN Merging synthetic and real embryo data for advanced AI predictions

Oriana Presacan^{1,9}, Alexandru Dorobanțiu^{2,9}, Vajira Thambawita³, Michael A. Riegler⁸, Mette H. Stensen⁵, Mario Iliceto⁵, Alexandru C. Aldea⁶ & Akriti Sharma^{4,7,9}

Accurate embryo morphology assessment is essential in assisted reproductive technology for selecting the most viable embryo. Artificial intelligence has the potential to enhance this process. However, the limited availability of embryo data presents challenges for training deep learning models. To address this, we trained two generative models using two datasets—one we created and made publicly available, and one existing public dataset—to generate synthetic embryo images at various cell stages, including 2-cell, 4-cell, 8-cell, morula, and blastocyst. These were combined with real images to train classification models for embryo cell stage prediction. Our results demonstrate that incorporating synthetic images alongside real data improved classification performance, with the model achieving 97% accuracy compared to 94.5% when trained solely on real data. This trend remained consistent when tested on an external Blastocyst dataset from a different clinic. Notably, even when trained exclusively on synthetic data and tested on real data, the model achieved a high accuracy of 92%. Furthermore, combining synthetic data from both generative models yielded better classification results than using data from a single generative model. Four embryologists evaluated the fidelity of the synthetic images through a Turing test, during which they annotated inaccuracies and offered feedback. The analysis showed the diffusion model outperformed the generative adversarial network, deceiving embryologists 66.6% versus 25.3% and achieving lower Fréchet inception distance scores.

Infertility affects millions as a disorder of the reproductive system in both men and women. According to the World Health Organization (WHO), approximately one in six people of reproductive age worldwide experience infertility¹. As part of treatment, Assisted Reproductive Technology (ART) procedures involve inseminating multiple eggs retrieved from the ovary and then selecting the most suitable embryo for implantation in the uterus. Although it is the leading solution for infertility, ART has an average success rate less than 30%². The success of ART depends directly on the quality of the embryo selected for transfer. However, visual evaluations by embryologists are subjective and prone to human error^{3,4}. Typically, ART data consists of images and videos. Artificial intelligence (AI) methods are commonly applied to tasks such as image and video classification and segmentation. Incorporating AI techniques can help embryologists add objectivity to the decision-making process and improve outcomes⁵. Furthermore, the number of ART treatments in Europe is increasing⁶, making the manual analysis of the growing volume of data increasingly time-consuming and resource-intensive⁷.

AI-based image and video analysis techniques have been applied in both research and industrial settings to perform embryo evaluations and annotations on time-lapse data^{8–13}. Indeed, the use of AI is gaining momentum as a promising technology to improve the efficiency of ART procedures¹⁴. According to the existing literature, convolutional neural network (CNN) and long short-term memory (LSTM) models are the most used AI approaches in ART. The application of CNNs for embryo assessments has been shown to outperform traditional methods¹⁵. However, advancements in AI now extend beyond CNNs and LSTMs¹⁶. To fully harness the advancements of AI in ART, embryo data is essential, with both its quantity (volume) and quality (variety and veracity) serving as key factors.

Embryo data encompasses images that capture the development of embryos at different cell stages and specific time points, highlighting the morphological changes throughout their developmental phases. The introduction of time-lapse imaging in ART laboratories has enabled comprehensive evaluation of embryo morphology by

¹Faculty of Electronics, Telecommunications, and Information Technology, National University of Science and Technology Politehnica Bucharest, 061071 Bucharest, Romania. ²Department of Computer Science and Electrical Engineering, Lucian Blaga University of Sibiu, 550024 Sibiu, Romania. ³Department of Holistic Systems, SimulaMet, 0170 Oslo, Norway. ⁴Department of Computer Science, Oslo Metropolitan University, 0167 Oslo, Norway. ⁵Volvat Spiren, 0176 Oslo, Norway. ⁶Faculty of Biotechnologies, University of Agronomic Sciences and Veterinary Medicine, 011464 Bucharest, Romania. ⁷Department of Validation Intelligence for Autonomous Software Systems, Simula Research Laboratory, 0164 Oslo, Norway. ⁸Cyber Security, Simula Research Laboratory, 0164 Oslo, Norway. ⁹Oriana Presacan, Alexandru Dorobanțiu and Akriti Sharma contributed equally to this work. ✉email: orianapresacan@gmail.com

providing detailed insights into embryo development¹⁷. Embryo development is characterized by cell divisions, with each division representing a distinct stage, ranging from the 2-cell stage to later phases that exhibit distinct morphological features such as the morula and blastocyst¹⁸. A morula is a stage in embryo development, typically observed around day four, and is characterized by a compact mass of 16–32 cells that are tightly packed and undifferentiated into layers. The increased cell-to-cell adhesion in morula sets the foundation for subsequent development into the blastocyst stage. The blastocyst is defined by a fluid-filled cavity (the blastocoel), an inner cell mass that later develops into the embryo, and an outer trophectoderm contributing to the placenta. Morula cells secrete fluid to form a small cavity, while the outer cells flatten and form tight junctions that separate the embryo's interior from its exterior. Membrane channels then increase the salt concentration for drawing in water osmotically to expand the cavity. Concurrently, ongoing cell division enlarges the embryo with thinning of the zona pellucida until the blastocyst hatches, forming the inner cell mass and an expanding fluid-filled cavity¹⁹. Embryologists assess embryo quality by examining the morphological characteristics and duration of each developmental stage, ultimately selecting an embryo for transfer at either the cleavage or blastocyst stage²⁰. A cleavage-stage transfer involves transferring embryos on day two or three of development. Embryos on day two typically consist of two to six blastomeres, while on day three, they generally contain six to eight or more blastomeres. Additionally, embryologists consider the synchrony of cell divisions, which refers to the sequence of simultaneous divisions of all cells in an embryo. Perfect synchrony results in the embryo progressing through the 2-cell, 4-cell, and 8-cell stages. Since synchronous development is regarded as a positive indicator of embryo quality²¹, this study focused exclusively on the 2-cell, 4-cell, and 8-cell stages for cleavage-stage transfers and the morula and blastocyst stages for blastocyst-stage transfers.

The primary challenge limiting the availability of embryo data arises from privacy and ethical concerns regarding data sharing. Nearly all existing studies in ART have not made their embryo data publicly available²². Furthermore, research frameworks often rely on customized datasets, resulting in subjective interpretations of ground truth that are specific to each framework²³. Study²⁴ introduced STORK, an AI-based method for automated blastocyst quality assessment; however, access to their dataset was limited. In contrast, in study²⁵, the authors publicly released their dataset of blastocyst and non-blastocyst embryos, captured using various imaging systems, from high-quality clinical time-lapse systems to lower-quality 3D-printed smartphone-based systems. Using an adaptive neural network, they demonstrated that a model trained on one dataset maintained high accuracy when tested on images from different datasets. Additionally, study²⁶ advanced AI for embryo selection by releasing a public dataset of static blastocyst images, annotated with Gardner criteria and clinical parameters, for deep learning model training. While these studies focused on the blastocyst stage, another study²⁷ provided a fully annotated public dataset of time-lapse embryo development videos, covering 16 developmental phases. Table 1 summarizes these publicly available embryo datasets and their characteristics.

While publicly available datasets are important for benchmarking and reproducibility in ART, their limited size can restrict the development of robust deep learning models. Various approaches, such as federated learning²⁸ and synthetic data generation²⁹, have been developed to address the challenges posed by limited data volume. Indeed, the use of synthetic data in the medical domain is not new; it has successfully provided solutions to overcome the legal, privacy, and security barriers that restrict access to medical data. Deep convolutional generative adversarial networks (GANs)³⁰ have been used to create synthetic images for skin lesion classification³¹ and to enhance the training of CNNs for improved pneumonia detection in chest X-ray images³². State-of-the-art generative techniques such as the Style-based GAN (StyleGAN)³³ and the latent diffusion model (LDM)³⁴ are extensively explored for synthetic data generation in the medical field. Research has shown that StyleGAN can generate synthetic brain tumor data³⁵ and 3D brain MRI images for enhancing diagnostic tools and medical research³⁶. Additionally, LDMs effectively generated synthetic brain images with conditioning on various brain attributes³⁷ and augmented biomedical datasets for microscopy³⁸.

We identified a few research studies that utilized generative AI to create synthetic images of human embryos. The first study, Human Embryo Image Generator Based on Generative Adversarial Networks (HEMIGEN)³⁹, used a GAN to generate synthetic images of human embryos at the 1-cell, 2-cell, and 4-cell developmental stages. The authors included a qualitative assessment in which embryologists were asked to identify the number of cells in the synthetic images. However, it did not evaluate the visual quality of the generated images. Additionally, neither the code nor the dataset used in the study was made publicly available, limiting the reproducibility and further exploration of their work. A more recent study⁴⁰ employed the StyleGAN model to produce synthetic images of human embryos at the blastocyst stage, achieving a Fréchet inception distance (FID) score of 15.2 by fine-tuning a pre-trained model. This study also included a Turing test, where the synthetic images fooled embryologists 44.3% of the time. Another study⁴¹ employed the LDM to generate embryo images at both the

Title	Size	Description
Adaptive adversarial neural networks for the analysis of lossy and domain-shifted datasets of medical images ²⁵	3,063 images	Annotated embryo images, classified into blastocyst and non-blastocyst categories, with quality levels labeled on a scale from 4 to 1
A time-lapse embryo dataset for morphokinetic parameter prediction ²⁷	704 videos	Annotated embryo images capturing 16 key developmental events, from polar body appearance to blastocyst hatching, with frames labeled by their post-fertilization timing
An annotated human blastocyst dataset to benchmark deep learning architectures for in vitro fertilization ²⁶	2,344 images	Annotated blastocyst images with expansion grade, inner cell mass, trophectoderm quality, expert agreement scores, and clinical data including age, oocyte metrics, pregnancy outcomes, etc
Ours	5,500 images	Annotated embryo images depicting the 2-cell, 4-cell, 8-cell, morula, and blastocyst developmental stages, supplemented by synthetic images generated using advanced generative models

Table 1. Overview of other existing public annotated embryo datasets.

blastocyst stage (day five) and the cleavage stage (day three). To evaluate the model, Turing tests compared real and synthetic images, resulting in accuracies of 0.57 for day three and 0.59 for day five embryos. For day three, the evaluation focused on cell number, evenness, and fragmentation percentage, while for day five, the inner cell mass, trophectoderm, and overall blastocyst morphology were assessed. The study relied solely on expert evaluation and did not use any quantitative assessment methods. Although these studies used synthetic datasets to address data scarcity, they covered limited embryo stages and did not assess whether the synthetic images actually improved classification performance, thereby lacking a comprehensive end-to-end solution. Classifying the morphological characteristics of embryo cell stages is a fundamental task that AI algorithms must be trained for, especially when applied to complex tasks such as assessing embryo quality or predicting pregnancy. Several studies have focused on single-frame embryo morphology, either targeting specific features like blastocyst formation⁴² or identifying multiple morphological events from the 1-cell to 4-cell stages⁴³. Recently, a study introduced a deep learning model capable of detecting 11 key embryo morphokinetic events from videos instead of individual images, enhancing embryo stage classification across various time-lapse platforms⁴⁴.

To provide a comprehensive end-to-end solution, our current work addresses these gaps by generating synthetic images across a broader range of embryo cell stages, including the 2-cell, 4-cell, 8-cell, morula, and blastocyst stages. These images were then used to train multiple classification models, tackling the critical aspect of data scarcity. Recent deep learning studies have demonstrated that augmenting real datasets with synthetic images can improve classification performance^{45–47}. To further enhance the diversity of the synthetic dataset, we incorporated data generated from two distinct models: a GAN and a diffusion model. Given that these models employ different generation processes, we hypothesized that their images would introduce unique and potentially complementary features. This increased diversity could help the classification models generalize better across various image features, leading to better predictions. This aligns with findings from⁴⁵, which demonstrated that combining synthetic data from different diffusion models led to improved classification performance, and⁴⁶, which highlighted that merging synthetic data from diverse generation techniques enhanced model robustness and generalization. Furthermore, to evaluate the quality of the generated images, we conducted a qualitative analysis using a custom-developed web application. Embryologists utilized the platform to classify images as 'real' or 'fake', provide comments, and highlight areas of low fidelity in the synthetic images. This approach aimed to validate that the generated images are indeed realistic, represent key morphological characteristics of embryo cell stages, and can be reliably used for training AI models, thereby aligning with the second key aspect of embryo data: quality.

Below, we summarize our main contributions as follows:

- Published a new open access annotated real embryo dataset (<https://zenodo.org/records/14253170>) to facilitate and advance research efforts in the field of ART.
- Employed two types of generative models, specifically the LDM and the StyleGAN, to generate synthetic images of embryos.
- Conducted a comprehensive qualitative assessment to validate the fidelity of the generated images, with evaluations performed by expert embryologists to ensure data quality.
- Improved the accuracy of embryo stage classification by integrating synthetic images into the training process of deep learning models, achieving higher accuracy even on an external dataset.
- Published a synthetic embryo dataset (https://huggingface.co/datasets/deepsynthbody/synembryo_latentdiffusion, https://huggingface.co/datasets/deepsynthbody/synembryo_stylegan) and the corresponding pre-trained generative models (https://huggingface.co/deepsynthbody/synembryo_ldm, https://huggingface.co/deepsynthbody/synembryo_stylegan) to augment existing data and support deep learning applications, thereby addressing data scarcity in ART.

Materials and methods

Time-lapse incubation

Embryologists evaluate embryo quality by assessing specific morphological parameters related to cell division and embryo stages⁴⁸. In this context, the morphological development present in embryo images captured during incubation provides a reference for embryo assessment⁴⁹. Additionally, stable culture conditions particularly concerning pH and temperature must be maintained during the incubation⁵⁰. The introduction of time-lapse incubator (TLI) systems in ART laboratories has ensured such an optimal environment where embryologists can assess the morphokinetics without removing embryos from the incubator or exposing them to unstable culture conditions⁵¹.

A TLI system consists of three components: a microscope, imaging software, and an incubator, which allow for continuous imaging and a non-invasive evaluation of the embryo's quality⁵². The system illuminates an embryo with a light source and magnifies embryo cells using an inverted microscope, capturing images with a digital camera at regular intervals across different focal planes. The software then displays a timed sequence of embryo development images, allowing embryologists to annotate morphology. The incubator also includes an 'EmbryoSlide', a dish with multiple wells for culturing individual embryos separately.

Data

The images used for training and evaluating the generative models originated from two sources. The first source was Volvat Spiren, a fertility clinic in Oslo, Norway. These images were extracted from embryo video frames recorded using a TLI system known as the Embryoscope™ (Vitrolife, Sweden), equipped with an EmbryoSlide containing 12 wells. The TLI system features a camera positioned beneath a 635 nm LED light source, which passes through Hoffman's contrast modulation optics. For each embryo, the system captured 8-bit images every 7–20 min across 3 to 5 focal planes, based on the settings of the older EmbryoScope system and EmbryoViewer

software. The frames were extracted and labeled according to the cell stage they depicted, as identified by embryologists' annotations. These annotations provide timings in hours post-fertilization, indicating the temporal position of each cell stage present in the video frames. Based on these timings, we selected frames corresponding to the specific cell stages of morula and blastocyst. The initial dataset consisted of 4,172 time-lapse videos capturing embryo morphokinetics. For our study, we excluded embryos with fragmentation rates above 15%. Fragmentation rates were available for 2,254 videos, of which 1,644 had fragmentation rates below 15%. Consequently, we used these 1,644 videos and included only frames from the central plane. Although the embryo is not always centered in the image and may be partially occluded due to its position in the EmbryoSlide well, the images in the dataset still accurately represent the embryo's developmental stage.

The second data source is the publicly available dataset 'Human Embryo Time-Lapse Video'²⁷, also generated using the Embryoscope™ system with similar imaging parameters: a 635 nm LED light source, Hoffman's contrast optics, captured every 10–20 min across 7 focal planes. We selected frames corresponding to the central plane only. In addition to the images, the dataset included class labels associating each image with the embryo developmental stages (cell stages, morula, and blastocyst). Since there was no information on fragmentation rates, we visually inspected the frames and selected those with low fragmentation for the 8-cell stage, as fragmentation cannot be assessed at the morula and blastocyst stages. The images also had issues with brightness changes and white spots caused by problems during image acquisition²⁷.

Our final dataset consists of 5,500 embryo images, representing five developmental stages, 2-cell, 4-cell, 8-cell, morula, and blastocyst, with 1,100 images per stage. Due to variations in the number of images from each data source for different stages, Table 2 provides a detailed breakdown of the image distribution from each source. For training the generative models (StyleGAN and LDM), 1,000 images per stage were used. The remaining 100 images per stage were reserved solely for testing the downstream classification models (VGG, ResNet, and ViT) and were not used in any way during the training of the generative models. The train/test split was performed at the sequence level, selecting a single representative frame for each stage of embryo development to ensure that frames from the same sequence did not appear in both sets. To expand the dataset, we augmented the real images with 5,000 synthetic images per stage generated using two different generative models, enhancing the training data for classification.

To ensure the robustness and generalizability of our results, we evaluated our models on a publicly available subset of an external Blastocyst dataset²⁴. This subset includes 98 labeled images, which were captured using the EmbryoScope time-lapse system (Vitrolife, Sweden). Images were acquired using a single red LED (635 nm) every 20 min and were annotated by embryologists with quality labels: good-quality, fair-quality, and poor-quality.

Ethical considerations

In this study, we used human embryo videos collected at the Volvat Spiren fertility clinic in Oslo between 2013 and 2019. The Regional Committee for Medical and Health Research Ethics–South East Norway (REK) approved the data collection and use of these videos for research purposes. All methods were carried out in accordance with relevant guidelines and regulations, including the Declaration of Helsinki and the General Data Protection Regulation (GDPR). All patients gave informed consent for the intended research use of their data. However, the data were fully anonymized. Specifically, all patient-identifying information was removed prior to data processing, ensuring compliance with REK regulations and GDPR.

System workflow

Figure 1 illustrates the workflow of our system. For each embryo class, 2-cell, 4-cell, 8-cell, morula, blastocyst, 1,000 real images were used to train two types of generative models: GAN and diffusion model. A separate model was trained for each stage, resulting in a total of ten models. The best-performing model checkpoints, based on the lowest FID scores, were selected for image generation. From each selected model, 5,000 synthetic images were generated. These synthetic images, combined with the original real images, were then used to train three classification models: VGG, ResNet, and ViT. Model performance was evaluated using a separate set of 100 real images per stage and an external blastocyst dataset to ensure the generalizability and robustness of our results. Furthermore, four embryologists reviewed the generated images via a web application, identifying real versus synthetic samples and offering insights to confirm the visual quality of the data.

Generative models

Two generative models were selected for synthesizing images: a diffusion model, namely the LDM³⁴ and a GAN, specifically StyleGAN⁵³. These models were chosen not only for their well-established reputations in generative

Embryo cell stage	Number of images from datasets	
	Volvat Spiren, Oslo	Human embryo time-lapse video ²⁷
2-cells	1100	0
4-cells	1100	0
8-cells	850	250
morula	660	440
blastocyst	600	500

Table 2. Distribution of embryo images across different developmental stages and data sources.

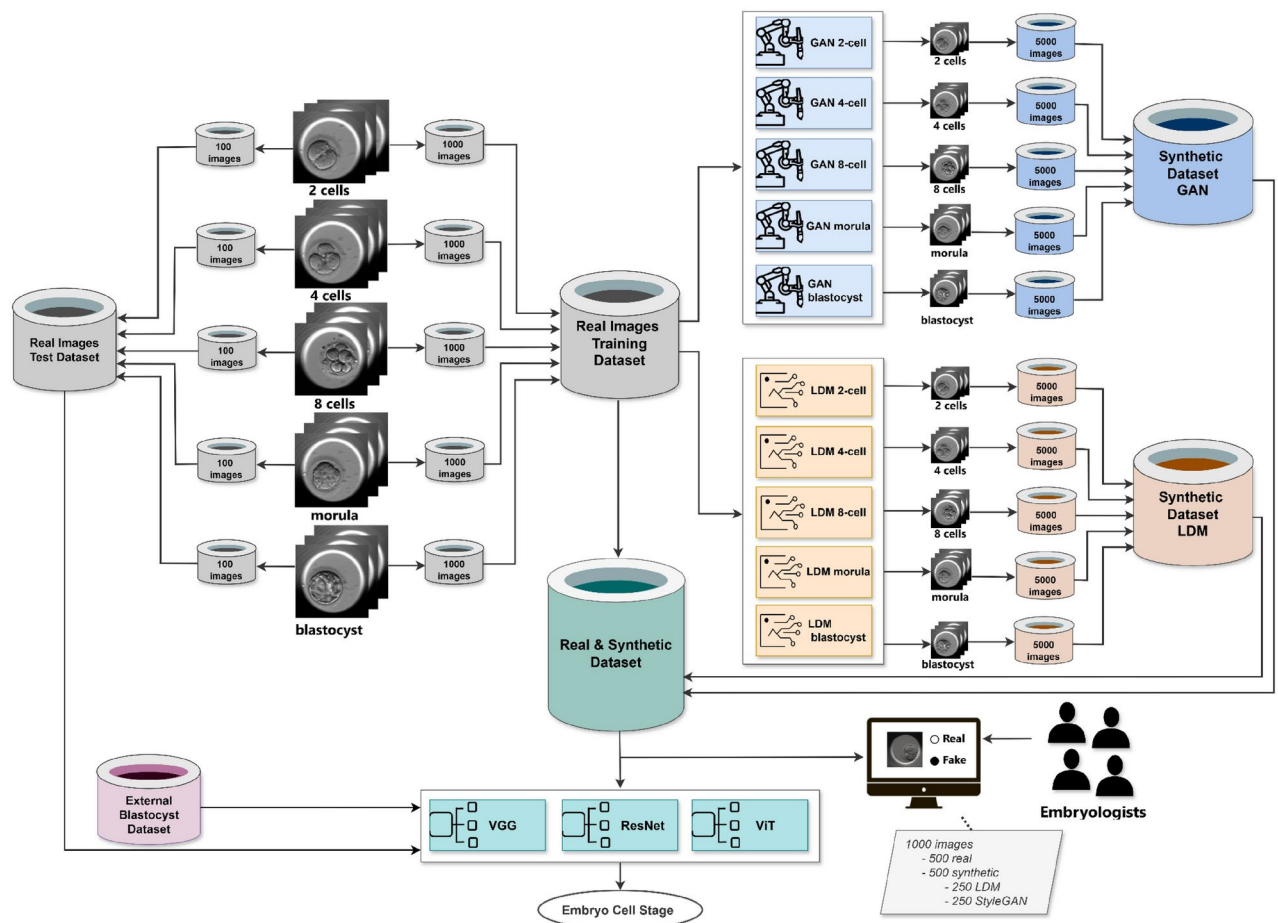


Fig. 1. Pipeline of our proposed system, encompassing the training of generative models, the generation of synthetic data, the training of classification models, and the qualitative assessment conducted by embryologists.

modeling but also for their distinct ways of generating images. The LDM transforms noise into high-quality images via a sequence of iterative steps, based on learned data distributions. On the other hand, the GAN uses an adversarial training methodology where two neural networks—the generator and discriminator—compete with each other: the generator creates images, while the discriminator evaluates them, driving the system towards producing precise results.

The LDM starts by reducing the dimensions of the input images using an autoencoder. This step improves computational efficiency while preserving the high quality of the images³⁴. Subsequently, it proceeds to the diffusion phase, systematically introducing noise into the latent space. This converts the data into a noisy state through a predefined series of steps. A U-Net⁵⁴ architecture is then used to perform the reverse process, which involves progressively removing the added noise at each step and transforming the noisy latent representation back into a clear image. During training, the model is exposed to examples of noisy data at various stages of the diffusion process, learning to predict earlier, less noisy states. Once trained, during the sampling phase, the model will be given pure noise from a normal distribution and will progressively transform this noise into a coherent image³⁴.

StyleGAN operates based on the foundational principle of GAN, utilizing two networks: a generator and a discriminator. The generator's role is to produce realistic images given random noise. It starts this process by transforming the noise into a latent space representation via a mapping network. This representation is then adjusted to manipulate the generated image's style at various resolutions, employing a sequence of convolutional layers to enhance the image's resolution. The discriminator evaluates whether an image is real (from the dataset) or produced by the generator. Both the generator and discriminator undergo simultaneous training in a competitive manner, where the generator aims to fool the discriminator while the discriminator strives to become better at distinguishing real images from fakes⁵³.

The generative models were trained on 1,000 real images per class, each with an original resolution of 256×256 pixels. Training of the StyleGAN model was conducted using an NVIDIA RTX 4090, whereas the LDM utilized an NVIDIA Tesla V100 GPU. The LDM underwent unconditional training for 1000 epochs with a batch size of 16 and a learning rate of 0.000002, as determined through repository code review and empirical testing. A linear noise scheduler was implemented for the diffusion process across 1000 steps. The pre-trained VQ-f4 autoencoder³⁴ was used to encode images into a $64 \times 64 \times 3$ latent space. For StyleGAN, we employed

transfer learning from a pre-trained model based on the Flickr-Faces-HQ Dataset (FFHQ)⁵⁵ after initial attempts at training from scratch yielded unsatisfactory results due to our limited dataset size (1,000 images per stage). The small dataset was insufficient for the StyleGAN model to learn the complex features and patterns needed for high-quality image generation when training from scratch. Data augmentation techniques were also used to enhance the training process. The model underwent training for 1500 steps with a batch size of 16, a learning rate of 0.002 for both the generator and discriminator, and an R1 regularization weight of 10. Detailed information about the configurations for both models used in this project is available in the project's GitHub repository, as well as links to the checkpoints of the models used to generate the final images.

To evaluate model performance throughout training, we calculated the FID every 50 epochs by generating 1,000 synthetic images and assessing their similarity to the 1,000 real training images. The training set was specifically used for FID calculation to track how well the model was learning the distribution of the training data. FID quantifies the difference in feature distributions between real and synthetic images, with lower scores indicating higher similarity. By monitoring FID scores, we can assess the authenticity of the synthetic images, ensuring they capture essential features from the real dataset. Ultimately, the model with the lowest FID score was selected for data generation. During sampling, random noise from a normal distribution was fed into the models, which converted them into images based on the learned patterns and features at various scales. A total of 5,000 synthetic images were generated for each class with each model and subsequently used for classification tasks and qualitative analysis.

Classification models

Three distinct models were employed for multi-class classification among five cell stages, including two CNNs—VGG and ResNet—and a transformer, ViT. These models were chosen for their established effectiveness in image classification tasks, with the goal of assessing whether a transformer-based architecture can outperform traditional CNN-based models. After experimenting with multiple model sizes, we identified VGG16, ResNet50, and ViT16 as the most effective architectures for our task. These models and their pre-trained versions were sourced from the *torchvision* library, specifically *vit_b_16*, *vgg16_bn*, *resnet50*. The training approach for all three models was uniform, utilizing a batch size of 32, an Adam optimizer with an initial learning rate of 0.0001, a learning rate scheduler, and cross-entropy loss. Training was halted if there was no reduction in validation loss for 30 epochs, with the model showing the lowest validation loss being kept. All three models were trained on an NVIDIA RTX 4090.

For image processing, normalization was conducted using the mean and standard deviation of the dataset, and the images were resized to 224×224 pixels to accommodate the requirements of the ViT. We experimented with various combinations of images, including solely synthetic images produced by StyleGAN and the LDM in amounts ranging from 250 to 5,000, as well as combinations of real and synthetic images. We used 1,000 real images to match the dataset size used to train the generative models, ensuring a consistent training process, and then progressively incorporated varying numbers of synthetic images to determine if this could enhance the outcomes. The models were then evaluated on a separate set of 100 real images reserved for testing.

Results

Data generation results

Table 3 presents the lowest FID scores achieved by StyleGAN and LDM models, which were trained for each embryo stage. The FID was calculated periodically throughout the training process, and the best checkpoints were selected based on the FID score. Overall, the LDM model demonstrates superior performance, as indicated by its lower FID scores. Figure 2 illustrates real embryo images alongside synthetic images produced by StyleGAN and LDM for each of the five developmental cell stages: 2-cell, 4-cell, 8-cell, morula, and blastocyst.

Classification results

Table 4 presents the classification results of the three models—VGG, ViT, and ResNet—fine-tuned on both real and synthetic data and evaluated on real data using various metrics, including accuracy, F1 score, precision, recall, and the Matthews correlation coefficient (MCC). For clarity, we report only the best result for each model along with the corresponding combination of training data used, comparing two scenarios: training with synthetic data alone versus training with a combination of real and synthetic data. The VGG model demonstrated the best overall performance, and therefore, our subsequent analysis focuses exclusively on it. Comprehensive results for all three models are available in the "Supplementary Material".

To ensure the robustness and reliability of our results, we trained the VGG model using five different random seeds for initialization. We then computed the mean and standard deviation of the performance metrics to account for variability. Complete results for VGG models trained from scratch versus fine-tuned are in Tables S3 and S4 of the "Supplementary Material". Figure 3 illustrates the mean performance of the VGG16 model when trained solely on synthetic data. We employed different combinations of synthetic images generated by

Model	2-cell	4-cell	8-cell	Blastocyst	Morula
LDM	24	28	31	10	14
StyleGAN	24	24	34	41	36

Table 3. FID scores for the synthetic data generated by the best generative models trained for the different classes.




Fig. 2. Comparison of real versus synthetic images generated with StyleGAN and LDM models for each of the five embryo classes: 2-cell, 4-cell, 8-cell, morula, and blastocyst.

Model	Real data	StyleGAN data	LDM data	Accuracy (%)	F1 Score	Precision	Recall	MCC
VGG	0	5000	5000	93	0.93	0.93	0.95	0.92
	1000	5000	5000	97	0.97	0.97	0.97	0.96
ResNet	0	1000	1000	89	0.89	0.90	0.89	0.87
	1000	4000	4000	97	0.97	0.97	0.97	0.96
ViT	0	2000	2000	91	0.91	0.91	0.91	0.89
	1000	4000	4000	93	0.93	0.94	0.93	0.92

Table 4. Best performance results for each of the three fine-tuned classification models—VGG, ResNet, and ViT—under two scenarios: training with only synthetic data and training with a combination of synthetic and real data. The table also shows the corresponding combinations of synthetic and real data used for training.

StyleGAN and LDM, varying these combinations to assess their impact on classification results. For results obtained from training with only LDM or only StyleGAN images, please refer to Tables S7 and S8 in the "Supplementary Material". Notably, training with 500 LDM images achieved an accuracy of 70%, compared to 47% when using 500 StyleGAN images alone. This suggests that the LDM model might have generated higher-quality images. However, combining 250 StyleGAN images with 250 LDM images—maintaining a total of 500 images—increased the accuracy to 74%. This improvement indicates that integrating data from different generative models may enhance classification performance within the tested sample size range (up to 10,000 total synthetic images). While our dataset size limited testing at larger scales, this finding aligns with previous work on synthetic data combinations. Further testing with larger datasets are needed to confirm if this trend continues at greater scales but is out of scope of this work. The highest accuracy (93.6%) was achieved by combining 4,000 StyleGAN images with 4,000 LDM images. Accuracy slightly decreased when using 5,000 GAN and 5,000 LDM images, possibly due to diminishing returns, where additional synthetic data introduces noise, subtle artifacts, or redundancy rather than meaningful variation. Figure 4 presents a similar results chart that includes training with 1,000 real images in addition to the synthetic ones. When training exclusively on real images, an accuracy of 92.7% was obtained. In contrast, training with 1,000 real images combined with 4,000 LDM images and 4,000 StyleGAN images increased the accuracy to 95.5%. These results demonstrate that




Fig. 3. Classification accuracy trends on test data (100 real images) for the VGG model, trained with various combinations of synthetic images generated by LDM and StyleGAN models.

incorporating both synthetic and real data, especially from multiple generative models, enhances classification performance. Figure 5 shows the accuracy gain from training from scratch versus fine-tuning pre-trained models for the same data combinations from Fig. 4. Overall, the pre-trained VGG achieved the highest accuracy of 97% when trained with 1,000 real images, plus 5000 synthetic images each from StyleGAN and LDM, compared to 94.5% when trained solely on 1000 real images. This shows that synthetic images can effectively supplement real images to enhance model performance, which is valuable when real data is scarce or hard to obtain. However, the performance differences between different synthetic data combinations (e.g., 1000-2000-2000 and 1000-5000-5000) were not statistically significant, leaving the optimal amount of synthetic data an open question. Although the improvement is modest, it remains significant because even a slight enhancement can increase the model's reliability and robustness, positively impacting user trust, qualities that are essential in the medical field.

Table 5 presents the mean accuracy, standard deviation, and confidence intervals for multiple model configurations evaluated on an external Blastocyst dataset. As expected, performance on the external dataset was lower than on our own dataset. However, we observed a substantial and statistically significant improvement when comparing the baseline configuration (1000-0-0) to those that included synthetic data. The best-performing synthetic data configuration (1000-4000-4000) achieved an accuracy of 84.69% on the external dataset, marking a notable 15.81% point increase over the baseline's 68.88%. This result highlights that our synthetic data approach not only enhances classification performance but also significantly improves model generalization to entirely unseen external datasets, an especially challenging task in medical image analysis. Moreover, all synthetic data configurations consistently outperformed the baseline, with the 1000-5000-5000 configuration demonstrating the most stable performance, as indicated by its narrower confidence interval ($\pm 2.84\%$) compared to the baseline's ($\pm 5.86\%$).

Qualitative analysis

We developed a web application to present the synthetic images to four human experts (embryologists), employing a Turing test framework. Each image was presented individually, and the embryologists had to determine whether it was real or fake. The goal was to evaluate the fidelity of the synthetic images and determine if they could mislead the experts. A total of 1,000 images were presented, comprising 500 real images (100 images per cell stage: 2-cell, 4-cell, 8-cell, morula, blastocyst) and 500 synthetic images. The synthetic images were evenly split between those generated by GAN (250 images) and LDM (250 images), with 50 images per cell stage. The application also allowed embryologists to mark areas of the images they considered unrealistic and included a text box for additional comments. A figure illustrating the application interface is provided in the "Supplementary Material".

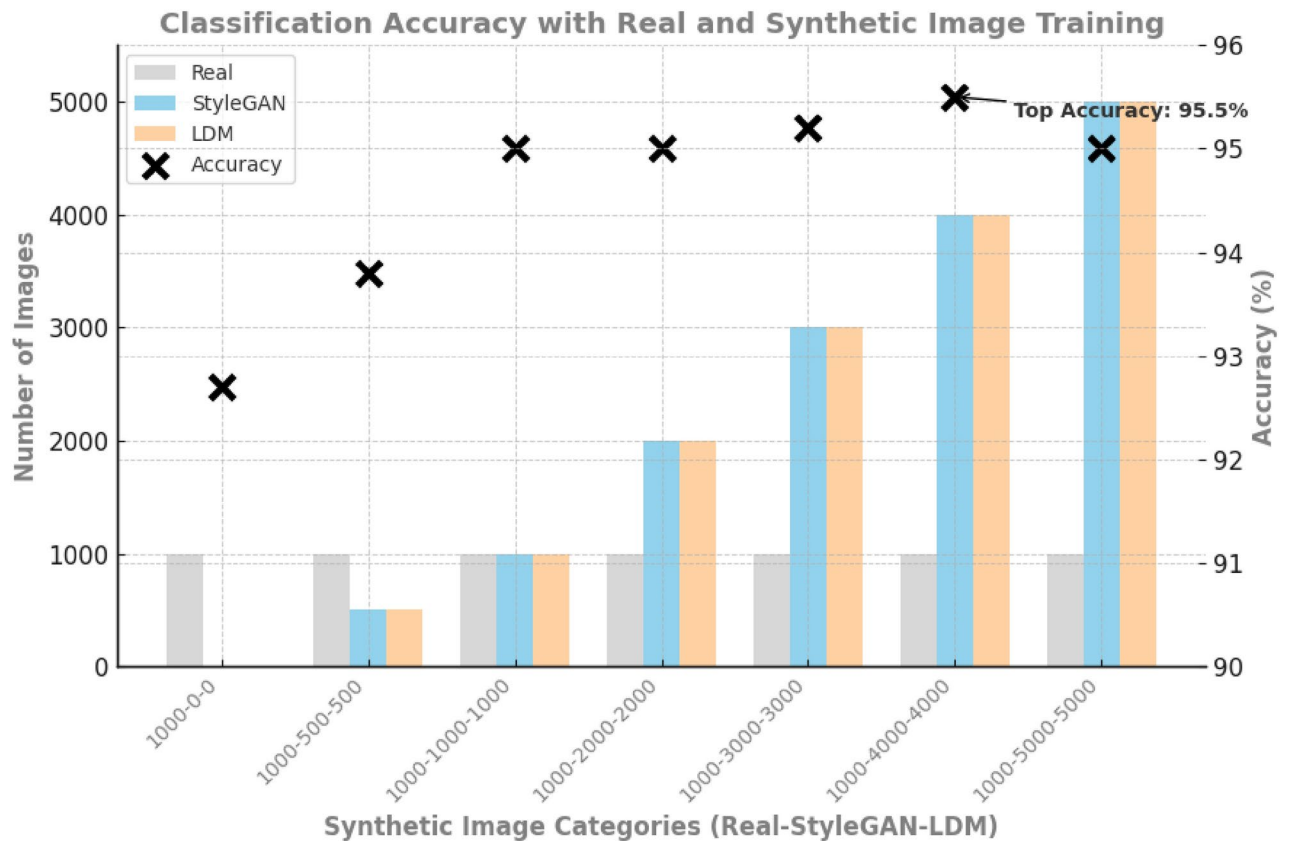


Fig. 4. Classification accuracy trends on test data (100 real images) trained with various combinations of real and synthetic data generated by LDM and StyleGAN models.

The averaged results of the Turing test conducted by the four embryologists are shown in Fig. 6. We divided the outcomes into two categories: real images and generated images. The first chart shows the proportion of real images accurately identified as real, at 81.3%, versus those mistakenly classified as generated. This result is noteworthy, as the embryologists had initially anticipated a balanced, 50–50 distribution, citing that they often felt as though they were merely guessing. The second chart offers a similar comparison for synthetic images, revealing that the embryologists correctly identified only 54.5% of these images. To assess the differences between the two generative models, we further divided the synthetic images into StyleGAN and LDM categories. This clearly shows that images generated by LDM appeared more realistic, as they were harder to classify correctly as “real” or “fake.” Specifically, the embryologists correctly identified only 34.4% of LDM-generated images, compared to 74.7% of StyleGAN images. For a more detailed overview, Table 6 shows the average accuracy with which the embryologists correctly identified real images as real and LDM-generated and StyleGAN-generated images as fake, across the five cell stages. However, no single category stands out as consistently performing best. Overall, we can see that while embryologists are generally good at identifying real embryos, their accuracy in detecting synthetic images varies significantly depending on the model used to generate these images.

Figure 7 showcases examples of annotations made by embryologists on embryo images they perceived as synthetic, marking specific image regions or morphological features they identified as “fake.” Several annotations focused around the ZP region, which is the outer layer surrounding the embryo. For instance, in Fig. 7a, despite the image depicting a real embryo, the embryologists could not clearly discern the boundary of the ZP, resulting in uncertainty about the image’s authenticity. Figure 7b,c display blastocyst-stage embryos. Figure 7b shows a real image that was mistakenly judged as fake because the trophectoderm—the outer cell layer that forms the placenta—was poorly visible due to darkness in the image, which misled the embryologists. In contrast, Fig. 7c is a synthetic image correctly identified as fake, as it lacked the end of the trophectoderm, a clear indicator for embryologists. Figure 7d,e show sperm sitting on ZP with abnormal morphological structure leading embryologists to identify images as synthetic. To investigate the root cause, further training of our models on embryo images with unfused sperm on the ZP is recommended followed by an evaluation of the generated synthetic data.

The participants also provided several comments, with some focusing on the presence of unusual dots within the embryos. All images containing these dots were indeed synthetic, generated by StyleGAN. One such image, Fig. 7f, was specifically highlighted due to a comment on these dots. The embryologist who identified the spots mentioned having previous experience with generative models, albeit not specifically with embryos, and recognized similar artifacts in other AI-generated images. This familiarity allowed him to recognize the dots as a sign of synthetic generation. In contrast, embryologists who were less familiar with synthetic images did not




Fig. 5. Accuracy differences between training the VGG model from scratch versus using a pre-trained model, based on the same data combinations as in Fig. 4.

Real	StyleGAN	LDM	Accuracy (95% CI) External dataset (Blastocyst)
1000	0	0	68.88% \pm 5.86 (63.137, 74.623)
1000	500	500	77.55% \pm 6.687 (71.004, 84.096)
1000	1000	1000	74.98% \pm 6.949 (68.179, 81.781)
1000	2000	2000	73.47% \pm 6.432 (67.169, 79.771)
1000	3000	3000	76.27% \pm 7.622 (68.803, 83.737)
1000	4000	4000	84.69% \pm 6.828 (78.007, 91.373)
1000	5000	5000	78.08% \pm 2.844 (75.297, 80.863)

Table 5. Classification results (mean accuracy \pm standard deviation) and confidence intervals (CI) for models trained on various real-synthetic data combinations and evaluated on an external Blastocyst dataset.

observe these abnormalities. Another remark was that some cells in the StyleGAN-generated images appeared flat, as seen in Fig. 7f. This contributed to the perception that these images were synthetic as real embryo cells are rounder, not flat.

Discussion

The primary goal of this study was to assess whether the use of synthetic data can address embryo data scarcity and improve classification accuracy, with a particular emphasis on ensuring data quality. Our findings confirm that this approach is indeed effective: training classifiers solely on synthetic data resulted in good accuracy, and incorporating synthetic images alongside real ones further slightly improved performance. Notably, this performance boost persisted even when tested on an external dataset, demonstrating the potential of synthetic data for enhancing model generalization.

To understand how synthetic images improved model performance, we examined samples misclassified by the real-data-only model but correctly classified when using both real and synthetic data. Figure 8 shows such embryo images, with the true class and the real-data-only misclassification noted below each image. In the first image, misclassification may stem from an indistinct membrane, resembling a morulated embryo. The second image could reflect fragmentation being mistaken for individual cells. In the third, a central structure within




Fig. 6. Average Turing test results conducted by embryologists, evaluating their accuracy in identifying real and synthetic images. The “Real Images” pie chart displays the proportion of images correctly identified as real versus incorrectly labeled as fake. The “Synthetic Images” chart (covering both StyleGAN and LDM-generated images) illustrates the proportion of fake images accurately recognized as fake versus those mistakenly considered real. The final two charts break down these results by StyleGAN and LDM-generated images, respectively.

	2 cell (%)	4 cell (%)	8 cell (%)	Morula(%)	Blastocyst(%)
Real Images	80	84	89	79	75
StyleGAN Images	76	77	73	71	86
LDM Images	37	40	29	38	27

Table 6. The average accuracy with which embryologists identified synthetic images as fake and real images as real, across the five cell classes.

the morula may have been misidentified as a membrane. The fourth image shows an embryo in early blastocyst formation, indicated by a small fluid cavity, rather than a fully developed blastocyst.

We trained separate generative models for each stage, rather than using a single conditional model, to produce higher-quality synthetic images. This decision was guided by empirical results, where separate LDM models outperformed the conditional approach in both FID scores and visual quality. However, this comparison was limited to LDMs, not GANs, which we acknowledge as a limitation. To evaluate the impact of these synthetic images on classification, we tested three models—VGG, ResNet, and ViT—observing distinct performance patterns. VGG achieved the best overall results, likely due to its simple architecture and use of small 3×3 kernels, which effectively capture local features in embryo images. In contrast, ResNet and ViT performed less favorably, particularly ViT when trained from scratch, as transformers typically require larger datasets to perform well. The more complex architectures of ResNet and ViT, which excel at capturing global dependencies, may better demonstrate their strengths with larger training datasets.

Furthermore, the synthetic images were convincing enough to lead embryologists to believe they were real, highlighting the high quality of the generated data. We also found that images generated by the LDM exhibit




Fig. 7. Annotations by embryologists highlighting perceived inaccuracies in images identified as synthetic. The first two images (a) and (b) depict real images, the following one (c) is generated using the LDM model, and the last three images (d), (e), and (f) are produced with the StyleGAN.




Fig. 8. Examples of embryos misclassified by the real-data-only model but correctly classified with real and synthetic data. True (black) and misclassified (red) classes are shown below each image.

higher fidelity compared to those produced by the GAN. This superiority was supported by our quantitative results, including lower FID scores and higher classification accuracies on the test data, as well as qualitative results from expert evaluations, where the LDM images fooled embryologists more frequently. Despite these positive outcomes, certain challenges were identified during the evaluation process.

One such challenge noted by embryologists was that some images appeared too dark, making it difficult to evaluate the embryo morphology and rendering them unreliable for accurate assessment. Interestingly, this issue was not exclusive to the synthetic images: twelve comments referred to the real images, while eight concerned

the LDM-generated images. This occasional darkness in real images may stem from a bubble obstructing the camera's direct view, floating between the camera and the EmbryoSlide well containing the embryo. This suggests that the LDM model learned and reproduced this characteristic from the real dataset. A possible solution to address this issue would be to preprocess the real images before training the generative models. There could be a quality control step, where only images meeting specific brightness and contrast thresholds are included in the training dataset. Very dark images could either be filtered out or adjusted using brightness correction techniques, ensuring that structural features essential for embryo assessment remain unaffected. This approach would ensure that the generative models are exposed only to high-quality examples, reducing the likelihood of producing too dark synthetic images. Addressing such data quality issues is crucial for further improving the effectiveness of the models.

In addition to these considerations, while our work demonstrates the potential of generative models to enhance embryo classification and overcome data scarcity, several limitations should be acknowledged. We did not assess the similarity between synthetic and real images, which could be measured using similarity metrics to rule out the risk of overfitting. Overfitting is a common problem when dealing with small datasets and powerful deep-learning models. It happens when models achieve very good results on the training data but struggle with new data (such as test data) because they have memorized specific patterns rather than learned to generalize effectively. In our case, overfitting would mean that the generative models 'copied' the real images rather than generated new, diverse ones. Since incorporating synthetic data led to improved classification performance, even on an external dataset, we believe it is unlikely that overfitting occurred. The improved performance suggests that the synthetic data adds value rather than merely replicating the training set. However, to fully rule out overfitting or mode collapse in the generative models, future work should quantify training-generated image similarity using appropriate metrics. Additionally, validating the classification models on embryo images from all developmental stages, sourced from different clinics and imaging systems, will be essential for assessing true generalization and ensuring the models do not overfit to the distribution of synthetic data from a single source.

Another limitation of this study is the selection of cell stages included in the dataset. We focused on the 2-, 4-, 8-cell, morula, and blastocyst stages while excluding intermediate stages, such as the 3- and 5-cell stages, which are clinically relevant despite being less frequent. These stages are often underrepresented because they typically last less than two hours, making them more challenging to capture. However, their scarcity in datasets makes them particularly valuable in the context of missing data. Additionally, our dataset primarily included embryos with minimal fragmentation. In clinical practice, embryos with higher fragmentation are common, making classification more challenging. Focusing on lower-fragmentation embryos may have simplified the task, potentially overestimating our model's performance. Future work should address these limitations by incorporating intermediate cell stages, such as 3- and 5-cell embryos, as well as embryos with higher fragmentation rates. Furthermore, expanding the dataset to include multiple focal planes would better capture the full spectrum of embryo variations observed in clinical practice, ultimately improving model robustness and clinical applicability.

Furthermore, embryologists analyzing the synthetic images were all from the same clinic, Volvat Spiren, where they are accustomed to viewing and analyzing embryo images from their lab only. Upon analyzing the qualitative results, we found that certain image regions, such as the dark background or the embryo's zona pellucida (ZP), sometimes misled embryologists into identifying real images as synthetic. These images were from the second dataset we used. Although images from this second dataset were captured with the same Embryoscope™, variations in lab settings may have impacted image quality, introducing potential bias into the evaluation process. To address this, future research could build on this work by involving a more diverse group of embryologists from different clinics to refine the process of embryo synthetic image evaluation. Another potential direction for future research could involve incorporating images of embryos undergoing necrosis into model training. Our real dataset contained a limited number of such images, which impacted the generative models' ability to produce high-fidelity embryo images under necrosis. Introducing synthetic embryo images under necrosis could prove valuable for future studies by enabling models to better recognize when an embryo is no longer viable. This advancement would aid embryologists in determining when an embryo is unsuitable for use, enhancing decision-making in clinical settings.

Data availability

The two real datasets used in this study are publicly available and can be accessed on Zenodo at <https://zenodo.org/records/6390798> and <https://zenodo.org/records/14253170>. The synthetic dataset we generated and used to train the classification models can be found at https://huggingface.co/datasets/deepsynthbody/synembryo_latentdiffusion and https://huggingface.co/datasets/deepsynthbody/synembryo_stylegan. The blastocyst dataset used for evaluation is available at <https://github.com/ih-lab/STORK>.

Code availability

All the code used in this study is publicly available in our GitHub repository <https://github.com/orianapresaca/EmbryoStageGen> to ensure reproducibility. Additionally, the repository includes links to the checkpoints of the best-performing models, which were used for image generation and classification tasks.

Received: 9 December 2024; Accepted: 17 March 2025

Published online: 21 March 2025

References

1. World Health Organization. Infertility. <https://www.who.int/news-room/fact-sheets/detail/infertility> (2024).

2. The European IVF Monitoring Consortium (EIM) for the European Society of Human Reproduction and Embryology (ESHRE) et al. ART in Europe, 2019: Results generated from European registries by ESHRE†. *Hum. Reprod.* **38**, 2321–2338, <https://doi.org/10.1093/humrep/dead197> (2023).
3. Sundvall, L., Ingerslev, H. J., Breth Knudsen, U. & Kirkegaard, K. Inter- and intra-observer variability of time-lapse annotations. *Hum Reprod.* **28**(12), 3215–3221. <https://doi.org/10.1093/humrep/det366> (2013).
4. Adolfsen, E. & Andershed, A. N. Morphology vs morphokinetics: A retrospective comparison of inter-observer and intra-observer agreement between embryologists on blastocysts with known implantation outcome. *JBRA Assist. Reprod.* **22**(3), 228–237. <https://doi.org/10.5935/1518-0557.20180042> (2018).
5. Riegler, M. A. et al. Artificial intelligence in the fertility clinic: Status, pitfalls and possibilities. *Hum. Reprod.* **36**, 2429–2442. <https://doi.org/10.1093/humrep/deab168> (2021).
6. Gliozheni, O. et al. ART in Europe results generated from European registries by ESHRE. *Hum. Reprod. Open.* <https://doi.org/10.1093/hropen/hoac022> (2018).
7. Sciorio, R., Aiello, R. & Janssens, R. Considerations on staffing levels for a modern assisted reproductive laboratory. *JBRA Assist. Reprod.* <https://doi.org/10.5935/1518-0557.20220048> (2022).
8. Thirumalaraju, P. et al. Evaluation of deep convolutional neural networks in classifying human embryo images based on their morphological quality. *Heliyon* **7**, e06298. <https://doi.org/10.1016/j.heliyon.2021.e06298> (2021).
9. Danardono, G. B. et al. A homogeneous ensemble of robust pre-defined neural network enables automated annotation of human embryo morphokinetics. *J. Reprod. Infertil.* **23**, 250–256 (2022).
10. Dimitriadis, I. et al. Deep convolutional neural networks (CNN) for assessment and selection of normally fertilized human embryos. *Fertility Steril.* **112**, e272. <https://doi.org/10.1016/j.fertnstert.2019.07.805> (2019).
11. Liu, Z. et al. Multi-task deep learning with dynamic programming for embryo early development stage classification from time-lapse videos. *IEEE Access* **7**, 122153–122163. <https://doi.org/10.1109/ACCESS.2019.2937765> (2019).
12. Lucio, C. M. et al. CHLOE (FAIRILITY) can automatically annotate images from time-lapse cultured embryos for Pronucleate (PN) count, Morphokinetics and Ranking according to ploidy and implantation potential, with a strong agreement compared to experienced embryologists. *Reprod. BioMed. Online* **45**, e34–e35. <https://doi.org/10.1016/j.rbmo.2022.08.058> (2022).
13. Theilgaard Lassen, J., Fly Kragh, M., Rimestad, J., Nygård Johansen, M. & Berntsen, J. Development and validation of deep learning based embryo selection across multiple days of transfer. *Sci. Rep.* <https://doi.org/10.1038/s41598-023-31136-3> (2023).
14. Jiang, V. Artificial intelligence in the IVF laboratory: A review of advancements over the last decade. *Fertility Steril.* <https://doi.org/10.1016/j.fertnstert.2023.05.149> (2023).
15. Bormann, C. et al. Consistency and objectivity of automated embryo assessments using deep neural networks. *Fertility Steril.* **113**, 781–787.e1. <https://doi.org/10.1016/j.fertnstert.2019.12.004> (2020).
16. Goled, S. Why Transformers Are Increasingly Becoming As Important As RNN And CNN? <https://analyticsindiamag.com/why-transformers-are-increasingly-becoming-as-important-as-rnn-and-cnn/> (2021).
17. Cruz, M. et al. Embryo quality, blastocyst and ongoing pregnancy rates in oocyte donation patients whose embryos were monitored by time-lapse imaging. *J. Assisted Reprod. Genetics* **28**, 569–573. <https://doi.org/10.1007/s10815-011-9549-1> (2011).
18. Gilbert, S. F. Early Mammalian Development (Sinauer Associates, 2000), 6th edn. Retrieved 13 May 2022.
19. European Society of Human Reproduction and Embryology (ESHRE). Atlas of human embryology: From oocytes to preimplantation embryos. <https://atlas.eshre.eu/> (2016).
20. Martins, W. P. et al. Blastocyst vs cleavage-stage embryo transfer: Systematic review and meta-analysis of reproductive outcomes. *Ultrasound Obstet. Gynecol.* **49**, 583–591. <https://doi.org/10.1002/uog.17327> (2017).
21. Lemmen, J., Agerholm, I. & Ziebe, S. Kinetic markers of human embryo quality using time-lapse recordings of ivf/icsi-fertilized oocytes. *Reprod. BioMed. Online* **17**, 385–391. [https://doi.org/10.1016/S1472-6483\(10\)60222-2](https://doi.org/10.1016/S1472-6483(10)60222-2) (2008).
22. Urcelay, L. et al. Exploring the Role of Explainability in AI-Assisted Embryo Selection (2023). 2308.02534
23. Salih, M. et al. Embryo selection through artificial intelligence versus embryologists: A systematic review. *Hum. Reprod. Open* **2023**, hoad031. <https://doi.org/10.1093/hropen/hoad031> (2023).
24. Khosravi, P. et al. Deep learning enables robust assessment and selection of human blastocysts after in vitro fertilization. *NPJ Digit. Med.* **2**, 21. <https://doi.org/10.1038/s41746-019-0096-y> (2019).
25. Kanakasabapathy, M. K. et al. Adaptive adversarial neural networks for the analysis of lossy and domain-shifted datasets of medical images. *Nat. Biomed. Eng.* **5**, 571–585. <https://doi.org/10.1038/s41551-021-00733-w> (2021).
26. Kromp, F. et al. An annotated human blastocyst dataset to benchmark deep learning architectures for in vitro fertilization. *Sci. Data* **10**, 271. <https://doi.org/10.1038/s41597-023-02182-3> (2023).
27. Gomez, T. et al. A time-lapse embryo dataset for morphokinetic parameter prediction. *Data Brief* **42**, 108258. <https://doi.org/10.1016/j.dib.2022.108258> (2022).
28. Sheller, M. J. et al. Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data. *Sci. Rep.* <https://doi.org/10.1038/s41598-020-69250-1> (2020).
29. Murtaza, H. et al. Synthetic data generation: State of the art in health care domain. *Comput. Sci. Rev.* **48**, 100546. <https://doi.org/10.1016/j.cosrev.2023.100546> (2023).
30. Radford, A., Metz, L. & Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial. *Networks* **1511**, 06434 (2016).
31. Behara, K., Bhero, E. & Agee, J. T. Skin lesion synthesis and classification using an improved dcgan classifier. *Diagnostics*. <https://doi.org/10.3390/diagnostics13162635> (2023).
32. Porkodi, S. P., Sarada, V. & Maik, V. Dcgan for data augmentation in pneumonia chest x-ray image classification. In Mahapatra, R. P., Peddoju, S. K., Roy, S. & Parwekar, P. (eds.) *Proceedings of International Conference on Recent Trends in Computing*, 129–137 (Springer Nature Singapore, Singapore, 2023).
33. Karras, T., Laine, S. & Aila, T. A Style-Based Generator Architecture for Generative Adversarial Networks (2019). 1812.04948.
34. Rombach, R., Blattmann, A., Lorenz, D., Esser, P. & Ommer, B. High-Resolution Image Synthesis with Latent Diffusion Models (2022). 2112.10752.
35. Abdalaziz, A. & Schwenker, F. *Generating synthetic brain tumor data using StyleGAN3 for lower class enhancement*, 73–87 (Springer Nature, Switzerland, 2024).
36. Hong, S. et al. 3D-StyleGAN: A style-based generative adversarial network for generative modeling of three-dimensional medical images (2021).
37. Pinaya, W. H. L. et al. Brain imaging generation with latent diffusion models. In *Deep Generative Models*, 117–126 (eds Mukhopadhyay, A. et al.) (Springer Nature, Cham, 2022).
38. Ivanovs, M. et al. Synthetic image generation with a fine-tuned latent diffusion model for organ on chip cell image classification. In *2023 Signal Processing: Algorithms, Architectures, Arrangements, and Applications (SPA)*, 148–153. <https://doi.org/10.23919/SPA59660.2023.10274460> (2023).
39. Dirvanauskas, D., Maskeliunas, R., Raudonis, V., Damaševičius, R. & Scherer, R. HEMIGEN: Human embryo image generator based on generative adversarial networks. *Sensors* **19**, 3578. <https://doi.org/10.3390/s19163578> (2019).
40. Cao, P. et al. Generative artificial intelligence to produce high-fidelity blastocyst-stage embryo images. *Hum. Reprod.* **39**, 1197–1207. <https://doi.org/10.1093/humrep/deae064> (2024).
41. Chae, H. J. et al. P-182 high-resolution synthetic embryo images generated by AI. *Hum. Reprod.* **39**, 108553. <https://doi.org/10.1093/humrep/deae108.553> (2024).

42. Liao, Q. et al. Development of deep learning algorithms for predicting blastocyst formation and quality by time-lapse monitoring. *Commun. Biol.* **4**, 415. <https://doi.org/10.1038/s42003-021-01937-1> (2021).
43. Raudonis, V., Paulauskaite-Taraseviciene, A., Sutiene, K. & Jonaitis, D. Towards the automation of early-stage human embryo development detection. *Biomed. Eng. Online* **18**, 120. <https://doi.org/10.1186/s12938-019-0738-y> (2019).
44. Canat, G., Duval, A., Gidel-Dissler, N. & Boussommier-Calleja, A. A novel deep learning approach to identify embryo morphokinetics in multiple time lapse systems. *Sci. Rep.* **14**, 29016. <https://doi.org/10.1038/s41598-024-80565-1> (2024).
45. Azizi, S., Kornblith, S., Saharia, C., Norouzi, M. & Fleet, D. J. Synthetic Data from Diffusion Models Improves ImageNet Classification (2023). 2304.08466.
46. Zhou, Y., Sahak, H. & Ba, J. Training on Thin Air: Improve Image Classification with Generated Data (2023). 2305.15316.
47. Al-Qerem, A. et al. Synthetic generation of multidimensional data to improve classification model validity. *J. Data Inf. Qual.* <https://doi.org/10.1145/3603715> (2023).
48. ESHRE Guideline Group on Good Practice in IVF Labs et al. Revised guidelines for good practice in IVF laboratories (2015)†. *Hum. Reprod.* **31**, 685–686. <https://doi.org/10.1093/humrep/dew016> (2016).
49. Montag, M., Liebenthron, J. & Köster, M. Which morphological scoring system is relevant in human embryo development?. *Placenta* **32**, S252–S256. <https://doi.org/10.1016/j.placenta.2011.07.009> (2011).
50. Lundin, K. & Park, H. Time-lapse technology for embryo culture and selection. *Upsala J. Med. Sci.* **125**, 1–8. <https://doi.org/10.1080/03009734.2020.1728444> (2020).
51. Sciorio, R. & Smith, G. Embryo culture at a reduced oxygen concentration of 5%: A mini review. *Zygote* **27**, 355–361. <https://doi.org/10.1017/s0967199419000522> (2019).
52. Louis, C. M. et al. Review of computer vision application in in vitro fertilization: The application of deep learning-based computer vision technology in the world of IVF. *J. Assist. Reprod. Genet.* **38**, 1627–1639. <https://doi.org/10.1007/s10815-021-02123-2> (2021).
53. Karras, T. et al. Training Generative Adversarial Networks with Limited Data (2020). 2006.06676.
54. Ronneberger, O., Fischer, P. & Brox, T. U-Net: Convolutional networks for biomedical image segmentation. In International Conference on Medical Image Computing and Computer-Assisted Intervention, 234–241 (Springer, 2015).
55. Karras, T., Laine, S. & Aila, T. Flickr-Faces-HQ Dataset (FFHQ). <https://github.com/NVLabs/ffhq-dataset> (2019).

Author contributions

O.P., A.S., V.T., A.D., and M.R. conceived the experiments and study design. O.P. conducted the experiments. A.S. conducted the background research on embryology and AI. M.S. and M.I. collected the data, while A.S. was responsible for the preprocessing. A.D. developed the web application. M.S. and M.I. performed qualitative evaluations of the synthetic images. O.P., V.T., M.R., A.A., and A.S. analyzed the results. All authors reviewed and contributed to the final manuscript.

Declarations

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-025-94680-0>.

Correspondence and requests for materials should be addressed to O.P.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025