



HAL
open science

Supervised contrastive learning for cell stage classification of animal embryos

Yasmine Hachani, Patrick Bouthemy, Elisa Fromont, Sylvie Ruffini, Ludivine Laffont, Alline de Paula Reis

► **To cite this version:**

Yasmine Hachani, Patrick Bouthemy, Elisa Fromont, Sylvie Ruffini, Ludivine Laffont, et al.. Supervised contrastive learning for cell stage classification of animal embryos. 2025. hal-04937720v2

HAL Id: hal-04937720

<https://hal.science/hal-04937720v2>

Preprint submitted on 13 Feb 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Supervised contrastive learning for cell stage classification of animal embryos

Yasmine Hachani^{1,*}, Patrick Bouthemy¹, Elisa Fromont^{1,2}, Sylvie Ruffini³, Ludivine Laffont³, and Aline de Paula Reis^{3,4}

¹Inria center at Rennes University, France

²University of Rennes, IRISA, France

³Paris-Saclay University, UVSQ, INRAE, BREED, France

⁴The National Veterinary School of Alfort (EnvA), France

*yasmine.hachani@inria.fr

ABSTRACT

Video microscopy, when combined with machine learning, offers a promising approach for studying the early development of *in vitro* produced (IVP) embryos. However, manually annotating developmental events, and more specifically cell divisions, is time-consuming for a biologist and cannot scale up for practical applications. We aim to automatically classify the cell stages of embryos from 2D time-lapse microscopy videos with a deep learning approach. We focus on the analysis of bovine embryonic development using video microscopy, as we are primarily interested in the application of cattle breeding, and we have created a Bovine Embryos Cell Stages (ECS) dataset. The challenges are three-fold: (1) low-quality images and bovine dark cells that make the identification of cell stages difficult, (2) class ambiguity at the boundaries of developmental stages, and (3) imbalanced data distribution. To address these challenges, we introduce CLEmbryo, a novel method that leverages supervised contrastive learning combined with focal loss for training, and the lightweight 3D neural network CSN-50 as an encoder. We also show that our method generalizes well. CLEmbryo outperforms state-of-the-art methods on both our Bovine ECS dataset and the publicly available NYU Mouse Embryos dataset.

Keywords: Cell-stage classification, Deep learning, Supervised contrastive learning, Time-lapse video, Embryonic development, Morphokinetics

Introduction

The analysis of embryonic development is of key interest to better understand biological mechanisms and to address a range of human and animal concerns. When studying *in vitro* produced (IVP) embryos, analysis can benefit from image sequences obtained from a video microscopy setup. In this paper, we will study the embryonic developmental stages of IVP embryos from time-lapse video sequences. The normal development of a mammalian embryo is characterized by a series of mitoses, in which each cell divides into two daughter cells (also known as cleavage). Successive developmental stages, or cell stages, are characterized by a given number of cells, starting from one cell. We formulate this analysis of embryonic development as a supervised classification problem. Each class corresponds to a cell stage, and is related to a number of cells. This problem is unbalanced, and a large majority of images show a number of cells that is of a power of two. In the sequel, we will call intermediate stages the other cell stages (3-cell, 5-cell, 6-cell, and 7-cell stages). We will not proceed to an explicit numbering of cells based on cell segmentation as done in¹, but to a direct classification from images based on a deep learning approach.

Observation of embryos by conventional microscopy requires them to be removed from the incubator, which leads to disturbances in temperature and pH of the culture medium that can be detrimental for embryo development². Video microscopy performed directly in the incubator provides a reliable way of capturing images of embryos at regular and short intervals. It enables embryologists to determine cell numbers almost in real time, without interfering with the culture conditions required for normal embryo development. It also has the advantage of being compatible with further transfer to establish a pregnancy³. However, manually determining the number of cells is a time-intensive task and limits the extent to which laboratories equipped with a video microscope can make more extensive use of this information. This argues for automation.

We will focus on the analysis of bovine embryonic development using video microscopy, as we are primarily interested in the application for cattle breeding. Its use in the bovine domain is recent but promising. More specifically, it allows one to assess whether embryos can be transferred to a cow uterus or not⁴. From an image processing perspective, images of bovine embryos are especially difficult to process due to their dark appearance as shown in Figure 1. To demonstrate the generality of our method, we will also apply it to the analysis of mouse embryonic development.




Figure 1. Sample images of the different cell stages of bovine (top row) and mouse (bottom row) embryo development. The bovine embryo only occupies a small part of the image, it is located in a micro-well (light grey) within the Petri dish (black). The bovine embryo is darker. The mouse embryo is more transparent. Both embryos are surrounded by the zona pellucida. Images are taken from the Bovine Embryo CS dataset and from the NYU Mouse Embryo dataset respectively.

Existing work dealing with cell stage classification and making use of deep learning techniques can be divided into two categories. The first category of work considers every frame of a time-lapse video independently and employs 2D convolutional neural networks (CNNs) to achieve the classification⁵. This is usually followed by a post-processing step that enforces a monotonic progression of the number of cells as done in^{6,7,8} and⁹, all using dynamic programming to apply that constraint. In⁹, the development stages are identified through several sub-steps that also take into account the rate of embryo fragmentation. In¹, the authors adopt an approach based on object detection using the YOLO-v5 network¹⁰ and perform cell counting. Very recently, the DLT-Embryo method was proposed¹¹. It combines dual-branch local feature fusion (DLF) modules and transformer encoder modules to extract both local and global features from each frame. These methods make limited use of temporal information.

Embryologists take into account dynamic information to decide the completion of the cell cleavage event. In addition, temporal context helps maintain consistency between time-lapse images. The second category of methods integrates temporal information in the architecture, and then, exploits it as soon as the training stage.

The ESOD method¹² employs a 2D-CNN followed by a Long Short-Term Memory (LSTM) network classifier, and involves a synergic loss to learn embryo-independent features. The CNNs-CRF method proposed in¹³ involves a conditional random field (CRF) to favor the monotonic progression of the number of cells. In addition, the neural network leverages temporal information through two streams with different input, one single frame for the first one and two concatenated consecutive frames for the second one. The R2D1 method¹⁴ considers the spatio-temporal video as a volume, uses 3D-CNN (i.e., 2D+t) to work directly on video sub-sequences, and applies the Viterbi algorithm to enforce a monotonic order constraint. The EmbryosFormer framework¹⁵ is the first to involve transformer architecture, building a three-headed encoder-decoder transformer inspired by the Deformable DETR model¹⁶. All of these methods, were evaluated on datasets of human and mouse embryos, except R2D1 which was evaluated only on humans.

As stated before, our objective is to classify the cell stages of an animal embryo from 2D time-lapse videos acquired by video-microscopy. This prediction will be based on a 3D-CNN architecture and will leverage supervised contrastive learning (SCL)¹⁷. We assume that the n -cell stage lasts until the $n + 1$ -cell stage appears. We formulate this problem as a multi-class supervised classification. Each frame of an input time-lapse video will be assigned a class from the one-cell stage to the nine-and-plus-cell stage.

The classification problem is challenging, even if it is supervised. The cell classes are heavily imbalanced. The intermediate cell-stages are largely underrepresented. Frames located just around the end of a cell cleavage, consequently, around a stage change, are particularly hard to classify. In fact, immediately before the end of the daughter-cell separation, two daughter cells can be seen despite the incomplete membrane separation. In addition, immediately after complete separation, cell movements due to rearrangement of newly created cells in the embryo can be confused with the movements required for cell separation. Embryo images usually show low contrast, significant noise, and intricate motion. Bovine embryos are harder to study using video microscopy than their human or mouse counterparts due to the darkness of their cells. Cell cleavage may last longer than for mouse embryos.

Work that combines video microscopy and machine learning for bovine embryos is rare. In¹⁸, a method was developed for the prediction of embryo transferability based on random forests. It requires very detailed manual annotations of each video. In⁴, we have designed a 3D-CNN called SFR that includes three paths and the focal loss¹⁹, to achieve transferability prediction directly from videos. To our knowledge, we are the first to propose a deep learning-based model to classify the stages of early development of bovine embryos and to leverage SCL for this kind of task.

SCL¹⁷ is a contrastive learning technique that uses labeled data to learn more robust and semantically meaningful representations. It extends the principle of contrastive learning, which aims to pull similar data points closer in the representation space and push dissimilar ones further apart. SCL incorporates label information to define similarity and dissimilarity between samples. SCL has been shown to be a promising approach in the field of medical image analysis. It has been applied to a variety of tasks, including medical image classification²⁰ and segmentation²¹.

Materials and methods

This work had several objectives: 1) to propose a new method for classifying bovine embryo cell stages, 2) to verify that it generalizes well to other mammal species, 3) to compare its performance with four other existing methods (ESOD, CNNs-CRF, EmbryosFormer, R2D1) commented in the Introduction section. The method was tested on our Bovine Embryos Cell Stage dataset (or Bovine ECS dataset for short) and the NYU mouse embryos dataset described below.

Datasets

Bovine Embryos Cell Stage dataset

We have acquired a video dataset of IVP bovine embryos that we call the Bovine ECS dataset. The embryo production process and the acquisition of embryo videos are described in this section. The main features of our dataset are also specified in the following.

The bovine embryos were obtained from oocytes recovered *post mortem* on slaughterhouse ovaries, fertilized *in vitro* and placed in the incubator in PrimoVision dishes containing 16 individual microwells filled with SOF (Synthetic Oviduct Fluid) medium under oil and cultured for eight days at 38.5°C in a 5%CO₂, 5%O₂ and 90%N₂ humidified atmosphere. The PrimoVision system is equipped with a transmission light microscope. Images of the Petri dishes were automatically taken every fifteen minutes throughout the embryo culture (from 22 to 192 hours of development). The resulting 2D time-lapse videos were subsequently divided into sixteen individual time-lapse videos. The different cell stages of each embryo were then annotated by an embryologist: the first image presenting two daughter cells and their entire membranes was considered as a stage change. These annotations were used for training and evaluation. The videos used for this study comprises 300 frames of 256 x 256 pixels. Knowing that the first frame was taken at the beginning of the embryo culture (at $t_0 = 22h$), and that the interval between two successive images was 15 minutes, the end of the 300-frame video corresponds to the fourth day of embryo development and between 8 and 16 cells for normal development in bovine species.

Our final data set consists of 485 videos, of which 345 are kept for training our models, 50 are kept for parameter validation, and 79 videos are kept for evaluating the performance (test). Details on the data distribution of each cell stage in the dataset are given in Table 1.

Mouse Embryos dataset

We also consider mouse embryos to evaluate the generality of our proposed method. The NYU Mouse Embryo dataset²² contains 100 videos of developing mouse embryos, originally created for the task of cell tracking. They were acquired with a Nikon Eclipse Ti inverted microscope and a heated stage-top Tokai incubator. Each time-lapse video consists of 480 × 480 pixel images captured every seven minutes, resulting in 314 images per sequence on average. The image capture frequency is higher than for bovine videos, as mouse embryos develop faster. After downloading the dataset, we had access to 99 videos, divided into 72 sequences for training, 8 for validation, and 19 for testing. We used the annotations provided by the authors of the CNNs-CRF¹³ method. Details on the data distribution of each cell stage in the dataset are also given in Table 1.

Description of our method (CLEmbryo)

We propose an original method, called CLEmbryo, to classify the cell stages that appear during embryonic development. It integrates three key components: (1) a modified version of the SCL framework, (2) a loss function that combines focal loss (FL) and supervised contrastive loss (called SupCon¹⁷), and (3) the Channel-Separated Convolutional Network (CSN)²³ 3D-CNN architecture as an encoder.

SCL framework

Inspired by the original paper on supervised contrastive learning¹⁷, our framework consists of four components: the data augmentation module $Aug(\cdot)$, the encoder network $Enc(\cdot)$, the projection network $Proj(\cdot)$ and the classification network $Class(\cdot)$. The overall framework is presented in Figure 2. In contrast to recent implementations of the SCL framework with two losses and a single training stage, we also carry out a single training stage with two losses but corresponding to two heads, classification head and projection head.

At training time, each video sequence $\mathbf{v} \in \mathbb{R}^{T \times H \times W \times C}$, where H and W respectively denote the height and width of every image, T the temporal length of the sequence, and C the number of image channels, passes through the data augmentation module that generates random augmentations $\mathbf{v}_i \in \mathbb{R}^{T \times H \times W \times C}$ representing different views with distinct information derived

Table 1. Distribution of cell stages across the Bovine ECS and Mouse Embryos datasets, for training, validation and test. We excluded the 8-cell stage frames of the Mouse Embryos dataset as they appear only in the last two frames of each video, making it extremely underrepresented and with a high risk to bias the model training phases.

Bovine					Mouse				
Cell stage	Train	Val	Test	Total	Cell stage	Train	Val	Test	Total
1	14780	2619	3005	20404	1	3496	485	998	4979
2	12267	1753	2973	16993	2	11971	1311	3090	16372
3	1139	184	273	1596	3	567	49	138	754
4	9391	1125	2601	13117	4	6353	715	1689	8757
5	1292	253	267	1812	5	354	31	76	461
6	2369	200	803	3372	6	472	54	155	681
7	3225	806	826	4857	7	620	39	96	755
8	34716	3045	7467	45228					
9+	24921	5015	9385	39321					
Total	204100	15000	27600	246100	Total	23833	2684	642	32759




Figure 2. The SCL framework of our CLEmbryo method with the data augmentation module $Aug(\cdot)$, the encoder network $Enc(\cdot)$, the projection network $Proj(\cdot)$ and the classification network $Class(\cdot)$. The SupCon loss \mathcal{L}_{SupCon} is computed using the output of $Proj(\cdot)$, while the focal loss \mathcal{L}_{FL} is computed with the output of $Class(\cdot)$. The data augmentation module generates two augmented version of the original sequence, and each version is processed independently. We set $d_{proj} = 128$ in all the experiments. The value of d_{emb} depends on the encoder CNN and is set to 512 when CSN-50 is involved.

from the original input. The list of the augmentations taken into account is provided in the dedicated section below. Then, each view \mathbf{v}_i is processed independently.

The encoder embeds each frame of each sequence in a vector, $Enc(\mathbf{v}_{i,t}) = \mathbf{r}_{i,t} \in \mathbb{R}^{d_{emb}}$, $t \in [1, T]$. Every representation vector is normalized to the unit hypersphere in $\mathbb{R}^{d_{emb}}$. The projection head and the classification head process the representation vectors simultaneously. The projection head maps $\mathbf{r}_{i,t}$ to a vector $\mathbf{z}_{i,t} = Proj(\mathbf{r}_{i,t}) \in \mathbb{R}^{d_{proj}}$ to reduce the number of operations required to compute the SupCon loss. $Proj(\cdot)$ is a single linear layer and all projection vectors are normalized to the unit hypersphere in $\mathbb{R}^{d_{proj}}$. Likewise, the classification head is a single linear layer that maps $\mathbf{r}_{i,t}$ to a classification vector $\mathbf{l}_{i,t} = Class(\mathbf{r}_{i,t}) \in \mathbb{R}^K$, with K denoting the number of classes, from which the focal loss is computed. Once training is achieved, both the data augmentation module and the projection head are removed, and inference is performed on the resulting architecture.

Loss function

The training loss involves two loss terms as illustrated in Figure 2, the SupCon loss and the classification loss.

The SupCon loss is built on the concept of contrastive loss extended to supervised scenarios by the authors of¹⁷. This loss adds information on whether two elements belong to the same class, thus helping to fix the limitations of self-supervised contrastive learning, which might push apart the representations of same class samples. By enforcing clustering within classes, the SupCon loss ensures that the encoder produces representations that are more robust and discriminative, ultimately to the benefit of downstream classification tasks. Let us consider N video sequences; the data augmentation module generates two augmented versions of each sequence. We note $\mathcal{S} = \{1, \dots, 2N\}$ the set of indexes of all the augmented videos \mathbf{v}_i and $\mathbf{y}_i \in [1, K]^T$ their associated vector of class labels. For a given frame $\mathbf{v}_{i,t}$, let us also consider $A(i, t) = \mathcal{S} \times [1, T] \setminus \{(i, t)\}$, the

set of indexes of all the augmented frames except $\mathbf{v}_{i,t}$ and $P(i,t) = \{(p,s) \in A(i,t), y_{p,s} = y_{i,t}\}$, the set of indexes of all the augmented frames different from $\mathbf{v}_{i,t}$ but with the same label. The SupCon loss writes:

$$\mathcal{L}_{SupCon} = \sum_{i \in \mathcal{I}} \sum_{t=1}^T \mathcal{L}_{SupCon}(\mathbf{v}_{i,t}, y_{i,t}) = \sum_{i \in \mathcal{I}} \sum_{t=1}^T \frac{1}{|P(i,t)|} \sum_{(p,s) \in P(i,t)} \log \left(\frac{\exp(\frac{\mathbf{z}_{i,t} \cdot \mathbf{z}_{p,s}}{\tau})}{\sum_{(a,x) \in A(i,t)} \exp(\frac{\mathbf{z}_{i,t} \cdot \mathbf{z}_{a,x}}{\tau})} \right), \quad (1)$$

with $\mathbf{z}_{i,t} = Proj(Enc(\mathbf{v}_{i,t}))$, and $\tau \in \mathbb{R}_+^*$ a scalar temperature parameter. As illustrated in²⁴, this loss can be used as an auxiliary objective during the training process.

We could adopt different loss functions for the classification objective. Since our data are unbalanced with respect to the cell stages, we have considered the focal loss¹⁹, initially introduced for the object detection task. The focal loss helps to mitigate this imbalance, while focusing on the most difficult examples. The focal loss writes:

$$\mathcal{L}_{FL} = \sum_{i \in \mathcal{I}} \sum_{t=1}^T \mathcal{L}_{FL}(\mathbf{v}_{i,t}, y_{i,t}) = \sum_{i \in \mathcal{I}} \sum_{t=1}^T - \sum_{k=1}^K \alpha_k (1 - \hat{p}(k|\mathbf{v}_{i,t}))^\gamma p(k|\mathbf{v}_{i,t}) \log \hat{p}(k|\mathbf{v}_{i,t}), \quad (2)$$

where k denotes one of the K classes, $\hat{p}(k|\mathbf{v}_{i,t})$ the predicted probability of having class k given video frame $\mathbf{v}_{i,t}$, and $p(k|\mathbf{v}_{i,t})$ the true one. The latter is equal to 1 if the right class k is assigned to $\mathbf{v}_{i,t}$, since we are dealing with supervised classification. In addition, α_k is the weight for class k , γ the focusing parameter. The larger γ , the less importance is given to well-classified samples.

Our final objective function combines the proposed SupCon loss for representation learning and the focal loss for classification, as follows:

$$\mathcal{L}_{train} = w \mathcal{L}_{SupCon} + (1 - w) \mathcal{L}_{FL}. \quad (3)$$

In all experiments, we set $w = 0.5$, $\tau = 0.5$ and $\gamma = 2$, this value for γ being recommended in²⁵ and⁴.

Encoder neural network

We believe that a 3D-CNN network is more adapted to properly capture the spatio-temporal features characterizing the embryonic development, than for instance a recurrent neural network as demonstrated previously^{4,14}. By 3D, we mean two spatial dimensions and the temporal one (2D+t). The 3D-CNN selected for the encoder must be both accurate and lightweight as 3D convolutions are computationally costly. For this reason, we chose the CSN²³, which effectively meets these requirements. The CSN architecture leverages group convolution to separate channel interactions from spatio-temporal interactions. This significantly reduces computational cost while preserving sufficient channel interactions to maintain high accuracy. We use CSN-50 with 50 layers and 13.6M parameters. For comparison, 3D-ResNet-50 approximately comprises 45M parameters. We have adapted the CSN architecture to avoid any reduction in time resolution. To ensure that a prediction is made for each frame in the sequence, we set pooling and stride to one in the temporal dimension.

Data augmentation

Data augmentation should introduce sufficient diversity in the versions of the original input without altering the biological meaning of the videos. Each augmentation is selected at random in a set of augmentation types, and is applied to all images of a given sequence. The set of augmentations includes random change in brightness and contrast, horizontal and vertical flipping, random rotation from 0 to 30 degrees, random translation, random cropping, and cutout²⁶ that occludes a square region, chosen at random, of the input video.

Implementation details

Each model was trained using the AdamW optimizer²⁷, with a learning rate of 5×10^{-4} and the other parameters kept at their default values. We applied a learning rate reduction by a factor of 0.1 in the plateau scheduler. When training CLEmbryo, a random subsequence of consecutive frames is selected from each video at each epoch. We prefer to use a subsequence rather than a single frame, as local temporal information is beneficial to classify the cell stage. As for the size of the subsequence, the larger the time window, the more temporal context is provided, but this leads to higher computational costs. We set the size at ten, this choice being motivated in the ablation study. We trained the models using mini-batches of 32 samples in total, meaning that within the SCL framework, 16 different subsequences are randomly selected. We applied the stochastic weight averaging (SWA)²⁸ technique, which improves the generalization of our models by averaging the network weights obtained at several well-chosen epochs. We used early stopping to end training, when the loss computed on the validation set increases ten epochs in a row. Then, we selected the model at the epoch with the lowest loss in the validation set, which is feasible for supervised training.

During inference of our method, we use an overlapping rolling window technique with ten consecutive frames and a step of four, meaning that for frames not located at the beginning or the end of the sequence, their classification vector is computed as the mean of three classification vectors from the different overlapping windows.

We also developed a simple post-processing step to apply a monotonic growth constraint on the number of cells over time. We denote our method augmented with this added post-processing step as CLEmbryo_PP, and its results are reported in the last column of Table 4 and Table 5. Our post-processing step consists of a simple temporal heuristic that maintains label consistency over time. We substitute, within a sequence of label k , short intervals of labels $k', k' \neq k$ for label k . Our extended method is hence able to detect and correct local errors, that is, some inconsistent transitions in CLEmbryo label output.

Evaluation metrics

To evaluate the performance of all methods, we consider three metrics: global accuracy, F1-score per class, and temporal accuracy. Global accuracy ϑ is the ratio of images correctly classified by the model, and is given by N_{corr}/N_{total} , where N_{corr} is the number of correctly classified images and N_{total} the total number of images. The F1-score per class is defined as the harmonic mean of precision (P) and recall (R) per class: $F1 = 2PR/(P+R)$. Precision (P) and recall (R) are respectively defined by $P = \frac{TP}{TP+FP}$, and $R = \frac{TP}{TP+FN}$, where TP, TN, FP, and FN represent the number of true positives, true negatives, false positives, and false negatives, respectively. The F1-score provides a balanced measure of the method performance. Introduced by R2D1¹⁴, the temporal accuracy ζ is defined as the average proportion of cell stage transitions predicted within a specified time range of the corresponding actual transitions. In this context, the time difference between the predicted and actual transitions is considered acceptable if it is less than a threshold that we set to two frames. This threshold accounts for variability in label annotations by different expert biologists. We have: $\zeta = \frac{\theta - \theta_{far}}{\theta}$, where θ is the total number of cell stage transitions and θ_{far} refers to the number of transitions predicted more than two frames away from the ground truth.

Results

The results of the ablation study performed on the Bovine ECS dataset is presented in Tables 2 and 3. The results of the comparison of the ESOD, CNNs-CRF, EmbryosFormer, R2D1 and our CLEmbryo without or with post-processing (_PP) on the Bovine ECS and the NYU mouse datasets are presented in Tables 4 and 5, respectively. The results obtained on bovine or mouse embryos with pre-training on the mouse or bovine embryos respectively are presented in Table 6.

Table 2. Ablation study on our Bovine ECS dataset in terms of per-class F1-score, global accuracy, and temporal accuracy obtained by our method with cross-entropy and without SCL, our method with cross-entropy, and our method with the R(2+1)D-18 architecture. we carried out five evaluations each time with different training seeds, and we provide the mean and standard deviation. The best scores are highlighted in bold.

Configuration		Cross-entropy w/o SCL	Cross-entropy	Cross-entropy and R(2+1)D-18	CLEmbryo
Metrics	1	86.36 ± 1.98	91.47 ± 0.72	90.94 ± 3.85	90.98 ± 0.67
	2	78.32 ± 3.61	84.99 ± 0.56	83.15 ± 6.42	85.85 ± 0.72
	3	9.86 ± 3.36	12.21 ± 6.88	1.83 ± 2.61	22.46 ± 5.49
	4	68.36 ± 4.46	75.76 ± 0.73	73.28 ± 7.17	77.43 ± 0.92
	5	4.65 ± 1.92	5.67 ± 3.13	0.92 ± 1.65	11.41 ± 4.78
	6	4.77 ± 0.79	7.54 ± 3.97	0.87 ± 1.13	16.17 ± 12.92
	7	1.47 ± 1.20	1.80 ± 1.51	0.93 ± 1.18	12.34 ± 6.70
	8	62.02 ± 2.75	65.78 ± 1.81	61.48 ± 7.36	65.41 ± 1.88
	9+	73.80 ± 4.98	78.46 ± 1.05	76.19 ± 5.14	78.10 ± 0.67
Accuracy (ϑ)		70.09 ± 1.05	73.77 ± 0.63	71.45 ± 5.43	73.57 ± 0.48
Temporal accuracy (ζ)		28.72 ± 4.24	40.42 ± 1.41	26.87 ± 9.60	45.48 ± 2.49
Number of parameters		13.6M	13.6M	31.5M	13.6M

Ablation study

We carried out an ablation study on the components of our method. All the results are collected in Table 2. First, we conducted an ablation experiment on the classification loss. We tested the cross-entropy loss instead of the focal loss. These two losses have different behaviors. The former is expected to provide better F1-scores on the most represented classes, which should result in a better global accuracy. The latter enhances the F1-scores on less-represented or underrepresented classes, improving the temporal accuracy. In our context, the behavior of the focal loss is *a priori* more desirable, as we value the detection

Table 3. Results on our Bovine ECS dataset, in terms of global accuracy and temporal accuracy, obtained by our CLEmbryo method trained with input sequence lengths of 3, 5, and 10 frames. we carried out five evaluations each time with different training seeds, and we provide the mean and standard deviation. The best scores are highlighted in bold.

	Sequence length		
Metrics	3	5	10
Accuracy (ϑ)	70.70 \pm 1.71	71.31 \pm 1.14	73.57 \pm 0.48
Temporal accuracy (ζ)	22.04 \pm 2.27	28.60 \pm 6.25	45.48 \pm 2.49

of cell stage changes, which is confirmed by the results reported in Table 2. The goal of the second ablation study was to verify if the supervised contrastive learning framework does improve performance. Indeed, training CSN-50 within our supervised contrastive learning framework yields accuracy and temporal accuracy gains of 3.68 and 11.7 respectively. We also observe less variability on most evaluation metrics. Our last ablation experiment dealt with the encoder neural network. We compared CSN-50 and R(2+1)D-18²⁹ that achieved better results than a standard 3D-ResNet-18¹⁴. All evaluation metrics display higher values with CSN-50, which justifies the choice of this CNN. In addition, CSN-50 is significantly lighter (13.6M parameters versus 31.5M). Incidentally, a transformer module is even heavier, for instance, it contains 41M parameters in the EmbryosFormer method.

Finally, we investigated the influence of the length of the input sequence on the performance of our CLEmbryo method. We trained our method with input sequence lengths of 3, 5, and 10 frames. We set the maximum sequence length to ten frames to allow a large enough batch size, which is very important for the contrastive learning-based method to converge¹⁷. The results are reported in Table 3. Global accuracy and temporal accuracy increase as the temporal window widens. This highlights the importance of providing sufficient temporal context to the contrastive learning-based method to converge.

Classification results

We carried out comparative experiments on the classification of the developmental stages of bovine and mouse embryos. We performed the cell stage classification with our CLEmbryo method and four existing state-of-the-art methods, including EmbryosFormer¹⁵, R2D1¹⁴, CNNs-CRF¹³ and ESOD¹². We were unable to compare our method with DLT-embryo as we did not have access to their code. All methods were evaluated on the same data split with five evaluations performed using the same set of training seeds. We pre-trained the 2D-ResNet-50 used by EmbryosFormer to extract visual features of each frame before letting the transformer process the temporal information, with a contrastive learning task based on the SimCLR framework³⁰, for both datasets. SimCLR has proven to be a good pre-training method to learn efficient representations, including for biomedical images³¹. We trained and tested R2D1 just as our method, using subsequences of ten consecutive frames.

Our CLEmbryo method outperforms the four other methods for both datasets. It achieves higher scores for all metrics and all stages except the five cell stage for the Bovine ECS dataset and for all metrics and all stages for the Mouse Embryos dataset.

The results of all the compared methods obtained on the Bovine ECS dataset are presented in Table 4. This dataset is highly unbalanced, making the correct classification of the intermediate stages difficult. This explains the low scores obtained for these classes. CLEmbryo achieves better performance for all the F1-scores per class, except the 5-cell class, where EmbryosFormer performs better. CLEmbryo reaches a global accuracy of 73.57% with high stability (standard deviation of 0.48). The temporal accuracy of our method is also considerably better with a gain of 18 points compared to the second-best score obtained by R2D1. ESOD could not handle the class imbalance at all and is unable to make a right prediction of the minority classes.

From Table 5, we can observe that the Mouse Embryos dataset is easier to work with than the Bovine one. Indeed, the global accuracy obtained by our CLEmbryo method is 97.73% for the mouse dataset, while it was 73.57% for the bovine dataset. CLEmbryo has the best results for all evaluation metrics for the mouse dataset. We did not include the mouse 8-cell stage as this dataset is very specific for this cell stage. The 8-class stage occurs only and systematically in the last two frames of each video. Then, this cell stage is largely under-represented and this constant configuration can bias training. Regarding EmbryosFormer, the results we obtained are significantly lower than those published by the authors in their paper¹⁵ (global accuracy of 98.4%). This may be due to the way the 2D-ResNet-50 module was pre-trained, and to the choice of data split. In addition, the authors reported results obtained from one single experiment, which may unintentionally hide a variability issue.

We also applied to both datasets the simple post-processing step of CLEmbryo_PP that is described in the section on implementation details. Results obtained with CLEmbryo_PP are slightly better than CLEmbryo ones for the Mouse Embryos dataset as shown in Table 5. For the Bovine ECS dataset, results are mixed. This suggests that on a difficult dataset, where more prediction errors may occur and errors of different types, our heuristic may not be able to mitigate all the errors.

Table 4. Comparison of results obtained on our Bovine ECS dataset. For each method, we carried out five evaluations, each time with different training seeds, and we provide the mean and standard deviation for the per-class F1-score, global accuracy, and temporal accuracy. The best scores are highlighted in bold, excluding CLEmbryo_PP. CLEmbryo_PP corresponds to our method with a post-processing step added.

Metrics \ Methods	ESOD	CNNs-CRF	EmbryosFormer	R2D1	CLEmbryo	CLEmbryo_PP	
	VGG-16	ResNet-50	ResNet-50	R(2+1)D-18	CSN-50	CSN-50	
F1-score per class	1	65.59 ± 36.70	84.60 ± 2.49	87.87 ± 1.87	89.90 ± 0.71	90.98 ± 0.67	90.66 ± 0.62
	2	62.53 ± 6.38	77.87 ± 3.88	81.55 ± 0.55	81.09 ± 0.68	85.85 ± 0.72	85.67 ± 0.50
	3	0 ± 0	4.67 ± 2.68	11.06 ± 3.93	9.08 ± 4.56	22.46 ± 5.49	21.73 ± 5.42
	4	9.79 ± 11.88	66.21 ± 2.39	72.99 ± 1.60	69.52 ± 2.57	77.43 ± 0.92	78.06 ± 2.22
	5	0 ± 0	7.92 ± 8.65	11.47 ± 2.65	5.91 ± 3.51	11.41 ± 4.78	10.52 ± 4.44
	6	0 ± 0	0.69 ± 0.39	5.67 ± 1.25	7.83 ± 7.04	16.17 ± 12.92	16.21 ± 14.92
	7	0 ± 0	9.23 ± 17.86	3.97 ± 1.35	5.34 ± 1.93	12.34 ± 6.70	13.56 ± 12.15
	8	21.57 ± 22.74	49.63 ± 6.27	54.30 ± 1.54	57.85 ± 1.49	65.41 ± 1.88	66.95 ± 2.28
	9+	64.48 ± 2.63	74.59 ± 1.15	68.16 ± 1.78	75.36 ± 1.49	78.10 ± 0.67	77.29 ± 2.44
Accuracy (ϑ)	51.54 ± 4.47	65.61 ± 1.89	65.04 ± 0.54	69.48 ± 0.33	73.57 ± 0.48	73.05 ± 1.12	
Temporal accuracy (ζ)	4.16 ± 1.61	26.10 ± 2.21	25.42 ± 2.65	27.26 ± 1.49	45.48 ± 2.49	46.35 ± 2.53	

Table 5. Comparison of results obtained on the Mouse Embryos dataset²². For each method, we carried out five evaluations, each time with different training seeds, and we provide the mean and standard deviation for the per-class F1-score, global accuracy, and temporal accuracy. The best scores are highlighted in bold, excluding CLEmbryo_PP. CLEmbryo_PP corresponds to our method with a post-processing step added.

Metrics \ Method	ESOD	CNNs-CRF	EmbryosFormer	R2D1	CLEmbryo	CLEmbryo_PP	
	VGG-16	ResNet-50	ResNet-50	R(2+1)D-18	CSN-50	CSN-50	
F1-score per class	1	98.26 ± 1.08	98.66 ± 2.08	98.43 ± 0.59	95.65 ± 1.34	99.61 ± 0.04	99.62 ± 0.03
	2	95.52 ± 2.66	98.33 ± 0.72	98.16 ± 0.55	97.43 ± 0.97	99.58 ± 0.10	99.58 ± 0.11
	3	0 ± 0	24.05 ± 37.38	70.68 ± 6.99	34.19 ± 23.97	90.20 ± 3.40	90.92 ± 3.09
	4	86.18 ± 4.33	90.44 ± 2.50	94.58 ± 1.31	94.14 ± 1.04	98.97 ± 0.87	99.50 ± 0.17
	5	0 ± 0	6.26 ± 10.86	28.62 ± 5.64	14.37 ± 11.17	70.36 ± 6.16	72.69 ± 3.63
	6	8.71 ± 19.48	6.19 ± 11.54	45.56 ± 8.85	19.56 ± 3.76	69.32 ± 10.54	81.50 ± 6.91
	7	0 ± 0	0.51 ± 1.02	61.48 ± 0.51	40.52 ± 20.76	72.71 ± 5.51	83.40 ± 7.43
Accuracy (ϑ)	89.57 ± 2.67	92.24 ± 1.51	94.54 ± 0.71	91.98 ± 0.88	97.73 ± 0.61	98.62 ± 0.44	
Temporal accuracy (ζ)	2.91 ± 3.61	38.07 ± 17.43	69.69 ± 1.69	43.30 ± 9.34	85.14 ± 3.00	85.69 ± 2.11	

Pretraining

Finally, we wanted to know if pretraining our method on the Mouse Embryos (respectively Bovine ECS) dataset could be beneficial on the Bovine ECS (respectively Mouse Embryos) dataset. Although the Bovine Embryos dataset is harder to process than its mouse counterpart, both datasets show similar events and patterns. It is now well-known that pretraining a deep learning model on another dataset is usually beneficial³². Results are reported in Table 6. In both cases, pretraining slightly improves performance as confirmed by almost all evaluation metrics. In particular, we achieve a significant gain in temporal accuracy.

Discussion

We developed a novel method for predicting the cell stages of animal embryos and assessed its performance on embryo video datasets for two different species: bovine and mouse²². We also elaborated a video dataset of annotated bovine embryos. This bovine biological material is particularly challenging because of the darkness of the cells due to the presence of high quantities of intracellular lipids^{33, 34, 35}. This results in a lack of detail in the bovine embryo cells compared to mouse and human embryos that present a higher degree of transparency.

In summary, we formulated the problem as a supervised multi-stage classification while overcoming imbalanced data distribution. Our experiments demonstrated the interest of introducing, for this classification task, supervised contrastive learning, focal loss training, and the lightweight 3D convolutional network CSN-50 as sub-sequence encoder. We have shown that our CLEmbryo method generalizes well. CLEmbryo was able to produce convincing results on embryos from different animal species and on videos acquired with different setups.

Table 6. Results on our Bovine ECS (respectively Mouse Embryos) dataset, in terms of per-class F1-score, global accuracy, and temporal accuracy, obtained by our CLEmbryo method pre-trained on the Mouse Embryos (respectively Bovine Embryos) dataset. We carried out five evaluations each time with different training seeds, and we provide the mean and standard deviation. The best scores are highlighted in bold.

Pretraining		Bovine		Mouse	
		No pretraining	Pretraining	No pretraining	Pretraining
F1-score per class	1	90.98 ± 0.67	92.18 ± 0.79	99.61 ± 0.04	99.65 ± 0.05
	2	85.85 ± 0.72	85.08 ± 1.16	99.58 ± 0.10	99.70 ± 0.11
	3	22.46 ± 5.49	28.60 ± 2.14	90.20 ± 3.40	93.98 ± 1.48
	4	77.43 ± 0.92	77.84 ± 0.98	98.97 ± 0.87	99.09 ± 0.27
	5	11.41 ± 4.78	15.83 ± 5.63	70.36 ± 6.16	69.11 ± 7.11
	6	16.17 ± 12.92	15.27 ± 2.84	69.32 ± 10.54	64.31 ± 9.62
	7	12.34 ± 6.70	24.05 ± 5.53	72.71 ± 5.51	74.75 ± 3.94
	8	65.41 ± 1.88	66.02 ± 1.40		
	9+	78.10 ± 0.67	78.68 ± 0.68		
Accuracy (ϑ)		73.57 ± 0.48	74.17 ± 0.41	97.73 ± 0.61	97.85 ± 0.34
Temporal accuracy (ζ)		45.48 ± 2.49	50.18 ± 2.82	85.14 ± 3.00	87.89 ± 2.19

We favorably compared our CLEmbryo method with other state-of-the-art methods on both the Bovine ECS dataset and the NYU Mouse Embryos dataset. CLEmbryo provides the best scores for almost all of the evaluation metrics, while remaining lightweight. This allowed us to efficiently and accurately predict the cell stages of animal embryos. The improved performance of CLEmbryo is first due to the supervised contrastive learning framework. Adding this auxiliary head to the network to bring the embedding vectors of the same class of images closer together, results in a more robust and accurate model. The focal loss allowed us to improve the results on the less represented and underrepresented classes, thereby increasing the temporal accuracy of our method. The choice of the network architecture is also impactful. Unlike LSTM used in¹², 3D-CNN focuses on local temporal view of the input sequence³⁶, which allows it to better capture short intervals of intermediate cell stages. The LSTM architecture is more sensitive to this class imbalance, since it leads to a more global temporal view of the input, and therefore, is unable to focus on the short appearance of intermediate classes. On our side, CLEmbryo does not require the integration of the monotonic growth constraint to get better results than the other methods. This makes our method more flexible. Nevertheless, the addition of a post-processing step may further improve results. In addition, our method provides results with low variability. This property of stability is of fundamental importance for biologists.

Although CLEmbryo performs very well on the Bovine ECS dataset, there is still room for extension. It is now known that IVP embryos can present variable cleavage durations and cleavage errors such as direct cleavage, i.e., division of a cell directly into three or four daughter cells, or reverse cleavage, i.e., cleavage followed by fusion of the daughter cells. These variations have a direct influence on embryo viability. Consequently, it is key to determine the number of cells in an embryo at any given time. The CLEmbryo correctly manages normal and direct cleavage but has not yet been tested on reverse cleavage. It should be able to handle it since it does not involve the monotonic growth constraint. The frequency of reverse cleavage is not widely documented but may reach up to 25% of transferable embryos in some studies and may reduce the quality of these embryos^{37,38}. Therefore, it deserves to be addressed in future work. This is challenging because reverse cleavage embryos have unstable developmental dynamics; after cell division, the number of cells can increase and then decrease to the initial number due to the fusion of two daughter cells. Consequently, our future work will also involve the extension of our dataset of bovine embryo videos to include examples of such cleavage behavior. More research is also needed to determine the exact timing of cell divisions. A cell division may last several minutes and be documented through successive frames. The transition from one cell stage label to the next between two frames is not sufficient to infer the precise timing of the start and end of cell division. This would require going beyond the cell stage classification process.

References

1. Sharma, A. *et al.* Detecting human embryo cleavage stages using YOLO V5 object detection algorithm. In *Nordic Artificial Intelligence Research and Development (NAIS)* (Oslo, 2022).
2. Nguyen, Q. *et al.* Effects of opening the incubator on morphokinetics in mouse embryos. *Eur. J. Obstet. & Gynecol. Reproductive Biol.* **229**, 64–69 (2018).
3. Magata, F. Time-lapse monitoring technologies for the selection of bovine *in vitro* fertilized embryos with high implantation potential. *J. Reproduction Dev.* **69**, 57–64 (2023).

4. Hachani, Y. *et al.* Early prediction of the transferability of bovine embryos from videomicroscopy. In *International Conference on Image Processing (ICIP)* (Abu Dhabi, 2024).
5. Dirvanauskas, D., Maskeliunas, R., Raudonis, V. & Damasevicius, R. Embryo development stage prediction algorithm for automated time lapse incubators. *Comput. Methods Programs Biomed.* **177**, 161–174 (2019).
6. McAuley, J., Ng, N. H., Lipton, Z. C., Gingold, J. A. & Desai, N. Predicting embryo morphokinetics in videos with late fusion nets & dynamic decoders. In *International Conference on Learning Representations (ICLR)* (Vancouver, 2018).
7. 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 (2019).
8. Malmsten, J., Zaninovic, N., Zhan, Q., Rosenwaks, Z. & Shuan, J. Automated cell division classification in early mouse and human embryos using convolutional neural networks. *Neural Comput. & Appl.* **33**, 2217–2228 (2020).
9. Leahy, B. D. *et al.* Automated measurements of key morphological features of human embryos for IVF. In *International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)* (Lima, 2020).
10. Redmon, J., Divvala, S., Girshick, R. & Farhadi, A. You only look once: Unified, real-time object detection. In *Conference on Computer Vision and Pattern Recognition (CVPR)* (Las Vegas, 2016).
11. Liu, X. *et al.* DLT-Embryo: A dual-branch local feature fusion enhanced transformer for embryo multi-stage classification. *Biomed. Signal Process. Control.* **102**, 107266 (2025).
12. Lockhart, L., Saeedi, P., Au, J. & Havelock, J. Automating embryo development stage detection in time-lapse imaging with synergic loss and temporal learning. In *International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)* (Strasbourg, 2021).
13. Lukyanenko, S. *et al.* Developmental stage classification of embryos using two-stream neural network with linear-chain conditional random field. In *International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)* (Strasbourg, 2021).
14. Gomez, T. *et al.* Towards deep learning-powered IVF: A large public benchmark for morphokinetic parameter prediction (2022).
15. Nguyen, T.-P. *et al.* EmbryosFormer: Deformable transformer and collaborative encoding-decoding for embryos stage development classification. In *Winter Conference on Applications of Computer Vision (WACV)* (Waikoloa, 2023).
16. Zhu, X. *et al.* Deformable DETR: Deformable transformers for end-to-end object detection. In *International Conference on Learning Representations (ICLR)* (Vienna, 2021).
17. Khosla, P. *et al.* Supervised contrastive learning. In *Advances in Neural Information Processing Systems (NeurIPS)* (Virtual site, 2020).
18. Reis, A. D. P. *et al.* Identification and mathematical prediction of different morphokinetic profiles of *in vitro* developed bovine embryos. In *34rd Meeting of the Association of Embryo Transfer in Europe*, hal-02737515 (Nantes, 2018).
19. Lin, T.-Y., Goyal, P., Girshick, R., He, K. & Dollár, P. Focal loss for dense object detection. In *International Conference on Computer Vision (ICCV)* (Venice, 2017).
20. Rahaman, M. M., Millar, E. K. A. & Meijering, E. Histopathology image classification using supervised contrastive deep learning. In *International Symposium on Biomedical Imaging (ISBI)* (Athens, 2024).
21. Lee, S., Lee, Y., Lee, G. & Hwang, S. Supervised contrastive embedding for medical image segmentation. *IEEE Access* **9**, 138403–138414 (2021).
22. Cicconet, M., Gutwein, M., Gunsalus, K. C. & Geiger, D. Label-free cell-tracking and division detection based on 2D time-lapse images for lineage analysis of early embryo development (2014).
23. Tran, D., Wang, H., Torresani, L. & Feiszli, M. Video classification with channel-separated convolutional networks. In *International Conference on Computer Vision (ICCV)* (2019).
24. Berg, P., Uzun, B., Pham, M.-T. & Courty, N. Multimodal supervised contrastive learning in remote sensing downstream tasks. *IEEE Geosci. Remote. Sens. Lett.* **21**, 1–5 (2024). 10.1109/LGRS.2024.3385995.
25. Berntsen, J., Rimestad, J., Lassen, J. T., Tran, D. & Kragh, M. F. Robust and generalizable embryo selection based on artificial intelligence and time-lapse image sequences. *PLOS One* (2022).
26. DeVries, T. & Taylor, G. W. Improved regularization of convolutional neural networks with cutout (2017).

27. Loshchilov, I. & Hutter, F. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)* (New Orleans, 2019).
28. Izmailov, P., Podoprikin, D., Garipov, T., Vetrov, D. & Wilson, A. G. Averaging weights leads to wider optima and better generalization. In *Conference on Uncertainty in Artificial Intelligence (UAI)* (Monterey, 2018).
29. Feichtenhofer, C., Pinz, A. & Wildes, R. P. Spatiotemporal residual networks for video action recognition. In *Advances in Neural Information Processing Systems (NeurIPS)* (Barcelona, 2016).
30. Chen, T., Kornblith, S., Norouzi, M. & Hinton, G. E. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning (ICML)* (Virtual site, 2020).
31. Aissam, D. *et al.* Tailored 3D CT contrastive pretraining to improve pulmonary pathology classification. In *International Conference on Signal Processing (ICSP)* (Beijing, 2022).
32. Hendrycks, D., Lee, K. & Mazeika, M. Using pre-training can improve model robustness and uncertainty. In *International Conference on Machine Learning (ICML)* (Long Beach, 2019).
33. Tsujii, H., Khandoker, M. Y. & ichi Hamano, K. Lipid in mammalian embryo development. *J. Mammalian Ova Res.* **18**, 73–80 (2001).
34. Abe, H., Yamashita, S., Satoh, T. & Hoshi, H. Accumulation of cytoplasmic lipid droplets in bovine embryos and cryotolerance of embryos developed in different culture systems using serum-free or serum-containing media. *Mol. Reproduction Dev.* **61**, 57–66 (2002).
35. Genicot, G., Leroy, J., Soom, A. V. & Donnay, I. The use of a fluorescent dye, Nile red, to evaluate the lipid content of single mammalian oocytes. *Theriogenology* **63**, 1181–1194 (2005).
36. Mänttari, J., Broomé, S., Folkesson, J. & Kjellström, H. Interpreting video features: A comparison of 3D convolutional networks and convolutional LSTM networks. In *Asian Conference on Computer Vision (ACCV)* (Virtual site, 2020).
37. Jin, L., Dong, X., Tan, W. & Huang, B. Incidence, dynamics and recurrences of reverse cleavage in aneuploid, mosaic and euploid blastocysts, and its relationship with embryo quality. *J. Ovarian Res.* **15**, 91 (2022).
38. Sugimura, S., Yao, T., Matoba, S., Yamagata, K. & Imai, K. Morphokinetic prediction of embryo viability in cattle. *Reproduction, Fertility Dev.* **37** (2025).

Acknowledgements

The authors would like to acknowledge the collaboration of Dr. Véronique Duranthon and Brigitte Marquant-LeGuienne for the experimental protocol design for the embryo production. The production of the original embryo data was funded by CRB-Anim. Yasmine Hachani's doctoral fellowship is funded by Inria-INRAE. Operation of the research project is also partly funded by the DIGIT-BIO program of INRAE.

Author contributions statement

Y.H., P.B., E.F., and A.D.P.R. conceived this research work and the experiments. Y.H. implemented the software, ran the experiments. A.D.P.R., S.R and L.L produced the bovine embryos and acquired the videos with the Primovision system. A.D.P.R. annotated the videos. Y.H., P.B., E.F., and A.D.P.R wrote the paper. Y.H. drew the figures.

Additional Information

Competing interests

The authors declare no competing interests.

Approval for animal experiments

This research study was conducted using data available in our laboratory. No live animals or euthanised animals were used to create the original data. The semen was acquired from a commercial company and the cumulus oocyte complexes were harvested from ovaries recovered *post-mortem* in a commercial slaughterhouse. Both these companies and our laboratory are based in France and state-approved. The necessary authorisations for the use of *post-mortem* biological material have been obtained from the responsible Ministry.