PREDICTING EMBRYO MORPHOKINETICS IN VIDEOS WITH LATE FUSION NETS & DYNAMIC DECODERS

Nathan H Ng & Julian McAuley
University of California, San Diego
{nhng, jmcauley}@ucsd.edu

Julian A Gingold & Nina Desai
Cleveland Clinic Foundation
{gingolj, desain}@ccf.org

Zachary C Lipton
Carnegie Mellon University
zlipton@cmu.edu

ABSTRACT

To optimize clinical outcomes many fertility clinics manually annotate time-lapse videos of developing embryos with their corresponding morphokinetics. The process is time-consuming and requires experienced embryologists. We propose late-fusion ConvNets with a dynamic programming-based decoder for automatically labeling these videos. Experiments address a dataset extracted from EmbryoScope incubators at the Cleveland Clinic Foundation Fertility Center. We focus on 6 developmental stages, demonstrating 87% per-frame accuracy.

1 INTRODUCTION

To select embryos for in vitro fertilization (IVF) (Cetinkaya and Kahraman, 2016), doctors typically rank them based on a morphological evaluation. Due to advancements in incubators with built-in time lapse monitoring (EmbryoScope), morphokinetics can be determined non-invasively (Chamayou et al., 2013). These incubators use a high powered microscope to capture images of a developing embryo every 15 minutes. Embryologists perform the analysis, viewing a sequence of photographs and manually annotating the time stamps at which each embryo achieves certain milestones. These scores are combined according to a heuristic formula to rank the embryos by their (believed) viability for transfer into a prepared endometrium.

We investigate deep learning techniques for automating the labeling of these transition times. Specifically, we consider different methods of incorporating the temporal context of the video into the model architecture by fusing frame information early vs later on. We also consider the use of a post-processing decoder which enforces monotonicity constraints.

Our dataset consists of 1309 EmbryoScope time lapse videos extracted from incubators at the Cleveland Clinic Foundation Fertility Center. Each frame is a grayscale 500×500 resolution image with a *well number* in the lower left corner and an hour time marker in the lower right corner, as seen in Figure 1a. The videos in our dataset span 113 different patients, each with up to 12 wells and corresponding videos. Videos begin around 18 hours after fertilization, and end around 140 hours after fertilization. For each video, we accessed manually labeled times at which the embryo was first observed at each developmental stage. These stages are monotonically increasing, meaning that once an embryo reaches a later stage, it cannot regress to an earlier one.

We focus on 6 stages of development for each embryo, cutting off each video at hour 70. The stages include the initial stage (`tStart`), the appearance and breakdown of the male and female pronucleus (`tPNf`), and the appearence of 2 through 4+ cells (`t2, t3, t4, t4+`). The class distribution of frames are 10.3%, 5.3%, 19.4%, 4.5%, 19.8%, and 40.7% in each stage, respectively. We split the dataset into training/validation/test sets by randomly selecting 93/10/10 patients and their respective wells. This gives us a total of 191449 / 21798 / 23348 frames in the respective splits, roughly 81%/9%/10%.

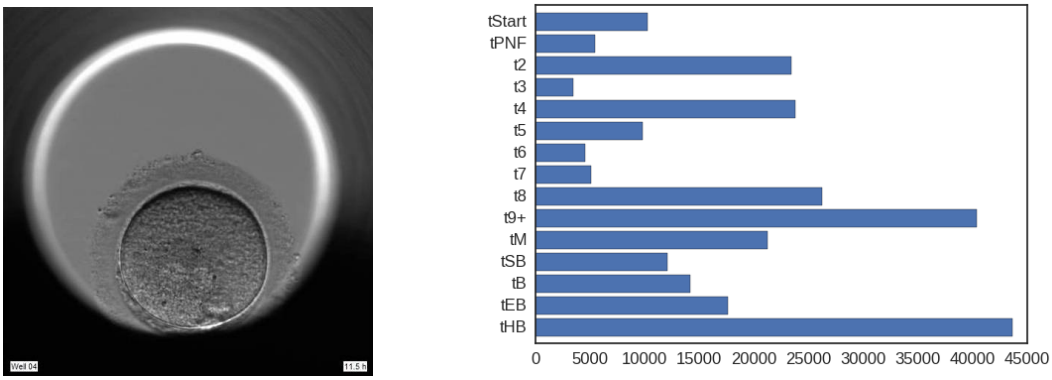

(a) (b)

Figure 1: Left (a) a sample frame from an EmbryoScope (image displayed is publicly available, not private medical data from our set) and right (b) summary statistics on the number of frames assigned to each stage of development in the ground truth human annotations.

2 METHODS

We cast predicting embryo morphokinetics as a multiclass classification problem, where the input is an EmbryoScope video, and the output is a label indicating which stage the embryo is currently in at each frame of the video. Because the frames arrive in sequence, and we have prior knowledge that human embryologists rely on contextual frames (several before, several after) to make their annotations, we investigate several methods of incorporating the context frames and temporal information. Among other techniques, we (i) consider the use of time as a feature and (ii) experiment with various architectures for incorporating context from temporally local frames.

We first consider a baseline of a single frame ResNet(He et al., 2016) architecture to understand how predictive static images are of morphokinetics. We then consider two methods of temporal fusion: *early fusion* and *late fusion* (Karpathy et al., 2014). In early fusion, data is combined temporally at a pixel level. This is accomplished by stacking surrounding frames prior to applying any convolutions and modifying the number of channels in the input of our single frame model. In late fusion, data is combined temporally after the convolutional layers have been applied by concatenating the resulting fully-connected layers. The same single frame model is applied on context frames.

Structured Decoding with Dynamic Programming The ground truth stages must be monotonically non-decreasing. A vanilla CNN does not have this constraint. We impose this inductive bias through a structured decoding approach using dynamic programming to enforce monotonicity of the final predictions. For any given frame, our model outputs a probability distribution p such that p_i is the probability that the frame is in stage i. We define a cost function for making a prediction of stage s for this frame as $E(s, p)$. Optimizing for total error subject to a monotonicity constraint can be solved efficiently using dynamic programming. We investigate two error functions, negative label likelihood and earthmover distance. We define negative label likelihood as $E(s, p) = -\log(p_s)$ and earthmover distance as $E(s, p) = |s - \sum_i p_i(i)|$.

3 EXPERIMENTS

We now present our basic experimental setup. For all experiments we use the same 81%/9%/10% training/validation/test split. We report the per frame accuracy of our raw predictions, as well as the per frame accuracy of our smoothed predictions. We also report the mean absolute error (MAE) and root mean squared error (RMSE) in the frame timing of predicted smoothed transitions.

For all convolutional models we use a ResNet50 architecture pretrained on ImageNet(He et al., 2016). For single and early fusion models we modify the number of input channels to take in varying amounts of temporal context. In our late fusion model we concatenate the flattened fully connected

Table 1: Results

Model	Frames	Raw Acc.	DP: label likelihood s.t. monotonicity			DP: earthmover's distance s.t. monotonicity		
			Accuracy	MAE	RMSE	Accuracy	MAE	RMSE
ResNet50	1	0.8200	0.8460	11.225	29.650	0.8320	11.417	29.580
Early Fusion	3	0.8237	0.8448	10.555	27.370	0.8308	11.014	28.892
Early Fusion	9	0.8252	0.8423	10.927	29.400	0.8287	11.101	28.362
Early Fusion	15	0.8182	0.8456	10.935	27.719	0.8100	12.442	29.465
Early Fusion+time	9	0.8343	0.8430	11.152	26.321	0.8292	11.760	26.173
Early Fusion+time	15	0.8420	0.8446	10.761	26.849	0.8373	10.901	26.429
Late Fusion	15	**0.8479**	**0.8676**	**8.963**	**24.756**	0.8452	10.112	27.254

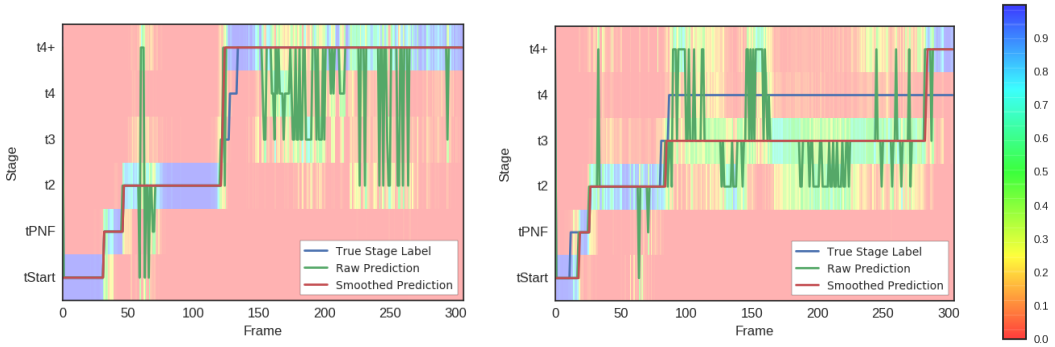

Figure 2: Smoothing reduces erratic changes in prediction caused when the model finds an ambiguous frame. On the left, the smoothed predictions match almost exactly the ground truth. On the right, smoothing makes clear the model's misprediction of stage t3 rather than stage t4.

layers together after applying the same convolutions separately. We also append the timestamp as a feature at the fully connected layer.

All models were trained using stochastic gradient descent with weight decay. We augment our dataset by applying random rotations and flips to images. Model weights are updated on the training set and we choose all non-differentiable hyper-parameters and architecture details based on validation set performance.

4 Related Work

The problem of predicting embryo annotations from time lapse videos has been addressed in the literature by Khan et al.. In that work they use an 8-layer convolutional network to count the number of embryos in an image. To incorporate temporal information they use conditional random fields and similarly use dynamic programming to enforce monotonicity constraints. To our knowledge, our work is the first to use deep learning to predict embryo morphokinetics and to model it using late fusion models and dynamic decoders. The idea of extending models to include temporal information has been explored extensively in recent years. Simonyan and Zisserman (2014) used a two-stream architecture applied to a single frame as well as multi-frame optical flow in order to combine spatial and temporal information. Karpathy et al. (2014) introduced the architecture of early, late, and slow fusion models in order to incorporate temporal information at different points in the model.

References

Caroline Pirkevi Cetinkaya and Semra Kahraman. Morphokinetics of embryos - where are we now? *Journal of Reproductive Biotechnology and Fertility*, pages 1–8, 2016.

Sandrine Chamayou, Pasquale Patrizio, Giorgia Storaci, Venera Tomaselli, Carmelita Alecci, Carmen Ragolia, Claudia Crescenzo, and Antonino Guglielmino. The use of morphokinetic param-

eters to select all embryos with full capacity to implant. *Assisted Reproductive Genetics*, pages 703–710, 2013.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, 2014.

Aisha Khan, Stephen Gould, and Mathieu Salzmann. Deep convolutional neural networks for human embryonic cell counting.

Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. 2014.