A Hierarchical Bayesian Approach for Unsupervised Cell Phenotype Clustering

Mahesh Venkata Krishna and Joachim Denzler

Computer Vision Group, Friedrich Schiller University Jena
Ernst-Abbe-Platz 2, 07743 Jena, Germany
{mahesh.vk,joachim.denzler}@uni-jena.de

Published at GCPR 2014. This is a pre-publish version of the paper. The final version of the paper will appear in the proceedings of the GCPR-2014.

Abstract. We propose a hierarchical Bayesian model - the *wordless Hierarchical Dirichlet Processes-Hidden Markov Model* (wHDP-HMM), to tackle the problem of unsupervised cell phenotype clustering during the mitosis stages. Our model combines the unsupervised clustering capabilities of the HDP model with the temporal modeling aspect of the HMM. Furthermore, to model cell phenotypes effectively, our model uses a variant of the HDP, giving preference to morphology over co-occurrence. This is then used to model individual cell phenotype time series and cluster them according to the stage of mitosis they are in. We evaluate our method using two publicly available time-lapse microscopy video datasets and demonstrate that the performance of our approach is generally better than the state-of-the-art.

Keywords: Hierarchical Bayesian methods; hidden Markov models; cell phenotypes; unsupervised clustering; mitosis phase modeling; time-lapse microscopy

1 Introduction

Machine analysis of time-lapse microscopy videos has become a very important application field of computer vision. Presence of a large amount of objects like cells, microbes etc. in these videos often make manual analysis cumbersome and prone to subjective errors. One of the main tasks for machine learning algorithms in this field is the analysis of time-lapse videos of cell culture, where *mitosis* events are happening.

Based on internal cell dynamics and morphology changes, biologists suggest that there are five main stages of mitosis: *prophase, prometaphase, metaphase, anaphase* and *telophase* ([10]: §18.6, pages 849-851). The stage between the mitosis events is called the *interphase*. These various stages of cell life cycle are shown in Fig. 1. Given a video, the problem is to classify each cell in it according to the stage of mitosis it is in.

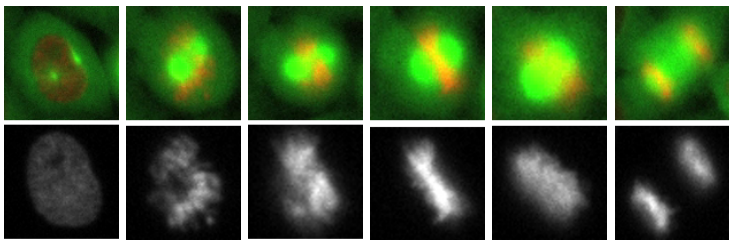

Fig. 1. The six main stages of cell life cycle: (a) *Interphase*, (b) *Prophase*, (c) *Prometaphase*, (d) *Metaphase*, (e) *Anaphase* and (f) *Telophase*. The second row in (a)-(f) show the corresponding phenotypes of the nuclei, shown in orange in the first row.

In biology research labs, often, different types of cells are analyzed, using different dyes and illumination methods. This makes the problem challenging, as a learning system trained on one type of cell with certain dye and illumination will not work well in other situations. To retrain the algorithm, an expert has to painstakingly label a new sequence and this is again time consuming, and undesirable.

Using cell population analysis tools such as the one in [2] to perform this task does not often suffice. Studying the temporal dynamics of cell phenotypes at a single cell level across a recording is an important aspect of the state-of-the-art biology research [4].

Thus, from a time-lapse microscopy video analysis system, we desire the following three properties:

1. it should be able segregate multiple stages of the cell life cycle.
2. it should be unsupervised.
3. it should model the cell phenotypes in tracks of single cells extracted from across the video.

For unsupervised clustering tasks, *Hierarchical Dirichlet Processes* (HDP) and their variants have been used before on tasks ranging from text analysis [13] to Traffic Scene Analysis [15]. Since the traditional HDP models lack the ability to handle temporal information, using HDP to provide prior distribution for a *Hidden Markov Model* (HMM) was proposed in [13]. This HMM will model the temporal changes in cell morphologies, resulting in unsupervised clustering. To model distances in feature spaces better compared to standard topic models, Rematas *et al.* [12], proposed a Kernel Density Estimate (KDE) based scheme for LDA models. Here, the dictionary-of-visual-words representation of standard HDP models is replaced with kernel densities, making the approach *wordless*. We combine the above ideas to formulate a *wordless Hierarchical Dirichlet Processes - Hidden Markov Model* (wHDP-HMM), derive the inference procedure for the model, and develop an unsupervised method based on it to cluster cells in a time-lapse video according to the stage of life cycle they are in.

This paper is arranged as follows. In Sect. 2 we provide a brief overview of the existing literature on the problem. Sect. 3 discusses our temporal HDP model, and inference procedure. Experiments conducted on two publicly available data-sets are described in Sect. 4, along with results and discussions. Sect. 5 covers some concluding remarks and ideas for future work.

2 Previous Works

Due to its immense potential in aiding biology research, the problem of cell life cycle modeling has been extensively dealt with in the past. The approach of Yang *et al.* [16] performs the tasks of segmentation, tracking and mitosis detection. They use level-set methods for segmentation and tracking and use image attributes like circularity, area, average intensity etc. to classify cells under mitosis. While they report good results, it is to be noted that their approach is only effective in detecting late anaphase/telophase. Furthermore, each different sequence has to be separately analyzed and parameters over attributes need to be readjusted.

Supervised cell life stage classification problem has been dealt with previously with considerable success. Online Support Vector Classifiers are used in [14], continuously retraining the model to accommodate changing experimental conditions. But they do not consider temporal information that can be an important influence on the performance. Liu *et al.* [9] and Huh *et al.* [6] use Hidden Conditional Random Fields (HCRF) and their variants to perform mitosis detection on four stages of mitosis. Harder *et al.* [4] performed mitosis stage classification through a finite state machine (FSM), also accounting for abnormal shapes. The authors use cell tracks to construct time series and traverse a FSM for each track. Thus, they prevent biologically impossible results and improve the performance. Further in the time series-based methods, the HMM-based approaches of Held *et al.* [5] and Gallardo *et al.* [3] were demonstrated to perform very well. In their approach, a HMM was trained with the mitosis stages as the hidden states of the Markov model. We use this idea in our model (*cf.* Sect. 3).

The recent *Temporally Constrained Combinatorial Clustering* (TC3) scheme of [17] tackles a scenario similar to ours, and reports the best results on a publicly available data-set, according to our knowledge. Hence we use this in our performance analysis. The TC3 scheme is a combinatorial clustering scheme with biological causality constraints like "no cell goes back to the previous state of its life". Similarly to the approaches of [4, 5], they segment the cells in each frame, extract various features such as shape, intensity etc., and construct multiple time series by tracking each cell using a nearest neighbor tracker. They then use the TC3 stage followed by Gaussian Mixture Model Clustering and then by a HMM stage to improve performance by correcting errors.

Whereas the approach of Zhong *et al.* [17] performs quite well, the clustering stages consider each individual cell time series and does not take into account the clustering in other time series. However, information regarding the causal progression of life stages in cell time series can be shared among the time series to

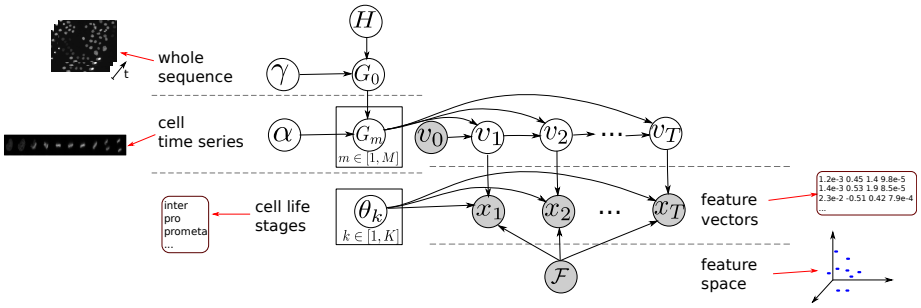

Fig. 2. Our wHDP-HMM model, with correspondences of the model components to data representation for microscopy videos.

improve the performance. This leads us to apply HDP models, as they inherently share information among various data groupings.

Unlike the various methods discussed above, our model has three important aspects, namely, (i) unsupervised clustering, (ii) temporal modeling and (iii) sharing information among different time series.

3 Wordless HDP-HMM

Figure 2 shows our wHDP-HMM model. It is to be noted that, unlike the HDP-HMM models of [8, 13], we limit our HDP-HMM model to single HMM. This is due to the fact that biologically, causal ordering in cell life cycle is fixed and one HMM is sufficient to model it.

3.1 Modeling Principles

We now describe the generative process within our model. As shown in Fig. 2, the topmost level is the whole sequence, which is the complete set of cell tracks, represented by a Dirichlet Process (DP), G_0. This DP is parametrized by the concentration parameter γ and the base distribution H. (a more detailed description of HDP is presented in [13]). The base distribution is set to be a Dirichlet distribution, with a hyper-parameter ζ. In the standard HDP terms, this level corresponds to the *data corpus*.

From the DP G_0, M number of cell tracks are sampled. And each cell time series is associated with a DP G_m ($m \in [1, M]$), with α as the concentration parameter and G_0 as base distribution. These correspond to *documents* in HDP models.

And these cell time series DPs provide priors for the HMM hidden states $v_{t,m}$ ($t \in [1, T]$, $m \in [1, M]$). These state variables v_t are indicator variables which denote which of the six states θ_k most likely generated the corresponding feature vectors x_t. The features are seen as originating from the feature space $\mathcal{F}$, which is a deviation from the usual HDP models where the feature words are

seen as samples from a dictionary with multinomial distribution for sampling. In the following, we omit the time series index m for variables v and x for lucidity. In HDP terms, v correspond to *topic mixtures* and x to *words*.

We have the following generative model:

$$\begin{aligned} G_0 \,|\, \gamma, H &\sim DP(\gamma, H) \\ G_m \,|\, \alpha, G_0 &\sim DP(\alpha, G_0) \quad \text{for } m \in [1, M] \end{aligned} \tag{1}$$

And for each time series m,

$$\begin{aligned} v_t \,|\, v_{t-1}, G_m &\sim G_m \quad \text{for } t \in [1, T] \\ x_t \,|\, v_t, \theta_k, \mathcal{F} &\sim F(\theta_{v_t}) \quad \text{where } k \in [1, K] \end{aligned} \tag{2}$$

Where $F(\cdot)$ is the prior feature distribution given the state. In line with the biological considerations as explained in Sect. 1, we set the number of states K for our problem of mitosis stage modeling to six.

3.2 Inference

Given the observed data x_t for all cell time series, to estimate the hidden states v_t, we perform Bayesian inference using the Markov Chain Monte Carlo procedure, specifically the Gibbs sampling scheme. Here, we iterate over the conditional distribution of the hidden states given the previous state and the feature vectors. This distribution can be factored as follows.

$$\begin{aligned} P(v_t = k \,|\, v_{t-1}, x_t) \propto \\ P(x_t \,|\, v_t = k, v_{t-1}, x_{1,\ldots,(t-1),(t+1),\ldots,T}) P(v_t = k \,|\, v_{t-1}) \end{aligned} \tag{3}$$

Here, the first term in right hand side is the probability of word distribution, and following the idea of [12], we evaluate it through Gibbs iteration as follows.

$$P(x_t \,|\, v_t = k, v_{t-1}, x_{1,\ldots,(t-1),(t+1),\ldots,T}) \propto \frac{1}{|\Phi_{k,t' \neq t}|} \sum_{t'}^{T} \Phi_{k,t'} \cdot K(x_t, x_{t'}) \tag{4}$$

where $K(\cdot)$ is the kernel function and Φ is the $|\mathcal{F}| \times K$ matrix of features and their corresponding topic allocations.

The second term in the right hand side of (3), $P(v_t \,|\, v_{t-1})$ is the transition probability of the HMM. This is evaluated by first initializing it to α and then iteratively evaluating

$$P^i(v_t = k \,|\, v_{t-1}) \propto \frac{P^{i-1}(v_t = k \,|\, v_{t-1}) + \alpha \cdot \frac{P^{i-1}(v_0 = k)}{\sum_k P^{i-1}(v_0) + \gamma}}{\sum_{k' \neq k} P^{i-1}(v_t = k' \,|\, v_{t-1} = k) + \alpha} \tag{5}$$

where i represents the iteration index. The derivation of (5) follows similarly to the one of Infinite HMM discussed in [1].

For the overall inference task, we alternate between the Gibbs sampling and transition probability iteration at each step. The hyper-parameters α and γ are also sampled, whose details are discussed in [1, 13].

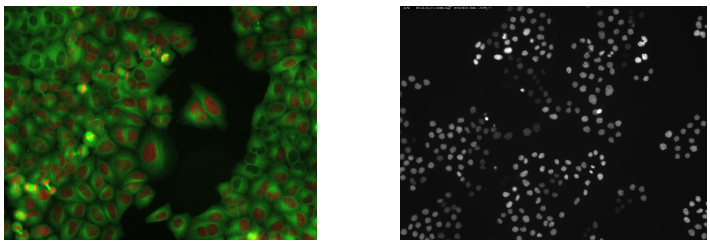

(a) *cellcognition* (b) *mitocheck*

Fig. 3. Example frames from the two data-sets used for experimental evaluation of our method.

4 Experiments and Results

4.1 Data-sets

The *cellcognition* data-set The *cellcognition* data-set from [5, 17] is a part of the *Cellcognition*[1] project. An example frame can be seen in Fig. 3(a). The reference data consists of 7 sequences of RNAi treated human HeLa Kyoto cells expressing fluorescent H2B-mCherry (orange) and α-tubulin (green). The frame resolution is 1392×1040 pixels. The data-set, in total, contains 363,120 individual cell objects, segmented and tracked using the *Cellcognition* framework. There are 9,078 cell division events in the whole data-set.

The data-set comes with ground-truth regarding the life cycle stage of each segmented cell. Furthermore, the authors provide 257-dimensional feature vectors they extracted for each cell. These features consist of various intensity, shape, and texture features. We use these features in our experiments to effectively compare our method with respect to theirs.

The *mitocheck* data-set The *mitocheck* data-set of [11] is a part of the *mitocheck* project[2]. Overall, the database contains more than 129,500 video sequences, each of approximately 32 seconds, recorded at 2 fps. The frame resolution is 1344×1024 pixels. Since the data-set is extremely large, to restrict the task of evaluation to reasonable limits without giving up significantly on generality, we randomly selected a subset of 100 sequences from them, each containing on an average 64 frames.

In the publicly available data-set videos, frames contain the mCherry expressing nuclear spindles, *i.e.* the microtubule structures formed during mitosis. Thus, the task is to segment the spindles, track and extract them, and to apply our wHDP-HMM to cluster them.

Segmentation for this data-set is performed using the *ilastik*[3] framework and tracking is done using the two-stage graph optimization method of Jiang *et al.* [7].

[1] http://www.cellcognition.org
[2] http://www.mitocheck.org
[3] http://www.ilastik.org

Since the ground-truth for the data-set has not been made public, we have manually marked the ground-truth. The ground-truth was marked by two non-experts independently and conflicts in markings were resolved in a second iteration by one of them. Owing to the fact that the ground-truth markings were not done by biology experts, class labels were limited to the clearly defined *telophase*. Thus, evaluation on this data-set is limited to two classes, *telophase* and *non-telophase*. The data-set contains 994,163 cell objects (returned by the segmentation step) and 16,491 cell division instances.

4.2 Experimental Set-up and Evaluation

We design our experiments in order to demonstrate our various modeling choices such as the use of KDEs as opposed to the standard co-occurrence statistics in HDP models. In the following, we describe the experimental set-up in each case.

Standard HDP The standard HDP model is used to analyze the video. As the model lacks the capability to handle temporal information, the representation of the data is changed for this case. Here, cells are detected from each frame and features for them are extracted. These features are then used to build a dictionary with 512 visual words, using a Gaussian Mixture Model clustering algorithm. These visual words from each frame form one document for the HDP.

Once we have this grouped representation, we analyze the data through Bayesian inference, with topics representing cell states. For our experiments, we set the three hyper-parameters to 0.5. These hyperparameters are resampled based on the data later and initial values do not affect the results by a large margin.

HDP-HMM Following the standard HDP, we test the HDP-HMM algorithm presented in [8,13]. Here, we analyze the cell time series data from tracking and make use of the temporal information. In Gibbs sampling, we use the co-occurrence statistics as described in [13]. For this, we construct a dictionary as with the standard HDP model. Thus, we see the effects of not using the cell morphology information explicitly in modeling, and only implicitly through dictionary construction. In this experiment, we set the hyper-parameter values as with the standard HDP.

wHDP-HMM Here, we use the cell time series information as for the HDP-HMM case. However, unlike before, we now use the kernel matrix of the features extracted in stead of the co-occurrence statistics, as discussed in Sect. 3. For our experiments, we used the Gaussian Radial Basis Function (RBF) kernel, with parameter $\sigma = 1.5$, determined through cross-validation.

Evaluation Metrics For evaluation, we follow the evaluation schemes of [17]. Thus, if TP represents true positives, TN- true negative, FP- false positive, and FN- false negative, then,

$$\begin{aligned} \text{Precision} &= \frac{TP}{TP+FP}, \quad \text{Recall} = \frac{TP}{TP+FN} \\ \text{F-score} &= \frac{2 \cdot (\text{Precision} \times \text{Recall})}{\text{Precision} + \text{Recall}} \end{aligned} \tag{6}$$

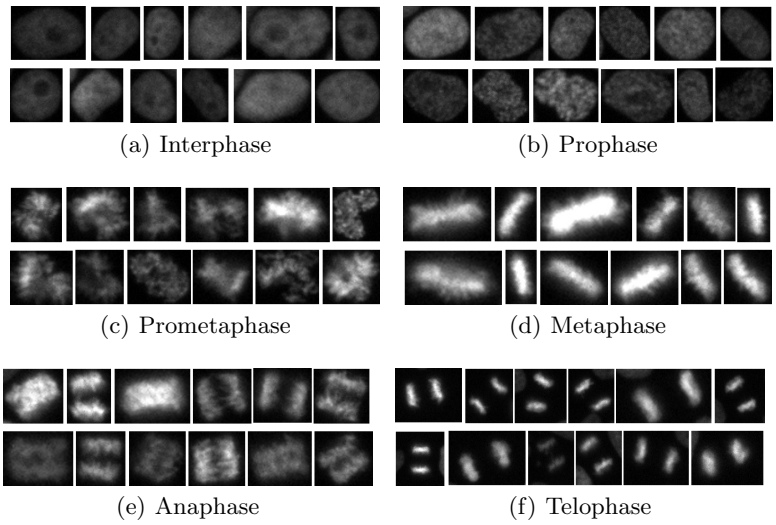

Fig. 4. Unsupervised clustering results using our wHDP-HMM model for the *cellcognition* data-set.

For both the data-sets, we average and report the results over all the sequences involved, and additionally report the maximum deviation from it.

4.3 Results and Discussion

Results for the *cellcognition* Data-set

For the experiment involving HDP, HDP-HMM and our wHDP-HMM models, we ran 1500 iterations each of the Gibbs sampler. Our unoptimized MATLAB® code for the wHDP-HMM model, the whole process took 20 minutes per video of approximately 200 frames, with precalculated feature kernels. In comparison, the method of [17] took 15 minutes with precomputed feature data.

Figure 4 shows some unsupervised clustering results. It shows some nuclear spindle images clustered into each group by our method. As can be seen, the algorithm preforms reasonably well in separating the cell life stages. Due to their closeness in morphology and temporal ordering, *anaphase* and *telophase* are often confused. This is reflected in the quantitative results as well, as we shall see in the following.

Table 1 shows the quantitative results for the *cellcognition* data-set. For [17], we use the results reported by the authors. As can be expected, owing to its being the first stage and morphologically separated from the rest of the stages, performance for the *interphase* is better than the rest. The performances for the following three phases are quite similar to one another. The last two stages, *anaphase* and *telophase* have relatively worse performance, as their morphological closeness results in increased confusion.

Table 1. Results for various experiments over the *cellcognition* data-set. The numbers represent the mean values over 7 sequences and maximum variations. " [17] TC3+" implies the results of the TC3+GMM+DiscreteHMM approach of Zhong et al., as reported in the cited paper.

Precision

Algorithm	Inter	Pro	Prometa	Meta	Ana	Telo
[17] TC3+	**95.97±0.83**	83.53 ± 2.07	**91.47±2.45**	**96.82±0.92**	80.57 ± 7.67	84.57 ± 5.28
HDP	94.72 ± 1.47	74.03 ± 1.79	69.12 ± 1.05	78.23 ± 2.02	58.02 ± 6.52	51.21 ± 3.49
HDP-HMM	93.45 ± 1.03	79.25 ± 1.93	76.24 ± 1.63	81.02 ± 1.95	72.12 ± 5.67	67.02 ± 2.89
wHDP-HMM	93.11 ± 0.98	**85.43±1.47**	88.94 ± 2.85	92.25 ± 1.37	**81.44±6.73**	**86.74±3.28**

Recall

Algorithm	Inter	Pro	Prometa	Meta	Ana	Telo
[17] TC3+	99.51 ± 0.32	82.75 ± 4.13	84.43 ± 2.96	88.24 ± 3.63	80.22 ± 6.24	79.50 ± 5.09
HDP	96.13 ± 0.75	71.32 ± 2.94	69.03 ± 1.89	64.57 ± 4.31	66.18 ± 4.21	64.14 ± 6.12
HDP-HMM	98.78 ± 1.26	80.37 ± 3.12	81.52 ± 2.41	83.31 ± 4.10	79.05 ± 4.69	69.31 ± 5.64
wHDP-HMM	**99.62±0.94**	**85.01±3.25**	**88.81±1.79**	**90.02±2.72**	**82.33±5.82**	**81.12±4.62**

F-score

Algorithm	Inter	Pro	Prometa	Meta	Ana	Telo
[17] TC3+	**97.69±0.36**	82.84 ± 2.62	87.64 ± 2.35	**92.05±2.00**	80.03 ± 6.79	81.51 ± 4.70
HDP	95.14 ± 0.99	72.36 ± 2.27	68.93 ± 1.35	70.53 ± 2.95	61.71 ± 5.26	58.08 ± 4.72
HDP-HMM	95.86 ± 1.12	79.55 ± 2.35	78.43 ± 1.94	82.04 ± 2.71	75.12 ± 4.93	68.05 ± 4.26
wHDP-HMM	96.03 ± 0.95	**84.89±2.48**	**88.43±2.03**	90.92 ± 1.98	**81.65±6.02**	**83.57±3.63**

Results for the *mitocheck* Data-set

Similar to the previous data-set, for the experiment involving HDP, HDP-HMM and our wHDP-HMM models, we ran 1500 iterations each of the Gibbs sampler, and one round of execution time for our wHDP-HMM was approximately 13 minutes for a video of 64 frames.

Figure 5 shows the unsupervised clustering for the *mitocheck* data-set. Again, it can be noted that there is greater confusion between *anaphase* and *telophase* than among other phases.

Table 2 shows the quantitative results for the *mitocheck* data-set. For [17], we use the program made public by the authors and use parameter settings used by them. Performance is measured only with respect to the *telophase*, since the ground-truth only involved the this stage. As can be seen, our method performs well compared to other schemes. Nevertheless, the added confusion in the Markov chain shows that low frame rates in video recordings of time-lapse microscopy can have detrimental effects on the performance of cell tracking based approaches like ours, due to gaps in time series.

5 Conclusions

We presented the wHDP-HMM model and as discussed Sect. 1, used it to perform the task of unsupervised mitosis stage modeling in time-lapse microscopy videos, represented in terms of temporal tracks. In our model, HDP provides

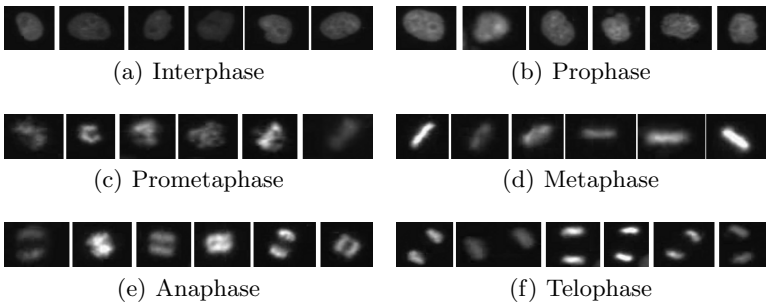

Fig. 5. Unsupervised clustering results using our wHDP-HMM model for the *mitocheck* data-set.

Table 2. Results for various experiments over the *mitocheck* data-set. The numbers represent the mean values over 100 sequences and maximum variations. " [17] TC3+" implies the results of the TC3+GMM+DiscreteHMM approach of Zhong et al., obtainable from the code made public by the authors.

Method	Precision	Recall	F-Score
[17] TC3+	83.19 ± 4.47	76.29 ± 3.52	79.41 ± 3.83
HDP	51.03 ± 3.63	62.86 ± 5.45	56.11 ± 4.45
HDP-HMM	66.47 ± 5.26	68.13 ± 5.23	67.04 ± 5.24
wHDP-HMM	$\mathbf{84.52 \pm 3.79}$	$\mathbf{79.48 \pm 4.71}$	$\mathbf{81.62 \pm 4.02}$

prior distributions for the HMM, making the system unsupervised and able to handle temporal information. To directly handle cell phenotypes, we replace the co-occurrence based sampling scheme of standard HDP- HMM models with one based on kernel density estimates.

We demonstrated the performance of our method using two publicly available data-sets: *cellcognition* and *mitocheck*. The results compared favorably with the state-of-the-art.

The Bayesian inference step is a time-consuming operation and in deployment scenarios, maybe undesirable. To handle this situation, one can use a two-stage approach, combining our generative model for training and a discriminative classifier for testing. Furthermore, it will be of immense practical use if it is possible to extend the model to jointly perform tracking and clustering. Another interesting idea for future research is to extend the model to allow a certain degree of user interaction, so that an expert user can provide a few inputs to improve system performance.

Acknowledgments The authors gratefully acknowledge financial support by ZEISS and would like to thank Christian Wojek and Stefan Saur (ZEISS Corporate Research and Technology) for helpful discussions and suggestions.

References

1. Beal, M.J., Ghahramani, Z., Rasmussen, C.E.: The infinite hidden markov model. In: Advances in Neural Information Processing Systems (NIPS). pp. 577–584 (2002)
2. Carpenter, A.E., Jones, T.R., Lamprecht, M.R., Clarke, C., Kang, I.H., Friman, O., Guertin, D.A., Chang, J.H., Lindquist, R.A., Moffat, J., Golland, P., Sabatini, D.M.: Cellprofiler: image analysis software for identifying and quantifying cell phenotypes. Genome Biology 7(10), R100 (2006)
3. Gallardo, G.M., Yang, F., Ianzini, F., Mackey, M., Sonka, M.: Mitotic cell recognition with hidden markov models. In: Proceedings of SPIE. vol. 5367, pp. 661–668 (2004)
4. Harder, N., Mora-Bermúdez, F., Godinez, W.J., Wünsche, A., Eils, R., Ellenberg, J., Rohr, K.: Automatic analysis of dividing cells in live cell movies to detect mitotic delays and correlate phenotypes in time. Genome Research 19(11), 2113–2124 (2009)
5. Held, M., Schmitz, M.H.A., Fischer, B., Walter, T., Neumann, B., Olma, M.H., Peter, M., Ellenberg, J., Gerlich, D.W.: Cellcognition: time-resolved phenotype annotation in high-throughput live cell imaging. Nature Methods 7(9), 747–754 (Sep 2010)
6. Huh, S., Chen, M.: Detection of mitosis within a stem cell population of high cell confluence in phase-contrast microscopy images. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1033 –1040 (2011)
7. Jiang, X., Haase, D., Körner, M., Bothe, W., Denzler, J.: Accurate 3d multi-marker tracking in x-ray cardiac sequences using a two-stage graph modeling approach. In: Computer Analysis of Images and Patterns. Lecture Notes in Computer Science, vol. 8048, pp. 117–125 (2013)
8. Kuettel, D., Breitenstein, M.D., Gool, L.V., Ferrari, V.: What is Going on? Discovering SpatioTemporal Dependencies in Dynamic Scenes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2010)
9. Liu, A.A., Li, K., Kanade, T.: Mitosis sequence detection using hidden conditional random fields. In: IEEE International Symposium on Biomedical Imaging: From Nano to Macro. pp. 580–583 (April 2010)
10. Lodish, H., Berk, A., Kaiser, C.A., Krieger, M., Bretscher, A., Ploegh, H., Amon, A., Scott, M.P.: Molecular Cell Biology. W.H.Freeman & Co Ltd, 7th international edition edn. (2013)
11. Neumann, B., Walter, T., Heriche, J.K., Bulkescher, J., Erfle, H., Conrad, C., Rogers, P., Poser, I., Held, M., Liebel, U., Cetin, C., Sieckmann, F., Pau, G., Kabbe, R., Wuensche, A., Satagopam, V., Schmitz, M.H.A., Chapuis, C., Gerlich, D.W., Schneider, R., Eils, R., Huber, W., Peters, J.M., Hyman, A.A., Durbin, R., Pepperkok, R., Ellenberg, J.: Phenotypic profiling of the human genome by time-lapse microscopy reveals cell division genes. Nature 464(7289), 721–727 (2010)
12. Rematas, K., Leuven, K., Fritz, M., Tuytelaars, T.: Kernel density topic models: Visual topics without visual words. In: Modern Non Parametric Methods in Machine Learning, NIPS Workshop (2012)
13. Teh, Y., Jordan, M., Beal, M., Blei, D.: Hierarchical Dirichlet Processes. Journal of the American Statistical Association pp. 1566–1581 (2006)
14. Wang, M., Zhou, X., Li, F., Huckins, J., King, R.W., Wong, S.T.: Novel cell segmentation and online svm for cell cycle phase identification in automated microscopy. Bioinformatics 24(1), 94–101 (2008)

15. Wang, X., Ma, X., Grimson, W.: Unsupervised activity perception in crowded and complicated scenes using hierarchical bayesian models. IEEE Transactions on Pattern Analysis and Machine Intelligence 31(3), 539 –555 (march 2009)
16. Yang, F., Mackey, M., Ianzini, F., Gallardo, G., Sonka, M.: Cell segmentation, tracking, and mitosis detection using temporal context. In: Medical Image Computing and Computer-Assisted Intervention MICCAI, pp. 302–309 (2005)
17. Zhong, Q., Busetto, A.G., Fededa, J.P., Buhmann, J.M., Gerlich, D.W.: Unsupervised modeling of cell morphology dynamics for time-lapse microscopy. Nature Methods 9(7), 711–713 (Jul 2012)