Understanding the Success of Semi-Supervised Learning: A Case Study of Mitotic Phase Classification Using Raman Imaging

Jurica Levatić[1], Marjan Stoimchev[1,2], Csaba Vörös[3], David Bauer[3], Peter Horvath[3], and Sašo Džeroski[1]

[1] Jožef Stefan Institute, Ljubljana, Slovenia
{firstname.lastname}@ijs.si
[2] Jožef Stefan International Postgraduate School, Ljubljana, Slovenia
[3] Biological Research Centre, Szeged, Hungary
voroscsaba.kanizsa@gmail.com, {lastname.firstname}@brc.hu

Abstract. Semi-supervised learning (SSL) offers the potential to improve predictive performance by exploiting unlabeled data alongside limited labeled examples.While SSL has demonstrated success in many applications, it is not guaranteed to outperform traditional supervised learning and can, in some cases, achieve worse predictive performance. The success of SSL methods relies on certain assumptions about the relationship between the descriptive attributes and corresponding labels, such as low-density separation and smoothness assumptions. Despite the risk of performance degradation, the conditions under which SSL performs well (including the validity of SSL assumptions in real-world data) are poorly understood and seldom investigated.

In this study, we investigate the viability of semi-supervised learning for mitotic phase classification using Raman spectroscopy data. Determining the mitotic phase of a cell has numerous important applications in biology, medicine, and other fields. This task is also particularly well-suited for semi-supervised learning, as obtaining labels involves laborious analysis of microscopy images by experts. We evaluate two SSL approaches: semi-supervised predictive clustering trees (SSL-PCTs) and semi-supervised masked autoencoders, and compare their performance against supervised baselines across varying amounts of labeled data.

Our initial findings indicate that these methods generally fail to outperform supervised methods. We assess the low-density separation and smoothness assumptions using inter- and intra-class distances and local class homogeneity, finding that our spectroscopy data mostly violate these assumptions. In experiments limited to biologically distinct mitotic phases that comply with SSL assumptions, SSL-PCTs outperformed supervised ones demonstrating the critical role of data structure in determining the success of semi-supervised methods.

Keywords: Semi-supervised learning · Raman imaging · Cell mitosis · Predictive Clustering Trees · Masked Autoencoders

1 Introduction

Traditional supervised machine learning algorithms rely solely on labeled data for model training. Consequently, to achieve satisfactory predictive performance, they often need a large number of labeled training examples. However, in many real-world domains, labeled data is scarce due to the expensive and/or time-consuming labeling procedure, which can severely limit the performance of supervised methods. To address this limitation, semi-supervised learning (SSL) techniques have emerged [1]. These methods use both labeled and unlabeled data to enhance predictive performance, based on the premise that unlabeled data is typically abundant and easily available, offering an opportunity to improve model training.

A wide variety of SSL algorithms and techniques have been developed, and SSL has been applied across diverse tasks [5]. These developments have led to some striking successes: In certain cases, state-of-the-art SSL methods achieve predictive performance nearly on par with fully-supervised learning even when only a fraction of the dataset is labeled [12]. However, despite its promise, SSL is not guaranteed to yield improvements in every situation. Researchers have long observed that incorporating unlabeled data can sometimes fail to help and may degrade a model's predictive performance [8, 5, 12].

Such negative outcomes can be expected when the assumptions of SSL algorithms are violated. SSL methods typically rely on some of the core SSL assumptions [1]: the *smoothness* assumption (if two samples x and x' are close in the descriptive space, their labels y and y' should be the same); the *low-density* separation assumption (the decision boundary should lie in a low-density region of the input space); and the *manifold* assumption (data points lying on the same low-dimensional manifold should have the same label).

Despite the risk of performance degradation in semi-supervised learning, the conditions under which SSL is effective are poorly understood and rarely investigated. Most studies focus on developing new SSL methods or demonstrating improvements on specific benchmarks, often assuming that key conditions (such as the smoothness or low-density separation assumptions) hold true. Very few works analyze negative outcomes of SSL in light of these assumptions, examine when they are valid, or propose approaches for such analysis. One likely reason for this is publication bias against negative findings [18].

In this study, we examine a case study particularly well-suited for semi-supervised learning: mitotic phase classification using Raman imaging (see Section 2 for details). We use this task to probe the effects of the low-density separation and smoothness assumptions on the performance of two recent SSL methods: semi-supervised predictive clustering trees [9] and semi-supervised masked autoencoders [14]. Initial results demonstrate that these methods largely fail to outperform their supervised counterparts. In further analysis, we show that the low-density separation and smoothness assumptions can be empirically assessed using inter- and intra- class distances and local class homogeneity, revealing that our spectroscopy data largely violates these assumptions. In subsequent experiments targeting biologically distinct mitotic phases that better comply with SSL

assumptions, semi-supervised predictive clustering trees consistently outperform their supervised counterparts.

This study underscores the importance of understanding how data structure aligns with the inner mechanics of SSL methods and proposes a way to estimate the compliance of data with the key SSL assumptions.

2 Case Study: Mitotic Phase Classification Using Raman Imaging

During mitosis, a cell undergoes five phases (pro-, prometa-, meta-, ana-, and telophase) during which the DNA of the cell's nucleus is split into two equal sets of chromosomes, creating two daughter cells [11]. Understanding the mitotic phase of a cell can be very useful for various biological, medical, and other applications. For instance, in cancer research, identifying cells in different stages of mitosis can help determine the aggressiveness of tumors and guide targeted therapies [4]. In agriculture, understanding meiosis (i.e., a process similar to mitosis that produces four non-identical daughter cells) is important for plant breeding, as it allows for the manipulation of plant cells to produce desirable traits [2].

Determining the mitotic phase of a cell typically requires analysis of microscope images by a human expert. The most common technique for obtaining these images is fluorescence microscopy, which involves staining cells with fluorescent dyes to visualize cellular structures and organelles. However, these dyes can be cytotoxic, potentially damaging the sample and confounding measurements [3]. A non-invasive, in vivo alternative for analyzing cell cycle dynamics is Raman spectroscopy, which does not require staining. This technique works by shining a laser onto a sample and measuring inelastic light scattering to provide insights into the molecular structure and composition of cells [15, 16].

Due to the large number of cells that can be present on a single microscope slide and the similarity between mitotic and normal nuclei, annotating cells with their mitotic phases remains a challenging task for humans. Consequently, machine learning systems have been developed to assist experts in mitosis classification [10]. However, due to the aforementioned laborious annotation procedures, the labeled training examples for mitotic phase classification are limited. This calls for machine learning approaches that can operate effectively in regimes with limited availability of labeled data, such as semi-supervised methods.

In this work, we investigate the viability of semi-supervised machine learning methods for classifying mitotic phases based on Raman spectroscopy data. Specifically, we consider semi-supervised predictive clustering trees and random forests [9], which were designed for tabular data, as well as semi-supervised deep neural networks based on masked autoencoders, which were designed specifically for imaging data [14].

3 Methods

For mitotic phase classification from Raman imaging data, we consider two recent state-of-the-art semi-supervised methods: semi-supervised predictive clustering trees (SSL-PCT) [9] and semi-supervised masked autoencoders (SSL-MAE) [14]. Below, we briefly describe these methods, while more details can be found in the original publications.

3.1 Semi-supervised Predictive Clustering Trees

Predictive clustering trees (PCTs) build upon traditional decision trees, enabling predictions of more complex and structured outputs. In PCTs, decision trees are considered as a hierarchy of clusters, with the root cluster containing all the data. As tree growth progresses from root to leaves, the data is recursively split into smaller clusters. The algorithm used to construct the tree evaluates potential splits based on a variance function. In the semi-supervised variant of PCTs, the variance function accounts for both descriptive and target attributes, allowing the model to use both labeled and unlabeled data:

$$Var_f = w \cdot Var_f(Y) + (1 - w) \cdot Var_f(X), \tag{1}$$

where $w \in [0, 1]$ controls the relative contributions of the target space (Y) and the descriptive space (X) to the variance function Var_f (with only the descriptive space available for unlabeled examples). The w parameter is crucial for the flexibility of the model, as it enables learning across a range from fully supervised ($w = 1$) to fully unsupervised ($w = 0$) scenarios. By adjusting the importance of unlabeled examples through the weight w, the level of supervision can be tailored for each dataset, helping to mitigate the risk of performance degradation when using unlabeled data. Semi-supervised random forests are built using SSL-PCTs as base learners. Note that

Semi-supervised PCTs and random forests are implemented in the freely available CLUSplus framework[4] [13].

3.2 Semi-supervised masked autoencoders

SSL-MAE is built on the Masked Autoencoders (MAE) framework, a popular self-supervised method that employs masked image modeling (MIM) as a pretext task for learning representations [17]. Unlike traditional two-stage self-supervised learning methods, SSL-MAE integrates both self-supervised and discriminative learning in a unified, end-to-end architecture. An overview of SSL-MAE is provided in Figure 1. The self-supervised component, inspired by SimMIM [17], uses unlabeled data for representation learning. Simultaneously, a discriminative classifier in the model extracts supervision from the limited labeled data. Both components share a common encoder, and the model is optimized with

[4] https://github.com/knowledge-technologies/clus

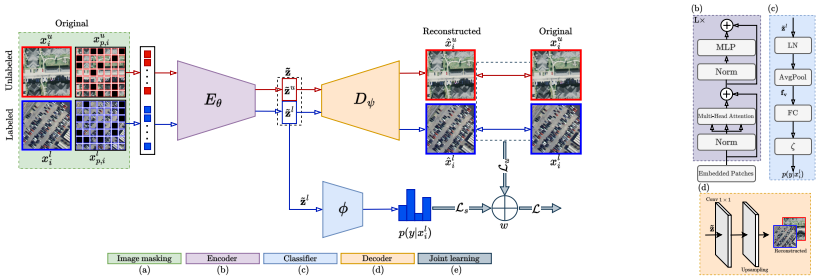

Fig. 1. Overview of SSL-MAE (image adapted from [14]). In (a), labeled (x^l) and unlabeled (x^u) images are divided into patches, and a fraction of these patches is randomly masked. In (b), a vision transformer encoder processes both labeled and unlabeled patches to produce latent representations. The supervised objective $\mathcal{L}_s$ in (c) is computed solely on labeled (and unmasked) patches for classification. Meanwhile, the unsupervised objective $\mathcal{L}_u$ in (d) reconstructs the masked regions of both labeled and unlabeled images via an ℓ_1 pixel-regression loss. In (e), the two losses are combined through a weight w, which regulates the degree of supervision during training. During inference, only the encoder (b) and classification head (c) are used on full (unmasked) test images.

a joint learning objective aimed at enhancing the discriminative power of the learned features, ultimately improving prediction accuracy. Similarly as in SSL-PCTs, a mechanism is incorporated into the joint loss to regulate the supervision level and provide flexibility during the learning process, resulting in the following loss function of SSL-MAE:

$$\mathcal{L} = w \cdot \mathcal{L}_s + (1 - w) \cdot \mathcal{L}_u, \tag{2}$$

where $\mathcal{L}_s$ and $\mathcal{L}_u$ are the supervised and the unsupervised loss terms, respectively, and $w \in [0, 1]$ is a weight that balances the amount of supervision to the overall learning objective.

4 Data

A dataset comprising of 298 cells with annotated mitotic phases was obtained from the Synthetic and Systems Biology Unit, Biological Research Centre, Szeged, Hungary. Below we provide a short description of the data, while more details can be found in the paper by Voros et al. [16].

Henrietta Lacks (HeLa) cells were cultured on glass slides and incubated for 24h. The nuclei of the cells were stained and widefield fluorescence images were obtained using a Leica TCS SP8 confocal laser scanning microscope. The mitotic cells were manually identified by an expert using the fluorescence microscope and annotated with five mitotic phases (pro-, prometa-, meta-, ana-, and telophase) or an interphase (i.e., a phase prior to mitosis). The final dataset consists of 53,

29, 46, 97, 19, and 54 cells annotated as inter-, pro-, prometa-, meta-, ana-, and telophase, respectively.

The positions of cells were transferred to a NT-MDT Spectrum Instruments Ntegra II Raman–AFM microscope system using registration points and a custom software transformation. Raman imaging was performed with a high-resolution system, capturing spectra from single cells. The Raman spectra were collected over a range from 439 cm^{-1} to 3228 cm^{-1}, with each spectrum acquired in 1.2 seconds. The spatial resolution was 1 µm, and each spectral map covered an area of 24 µm × 24 µm, which was then cropped to exclude neighboring cells and minimize background interference.

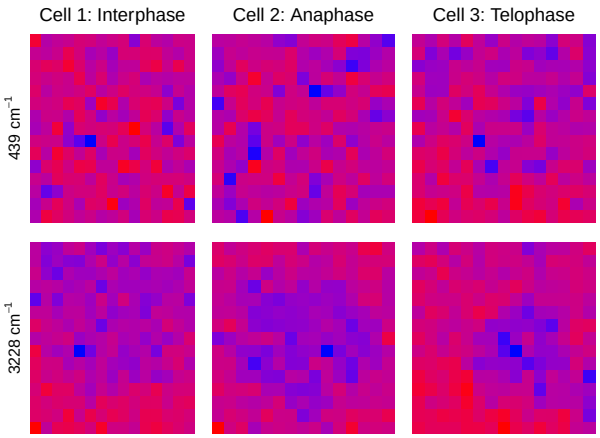

Fig. 2. Examples of Raman images of three cells in different phases. The first row shows images at the lowest wavenumber, while the second row shows images at the highest wavenumber considered.

The data was then preprocessed by removing cosmic rays, normalizing intensity, and fitting the spectra to a reference spectrum using a custom algorithm (for details, see Voros et al. [16]). Each cell was imaged with 1600 different wavenumber, in NxM spatial positions, where N and M vary between 15 and 24. To make the spatial dimensions consistent, we resized each grid to 15x15 using linear. The examples of Raman images are shown in Figure 2.

Since SSL-PCTs are designed for tabular data, they cannot directly process 3D imaging data (with dimensions 1600x15x15). As a result, we flattened the data for each cell into a 1D vector of length 360,000. To address the issue of high dimensionality, we applied Principal Component Analysis (PCA) and retained the first 263 principal components, which account for 95% of the variance in the data. In contrast, the SSL-MAE method was directly fed with the 3D imaging data, preserving the spatial information.

5 Experimental design

Mitotic phase prediction was formulated as a multi-class classification problem with 6 classes: inter-, pro-, prometa-, meta-, ana-, and telophase. We were primarily interested in investigating if unlabeled data can help semi-supervised methods outperform the classical supervised methods. To this end, we compare semi-supervised predictive clustering trees (SSL-PCT), semi-supervised random forests (SSL-RF), and semi-supervised masked autoencoders (SSL-MAE) to their supervised counterparts, SL-PCT, SL-RF, and SL-MAE, respectively.

To assess the impact of unlabeled data on classification performance, we used different amounts of labeled data: 5%, 10%, 25%, and 50%. The unlabeled data was extracted from the original labeled dataset by randomly selecting examples and temporarily removing their labels. Algorithm performance was evaluated on a test set consisting of these unlabeled examples, with their labels restored. The random sampling process was repeated five times with different initializations, and the final performance metrics were reported as the average across the five runs.

The weight parameter w for SSL-PCTs and SSL-RFs was optimized via internal 5-fold cross validation on the labeled part of the training set while using the available unlabeled data. The weight parameter of SSL-MAE was set to 0.5, since this value provides generally good performance [14]. SSL-PCTs and SSL-RFs are, on the other hand, more sensitive to the w parameter, hence it was optimized. The supervised counterparts of these methods were achieved by (i) supplying the methods only with the labeled data and (ii) setting the w parameter to 1, effectively cancelling out the contribution of the descriptive attributes in variance calculation for SSL-PCTs and unsupervised loss for SSL-MAE. The SSL-PCTs were pruned with the M5 pruning technique, while SSL-RFs were built with 100 unpruned trees. SSL-MAE was trained for 100 epochs by using the AdamW optimizer with a base learning rate of 1×10^{-3} using a *cosine* learning rate scheduler, a mini-batch of 64, and a mixed precision of training to speed up the training process and reduce memory consumption.

The performance of the algorithms was evaluated using accuracy and one-miss accuracy. The latter metric considers a prediction successful if the predicted class is at most one phase away from the actual mitotic phase. This metric was chosen because consecutive mitotic phases are often similar, and most misclassifications occur between adjacent phases.

6 Results

6.1 Predictive performance of the methods

The results of the experimental evaluation are presented in Table 1. In most cases, the semi-supervised algorithms failed to outperform their supervised counterparts, with one exception: for 5% labeled data, SSL-PCT outperformed SL-PCT and achieved the best overall accuracy for that fraction. However, in terms of one-miss accuracy, which allows for a single mitotic phase deviation, none

of the semi-supervised methods improved over supervised learning. Overall, supervised random forests tend to perform best with 25% or more labeled data, while for smaller fractions, the best performance is divided between SSL-PCT and SL-MAE. The SL-MAE and SSL-MAE methods showed some instability, with deteriorating accuracy with 50% of the labeled data.

Table 1. Comparison of predictive performance of supervised and semi-supervised methods across different fractions of labeled data. For each fraction, the best result is marked in bold.

	Accuracy				One-miss Accuracy			
Method	5%	10%	25%	50%	5%	10%	25%	50%
SL-MAE	0.208	**0.317**	0.317	0.150	**0.542**	**0.533**	0.533	0.567
SSL-MAE	0.142	0.208	0.317	0.250	0.358	0.483	0.533	0.408
SL-PCT	0.206	0.236	0.251	0.339	0.466	0.452	0.500	0.572
SSL-PCT	**0.214**	0.229	0.231	0.291	0.477	0.473	0.536	0.592
SL-RF	0.200	0.265	**0.364**	**0.395**	0.505	0.497	**0.583**	**0.632**
SSL-RF	0.193	0.261	0.334	0.360	0.493	0.485	0.550	0.607

These results suggest that, in the context of the data at hand, semi-supervised learning may not be suitable for mitotic phase classification using Raman spectroscopy. A possible reason is the small size of our dataset, which contains only 298 examples. Consequently, there are few labeled examples to learn from, and the pool of unlabeled examples is also limited. Additionally, the overall accuracy across all methods is low, suggesting that Raman spectroscopy may not provide sufficient information for accurate mitotic phase prediction in this case. Raman spectroscopy is known to be noisy particularly when applied to biological samples [6]. If the spectral data does not contain clear discriminative features related to mitotic phases, both supervised and semi-supervised models may struggle to extract meaningful information from the data. Additionally, SSL algorithms rely heavily on the assumptions that unlabeled data can provide useful information to reinforce the decision boundary. The two basic assumptions of SSL are that data points with similar features should have similar labels (i.e., the smoothness assumption) or that different classes are well-separated (i.e., the low-density separation assumption). However, if the spectral differences between mitotic phases are subtle, these assumptions may not hold due to overlapping spectral features and the additional unlabeled data may contribute little or even introduce noise.

6.2 Analysis of SSL assumptions

In an attempt to shed more light on the performance of the methods, we investigate whether the previously mentioned SSL assumptions – namely, the smoothness and the low-density separation assumptions – are satisfied in our data. To this end, we calculated intra- and inter-class Euclidean distances on principal components used as input features for PCTs (Figure 3A). If the low-density

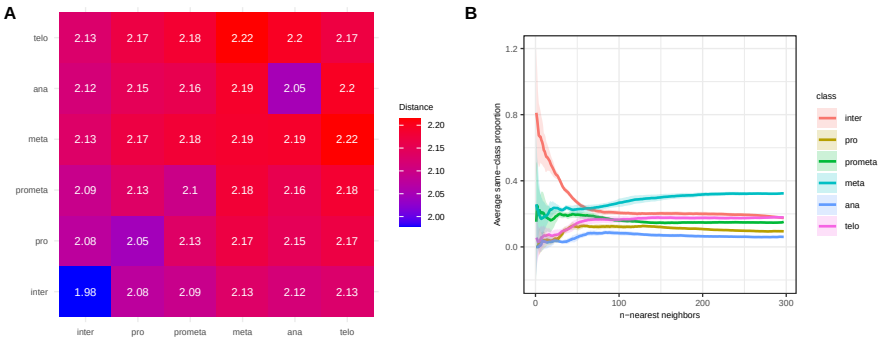

Fig. 3. A) Average intra- and inter-class distances. The input features were normalized before calculation of the distances. B) Proportion of samples of the same class within each sample's n-nearest neighborhood ($\forall n \in \{1, \ldots, 298\}$), averaged across all samples of the same class for each n. Shaded areas represent the standard deviation

separation assumption holds, we would expect smaller intra-class distances (indicating greater similarity among samples of the same class) and larger inter-class distances (indicating dissimilarity among samples of different class). However, our results show that this is only partially true. Most notably, the interphase is the most compact class, exhibiting the greatest difference between intra- and inter-class distances. This observation is consistent with biological expectation, as interphase – preceding mitosis – is not part of the mitotic phases. To some degree, similarly, anaphase appears markedly more compact than the other mitotic phases, aligning with its distinct biologocal characteristics. Anaphase is considered the most distinct mitotic phase due to dramatic structural changes in the cell: Sister chromatids are actively pulled apart to opposite poles of the cell, which is visually distinct event compared to the other phases and it is the only phase where chromosomes move rapidly towards poles, driven by shortening microtubules [7].

The other mitotic phases are harder to distinguish, having similar intra- and inter-class distances, suggesting violation of the SSL low-density separation assumption. More specifically, the two consecutive phases, pro and prometa, are similar to each other (and somewhat also to the interphase) making it more difficult for classifiers to establish firm decision boundaries. In these phases chromatin begins to condense and the nucleolus fades, which is a gradual rather than abrupt transition, which can explain the observed similarity between these phases. The meta and telo phases clearly lack distinct boundaries to other classes.

We next explore the validity of the smoothness assumption. For this purpose, we calculate the local class homogeneity, i.e., the proportion of samples of the same class within each sample's n-nearest neighborhood, where size of the neighborhood varies from 1 to 298 (i.e., the total number of samples). For each n, the proportions are averaged across all samples of the same class (Figure 3B). If the smoothness assumption holds, we would expect to observe mostly sam-

ples of the same class among close neighbors, followed by a gradual decrease in the same-class proportion as we move to more distant neighbours. Remarkably, this behaviour is observed only for interphase (Figure 3B), which we identified as the most compact among the classes, with the smallest intra-class distances. The other classes clearly violate the smoothness assumption, as the same-class proportion only slightly increases in closer neighbourhoods, or even increases in more distant neighbourhoods, as seen in metaphase.

This analysis provides evidence that our Raman spectroscopy data largely does not comply with the SSL assumptions, which may explain the observed underperformance of the SSL methods (Table 1). However, certain classes, anaphase and especially interphase, may be more suitable for SSL.

6.3 Predictive performance with better class separation

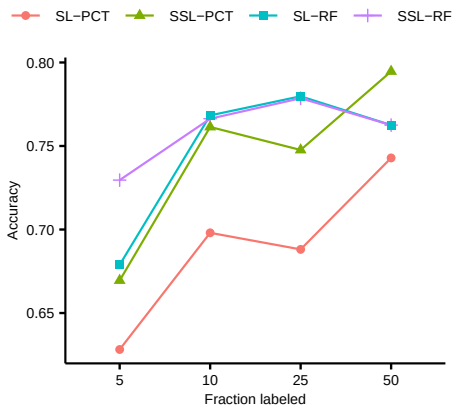

Fig. 4. Predictive performance in classifying interphase versus all other phases. Samples belonging to pro and prometa classes were removed to enhance class separability.

To test whether semi-supervised methods perform more favourably when data complies with SSL assumptions and has better class separability, we conduct experiments to classify interphase versus all other phases. To further enhance class separability, we remove the pro and prometa classes, which resemble interphase (Figure 3A). For this experiment, we consider only semi-supervised PCTs and random forests, as they demonstrated more stable performance than SSL-MAE. Additionally, the analysis of SSL assumptions was based on principal components, which are only used as input features by PCTs.

The results demonstrate that, under these conditions, SSL-PCTs outperform their supervised counterpart by a large margin across all fractions of labeled data (Figure 4). SSL-RF, on the other hand, outperformed supervised random forests only for 5% of labeled data, while for other fractions the performance is

very similar to the one of SL-RF. This aligns with previous findings that semi-supervised random forests are typically most beneficial in scenarios with few labeled examples [9].

These results underscore the crucial role of the interplay between data structure and SSL assumptions in determining the performance of SSL methods. When classes are well separated, unlabeled data has a better chance of enhancing decision boundaries and improving predictive performance. On the other hand, in situations with fuzzy class boundaries, unlabeled data may degrade predictive performance.

7 Conclusion

In this study, we explored the utility of semi-supervised learning (SSL) for mitotic phase prediction from Raman spectroscopy data. Determining a cell's mitotic phase has numerous biological, medical, and research applications, while Raman spectroscopy offers a non-invasive, in-vivo imaging technique. In the initial experiments, semi-supervised predictive clustering trees and semi-supervised masked autoencoders largely failed to outperform their supervised counterparts. In further analysis, we show that the low-density separation and smoothness assumptions can be assessed using inter- and intra- class distances and local class homogeneity. These metrics reveal that our spectroscopy data violates the key SSL assumptions – low density and smoothness assumptions – which are essential for SSL performance.

Consecutive mitotic phases closely resemble each other, making it difficult to establish clear decision boundaries that can be reinforced by unlabeled data. Notably, the biologically distinct interphase, which precedes mitosis, was reflected in the spectroscopy data and exhibited better class separation than other phases. Semi-supervised models developed specifically for classifying interphase achieved superior predictive performance compared to supervised models, highlighting the crucial role of data structure in the success of semi-supervised methods.

This study is clearly limited by the small dataset size and the narrow selection of semi-supervised learning methods. Future research should more systematically investigate the interplay between data structure, SSL assumptions, and method performance – using larger datasets and a broader range of SSL techniques – to develop data-driven, verifiable guidelines for the effective application of semi-supervised learning.

Acknowledgments. We acknowledge the support of the Slovenian Research and Innovation Agency via the research projects N2-0236, GC-0001, and J7-4636.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Chapelle, O., Schölkopf, B., Zien, A.: Semi-supervised Learning. MIT Press (2006)

2. Chen, L., Wang, K., Wang, C.: Meiosis in plants: from understanding to manipulation. New Crops p. 100055 (2024)
3. Dixit, R., Cyr, R.: Cell damage and reactive oxygen species production induced by fluorescence microscopy: effect on mitosis and guidelines for non-invasive fluorescence microscopy. The Plant Journal **36**(2), 280–290 (2003)
4. Dominguez-Brauer, C., Thu, K.L., Mason, J.M., Blaser, H., Bray, M.R., Mak, T.W.: Targeting mitosis in cancer: emerging strategies. Molecular cell **60**(4), 524–536 (2015)
5. van Engelen, J.E., Hoos, H.H.: A survey on semi-supervised learning. Machine Learning **109**(2), 373–440 (2020). https://doi.org/10.1007/s10994-019-05855-6, https://doi.org/10.1007/s10994-019-05855-6
6. Fang, S., Wu, S., Chen, Z., He, C., Lin, L.L., Ye, J.: Recent progress and applications of raman spectrum denoising algorithms in chemical and biological analyses: A review. TrAC Trends in Analytical Chemistry **172**, 117578 (2024)
7. de Gramont, A., Cohen-Fix, O.: The many phases of anaphase. Trends in biochemical sciences **30**(10), 559–568 (2005)
8. Guo, Y., Niu, X., Zhang, H.: An extensive empirical study on semi-supervised learning. In: 2010 IEEE International Conference on Data Mining. pp. 186–195. IEEE (2010)
9. Levatić, J., Ceci, M., Kocev, D., Džeroski, S.: Semi-supervised predictive clustering trees for (hierarchical) multi-label classification. International Journal of Intelligent Systems **2024**(1), 5610291 (2024)
10. Nofallah, S., Mehta, S., Mercan, E., Knezevich, S., May, C.J., Weaver, D., Witten, D., Elmore, J.G., Shapiro, L.: Machine learning techniques for mitoses classification. Computerized Medical Imaging and Graphics **87**, 101832 (2021)
11. O'Connor, C.: Cell division: Stages of mitosis. Nature Education **1**(1), 188 (2008)
12. Oliver, A., Odena, A., Raffel, C.A., Cubuk, E.D., Goodfellow, I.: Realistic evaluation of deep semi-supervised learning algorithms. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) Advances in Neural Information Processing Systems. vol. 31. Curran Associates, Inc. (2018)
13. Petković, M., Levatić, J., Kocev, D., Breskvar, M., Džeroski, S.: Clusplus: A decision tree-based framework for predicting structured outputs. SoftwareX **24**, 101526 (2023)
14. Stoimchev, M., Levatić, J., Kocev, D., Džeroski, S.: Ssl-mae: Adaptive semisupervised learning framework for multilabel classification of remote sensing images using masked autoencoders. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing **18**, 14882–14896 (2025)
15. Swain, R.J., Jell, G., Stevens, M.M.: Non-invasive analysis of cell cycle dynamics in single living cells with raman micro-spectroscopy. Journal of cellular biochemistry **104**(4), 1427–1438 (2008)
16. Voros, C., Bauer, D., Migh, E., Grexa, I., Végh, A.G., Szalontai, B., Castellani, G., Danka, T., Dzeroski, S., Koos, K., et al.: Correlative fluorescence and raman microscopy to define mitotic stages at the single-cell level: Opportunities and limitations in the ai era. Biosensors **13**(2), 187 (2023)
17. Xie, Z., Zhang, Z., Cao, Y., Lin, Y., Bao, J., Yao, Z., Dai, Q., Hu, H.: Simmim: A simple framework for masked image modeling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9653–9663 (June 2022)
18. Zhu, X.: Semi-supervised learning literature survey. Tech. rep., Computer Sciences, University of Wisconsin-Madison (2008)