



# Using Whole Slide Image Representations from Self-supervised Contrastive Learning for Melanoma Concordance Regression

Sean Grullon<sup>1</sup>(✉), Vaughn Spurrier<sup>1</sup>, Jiayi Zhao<sup>1</sup>, Corey Chivers<sup>1</sup>, Yang Jiang<sup>1</sup>, Kiran Motaparathi<sup>2</sup>, Jason Lee<sup>3</sup>, Michael Bonham<sup>1</sup>, and Julianna Ianni<sup>1</sup>

<sup>1</sup> Proscia, Inc., Philadelphia, USA

[sean.grullon@proscia.com](mailto:sean.grullon@proscia.com)

<sup>2</sup> Department of Dermatology, University of Florida College of Medicine, Gainesville, USA

<sup>3</sup> Department of Dermatology, Sidney Kimmel Medical College at Thomas Jefferson University, Philadelphia, USA

**Abstract.** Although melanoma occurs more rarely than several other skin cancers, patients' long term survival rate is extremely low if the diagnosis is missed. Diagnosis is complicated by a high discordance rate among pathologists when distinguishing between melanoma and benign melanocytic lesions. A tool that provides potential concordance information to healthcare providers could help inform diagnostic, prognostic, and therapeutic decision-making for challenging melanoma cases. We present a melanoma concordance regression deep learning model capable of predicting the concordance rate of invasive melanoma or melanoma in-situ from digitized Whole Slide Images (WSIs). The salient features corresponding to melanoma concordance were learned in a self-supervised manner with the contrastive learning method, SimCLR. We trained a SimCLR feature extractor with 83,356 WSI tiles randomly sampled from 10,895 specimens originating from four distinct pathology labs. We trained a separate melanoma concordance regression model on 990 specimens with available concordance ground truth annotations from three pathology labs and tested the model on 211 specimens. We achieved a Root Mean Squared Error (RMSE) of  $0.28 \pm 0.01$  on the test set. We also investigated the performance of using the predicted concordance rate as a malignancy classifier, and achieved a precision and recall of  $0.85 \pm 0.05$  and  $0.61 \pm 0.06$ , respectively, on the test set. These results are an important first step for building an artificial intelligence (AI) system capable of predicting the results of consulting a panel of experts and delivering a score based on the degree to which the experts would agree on a particular diagnosis. Such a system could be used to suggest additional testing or other action such as ordering additional stains or genetic tests.

**Keywords:** Self supervised learning · Contrastive learning · Melanoma · Weak supervision · Multiple instance learning · Digital pathology

## 1 Introduction

More than 5 million diagnoses of skin cancer are made each year in the United States, about 106,000 of which are melanoma of the skin [1]. Diagnosis requires microscopic examination of hematoxylin and eosin (H&E) stained, paraffin wax embedded biopsies of skin lesion specimens on glass slides. These slides can be manually observed under a microscope, or digitally on a Whole Slide Image (WSI) scanned on specialty hardware. The 5-year survival rate of patients with metastatic malignant melanoma is less than 20% [15]. Melanoma occurs more rarely than several other types of skin cancer, and its diagnosis is challenging, as evidenced by a high discordance rate among pathologists when distinguishing between melanoma and benign melanocytic lesions (~40% discordance rate; e.g. [7, 10]). The high discordance rate highlights that greater scrutiny is likely needed to arrive at an accurate melanoma diagnosis, however patients receive diagnoses only from a single dermatopathologist in many instances. This tends to increase the probability of misdiagnosis, where frequent over-diagnosis of melanocytic lesions results in severe costs to a clinical practice and additional costs and distress to patients. [26]. In this scenario, the decision-making of the single expert would be further informed by knowledge of a likely concordance level among a group of multiple experts in a given case under consideration. Additional methods of providing concordance information to healthcare providers could help further inform diagnostic, prognostic, and therapeutic decision-making for challenging melanoma cases. A method capable of predicting the results of consulting a panel of experts and delivering a score based on the degree to which the experts would agree on a particular diagnosis would help reduce melanoma misdiagnosis and subsequently improve patient care.

The advent of digital pathology has brought the revolution in machine learning and artificial intelligence to bear on a variety of tasks common to pathology labs. Campanella *et al.* [2] trained a model in a weakly-supervised framework that did not require pixel-level annotations to classify prostate cancer and validated on ~10,000 WSIs sourced from multiple countries. This represented a considerable advancement towards a system capable of use in clinical practice for prostate cancer. However, some degree of human-in-the-loop curation was performed on their data set, including manual quality control such as post-hoc removal of slides with pen ink from the study. Pantanowitz *et al.* [17] used pixel-wise annotations to develop a model trained on ~550 WSIs that distinguishes high-grade from low-grade prostate cancer. In dermatopathology, the model developed in [22] classified skin lesion specimens between six morphology-based groups (including melanoma), was tested on ~5099 WSIs, provided automated quality control to remove WSI patches with pen ink or blur, and also demonstrated that use of confidence thresholding could provide a high accuracy.

The recent application of deep learning to digital pathology has predominately leveraged the use of pre-trained WSI tile representations, usually obtained by using feature extractors pre-trained on the ImageNet [6] data set. The features learned by such pre-training is dominated by features present in natural-scene images, which are not guaranteed to generalize to histopathology images. Such

representations can limit the reported performance metrics and affect model robustness. It has been shown in [13] that self-supervised pre-training on WSIs improved the downstream performance in identifying metastatic breast cancer.

In this work, we present a deep learning regression model capable of predicting from WSIs the concordance rate of consulting a panel of experts on rendering a case diagnosis of invasive melanoma or melanoma in-situ. The deep learning model learns meaningful feature representations from WSIs through self-supervised pre-training, which are used to learn the concordance rate through weakly-supervised training.

## 2 Methods

### 2.1 Data Collection and Characteristics

The melanoma concordance regression model was trained and evaluated on 1,412 specimens (consisting of 1,722 WSIs) from three distinct pathology labs. The first lab consists for 611 suspected melanoma specimens from a leading dermatopathology lab in a top academic medical center (Department of Dermatology at University of Florida College of Medicine), denoted as *University of Florida*. The second lab consisted of 605 suspected melanoma specimens distributed across North America, but re-scanned at The Department of Dermatology at University of Florida College of Medicine, denoted as *Florida - External*. The third lab consisted of 319 suspected specimens from Jefferson Dermatopathology Center, Department of Dermatology & Cutaneous Biology, Thomas Jefferson University denoted as *Jefferson*. The WSIs consisted exclusively of H&E-stained, formalin-fixed, paraffin-embedded dermatopathology tissue and were all scanned using a 3DHistech P250 High Capacity Slide Scanner at an objective power of 20X, corresponding to  $0.24\ \mu\text{m}/\text{pixel}$ . The diagnostic categories present in our data set are summarized in Table 1.

The annotations for our data set were provided by at least three board-certified pathologists who reviewed each melanocytic specimen. The first review was the original specimen diagnosis made via glass slide examination under a microscope. At least two and up to four additional dermatopathologists independently reviewed and rendered a diagnosis digitally for each melanocytic specimen. The patient's year of birth and gender were provided with each specimen upon review. Two dermatopathologists from the United States reviewed all 1,412 specimens in our data set and up to two additional dermatopathologists reviewed a subset of our data set. A summary of the number of concordant reviews in this study is given in Table 2.

The concordance reviews are converted to a concordance rate by calculating the fraction of dermatopathologists who rendered a diagnosis of melanoma in-situ or invasive melanoma. The concordance rate runs from 0.0 (No dermatopathologist rendered a melanoma in-situ/invasive melanoma diagnosis) to 1.0 (all dermatopathologists rendered a melanoma in-situ/invasive melanoma diagnosis). It has been previously noted [20] that the concordance rate itself is correlated with the likelihood that a specimen is malignant. The concordant labels present in

**Table 1.** Specimen counts of each of the pathologies in the data set, broken-out into specific diagnostic categories.

Diagnostic morphology	Number of specimens
Melanoma In Situ	607
Invasive Melanoma	306
Mild-to-moderate Dysplastic Nevus	209
Conventional Melanocytic Nevus	123
Severe Dysplastic Nevus	49
Spitz Nevus with Spindle Cell Morphology	43
Spitz Nevus	36
Junctional Nevus	27
Dermal Nevus	23
Blue Nevus	20
Halo Nevus	11
Total	1412

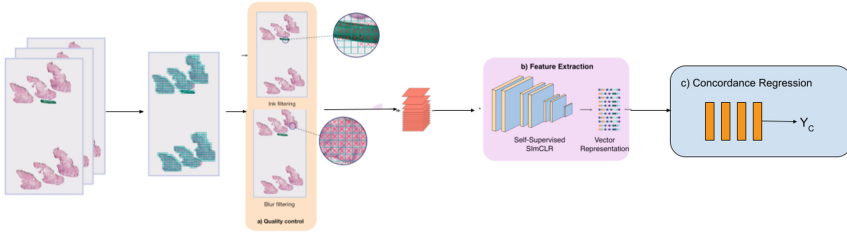
our data set include 0.0, 0.25, 0.33, 0.5, 0.67, 0.75, and 1.0. For training, validating, and testing, we divided this data set into three partitions by sampling at random without replacement with 70% of specimens used for training, and 15% used for each of validation and testing. 990 specimens were used for training, 211 specimens were used for validation, and 211 specimens were used for testing.

**Table 2.** Number of concordant reviews in our data set.

Number of dermatopathologists	Number of specimens
Three	687
Four	216
Five	509
Total	1412

## 2.2 Melanoma Concordance Regression Deep Learning Architecture

The Melanoma Concordance Regression deep learning pipeline consists of three main components: quality control, feature extraction and concordance regression. A diagram of the pipeline is shown in Fig. 1. Each specimen was first segmented into tissue-containing regions through Otsu’s method [16], subdivided into  $128 \times 128$  pixel tiles, and extracted at an objective power of 10X. Each tile was passed through the quality control and feature extraction components of the pipeline.



**Fig. 1.** The stages of the Melanoma Concordance Regression pipeline are: Quality Control, Feature Extraction, and Regression. All single specimen WSIs were first passed through the tiling stage, then the quality control stage consisting of ink and blur filtering. The filtered tiles were passed through the feature extraction stage consisting of a self-supervised SimCLR network pre-trained on WSIs with a ResNet50 backbone to obtain embedded vectors. Finally, the vectors were propagated through the regression stage consisting of fully connected layers to obtain a concordance prediction.

**Quality Control.** Quality control consisted of ink and blur filtering. Pen ink is common in labs migrating their workload from glass slides to WSIs where the location of possible malignancy was marked. This pen ink represented a biased distractor signal in training that is highly correlated with melanoma. Tiles containing pen ink were identified by a weakly supervised model trained to detect inked slides. These tiles were removed from the training and validation data and before inference on the test set. We also sought to remove areas of the image that were out of focus due to scanning errors by setting a threshold on the variance of the Laplacian over each tile [18, 19].

**Self-supervised Feature Extraction.** The next component of the Melanoma Concordance Regression pipeline extracted informative features from the quality controlled tiles. To capture higher-level features in these tiles, we trained a self-supervised feature extractor based on the contrastive learning method proposed in [3, 4] known as SimCLR. SimCLR relies on maximizing agreement in the latent space between representations of two augmented views of the same image. In particular, we maximized agreement between two augmented views of  $128 \times 128$  pixel tiles in our data set.

In order to capture as much variety in real-world WSIs as possible, we trained a dermatopathology feature extractor neural network with the SimCLR contrastive learning strategy using skin specimens from four labs and three

distinct scanners. The skin specimens originated from both from a sequentially-accessioned workflow from these labs as well as a curated data set to capture the morphological diversity of different skin pathologies. The curated data set included various types of basal and squamous cell carcinomas, benign to moderately atypical melanocytic nevi, atypical melanocytic nevi, melanoma in-situ, and invasive melanoma. We note that the feature extractor training set consisted of wider variety of skin pathologies than the concordance regression model, which was trained and tested only on the skin pathologies outlined in Table 1. We included WSIs from the University of Florida and Jefferson that were scanned with 3D Histech P250 scanners. We also included WSIs from another top medical center, the Department of Pathology and Laboratory Medicine at Cedars-Sinai Medical Center, which were scanned with a Ventana DP 200 scanner. We finally included WSIs from an undisclosed partner lab in western Europe that were scanned with a Hamamatsu NanoZoomer XR scanner. We note that WSIs from Cedars-Sinai and the undisclosed partner lab were only used for training the feature extractor and not the concordance regression model, as concordance review annotations were not available for these labs.

We randomly sampled 26,209 tiles from Florida and 57,147 tiles from the remaining three labs for a total of 83,356 tiles randomly sampled from 10,895 specimens for use during training of the feature extraction network. Each tile was sampled from WSIs extracted at an objective power of 10x. We set the temperature hyperparameter, ( $\tau$ ) to  $\tau = 0.1$ , batch size = 128, and the learning rate = 0.001 during training. (We note that the feature extractor is trained separately from the Melanoma Concordance Regression model described in Sect. 2.2.) We randomly divided 80% of the tiles for training, and 20% for validation. We used the ResNet50 [11] backbone for training, and the Normalized Temperature-scaled Cross Entropy (NT-Xent) loss function as in [3]. We followed the same augmentation strategies in [4] applied to the tiles. We used a temperature hyperparameter value of 0.1. We achieved a minimum NT-Xent loss value on the validation set of 2.6. As a point of comparison, we note that the minimum validation NT-Xent loss when training with the ImageNet data set with a ResNet50 backbone is 4.4 [8]. However, given that tiles from WSIs are vastly different from landscape images in ImageNet, this may not be meaningful. Once the feature extractor was trained, it was deployed to the Melanoma Concordance Regression pipeline in order to embed each tile from the melanoma consensus regression data set (consisting of 1,722 WSIs) into a latent space of 2048-channel vectors.

**Melanoma Concordance Regression.** The melanoma concordance regression model predicted a value representing the fraction of dermatopathologists that is concordant with a diagnosis of Melanoma In-Situ or Invasive Melanoma. The model consisted of four fully-connected layers (two layers of 1024 channels each, followed by two of 512 channels each). Each neuron in these four layers was ReLU activated. The model was trained under a weakly-supervised multiple-instance learning (MIL) paradigm. Each embedded tile propagated through the feature extractor described in Sect. 2.2 was treated as an instance of a bag con-

taining all quality-assured tiles of a specimen. Embedded tiles were aggregated using sigmoid-activated attention heads [12]. The final layer after the attention head was a linear layer that takes 512 channels as input and outputs a concordance rate prediction. To help prevent over-fitting, the training data set consisted of augmented versions (inspired by Tellez *et al.* [25]) of the tiles. Augmentations were generated with the following augmentation strategies: random variations in brightness, hue, contrast, saturation, (up to a maximum of 15%), Gaussian noise with 0.001 variance, and random 90° image rotations. We trained the melanoma concordance regression model with the Root mean squared error (RMSE) loss function. We regularized model training by using the dropout method [24] with 20% probability.

### 3 Results



**Fig. 2.** Melanoma Concordance Regression Training and Validation Loss Curves. The minimum validation RMSE loss was found to be 0.30. The RMSE loss on the test set was found to be  $0.28 \pm 0.01$ .

To demonstrate the performance of the Melanoma Concordance Regression model, we first show the training and validation loss curves in Fig. 2. The minimum validation loss was found to be 0.30. We then calculated the RMSE and  $R^2$  on the test set both across laboratory sites and for individual laboratory sites with 90% confidence intervals derived from bootstrapping with replacement. The RMSE across laboratory sites on the test set was calculated to be  $0.28 \pm 0.01$  and the  $R^2$  was found to be  $0.51 \pm 0.05$ . The lab-specific regression performance of the model on the test set is summarized in Table 3. We note that the RMSE performance is consistent across sites and are within the error bars derived from bootstrapping within replacement. The correlation between

**Table 3.** Regression metrics across the individual laboratory sites in our test data set. Errors are 90% confidence intervals derived through bootstrapping.

Metric	All sites	University of Florida	Florida - External	Jefferson
RMSE	$0.28 \pm 0.01$	$0.27 \pm 0.02$	$0.26 \pm 0.02$	$0.3 \pm 0.03$
$R^2$	$0.51 \pm 0.05$	$0.49 \pm 0.09$	$0.52 \pm 0.08$	$0.36 \pm 0.13$

the melanoma concordance predictions and the ground truth labels is shown in Fig. 3.

We next assessed the goodness-of-fit of the regression model by calculating the standardized residuals:

$$e_i = \frac{y_{pred,i} - y_{true,i}}{\sigma} \quad (1)$$

where  $i$  is the  $i$ th data point, and  $\sigma$  is the standard deviation of the residuals. The P-P plot comparing the cumulative distribution functions of the standardized residuals to a Normal distribution is shown in Fig. 4. The standardized residuals were found to be consistent with a Normal distribution by performing the Shapiro-Wilk test [23], where the p-value was found to be 0.97.

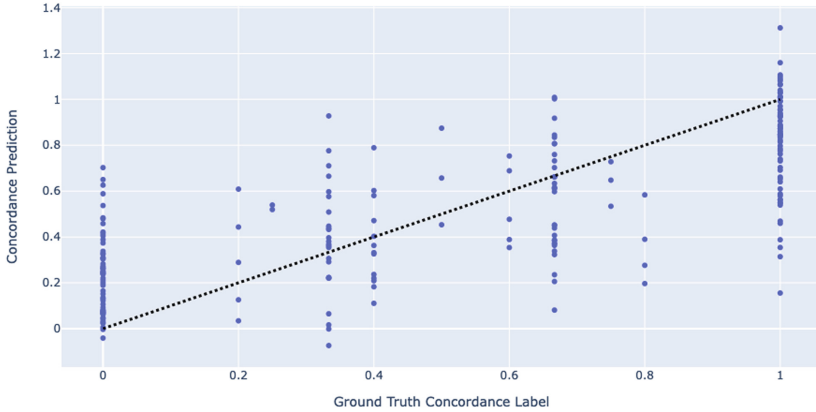
### 3.1 Malignant Classification

**Table 4.** Classification metrics when using the predicted concordance value as a malignant classifier across the individual laboratory sites in our test data set. Errors are 90% confidence intervals derived through bootstrapping with replacement.

Metric	All sites	University of Florida	Florida - External	Jefferson
AUC	$0.89 \pm 0.02$	$0.86 \pm 0.04$	$0.9 \pm 0.04$	$0.85 \pm 0.06$
Precision	$0.86 \pm 0.05$	$0.89 \pm 0.08$	$1.0 \pm 0.0$	$0.7 \pm 0.15$
Recall	$0.61 \pm 0.06$	$0.65 \pm 0.09$	$0.5 \pm 0.1$	$0.74 \pm 0.12$
Specificity	$0.97 \pm 0.02$	$0.94 \pm 0.04$	$1.0 \pm 0.0$	$0.9 \pm 0.05$

As mentioned in Sect. 2.1, increased inter-pathologist agreement on melanoma correlates with malignancy. We therefore investigated the performance of using the predicted concordance rate of the melanoma concordance model as a binary classifier to classify malignancy. We derived malignancy binary ground truth labels by defining a threshold value to binarize the ground truth concordance label, where a specimen with an observed concordance rate above this threshold value received a label of malignant, and below this threshold value received a label of not malignant. We performed a grid search on possible ground truth thresholds, and chose the threshold that maximized the Area Underneath





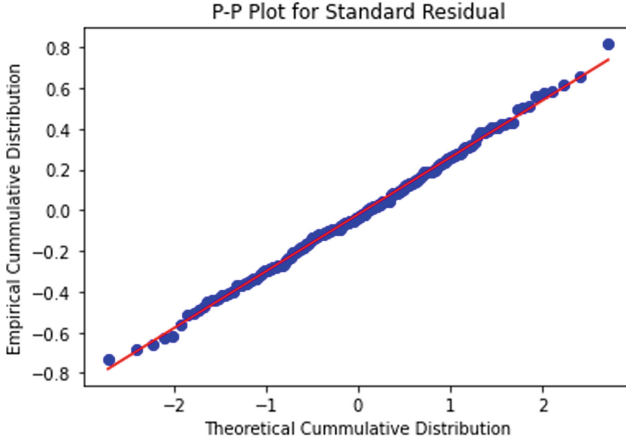
**Fig. 3.** The correlation of melanoma concordance predictions against the ground truth concordance label.

the Receiver Operating Characteristic (ROC) Curve (AUC). We also found the same ground truth threshold maximized the average precision of the precision-recall curve. We found that a threshold of 0.85 on the ground truth concordance rate to label malignancy maximizes both AUC and average precision, and yielded an AUC value of  $0.89 \pm 0.02$  and an average precision of  $0.81 \pm 0.04$ . The classification metrics with this threshold are shown both across laboratory sites and individual laboratory sites in Table 4. The ROC and precision-recall curves of using the melanoma concordance prediction as a malignant classifier is shown in Fig. 5 both across laboratory sites as well as individual laboratory sites. We note that the classification performance are within the error bars derived from bootstrapping with replacement across sites, with the exception of the low false positive rate at Florida - External, which exhibits very high precision.

### 3.2 Ablation Studies

**Binary Classifier Ablation Study.** We compared the performance of using the concordance rate as a malignancy classifier with a model trained to perform a binary classification task. In particular, we configured the final linear layer of the deep learning model to predict one of two classes: malignant or not malignant. We subsequently trained a binary classification model with the cross entropy loss function. We used a threshold of 0.85 on the ground truth to annotate malignancy derived from Sect. 3.1. We performed this ablation study specifically on data from the University of Florida, and the results are shown in Table 5. It can be seen that using the concordance rate to classify malignancy yields better performance than training a dedicated binary classifier.

**Feature Extractor Ablation Study.** We mentioned in Sect. 1 that features learned by pre-training on ImageNet are not guaranteed to generalize to

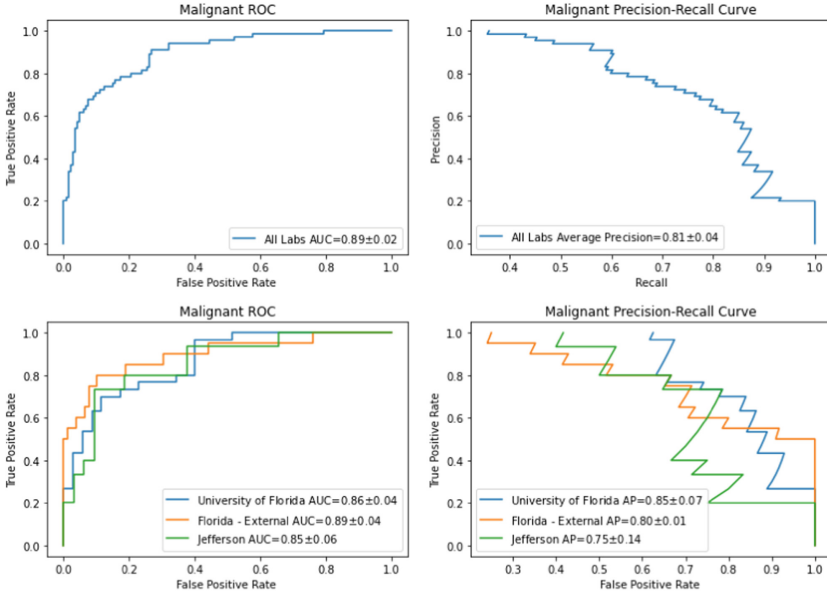


**Fig. 4.** A Probability-Probability plot is used as a goodness-of-fit in order to compare the cumulative distribution function of the standardized residuals of the melanoma concordance regression model to a Gaussian distribution. The P-P plot demonstrates the standardized residuals of melanoma concordance regression are normally distributed. The residuals were found to be consistent with a Normal distribution by performing the Shapiro-Wilk test on the standardized residuals. The p-value of the test was found to be 0.97.

histopathology image and could limit performance metrics and model robustness. To test this claim, we first visually examined SimCLR feature vectors to investigate visual coherence with morphological features. In particular, we projected the SimCLR feature vectors of 1,440 selected tiles to two dimensions with the Uniform Manifold Approximation and Projection (UMAP [14]) algorithm. We then arranged the tiles on a grid where the two-dimensional UMAP coordinates are mapped to the nearest grid coordinate, which arranged the tiles onto an evenly spaced plane. The arranged tiles are shown in Fig. 6. It can be seen that the learned feature vectors displays a coherence with respect to morphological features.

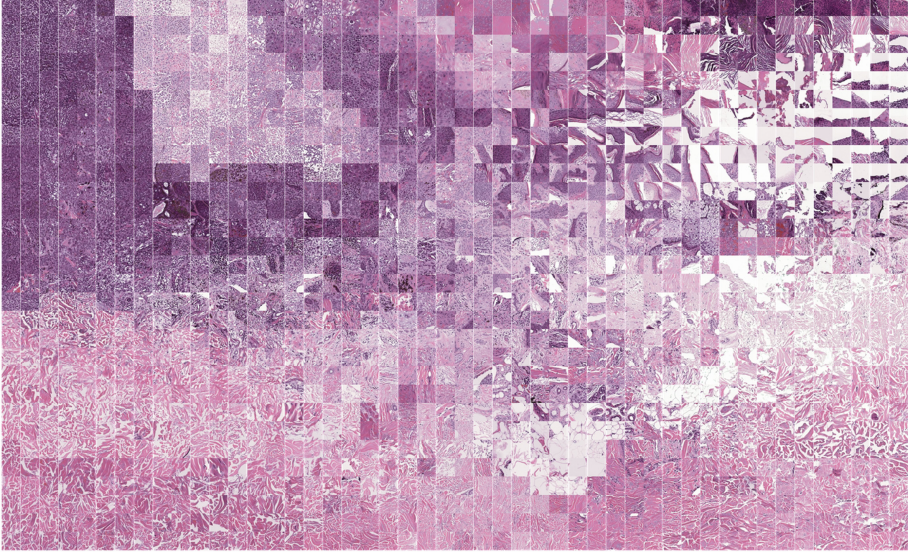
**Table 5.** Results of the binary classifier ablation experiment. A binary malignancy classifier was trained with the cross entropy loss function on data from the University of Florida, and the results are compared to using the predicted concordance rate from the regression model to classify malignancy.

Metric	Melanoma concordance regression	Binary classifier
AUC	0.86 ± 0.04	0.67 ± 0.03
Precision	0.89 ± 0.08	0.57 ± 0.07
Recall	0.62 ± 0.08	0.51 ± 0.07
Specificity	0.97 ± 0.02	0.75 ± 0.04



**Fig. 5.** Receiver Operating Characteristic (ROC) and Precision-Recall curves for using the melanoma concordance prediction as a malignant classifier on the test data set. The top row shows the curves for all laboratory sites, and the bottom row shows the curves for individual sites. The AUC value for malignant classification was found to be  $0.89 \pm 0.02$  and the average precision (AP) was found to be  $0.81 \pm 0.04$ . The bottom row demonstrates the curves for individual laboratory sites.

We next performed an ablation experiment to quantify the impact that our self-supervised feature extractor has on model performance. In particular, we ablated the SimCLR embedder by propagating the WSI tiles specifically from the University of Florida through a ResNet50 [11] pre-trained on the ImageNet [6] data set to embed each input tile into 2048 channel vectors. The embedded vectors were used to train a subsequent melanoma concordance regression model. The resulting model performance for this feature ablation experiment is summarized in Table 6 for both the raw regression metrics as well as the malignant classification task. (The same 0.85 threshold on the ground truth concordance rate defined in Sect. 3.1 was used to define malignancy) It can be seen that SimCLR results in higher performance in both the concordance regression and malignant classification tasks. In particular, there is a 14% improvement in RMSE in using the SimCLR embedder over imagenet for the regression task, and a 27.5% improvement in recall for the malignant classification task.



**Fig. 6.** Grid of SimCLR feature vectors arranged by their embedded vector values. First, SimCLR feature vectors of 1,440 selected tiles were projected onto two dimensions with the UMAP algorithm. We next arranged the WSI tiles on a grid where the two-dimensional UMAP coordinates are mapped to the nearest grid coordinate, which arranged the tiles onto an evenly spaced plane.

**Table 6.** Feature Extractor ablation study results from the University of Florida. feature vectors derived from SimCLR with a ResNet50 model backbone trained on Whole Slide Images were compared to feature vectors derived from a pre-trained ResNet50 model on the ImageNet data set. Concordance regression metrics and classification metrics are shown. Classification metrics were derived by utilizing the predicted concordance value as a malignant classifier.

Metric	SimCLR trained with WSIs	ResNet50 trained with ImageNet
RMSE	$0.27 \pm 0.02$	$0.31 \pm 0.02$
$R^2$	$0.49 \pm 0.09$	$0.15 \pm 0.15$
AUC	$0.86 \pm 0.04$	$0.75 \pm 0.06$
Precision	$0.89 \pm 0.08$	$0.76 \pm 0.07$
Recall	$0.62 \pm 0.08$	$0.45 \pm 0.08$
Specificity	$0.97 \pm 0.02$	$0.71 \pm 0.08$

## 4 Conclusions

We presented in this paper a melanoma concordance regression model that was trained across three laboratory sites that demonstrates regression performance of  $0.28 \pm 0.01$  on the test set across the three laboratory sites in our data set.

The performance of the melanoma concordance regression model was limited by the number of dermatopathologists performing concordant reviews on our data set. From Table 2, the average number of dermatopathologists reviewing a case is 3.7 pathologists, resulting in a concordance rate resolution of 0.27. The melanoma concordance regression model that we built is at the limit of this resolution. The regression performance could be further improved by collecting more concordance reviews from dermatopathologists. We maximized the melanoma concordance regression performance by learning informative features through self-supervised training with SimCLR. We also demonstrated that the predicted melanoma concordance rate can be used as a malignant classifier, as the concordance rate is correlated with the likelihood of malignancy. We adjusted the threshold of malignancy to maximize both AUC and average precision, where we found the AUC value to be  $0.89 \pm 0.02$  and the average precision to be  $0.81 \pm 0.04$ .

These results are an important first step for building an AI system capable of predicting the results of consulting a panel of experts and delivering a score based on the degree to which the experts would agree on a particular diagnosis or opinion. Upon further improvement, a concordance score reliably representing a panel of experts can additionally be used as a diagnostic assist to suggest additional testing or other action, for example the ordering of additional stains or a genetic test. Additionally, dermatopathologists utilizing a concordance score as a diagnostic assist can empower them to speed up malignancy diagnosis, increase confidence in case sign out, and lower the misdiagnosis rate, thereby increasing sensitivity and improving patient care. The possibility also exists that a concordance score could be fed into a combined-test that incorporates the results of a multi-gene assay (e.g., Castle MelanomaDx [9]) in order to enhance the performance of the test. Finally, such a concordance AI system can be extended to any pathology that exhibits a high discordance rate, such as breast cancer staging [21] and Gleason grading of prostate cancer [5].

**Acknowledgments.** The authors thank the support of Jeff Baatz, Ramachandra V. Chamarthi, Nathan Langlois, and Liren Zhu at Proscia for their engineering support; Theresa Feeser, Pratik Patel, and Aysegul Ergin Sutcu at Proscia for their data acquisition and Q&A support; and Dr. Curtis Thompson at CTA and Dr. David Terrano at Bethesda Dermatology Laboratory for their consensus annotation support.

## References

1. American Cancer Society. Cancer facts and figures 2021. American Cancer Society, Inc. (2021)
2. Campanella, G., et al.: Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nat. Med.* **25**(8), 1301–1309 (2019)
3. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. arXiv preprint [arXiv:2002.05709](https://arxiv.org/abs/2002.05709) (2020)
4. Chen, T., Kornblith, S., Swersky, K., Norouzi, M., Hinton, G.: Big self-supervised models are strong semi-supervised learners. arXiv preprint [arXiv:2006.10029](https://arxiv.org/abs/2006.10029) (2020)

5. Coard, K.C., Freeman, V.L.: Gleason grading of prostate cancer: level of concordance between pathologists at the University Hospital of the West Indies. *Am. J. Clin. Pathol.* **122**(3), 373–376 (2004)
6. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: a large-scale hierarchical image database. In: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255 (2009)
7. Elmore, J.G., et al.: Pathologists' diagnosis of invasive melanoma and melanocytic proliferations: observer accuracy and reproducibility study. *BMJ* **357** (2017)
8. Falcon, W., Borovec, J., Brundyn, A., Harsh Jha, A., Koker, T., Nitta, A.: Pytorch lightning bolts self supervised pre-trained models documentation (2022). [https://pytorch-lightning-bolts.readthedocs.io/en/latest/self\\_supervised\\_models.html](https://pytorch-lightning-bolts.readthedocs.io/en/latest/self_supervised_models.html)
9. Fried, L., Tan, A., Bajaj, S., Liebman, T.N., Polsky, D., Stein, J.A.: Technological advances for the detection of melanoma: advances in molecular techniques. *J. Am. Acad. Dermatol.* **83**(4), 996–1004 (2020)
10. Gerami, P., et al.: Histomorphologic assessment and interobserver diagnostic reproducibility of atypical spitzoid melanocytic neoplasms with long-term follow-up. *Am. J. Surg. Pathol.* **38**(7), 934–940 (2014)
11. He, K., Xiangyu, Z., Ren, S., Sun, J.: Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385* (2015)
12. Ilse, M., Tomczak, J., Welling, M.: Attention-based deep multiple instance learning. In: *Proceedings of the 35th International Conference on Machine Learning*, pp. 2127–2136 (2018)
13. Li, B., Li, Y., Eliceiri, K.W.: Dual-stream multiple instance learning network for whole slide image classification with self-supervised contrastive learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14318–14328 (2021)
14. McInnes, L., Healy, J., Saul, N., Großberger, L.: UMAP: uniform manifold approximation and projection. *J. Open Source Softw.* **3**(29), 861 (2018). <https://doi.org/10.21105/joss.00861>
15. Noone, A.M., et al.: Cancer incidence and survival trends by subtype using data from the surveillance epidemiology and end results program, 1992–2013. *Cancer Epidemiol. Prev. Biomark.* **26**(4), 632–641 (2017)
16. Otsu, N.: A threshold selection method from gray-level histograms. *IEEE Trans. Syst. Man Cybern.* **9**(1), 62–66 (1979). <https://doi.org/10.1109/TSMC.1979.4310076>
17. Pantanowitz, L., et al.: An artificial intelligence algorithm for prostate cancer diagnosis in whole slide images of core needle biopsies: a blinded clinical validation and deployment study. *Lancet Digit. Health* **2**(8), e407–e416 (2020). [https://doi.org/10.1016/S2589-7500\(20\)30159-X](https://doi.org/10.1016/S2589-7500(20)30159-X). <https://www.sciencedirect.com/science/article/pii/S258975002030159X>
18. Pech-Pacheco, J.L., Cristóbal, G., Chamorro-Martinez, J., Fernández-Valdivia, J.: Diatom autofocusing in brightfield microscopy: a comparative study. In: *Proceedings 15th International Conference on Pattern Recognition, ICPR-2000*, vol. 3, pp. 314–317. IEEE (2000)
19. Pertuz, S., Puig, D., García, M.Á.: Analysis of focus measure operators for shape-from-focus. *Pattern Recogn.* **46**, 1415–1432 (2013)
20. Piepkorn, M.W., et al.: The MPATH-Dx reporting schema for melanocytic proliferations and melanoma. *J. Am. Acad. Dermatol.* **70**(1), 131–141 (2014)
21. Plichta, J.K., et al.: Clinical and pathological stage discordance among 433,514 breast cancer patients. *Am. J. Surg.* **218**(4), 669–676 (2019)

22. Sankarapandian, S., et al.: A pathology deep learning system capable of triage of melanoma specimens utilizing dermatopathologist consensus as ground truth. In: Proceedings of the ICCV 2021 CDpath Workshop (2021)
23. Shapiro, S.S., Wilk, M.B.: An analysis of variance test for normality (complete samples). *Biometrika* **52**(3–4), 591–611 (1965). <https://doi.org/10.1093/biomet/52.3-4.591>
24. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **15**(56), 1929–1958 (2014). <http://jmlr.org/papers/v15/srivastava14a.html>
25. Tellez, D., et al.: Quantifying the effects of data augmentation and stain color normalization in convolutional neural networks for computational pathology. *Med. Image Anal.* **58**, 101544 (2019)
26. Welch, H.G., Mazer, B.L., Adamson, A.S.: The rapid rise in cutaneous melanoma diagnoses. *N. Engl. J. Med.* **384**(1), 72–79 (2021)