

# How Many Annotators Do We Need?

## A Study on the Influence of Inter-Observer Variability on the Reliability of Automatic Mitotic Figure Assessment

Frauke Wilm<sup>1</sup>, Christof A. Bertram<sup>2</sup>, Christian Marzahl<sup>1</sup>, Alexander Bartel<sup>3</sup>,  
Taryn A. Donovan<sup>4</sup>, Charles-Antoine Assenmacher<sup>5</sup>, Kathrin Becker<sup>6</sup>,  
Mark Bennett<sup>7</sup>, Sarah Corner<sup>8</sup>, Brieuc Cossic<sup>9</sup>, Daniela Denk<sup>10</sup>,  
Martina Dettwiler<sup>11</sup>, Beatriz Garcia Gonzalez<sup>7</sup>, Corinne Gurtner<sup>11</sup>,  
Annabelle Heier<sup>12</sup>, Annika Lehmbeker<sup>12</sup>, Sophie Merz<sup>2,12</sup>, Stephanie Plog<sup>7</sup>,  
Anja Schmidt<sup>12</sup>, Franziska Sebastian<sup>12</sup>, Rebecca C. Smedley<sup>8</sup>, Marco Tecilla<sup>13</sup>,  
Tuddow Thaiwong<sup>8</sup>, Katharina Breining<sup>1</sup>, Matti Kiupel<sup>8</sup>, Andreas Maier<sup>1</sup>,  
Robert Klopffleisch<sup>2</sup>, Marc Aubreville<sup>1,14</sup>

<sup>1</sup>Pattern Recognition Lab, Computer Sciences, Friedrich-Alexander-Universität  
Erlangen-Nürnberg, Erlangen, Germany

<sup>2</sup>Institute of Veterinary Pathology, Freie Universität Berlin, Germany

<sup>3</sup>Institute for Veterinary Epidemiology and Biostatistics, Freie Universität Berlin

<sup>4</sup>Department of Anatomic Pathology, Animal Medical Center, New York, USA

<sup>5</sup>Department of Pathobiology, University of Pennsylvania, Philadelphia, USA

<sup>6</sup>Department of Pathology, University of Veterinary Medicine Hannover, Germany

<sup>7</sup>Synlab's VPG Histology, Bristol, UK

<sup>8</sup>Veterinary Diagnostic Laboratory, Michigan State University, Lansing, USA

<sup>9</sup>Pharmacology & Preclinical Development, Idorsia Pharmaceuticals Ltd, Switzerland

<sup>10</sup>International Zoo Veterinary Group, Keighley, UK

<sup>11</sup>Institute of Animal Pathology, Vetsuisse Faculty, University of Bern, Switzerland

<sup>12</sup>IDEXX Vet Med Labor GmbH, Kornwestheim, Germany

<sup>13</sup>F. Hoffmann-La Roche Ltd, Basel, Switzerland

<sup>14</sup>Technische Hochschule Ingolstadt, Ingolstadt, Germany

`frauke.wilm@fau.de`

**Abstract.** Density of mitotic figures in histologic sections is a prognostically relevant characteristic for many tumours. Due to high inter-pathologist variability, deep learning-based algorithms are a promising solution to improve tumour prognostication. Pathologists are the gold standard for database development, however, labelling errors may hamper development of accurate algorithms. In the present work we evaluated the benefit of multi-expert consensus ( $n = 3, 5, 7, 9, 11$ ) on algorithmic performance. While training with individual databases resulted in highly variable  $F_1$  scores, performance was notably increased and more consistent when using the consensus of three annotators. Adding more annotators only resulted in minor improvements. We conclude that databases by few pathologists and high label precision may be the best compromise between high algorithmic performance and time investment.

---

F. Wilm and C. A. Bertram contributed equally to this work

## 1 Introduction

Histologic examination of tumour specimens is used to derive important information with regards to patient prognosis and selection of appropriate treatment. For numerous tumour types, including canine mast cell tumours, cellular proliferation is one of the most meaningful prognostic parameters. As part of the recommended grading schemes, cells undergoing division (mitotic figures) must be counted in histologic sections. However, identification of mitotic figures has a high degree of inter-observer variability due to inconsistent classification of mitotic figures (as opposed to mitotic-like impostors) or overlooking/omitting mitotic figure candidates [1,2]. In order to improve reproducibility and accuracy of enumerating mitotic figures, promising deep learning-based algorithms for the automated analysis of digitised histologic sections have been developed [1,3,4,5]. However, as pathologists are the gold standard for dataset development, visual and cognitive limitations of human experts may hamper the consistency of datasets and subsequently algorithmic performance [2].

All available datasets on mitotic figures from human and canine tumours have used not more than two pathologists as annotators for initial labelling in histologic images [2,3,4,5,6]. Disagreement between the labels of these two pathologists was reported in up to 68.2% [3]. Divergent labels between these two pathologists were reviewed for final consensus by one [5] or two [3,4] additional pathologists or by reassessment of the same experts [2,6]. Although consensus by multiple pathologists is expected to counterbalance the high inter-rater variability, influence on algorithmic performance has not been examined to date.

This study aims to evaluate the ideal number of expert opinions required for the development of deep learning-based mitotic figure detection algorithms with high accuracy. For this purpose, a well established object detection algorithm was trained with labels derived from a consensus of a range of pathologists ( $n = 1-11$ ) and evaluated against reference annotations derived from the total consensus of twelve pathologists.

## 2 Material and methods

For this investigation, datasets were created from 50 histologic images of canine mast cell tumours from 50 patients. Use of these samples was approved by the local governmental authorities (State Office of Health and Social Affairs of Berlin, approval ID: StN 011/20). Histologic sections were created with routine methods (haematoxylin and eosin stain) and digitised at a resolution of  $0.25 \mu\text{m}$  per pixel (400x magnification) using an Aperio ScanScope CS2 (Leica, Germany) scanner.

### 2.1 Creation of databases

Independent databases were created by twelve veterinary pathologists (each at least three years of experience in histopathology) using the SlideRunner annotation software (<https://github.com/DeepPathology/SlideRunner>). For each of the

50 slides, a field of interest in the tumour area with a standard size of  $2.37 \text{ mm}^2$  was selected and annotators marked centroid coordinates of all mitotic figures recognised in these tumour regions. Each pathologist identified between 1,324 and 4,412 mitotic figures (total number of annotations: 32,917). The dataset was split into 35 training, 5 validation and 10 test images with similar variability in mitotic figure density in the selected regions.

## 2.2 Deep learning-based mitotic figure detection

For mitotic figure detection we customised a publicly available RetinaNet implementation [1] with a ResNet18 stem. For network training, 2,500 patches ( $1,024 \times 1,024$  pixels at highest resolution) each containing at least one mitotic figure were randomly drawn from the 35 training images. The network was trained with a variable number of databases created by the pathologists. The training and validation reference was individually defined as the majority vote of this subgroup. The training was split into two phases. First, only the randomly initialised network heads were trained (batch size: 12) for five epochs using a maximal learning rate of  $10^{-3}$ . Afterwards, the complete model was trained for an additional ten epochs and a maximal learning rate of  $10^{-4}$ . During this second phase, 1,500 patches from the validation set were used for model selection.

Four different training set-ups were evaluated: (1) The network was trained with each individual pathologist’s database and the model performance was compared to the annotator’s performance on the test set. To ensure stability, the  $F_1$  score was computed as median of three independent training runs. (2) The network was trained on the consensus of an increasingly larger, randomly chosen subset of pathologists. In order to obtain unambiguous agreement, these increases consisted only of odd increments. For each addition, ten training runs were averaged to determine the influence of the random selection of annotators. For the last two experiments (3) and (4) the annotators were sorted in descending order by their label agreement compared to the majority vote measured by the  $F_1$  score. The model was then trained on (3) the first  $n$  pathologists ( $\hat{=}$  highest agreement) and (4) on the last  $n$  pathologists ( $\hat{=}$  lowest agreement).

By using the consensus, we only trained with mitotic figures with at least 50% agreement. We modified the loss computation by weighting the sample with the percentage of pathologists that agreed upon the respective sample ( $> 50\text{-}100\%$ ).


## 2.3 Model evaluation

The test set was generated from ten test slides and comprised all mitotic figure labels upon which the majority of all twelve pathologists (i.e. at least seven pathologists) agreed. For model performance evaluation, the  $F_1$  score was computed. As previously defined by Bertram et al. [6], a mitotic figure detection was counted as true positive if the Euclidean distance of the reference and predicted bounding box centroids was at most 25 pixels ( $\hat{=}$   $6.25 \mu\text{m}$ , i.e. approximately the average cell radius of neoplastic mast cells). The detector confidence threshold was chosen based on the highest performance on the validation set.

### 3 Results


Label accuracy of the twelve pathologists compared to the majority vote was highly variable with  $F_1$  scores ranging from 0.64 to 0.85 (median: 0.77) for the whole dataset and from 0.68 to 0.86 (median: 0.77) for the test cases. In Fig. 1 the labelling performance of each annotator is compared to the  $F_1$  scores of the RetinaNet, which was trained on mitotic figure labels of the same pathologists. For the five pathologists with the lowest label agreement on the reference annotations, the algorithmic approach showed similar performance on the test dataset. Regardless of the label performance of the annotators, the algorithmic  $F_1$  score was capped at around 0.74 for the network architecture used. A more detailed analysis showed that higher algorithmic performance was associated with high annotation precision, while labelling sensitivity seemed to have a negligible influence.

**Fig. 1.** Labelling performance ( $F_1$  score) of individual annotator compared to the performance of networks trained with annotations from each respective pathologist (median and range of three training runs) evaluated on the majority vote test set.



Tab. 1 summarises the network performance when training the algorithm with databases of an increasingly larger subset of annotators. Compared to a single annotator, models trained with multi-expert labels resulted in an overall higher performance and a lower variability as measured by the interquartile range. Fig. 2 shows that training with databases that have the highest agreement with the majority vote yielded the best results. However, this arrangement benefited the least from higher numbers of annotators (as opposed to using the databases with the lowest agreement).

**Fig. 2.** Performance ( $F_1$  score) of models trained with different combinations of annotators. Whiskers represent the minimum and maximum  $F_1$  score from the training runs.



**Table 1.** Comparison of  $F_1$  scores of models trained with different reference annotations based upon the majority vote from different combinations of annotators.

Database	Number of annotators	Training Runs	Median	Minimum	Maximum	Interquartile Range
Each once	1	18 x 3	0.726	0.625	0.763	0.038
Random	3	10	0.75	0.728	0.769	0.016
Random	5	10	0.749	0.729	0.777	0.017
Random	7	10	0.756	0.730	0.765	0.008
Random	9	10	0.763	0.745	0.771	0.008
Random	11	10	0.763	0.746	0.769	0.010
Highest agreement	3	5	0.751	0.73	0.762	0.010
Highest agreement	5	5	0.759	0.748	0.770	0.007
Highest agreement	7	5	0.767	0.75	0.776	0.003
Highest agreement	9	5	0.768	0.749	0.771	0.016
Lowest agreement	3	5	0.756	0.753	0.76	0.003
Lowest agreement	5	5	0.74	0.720	0.745	0.018
Lowest agreement	7	5	0.754	0.736	0.761	0.017
Lowest agreement	9	5	0.756	0.749	0.768	0.012

## 4 Discussion

The high variability between annotators in the present study is consistent with previous studies [2,3,7] and warrants a detailed evaluation of label consistency and the impact on algorithmic performance, as was the goal of the present study. Generally, our results show that the use of a consensus of a higher number of pathologists for training the algorithm yields better and more consistent results. In particular,  $F_1$  scores were noticeably variable when training the algorithm with single annotators (interquartile range: 0.038). A consensus by three pathologists was highly beneficial for more consistent training results (interquartile range: 0.016). Increasing the number of randomly selected annotators further only improved the median  $F_1$  score by a small amount (+0.013). Nevertheless, we have also shown that even training with a single annotator can result in high performance. Of note, a high annotation precision of individual pathologists, i.e. few false positive labels, seemed to be more important than high labelling sensitivity, i.e. few false negatives labels, for training of high-performing models. This analysis, however, might have been biased by the ground truth definition as consensus of the majority of pathologists. Our results emphasise that high annotation accuracy and consensus by a small number of pathologists may result in the best trade-off between algorithmic performance and labour intensity of dataset development. Further enhancement of label consistency may be achieved with repeated screening of images or algorithmically augmented labelling [6]. An interesting approach to reduce subjectivity for future mitotic figures datasets was recently introduced by Tellez et al. [8]. In this study, a specific immunohis-

tochemical marker of mitotic figures was used to derive object labels and labels were assigned to images with standard histologic stain via image registration.

The major limitation of the conducted experiments was the pathologist-defined ground truth. Although pathologists are the current gold standard for labelling histologic images, they have high inter-rater variability which hampers not only training of data-driven algorithms (as proven in the present study) but also biases performance evaluation. The finding that algorithms trained with multi-expert databases outperformed many pathologists on the test set was attributed to the fact that algorithms yielded high sensitivity. Furthermore, the experiments in the present work were limited to a standard object detection architecture with relatively low complexity. Compensation for noisy labels and higher  $F_1$  scores may be achieved with a more complex model, such as by adding a second classification stage [1,6], and larger training datasets. Further improvement of data-derived algorithms may be accomplished with advanced deep learning methods that incorporate label accuracy during training.

**Acknowledgement.** FW gratefully acknowledges financial support received by Merck KGaA. CAB gratefully acknowledges financial support received from the Dres. Jutta & Georg Bruns-Stiftung für innovative Veterinärmedizin.

## References

1. Aubreville M, Bertram CA, Marzahl C, et al. Deep learning algorithms out-perform veterinary pathologists in detecting the mitotically most active tumor region. *Sci Rep.* 2020;10(16447):1–11.
2. Bertram CA, Veta M, Marzahl C, et al. Are pathologist-defined labels reproducible? Comparison of the TUPAC16 mitotic figure dataset with an alternative set of labels. In: *Interpretable and Annotation-Efficient Learning for Medical Image Computing. IMIMIC 2020, MIL3ID 2020, LABELS 2020. Lecture Notes in Computer Science.* Springer; 2020. p. 204–213.
3. Veta M, Van Diest PJ, Willems SM, et al. Assessment of algorithms for mitosis detection in breast cancer histopathology images. *Med Image Anal.* 2015;20(1):237–248.
4. Veta M, Heng YJ, Stathonikos N, et al. Predicting breast tumor proliferation from whole-slide images: the TUPAC16 challenge. *Med Image Anal.* 2019;54:111–121.
5. Roux L, Racoceanu D, Capron F, et al. MITOS & ATYPICIA-Detection of mitosis and evaluation of nuclear atypia score in breast cancer histological images. *IPAL, Agency Sci. Technol Res Inst Infocom Res, Singapore, Tech Rep.* 2014;.
6. Bertram CA, Aubreville M, Marzahl C, et al. A large-scale dataset for mitotic figure assessment on whole slide images of canine cutaneous mast cell tumor. *Scientific Data.* 2019;6(1):1–9.
7. Veta M, Diest PJV, Jiwa M, et al. Mitosis counting in breast cancer: Object-level interobserver agreement and comparison to an automatic method. *PLoS one.* 2016;11(8).
8. Tellez D, Balkenhol M, Otte-Höller I, et al. Whole-slide mitosis detection in H&E breast histology using PHH3 as a reference to train distilled stain-invariant convolutional networks. *IEEE Trans Med Imaging.* 2018;37(9):2126–2136.