

MIDOG 2025
Mitosis Domain Generalization Challenge

MICCAI 2025
Daejeon
REPUBLIC OF KOREA

Mix, Align, Distil: Reliable Cross-Domain Atypical Mitosis Classification

Kaustubh Atey Sameer Anand Jha Gouranga Bala Amit Sethi

Indian Institute of Technology Bombay, India

September 23, 2025

MeDAL

Mix-Align-Distil: Method Overview

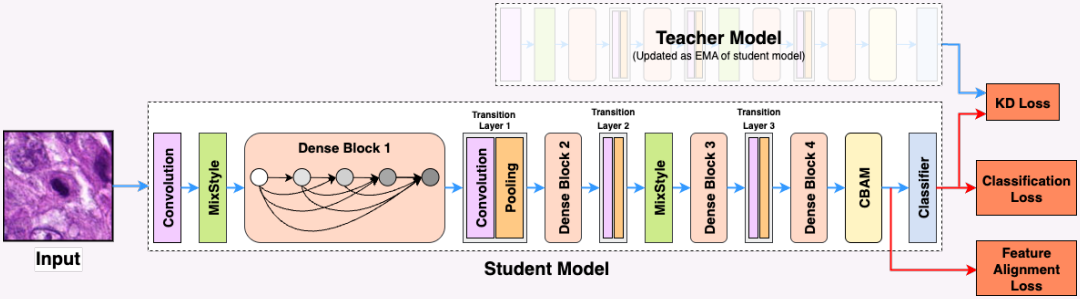


Figure 1: Training-time framework for cross-domain robustness.

Mix

- MixStyle [1] simulates stain and contrast shifts in feature space.
- Encourages structure-focused features over color.

Align

- Domain-wise channel variance loss on backbone features.
- Makes features invariant to domain shifts.

Distil

- Knowledge distillation [2] using an EMA teacher.
- Stabilizes training and predictions.

Our method diversifies styles, aligns domain statistics, and transfers stable predictions: boosting cross-domain generalization with negligible inference overhead.

Implementation, Results and Discussion

Implementation details

- **Datasets & splits:** MIDOG 2025 Atypical [3] & MIDOG 2021 AMi-Br subset [4]; 75/25 train-val (leakage-free). Held-out evaluation on OMG-Octo Atypical [5] & TUPAC16 AMi-Br [6] subset.
- **Training:** DenseNet-121 backbone with MixStyle (early, mid), CBAM [7] before GAP; strong augmentations.
- **Optimization:** AdamW, ReduceLROnPlateau, Gradient Clipping; EMA-teacher KD [8] (T=2); batch size 512.

Method	Bal. Acc	Sens.	Spec.	ROC AUC
Baseline	0.7933	0.9014	0.6851	0.8859
Ours	**0.8762**	**0.8873**	**0.8651**	**0.9499**

Table 1: Preliminary leaderboard results (Track 2).

Final discussion

- Tackles domain shift with three complementary, train-time components: MixStyle (style perturbation), feature alignment, and EMA-teacher distillation.
- Robust on held-out domains with strong, balanced leaderboard metrics (Table 1), while remaining lightweight ($\sim$ 28 MB) with minimal test-time overhead.
- Next steps: Adopt foundational model backbones (e.g. UniV2 [9], DINOv3 [10]) to further improve accuracy and robustness.

References i


Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang.
Domain generalization with mixstyle.
arXiv preprint arXiv:2104.02008, 2021.


Geoffrey Hinton, Oriol Vinyals, and Jeff Dean.
Distilling the knowledge in a neural network.
arXiv preprint arXiv:1503.02531, 2015.


Viktoria Weiss, Sweta Banerjee, Taryn Donovan, Thomas Conrad, Robert Klopfleisch, Jonas Ammeling, Christopher Kaltenecker, Dominik Hirling, Mitko Veta, Nikolas Stathonikos, Peter Horvath, Katharina Breininger, Marc Aubreville, and Christof Bertram.
A dataset of atypical vs normal mitoses classification for midog - 2025, April 2025.


Christof A Bertram, Viktoria Weiss, Taryn A Donovan, Sweta Banerjee, Thomas Conrad, Jonas Ammeling, Robert Klopfleisch, Christopher Kaltenecker, and Marc Aubreville.
Histologic dataset of normal and atypical mitotic figures on human breast cancer (ami-br).
In *BVM Workshop*, pages 113–118. Springer, 2025.


Zhuoyan Shen, Mikaël Simard, Douglas Brand, Vanghelita Andrei, Ali AlKhader, Fatine Oumlil, Katherine Trevers, Thomas Butters, Simon Haefliger, Eleanna Kara, Fernanda Amary, Roberto Tirabosco, Paul Cool, Gary Royle, Maria A. Hawkins, Adrienne M. Flanagan, and Charles-Antoine Collins Fekete.
Omg-octo: Uniformised large scale database of mitotic figures in haematoxylin and eosin-stained slides, August 2024.

References ii

Mitko Veta, Josien P. W. Pluim, Paul J. van Diest, and Santo Santorcus.
Tupac16: Tumor proliferation assessment challenge 2016.
https://tupac.tue-image.nl/, 2016.
Accessed: 2023-10-15.

Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon.
Cbam: Convolutional block attention module.
In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018.

Antti Tarvainen and Harri Valpola.
Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results.
Advances in neural information processing systems, 30, 2017.

Richard J Chen, Tong Ding, Ming Y Lu, Drew FK Williamson, Guillaume Jaume, Bowen Chen, Andrew Zhang, Daniel Shao, Andrew H Song, Muhammad Shaban, et al.
Towards a general-purpose foundation model for computational pathology.
Nature Medicine, 2024.

Oriane Siméoni, Huy V Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, et al.
Dinov3.
arXiv preprint arXiv:2508.10104, 2025.

Thank You!