

Genetically Aligned Patient Representations Improve Hematological Diagnosis

Muhammed Furkan Dasdelen^{1,2*}, Fatih Ozlugedik^{1*}, Ilaria Looser¹, Rao Muhammad Umer¹, Christian Pohlkamp³, and Carsten Marr^{1,4,5,6,7**}

¹ Institute of AI for Health, Helmholtz Munich, Germany

² International School of Medicine, Istanbul Medipol University, Türkiye

³ Munich Leukemia Laboratory, Germany

⁴ Department of Medicine III, Ludwig-Maximilian-University Hospital, Germany

⁵ Department of Physics, University of Munich, Germany

⁶ Munich Center for Machine Learning (MCML), Germany

⁷ DKTK, German Cancer Consortium, Germany

Abstract. Multimodal alignment of histopathology encoders with transcriptomic and genomic data has been shown to significantly improve performance in downstream diagnostic tasks. Hematological cytology is unique in that visual single-cell evaluation is often paired with cytogenetics and molecular genetics for blood cancer diagnosis. In this study, we present a framework to align single white blood cell images with chromosomal aberrations (karyotype) and somatic mutations from targeted gene panels. Our training strategy follows a two-stage approach: (i) self-supervised, vision-only pretraining of a transformer aggregator using an iBOT head on a cohort of over 1500 patients, and (ii) genetic alignment via supervised contrastive loss on acute myeloid leukemia patients. Our genetically aligned patient encoder improves hematological diagnostic tasks, outperforming slide-level histopathology foundation models. Additionally, the model provides off-the-shelf retrieval capabilities for diseases and genetic alterations. Incorporating genetic data into patient encoders increases the quality of patient representations, providing a framework that aligns with clinical diagnostic workflows and paves the way for future multimodal hematology-specific AI. Source code and model weights are available at <https://github.com/marrlab/GenBloom>.

Keywords: cytology · karyotype · multimodal alignment · genetic.

1 Introduction

Hematological diagnosis fundamentally relies on integrating microscopic morphology with cytogenetic and molecular profiling for precise tumor classification, risk stratification, and treatment selection [10,8,18]. While evaluating peripheral

* These authors contributed equally.

** Corresponding author: carsten.marr@helmholtz-munich.de

blood and bone marrow smears provides a rapid initial assessment, genetic profiling is essential to formally classify the disease. This tight clinical coupling of modalities motivates the need for joint computational modeling.

In computational pathology, deep learning foundation models trained on whole-slide images (WSIs) have shown strong performance on downstream tasks like tumor subtyping and survival prediction [20,15,3]. Recent multimodal frameworks further demonstrate that aligning histopathology representations with molecular data improves predictive performance, generalization, and the biological relevance of visual embeddings [17,19]. These findings confirm that morphology and molecular features share an exploitable latent space.

Blood smears provide information at single-cell resolution—the scale where morphological variation most directly reflects underlying genetic alterations. Existing hematological models predominantly analyze morphology alone [7,2,16,5], failing to capture molecular heterogeneity. Furthermore, the scarcity of large-scale paired datasets restricts the development of models.

To bridge this gap, we present GenBloom, the first genetically aligned, slide-level blood model tailored to hematology. GenBloom integrates single white blood cell images with cytogenetic abnormalities and somatic mutations to learn comprehensive patient embeddings. Our two-stage training paradigm involves: (i) large-scale self-supervised pretraining on cellular morphology to extract robust visual features, and (ii) supervised multimodal alignment to anchor patient embeddings within the genetic space. This design successfully captures clinically meaningful relationships between cellular morphology and molecular disease drivers.

2 Methods

2.1 Pretraining dataset

For image pretraining, we used an in-house peripheral blood smear dataset (collected at Munich Leukemia Laboratory), which contains single-cell images from 1,634 patients spanning a range of hematological diseases (acute leukemia, myelodysplastic syndrome (MDS), myeloproliferative neoplasm (MPN), overlap syndromes, lymphoma, multiple myeloma, and reactive changes) as well as healthy controls. It includes 794,527 single-cell images for pretraining (Fig. 1a). For genetic alignment, we used the AML-Hehr dataset [7], which comprises peripheral blood smear images from 189 acute myeloid leukemia (AML) patients and a healthy cohort, paired with molecular and cytogenetic profiles. The dataset includes 37 unique somatic mutations and 52 distinct karyotypes (Fig. 1b, c). We held out a test set ($n=43$) and contrastively aligned the remaining 146 patients' images with their genetic information.

2.2 Evaluation tasks and dataset

For downstream evaluation, we considered three publicly available patient-level classification tasks. First, we performed AML genetic subtyping on the held-out AML-Hehr test set, focusing on *PML::RARA* fusion (train/test $n = 18/6$),

CBFB::MYH11 fusion (train/test $n = 28/9$), *NPM1* mutation (train/test $n = 28/8$), and *RUNX1::RUNX1T1* fusion (train/test $n = 25/7$), and controls (train/test $n = 47/13$). The other two datasets are out-of-domain and were used to test generalizability: the APL-AML dataset includes acute promyelocytic leukemia (APL, train/test $n = 22/12$) and other acute myeloid leukemias (AML, train/test $n = 60/12$), and AMH was curated from the cAItomorph test set [2], including AML (train/test $n = 79/20$) and healthy individuals (train/test $n = 29/8$). In all experiments, we performed 5-fold cross-validation while keeping the test set fixed.

We also performed retrieval analysis. Here, we embedded each patient’s slide and genomic profile (karyotype or mutation) into a shared representation space and performed cosine similarity to rank candidates. We evaluated genomics→slide by retrieving the correct slide given a genomic query, slide→genomics by retrieving the correct genomic profile given a slide query, and slide→slide by retrieving slides from the same disease given a slide query.

2.3 Data processing

Single cell images were reshaped to 224×224 , normalized with Imagenet statistics. We used the DinoBloom-B [11] hematology image encoder to create single cell representations. DinoBloom was frozen throughout all experiments.

Structural abnormalities were characterized through cytogenetics and fluorescence in situ hybridisation (FISH), with karyotyping by chromosome banding analysis documented according to the International System for Human Cytogenetic Nomenclature (ISCN) standards [14,12]. Karyotype data were processed using CytoGPS [1], which converts ISCN text strings into patient-level binary indicators of chromosomal loss, gain, and fusion. Encoding three indicators per cytoband (368 bands) yielded a 1,104-dimensional input for the cytogenetics branch of the model (Fig. 1c,e).

At diagnosis, patients also underwent targeted molecular genetics assessment following the protocol previously described [6]. Pathogenic variants were aggregated at the gene level, collapsing all alterations of a given gene into a single binary indicator. We retained features with recorded measurements for at least 30 patients and with both positive and negative labels present, yielding 25 binary gene-level mutation features for the molecular genetics branch (Fig. 1b,e).

2.4 Image pretraining

GenBloom is a patient-level transformer aggregator [4] (vision transformer, ViT) with $L=6$ layers, $H=12$ heads, embedding dimension $D=768$ (MLP hidden dim=3072). This small ViT has been shown to effectively encode slide-level information in histopathology images [3]. GenBloom operates on an unordered set of single-cell embeddings; therefore, we removed patchification (replaced by an MLP) and positional encodings of ViT to enforce permutation invariance. For each patient with up to 500 cells, we extracted per-cell embeddings using a




Fig. 1. GenBloom pretraining and genetic alignment. (a) Image pretraining cohort contains >1,500 patients and >700k single-cell images spanning major hematologic entities and cell lineages. (b) AML-Hehr mutation frequencies for genes used in the alignment. (c) Distribution of loss, gain and fusion events across chromosomal arms in AML-Hehr. (d) UMAP of GenBloom-G embeddings on AML-Hehr training patients colored by recurrent cytogenetic/molecular subtypes. (e) DINOv2-adapted image pretraining and supervised contrastive genetic alignment for GenBloom.

frozen DinoBloom-B [11] encoder and aggregated them with GenBloom, using the [CLS] token as the patient representation.

We adapted DINOv2/iBOT pretraining [13,21] to patient bags (or slides) via multi-crop subsampling. Our training pipeline operates in embedding space rather than on raw images (Fig. 1d). For each patient, we sampled $K_g=2$ global bags and $K_\ell=8$ local bags by randomly selecting 70% (≈ 350 cells) and 20% (≈ 100 cells) of cell embeddings, respectively. We trained a student (s)-teacher (t) model, where the teacher is an exponential moving average (EMA) of the student. The objective combines (i) DINO-style [CLS] alignment across views and (ii) iBOT masked embedding prediction on randomly masked cell embeddings:

$$\mathcal{L}_{img} = \mathcal{L}_{DINO} + \lambda \mathcal{L}_{iBOT},$$

$$\mathcal{L}_{DINO} = \frac{1}{|\mathcal{P}|} \sum_{(g,k) \in \mathcal{P}} \text{CE}(p_t^{(g)}, p_s^{(k)}), \quad \mathcal{L}_{iBOT} = \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \text{CE}(q_i^t, q_i^s),$$




Fig. 2. GenBloom outperforms histopathology slide encoders on hematology tasks. (a) Performance on AML-Hehr (5-class), APL-AML (2-class), and AMH (2-class) using k -NN probing ($k=5$), logistic regression, and patient retrieval (mAP@3). (b) Average performance versus pretraining scale (number of training slides; marker size indicates parameter count)

where $\mathcal{P}$ denotes all (teacher global, student view) pairs, $\mathcal{M}$ the set of masked tokens, and $\text{CE}(a, b) = -\sum_c a_c \log b_c$. This pretraining encourages GenBloom to learn robust patient representations and to model cell-composition statistics under realistic subsampling.

We trained GenBloom on a single NVIDIA H100 80GB GPU for 100 epochs with 64 batch size (≈ 7.26 kg CO₂eq emitted for all experiments).

2.5 Genetic alignment

After image pretraining, we performed supervised genetic alignment in embedding space to couple morphology with cytogenetics and molecular genetics (Fig. 1d). For each patient p , GenBloom produced a slide-level embedding from the [CLS] token, $s_p \in \mathbb{R}^{768}$. We encoded cytogenetics (karyotype) and molecular genetics (mutations) as binary vectors $y_p^k \in \{0, 1\}^{d_k}$ and $y_p^m \in \{0, 1\}^{d_m}$, respectively ($d_k = 1,104$, $d_m = 25$).

Modality-specific MLP projection heads mapped all modalities into a shared 128-dimensional ℓ_2 -normalized space:

$$z_p^s = \frac{\phi_s(s_p)}{\|\phi_s(s_p)\|}, \quad z_p^k = \frac{\phi_k(y_p^k)}{\|\phi_k(y_p^k)\|}, \quad z_p^m = \frac{\phi_m(y_p^m)}{\|\phi_m(y_p^m)\|}, \quad z_p \in \mathbb{R}^{128}.$$

We aligned modalities using a cross-modal supervised contrastive objective [9]. For anchor modality a and target modality b , let $\mathcal{P}(p) = \{j \neq p : c_j = c_p\}$ denote the set of samples in the batch (with batch size B) sharing the same class label (c_p) as patient p , excluding the anchor itself. The unidirectional loss is:

$$\mathcal{L}_{a \rightarrow b} = -\frac{1}{B} \sum_{p=1}^B \frac{1}{|\mathcal{P}(p)|} \sum_{j \in \mathcal{P}(p)} \log \frac{\exp(z_p^a \top z_j^b / \tau)}{\sum_{q=1}^B \exp(z_p^a \top z_q^b / \tau)},$$

where τ is a temperature parameter and the dot product equals cosine similarity due to ℓ_2 -normalization. Each alignment is trained symmetrically: $\mathcal{L}_{a \leftrightarrow b} = \frac{1}{2}(\mathcal{L}_{a \rightarrow b} + \mathcal{L}_{b \rightarrow a})$.

Table 1. In-domain cross-modal retrieval performance of GenBloom-G on the AML-Hehr test set. We report top-1, top-5 accuracy and mean reciprocal rank (MRR) for slide $\leftrightarrow$ karyotype (S $\leftrightarrow$ K) and slide $\leftrightarrow$ mutation (S $\leftrightarrow$ M) retrieval. * $p < 0.001$.

Metric	Direction	S $\leftrightarrow$ K		S $\leftrightarrow$ M	
		GenBloom-G	Random	GenBloom-G	Random
Top-1 Acc	$\rightarrow$	0.09 $\pm$ 0.04*	0.05 $\pm$ 0.04	0.12 $\pm$ 0.04*	0.06 $\pm$ 0.05
	$\leftarrow$	0.14 $\pm$ 0.03*	0.05 $\pm$ 0.04	0.12 $\pm$ 0.02*	0.06 $\pm$ 0.05
Top-5 Acc	$\rightarrow$	0.62 $\pm$ 0.06*	0.23 $\pm$ 0.09	0.46 $\pm$ 0.06*	0.29 $\pm$ 0.08
	$\leftarrow$	0.55 $\pm$ 0.04*	0.23 $\pm$ 0.09	0.66 $\pm$ 0.08*	0.30 $\pm$ 0.11
MRR	$\rightarrow$	0.41 $\pm$ 0.04*	0.20 $\pm$ 0.05	0.30 $\pm$ 0.03*	0.21 $\pm$ 0.05
	$\leftarrow$	0.33 $\pm$ 0.03*	0.17 $\pm$ 0.04	0.33 $\pm$ 0.02*	0.20 $\pm$ 0.05

To preserve modality-specific biological information and reduce representational collapse, we attached lightweight decoders to the bottleneck representations. These decoders reconstructed the original binary genetic vectors, where each entry indicates the presence or absence of a cytogenetic or molecular alteration. Reconstruction was supervised with a binary cross-entropy loss $\mathcal{L}_{\text{BCE}}$, which independently penalizes incorrect predictions for each genetic feature. The total objective was:

$$\mathcal{L}_{\text{gen}} = \mathcal{L}_{s \leftrightarrow k} + \mathcal{L}_{s \leftrightarrow m} + \lambda_r \mathcal{L}_{\text{BCE}},$$

where λ_r controls the reconstruction regularizer strength. This training aligned slide, karyotype, and mutation representations in a shared space, enabling retrieval and downstream prediction from any modality.

2.6 Baseline models

We compared GenBloom against 3 histopathology slide encoders—GigaPath [20], PRISM [15], and TITAN [3]—using linear probing. We also added another baseline by simply averaging DinoBloom embeddings (mean pooling). GigaPath (86.3M param.) was trained on 171,189 H&E-stained whole-slide images (WSIs). TITAN (42.1M param.) was trained on 335,645 H&E- and IHC-stained WSIs with report alignment. PRISM (99M param.) was trained on 587,196 WSIs with report alignment. For linear probing, we used logistic regression (*lbfgs* solver, regularization coefficient $C=1$) and k -NN implemented in *scikit-learn*.

2.7 Statistical analysis and metrics

Balanced accuracy (bAcc) was used for classification tasks, mean average precision (mAP) at $k = 3$ for slide-to-slide retrieval, and the F1 score for cross-modal retrieval. We performed 1,000 bootstrap iterations, randomly resampling the test set for retrieval tasks, and used the Wilcoxon signed-rank test for statistical comparisons, with a Bonferroni correction. Corrected p -values below 0.05 were considered statistically significant.

Table 2. Out-of-domain per-gene retrieval F1 score on the cAItomorph cohort for slide-to-mutation (S→M) and mutation-to-slide (M→S) directions. N : number of positive patients. Fold difference: ratio of GenBloom F1 to random F1. * $p < 0.001$.

Gene	N	S→M			M→S		
		GenBloom-G	Random	Fold diff.	GenBloom-G	Random	Fold diff.
<i>ASXL1</i>	81	0.23±0.08*	0.16±0.04	1.47	0.41±0.14*	0.16±0.04	2.56
<i>DNMT3A</i>	96	0.29±0.03*	0.19±0.04	1.53	0.23±0.11*	0.19±0.04	1.23
<i>FLT3-ITD</i>	11	0.08±0.07*	0.02±0.05	4.10	0.01±0.03	0.02±0.05*	0.67
<i>IDH2</i>	22	0.14±0.03*	0.04±0.05	3.16	0.19±0.08*	0.04±0.05	4.48
<i>JAK2</i>	213	0.57±0.04*	0.42±0.03	1.36	0.55±0.23*	0.42±0.03	1.33
<i>NPM1</i>	12	0.08±0.04*	0.02±0.05	3.30	0.10±0.08*	0.02±0.05	4.10
<i>NRAS</i>	17	0.10±0.05*	0.03±0.05	2.95	0.35±0.19*	0.03±0.05	10.63

3 Results

3.1 GenBloom improves hematological classification

We first evaluated GenBloom on downstream classification tasks using k -NN ($k=5$) and logistic regression (Fig. 2a). We denote the model after image pre-training as GenBloom-V and the model after genetic alignment as GenBloom-G.

On the AML-Hehr genetic subtyping task (5 classes), GenBloom-G achieved the highest balanced accuracy, outperforming the second best model TITAN by 38% in k -NN and 5% in logistic regression. On APL-AML dataset (2 classes), GenBloom-G outperformed the second best histopathology slide encoder PRISM by 15% in k -NN and by 45% in logistic regression. On AML vs. healthy (AMH), GenBloom-V was marginally better than TITAN in k -NN, with no difference in logistic regression. Overall, GenBloom-G ranked significantly higher than the other models across datasets for both k -NN and logistic regression (Friedman test, $p < 0.001$ for both).

We also performed a retrieval analysis to assess how well the model retrieves clinically relevant patients with the same diagnosis for a given query patient, using $k = 3$. GenBloom-G achieved the highest mAP@3, followed by GenBloom-V, on the AML-Hehr test set and the APL-AML dataset, outperforming the next-best histopathology slide encoder (TITAN) by $\approx 31\%$. On the AMH dataset, GenBloom-V achieved the best performance, followed by TITAN.

Notably, GenBloom achieved these results with substantially fewer parameters and less training data than the other slide-level models (Fig. 2b), highlighting the importance of domain-specific pretraining. Surprisingly, mean pooling of DinoBloom embeddings achieved the second- or third-best performance in linear probing and retrieval tasks on average, showing the advantage of using a hematology-specific image encoder.

3.2 Genetic alignment enables cross-modal retrieval

To assess whether GenBloom-G learns a shared embedding space across modalities, we evaluated cross-modal retrieval between slide embeddings (S), karyotype

Table 3. Ablation studies on the AML-Hehr test set. The best results per ablation are in **bold**.

Ablation	Setting	LogReg bAcc	S→K MRR	K→S MRR
Aggregator	Transformer (GenBloom-V init.)	0.83±0.03	0.38±0.10	0.22±0.03
	Transformer (Random init.)	0.82±0.03	0.36±0.09	0.21±0.05
	Mean pooling	0.78±0.04	0.28±0.03	0.19±0.02
Karyotype resolution	Band-level	0.83±0.03	0.38±0.10	0.22±0.03
	Arm-level	0.81±0.04	0.35±0.08	0.23±0.04
λ_r	$\lambda_r = 1.0$	0.83±0.03	0.38±0.10	0.22±0.03
	$\lambda_r = 0.1$	0.80±0.03	0.25±0.07	0.16±0.03
	$\lambda_r = 0.0$	0.81±0.06	0.28±0.10	0.19±0.04

embeddings (K), and mutation embeddings (M) on the held-out AML-Hehr test set (in-domain) and the cAItomorph cohort (out-of-domain). We report top-1, top-5 accuracy and mean reciprocal rank (MRR), comparing GenBloom-G against a random baseline across 1,000 bootstrap iterations (Table 1).

GenBloom-G significantly outperformed the random baseline across all retrieval directions and metrics ($p < 0.001$). For slide-to-karyotype retrieval (S→K), the model achieved a top-5 accuracy of 0.62 ± 0.06 , compared to 0.23 ± 0.09 for the random baseline, representing a $2.7\times$ improvement. Karyotype-to-slide retrieval (K→S) showed a similar pattern, with MRR increasing from 0.17 to 0.33. Slide-to-mutation (S→M) and mutation-to-slide (M→S) retrieval followed a comparable trend, with top-5 accuracy reaching 0.46 and 0.66, respectively.

In the out-of-domain setting (Table 2), we evaluated recurrent genetic alterations with clinical or biological relevance across hematologic malignancies: *ASXL1*, *DNMT3A*, *FLT3*-ITD, *IDH2*, *JAK2*, *NPM1*, and *NRAS*. GenBloom successfully retrieved relevant embeddings in both directions (mutation-to-slide and slide-to-mutation), significantly outperforming the random baseline in 13/14 cases. Notably, 186 of the 213 patients with *JAK2* mutations were diagnosed with myeloproliferative neoplasms.

3.3 Ablation studies

We ablated three design choices in the genetic alignment stage—the vision aggregator, karyotype encoding resolution, and reconstruction loss weight for genetic input λ_r —on the AML-Hehr test set (Table 3). Finetuning the pretrained transformer (GenBloom-V initialized) achieved the best overall performance (bAcc 0.83, S→K MRR 0.38, K→S MRR 0.22) compared to a randomly initialized transformer and mean pooling. Band-level karyotype encoding consistently outperformed the coarser arm-level representation (bAcc 0.83 vs. 0.81; S→K MRR 0.38 vs. 0.35), indicating that finer cytogenetic resolution provides a more informative alignment signal. Finally, the reconstruction objective improved S→K MRR by 36%, K→S MRR by 16% and logistic regression bAcc by 3%.

4 Conclusion

GenBloom is a hematology-specific slide-level encoder that learns genetically aligned patient representations. It outperforms large-scale histopathology foundation models on downstream hematology tasks, despite being trained on substantially less data, showing the benefits of domain-specific pretraining. Aligning morphology with karyotype and mutation profiles enable cross-modal retrieval between images and genetic information. This genetic alignment opens a path toward genetics-aware smear analysis, supporting faster triage, prioritization for confirmatory testing, and more informed treatment decisions in clinical workflows.

Author contributions Conceptualization: CM, MFD, FO; Data curation: CP, IL, MFD; Methodology and software: FO, MFD, IL, RMU; Writing-original draft: MFD, FO, IL; Writing-editing: RMU, CM.

Acknowledgments. CM received funding from the European Research Council under the European Union’s Horizon 2020 Research and Innovation Programme (grant agreements 866411, 101113551, and 101213822) and support from the Hightech Agenda Bayern.

Disclosure of Interests. The authors declare no competing interests.

References

1. Abrams, Z.B., Zhang, L., Abruzzo, L.V., Heerema, N.A., Li, S., Dillon, T., Rodriguez, R., Coombes, K.R., Payne, P.R.: Cytogps: a web-enabled karyotype analysis tool for cytogenetics. *Bioinformatics* **35**(24), 5365–5366 (2019)
2. Dasdelen, M.F., Kukuljan, I., Lienemann, P., Ozlugedik, F., Sadafi, A., Hehr, M., Spiekermann, K., Pohlkamp, C., Marr, C.: AI-based hematological malignancy prediction from peripheral blood smears in a large diagnostic laboratory cohort. *Leukemia* pp. 1–5 (2026)
3. Ding, T., Wagner, S.J., Song, A.H., Chen, R.J., Lu, M.Y., Zhang, A., Vaidya, A.J., Jaume, G., Shaban, M., Kim, A., et al.: A multimodal whole-slide foundation model for pathology. *Nature medicine* pp. 1–13 (2025)
4. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020)
5. Eckardt, J.N., Middeke, J.M., Riechert, S., Schmittmann, T., Sulaiman, A.S., Kramer, M., Sockel, K., Kroschinsky, F., Schuler, U., Schetelig, J., et al.: Deep learning detects acute myeloid leukemia and predicts *npml* mutation status from bone marrow smears. *Leukemia* **36**(1), 111–118 (2022)
6. Fuhrmann, I., Lenk, M., Haferlach, T., Stengel, A., Hutter, S., Baer, C., Meggen-dorfer, M., Kern, W., Haferlach, C.: Aml, nos and aml-mrc as defined by multilineage dysplasia share a common mutation pattern which is distinct from aml-mrc as defined by mds-related cytogenetics. *Leukemia* **36**(7), 1939–1942 (2022)

7. Hehr, M., Sadafi, A., Matek, C., Lienemann, P., Pohlkamp, C., Haferlach, T., Spiekermann, K., Marr, C.: Explainable ai identifies diagnostic cells of genetic aml subtypes. *PLOS Digital Health* **2**(3), e0000187 (2023)
8. Jaffe, E.S.: Pathology and genetics of tumours of haematopoietic and lymphoid tissues, vol. 3. Iarc (2001)
9. Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. In: *Advances in Neural Information Processing Systems (NeurIPS)*. vol. 33, pp. 15871–15882 (2020)
10. Khoury, J.D., Solary, E., Abal, O., Akkari, Y., Alaggio, R., Apperley, J.F., Bejar, R., Berti, E., Busque, L., Chan, J.K., et al.: The 5th edition of the world health organization classification of haematolymphoid tumours: myeloid and histiocytic/dendritic neoplasms. *leukemia* **36**(7), 1703–1719 (2022)
11. Koch, V., Wagner, S.J., Kazemina, S., Sancar, E., Hehr, M., Schnabel, J.A., Peng, T., Marr, C.: Dinobloom: a foundation model for generalizable cell embeddings in hematology. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 520–530. Springer (2024)
12. McGowan-Jordan, J., Hastings, R.J., Moore, S.: *Iscn 2020: An International System for Human Cytogenomic Nomenclature* (2020). Reprint of 'Cytogenetic and Genome Research 2020, Vol. 160, No. 7-8'. Karger Medical and Scientific Publishers (2020)
13. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193* (2023)
14. Schoch, C., Schnittger, S., Bursch, S., Gerstner, D., Hochhaus, A., Berger, U., Hehlmann, R., Hiddemann, W., Haferlach, T.: Comparison of chromosome banding analysis, interphase-and hypermetaphase-fish, qualitative and quantitative pcr for diagnosis and for follow-up in chronic myeloid leukemia: a study on 350 cases. *Leukemia* **16**(1), 53–59 (2002)
15. Shaikovski, G., Casson, A., Severson, K., Zimmermann, E., Wang, Y.K., Kunz, J.D., Retamero, J.A., Oakley, G., Klimstra, D., Kanan, C., et al.: Prism: A multi-modal generative foundation model for slide-level histopathology. *arXiv preprint arXiv:2405.10254* (2024)
16. Sidhom, J.W., Siddharthan, I.J., Lai, B.S., Luo, A., Hambley, B.C., Bynum, J., Duffield, A.S., Streiff, M.B., Moliterno, A.R., Imus, P., et al.: Deep learning for diagnosis of acute promyelocytic leukemia via recognition of genomically imprinted morphologic features. *NPJ precision oncology* **5**(1), 38 (2021)
17. Vaidya, A., Zhang, A., Jaume, G., Song, A.H., Ding, T., Wagner, S.J., Lu, M.Y., Doucet, P., Robertson, H., Almagro-Perez, C., et al.: Molecular-driven foundation model for oncologic pathology. *arXiv preprint arXiv:2501.16652* (2025)
18. Vardiman, J.W., Harris, N.L., Brunning, R.D.: The world health organization (who) classification of the myeloid neoplasms. *Blood, The Journal of the American Society of Hematology* **100**(7), 2292–2302 (2002)
19. Wang, W., Zhang, X., Xiong, Y.: Transcriptomic-guided whole-slide image classification for molecular subtype identification. *PLOS Computational Biology* **22**(2), e1013950 (2026)
20. Xu, H., Usuyama, N., Bagga, J., Zhang, S., Rao, R., Naumann, T., Wong, C., Gero, Z., González, J., Gu, Y., et al.: A whole-slide foundation model for digital pathology from real-world data. *Nature* **630**(8015), 181–188 (2024)
21. Zhou, J., Wei, C., Wang, H., Shen, W., Xie, C., Yuille, A., Kong, T.: ibot: Image bert pre-training with online tokenizer. *arXiv preprint arXiv:2111.07832* (2021)