Unsupervised Anomaly Detection on Preclinical Liver H&E Whole Slide Images using Graph based Feature Distillation

Lin Li[1] ⋆, Lillie Shelton[1], Thomas Forest[2] ,Kyathanahalli Janardhan[2], Tiffany Jenkins[2], Michael J Napolitano[2], Roujia Wang[2], David Leigh[1], Tosha Shah[1],Grady Earl Carlson[1], Rajath Soans[1], Antong Chen[1]

[1]RaDS IT, Merck & Co., Inc., Rahway, NJ, USA
[2]NonClinical Drug Safety, Merck & Co., Inc., Rahway, NJ, USA
⋆ lin.li23@merck.com

Abstract. Toxicity assessment of candidate compounds is an essential part of safety evaluation in the preclinical stage of drug development. Traditionally, drug safety evaluations depend on manual histopathological examinations of tissue sections from animal subjects, often leading to significant effort in evaluating normal tissues. Moreover, the collection of abnormality samples poses significant challenges due to the rarity and diversity of various types of abnormalities. This makes it impractical to develop a comprehensive training dataset that encompasses all potential anomalies, particularly those that are underrepresented. Consequently, traditional supervised learning methods may face difficulties, leading to a growing interest in unsupervised approaches for anomaly detection. In this study, we present GraphTox, a multi-resolution graph-based anomaly detector designed to assess hepatotoxicity in *Rattus norvegicus* liver tissues. GraphTox is built upon a novel resolution-aware foundation model pre-trained on 2.7 million liver tissue patches. Additionally, GraphTox employs graph-based feature distillation on normal liver whole slide images (WSIs) to identify hepatotoxicity. Our results demonstrate that GraphTox achieves an 11.1% improvement in area under the receiver operating characteristic curve (AUC) on an independent testing set compared to the best-performing non-graph-based anomaly detection models, and an 8.1% improvement over a graph-based model derived from a resolution-agnostic foundation model $UNIv2$. These findings highlight that GraphTox effectively leverages the resolution-aware digital pathology foundation model to capture multi-scale tissue characteristics within the local tissue graphs, thereby enhancing anomaly detection across various scales[1].

Keywords: Unsupervised · graph · foundation model · digital pathology · anomaly detection · preclinical safety evaluation · drug development

[1] Our code is available at https://linlilamb.github.io/GraphTox-project-page/

1 Introduction

In drug development, the evaluation of histology slides from animal subjects, such as *Rattus norvegicus*, is vital for safety assessments [1]. Analyzing these histology slides is crucial for identifying potential adverse effects of new compounds on biological tissues, ensuring that only safe and effective drugs advance to clinical trials [1]. However, the manual analysis of histological slides is time-consuming and prone to human error, which may prolong drug development timelines and increase the risk of missing significant pathological changes [2].

To address this challenge, our study explores automated methodologies for detecting tissue anomalies, thereby accelerating the drug development process. In preclinical toxicity evaluation, gathering data on abnormal findings is complex due to the inherent rarity and heterogeneity of these abnormalities. This makes it impractical to create a comprehensive training dataset for all potential anomaly classes, particularly those that are underrepresented. Consequently, traditional supervised learning methods may face difficulties, leading us to focus on unsupervised methods for anomaly detection.

Unsupervised anomaly detection models aim to identify patterns and signals that diverge from the established distribution of normal samples [3]. Both image reconstruction and feature distillation (FD) approaches have been investigated for anomaly detection in medical imaging applications [3]. Image reconstruction approaches focus on identifying potential anomalies by examining the discrepancies between reconstructed images and original images. For example, s2-AnoGAN [14] leverages generative adversarial network (GAN) to learn how to reconstruct normal images, subsequently flagging abnormal samples when reconstruction fails. In contrast, FD-based methods train a student encoder to generate embeddings that closely resemble those produced by a pre-trained teacher encoder, using only normal samples [4, 5]. As a result, the student encoder may find it difficult to replicate the teacher encoder's output for abnormal samples, which ultimately triggers alerts for potential anomalies. Furthermore, the recent emergence of digital pathology (DP) foundation models shed lights for further advanced FD-based anomaly detection in H&E whole slide images (WSIs) by providing encoders pre-trained on diverse and large-scale datasets [6–8].

Given the gigapixel resolution of WSIs and their hierarchical organization at various magnifications, e.g. 5× (2 μm/pixel), 10× (1 μm/pixel) and 20× (0.5 μm/pixel), it is necessary to generate smaller tissue patches (e.g., 256 × 256 pixels) at selected magnifications before inputting them into FD models. This design narrows the focus of FD to extracting tissue information at specific magnifications. However, histopathological anomalies can appear across various magnifications, and pathologists often rely on the multi-resolution tissue morphology to inform their decisions (Fig. 1). For instance, severe necrosis may be readily observed at 5×, whereas vacuolation, the textural change with bubbles present, requires higher magnification (e.g., 10×) for effective evaluation (Fig. 1). As current DP foundation models are typically resolution-agnostic, which are trained on tissue patches of a single magnification or mixed patches of varying magni-

fications [6–8], their ability to capture resolution-specific tissue features may be compromised.

To address these limitations, we propose GraphTox, a graph-based FD model for unsupervised anomaly detection in liver WSIs. Our contributions include:

(1) GraphTox incorporates a resolution-aware DP foundation model pretrained on 2.7 million liver tissue patches across 5×, 10×, and 20× to provide multi-scale tissue embeddings for graph-based FD.

(2) GraphTox creates multi-resolution tissue patch graphs to integrate tissue information across various resolutions, rather than distilling tissue morphology without considering resolution.

(3) We conduct experiments for GraphTox on hepatotoxicity, which is a leading cause of adverse drug reactions. Our results demonstrate that GraphTox effectively harnesses information across multiple resolutions in WSIs, enabling it to capture anomalies at various scales.

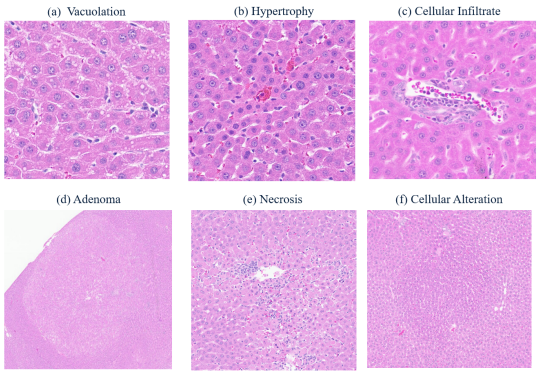

Fig. 1. Anomaly examples. (a) Vacuolation: Textural change with bubbles, best observed at 20×. (b) Hypertrophy: Size variation in hepatocytes, easier to detect at 10×. (c) Cellular infiltrate: Inflammatory cells visible at 10×. (d) Adenoma: Expansile structure compressing surrounding tissue, observable at 5×. (e) Necrosis: Pinkish regions infiltrated with inflammatory cells, detectable at 5×. (f) Cellular alteration: Smaller, bluer hepatocytes, slightly expansile, challenging to observe by human eye.

2 Methods

Our proposed model, GraphTox, is developed based on two main components: (1) a resolution-aware DP foundation model (Fig. 2 (a)-(b)) and (2) a graph based FD for anomaly detection (Fig. 2(c)-(d)).

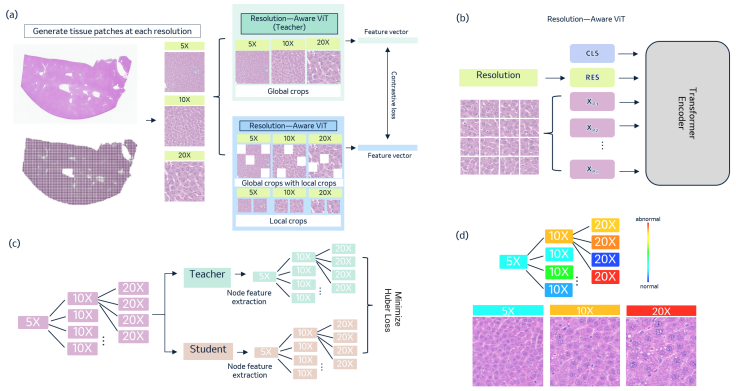

Fig. 2. Overview of GraphTox: (a) A resolution-aware digital pathology foundation model; (b) A resolution-aware vision transformer; (c) GraphTox training, where the teacher model is a pre-trained resolution-aware DP foundation model and the student model, sharing the same architecture, is trained exclusively on normal tissue samples. GraphTox inputs are local multi-resolution tissue graphs. (d) GraphTox inference enables anomaly detection across multiple resolutions at the same location. For example, tissue morphology observed at 5× may not indicate significant abnormalities, while 10× and 20× can reveal potential anomalies like hypertrophy and vacuolation.

2.1 Resolution-aware Foundation model

DP foundation models have been trained on millions of small local tissue regions, or called patches, at selected magnifications, e.g. 20×, using self-supervised learning [6–8]. However, tissue morphology varies at different resolutions and contributes differently to the decisions made by the pathologists (Fig.1) . Therefore, instead of treating the patches independently of their resolution, we developed a resolution-aware foundation model (F_{RA}) to extract tissue features conditioned on the resolution of the patches. As most recent DP foundation models [6–8] are built on DINOv2 [12], a vision transformer (ViT) based knowledge distillation architecture, we replace the traditional ViT with a resolution-aware ViT by adding a learnable resolution embedding to condition the feature learning with resolution information (Fig.2).

Specifically, each input 2D tissue patch $x \in \mathbb{R}^{H\times W\times C}$ is reshaped into a sequence of flattened 2D patches $t_p \in \mathbb{R}^{n\times(P^2C)}$, where H, W refers to the height and width of the image, C is the number of channels, P is the ViT patch size, and $n = HW/P^2$ referring to the number of ViT patches. Then, t_p are mapped to a embedding space with dimensions as D using a trainable linear projection, producing patch embeddings $x_p \in \mathbb{R}^{n\times D}$. The resolution embedding $x_r \in \mathbb{R}^D$ is determined by a learnable dictionary $r = [x_0, x_1, ..., x_m], r \in \mathbb{R}^{m\times D}$, where m is the number of resolutions. In addition, we include a learnable embedding as a class token $x_{class} \in \mathbb{R}^D$. Position embeddings $E_{pos} \in \mathbb{R}^{(n+2)\times D}$ are added to inform the positional information in the transformer encoder (Eq. 1). z_0 denotes the first layer in the transformer encoder which is followed up with alternating

layers of multiheaded self-attention, layernorm, and multi-layer perceptron as described in ViT [17].

$$z_0 = [x_{class}, x_r, x_p^1, x_p^2, ..., x_p^n] + E_{pos} \tag{1}$$

Similar to DINOv2 [12] , F_{RA} is trained with a discriminative self-supervised method using loss function shown in Eq.2. The p_t and p_s refer to the teacher and student class token respectively, p_{ti} and p_{si} are the teacher and student patch tokens and i are patch indices for masked tokens. The $d_{n,i}$ is a distance measure to evaluate the embedding distance between nth and ith sample in a batch, where $n \neq i$.

$$L = -(\sum p_t \log(p_s) + \sum_i p_{ti} \log(p_{si}) + \frac{1}{N} \sum_{i=1}^{n} \log(d_{n,i})) \tag{2}$$

2.2 Multi-resolution Graph-based Feature Distillation

After pre-training F_{RA}, we further develop GraphTox using FD. FD-based anomaly detection models train a student encoder to learn to mimic teacher encoder embeddings solely from normal samples. Consequently, the student encoder may struggle to match the teacher's outputs when faced with abnormal samples. Since anomalies in WSIs often span multiple scales, it is crucial to train the FD-based anomaly detection model with patches from various resolutions. Additionally, while tissue morphology varies by resolution, features in patches at the same spatial locations are interrelated rather than independent. Therefore, during the FD process, it is vital to integrate information across all resolutions at the same spatial locations in a learnable manner. To achieve this, we first construct multi-resolution tissue graphs locally and then optimize the FD by minimizing the difference between the two multi-resolution graphs at the same location, using student and teacher embeddings as the node features respectively. To create the local multi-resolution tissue graph, we generated non-overlapping tissue patches at 5×. For each 5× patch, we identify four (2-by-2) nearest-neighbor patches at 10× and sixteen (4-by-4) nearest-neighbor patches at 20×. All the patches are connected to its corresponding lower resolution node as illustrated in Fig.2 (c). Consequently, the loss function of the GraphTox can be illustrated as Eq. 3.

$$L_{FD} = \sum_{i}^{n} S(G_{x_s^i}, G_{x_t^i}). \tag{3}$$

The $G_{x_s^i}$ denotes the ith node's student embedding and $G_{x_t^i}$ is the ith node's teacher embedding. S denotes the distance measurement and we use Huber loss [13] as S in this study. L_{FD} denotes the loss function to be minimized during training. Only ViT patch embeddings are used as the teacher and student embedding during optimization. In addition, L_{FD} will be used as an anomaly score during inference in the testing set. The larger the anomaly score of a patch, the higher the likelihood that it contains abnormal signals.

3 Experiments and Results

Our DP foundation model F_{RA} was trained on 2,701,608 patches sampled from 819 *Rattus norvegicus* subjects. These patches were equally distributed across 5×, 10× and 20×. The backbone of F_{RA} is ViT-L/16 with an additional resolution token. The pre-trained F_{RA} is used as the teacher encoder in GraphTox, where the architecture of the student encoder is the same as the teacher encoder, a resolution-aware ViT-L/16. The anomaly detection training dataset contains 566,784 tissue patches sampled from 40 normal liver WSIs across 5×, 10× and 20×. In addition, 28,738 local multi-resolution graphs were created in the training set to establish graph-based FD model development (Fig. 2 (c)). We also collect 158 liver WSIs as an independent testing set with 74 WSIs marked as anomaly and 84 WSIs without remarkable anomaly findings. The tissue patches with size 256×256 pixels will be resized to 224 × 224 as needed for DP foundation model feature extraction.

Table 1. Mean AUC and 95% confidence interval of anomaly detection models on independent testing set. The 'combined' results use max pooling to select the highest anomaly score across three single-resolution FD models, meaning different resolutions may be chosen per WSI.

Types	Models	5×	10×	20×	Combine
FD	$UNI_{res_{single}}$	0.64 (0.60 - 0.69)	0.71 (0.67 - 0.75)	0.55 (0.50 - 0.60)	0.55 (0.50 - 0.6)
	$UNI_{res_{mix}}$	0.71 (0.66 - 0.75)	0.69 (0.64 - 0.73)	0.65 (0.61 - 0.69)	0.70 (0.66 - 0.74)
	$UNIv2_{res_{mix}}$	0.73 (0.69 - 0.77)	0.73 (0.69 - 0.78)	0.70 (0.66 - 0.74)	0.72 (0.68 - 0.77)
	$RA_{res_{mix}}$	0.73 (0.69 - 0.77)	0.67 (0.63 - 0.72)	0.69 (0.65 - 0.74)	0.71 (0.67 - 0.76)
GAN	s2-AnoGAN	-	-	-	0.61 (0.56-0.65)
Graph-FD	$UNIv2_{graph}$	-	-	-	0.74 (0.70 - 0.78)
	$RA_{transformer}$	-	-	-	0.74 (0.70 - 0.78)
	GraphTox	-	-	-	**0.80(0.76 - 0.84)**

3.1 FD and GAN based anomaly detection models

To evaluate the contribution of the multi-scale tissue information for anomaly detection, we implemented UNI [6], $UNIv2$ [6] and our DP foundation model F_{RA} as teacher encoders, training student encoders using the same data without constructing local multi-resolution tissue graphs. The UNI model utilizes ViT-L/16 architecture, while $UNIv2$ employs ViT-H/14-reg8. For both UNI-based models, the student encoders are of ViT-L/16 and ViT-L/14, respectively.

First, we trained the anomaly detection model ($UNI_{res_{single}}$) using UNI as the teacher encoder with tissue patches solely at 10×. We then validated $UNI_{res_{single}}$ on an independent testing set using tissue patches at 5×, 10×, and 20×, combining anomaly scores from all resolutions through max pooling. Next, we trained the anomaly detection model ($UNI_{res_{mix}}$) with UNI as the teacher encoder on all tissue patches from 5×, 10×, and 20× in the training set, following the same validation strategy. Additionally, with the recent release of $UNIv2$,

we applied the same training approach to develop $UNIv2_{res_{mix}}$. Finally, we employed pre-trained F_{RA} as the teacher encoder to create the anomaly detection model ($RA_{res_{mix}}$), assessing whether resolution-conditioned tissue characteristics enhance anomaly detection across different resolutions.

As shown in Table 1, all anomaly detection models trained with mix resolution tissue patches led to a larger WSI-level area under the receiver operating characteristic curve (AUC) than $UNI_{res_{single}}$ when combining all prediction results at all resolutions. This result suggests that anomaly detection on WSIs requires multi-scale tissue information. Since $UNI_{res_{single}}$ was trained on the tissue patches at 10×, it performed best on 10× compared to the prediction results at 5× and 20×. In addition, $UNIv2_{res_{mix}}$ demonstrated better results in tissue patches in all resolutions compared to $UNI_{res_{mix}}$, indicating that $UNIv2$ has better capacity to capture multi-scale tissue information than UNI. Additionally, $UNIv2_{res_{mix}}$ outperformed $UNI_{res_{mix}}$ across all resolutions, highlighting $UNIv2$'s superior ability to capture multi-scale tissue information. Meanwhile, $RA_{res_{mix}}$ showed comparable results to $UNIv2_{res_{mix}}$ across all resolutions but had a lower AUC at 10×. Despite leveraging multi-scale tissue information by mixing patches during training and selecting the highest anomaly score across resolutions, the ability to fully utilize multi-scale tissue information remains limited. Therefore, the next section will discuss our efforts in enhancing multi-scale tissue information usage through graph-based FD.

In addition to FD-based methods, we evaluated s2-AnoGAN [14] using the same training and testing sets. It has been shown that s2-AnoGAN exhibits superior performance in detecting DP anomalies compared to f-AnoGAN [9] and pg-AnoGAN [15]. Additionally, our previous study [16] suggests that using the structural similarity index measure (SSIM) to quantify reconstruction quality enhances DP anomaly detection. Therefore, we used SSIM as the anomaly score for performance evaluation on the testing set. As shown in Table 1, FD-based methods outperformed GAN-based method, indicating the significant contribution of the DP foundation models.

3.2 Graph-FD-based anomaly detection models

Instead of treating tissue patches at various resolutions as independent entities, we proposed the creation of local multi-resolution graphs and distilling the features of the teacher encoder to the student encoder at the graph level. GraphTox utilizes F_{RA} as the teacher encoder and trains the student encoder on the local multi-resolution graphs. Our results demonstrate that GraphTox enhances anomaly detection performance compared to $RA_{res_{mix}}$ with 12.7% increase in AUC on testing set (Table 1), highlighting the effectiveness of the graph-based FD. Moreover, GraphTox has the capability to identify the anomaly with reflecting the corresponding resolutions (Fig. 2 (d)). Additionally, we investigated the use of $UNIv2$ as the teacher encoder for the graph-based FD ($UNIv2_{graph}$). The results indicate that GraphTox outperforms $UNIv2_{graph}$ on the testing set, suggesting that F_{RA} enables the student encoder to capture resolution-specific

features, thus enhancing its ability to detect multi-scale anomalies. Both GraphTox and $UNIv2_{graph}$ demonstrate improved WSI-level anomaly detection compare to the $RA_{res_{mix}}$ and $UNIv2_{res_{mix}}$, which are trained purely with mixed resolution tissue patches without using a graph model (Table 1). Fig. 3 illustrates an WSI containing multiple anomaly regions. GraphTox and $UNIv2_{graph}$ show superior localization of anomalies compared to the mixed-resolution models (Fig. 3 (b)-(e)). Notably, GraphTox detects a greater number of anomalies than $UNIv2_{graph}$, e.g. vacuolation (Fig. 3 (f) orange box). Furthermore, we explored the inclusion of a graph transformer in the student model ($RA_{transformer}$) to facilitate message passing across nodes of local tissue graphs during training. This approach improved the anomaly detection performance on the testing set compared to $RA_{res_{mix}}$. However, GraphTox demonstrated a significant improvement over $RA_{transformer}$ by removing the graph transformer in the student model. This finding suggests that the graph transformer may inadvertently dilute the anomaly signal flagged at individual patches by learning to mimic the teacher's output for the current node based on information from surrounding nodes.

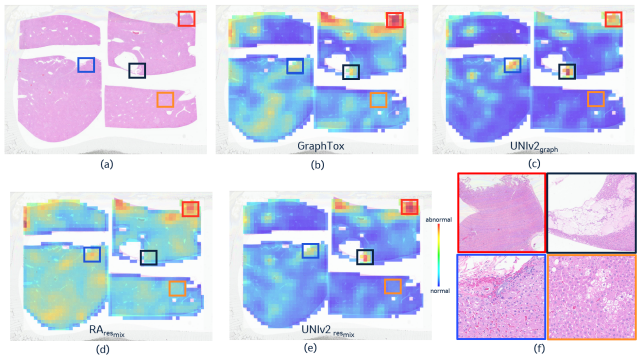

Fig. 3. Anomaly Detection on a whole slide image. (a) A liver whole slide image of a *Rattus norvegicus* study. (b) Anomaly detection heatmap of GraphTox; (c) Heatmap of $UNIv2_{graph}$. (d) Heatmap of $RA_{res_{mix}}$. (e) Heatmap of $UNIv2_{res_{mix}}$. (f) The zoomed in visualizations of the anomaly regions. Red: dried-out artifact; Black: dilatation; Blue: vacuolation; Orange: vacuolation. Both $UNIv2_{graph}$ and $UNIv2_{res_{mix}}$ missed the vaculation at the region highlighted in orange box.

4 Conclusion

In summary, this paper proposes an innovative approach for unsupervised anomaly detection on preclinical liver WSIs. Our model, GraphTox, achieving a novel resolution-aware digital pathology foundation model and graph-based feature distillation, illustrating state-of-the-art performance on hepatotoxicity detection. GraphTox achieves an 11.1 % improvement in AUC on an independent testing set compared to the best-performing non-graph models , and an 8.1% improvement

over the UNI v2 graph-based anomaly detection model. These results demonstrate that our resolution-aware digital pathology foundation model effectively utilizes the multi-scale tissue characteristics present in the local multi-resolution tissue graph, thereby enhancing anomaly detection across various scales. Our method has only been experimented to a single organ type. Future work will aim to extend anomaly detection capabilities to multiple organ types.

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

References

1. Waring, Michael J and Arrowsmith, John and Leach, Andrew R and Leeson, Paul D and Mandrell, Sam and Owen, Robert M and Pairaudeau, Garry and Pennie, William D and Pickett, Stephen D and Wang, Jibo and others, "An analysis of the attrition of drug candidates from four major pharmaceutical companies", *Nature reviews Drug discovery*, vol. 14, no. 7, pp. 475–486, 2015.
2. Weaver, Richard J and Blomme, Eric A and Chadwick, Amy E and Copple, Ian M and Gerets, Helga HJ and Goldring, Christopher E and Guillouzo, Andre and Hewitt, Philip G and Ingelman-Sundberg, Magnus and Jensen, Klaus Gjervig and others, "Managing the challenge of drug-induced liver injury: a roadmap for the development and deployment of preclinical predictive models", *Nature Reviews Drug Discovery*, vol. 19, no. 2, pp. 131–148, 2020.
3. Bao, Jinan and Sun, Hanshi and Deng, Hanqiu and He, Yinsheng and Zhang, Zhaoxiang and Li, Xingyu, "Bmad: Benchmarks for medical anomaly detection", *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4042–4053, 2024.
4. Bergmann, Paul and Fauser, Michael and Sattlegger, David and Steger, Carsten, "Uninformed students: Student-teacher anomaly detection with discriminative latent embeddings", *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 4183–4192, 2020.
5. Cao, Yunkang and Wan, Qian and Shen, Weiming and Gao, Liang, "Informative knowledge distillation for image anomaly segmentation", *Knowledge-Based Systems*, vol. 248, pp. 108846, 2022.
6. Chen, Richard J and Ding, Tong and Lu, Ming Y and Williamson, Drew FK and Jaume, Guillaume and Chen, Bowen and Zhang, Andrew and Shao, Daniel and Song, Andrew H and Shaban, Muhammad and others, "Towards a General-Purpose Foundation Model for Computational Pathology", *Nature Medicine*, 2024.
7. Xu, Hanwen and Usuyama, Naoto and Bagga, Jaspreet and Zhang, Sheng and Rao, Rajesh and Naumann, Tristan and Wong, Cliff and Gero, Zelalem and González, Javier and Gu, Yu and Xu, Yanbo and Wei, Mu and Wang, Wenhui and Ma, Shuming and Wei, Furu and Yang, Jianwei and Li, Chunyuan and Gao, Jianfeng and Rosemon, Jaylen and Bower, Tucker and Lee, Soohee and Weerasinghe, Roshanthi and Wright, Bill J. and Robicsek, Ari and Piening, Brian and Bifulco, Carlo and Wang, Sheng and Poon, Hoifung, "A whole-slide foundation model for digital pathology from real-world data", *Nature*, 2024.
8. Vorontsov, Eugene and Bozkurt, Alican and Casson, Adam and Shaikovski, George and Zelechowski, Michal and Severson, Kristen and Zimmermann, Eric and Hall,

James and Tenenholtz, Neil and Fusi, Nicolo and others, "A foundation model for clinical-grade computational pathology and rare cancers detection", *Nature medicine*, pp. 1–12, 2024.
9. Schlegl, Thomas and Seeböck, Philipp and Waldstein, Sebastian M and Langs, Georg and Schmidt-Erfurth, Ursula, "f-AnoGAN: Fast unsupervised anomaly detection with generative adversarial networks", *Medical image analysis*, vol. 54, pp. 30–44, 2019.
10. Zavrtanik, Vitjan and Kristan, Matej and Skočaj, Danijel, "Draem-a discriminatively trained reconstruction embedding for surface anomaly detection", *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 8330–8339, 2021.
11. Chen, Liyang and You, Zhiyuan and Zhang, Nian and Xi, Juntong and Le, Xinyi, "UTRAD: Anomaly detection and localization with U-transformer", *Neural Networks*, vol. 147, pp. 53–62, 2022.
12. Oquab, Maxime and Darcet, Timothée and Moutakanni, Théo and Vo, Huy and Szafraniec, Marc and Khalidov, Vasil and Fernandez, Pierre and Haziza, Daniel and Massa, Francisco and El-Nouby, Alaaeldin and others, "Dinov2: Learning robust visual features without supervision", *arXiv preprint arXiv:2304.07193*, 2023.
13. Huber, Peter J, "Robust estimation of a location parameter", *Breakthroughs in statistics: Methodology and distribution*, pp. 492–518, 1992.
14. Pocevičiūtė, Milda and Eilertsen, Gabriel and Lundström, Claes, "Unsupervised anomaly detection in digital pathology using GANs", *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pp. 1878–1882, 2021.
15. Karras, Tero and Aila, Timo and Laine, Samuli and Lehtinen, Jaakko, "Progressive growing of gans for improved quality, stability, and variation", *arXiv preprint arXiv:1710.10196*, 2017.
16. Shelton, Lillie and Soans, Rajath and Shah, Tosha and Forest, Thomas and Janardhan, Kyathanahalli and Napolitano, Michael and Gonzalez, Raymond and Carlson, Grady and Shah, Jyoti K and Chen, Antong, "Automated anomaly detection in histology images using deep learning", *SPIE Medical Imaging 2024: Digital and Computational Pathology*, vol. 12933, pp. 152–158, 2024.
17. Dosovitskiy, Alexey and Beyer, Lucas and Kolesnikov, Alexander and Weissenborn, Dirk and Zhai, Xiaohua and Unterthiner, Thomas and Dehghani, Mostafa and Minderer, Matthias and Heigold, Georg and Gelly, Sylvain and others, "An image is worth 16x16 words: Transformers for image recognition at scale", *arXiv preprint arXiv:2010.11929*, 2020.