Review

Deep Learning and Its Applications in Computational Pathology

Runyu Hong *[ID] **and David Fenyö** *[ID]

Institute for Systems Genetics, NYU Grossman School of Medicine, New York, NY 10016, USA
* Correspondence: Runyu.Hong@nyu.edu (R.H.); david@fenyolab.org (D.F.)

Abstract: Deep learning techniques, such as convolutional neural networks (CNNs), generative adversarial networks (GANs), and graph neural networks (GNNs) have, over the past decade, changed the accuracy of prediction in many diverse fields. In recent years, the application of deep learning techniques in computer vision tasks in pathology has demonstrated extraordinary potential in assisting clinicians, automating diagnoses, and reducing costs for patients. Formerly unknown pathological evidence, such as morphological features related to specific biomarkers, copy number variations, and other molecular features, could also be captured by deep learning models. In this paper, we review popular deep learning methods and some recent publications about their applications in pathology.

Keywords: deep learning; machine learning; histopathology; computational pathology; convolutional neural networks; generative adversarial networks

1. Introduction

With the development of artificial intelligence and machine learning techniques in the past decade, many deep-learning-based computer vision models are playing important roles in daily life, and have revolutionized various industries through their superior performance and efficiency in prediction tasks, such as autopilot, machine translation, electronic sports, and biometry [1–6]. Recently, these technologies have also shown their extraordinary potential and capabilities in solving many complicated questions in the biomedical field by analyzing massive amounts of biomedical data, such as protein structural predictions with Alphafold which outperforms experimental results [7,8], and tumor segmentation in MRI scans [1]. In particular, computational pathology, a discipline that involves the effort of both pathologists and informaticians, has especially benefitted from the advancement of deep learning in recent years [9–12]. Several models have also been demonstrated to be useful in clinical diagnoses based on histopathology images [13–15]. In addition, some model-extracted morphological features show correlations with features at a molecular level, including single mutations and subtypes, most of which are previously unknown to human pathologists and clinicians [14]. Here, we discuss these deep learning methods from a technical perspective and summarize their successful applications in pathology from recent publications.

2. Deep Learning Techniques

2.1. Convolutional Neural Networks

Deep learning is a type of machine learning method that is using a multi-layer perceptron called artificial neural networks (ANN) [1,2,16]. Training a deep learning model involves designing and selecting a neural network architecture, loss functions, and evaluation metrics, as well as tuning the hyperparameters of batch size, step size, and regularization methods [1,2,17,18]. Convolutional neural networks (CNNs), variants of ANNs, have proved their power in tackling various computer vision tasks, such as image classification, segmentation, and object detection [19–23]. The first modern CNN architecture, LeNet5, was introduced by Yann LeCun et al. in 1998 [24]. This gradient-based six-layer

convolutional neural network shows its power in recognizing hand-written digits and characters [24]. However, the development of CNNs was restricted by limited computational compacities and resources for over a decade. The advancement of computational hardware in recent years, especially graphical processing units (GPUs) and tensor processing units (TPUs), empowers the development of deep neural networks. Many CNN architectures, such as AlexNet [25], VGG [26], InceptionNet [20], and ResNet [27], can be trained into models that even outperform human beings in a computer vision classification challenge called ImageNet, which contains 1.2 million high-resolution images of more than 1000 classes [19,27–30].

AlexNet was introduced in 2012. This architecture is much larger than the previous LeNet5, with 650,000 neurons and 60 million trainable parameters packed into this design, with 5 convolutional and 3 fully connected layers [25]. Overlapping max pooling, ReLU nonlinearity, and dropout regularization are also incorporated. Due to the development of hardware, AlexNet at that time could only be trained and run on two GPUs [25]. It achieved a top-five test error rate of 15.3%, which made it the winner of the ILSVRC-2012 competition and outperformed the second-best model by more than 10% [25]. The overwhelming success of AlexNet drew people's attention back to CNNs, and numerous other new architectures, including the VGG, InceptionNet and ResNet, were developed in the following years.

VGG architecture was introduced in 2014, when it won the ImageNet challenge [26]. Compared with AlexNet, VGG increases the depth of the model by adding more convolutional layers with smaller convolutional filters [26]. However, with the introduction of newer architectures in the following years, VGG architecture has lost its popularity due to the gigantic size, high complexity to train, and less accurate performance.

InceptionV1 architecture was announced in 2015, with the name GoogLeNet, which is a 22-layer deep CNN (Figure 1) [20]. The two key innovations that make Inception architectures outstanding are the inception module and auxiliary classifier. The inception module consists of multiple convolutional kernels with different sizes on the same layer [20]. This design allows the model to capture similar features of various sizes. Deep CNNs are prone to overfitting and passing gradient updates through the entire network is hard, which is often referred to as the vanishing gradient problem. By adding auxiliary classifiers in the middle of the network, the auxiliary loss from the middle of the model is taken in the final loss calculation, so that the gradients also represent the middle part of the network [20]. InceptionV2 and InceptionV3 were introduced in 2016, which modified the inception module by factorizing the larger kernels into a stack of smaller kernels to make the architecture more computationally efficient [28]. In addition, InceptionV3 uses the RMSProp optimizer and adds batch normalization into the auxiliary classifiers, significantly improving the performance, with a top-five error of 3.57% and top-one error of 17.2% on ImageNet, much better than a human [28]. InceptionV4 further refined the architecture by adding reduction blocks and unifying the inception modules [29].

A major competitor of InceptionNet is Resnet, which applies the idea of residual connection (Figure 2) [27]. In this architecture, each layer learns the residuals from the previous layer with reference to the layer inputs [27]. The top-five error on ImageNet is 3.57%, which is similar to the performance of InceptionV3 [27]. Interestingly, this residual connection idea was later adapted by the InceptionNet team to develop InceptionResNetV1, a modified version of InceptionV3, and InceptionResNetV2, a modified version of InceptionV4 [29]. Using the residual connection, InceptionResNetV2 achieved a markedly improved 3.1% top-five error on ImageNet [29].
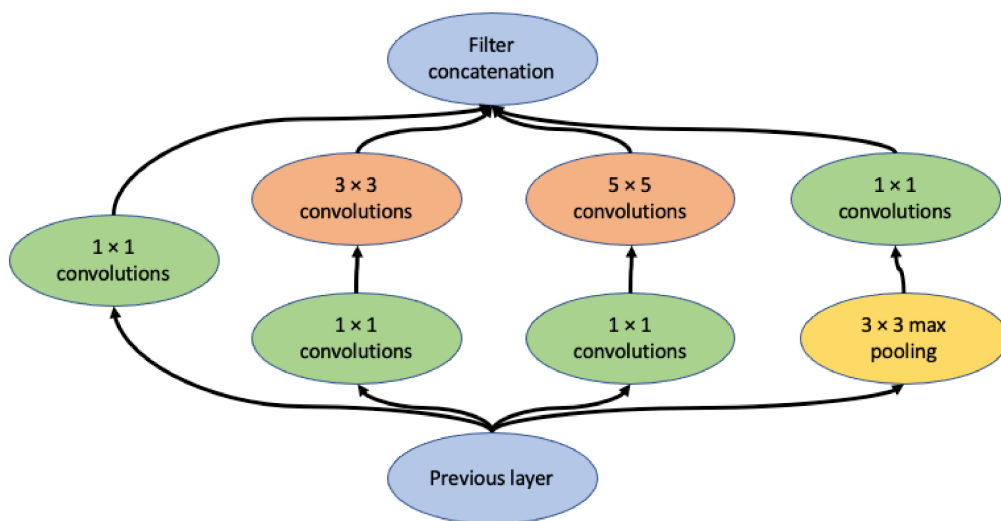
Figure 1. Diagram of the inception module [19], containing 4 branches with convolutional kernels of different sizes.
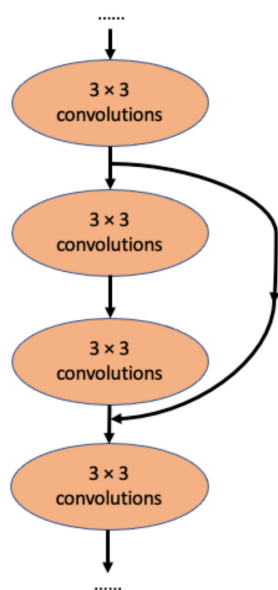
Figure 2. The concept of residual connection in ResNet [27]. Here, the middle 2 convolutional layers are skipped by the residual connection.

2.2. Visualization of CNN Models

Ever since the introduction of deep neural networks, people have been eager to know what their models have learned [31]. For image-based tasks with CNN models, visualizing the captured features is the most straightforward way. Class activation mapping (CAM) and saliency maps are two simple ways to visualize the learned features by projecting the weights and gradients of the output layer back to the input image [32–34]. However, these visualization methods are image-specific and will only roughly imply where the models are focusing. In addition, many saliency methods have been criticized recently for giving misleading visualization interpretations, and researchers are advised to use them with caution [35]. To unveil the CNN models further, direct deconvolution and indirect optimization are the two major approaches [36]. Deconvolution starts with finding an image from the dataset that triggers high activity to the neuron of interest and the gradient of neuron activity is calculated [36]. In general, a deconvolutional network is a reversed convolutional network, which maps features back to pixels [37]. However, deconvolution

visualization can be noisy and may contain features that are not easy to interpret [38]. The indirect optimization approach can provide more accurate visualization than that from deconvolution [39]. The algorithms optimize the colors of the pixels of an image to maximize activation of the neuron of interest [38–40]. Once a set of optimized images for many neurons has been obtained, a dimensional reduction visualization method, such as UMAP and tSNE, can create an atlas that systematically displays the correlations of features captured by different neurons at the same layer [38,41–43].

2.3. Graph Neural Networks

Graph neural networks (GNNs) are a type of neural network which deal with data consisting of relational information [44]. Data with a non-Euclidean structure of information, such as particle interactions, molecular structures, and object relationships in images, could be modeled by GNNs [45]. In general, GNNs can be further classified into four categories: recurrent GNNs, convolutional GNNs, graph autoencoders, and spatial–temporal GNNs [45].

2.4. Generative Adversarial Networks

Generative adversarial networks (GANs) are a type of neural network consisting of two networks that are trained at the same time [46]. The generator part is trained to create fake images which tries to fool the discriminator, while the discriminator, trained with both real and generated fake images, is able to distinguish them [46]. Many variants of GANs have been applied to different tasks, such as style transfer, the visualization of neural networks, and object segmentation [46–53]. Cycle-GAN, a GAN variant using cycle-consistence loss to train two pairs of generators and discriminators simultaneously, has become increasingly popular for image-to-image translation tasks [50]. Unlike conditional GANs, which require two styles of paired data, cycle-GANs only need two sets of images of two styles, which significantly lowers the data requirements while preserving the quality of style transfer [50].

3. Applications in Computational Pathology

3.1. Classification and Feature Prediction

With the success of CNN models in various real-world computer-vision classification tasks, researchers and scientists have also trained and tested these models in case scenarios in biomedical fields, including pathology. These studies may involve training an existing CNN architecture from scratch. However, it requires more data, and the data augmentation techniques may not always be suitable for biomedical images. Alternatively, transfer learning techniques, which freeze most of the parameters from a model often pre-trained on ImageNet, have more advantages in terms of the data size requirements. For example, an InceptionV3-based ImageNet pre-trained CNN model can achieve a high level of accuracy in determining skin lesion malignancy and the possibility of melanoma [13,54].

In clinical settings, pathologists typically examine histopathology slides under microscopes to provide diagnosis or other clinical information. Due to the development of digital pathology equipment, digitizing histopathology slides is cheaper and more accessible. As a result, more and more deidentified digital histopathology slide images have become available in many databases. These images, often with extremely large dimensions, are saved in special image file formats (e.g., .svs or .scn), which is a tuple of the same image with different resolutions [55]. Thus, in order to fit these digital histopathology images into CNN architectures, people usually develop their own customized pipelines with commonly used techniques, such as tiling the whole slide images (WSI) or sampling regions of interest (ROIs) (Figure 3) [56]. In the past few years, classification CNN models trained on histopathology images have shown phenomenally high performance and promising clinical potential in predicting both morphological features and molecular features. The visualization techniques also reveal results that often match pathologists' expectations and many models are generalizable to independent real-world clinical images. For example,

Inception and InceptionResNet architectural models demonstrate high accuracy and other statistical metrics in predicting subtypes and key biomarker mutations, such as STK11 and EGFR, in non-small-cell lung cancer histopathology slides [14,57,58]. With the integration of other critical clinical variables and images, immune response, G-CIMP, and telomere length can be predicted in glioblastoma patients [59]. BRAF mutation, a well-known biomarker in malignant melanoma, can also be accurately predicted with a CNN-based model [60]. Other molecular and genomic features, such as microsatellite instability (MSI), can be predicted from histopathology slides with a reasonable accuracy as well [15]. The critical gene expression level could also be inferred by applying these CNN classification models to WSI [61]. Some contemporary models also show successful classification results in the histopathology images of multiple tissue types [62,63]. These successful cases indicate that CNNs represent a suitable approach to study the correlation between molecular features and morphological features in histopathology slides, some of which may be undetectable or often ignored by human pathologists.
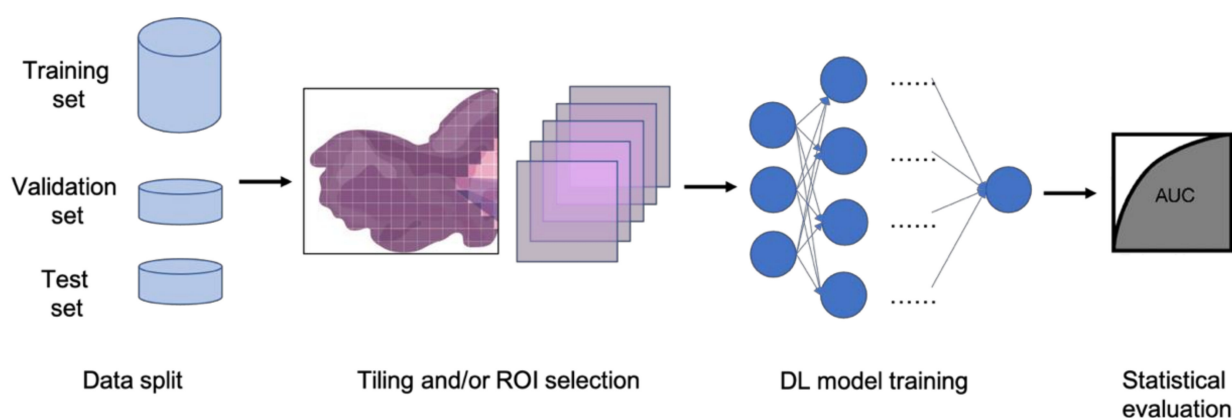
Figure 3. Typical classification model pipeline for histopathology images.

However, histopathology images are quite different from the images in the ImageNet because of their extremely large sizes, higher resolution, and sparser useful feature distributions [11,55,64]. Deep learning architectures that could take advantage of these characteristics are very likely to achieve better results, unveiling more interesting hidden features in histopathology image classification tasks. For instance, a multi-resolution CNN model, which takes advantage of the data structure of .svs and .scn image files, achieves higher performance in classifying endometrial cancer molecular features than its single resolution counterparts [65]. Weakly supervised techniques, such as multiple instance learning, also demonstrate decent performance in classification tasks of histopathology images, and have gained popularity in recent years [64,66,67]. The innovative idea of bringing GNN models into solving histopathology classification problems develops greater capacity in understanding the subtle relationships between features of different tissue structures and at different locations on giant digital histopathology slides [68,69].

3.2. Segmentation

In addition to classification tasks, CNN models are also capable of segmenting cells or tissue in histopathology slides [9,55]. The segmented cells or tissue could then be used to train classification models for different prediction tasks, including the recurrence of non-small-cell lung cancer [70] and endometrial tissue types [71]. A popular segmentation CNN architecture used in the biomedical field is U-net, which has a similar structure to an autoencoder [72]. A 3D version of U-net, which has 3D convolutional layers instead of 2D convolutional layers, is capable of segmenting volumetric images [22]. Modified U-net architectures, such as USE-net [73], Focus U-net [74], and U-net with attention gate [75], have achieved even better performance in various biomedical image segmentation tasks than vanilla U-net. Other autoencoder-based methods have also achieved promising

results in segmentation tasks of histopathology images, such as highlighting tumor regions in liver cancer WSI [76]. Well-trained style transfer models are also viable options for segmentation tasks [48]. With the introduction of GAN, using conditional-GAN or cycle-GAN models and in combination with CNN models for segmentation problems is also shown to be viable, with less stringent training data requirements [46,53,77]. Unlike most classification models, the segmentation models can be more adaptive to different types of tissues due to the similarities of the stained features and textures of the histopathology slides [78]. Additionally, the evaluation metrics of these classification models can be drastically different from those of the classification models. The segmentation labels are also usually images; therefore, it is not easy to determine a binary prediction or even a prediction score at the per-image level. Hence, typical statistical metrics, such as AUROC or precision and recall, are often not capable of fairly evaluating segmentation tasks. Pixel-level metrics, such as intersection over the union (IoU), also pose weaknesses because it cannot objectively give relative importance to pixels of different regions. Object-level metrics can be an optimal alternative, but the requirement of identifying all objects on the label images prohibits its adoption in real-world model evaluation. Therefore, researchers often use customized evaluation metrics with a combination of customized pixel weights, dice loss, and IoU with specific thresholds [79,80].

4. Summary

In this review paper, we have introduced popular deep learning algorithms, CNN, GNN, and GAN, and also highlighted mechanisms of how they work and how they can be applied to solve clinical and scientific questions in pathology. We have also discussed recent publications which show that these deep learning techniques have the potential to be useful in classifying or segmenting histopathology imaging data. With the continuing advancement of machine learning and deep learning techniques and the development of hardware and software, it is realistic to believe that the integration of artificial intelligence and pathology will become an even more attractive field to explore. Compared with conventional computational methods, deep learning techniques generally run faster and have much better performance in pathology tasks. Although one has to be rigorous and ethical about translating these AI-based technologies into clinical settings, we still hold an optimistic view that they will eventually revolutionize medical diagnosis processes and really push the development of precision medicine forward. Above all, the ultimate goal of introducing AI into pathology and biomedicine in general is to make healthcare more accessible, affordable, and agreeable.

Nevertheless, there are still a number of limitations of the current studies and potential obstacles which prevent these models from implementation in contemporary real-world clinical settings. For example, only patterns with prior understanding from pathologists can be used as reliable evidence for prediction, which significantly limits the tasks for which deep learning models can be applied. In addition, the patient samples that can be used as training, validation, and test sets are also very limited for each of the specific tasks of interests. Moreover, detailed labels of medical images are often not available, and the labeling standards among clinicians also vary significantly in different countries. Additionally, the interpretability of deep learning models applied to histopathology images remains debatable, especially among clinicians. More advanced self-supervised or semi-supervised methods may solve some of these problems from a technical perspective in the future.

References

1. LeCun, Y.; Bengio, Y.; Hinton, G. Deep learning. *Nature* **2015**, *521*, 436–444. [CrossRef] [PubMed]
2. Schmidhuber, J. Deep learning in neural networks: An overview. *Neural Netw.* **2015**, *61*, 85–117. [CrossRef] [PubMed]
3. Barbosa, I.B.; Cristani, M.; Caputo, B.; Rognhaugen, A.; Theoharis, T. Looking beyond appearances: Synthetic training data for deep CNNs in re-identification. *Comput. Vis. Image Underst.* **2018**, *167*, 50–62. [CrossRef]
4. Silver, D.; Huang, A.; Maddison, C.J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. Mastering the game of Go with deep neural networks and tree search. *Nature* **2016**, *529*, 484–489. [CrossRef]
5. Artetxe, M.; Labaka, G.; Agirre, E.; Cho, K. Unsupervised Neural Machine Translation. In Proceedings of the 6th International Conference on Learning Representations, ICLR 2018; Conf. Track Proc. 2017, Vancouver, BC, Canada, 30 April–3 May 2018.
6. Militello, C.; Rundo, L.; Vitabile, S.; Conti, V. Fingerprint Classification Based on Deep Learning Approaches: Experimental Findings and Comparisons. *Symmetry* **2021**, *13*, 750. [CrossRef]
7. Senior, A.W.; Evans, R.; Jumper, J.; Kirkpatrick, J.; Sifre, L.; Green, T.; Qin, C.; Žídek, A.; Nelson, A.W.R.; Bridgland, A.; et al. Improved protein structure prediction using potentials from deep learning. *Nature* **2020**, *577*, 706–710. [CrossRef]
8. Jumper, J.; Evans, R.; Pritzel, A.; Green, T.; Figurnov, M.; Ronneberger, O.; Tunyasuvunakool, K.; Bates, R.; Žídek, A.; Potapenko, A.; et al. Highly accurate protein structure prediction with AlphaFold. *Nature* **2021**, *596*, 583–589. [CrossRef]
9. Madabhushi, A.; Lee, G. Image analysis and machine learning in digital pathology: Challenges and opportunities. *Med. Image Anal.* **2016**, *33*, 170–175. [CrossRef]
10. Louis, D.N.; Feldman, M.; Carter, A.B.; Dighe, A.S.; Pfeifer, J.D.; Bry, L.; Almeida, J.S.; Saltz, J.; Braun, J.; Tomaszewski, J.E.; et al. Computational Pathology: A Path Ahead. *Arch. Pathol. Lab. Med.* **2016**, *140*, 41–50. [CrossRef]
11. Komura, D.; Ishikawa, S. Machine Learning Methods for Histopathological Image Analysis. *Comput. Struct. Biotechnol. J.* **2018**, *16*, 34–42. [CrossRef]
12. Nawaz, S.; Yuan, Y. Computational pathology: Exploring the spatial dimension of tumor ecology. *Cancer Lett.* **2016**, *380*, 296–303. [CrossRef] [PubMed]
13. Esteva, A.; Kuprel, B.; Novoa, R.A.; Ko, J.; Swetter, S.M.; Blau, H.M.; Thrun, S. Dermatologist-level classification of skin cancer with deep neural networks. *Nature* **2017**, *542*, 115–118. [CrossRef] [PubMed]
14. Coudray, N.; Ocampo, P.S.; Sakellaropoulos, T.; Narula, N.; Snuderl, M.; Fenyö, D.; Moreira, A.L.; Razavian, N.; Tsirigos, A. Classification and mutation prediction from non–small cell lung cancer histopathology images using deep learning. *Nat. Med.* **2018**, *24*, 1559–1567. [CrossRef] [PubMed]
15. Kather, J.N.; Pearson, A.T.; Halama, N.; Jäger, D.; Krause, J.; Loosen, S.H.; Marx, A.; Boor, P.; Tacke, F.; Neumann, U.P.; et al. Deep learning can predict microsatellite instability directly from histology in gastrointestinal cancer. *Nat. Med.* **2019**, *25*, 1054–1056. [CrossRef]
16. Bengio, Y.; Courville, A.; Vincent, P. Representation Learning: A Review and New Perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.* **2013**, *35*, 1798–1828. [CrossRef]
17. Ioffe, S.; Szegedy, C. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. *arXiv* **2015**, arXiv:1502.03167.
18. Zhang, H.; Cisse, M.; Dauphin, Y.N.; Lopez-Paz, D. mixup: Beyond Empirical Risk Minimization. *arXiv* **2017**, arXiv:1710.09412.
19. Bui, H.M.; Lech, M.; Cheng, E.; Neville, K.; Burnett, I.S. Object Recognition Using Deep Convolutional Features Transformed by a Recursive Network Structure. *IEEE Access* **2016**, *4*, 10059–10066. [CrossRef]
20. Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; Rabinovich, A. Going Deeper with Convolutions. *arXiv* **2015**, arXiv:1409.4842.
21. Fu, C.; Ho, D.J.; Han, S.; Salama, P.; Dunn, K.W.; Delp, E.J. Nuclei Segmentation of Fluorescence Microscopy Images using Convolutional Neural Networks. In Proceedings of the 2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017), Melbourne, Australia, 18–21 April 2017; IEEE; 2017; pp. 704–708.
22. Çiçek, Ö.; Abdulkadir, A.; Lienkamp, S.S.; Brox, T.; Ronneberger, O. *3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation*; Springer: Cham, Switzerland, 2016; pp. 424–432.
23. Machado, S.; Mercier, V.; Chiaruttini, N. LimeSeg: A coarse-grained lipid membrane simulation for 3D image segmentation. *BMC Bioinform.* **2019**, *20*, 2. [CrossRef]
24. Lecun, Y.; Bottou, L.; Bengio, Y.; Haffner, P. Gradient-based learning applied to document recognition. *Proc. IEEE* **1998**, *86*, 2278–2324. [CrossRef]
25. Krizhevsky, A.; Sutskever, I.; Hinton, G.E. ImageNet Classification with Deep Convolutional Neural Networks. *Adv. Neural Inf. Process. Syst.* **2012**, *25*. [CrossRef]

26. Simonyan, K.; Zisserman, A. Very Deep Convolutional Networks for Large-Scale Image Recognition. In Proceedings of the 3rd International Conference on Learning Representations, ICLR 2015—Conference Track Proceedings, San Diego, CA, USA, 7–9 May 2015.

27. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 27–30 June 2016; pp. 770–778. [CrossRef]

28. Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; Wojna, Z. Rethinking the Inception Architecture for Computer Vision. *arXiv* **2016**, arXiv:1512.00567.

29. Szegedy, C.; Ioffe, S.; Vanhoucke, V.; Alemi, A.A. Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning. In Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, San Francisco, CA, USA, 4–10 February 2017.

30. Springenberg, J.T.; Dosovitskiy, A.; Brox, T.; Riedmiller, M. Striving for Simplicity: The All Convolutional Net. *arXiv* **2014**, arXiv:1412.6806.

31. Mahendran, A.; Vedaldi, A. Understanding Deep Image Representations by Inverting Them. *arXiv* **2015**, arXiv:1412.0035.

32. Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; Torralba, A. Learning Deep Features for Discriminative Localization. *arXiv* **2016**, arXiv:1512.04150.

33. Simonyan, K.; Vedaldi, A.; Zisserman, A. Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps. *arXiv* **2013**, arXiv:1312.6034.

34. Selvaraju, R.R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; Batra, D. Grad-CAM: Visual Explanations From Deep Networks via Gradient-Based Localization. *arXiv* **2017**, arXiv:1610.02391.

35. Adebayo, J.; Gilmer, J.; Muelly, M.; Goodfellow, I.; Hardt, M.; Kim, B.; Brain, G. Sanity Checks for Saliency Maps. *arXiv* **2018**, arXiv:1810.03292.

36. Nekhaev, D.; Demin, V. Visualization of maximizing images with deconvolutional optimization method for neurons in deep neural networks. *Procedia Comput. Sci.* **2017**, *119*, 174–181. [CrossRef]

37. Zeiler, M.D.; Fergus, R. *Visualizing and Understanding Convolutional Networks*; Springer: Cham, Switzerland, 2014; pp. 818–833.

38. Nguyen, A.; Yosinski, J.; Clune, J. Multifaceted Feature Visualization: Uncovering the Different Types of Features Learned by Each Neuron in Deep Neural Networks. *arXiv* **2016**, arXiv:1602.03616.

39. Courville, A.; Erhan, D.; Bengio, Y.; Vincent, P. *Visualizing Higher-Layer Features of a Deep Network Visualizing Higher-Layer Features of a Deep Network*; Département d'Informatique et Recherche Opérationnelle: Montreal, QC, Canada, 2009.

40. Yosinski, J.; Clune, J.; Nguyen, A.; Fuchs, T.; Lipson, H. Understanding Neural Networks Through Deep Visualization. *arXiv* **2015**, arXiv:1506.06579.

41. McInnes, L.; Healy, J.; Melville, J. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *arXiv* **2018**, arXiv:1802.03426.

42. Carter, S.; Armstrong, Z.; Schubert, L.; Johnson, I.; Olah, C. Activation Atlas. *Distill* **2019**, *4*, e15. [CrossRef]

43. Van der Maaten, L.; Hinton, G. Visualizing Data using t-SNE. *J. Mach. Learn. Res.* **2008**, *9*, 2579–2605.

44. Zhou, J.; Cui, G.; Zhang, Z.; Yang, C.; Liu, Z.; Wang, L.; Li, C.; Sun, M. Graph Neural Networks: A Review of Methods and Applications. *arXiv* **2018**, arXiv:1812.08434. [CrossRef]

45. Wu, Z.; Pan, S.; Chen, F.; Long, G.; Zhang, C.; Yu, P.S. A Comprehensive Survey on Graph Neural Networks. *arXiv* **2019**, arXiv:1901.00596. [CrossRef]

46. Yi, X.; Walia, E.; Babyn, P. Generative Adversarial Network in Medical Imaging: A Review. *arXiv* **2018**, arXiv:1809.07294. [CrossRef]

47. Bau, D.; Zhu, J.-Y.; Strobelt, H.; Zhou, B.; Tenenbaum, J.B.; Freeman, W.T.; Torralba, A. GAN Dissection: Visualizing and Understanding Generative Adversarial Networks. *arXiv* **2018**, arXiv:1811.10597.

48. Hollandi, R.; Szkalisity, A.; Toth, T.; Tasnadi, E.; Molnar, C.; Mathe, B.; Grexa, I.; Molnar, J.; Balind, A.; Gorbe, M.; et al. A deep learning framework for nucleus segmentation using image style transfer. *bioRxiv* **2019**, 580605. [CrossRef]

49. Gatys, L.A.; Ecker, A.S.; Bethge, M. A Neural Algorithm of Artistic Style. *arXiv* **2015**, arXiv:1508.06576. [CrossRef]

50. Zhu, J.-Y.; Park, T.; Isola, P.; Efros, A.A. Unpaired Image-To-Image Translation Using Cycle-Consistent Adversarial Networks. *arXiv* **2017**, arXiv:1703.10593.

51. Isola, P.; Zhu, J.-Y.; Zhou, T.; Efros, A.A. Image-To-Image Translation With Conditional Adversarial Networks. *arXiv* **2017**, arXiv:1611.07004.

52. Radford, A.; Metz, L.; Chintala, S. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *arXiv* **2015**, arXiv:1511.06434.

53. Fu, C.; Lee, S.; Joon Ho, D.; Han, S.; Salama, P.; Dunn, K.W.; Delp, E.J. Three Dimensional Fluorescence Microscopy Image Synthesis and Segmentation. *arXiv* **2018**, arXiv:1801.07198.

54. Puri, M.; Hoover, S.B.; Hewitt, S.M.; Wei, B.-R.; Adissu, H.A.; Halsey, C.H.C.; Beck, J.; Bradley, C.; Cramer, S.D.; Durham, A.C.; et al. Automated Computational Detection, Quantitation, and Mapping of Mitosis in Whole-Slide Images for Clinically Actionable Surgical Pathology Decision Support. *J. Pathol. Inform.* **2019**, *10*, 4. [CrossRef]

55. Cooper, L.A.; Demicco, E.G.; Saltz, J.H.; Powell, R.T.; Rao, A.; Lazar, A.J. PanCancer insights from The Cancer Genome Atlas: The pathologist's perspective. *J. Pathol.* **2018**, *244*, 512–524. [CrossRef] [PubMed]

56. Mobadersany, P.; Yousefi, S.; Amgad, M.; Gutman, D.A.; Barnholtz-Sloan, J.S.; Velázquez Vega, J.E.; Brat, D.J.; Cooper, L.A.D. Predicting cancer outcomes from histology and genomics using convolutional networks. *Proc. Natl. Acad. Sci. USA* **2018**, *115*, E2970–E2979. [CrossRef]

57. Gillette, M.A.; Satpathy, S.; Cao, S.; Dhanasekaran, S.M.; Vasaikar, S.V.; Krug, K.; Petralia, F.; Li, Y.; Liang, W.W.; Reva, B.; et al. Proteogenomic Characterization Reveals Therapeutic Vulnerabilities in Lung Adenocarcinoma. *Cell* **2020**, *182*, 200–225.e35. [CrossRef]

58. Hong, R.; Liu, W.; Fenyö, D. Predicting and Visualizing STK11 Mutation in Lung Adenocarcinoma Histopathology Slides Using Deep Learning. *BioMedInformatics* **2021**, *2*, 6. [CrossRef]

59. Wang, L.B.; Karpova, A.; Gritsenko, M.A.; Kyle, J.E.; Cao, S.; Li, Y.; Rykunov, D.; Colaprico, A.; Rothstein, J.H.; Hong, R.; et al. Proteogenomic and metabolomic characterization of human glioblastoma. *Cancer Cell* **2021**, *39*, 509–528.e20. [CrossRef]

60. Kim, R.H.; Nomikou, S.; Coudray, N.; Jour, G.; Dawood, Z.; Hong, R.; Esteva, E.; Sakellaropoulos, T.; Donnelly, D.; Moran, U.; et al. Deep learning and pathomics analyses reveal cell nuclei as important features for mutation prediction of BRAF-mutated melanomas. *J. Investig. Dermatol.* **2021**, *in press*. [CrossRef] [PubMed]

61. Schmauch, B.; Romagnoni, A.; Pronier, E.; Saillard, C.; Maillé, P.; Calderaro, J.; Kamoun, A.; Sefta, M.; Toldo, S.; Zaslavskiy, M.; et al. A deep learning model to predict RNA-Seq expression of tumours from whole slide images. *Nat. Commun.* **2020**, *11*, 3877. [CrossRef] [PubMed]

62. Kather, J.N.; Heij, L.R.; Grabsch, H.I.; Loeffler, C.; Echle, A.; Muti, H.S.; Krause, J.; Niehues, J.M.; Sommer, K.A.J.; Bankhead, P.; et al. Pan-cancer image-based detection of clinically actionable genetic alterations. *Nat. Cancer* **2020**, *1*, 789–799. [CrossRef]

63. Fu, Y.; Jung, A.W.; Torne, R.V.; Gonzalez, S.; Vöhringer, H.; Shmatko, A.; Yates, L.R.; Jimenez-Linan, M.; Moore, L.; Gerstung, M. Pan-cancer computational histopathology reveals mutations, tumor composition and prognosis. *Nat. Cancer* **2020**, *1*, 800–810. [CrossRef]

64. Campanella, G.; Hanna, M.G.; Geneslaw, L.; Miraflor, A.; Werneck Krauss Silva, V.; Busam, K.J.; Brogi, E.; Reuter, V.E.; Klimstra, D.S.; Fuchs, T.J. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nat. Med.* **2019**, *25*, 1301–1309. [CrossRef]

65. Hong, R.; Liu, W.; DeLair, D.; Razavian, N.; Fenyö, D. Predicting endometrial cancer subtypes and molecular features from histopathology images using multi-resolution deep learning models. *Cell Rep. Med.* **2021**, *2*, 100400. [CrossRef]

66. Lu, M.Y.; Chen, T.Y.; Williamson, D.F.K.; Zhao, M.; Shady, M.; Lipkova, J.; Mahmood, F. AI-based pathology predicts origins for cancers of unknown primary. *Nature* **2021**, *594*, 106–110. [CrossRef]

67. Cinbis, R.G.; Verbeek, J.; Schmid, C. Weakly Supervised Object Localization with Multi-Fold Multiple Instance Learning. *IEEE Trans. Pattern Anal. Mach. Intell.* **2017**, *39*, 189–203. [CrossRef]

68. Anand, D.; Gadiya, S.; Sethi, A. Histographs: Graphs in histopathology. *arXiv* **2020**, arXiv:1908.05020.

69. Gao, Z.; Shi, J.; Wang, J. GQ-GCN: Group Quadratic Graph Convolutional Network for Classification of Histopathological Images. In Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention, Strasbourg, France, 27 September–1 October 2021; LNCS 2021. Volume 12908.

70. Wang, X.; Velcheti, V.; Vaidya, P.; Bera, K.; Madabhushi, A.; Khunger, A.; Patil, P.; Choi, H. RaPtomics: Integrating Radiomic and Pathomic Features for Predicting Recurrence in Early Stage Lung Cancer. In *Medical Imaging 2018: Digital Pathology, Proceedings of the*; Gurcan, M.N., Tomaszewski, J.E., Eds.; SPIE: Bellingham, WA, USA, 2018; Volume 10581, p. 21.

71. Sun, H.; Zeng, X.; Xu, T.; Peng, G.; Ma, Y. Computer-Aided Diagnosis in Histopathological Images of the Endometrium Using a Convolutional Neural Network and Attention Mechanisms. *IEEE J. Biomed. Heal. Inform.* **2020**, *24*, 1664–1676. [CrossRef]

72. Ronneberger, O.; Fischer, P.; Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. *Lect. Notes Comput. Sci. (Incl. Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinform.)* **2015**, *9351*, 234–241.

73. Rundo, L.; Han, C.; Nagano, Y.; Zhang, J.; Hataya, R.; Militello, C.; Tangherloni, A.; Nobile, M.S.; Ferretti, C.; Besozzi, D.; et al. USE-Net: Incorporating Squeeze-and-Excitation blocks into U-Net for prostate zonal segmentation of multi-institutional MRI datasets. *Neurocomputing* **2019**, *365*, 31–43. [CrossRef]

74. Yeung, M.; Sala, E.; Schönlieb, C.B.; Rundo, L. Focus U-Net: A novel dual attention-gated CNN for polyp segmentation during colonoscopy. *Comput. Biol. Med.* **2021**, *137*, 104815. [CrossRef] [PubMed]

75. Schlemper, J.; Oktay, O.; Schaap, M.; Heinrich, M.; Kainz, B.; Glocker, B.; Rueckert, D. Attention gated networks: Learning to leverage salient regions in medical images. *Med. Image Anal.* **2019**, *53*, 197–207. [CrossRef] [PubMed]

76. Roy, M.; Kong, J.; Kashyap, S.; Pastore, V.P.; Wang, F.; Wong, K.C.L.; Mukherjee, V. Convolutional autoencoder based model HistoCAE for segmentation of viable tumor regions in liver whole-slide images. *Sci. Rep.* **2021**, *11*, 139. [CrossRef]

77. Ho, D.J.; Fu, C.; Salama, P.; Dunn, K.W.; Delp, E.J. Nuclei Segmentation of Fluorescence Microscopy Images Using Three Dimensional Convolutional Neural Networks. In Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), Honolulu, HI, USA, 21–26 July 2017.

78. Greenwald, N.F.; Miller, G.; Moen, E.; Kong, A.; Kagel, A.; Camacho Fullaway, C.; Mcintosh, B.J.; Leow, K.; Schwartz, M.S.; Dougherty, T.; et al. Whole-cell segmentation of tissue images with human-level performance using large-scale data annotation and deep learning. *bioRxiv* **2021**. [CrossRef]

79. Caicedo, J.C.; Roth, J.; Goodman, A.; Becker, T.; Karhohs, K.W.; Broisin, M.; Csaba, M.; McQuin, C.; Singh, S.; Theis, F.; et al. Evaluation of Deep Learning Strategies for Nucleus Segmentation in Fluorescence Images. *bioRxiv* **2019**, 335216. [CrossRef]
80. Sudre, C.H.; Li, W.; Vercauteren, T.; Ourselin, S.; Jorge Cardoso, M. *Generalised Dice Overlap as a Deep Learning Loss Function for Highly Unbalanced Segmentations*; Springer: Cham, Switzerland, 2017; pp. 240–248.