Unsupervised Machine Learning in Pathology
The Next Frontier

Adil Roohi, BSc[a,b,1], Kevin Faust[b,c,1], Ugljesa Djuric, PhD[d], Phedias Diamandis, MD, PhD[d,e,f,*]

KEYWORDS

• Pathology • Artificial intelligence • Deep learning • Unsupervised learning • Machine learning • Neuropathology

Key points

- Artificial intelligence and deep learning are increasingly prevalent in pathology and used to process large amounts of data.
- Unsupervised learning allows computational networks to discover patterns in data without significant training.
- This emerging type of machine learning facilitates a more human-like analytical approach, allowing for nuanced conclusions to be made without the need for specific pre-defined direction.

ABSTRACT

Applications of artificial intelligence and particularly deep learning to aid pathologists in carrying out laborious and qualitative tasks in histopathologic image analysis have now become ubiquitous. We introduce and illustrate how unsupervised machine learning workflows can be deployed in existing pathology workflows to begin learning autonomously through exploration and without the need for extensive direction. Although still in its infancy, this type of machine learning, which more closely mirrors human intelligence, stands to add another exciting layer of innovation to computational pathology and accelerate the transition to autonomous pathologic tissue analysis.

Non-standard abbreviations

CNN	Convolutional neural network
TGCA	The Cancer Genome Atlas
t-SNE	t-Distributed Stochastic Neighbor Embedding
WHO	World Health Organization

Most of human and animal learning is unsupervised. If intelligence was a cake, unsupervised learning would be the cake and supervised learning would be the icing on the cake.
—Yann LeCun, Deep Learning Pioneer

[a] Harvard Extension School, 51 Brattle Street, Cambridge, MA 02138, USA; [b] Princess Margaret Cancer Centre, 101 College Street, Toronto, Ontario M5G 1L7, Canada; [c] Department of Computer Science, University of Toronto, 40 St. George Street, Toronto, Ontario M5S 2E4, Canada; [d] Laboratory Medicine Program, University Health Network, 200 Elizabeth Street, Toronto, Ontario M5G 2C4, Canada; [e] Department of Laboratory Medicine and Pathobiology, University of Toronto, Toronto, Ontario M5S 1A8, Canada; [f] Department of Medical Biophysics, University of Toronto, Toronto, Ontario, Canada
[1]Equal contribution.
* Corresponding author. Laboratory Medicine Program, University Health Network, 200 Elizabeth Street, Toronto, Ontario M5G 2C4, Canada.
E-mail address: p.diamandis@mail.utoronto.ca

Surgical Pathology 13 (2020) 349–358
https://doi.org/10.1016/j.path.2020.01.002
1875-9181/20/© 2020 Elsevier Inc. All rights reserved.

surgpath.theclinics.com

THE NEED FOR ARTIFICIAL AGENTS TO ANALYZE MEDICAL IMAGING DATA

There is a growing appreciation for the need to accelerate the transition of microscopic tissue analysis from one centered around glass slides, toward a more digital, automated, and quantitative discipline.[1,2] Even though this evolution has been met with both fear and excitement by pathologists, the adoption of machine learning is a trend that stretches far beyond pathology and is rather becoming a common theme across all of health care and society.[3,4] This revolution is perhaps partly driven by transformative changes in the ability to generate, digitally store, and process data. The cost of digital storage, for example, has decreased by almost 6 orders of magnitude over the past 4 decades, making it less of a barrier for data collection. This has been especially true for domains like pathology, where file sizes have been comparatively large with single glass slides often requiring a gigabyte of data for digital storage.

Regardless of the underlying reasons for this exponential decrease in the cost of data storage, these types of innovations are already changing how we use medical infrastructure and resources. For example, these technological improvements have led to the increasing reliance on advanced medical imaging (eg, computed tomography, MRIs) over the past 40 years and will likely continue to do so for the foreseeable future. As the medical community's ability to generate and store high-quality data continues to grow, the demand for highly specialized humans (pathologists and radiologists) needed to interpret the data will likely not be able to keep up. Specifically, this is becoming a reality in pathology, as the discipline's ability to digitally scan and store whole slide images at qualities suitable for diagnostic practice is also rapidly growing.[5] Moreover, the concurrent advent of artificial intelligence and deep learning could transform how histologic data are analyzed and used from each patient's specimen.[6,7]

THE PERILS AND PROMISES OF DEEP LEARNING FOR PATHOLOGIC IMAGE ANALYSIS

Although pattern recognition can be said to be an innate skill for the human brain, it is common practice to convert this diagnostic "art" into a set of teachable, highly reproducible, and objective set of rules to improve consistency among observers. This strategy has been particularly important for high-stakes decision in areas like medicine. For example, in pathology, trainees are encouraged to arrive at diagnoses by making a set of progressive decisions geared at carefully narrowing the differential diagnosis to a single or small set of likely diseases (eg, lesion vs no lesion; infectious vs neoplastic process). Development of such sets of reliable rules and robust decision tree-type frameworks not only improves reproducibility among humans but has also made this process highly conducive to automation by machines.

For example, in traditional machine learning approaches, computer engineers leverage these rules and frameworks to develop surrogate handcrafted computer features that attempt to mimic or parallel the morphologic features (eg, necrosis, mitoses) that pathologists use to arrive to a diagnosis. Once digitized, the quantitative value of these multiple features can serve as objective morphologic signatures for image classification. Although these have been quite effective,[7] one limitation of this traditional approach is that the relatively small number of manually engineered patterns (100–1000s), are often unable to capture complex positional information that humans innately use to carry out specialized pattern recognition tasks. Recently however, transformative innovations in computer vision, particularly a form of artificial intelligence (AI) known as deep learning, has helped overcome this.[8] Particularly, a specific type of deep learning algorithm, known as convolutional neural networks (CNN), attempts to mimic how the human brain processes visual information and has allowed scientists to transition away from these laborious manual hand crafted features and instead rely on data to drive feature design.[9] Like the brain's visual cortex, the multilayered CNN architecture first detects elementary features (eg, color, shapes, edges) within an image and sequentially aggregates different combinations of the patterns to generate millions of advanced spatially dependant features of higher classification value.[8] Importantly, these features are computationally designed and selected and not reliant on humans for their generation. Once these features are developed and trained on a particular task, test images can be introduced to these networks to quantify the presence of these complex features and then use them to carry out classification tasks. With sufficient data and processing power that has recently become available, these novel computational tools were theoretically predicted to surpass the performance of traditional tools for classification.

This hypothesis was ultimately proven to be correct during the 2012 ImageNet computer vision competition.[10,11] Here, computer scientists annually compete to determine the most effective

approaches and algorithms for classifying images that span 1000 different classes of common objects (eg, cats, dogs, planes). When this competition first started, the ImageNet winners used algorithms that were designed around the traditional hand-crafted feature approaches. These approaches had classification error rates of approximately 27%; much higher than a human given this same task (~3% error) and perhaps too high for any practical application. In 2012 however, AI pioneer Geoffrey Hinton and his team introduced, the first CNN-based algorithm in this image recognition competition and substantially reduced the winning error rate to 16%.[10,11] Every year since, sequential modifications have improved on his innovation, with state-of-the-art CNNs now equaling, and even surpassing, humans at classifying the diverse image types found in this competition. Although an impressive feat, many critics have rightfully pointed out that this competition represents a highly controlled environment and task that does not fully capture the dynamic decision-making capabilities of human observers in the real-world environment.

In recent years, these breakthroughs in deep learning approaches to pattern recognition have infiltrated the medical field.[12–16] Numerous studies have now shown that deep learning can perform many visual diagnostic tasks at levels that match or even exceed human experts in both primary care settings and more definitive pathologic analyses.[12,15] These systems are so robust that the tools have largely now become democratized, allowing anyone with access to data to develop and use them without the need for significant training in the computer sciences. This has been particularly effective in pathology, where repetitive patterns can be used to generate massive datasets that the algorithms can use to learn and classify future images (**Fig. 1**). However, despite these heralded successes, many surveys across all fields (eg, medicine, finance) find that most implementations of AI initiatives are met with significant challenges and failure. Arguably, a major component of these failures is the current inability for supervised learning approaches to fully mimic the diverse complement of skills and inference capabilities of the human observer. In this perspective, we wish to introduce the concept of unsupervised learning and how it differs from the supervised tasks that have now become ubiquitous in histopathologic image analysis. We supplement some of these theoretic concepts with examples from our own work to illustrate how unsupervised deep learning approaches can be introduced into pathology workflows to overcome some of the existing barriers we believe are preventing the deployment of AI in pathology.

UNSUPERVISED LEARNING

Unsupervised learning fundamentally differs from supervised learning because it does not rely on specific classification instructions. Instead, it relies on autonomously grouping objects through exploration and discovery of the underlying pattern and structures in data. As alluded to in the quote at the beginning of this piece, this approach is viewed as the primary way we as humans collect information and build knowledge about the world around us. Although this undirected approach may not afford unsupervised learning the ability to resolve the same level of detail for specific user-defined tasks, it provides a highly flexible approach to define patterns in data that do not need to be predicted or anticipated. Unsupervised learning also affords humans with the ability to handle unanticipated changes in conditions and extend knowledge outside the training parameters (**Fig. 2**). The regular implementation of such strategies into machine learning workflows could help broaden AI-workflow performances in a wider spectrum of scenarios without the constant and infeasible task of continual needing to update supervised training parameters as knowledge and diagnostic goals evolve.

To illustrate the differences between supervised and unsupervised learning, we consider both continuous and discrete outputs generated by both approaches (see **Fig. 2**). A classic supervised classification task that uses continuous data would involve investigating if the aggressiveness of meningiomas (or another tumor type) can be predicted as a continuous function of the Ki-67 index (regression analysis). Conversely, the World Health Organization (WHO) current grading system represents a more discrete supervised classification task in which the combined presence of specific features (eg, necrosis, nucleoli), at varying amounts, could be sufficient to warrant a specific WHO grades (I vs II). Although these techniques are highly effective, there could, of course, be other ways to better organize the patterns, combinations, and respective amounts of features, like mitoses and necrosis. Importantly, although some of these patterns may not correlate with a desired task (aggressiveness), reoccurring patterns may provide important insight and have other important implications to the underlying biology of these tumors and help guide more personalized therapies. These are the types of “serendipitous” or “anomalous” patterns that are

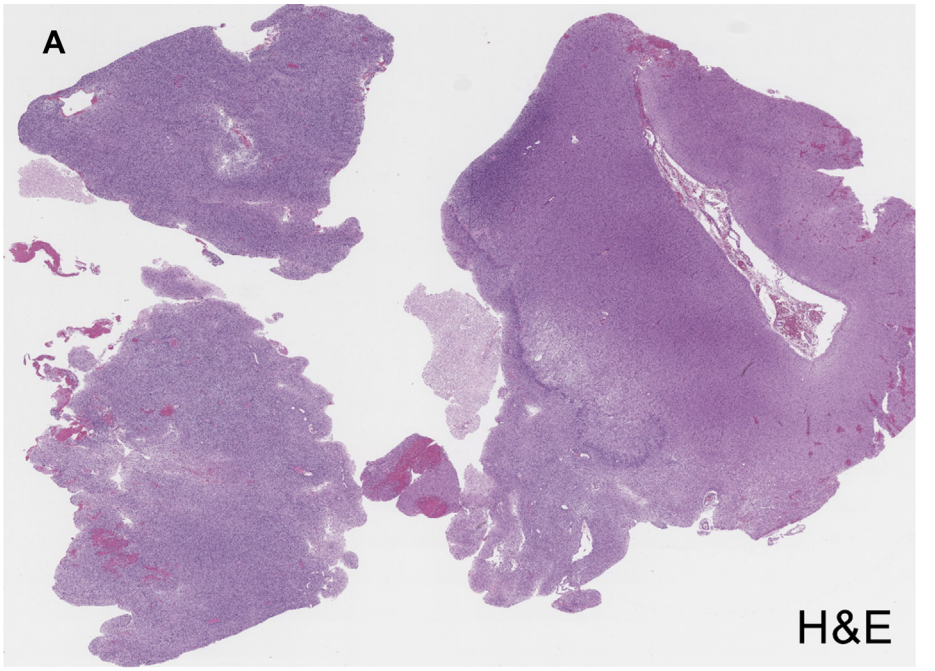

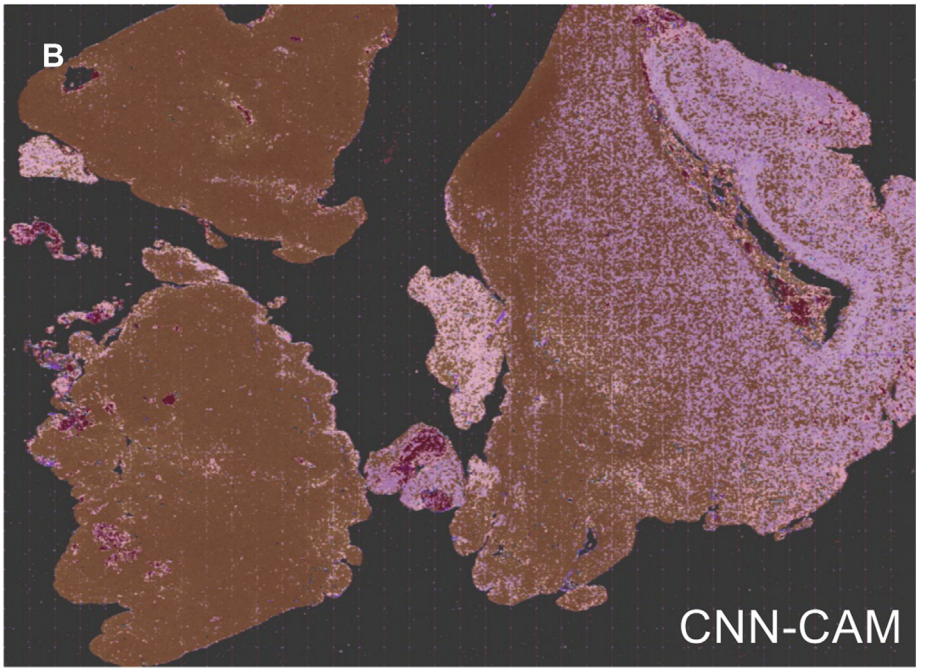

Blank Space: 29.92% (Black)
Grey Matter: 9.26% (Pink)
Meningioma: 0.51% (Brown)
Lymphoma: 1.29% (Brown)
Blood: 0.84% (Red)
Schwannoma: 0.32% (Brown)
Cerebellum: 0.10% (Pink)
White Matter: 0.61% (Pink)
Glioma: 52.83% (Brown)
Metastasis: 0.14% (Brown)
Necrosis: 3.92% (Light Pink)
Surgical: 0.09% (Magenta)
Dura: 0.18% (Dark Orchid)

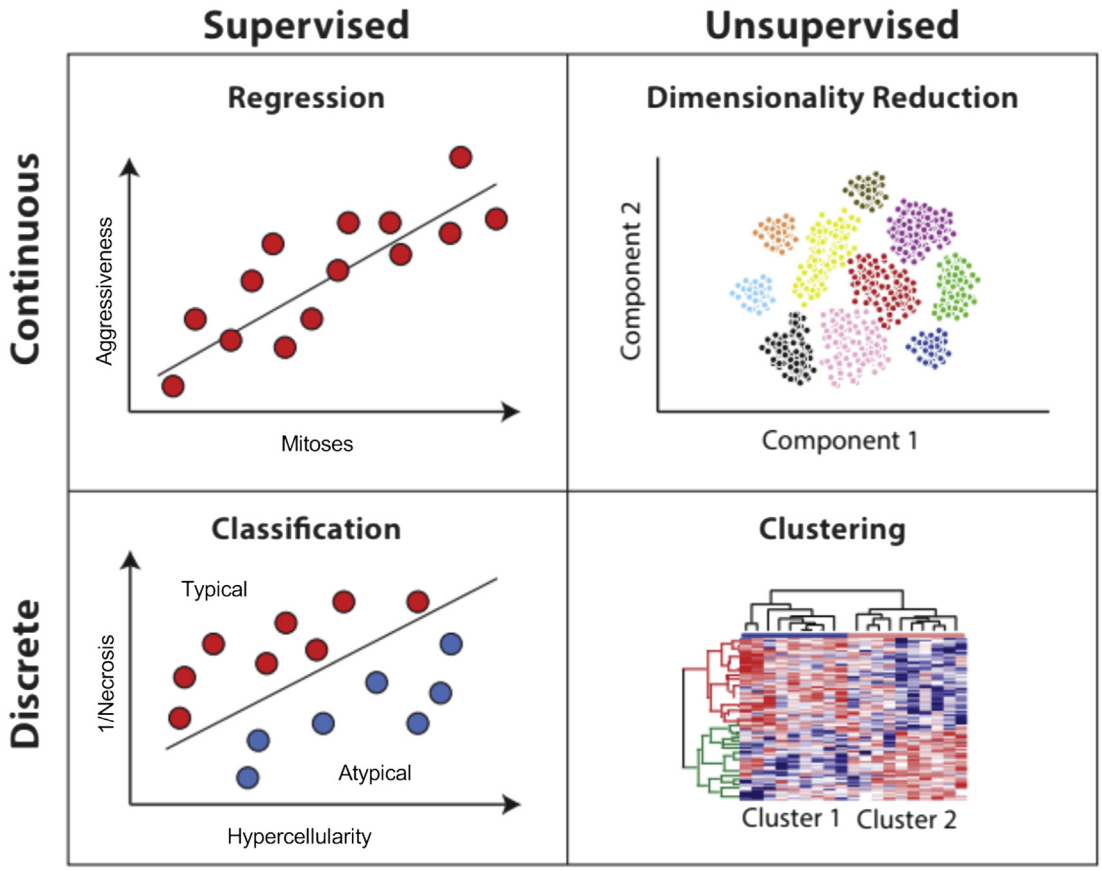

Fig. 2. Differences between machine learning methods. Illustrated are the distinct outputs of the 2 main groups of machine learning methods. In supervised learning, labeled samples with both input and output data are used to develop a model that can approximate the relationship between those values. The developed model can then be used to make future prediction of a specific output given a set of input values. In regression analysis, the output value is continuous, whereas for classic classification tasks, the output categorizes the test case into one specific class. In unsupervised learning, the algorithm uses only input data to propose the natural structure globally present within the data points, without the use of a specific output for guidance. Similar to supervised learning outputs, the outputs can be both continuous or grouped into more discrete clusters.

often discovered by expert pathologists who have seen high volumes of cases in their careers. Conversely, these are also the main types of patterns central to unsupervised analyses.

Dimensionality reduction techniques, such as principal component analysis, and t-SNE (t-distributed Stochastic Neighbor Embedding) provide data visualization outputs that allow exploration into local and global patterns in continuous data. This information also can be discretized using hierarchical clustering, which facilitates the grouping of data into different clusters with an ontological tree highlighting the distance similarity of each data point to one another. Although these spatial and cluster associations are "unlabeled," they provide an initial starting point for the human

Fig. 1. Supervised annotation of a digital whole slide image using CNNs. (*A*) Hematoxylin-eosin (H&E)-stained digital slide image of an infiltrating glioblastoma showing the typical heterogeneous mixture of tumor, necrosis, brain tissue, and blood. (*B*) Corresponding CNN-generated class activation map (CAM) of the entire slide. The probability scores for each of the 13 trained classes of this CNN listed below the image provide a global overview of the tumor and tissue patterns detected. The most likely diagnosis of the lesion represented in this case based on these values is a glioma. Lesional areas are collectively depicted with a brown color on the CAM, whereas other tissue components are a different color. This depicts the powerful supervised capabilities of supervised classification tasks by CNN in pathology.

observer to find clinical or other relevant biological correlates that could explain why these associations arose. In this article, we provide some examples from our work to illustrate how we have used unsupervised approaches to help automate "anomaly" detection during routine classification tasks and also for automated ontological organization of tumor types without explicit instruction. Further development of such tools could serve to complement more supervised approaches by providing access to more advanced decision-making capabilities needed to better automate the diagnostic pathology workflow.

DIMENSIONALITY REDUCTION FOR HIGHLIGHTING ANOMALOUS CASES AND REDUCING ERRORS

Most current studies that use deep learning for computer vision in pathology involve highly focused and controlled classification tasks. Unfortunately, in the real world, there often can be extreme biological variability from case to case, even for common lesions like glioblastoma (eg, gliosarcomatous, small cell and epithelioid variants).[17] Similarly, tumor classification schemes are still evolving, making it cumbersome to continually tune and authoritatively validate complex machine learning classifiers.[18] As a result, unexpected classification errors can occur in these situations (eg, mistaking a small cell glioblastoma for a diffuse large B-cell lymphoma) or even when more elementary artifacts are encountered that were not comprehensively included in the training data. This includes differences in cellularity from edema, changes in the intensity of staining, and folds in tissue. In our hands, this phenomenon is perhaps one of the largest barriers to widespread adoption of these tools in pathology. Although this obstacle can theoretically be overcome with increasing amounts of data, the unpredictability of how tumor phenotypes and groupings will continue to change in the era of personalized medicine, where patients may receive a myriad of individualized therapies, can make such efforts extremely difficult to standardize (or "supervise").

To highlight how unsupervised learning approaches can help overcome these limitations, we recently explored how anomalies and artifacts can be effectively detected and flagged using the dimensionality reduction technique known as t-SNE.[19] Specifically, we generated approximately 80,000 training images spanning approximately 100 surgical cases to train and create a multiclass classifier that could annotate 13 common tissue classes encountered in neuropathology (eg, white matter, dura, lymphoma, gliomas, blood, meningioma, metastasis) (see **Fig. 1**). When deploying our classifier in a traditional supervised manner to classify a set of 123 testing cases, it was able to correctly diagnose 86% of the test cases (14% errors). Interestingly, most errors arose from "untrained" and relatively "rare" tumor subtypes (eg, gliosarcoma, hemangioblastoma) that were not included in the original training set.

To address this issue, without having to develop additional training examples, we instead used unsupervised techniques to determine if outlier test cases could be efficiently detected as anomalous and prevent them from being classified incorrectly.[19] Toward this, we used the same CNN used to classify images shown in **Fig. 1**, to generate a t-SNE plot and visualize the learned representations within the network (**Fig. 3**). On this 2-dimensional (2D) grid, the proximity of 2 images (represented as individual dots) or groups of images (clusters of similarly colored dots) indicates their degree of similarity. Although the close proximity of the similarly colored dots represents the supervised component of learning, the distances between different classes is largely driven by unsupervised tissue patterns independently learned by the computer.

By closely examining the organization, many "intelligent" associations are evident. For example, the network grouped cellular tissue elements close together on the right half of the 2D plot. Similarly, the glioma image cluster (cluster of blue dots) is positionally closest to normal glial tissue (yellow/green/dark green dots), suggesting that the CNN could have learned some implicit patterns of similarities in these glial tissue elements. Moreover, more cohesive tumors (meningiomas, metastases) also appear close to one another on the plot (orange and purple dots). This organization of tissue classes demonstrates intelligible unsupervised learning that has been patterned within CNNs. From a practical perspective, this rich positional information on a plot can be used to automate the detection of anomalous cases into histopathologic image analysis. When a test image (a slide the computer has not seen before) is regionally sampled and overlaid onto these plots (depicted as red diamonds in **Fig. 3**), the overlap between test images and those found in the training set suggests high similarities in the tissue architecture (see **Fig. 3**A). However, when the CNN is presented with a previously unencountered class (eg, hemangioblastoma) or variant (eg, gliosarcoma), the vast differences in the histologic patterns (and imaging data structures) between the rare/untrained testing cases and stereotypical training datasets become

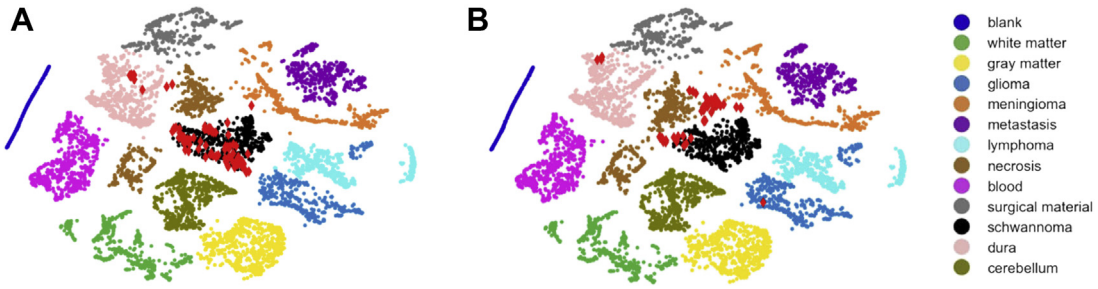

Fig. 3. Dimensionality reduction for anomaly detection. Planar representation of the internal high-dimensional data structure of the final layer of the 13-class CNN used to generate the CAM in Fig. 1. Each colored dot represents an ~512 μm^2 histologic image. During supervised training, images belonging to the same class allow the CNN to optimize features for robust classification. The t-SNE plot of these data allow us to understand the unsupervised association of the global data structure. In (*A*), images patched from a test image (red diamonds, glioblastoma) are overlaid onto the t-SNE plot. The overlap of the test and training images suggests that the data structure of the test image is similar to the training image. In (*B*), image patches from the test image of an untrained class (hemangioblastoma) show a distinct data structure compared with the training images. This suggests an anomalous or unlearned tissue class. (*Adapted from* Faust K, Xie Q, Han D, et al. Visualizing histopathologic deep learning classification and anomaly detection using nonlinear feature space dimensionality reduction. *BMC Bioinformatics*. 2018;19(1):173; with permission.)

immediately evident. These differences are depicted by the images grouping in a unique and nonoverlapping area of the t-SNE plot (see Fig. 3B). We leverage this property to objectively define untrained/anomalies in data and prevent classification errors in these challenging circumstances. In fact, even without any modifications to the network or use of additional data, we were able to reduce errors by more than 60% using this alternative approach.[19] As a result, unsupervised visualization of data structures provides immediately available approaches to monitor unexpected pathologies. This approach can potentially not only help reduce errors but also excitingly potentially detect common tumor with anomalous biology (ie, significant infiltration by lymphocytes) with important implications for emerging treatments, like immunotherapy.

CLUSTERING TO AUTOMATE ONTOLOGICAL ARRANGEMENT OF MICROSCOPIC PATTERNS ON A LARGE SCALE

Although the qualitative detection of novel/anomalous cases is an important component of the microscopic examination, the ability to link recurring patterns into potentially relevant subgroups also represents an advanced and important skill many pathologists contribute to the ontological organization of tumors through their diagnostic practice. Although the baseline tumor subtypes have already been largely defined, this skill will likely continue to help in characterizing new diseases, outbreaks, and differential responses to emerging treatment regimens. This type of specialized pattern grouping task also can be mirrored using unsupervised clustering algorithms that ontologically arrange cases into discrete groups based on shared features.

To highlight how this tasked could be mirrored in silico, we recently also developed a 74-class tissue classifier mostly composed of tumors outlined in the WHO classification guide for brain tumors. This diagnostic manual, developed from input from many international experts, provides a consensus approach to how tumors should be organized based on their microscopic features, molecular characteristics, and understood biology (gliomas, meningiomas, with different subtypes within each class).[20] Notably, some microscopic features, like high nuclear-to-cytoplasmic ratio, can be shared between biologically unrelated classes. As such, this manual also provides guidelines to help the practicing pathologist avoid known diagnostic pitfalls and effectively exclude entities that exist on the differential diagnosis.

Toward potentially being able to automate such large-scale classification schemes, we recently used 1691 whole slide images to generate approximately 850,000 images spanning 74 different tissues classes to train a CNN.[21] By clustering the classes based on 512 high-level features optimized during the training of the neural network, we found that indeed the unsupervised hierarchical arrangement of classes, based solely on their histologic patterns,

created a framework similar to that proposed in the WHO (**Fig. 4**). Importantly, these unsupervised approaches allow for examination of the individual features that the computer optimized and used to group the tumor classes. This exercise revealed the computer can autonomously optimize and use many of the same morphologic features (eg, perinuclear halos, mucin, epithelium, luminal structures) that pathologists use to guide grouping.[21] Interestingly, even "errors" appeared to be easily explainable. This included grouping of highly vascular tumors (eg, hemangioblastoma and angiomatous meningioma, see **Fig. 4**C) or lesions with a high nuclear-to-cytoplastic ratio (eg, glioblastoma and diffuse large B-cell lymphoma; see **Fig. 4**B). Together, this highlights the ability of unsupervised machine learning approaches to automate advanced organizational tasks and propose a complex grouping of histologic information without the need for direct instruction.

SUMMARY AND OUTLOOK

Although much of pathology can be taught through concrete knowledge and patterns described in textbooks, it is not uncommon for even seasoned pathologists to encounter lesions they have never previously seen. These could be due to extremely rare lesions that arise only a handful of times over a pathologist's career, newly emerging diseases (eg, microcephaly induced by Zika virus infection), changes to therapeutic management of common diseases, or even artifacts that arise during the slide preparation process. Although each of these may be rare events, serendipitous observations and a constant sense of uncertainty is a regular part of everyday pathologic analysis. As increasing amounts of data are generated in an era in which personalized therapies are on the horizon, these rare examples will likely increase in frequency and could provide clues to subgroups of patients who experience durable and exceptional

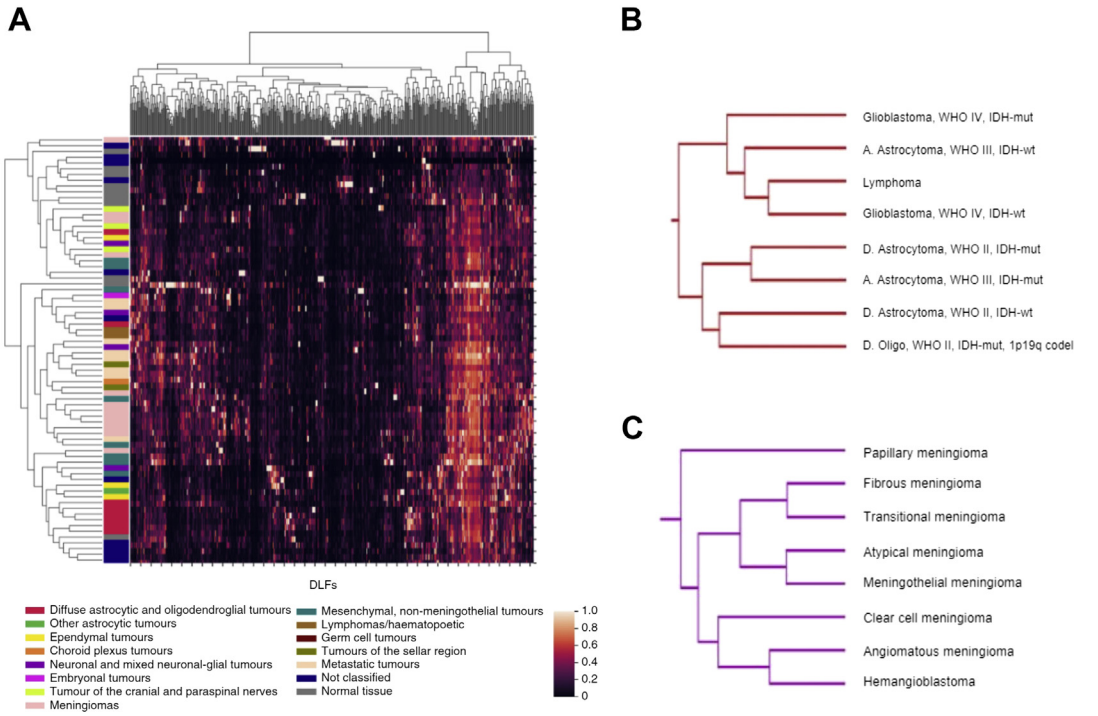

Fig. 4. Clustering for unsupervised ontological arrangement of morphologic patterns by deep learning. (*A*) Hierarchical arrangement of 74 trained classes (vertical axis) based on 512 deep learning derived features (DLF, horizontal axis). Many tumors belonging to the same board tumor class (eg, metastasis, diffuse gliomas, meningiomas) are grouped together. Other tumors were grouped based on overlapping morphologic features (mucin in chordoma and chondrosarcoma). (*B*) and (*C*) provide representative subtrees of the overall dendrogram, highlighting the unsupervised clustering/relationships detected between tumor subtypes based on the global patterns of deep learning feature activations. (*Adapted from* Faust K, Bala S, van Ommeren R, et al. Intelligent feature engineering and ontological mapping of brain tumour histomorphologies by deep learning. *Nat Mach Intell*. 2019;1(7):316-321; with permission.)

responses to treatment. AI will need to extend beyond traditional supervised approaches and develop more humanlike unsupervised approaches to pattern recognition in histologic image analysis to be a viable and useful tool for human pathologists.

In this article, we presented and discussed some of our own experiences with trying to overcome the challenges facing machine learning by shifting toward more unsupervised learning techniques and outputs. Although there could be concerns that these tools will lead to the eventual replacement of human pathologists, we hope the provided examples highlight that even after these "intelligent" outputs are generated, these automated observations require significant human insights to properly vet the findings in the context of other molecular and clinical information. In an era in which digital data continue to grow at an exponential rate, it is likely that these tools will serve to augment the productivity of pathologists and allow them to serve more important integrative roles in both the clinical and investigational components of their profession.

AUTHOR CONTRIBUTIONS

All authors contributed to the article equally.

ACKNOWLEDGMENTS

Funding support for the Diamandis Lab and trainees is provided by the Princess Margaret Cancer Foundation, University of Toronto's Department of Laboratory Medicine and Pathobiology, American Society of Clinical Oncology Career Development Award (ASCO-CDA), and The Brain Tumour Charity Expanding Theories Research Grant (GN-000560). U. Djuric is supported through a University of Toronto Department of Laboratory Medicine and Pathobiology research fellowship award. The authors thank Ms Clare Fiala for constructive feedback.

COMPETING INTERESTS

The authors declare no conflicts of interest.

REFERENCES

1. Djuric U, Zadeh G, Aldape K, et al. Precision histology: how deep learning is poised to revitalize histomorphology for personalized cancer care. NPJ Precis Oncol 2017;1. https://doi.org/10.1038/s41698-017-0022-1.
2. Jha S, Topol EJ, K.C., et al. Adapting to artificial intelligence. JAMA 2016;316(22):2353.
3. Sarwar S, Dent A, Faust K, et al. Physician perspectives on integration of artificial intelligence into diagnostic pathology. NPJ Digit Med 2019;2(1):28.
4. Lakhani P, Sundaram B. Deep learning at chest radiography: automated classification of pulmonary tuberculosis by using convolutional neural networks. Radiology 2017;284(2):574–82.
5. Campanella G, Hanna MG, Geneslaw L, et al. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. Nat Med 2019;25(8):1301–9.
6. Coudray N, Ocampo PS, Sakellaropoulos T, et al. Classification and mutation prediction from non-small cell lung cancer histopathology images using deep learning. Nat Med 2018;24(10):1559–67.
7. Yu K-H, Zhang C, Berry GJ, et al. Predicting non-small cell lung cancer prognosis by fully automated microscopic pathology image features. Nat Commun 2016;7:12474.
8. LeCun Y, Bengio Y, Hinton G. Deep learning. Nature 2015;521(7553):436–44.
9. Madabhushi A, Lee G. Image analysis and machine learning in digital pathology: challenges and opportunities. Med Image Anal 2016;33:170–5.
10. Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. 2014. Available at: https://scholar.google.com/scholar?q=Very%20Deep%20Convolutional%20Networks%20for%20Large-Scale%20Image%20Recognition.%20arXiv%202015.
11. Deng J, Dong W, Socher R, et al. ImageNet: A Large-Scale Hierarchical Image Database. Available at: http://www.image-net.org/papers/imagenet_cvpr09.pdf. Accessed December 30, 2017.
12. Mobadersany P, Yousefi S, Amgad M, et al. Predicting cancer outcomes from histology and genomics using convolutional networks. Proc Natl Acad Sci U S A 2018;115(13):E2970–9.
13. Ebert LC, Heimer J, Schweitzer W, et al. Automatic detection of hemorrhagic pericardial effusion on PMCT using deep learning - a feasibility study. Forensic Sci Med Pathol 2017;13(4):426–31.
14. Esteva A, Kuprel B, Novoa RA, et al. Dermatologist-level classification of skin cancer with deep neural networks. Nature 2017;542(7639):115–8.
15. Xie Q, Faust K, Van Ommeren R, et al. Deep learning for image analysis: personalizing medicine closer to the point of care. Crit Rev Clin Lab Sci 2019;56(1):61–73.
16. Litjens G, Sánchez CI, Timofeeva N, et al. Deep learning as a tool for increased accuracy and efficiency of histopathological diagnosis. Sci Rep 2016;6(1):26286.

17. Diamandis P, Aldape KD. Insights from molecular profiling of adult glioma. J Clin Oncol 2017;35(21):2386–93.
18. Diamandis P, Aldape K. World Health Organization 2016 classification of central nervous system tumors. Neurol Clin 2018;36(3):439–47.
19. Faust K, Xie Q, Han D, et al. Visualizing histopathologic deep learning classification and anomaly detection using nonlinear feature space dimensionality reduction. BMC Bioinformatics 2018;19(1):173.
20. Louis DN, Perry A, Reifenberger G, et al. The 2016 World Health Organization Classification of Tumors of the Central Nervous System: a summary. Acta Neuropathol 2016;131(6). https://doi.org/10.1007/s00401-016-1545-1.
21. Faust K, Bala S, van Ommeren R, et al. Intelligent feature engineering and ontological mapping of brain tumour histomorphologies by deep learning. Nat Mach Intell 2019;1(7):316–21.