



AI Support for Accelerating Histopathological Slide Examinations of Prostate Cancer in Clinical Studies

Mauro Del Rio¹, Luca Lianas¹, Oskar Aspegren^{2,4},
Giovanni Busonera¹, Francesco Versaci¹, Renata Zelic^{3,4},
Per H. Vincent^{3,4}, Simone Leo¹, Andreas Pettersson^{3,4}, Olof Akre^{3,4},
and Luca Pireddu¹

¹ Visual and Data-Intensive Computing, CRS4, 09050 Pula (CA), Italy
{mauro.delrio,luca.lianas,giovanni.busonera,francesco.versaci,
simone.leo,luca.pireddu}@crs4.it

² Department of Clinical Pathology and Cancer Diagnostics,
Karolinska University Hospital, 17176 Stockholm, Sweden

³ Department of Pelvic Cancer, Karolinska University Hospital,
17176 Stockholm, Sweden

⁴ Department of Molecular Medicine and Surgery, Karolinska Institutet,
17176 Stockholm, Sweden KI, Sweden

Abstract. While studies in pathology are essential for the progress in the diagnostic and prognostic techniques in the field, pathologist time is becoming an increasingly scarce resource, and can indeed become the limiting factor in the feasibility of studies to be performed. In this work, we demonstrate how the Digital Pathology platform by CRS4, for supporting research studies in digital pathology, has been augmented by the addition of AI-based features to accelerate image examination to reduce the pathologist time required for clinical studies. The platform has been extended to provide computationally generated annotations and visual cues to help the pathologist prioritize high-interest image areas. The system includes an image annotation pipeline with DeepHealth-based deep learning models for tissue identification and prostate cancer identification. Annotations are viewed through the platform's virtual microscope and can be controlled interactively (e.g., thresholding, coloring). Moreover, the platform captures inference provenance information and archives it as RO-Crate artifacts containing data and metadata required for reproducibility. We evaluate the models and the inference pipeline, achieving AUC of 0.986 and 0.969 for tissue and cancer identification, respectively, and verifying linear dependence of execution speed on image tissue content. Finally, we describe the ongoing clinical validation of the contribution, including preliminary results, and discuss feedback from clinical professionals regarding the overall approach.

Keywords: Digital pathology · Artificial intelligence · Workflows

M. Del Rio and L. Lianas—These authors contributed equally to this work.

© The Author(s) 2022

P. L. Mazzeo et al. (Eds.): ICIAP 2022 Workshops, LNCS 13373, pp. 545–556, 2022.

https://doi.org/10.1007/978-3-031-13321-3_48

1 Introduction

The examination and annotation of images of tissue samples is a critical part of clinical studies in the field of pathology – that is, the medical discipline that diagnoses disease from tissue and relevant clinical data. In such studies, tissue characteristics are identified and analyzed, along with other patient traits, in the search for information useful to improve clinical practice. The tissue samples are collected by either surgery or biopsy. Slices of these specimens are then mounted on glass slides, stained (coloured) and, finally, scanned by special-purpose digital optical microscopes. The resulting digital pathology images – often referred to as slides, like the glass slides from which they are scanned – are very high in resolution (typically 40× magnification, which results in images with a resolution in the order of tens of Gigapixels and a size 1–5 GB per compressed image). The examination and image annotation required of pathologists to conduct such clinical studies is defined in detail by the study protocol, and it is generally much more time-consuming than examining slides for regular clinical diagnostic practice. For instance, study protocols may require the annotation and measurement of additional exploratory features that are not currently part of any clinical diagnostic or prognostic protocol. Therefore, pathologist time tends to become the limiting factor in the feasibility of the studies – a situation worsened by the declining number of pathologists and the rise in biopsy volumes. Computer-aided pathology holds a lot of potential to ameliorate the situation by accelerating the slide examination and annotation process and reducing the required human effort through the use of AI methodologies [21].

In this setting, the support of bespoke software can be very valuable to facilitate the rigorous application of the study protocol and the correct management of such a study: the Digital Pathology (DP) Platform by CRS4 [27] is a system exactly for this purpose. It provides a platform for configuring vertical applications for managing, examining and annotating digital pathology images within the context of clinical research. It supports all major virtual slide formats and provides users with a web application that includes a virtual microscope, allowing the on-line interactive remote visualization and annotation of very large virtual slides, without loss of quality with respect to optical microscopes [33]. The platform also provides a palette of tools to accurately draw and measure regions of interest (ROIs) following irregular tissue contours. Using the platform, a vertical application has been created for studies on prostate cancer. In addition to general annotations, the application adds specifically detailed prostate cancer ROI annotation labels. A minimally customized variant of this configuration of the platform has been used to conduct the Prognostic Factors for Mortality in Prostate Cancer (ProMort) study [34], and another is currently being used in a study aiming to improve prostate cancer prognosis through the integration of advanced statistical modeling and the inclusion of new prognostic variables [28].

In this work, we describe how we extended the DP Platform with AI-based functionality to better support the examination and annotation of digital pathology slides in the context of clinical studies. The remainder of the manuscript is structured as follows: Sect. 2 provides relevant background; Sect. 3 describes the

main contributions, including the inference pipeline, the deep learning models, the visualization strategy and the collection of provenance information; Sect. 4 describes the evaluation and discusses the results; finally, Sect. 5 concludes the work.

2 Background and Related Work

The DeepHealth Toolkit. The DeepHealth toolkit [8] is an open source deep learning framework. It is specifically tailored to be easily applied to biomedical data (for instance, it includes functionality to read digital pathology images) and to leverage heterogeneous computing resources (e.g., high-performance computing clusters, cloud computing platforms, GPU and FPGA accelerators). The toolkit’s deep learning (EDDL) and computer vision (ECVL) libraries have been used (through their respective Python APIs) to implement the AI functionality described in this work.

Related Work. To the best of the authors’ knowledge, no other software platform has been published that aims to support the execution of clinical studies in digital pathology in a way analogous to the DP Platform. However, much work has been published in related fields. For instance, OMERO provides whole-slide image (WSI) data management functionality, and it is also one of the key components of the DP Platform; it provides generic key-value image annotation that is not specialized on any particular domain. Other tools like QuPath [5], Orbit [25], FastPathology [20], ASAP [4], PathML [6] provide functionality to view and annotate slides; some even support machine/deep learning based segmentation and classification. However, these tools do not aim to support the execution of clinical study protocols, nor to provide domain-specific annotation tools. Recently, the PANDA challenge [7] has catalyzed efforts on the application of deep learning techniques to prostate cancer histopathology; however, PANDA focuses on prostate cancer prognosis, rather than identifying cancer tissue in prostate tissue images. Work has also been done on characterizing the aggressiveness of cancer tissue based on its appearance [26].

3 Slide Examination Support System

The DP Platform is a multi-component system consisting of an image repository/server, an annotation management service and a web application. The image repository is based on the OpenMicroscopy OMERO server [1] – the same system that is behind some large-scale public image repositories [13, 30] – which has been extended with the purpose-built `ome.seadragon` software component [17] to add support for the Deep Zoom Image format [10] and web-based viewers for high-resolution zoomable images, such as the OpenSeadragon [18]. The user-facing web application interacts with the image and annotation services to provide functionality such as the virtual microscope, the annotation tools, etc.

In this work, the DP Platform has been augmented with new functionality to support the examination of histopathological slides by pathologists through computational annotation of the images. This functionality has been achieved by implementing a computational annotation pipeline that integrates multiple specialized deep learning models for image analysis, as well as custom visualization and examination tools that leverage the model predictions. In addition, to enhance the reproducibility of the computational results, the provenance information of the predictions is captured in RO-Crate artifacts [24] and stored with the annotations. The extended platform’s architecture is illustrated in Fig. 1. The following subsections describe the components in more detail.

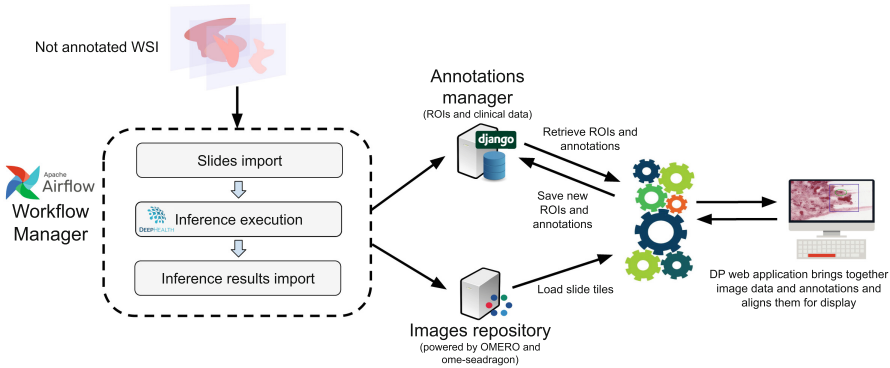


Fig. 1. The extended DP platform software architecture

3.1 Computational Annotation Pipeline

Computational annotation of digital pathology images has been integrated into the DP Platform through its slide import and pre-processing workflow. The workflow has been extended to apply AI-based analyses to digital pathology images at this stage. The added complexity of the process motivated the integration of the Apache Airflow workflow manager [3] into the platform as a process automation subsystem, providing sophisticated workflow execution and monitoring functionality (both programmatic and graphical).

As illustrated in Fig. 1, the image import pipeline is composed of three main stages: the first and third perform platform-specific data and metadata management; the “Inference execution” stage uses AI models to automatically analyze images. The computational pipeline of this stage is defined using the Common Workflow Language (CWL) [2] and executed using CWL-Airflow [16] – an implementation of CWL for Airflow. The choice of a standard workflow language improves workflow portability and facilitates the integration of novel data provenance approaches (see Sect. 3.4). The inference stage performs the following steps on the digital pathology image: a low-resolution mask of the tissue is inferred

(downsample factor of 2^9) and used to select the areas of the image that warrant further processing; the tissue mask is refined by repeating tissue inference at a higher magnification level (downsample factor of 2^4) on tissue areas recognized in step 1; cancer inference is performed at high magnification level (downsample factor of 2) on tissue areas recognized in step 1.

Each deep learning model is packaged as an executable Docker container image that provides inference functionality through a common interface – these images are based on pre-built DeepHealth toolkit images [11]. The abstract interface exposes the whole slide to the model, allowing it to support any kind of ensemble or complex model design. The generated annotations are stored using either Zarr [32] or TileDB [19]: both are modern formats for storing large N-dimensional arrays in a chunked, compressed and cloud-friendly data structure (at the moment the workflow implementation can be configured to use either solution). The import workflow invokes the inferencing containers, makes the data and previous predictions in the pipeline accessible to them (by mounting appropriate data volumes on the container), retrieves the resulting annotations and imports them into the annotations manager. The DP import workflow can be configured to run any number of annotating container images.

3.2 Deep Learning Image Annotation Models

Identification of Tissue. We have created a neural network model with the PyEDDL and PyECVL libraries for the identification of tissue areas on histological images. This model allows the platform to completely automate the tissue annotation phase in the clinical annotation process. To create the model, training, validation, and test sets were generated by manual annotation of background and tissue areas from a set of sample slides. Selected samples also contained different kinds of objects that can be found in the background – for example, markers and glue – which can easily generate errors in simple automated tissue recognition approaches, and which were also problematic for preliminary versions of the model. For this particular task, we defined a pixel-based model architecture made of dense layers which classifies input RGB pixels one-at-a-time, as opposed to patch-based architectures that are frequently used for image segmentation problems – such as U-Net [22] – which produce a single output for a given patch of pixels in input. This approach made it very simple to generate training data for the model by selecting blocks of pixels from tissue and background image areas, with no need to annotate samples of the tissue-background interface.

Identification of Prostate Cancer Areas. A second deep learning model has been developed to recognize adenocarcinoma of the prostate. Based on initial experiments, a patch classification approach was adopted: we decompose the histopathology slides into patches of size 256×256 pixels and classify each tissue patch as a whole as benign/malignant. Moreover, initial exploratory work brought us to select the VGG-16 network architecture [23]. The convolutional

layers of the network are pre-trained with the ImageNet dataset [12] (without freezing any layer parameter), improving convergence speed and overall accuracy with respect to starting with models that were either completely untrained or using the Glorot [14] or He [15] initializations. To mitigate the effects of overfitting, dropout layers were added to the classification part of the network. Other solutions were also tested during model development, ranging from regularization techniques to data augmentation, as well as reducing the number of parameters of the fully connected classifier and using staining normalization – based on stain separation with preservation of the sample structure. These did not provide significant improvements in inference performance.

To create the model, a set of 417 slide images were scanned using a 3DHi-tech Panoramic 250 Flash II at $40\times$ magnification (pixel size of approximately $0.1945\ \mu\text{m}/\text{pixel}$). The slide images were examined and annotated by pathologists (using the DP Platform). From these, a training dataset of about 123K patches was generated (38K normal and 85K cancer patches) and saved to a scalable Cassandra-based patch repository [29]. Leveraging our custom Cassandra-based data loader, we created two kinds of balanced split configurations: the first, composed of two splits (80% training and 20% validation), was used for rapid evaluation of different training hyperparameter configurations; the second, composed of five equally sized splits, was used to more robustly evaluate promising models through cross-validation. Like the tissue identification model, this one was also developed with the PyEDDL and PyECVL libraries.

3.3 Visualization of Computational Annotations

The DP Platform has been extended with functionality to visualize the results of inference by deep learning models and use them to provide visual cues to assist the end user in the examination of the image. The outputs of the models developed for this use case are 2D arrays of “scores” at the pixel or patch level: in particular, the tissue detection model classifies single pixels, while cancer prediction works on whole patches. We have implemented two alternative visualization methods for these computational annotations: the heatmap and the vectorial ROI. These methods have both been integrated into the DP Platform’s virtual microscope and produce visual artifacts that can be dynamically controlled by the user and overlaid onto the image (see Fig. 2).

The heatmap visualization (Fig. 2a), which is used for the cancer detection feature, is used to focus the attention of the pathologist on potentially relevant regions (like specific cancer patterns). To render the model outputs as heatmaps, the results produced by the models are registered into the DP Platform as new annotations of the slide and stored as arrays; they are then rendered at run time by a dedicated module by appropriately slicing the data and applying a color map. The web application includes controls to dynamically adjust the heatmap’s opacity and to specify a threshold to cut off non-significant values.

On the other hand, the visualization of vectorial ROIs, which are used for the tissue detector output, requires additional post-processing on the model’s 2D matrix output. Specifically, the continuous value matrix is transformed into

a boolean mask by applying a configured threshold value. The mask is further processed to identify clusters of “true” values bigger than a set threshold (smaller clusters are not useful for the clinical review) and compute the enclosing geometry defining the ROI (Fig. 2b).

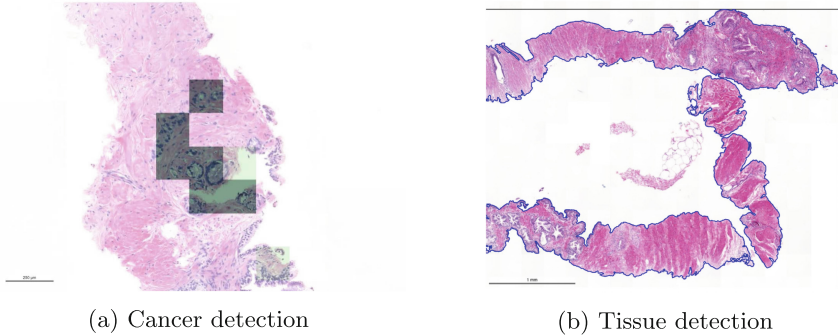


Fig. 2. Heatmap visualizations and tissue ROIs rendered in the Virtual Microscope, generated respectively from the outputs of the cancer and tissue identification models.

3.4 Prediction Provenance with RO-Crate

At the conclusion of the inference workflow execution for a given slide, the DP Platform generates an RO-Crate object (with the `ro-crate-py` library [9]) to capture the provenance data of the tissue and cancer predictions generated by the inference process. The RO-Crate references the slide, the predictions, the CWL inference workflow definition and its input parameters (including the model’s specific Docker image tag). The RO-Crates thus created serve as archivable snapshots of prediction runs, and they are expected to greatly enhance the reproducibility of the predictions; they can even be a part of a more articulate graph that also documents the provenance of the slides by using the Common Provenance Model [31].

4 Evaluation and Discussion

Evaluation of the Tissue Predictive Model. The tissue predictive model was trained using a balanced dataset of about 30M pixels (i.e., 30M RGB triplets). The training took about 18.5 min (30 epochs, about 37 s for each epoch) running on a Nvidia RTX 2080Ti GPU with 12 GB RAM. We evaluated the model’s accuracy with a balanced test set of 6M pixels. A maximum accuracy of 0.96 was achieved by setting the classification threshold (to assign the class of a pixel starting from its prediction score) to 0.65. The ROC curve (Fig. 3a) shows that the model can have its maximum sensitivity without increasing the false positive rate beyond 10%. The AUC is 0.986; the F1-score, which is 0.96 when using the threshold classification value of 0.65.

Evaluation of Cancer Predictive Model. The cancer prediction model was trained starting from a dataset composed of 384 annotated slides (taken from 200 different cases) resulting in 123,148 patches measuring 256×256 pixels each. The training was performed on a single GPU (Nvidia Titan RTX with 24 GB RAM) and the epoch time was on average 467 s (std 3.8 s). The minimum number of epochs needed to get the best validation accuracy, including 20 epochs for early stopping, ranged from 35 to 73.

The final evaluation was performed on a test set generated from 149 slides (taken from 109 cases), from which 66,698 patches were extracted. We characterized model performance through a five-fold cross validation procedure, from which we produced five different models. All test patches were classified with each of the five models. The predictions were averaged to have a single score that could be compared with the true labels with a soft voting approach. The resulting maximum accuracy is 0.91 with a classification threshold set to 0.58. With the same threshold, the F1-score is 0.9. Figure 3b shows the ROC curve, with an AUC equal to 0.969. All training, validation and test sets were generated to be balanced with respect to the two classes.




Fig. 3. ROC curves for the image analysis models

Computational Inference Pipeline Performance. We tested the inference pipeline described in Sect. 3.1 to characterize its overall speed and verify its linear progression with respect to the slide's tissue content. Experiments were conducted on node equipped with an Intel(R) Xeon(R) W-2145 CPU @ 3.70 GHz (8 cores, 16 threads), 128 GB RAM, a NVIDIA Quadro RTX 5000 (16 GB RAM), and a 3-HDD ZFS. The test set was composed of 203 WSI. Each slide was captured with a 3DHistech Panoramic 250 Flash II scanner at a magnification of $40 \times$ (pixel size of approximately $0.1945 \mu\text{m}/\text{pixel}$) and has a resolution of $112,908 \times 265,513$ pixels; the average tissue coverage per slide is 1.5%. Note that

all prediction times are comprehensive of I/O operations for reading slides and writing results.

We report wall clock execution times of the inference pipeline in Fig. 4a. The various pipeline stages are executed serially and each occupies the single GPU on the test system. The difference between the execution time of the whole pipeline and the sum of the individual steps is used for the data management work performed by the pipeline, such as packaging the output of the models as Zarr arrays and registering them with the platform.

The execution time of the tissue mask refinement and cancer identification steps is dependent on the amount of tissue on the image, which explains the increased variability in the execution times for these steps. Our experiments show that the execution time of the pipeline’s application of these models grows linearly with the image’s tissue coverage (see Fig. 4b).

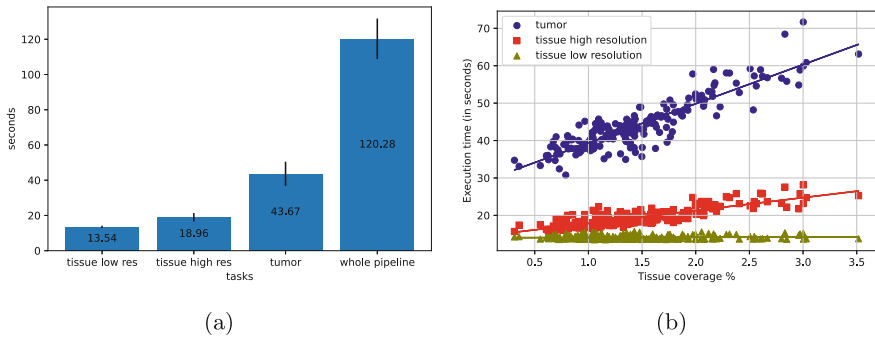


Fig. 4. (a) Average execution times of the prediction pipeline and its tasks. (b) Execution time of inference steps with respect to tissue coverage.

Clinical Evaluation. The AI-based slide examination functionality that has been added to the DP Platform is undergoing clinical evaluation based on measuring the time of annotation and assessment of the slides by the pathologist; this evaluation parameter is in line with the main goal of the platform’s AI-based features to support the examination by an expert, rather than fully automate the process. While the full evaluation is not yet complete, in this work we report the preliminary results, as well as the feedback collected from the evaluating pathologist.

The review protocol followed for the clinical evaluation involves two steps: identification and measurement of the ROIs of relevance in the slide (in particular, identification of tissue cores and invasive cancerous areas within them); for each identified ROI, annotation of relevant clinical features (e.g., cancerous or inflammatory patterns, cancer staging scores, etc.). The evaluation is being conducted in three different stages: 1) Baseline, without any AI support tools; 2) Intermediate, with the examination supported by the tissue recognition model;

3) Final: with the examination supported by both the tissue and cancer identification models.

We present results for the Baseline and Intermediate evaluation stages. Annotation times were measured on 139 slides for the Baseline stage and 131 slides for the Intermediate stage. The clinical examination focused on the single cancerous tissue core per slide considered most relevant and sufficient for diagnostic purposes. For the Baseline stage, we measured an average annotation time of 11 min. For the Intermediate stage, the average annotation time per slide was 9 min. At the time of writing, the Final evaluation phase is in progress and no data on annotation time is available yet.

These preliminary evaluations suggest that the AI-based support tools presented here successfully reduce the time required for cancer slide examination and annotation by a pathologist. We expect the addition of sufficiently accurate cancer-identification functionality to further accelerate annotation in the Final stage, as the pathologist will no longer need to examine the full area of each tissue core in search for cancer areas, but only those highlighted by the model. A preliminary evaluation by the pathologist of the predictions of invasive prostate adenocarcinoma revealed a very low false-negative error rate and a false-positive error rate that needs to be addressed and improved. The overall impression regarding the predictions of prostate cancer is quite promising.

5 Conclusion

The new AI-based functionality added to the DP Platform has shown promising results in the reduction of pathologist time required to examine and annotate digitized prostate cancer slides in the context of clinical studies. The accuracy and computational performance of these new tools has been experimentally evaluated and their clinical evaluation is in progress. The software used for this work has been released under the MIT open source license and the code is available on GitHub (https://github.com/deephealthproject/promort_pipeline, <https://github.com/crs4/slaid>, <https://github.com/crs4/deephealth-pipelines>).

Acknowledgements. This work has been supported by the DeepHealth and EOSC-Life projects, both funded by the European Union's Horizon 2020 research and innovation programme (respectively under grant agreement numbers 825111 and 824087), and by the SVDC project funded by the Sardinian Regional Authority.

References

1. Allan, C., et al.: OMERO: flexible, model-driven data management for experimental biology. *Nat. Methods* **9**(3), 245–253 (2012)
2. Amstutz, P., et al.: Common workflow language, v1. 0 (2016)
3. Apache Airflow (2022). <https://airflow.apache.org>
4. ASAP. <https://github.com/computationalpathologygroup/ASAP>
5. Bankhead, P., et al.: QuPath: open source software for digital pathology image analysis. *Sci. Rep.* **7**(1) (2017)

6. Berman, A.G., Orchard, W.R., Gehrung, M., Markowetz, F.: PathML: a unified framework for whole-slide image analysis with deep learning. medRxiv (2021)
7. Bulten, W., et al.: Artificial intelligence for diagnosis and Gleason grading of prostate cancer: the PANDA challenge. *Nat. Med.* **28**, 154–163 (2022)
8. Cancilla, M., et al.: The DeepHealth toolkit: a unified framework to boost biomedical applications. In: 2020 25th International Conference on Pattern Recognition (ICPR), pp. 9881–9888 (2021)
9. De Geest, P., et al.: ro-crate-py (2022). <https://zenodo.org/record/6594974#.YuOBJEzhXGg>
10. Deep Zoom Image format. [https://docs.microsoft.com/en-us/previous-versions/windows/silverlight/dotnet-windows-silverlight/cc645050\(v=vs.95\)?redirectedfrom=MSDN](https://docs.microsoft.com/en-us/previous-versions/windows/silverlight/dotnet-windows-silverlight/cc645050(v=vs.95)?redirectedfrom=MSDN)
11. Deephealth dockerhub organization. <https://hub.docker.com/u/dhealth>
12. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: a large-scale hierarchical image database. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition, pp. 248–255. IEEE (2009)
13. Ellenberg, J., et al.: A call for public archives for biological image data. *Nat. Methods* **15**(11), 849–854 (2018)
14. Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. In: Teh, Y.W., Titterton, M. (eds.) Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics. Proceedings of Machine Learning Research, vol. 9, pp. 249–256. PMLR, 13–15 May 2010
15. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: surpassing human-level performance on ImageNet classification. In: 2015 IEEE International Conference on Computer Vision (ICCV), pp. 1026–1034 (2015). <https://doi.org/10.1109/ICCV.2015.123>
16. Kotliar, M., Kartashov, A.V., Barski, A.: CWL-airflow: a lightweight pipeline manager supporting common workflow language. *Gigascience* **8**(7), giz084 (2019)
17. ome-seadragon (2022). https://github.com/crs4/ome_seadragon
18. OpenSeadragon (2022). <https://openseadragon.github.io>
19. Papadopoulos, S., Datta, K., Madden, S., Mattson, T.: The TileDB array data storage manager. *Proc. VLDB Endowment* **10**(4), 349–360 (2016)
20. Pedersen, A., Valla, M., Bofin, A.M., De Frutos, J.P., Reinertsen, I., Smistad, E.: FastPathology: an open-source platform for deep learning-based research and decision support in digital pathology. *IEEE Access* **9**, 58216–58229 (2021)
21. Regitnig, P., Müller, H., Holzinger, A.: Expectations of artificial intelligence for pathology. In: Holzinger, A., Goebel, R., Mengel, M., Müller, H. (eds.) Artificial Intelligence and Machine Learning for Digital Pathology. LNCS (LNAI), vol. 12090, pp. 1–15. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-50402-1_1
22. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
23. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: Bengio, Y., LeCun, Y. (eds.) 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, 7–9 May 2015, Conference Track Proceedings (2015). <http://arxiv.org/abs/1409.1556>
24. Soiland-Reyes, S., et al.: Packaging research artefacts with RO-Crate. *Data Sci.* **5**(2), 97–138 (2022). <https://zenodo.org/record/6594974#.YuOBJEzhXGg>
25. Stritt, M., Stalder, A.K., Vezzali, E.: Orbit image analysis: an open-source whole slide image analysis tool. *PLoS Comput. Biol.* **16**(2), e1007313 (2020)

26. Ström, P., Kartasalo, K., Olsson, H., Solorzano, L., Delahunt, B., Berney, D.M., Bostwick, D.G., Evans, A.J., Grignon, D.J., Humphrey, P.A., et al.: Artificial intelligence for diagnosis and grading of prostate cancer in biopsies: a population-based, diagnostic study. *Lancet Oncol.* **21**(2), 222–232 (2020)
27. The Digital Pathology platform. <https://github.com/crs4/ProMort>
28. Turin prostate cancer prognostication study. <https://sites.google.com/view/studio-tpcp/>
29. Versaci, F., Busonera, G.: Scaling deep learning data management with Cassandra DB. In: 2021 IEEE International Conference on Big Data (Big Data), December 2021. <https://doi.org/10.1109/BigData52589.2021.9672005>
30. Williams, E., et al.: Image data resource: a bioimage data integration and publication platform. *Nat. Methods* **15**(11), 984–984 (2018). (vol 14, pg 775, 2017)
31. Wittner, R., et al.: ISO 23494: biotechnology – provenance information model for biological specimen and data. In: Glavic, B., Braganholo, V., Koop, D. (eds.) IPAW 2020-2021. LNCS, vol. 12839, pp. 222–225. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-80960-7_16
32. Zarr. <https://zarr.readthedocs.io/en/stable/>
33. Zelic, R., et al.: Interchangeability of light and virtual microscopy for histopathological evaluation of prostate cancer. *Sci. Rep.* **11**(1), 1–11 (2021)
34. Zelic, R., et al.: Estimation of relative and absolute risks in a competing-risks setting using a nested case-control study design: example from the promort study. *Am. J. Epidemiol.* **188**(6), 1165–1173 (2019)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

