



Developments in AI and Machine Learning for Neuroimaging

Shane O'Sullivan¹(✉) , Fleur Jeanquartier² , Claire Jean-Quartier² ,
Andreas Holzinger² , Dan Shiebler³, Pradip Moon⁴ ,
and Claudio Angione^{4,5} 

¹ Department of Pathology, Faculdade de Medicina Universidade de Sao Paulo, São Paulo, Brazil

doctorshaneosullivan@gmail.com

² HCI-KDD, Holzinger Group, Institute for Medical Informatics and Statistics, Medical University of Graz, Graz, Austria

{f.jeanquartier,c.jeanquartier}@hci-kdd.org

³ Department of Cognitive Linguistic and Psychological Sciences, Carney Institute for Brain Science, Brown University, Providence, RI, USA

danshiebler@gmail.com

⁴ Department of Computer Science and Information Systems, Teesside University, Middlesbrough, UK

{P.Moon,C.Angione}@tees.ac.uk

⁵ Healthcare Innovation Centre, Teesside University, Middlesbrough, UK

Abstract. This paper reviews guidelines on how medical imaging analysis can be enhanced by Artificial Intelligence (AI) and Machine Learning (ML). In addition to outlining current and potential future developments, we also provide background information on chemical imaging and discuss the advantages of Explainable AI. We hypothesize that it is a matter of AI to find an invariably recurring parameter that has escaped human attention (e.g. due to noisy data). There is great potential in AI to illuminate the feature space of successful models.

Keywords: Explainable AI · Stereology · Neurodegenerative diseases · Neuroimaging · Disector · 7T post-mortem MRI · Brain mapping

1 Introduction and Motivation

In this paper, our approach is to discriminate between: general challenges to AI/ML image analysis (irrespective of the specific approach); challenges specific to Black Box methods; challenges that Explainable AI alone can help to overcome; and challenges that Black Box methods together with Explainable AI can help to overcome. In addition to outlining the challenges and our hypotheses, we also include an extensive review to assess the use of AI and ML-enhanced

We thank the organization CNPQ (Brazilian National Council for Scientific and Technological Development). This entity provided support that was invaluable to our research.

© Springer Nature Switzerland AG 2020

A. Holzinger et al. (Eds.): AI & ML for Digital Pathology, LNAI 12090, pp. 307–320, 2020.

https://doi.org/10.1007/978-3-030-50402-1_18

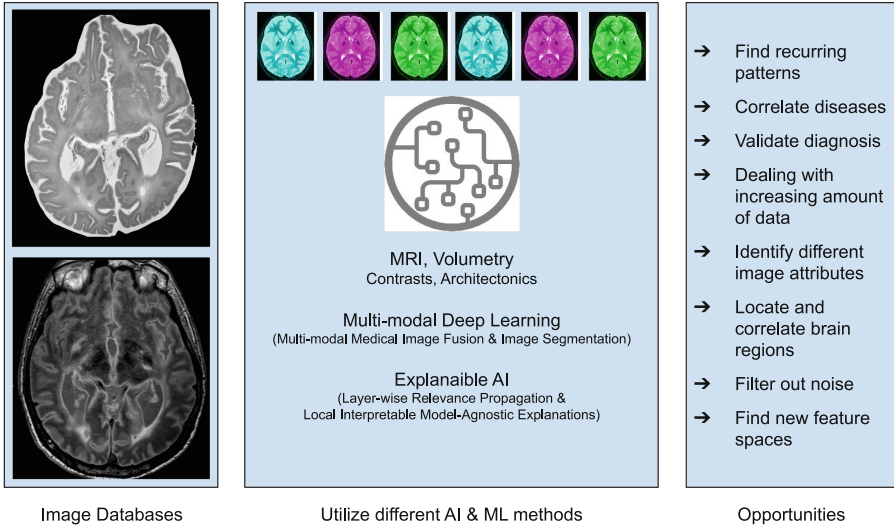


Fig. 1. Graphical abstract (brain images from [52])

magnetic resonance imaging (MRI) for the evaluation of neurodegenerative diseases in comparison to the use of histological information (Fig. 1).

2 Glossary

The following abbreviations and terms are used in this manuscript:

AD	Alzheimer Disease
AI	Artificial Intelligence
CAG	Cytosine-adenine-guanine
CV	Cross Validation
DL	Deep Learning
Explainable AI	Explaining the behavior of otherwise black-box models
fMRI	functional Magnetic Resonance Imaging
HD	Huntington’s Disease
LDA	Linear Discriminant Analysis
ML	Machine Learning
MRA	Magnetic Resonance Angiography
MRI	Magnetic Resonance Imaging
Neurodegenerative diseases	Diseases caused by the death of neurons and/or other cells in the brain
NMR	Nuclear Magnetic Resonance
PCA	Principal Component Analysis
PD	Parkinson’s disease
RF	Random Forests

RFDA	Regularized Fisher Discriminant Analysis
SC-CNN	Spatially Constrained Convolutional Neural Network
SPECT	Single Photon Emission Computed Tomography
Virtual Autopsy	The use of imaging to diagnose cause of death
USG	Ultrasonography

3 State-of-the-Art

3.1 Role of AI/ML

We define AI as the whole field working on understanding intelligence towards context-adaptive systems; the backbone, however, is ML as a methodological subset of AI, whilst e.g. Deep Learning (DL) is just one specific methodological subset of ML [24]. In the following we only briefly point to some of the manifold possibilities which AI/ML offer for the topic of this paper. The most important methods for classification of brain states include Linear Discriminant Analysis (LDA) and regularized Fisher Discriminant analysis (RFDA) to get a classifier, Independent Component Analysis (ICA) and Principal Component Analysis (PCA) for dimensionality reduction of the input data, several cross-validation (CV) schemes for model evaluation/selection, however, keeping potential pitfalls in mind [35, 52].

The human brain has a significantly high individual variability with substantial amount of information that needs to be analysed. To count neurons manually can be an extremely time-consuming process to perform by human eye observations alone. Non-homogeneous brain shrinkage is one major challenge where, in combination, imaging and AI/ML may be effective. By providing a huge amount of information for histology, several ML methods can help to validate in vivo MRI findings and MRI limitations. This leads to improving and enhancing the capabilities of MRI through the identification of even more alterations. Moreover, in brain informatics problems include the detection and interpretation of volume changes in neurodegenerative diseases and more global aspects such as brain volume, gyration, cortical thickness, hippocampal shape changes and allometric studies correlating, for example, cortical volume with hippocampal and subcortical nuclear volume in schizophrenics.

As a first step, by using different contrasts in MRI, one can obtain more information, not only regarding volume quantification, but also iron deposition, myelin quantification and brain parcellation with myeloarchitectonics [39] and other sequences. Architectonics is considered to reflect the functional properties of the brain and its surface. The manifold problems with such approaches include the high-dimensionality of the data, which would make standard approaches of AI awkward, yet impossible to use, therefore, it is necessary to harness the full potential of current AI/ML. As a next step, one can attempt to quantify the architectonics, the size of defined fields, and correlate with their normative data for different diseases and normal ageing. Due to the high individual variability of the brain, these studies will require investigating extensive data. For this step, accessing and analysing the large database using simple methods will most

likely lead to a null finding. We hypothesize that this is a matter of AI to find an invariably recurring parameter that has escaped human attention (e.g. due to noisy data).

However, one must take care when discussing AI in the context of initial/early stages of neurodegenerative or neuropsychiatric diseases. It is necessary to perform both MRI and volumetry when there is no access to the tissue. It is also important to take into consideration that changes other than visible changes (e.g. volume) are causing symptoms in the early and later stages of schizophrenia or depression. In addition, from a clinical and neuropsychiatric view, schizophrenia itself is very challenging and complicated. It is the same matter as assessing proficiency based on the size of the brain. Many would assume that bigger brains are more proficient than smaller brains, however, the contrary can be true because one can find small brains in highly intelligent people and vice versa.

Generally, next to imaging there are genetic biomarkers for disease classification but these are applicable only in specific cases. Established biomarkers for Alzheimer's disease (AD), Parkinson's disease (PD) and Huntington's disease (HD) are based on visual analysis including MRI. Genetic variants of apolipoprotein E, tau as well as amyloid precursor protein or presenilin 1 and 2 are significant for AD, however, they are not imperative for the development of the disease, thus, are not reliably applicable for diagnosis. So far, AD is diagnosed clinically by a patient's detailed history and mental state and finally determined pathologically by brain autopsy [10,70].

The use of imaging to diagnose the cause of death (virtual autopsy) has already proven its value, although it has limitations, especially regarding microscopic changes [52]. The possibility of having postmortem MRI in a high number of cases, with histology included, is an invaluable combination to validate the usefulness of MRI in this setting. This is in regards to, not only cases, but also healthy controls, and it provides more understanding to populational variability. Imaging yields an immense number of data and data needs to be interpreted. There is a region of transition between pure statistical analysis of data and interpretation by AI. One strategy is the investigation of brains from persons with well-characterized disabilities. A good example is Alzheimer's Disease, where the visual inspection of neuroimaging is susceptible to limitations of human vision; here AI methods have shown to be equally or even more effective than human clinicians in diagnosing dementia from neuroimages [16].

Pathognomonic indicators include cerebral atrophy and neurofibrillary tangle and amyloid plaque pathology. The likelihood and/or course of HD can be genetically tested by the determination of the number of cytosine-adenine-guanine (CAG)-repeats [12] of Huntingtin; the gene product of the affected gene on chromosome 4. The neuropathological interaction between early striatal and cortical atrophy proved to be puzzling [17] and imaging of prodromal cerebral cortical changes is difficult to detect [48]. An exemplary novel approach of disease characterization has been recently described for AD. The technique is based on imaging of brain structural connectivity atrophy in combination with a multiplex network for generating a classification score [1]. Automated differentiation of PD

has been described also based on various ML derived classification algorithms using quantitative MRI data [19].

3.2 AI/ML Methods

A myriad of different AI and ML methods exist, so we scratch the surface here and focus our description only on Deep Learning (DL), which we recommend for these particular studies. We refer the reader to [21] for an overview of ML in general, and to [6, 14, 20] for more specific details. AI-aided diagnosis has been used as a supporting tool for physicians for a long time [5, 63]. Due to the increasing computational power and available storage capacities, many methods, which proved to be computationally-demanding in the past, can now sufficiently be used in daily routine. Standard examples of state-of-the-art methods today, include not only DL [2], but also Support Vector Machines (SVM) [11] and Random Forests (RF) [7].

Among these methods, DL is rapidly proving to be the state-of-the-art foundation, leading to improved accuracy. It has also opened up new frontiers in biomedical data analysis generally, and in clinical medicine specifically, with astonishing rates of progress [15]. In favour of AI-aided diagnosis, quantitative changes could be evaluated in combination with functional neuroimaging and interpretation of big data in a longitudinal setting.

Recent work has been performed relating to automatic detection and classification of cell nuclei in histopathological images of cancerous tissue [61]. The authors applied DL [33] and produced encouraging results by applying a so-called Spatially Constrained Convolutional Neural Network (SC-CNN) [53] to perform nucleus detection. SC-CNN regresses the likelihood of a pixel being the centre of a nucleus, where high probability values are spatially constrained to locate in the vicinity of the centers of nuclei.

However, the current approaches in ML and neuroimaging do not facilitate essential mechanistic investigations validation by way of histology. Rather than only showing the ability to detect patterns of brain alterations, ML can also benefit from improving knowledge about algorithm choices and particular characteristics of precision power related to specific disease mechanisms.

3.3 Multimodal Deep Learning in Medical Imaging

In the machine learning context, algorithms dealing with data from multiple heterogeneous sources are referred to as “multimodal” or “multi-view” learning algorithms [47]. The advantages of using multimodal deep learning in the biomedical context are: (i) they require little or no pre-processing input data, because both features and fused representations are learned from data; (ii) they perform implicit dimensionality reduction within the architecture, which is a desired property in feature-rich biomedical datasets; (iii) they support early, late, or intermediate fusion [69]. However, they usually require powerful graphics processing units (GPUs) for reasonable training time.

Multimodal deep learning can be used to solve complex machine learning problems in areas of high dimensional unstructured data like computer vision, speech and natural language processing. The main advantage of deep learning is that, it automatically learns hierarchical representation for each modality instead of manually designing modality-specific features that are then fed with machine learning algorithm. In medical image analysis, the medical expert can use multiple image modalities information, e.g. computed tomography (CT), magnetic resonance imaging (MRI), and ultrasound imaging for diagnosis and treatment. Therefore, multimodal deep learning is suitable for medical applications issues like tissue and segmentation, multimodal medical image retrieval and computer-aided diagnosis. There are however two significant challenges faced by the medical applications community when using multimodal deep learning, namely the difficulty in obtaining sufficient labelled data, and class imbalance [57].

Multimodal deep learning is widely used for brain imaging studies. Collecting data of magnetic resonance imaging (MRI) of multiple modalities of the same individual is popular in brain imaging studies. Multimodal brain imaging study can provide a more comprehensive understanding of the brain and its disorders. For instance, it can inform us about how brain structure shapes brain function, in which way they are impacted by psychopathology, and which structural aspects of physiology could drive human behaviour and cognition [9].

Multimodal medical image fusion techniques are the most significant methods to identify and investigate disease to provide complementary information from different multimodalities. Multimodal medical images can be categorized into several types, which include computed tomography (CT), magnetic resonance angiography (MRA), magnetic resonance imaging (MRI), positron emission tomography (PET), ultra sonography (USG), nuclear magnetic resonance (NMR) spectroscopy, single photon emission computed tomography (SPECT), X-rays, visible, infrared and ultraviolet. Structural therapeutic images are MRI, CT, USG and MRA, which provide high-resolution images. Functional therapeutic images are PET, SPECT and functional MRI (fMRI) which provide low-spatial resolution images with useful information.

Multimodal medical image fusion increases the effectiveness of image-guided disease analysis, diagnosis and assessment of medical problems. Image fusion has several applications like medical imaging, biometrics, automatic change detection, machine vision, navigation aid, military applications, remote sensing, digital imaging, aerial and satellite imaging, robot vision, multi focus imaging, microscopic imaging, digital photography and concealed weapon detection [55]. Due to their versatility, multimodal algorithms can be used in wider biomedical applications involving genomics, proteomics, metabolomics and other types of omics data. Interestingly, they have been successfully used when the features originate from different domains, and some of them are generated by mechanistic models [67]. In this case, preprocessing of the features, coupled with late or intermediate fusion should be preferred to early fusion. This approach, based on computational systems biology and machine learning, could provide key mechanistic insights into neurological disorders [60].

Medical image segmentation is a challenging task in medical image analysis. Multimodal deep learning has been used in medical imaging, especially for providing multi-information about the target (tumour, organ or tissue). Segmentation using multimodal has been implemented as a fusion of multi-information to improve segmentation [36]. Deep learning provides state-of-the-art performance in image classification, segmentation, object detection and tracking tasks. Recently, deep learning has gained interest in multimodal image segmentation because of its self-learning and simplification ability over a large amount of data [71].

3.4 Databases for AI/ML

Reproducibility, validation and prediction benefit from existing imaging data and related information. There already exist some databases that provide open data for modeling, testing and inferring [28, 41, 54]. Thereupon these web resources, some information is provided about AI or ML approaches. However, both reliable data containing patient data as well as ML performance depend largely on the studied disease and its features. Some example of data from animal models include the Cambridge MRI database (this provides open phenotypic data for animal models of HD [59]), and the Mouse Tumor Biology Database (this provides different kinds of information on tumors in mice [8]). Moreover, the project BRAINS provides anonymised images and related clinical information from healthy subjects across human life span via data request and access agreement [29]. The Open Access Series of Imaging Studies provides MRI data sets of subjects clinically diagnosed with Alzheimer's disease [37]. Functional MRI data from subjects with Huntington's disease can be found within Track-HD study [64]. The Parkinson's Disease Biomarkers Program provides access to brain scans and related information for researchers [49].

Classification of disease subtypes, subjects, brain regions, and gradings are often based on ML approaches via automatically segmenting brain MRI data [23]. Making use of such databases, ML not only helps in (semi-)automatically segmenting images, but it is also a tool for trying to answer several research questions, for example predicting tumor growth [27] or investigating minimal tumor burden and therapy resistance by cancer patients [50, 51]. Some case reports also show AI outperforming human domain experts [18, 31, 56]. Recent advances already try to bridge imaging and genetic studies. Imaging genetic studies combine investigations of genotype and imaging phenotype to better understand brain structure, function and the further cause and effects of a specific disease. Imaging genetics studies improve our understanding of pathways that are related to the cause or effect in cerebral disorders [30]. Genome-wide association studies suggest genetic relationships for structural as well as functional measures among family members [65].

3.5 AI-Aided Disease Classification Using Chemical Imaging

Further use of chemicals, such as metabolic biomarkers, on imaging basis could be included, for example, in therapy monitoring on a cellular basis [46]. AI-aided diagnosis for clinical purposes and computational models for prediction could involve quantitative changes next to functional imaging built upon brain MRI data. In this regard, molecular fMRI techniques exhibit the specificity for neural pathways or signaling components at cellular-level specificity [4]. The method of fMRI has been used to study time-resolved volumetric measurements of dopamine release [34]. Chemical exchange saturation transfer allows for signal-amplification of, for example, deoxyglucose and its phosphorylated metabolite in order to image glucose uptake [45]. This technique of MR imaging has been used to image glucose uptake in head and neck cancer [66]. Proton MRS is commonly used for studies on brain metabolites, including the marker of neuroaxonal integrity N-acetylaspartate, cholin for membrane turnovers, (phospho) creatine for energy metabolism and myo-inositol for astroglial activation [44].

MR techniques for imaging brain metabolism can assist in the studying of brain disorders, in aid of novel MRI contrasts for visualizing neuronal firing across brain regions, pH imaging of glioma and both glutamatergic neurotransmission and cell-specific energetics [26]. The understanding of oxidative metabolism plays a fundamental role in many diseases, which supports demand for the development of non-invasive methods for routine analyses [42].

3.6 Explainable AI

There are several methods which are relevant for further studies and for testing whether and to what extent they can be useful to contribute towards the aforementioned use-cases. Six of the most relevant include: BETA, LRP, LIME as well as GAMs, Bayesian Rule Lists and Hybrid Models, particularly with a human-in-the-loop. BETA (Black Box Explanations through Transparent Approximations) is a model-agnostic framework for explaining the behavior of black-box classifiers by instantaneously optimizing for fidelity to the original model and interpretability of the explanation [32]. There is also a more general solution to the problem of understanding classification decisions by pixel-wise decomposition of nonlinear classifiers which allows visualizing the contributions of single pixels to predictions for kernel-based classifiers over bag of words features and for multilayered neural networks [3].

LRP (Layer-Wise Relevance Propagation) is another general solution for understanding classification decisions via pixel decomposition of nonlinear classifiers, which allows running the “thought processes” backwards [3, 43]. This enables to retrace which input had which influence on the respective result. In individual cases, this lets us understand how a deep learning method has come to a certain medical diagnosis or a risk assessment. LIME (Local Interpretable Model-Agnostic Explanations) developed by Ribeiro, Singh and Guestrin [58] is a model-agnostic system, where $x \in \mathbb{R}^d$ is the original representation of an instance

being explained, and $x' \in \mathbb{R}^{d'}$ is used to denote a vector for its interpretable representation (e.g. x' may be a feature vector containing word embeddings, with x' being the bag of words). The goal is to identify an interpretable model over the interpretable representation that is locally faithful to the classifier, i.e.

$$g : \mathbb{R}^{d'} \rightarrow \mathbb{R}, g \in G,$$

where G is a class of potentially interpretable models, such as linear models, decision trees, or rule lists etc.; given a model $g \in G$, it can be visualized as an explanation to the human expert in $\mathbb{R}^{\mathcal{F}}$. LIME works separately with each instance, they are permuted and a measure of similarity to the original instances is calculated. Consequently, the complex model provides predictions for each of these permuted instances and the influences of the alterations can be understood for each instance. In this way, for example, a medical doctor can check whether and to what extent results can be realistic. All these models cannot explain why a certain decision has been made, which is a goal of current research to find out in the context of the aforementioned use cases.

In medical domains, the explainability and interpretability of algorithms are as critical as their performances [25]. It can be nearly impossible for a doctor or medical professional to effectively integrate their expert knowledge with a model's output unless they can interpret why that model made the decision that it did [22]. Over the past several decades, AI researchers have developed a wide range of techniques for interpretable and explainable classification. These techniques fall into four general categories: sensitivity analysis, linear approximation, rule-based decompositions, and models of causality.

Sensitivity analysis techniques attempt to model which regions of the input space are most important for the classification decision. For a neural network, the simplest sensitivity analysis technique is the "input gradient technique," which involves taking the (smoothed) gradient of the input features with respect to the model loss function [62]. Several sensitivity analysis techniques, including the above-mentioned LRP [43] and LIME [13, 58] use a linear model to (locally) approximate a complex classifier, since linear models can be easily interpreted based on feature weights. LRP directly decomposes the model output on any training sample into the weighted sum of the model features, and LIME builds a linear classifier to approximate model behavior in the region of a particular training sample.

Rule-based algorithms represent a classification problem as a set of rules on the input features. These include algorithms like Decision Trees and Bayesian Rule Lists [68]. The most straightforward way to build these algorithms is to assemble them directly from the training data, but this approach can have extremely high variance and is often insufficient for modern applications. A more modern approach is to use rule-based algorithms to approximate pre-trained classifiers, similarly to how LIME and LRP approximate complex algorithms with linear functions. One example of this is the above-mentioned BETA algorithm, which builds a rule set to approximate a black box classifier [32]. The most direct way to build an explainable classifier is to directly model the causal relationships

between the features and the classification output. The classic way to do this is to use a Bayesian Network [38] to model the conditional independences of features and latent factors, but this approach can be challenging to scale to the size of modern datasets and feature spaces.

4 Open Problems and Future Outlook

This paper entails exploring an intriguing subject. Our understanding of brain disease comes from different sources, but pathology remains to be one of the most important. However, it is a very time-consuming process because it requires manually performed tasks. On the other hand, MRI provides a lot of information on a larger scale, and has already seen major transformations with its use of AI. For cerebral disorders, we discussed in what way can AI generally, and ML specifically, contribute.

We hypothesized that AI can help to find invariably recurring parameters that have escaped human attention (e.g. due to noisy data) to validate diagnosis. In addition, AI helps to deal with an ever increasing amount of data that would take much longer to be analyzed manually. Several ML methods can help to identify and provide more meaningful information regarding the signals of different contrasts, location (with high resolution 7 T MRI), texture, size, dimension, patient information and specific patterns. Moreover, AI could be used to locate, correlate and compare all brain regions, in order to study the high individual variability of human gyri and sulci, signal variability, normal ageing process, clinical records and neurodegenerative diseases.

AI has the potential to go beyond helping filter out noise. One reason humans may be limited is not only due to the noise, but also due to wrong decisions on the feature space wrong. For example, in structural MRI, we automatically make the feature space a voxel. However, that unit results from the measurement technique, rather than any hypothesis or regularity about the brain or disease. There is great potential for AI to reveal other meaningful feature spaces, such as volume, heterogeneity, variability. It is possible that a meaningful feature space involves voxel-to-voxel relationships, or interdependencies, for example. Therefore, AI can help us elucidate the right feature spaces, eliminating or reducing human bias.

References

1. Amoroso, N., et al.: Brain structural connectivity atrophy in Alzheimer's disease. arXiv (2017)
2. Arel, I., Rose, D.C., Karnowski, T.P.: Deep machine learning - a new frontier in artificial intelligence research - research frontier. *IEEE Comput. Intell. Mag.* **5**, 4 (2010). <https://doi.org/10.1109/MCI.2010.938364>
3. Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.R., Samek, W.: On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PloS one* **10**(7), e0130140 (2015)

4. Bartelle, B.B., Barandov, A., Jasanoff, A.: Molecular fMRI. *J. Neurosci.* **36**(15), 4139–4148 (2016)
5. Baxt, W.G.: Use of an artificial neural network for the diagnosis of myocardial infarction. *Ann. Intern. Med.* **115**, 11 (1991)
6. Bishop, C.M.: *Pattern Recognition and Machine Learning*. Springer, Heidelberg (2006)
7. Breiman, L.: Random forests. *Mach. Learn.* **45**, 1 (2001). <https://doi.org/10.1023/A:1010933404324>
8. Bult, C.J.: Mouse tumor biology (MTB): a database of mouse models for human cancer. *Nucleic Acids Res.* **43**, D818–D824 (2014). <https://doi.org/10.1093/nar/gku987>
9. Calhoun, V.D., Sui, J.: Multimodal fusion of brain imaging data: a key to finding the missing link(s) in complex mental illness. *Biol. Psychiatry Cogn. Neurosci. Neuroimaging* **1**(3), 230–244 (2016)
10. Daffner, K.: Current approaches to the clinical diagnosis of Alzheimer’s disease. In: Scinto, L.F.M., Daffner, K.R. (eds.) *Early Diagnosis of Alzheimer’s Disease. Current Clinical Neurology*, pp. 29–64. Humana Press, Totowa (2000). https://doi.org/10.1007/978-1-59259-005-6_2
11. Furey, T.S., Cristianini, N., Duffy, N., Bednarski, D.W., Schummer, M., Haussler, D.: Support vector machine classification and validation of cancer tissue samples using microarray expression data. *Bioinformatics* **16**, 10 (2000)
12. Gardiner, S.L., van Belzen, M.J., Boogaard, M.W., et al.: Huntingtin gene repeat size variations affect risk of lifetime depression. *Transl. Psychiatry* **7**, 1277 (2017). <https://doi.org/10.1038/s41398-017-0042-1>
13. Gilpin, L.H., Bau, D., Yuan, B.Z., Bajwa, A., Specter, M., Kagal, L.: Explaining explanations: an overview of interpretability of machine learning. *arXiv preprint arXiv:1806.00069* (2018)
14. Goodfellow, I., Bengio, Y., Courville, A.: *Deep Learning*. MIT Press, Cambridge (2016)
15. Greenspan, H., Van Ginneken, B., Summers, R.M.: Guest editorial deep learning in medical imaging: overview and future promise of an exciting new technique. *IEEE Trans. Med. Imaging* **35**, 5 (2016)
16. Gupta, A., Ayhan, M., Maida, A.: In natural image bases to represent neuroimaging data. In: *International Conference on Machine Learning*, pp. 987–994 (2013)
17. Hadzi, T.C., et al.: Assessment of cortical and striatal involvement in 523 Huntington disease brains. *Neurology* **79**, 1708–1715 (2012)
18. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: surpassing human-level performance on ImageNet classification. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1026–1034 (2015)
19. Heim, B., et al.: Magnetic resonance imaging for the diagnosis of Parkinson’s disease. *J. Neural Transm.* **124**, 8 (2017)
20. Holzinger, A.: Introduction to machine learning and knowledge extraction (make). *Mach. Learn. Knowl. Extr.* **1**(1), 1–20 (2017)
21. Holzinger, A., et al.: Machine learning and knowledge extraction in digital pathology needs an integrative approach. In: Holzinger, A., Goebel, R., Ferri, M., Palade, V. (eds.) *Towards Integrative Machine Learning and Knowledge Extraction. LNCS (LNAI)*, vol. 10344, pp. 13–50. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-69775-8_2
22. Holzinger, A., Carrington, A., Müller, H.: Measuring the quality of explanations: The system causability scale (scs). Comparing human and machine explanations.

- KI - Künstliche Intelligenz (German J. Artif. Intell.). Special Issue on Interactive Machine Learning, Edited by Kristian Kersting, TU Darmstadt **34**(2) (2020). <https://doi.org/10.1007/s13218-020-00636-z>
23. Holzinger, A., Goebel, R., Palade, V., Ferri, M.: Towards integrative machine learning and knowledge extraction. In: Holzinger, A., Goebel, R., Ferri, M., Palade, V. (eds.) *Towards Integrative Machine Learning and Knowledge Extraction*. LNCS (LNAI), vol. 10344, pp. 1–12. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-69775-8_1
 24. Holzinger, A., Kieseberg, P., Weippl, E., Tjoa, A.M.: Current advances, trends and challenges of machine learning and knowledge extraction: from machine learning to explainable AI. In: Holzinger, A., Kieseberg, P., Tjoa, A.M., Weippl, E. (eds.) *CD-MAKE 2018*. LNCS, vol. 11015, pp. 1–8. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-99740-7_1
 25. Holzinger, A., Langs, G., Denk, H., Zatloukal, K., Mueller, H.: Causability and explainability of artificial intelligence in medicine. *Wiley Interdisc. Rev. Data Min. Knowl. Disc.* **9**(4), e1312 (2019). <https://doi.org/10.1002/widm.1312>
 26. Hyder, F., Rothman, D.: Advances in imaging brain metabolism. *Ann. Rev. Biomed. Eng.* **19**, 485–515 (2017)
 27. Jeanquartier, F., et al.: Machine learning for *in silico* modeling of tumor growth. In: Holzinger, A. (ed.) *Machine Learning for Health Informatics*. LNCS (LNAI), vol. 9605, pp. 415–434. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-50478-0_21
 28. Jeanquartier, F., Jean-Quartier, C., Schreck, T., Cemernek, D., Holzinger, A.: Integrating open data on cancer in support to tumor growth analysis. In: Renda, M.E., Bursa, M., Holzinger, A., Khuri, S. (eds.) *ITBAM 2016*. LNCS, vol. 9832, pp. 49–66. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-43949-5_4
 29. Job, D.E., et al.: A brain imaging repository of normal structural MRI across the life course: brain images of normal subjects (brains). *NeuroImage* **144**, 299–304 (2017)
 30. Klein, M., et al.: Brain imaging genetics in ADHD and beyond- mapping pathways from gene to disorder at different levels of complexity. *Neurosci. Biobehav. Rev.* **80**, 115–155 (2017). <https://doi.org/10.1016/j.neubiorev.2017.01.013>
 31. Klöppel, S., Abdulkadir, A., Jack, C.R., Koutsouleris, N., Mourão-Miranda, J., Vemuri, P.: Diagnostic neuroimaging across diseases. *Neuroimage* **61**(2), 457–463 (2012)
 32. Lakkaraju, H., Kamar, E., Caruana, R., Leskovec, J.: Interpretable and explorable approximations of black box models. *arXiv preprint arXiv:1707.01154* (2017)
 33. Lecun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**, 7553 (2015). <https://doi.org/10.1038/nature14539>
 34. Lee, T., Cai, L.X., Lelyveld, V.S., Hai, A., Jasanoff, A.: Molecular-level functional magnetic resonance imaging of dopaminergic signaling. *Science* **344**(6183), 533–535 (2014)
 35. Lemm, S., Blankertz, B., Dickhaus, T., Mueller, K.R.: Introduction to machine learning for brain imaging. *Neuroimage* **10**(1016), 387–399 (2011)
 36. Li, C., Sun, H., Liu, Z., Wang, M., Zheng, H., Wang, S.: Learning cross-modal deep representations for multi-modal MR image segmentation. In: Shen, D., et al. (eds.) *MICCAI 2019*. LNCS, vol. 11765, pp. 57–65. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32245-8_7

37. Marcus, D.S., Wang, T.H., Parker, J., Csernansky, J.G., Morris, J.C., Buckner, R.L.: Open access series of imaging studies (OASIS): cross-sectional mri data in young, middle aged, nondemented, and demented older adults. *J. Cogn. Neurosci.* **19**(9), 1498–1507 (2007)
38. Margaritis, D.: Learning Bayesian network model structure from data. Ph.D. thesis, Carnegie-Mellon University Pittsburgh PA School Of Computer Science (2003)
39. Marques, J.P., et al.: Studying cyto and myeloarchitecture of the human cortex at ultra-high field with quantitative imaging: R1, R2(*) and magnetic susceptibility. *Neuroimage* **147**, 152 (2017)
40. Martino, D., et al.: The differential diagnosis of Huntington’s disease-like syndromes: ‘red flags’ for the clinician. *J. Neurol. Neurosurg. Psychiatry* **84**, 650–656 (2013)
41. Di Martino, A., et al.: The autism brain imaging data exchange: towards a large-scale evaluation of the intrinsic brain architecture in autism. *Mol. Psychiatry* **19**(6), 659–667 (2014)
42. Mellon, E.A., Beesam, R.S., Elliott, M.A., Reddy, R.: Mapping of cerebral oxidative metabolism with MRI. *Proc. Nat. Acad. Sci.* **107**(26), 11787–11792 (2010)
43. Montavon, G., Binder, A., Lapuschkin, S., Samek, W., Müller, K.-R.: Layer-wise relevance propagation: an overview. In: Samek, W., Montavon, G., Vedaldi, A., Hansen, L.K., Müller, K.-R. (eds.) *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*. LNCS (LNAI), vol. 11700, pp. 193–209. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-28954-6_10
44. Mormina, E., et al.: Cerebellum and neurodegenerative diseases: beyond conventional magnetic resonance imaging. *World J. Radiol.* **9**(10), 371–388 (2017). <https://doi.org/10.4329/wjr.v9.i10.371>
45. Nasrallah, F.A.: Imaging brain deoxyglucose uptake and metabolism by gluco-CEST MRI. *J. Cereb. Blood Flow Metab.* **33**(8), 1270–1278 (2013)
46. Ngen, E.J., Artemov, D.: Advances in monitoring cell-based therapies with magnetic resonance imaging: future perspectives. *Int. J. Mol. Sci.* **18**, 1 (2017)
47. Ngiam, J., Khosla, A., Kim, M., Nam, J., Lee, H., Ng, A.Y.: Multimodal deep learning. In: *Proceedings of the 28th International Conference on Machine Learning (ICML-2011)*, pp. 689–696 (2011)
48. Nopoulos, P.C., et al.: Cerebral cortex structure in prodromal Huntington disease. *Neurobiol. Dis.* **40**, 544–554 (2010). <https://doi.org/10.4322/acr.2018.003>
49. Ofori, E., Du, G., Babcock, D., Huang, X., Vaillancourt, D.E.: Parkinson’s disease biomarkers program brain imaging repository. *Neuroimage* **124**, 1120–1124 (2016). <https://doi.org/10.1186/s40708-019-0096-3>
50. O’Sullivan, S., Holzinger, A., Zatloukal, K., Saldiva, P., Sajid, M.I., Wichmann, D.: Machine learning enhanced virtual autopsy. *Autops. Case Rep.* **7**(4), 3–7 (2017). <https://doi.org/10.4322/acr.2017.037>
51. O’Sullivan, S., Holzinger, A., Wichmann, D., Saldiva, P., Sajid, M., Zatloukal, K.: Virtual autopsy: machine learning and artificial intelligence provide new opportunities for investigating minimal tumor burden and therapy resistance by cancer patients. *Autops. Case Rep.* **8**, 1 (2018). <https://doi.org/10.4322/acr.2018.003>
52. O’Sullivan, S., et al.: The role of artificial intelligence and machine learning in harmonization of high-resolution post-mortem MRI (virtopsy) with respect to brain microstructure. *Brain Inform.* **6**(1), 3 (2019). <https://doi.org/10.1186/s40708-019-0096-3>
53. Pathak, D., Krahenbuhl, P., Darrell, T.: Constrained convolutional neural networks for weakly supervised segmentation. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1796–1804 (2015)

54. Poldrack, R.A., Gorgolewski, K.J.: Making big data open: data sharing in neuroimaging. *Nature Neurosci.* **17**(11), 1510–1517 (2014)
55. Rajalingam, B., Priya, R.: Multimodal medical image fusion based on deep learning neural network for clinical treatment analysis. *Int. J. Chem. Tech. Res. CODEN (USA) IJCRGG* **11**, 0974–4290 (2018). ISSN
56. Rajpurkar, P., et al.: Chexnet: radiologist-level pneumonia detection on chest x-rays with deep learning. arXiv preprint [arXiv:1711.05225](https://arxiv.org/abs/1711.05225) (2017)
57. Ramachandram, D., Taylor, G.W.: Deep multimodal learning: a survey on recent advances and trends. *IEEE Signal Process. Mag.* **34**(6), 96–108 (2017)
58. Ribeiro, M.T., Singh, S., Guestrin, C.: Why should i trust you?: Explaining the predictions of any classifier. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1135–1144. ACM (2016)
59. Sawiak, S.J., Morton, A.J.: The cambridge MRI database for animal models of Huntington disease. *NeuroImage* **124**, 1260–1262 (2016)
60. Sertbas, M., Ulgen, K.O.: Unlocking human brain metabolism by genome-scale and multiomics metabolic models: relevance for neurology research, health, and disease. *OMICS* **22**(7), 455–467 (2018)
61. Sirinukunwattana, K., Raza, S.E.A., Tsang, Y.W., Snead, D.R., Cree, I.A., Rajpoot, N.M.: Locality sensitive deep learning for detection and classification of nuclei in routine colon cancer histology images. *IEEE Trans. Med. Imaging* **35**, 5 (2016)
62. Smilkov, D., Thorat, N., Kim, B., Viégas, F., Wattenberg, M.: SmoothGrad: removing noise by adding noise. arXiv preprint [arXiv:1706.03825](https://arxiv.org/abs/1706.03825) (2017)
63. Szolovits, P., Patil, R.S., Schwartz, W.B.: Artificial intelligence in medical diagnosis. *Ann. Intern. Med.* **108**, 1 (1988)
64. Tabrizi, S.J., et al.: Biological and clinical manifestations of Huntington's disease in the longitudinal TRACK-HD study: cross-sectional analysis of baseline data. *Lancet Neurol.* **8**(9), 791–801 (2009). [https://doi.org/10.1016/S1474-4422\(09\)70170-X](https://doi.org/10.1016/S1474-4422(09)70170-X)
65. Thompson, P.M., et al.: The enigma consortium: large-scale collaborative analyses of neuroimaging and genetic data. *Brain Imaging Behav.* **8**(2), 153–182 (2014)
66. Wang, J., et al.: Magnetic resonance imaging of glucose uptake and metabolism in patients with head and neck cancer. *Sci. Rep.* **6**, 30618 (2016)
67. Yaneske, E., Angione, C.: The poly-omics of ageing through individual-based metabolic modelling. *BMC Bioinform.* **19**(14), 415 (2018)
68. Yang, H., Rudin, C., Seltzer, M.: Scalable Bayesian rule lists. arXiv preprint [arXiv:1602.08610](https://arxiv.org/abs/1602.08610) (2016)
69. Zampieri, G., Vijayakumar, S., Yaneske, E., Angione, C.: Machine and deep learning meet genome-scale metabolic modeling. *PLoS Comput. Biol.* **15**(7), e1007084 (2019)
70. Zhang, X., et al.: PET/MR imaging: new frontier in Alzheimer's disease and other dementias. *Front. Mol. Neurosci.* **10**, 343 (2017)
71. Zhou, T., Ruan, S., Canu, S.: A review: deep learning for medical image segmentation using multi-modality fusion. *Array* **3–4**, 100004 (2019)