

CAPSULE NETWORKS FOR BRAIN TUMOR CLASSIFICATION BASED ON MRI IMAGES AND COARSE TUMOR BOUNDARIES

Parnian Afshar[†], Konstantinos N. Plataniotis[‡], and Arash Mohammadi[†]

[†]Concordia Institute for Information Systems Engineering, Concordia University, Montreal, QC, Canada

[‡]Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON, Canada

Emails: {p_afs, arashmoh}@encs.concordia.ca; kostas@ece.utoronto.ca

ABSTRACT

According to official statistics, cancer is considered as the second leading cause of human fatalities. Among different types of cancer, brain tumor is seen as one of the deadliest forms due to its aggressive nature, heterogeneous characteristics, and low relative survival rate. Determining the type of brain tumor has significant impact on the treatment choice and patient's survival. Human-centered diagnosis is typically error-prone and unreliable resulting in a recent surge of interest to automatize this process using convolutional neural networks (CNNs). CNNs, however, fail to fully utilize spatial relations, which is particularly harmful for tumor classification, as the relation between the tumor and its surrounding tissue is a critical indicator of the tumor's type. In our recent work, we have incorporated newly developed CapsNets to overcome this shortcoming. CapsNets are, however, highly sensitive to the miscellaneous image background. The paper addresses this gap. The main contribution is to equip CapsNet with access to the tumor surrounding tissues, without distracting it from the main target. A modified CapsNet architecture is, therefore, proposed for brain tumor classification, which takes the tumor coarse boundaries as extra inputs within its pipeline to increase the CapsNet's focus. The proposed approach noticeably outperforms its counterparts.

Index Terms: Brain Tumor Classification, Capsule Networks, Tumor Boundary, Convolutional Neural Networks.

1. INTRODUCTION

According to world health organization's statistics, cancer is considered as the second leading cause of human fatalities across the world, being responsible for an estimated 9.6 million deaths in this year. Among different type of cancers, brain tumor is widely seen [1] as one of the deadliest cancers due to its aggressive nature, heterogeneous characteristics (types), and low relative survival rate (e.g., in US relative survival rate following a diagnosis of a primary malignant brain tumor is around 35%). This cancer can drastically influence the quality of life, for both patients and their families. The key factor in treating brain cancer and increasing its survivability rate is early diagnosis and correctly determining its type. Brain tumor can have different types (e.g., Meningioma, Pituitary, and Glioma [2]) depending on several factors such as the shape, texture, and location of the tumor. Determining the correct type of brain tumor is of paramount importance, as it can significantly influence the choice of treatment and predicting patient's survival.

Medical screening is considered as one of the most common and accurate techniques for cancer type classification [3], and by being non-invasive, is drawing more and more attention. Among different screening technologies, Magnetic Resonance Imaging (MRI) is, typically, selected as the utilized technique for brain tumor classification, due to the high resolution images it can provide on brain tissue. However, cancer type recognition based on MRI images is a challenging, error-prone, and time-consuming procedure, as it highly depends on the experience of the radiologist, and more importantly, there may not be enough visible landmarks in the image to contribute to an accurate decision. This necessitates an urgent quest to develop and design new and innovative brain tumor classification techniques, which is the focus of the paper.

Prior Work: Considering the aforementioned problems with a human-centered cancer diagnosis, there has been a recent surge of interest [4–7] in development of autonomous processing systems for brain cancer diagnosis. The conventional workflow, for developing an automatic or semi-automatic system, is to first segment the tumor from the MRI images, with the aim of extracting quantitative features, referred to as “Radiomics” [8]. The extracted Radiomics, typically, contain a wide variety of feature categories [9] including but not limited to shape (quantifying the tumor geometric pattern), intensity (derived from the tumor region histogram), and texture features (concerning the relations between pixels to capture intra-tumor heterogeneity) to name a few. Radiomics are then utilized to train a predictive/survival model for cancer classification. For instance, Aerts *et al.* [10] have extracted 400 features from the segmented tumor to study the relations between the image-based features and clinical outcomes. Consequently, the association between these features and patients' survival were analyzed using different statistical models. Nevertheless, it is concluded that a strong correlation exists between the tumor annotation and the extracted features, i.e., Radiomics features are highly sensitive to inter-observer variability in segmenting the tumor. In other words, the hand-deigned features are not stable [11], which highly reduces the model reliability and applicability. More importantly this pipeline needs a prior knowledge on what types of features to extract, which is not always available.

The shortcomings of the conventional Radiomics workflow have resulted in a trend towards the use of deep learning, in particular Convolutional Neural Networks (CNNs) [12], for cancer diagnosis and classification. For instance, Li *et al.* [13] have used a 6 layer CNN to extract features from the brain images, with the ultimate goal of classifying brain tumors. CNNs do not need any prior knowledge on the type of features, and can be trained in an end-to-end manner without necessarily requiring the segmented tumor. Although these networks have extensive learning capacity, they suffer from some key drawbacks [14, 15], e.g., being incapable of considering

This work was partially supported by the Fonds de Recherche du Québec Nature et Technologies (FRQNT) Grant 2018-NC-206591.

the spatial relation between objects in the image, which in turn results in lack of robustness to rotation and affine transformation. Besides, huge amount of data is needed to improve the robustness of CNNs, which is not always available especially for brain tumor classification problem. To overcome the aforementioned shortcomings of the CNNs, Capsule Networks (CapsNets) [14, 15] are recently proposed, and are armed with a technique that enhances their robustness to transformations. Each capsule is a group of neurons, which can represent different instantiation parameters associated with different objects, as well as the probability of their existence. Since their introduction, there has been a great surge of interest in using CapsNets in different application domains [16–18], and also on development of different variations of CapsNets [19–23].

Contributions: An important property of CapsNets, which has made them potentially better models for handling transformations, is their “Routing by Agreement” process, during which capsules in lower levels predict the outcome of their parent capsules. Consequently, parent capsules are activated only if the predictions agree. In our recent work [16], we have shown that CapsNets can outperform CNNs for the task of brain tumor classification. However, CapsNets are highly sensitive to image background, and as such provide higher accuracy for classifying segmented tumors (i.e., scenarios where the input to the CapsNet is the segmented tumor region), compared to the scenarios where the whole brain image is provided as the input. Nevertheless, needing the segmented tumor has two major problems: (i) First, segmenting the tumor is a time-consuming task and can only be provided by experts, and; (ii) Second, the tumor surrounding tissue contains valuable information, which is not accessible, when the network is fed with only the segmented region.

The paper addresses the aforementioned issue, in particular, *giving a CapsNet the access to the tumor surrounding tissues, without distracting it from the main target, and not requiring the tumor detailed annotation are the main motivations/contributions of this work*. More specifically, to help the CapsNet to focus on the main region and, at the same time use the information from the surrounding tissues, we provide the network with the tumor coarse boundaries. This information is fed to the CapsNet at the last layer, before going through a set of fully connected layers and the final Softmax layer that makes the decision. Our results indicate that the proposed approach can outperform a CapsNet that is only fed with a brain or a segmented tumor image. Furthermore, since our proposed approach does not need any detailed annotations, it is more time-efficient and can take the burden of manual delineation off the experts/radiologists.

The rest of this paper is organized as follows: Section 2 describes required mathematical background for CNNs and Capsule networks. In section 3, we present our proposed approach followed by experimental results in section 4. Finally, Section 5 concludes the paper.

2. PROBLEM FORMULATION

In this section, the CNNs are briefly discussed, followed by intuitively describing the reasons that these networks may fail in the absence of enough training datasets. CapsNets that are proposed to solve this problem will be explained afterwards.

2.1. Convolutional Neural Networks

CNNs [24], which are basically the stack of convolutional, pooling and, sometimes, fully connected layers, benefit from the fact that weights are shared over the entire input, significantly reducing the computational cost, and allowing the network to extract elementary

and higher order local features. The fact that these networks do not need any prior knowledge on the types of the features to extract, has made them popular architectures in medical image processing [25]. Generally speaking, in a CNN with N layers, the output $\mathbf{Y}^{(l-1)}$ of layer $l-1$, for $(2 \leq l \leq N)$, is the input to the layer l resulting in the associated output \mathbf{Y}^l given by

$$\mathbf{X}_{i,j}^{(l)} = \sum_{a=0}^{M-1} \sum_{b=0}^{M-1} \mathbf{W}_{ab} \mathbf{Y}_{i+a,j+b}^{(l-1)} \quad (1)$$

$$\mathbf{Y}_{i,j}^{(l)} = \sigma(\mathbf{X}_{i,j}^{(l)}), \quad (2)$$

where $\mathbf{X}^{(l)}$ is the pre-activation output, M is the size of kernels, \mathbf{W} is the kernel matrix containing the CNN weights to be learned during the back propagation, and $\sigma(\cdot)$ denotes the activation function.

Sub-sampling or pooling layers, in a CNN, are incorporated to not only reduce the number of parameters, but also to make the network translation invariant. However, these layers, lose the information about the exact location of the feature detectors, which makes them unable to recognize objects when they are subject to rotation or some other kinds of transformations. This issue remains unsolved unless all possible situations are included in the training data, which is not possible in practice. CapsNets, presented next, are introduced to help with solving this problem.

2.2. Capsule Networks (CapsNets)

Each capsule in a CapsNet is responsible for capturing the probability of a specific object being present, and consists of several neurons that present different instantiation parameters, such as rotation and size, associated with the underlying object. In other words, a capsule is a vector of several features, and the length of the vector serves as the probability of the existence of the object that the capsule is representing. To make the length of the vector smaller than one, a squashing function is normally applied.

To solve the problems associated with the pooling layers in the CNNs, these layers are replaced with a procedure called “Routing by Agreement”, in which, instead of sub-sampling the feature maps, negligently, the contribution of capsules depend on how well they can predict the output of their consequent capsules. In other words, each capsule in a lower level, tries to predict the output of the parent capsules, and the parent capsules take their lower ones into consideration, only if they have been able to provide correct predictions. More specifically, let’s define u_i as the output of the lower level capsule i , $\hat{u}_{j|i}$ as its prediction for a higher level capsule j , and W_{ij} as the weights connecting them, which have to be learned through the back propagation. Armed with this notation, $\hat{u}_{j|i}$ can be calculated as

$$\hat{u}_{j|i} = W_{ij} u_i. \quad (3)$$

The strength of the connection between capsules and their parents depends on how much they agree on the actual output of the parent. In other words, this agreement between $\hat{u}_{j|i}$ and the actual output of the parent capsule j , denoted by s_j , determines the coupling coefficient c_{ij} . Then capsule i sends its output to capsule j as follows

$$s_j = \sum_i c_{ij} \hat{u}_{j|i}. \quad (4)$$

The log probability of whether capsule i should be coupled with capsule j is denoted by b_{ij} , which has to be learned in the “Routing by Agreement” process, and it is set to 0 at the initialization step. During the “Routing by Agreement” process, this probability is updated




Fig. 1. Defining the tumor boundary box.

based on the similarity between s_j and $\hat{u}_{j|i}$. One basic approach to compute this similarity is to take the inner product of the two underlying vectors. In other words, the agreement a_{ij} , which will be added to b_{ij} in each step, is calculated as follows

$$a_{ij} = s_j \cdot \hat{u}_{j|i}. \quad (5)$$

Coupling coefficients are set via the following *softmax* function

$$c_{ij} = \frac{\exp(b_{ij})}{\sum_{k=1}^K \exp(b_{ik})}, \quad (6)$$

where K is the number of capsules in the output layer. Each capsule j , for $(1 \leq j \leq K)$, in the last layer (classification part) is associated with a loss function l_j , which is designed to put high losses on capsules that have come up with large (in the sense of their norm values) instantiation vectors while their associated objects do not actually exist. The loss function l_j is computed as

$$l_j = T_j \max(0, m^+ - \|s_j\|)^2 + \lambda(1 - T_j) \max(0, \|s_j\| - m^-)^2. \quad (7)$$

Term T_j is 1 whenever the class j is present, and is 0 otherwise. Term m^+ , m^- and λ are hyper parameters to be indicated before the learning process. The total loss is the sum over the losses of all output capsules. The original Capsule network has also a set of fully connected layers, referred to as the decoder part, that takes the final instantiation parameters of the true classes as inputs, and try to reconstruct the original image, with the aim of forcing the network to capture real representative features. The decoder loss is defined as a simple squared error and contributes to the final error with a smaller weight, compared to the loss of the capsules. This is done to avoid distracting the network from its main target, which is classifying the objects. This completes a brief introduction to CNNs and CapsNets. Next, we present the proposed framework for tumor classification.

3. THE PROPOSED FRAMEWORK

As stated previously, the goal of this work is to classify brain tumors into three categories of Meningioma, Pituitary, and Glioma, using the MRI images, which are the most widely used modalities for detecting brain-related diseases. Type of the tumor depends on several properties of the tumor itself, as well as its surrounding tissue. For example Meningiomas, typically, originate in the tissue between the skull and the brain, and Gliomas are normally located in the substance of the brain. Automatic brain tumor classification is commonly based on the CNNs, which are powerful tools for extracting high-level and low-level features, without any prior knowledge to be provided by experts. However, as stated previously CNNs have a major drawback limiting their applicability in real-world problems, i.e., they fail to fully consider the exact spatial relations between objects, which is caused by the information loss in the pooling layers. On the other hand, the pooling layers can not be removed, as

without these layers, the network will be highly sensitive to slight translations of the image. The spatial information, which is lost in the CNNs, is of high importance in the problem of brain tumor classification, because the location of the tumor and its relation with the surrounding tissues can highly influence the type of the tumor.

The newly proposed CapsNets introduced in Section 2.2 have the potential to preserve the spatial relations, due to their Routing by Agreement process, and are, therefore, more suitable models for brain tumor classification. In our previous work [16], we have shown that CapsNets overcome CNNs in this problem. However, these networks are sensitive to the image background and try to account for everything in the image. Considering the detailed brain MRI images, this property can negatively affect the network performance. As such, based on our previous results, CapsNet has a higher accuracy, when fed with the segmented tumors, instead of the whole brain image. However, the tumor surrounding tissues contain valuable information, and should not be ignored, when determining the type of the tumor. Furthermore, annotating the brain images is time-consuming and not always feasible.

Motivated by the aforementioned issues, in this paper we have designed a CapsNet architecture that takes the brain images as inputs, however, it is also provided with the tumor coarse boundary, to make it pay more attention to the main target and not to get distracted by every single detail. The vector containing the tumor boundary, shown in the Fig. 1, is concatenated with the output of the capsule layer, and goes through a set of fully connected layers, in order to make the final decision, which is the type of the tumor. The detail of the proposed architecture, shown in Fig. 2, is as follows:

- The inputs to the network are brain MRI images which are downsampled to 128×128 from 512×512 .
- Second layer is a convolutional layer, with a total of 64 feature maps. The size of the filters is 9×9 with stride one.
- The third layer is a capsule layer resulted from 9×9 convolutions. This layer contains 32 capsules of dimension 8.
- The last capsule layer, which contains one capsule for each brain tumor type, determines the most probable class, along with its instantiation parameters. Outputs from this layer are masked based on the detected class, i.e., all capsules, but the winner, are set to 0.
- The tumor boundary box is concatenated with the obtained masked vector and goes through two fully connected layers, with 512 and 1024 neurons, respectively.
- The last layer is a Softmax layer that outputs the probability of each class being present.

Loss Function: The loss for the output of the capsule layers, as defined in Eq. (7), should be added to the Softmax layer loss, which we have defined as a cross entropy loss, with a smaller weight γ , not to dominate the final loss. As such, we have defined the final loss as

$$\text{Loss} = \underbrace{\gamma \times \sum_{j=1}^K l_j}_{\text{Capsule Loss}} - \underbrace{\sum_{j=1}^K y_j \log(p(y_j))}_{\text{Cross Entropy Loss}}, \quad (8)$$

where y_j is a binary variable indicating whether class j is present or not; Term $p(y_j)$ is the probability of this class being present, which is determined by the network, and; K is the number of output classes (types of the tumor). This loss is back propagated through the whole network, including both capsule and fully connected layers. This completes description of the proposed CapsNet architecture for brain tumor classification problem. Next, we present our experimental results to evaluate the effectiveness of the proposed architecture.




Fig. 2. Proposed CapsNet architecture for brain tumor classification. The proposed architecture takes the tumor coarse boundary into consideration, before making the final decision.

Table 1. Training hyper-parameters used for brain tumor classification via Adam [28] optimizer.

Hyper-parameter	Optimized Value
Optimizer	Adam [28]
Number of Epochs	50
Batch size	16
Routing iteration	3
Learning rate	0.01
Learning rate decay	0.9
γ (in Eq. (8))	0.1
λ (in Eq. (7))	0.5
m^+ (in Eq. (7))	0.9
m^- (in Eq. (7))	0.1

4. EXPERIMENTAL SETUP

To test our proposed CapsNet architecture (shown in Fig. 2), we have used a brain MRI dataset [26], consisting of 3064 images from 233 patients, diagnosed with one of the three brain tumor types. Table 1 shows the values of hyper-parameters used to train the proposed CapsNet architecture. Our model is implemented on Python 2.7, using the Keras library [29]. As shown in Table 2, the proposed CapsNet architecture is compared with different alternative scenarios where the network is fed with either the brain or the segmented tumor image. Furthermore, we have evaluated our method based on a CNN that is proposed in Reference [27] for the same problem of brain tumor type classification. The architecture of the CNN proposed in [27] is as follows:

- The two first layers are convolutional layers with 5×5 filters, outputting 64 feature maps. Each of these two layers are followed by 2×2 pooling layer.
- The third and forth layers are fully connected ones, with 800 neurons.
- The last layer is a softmax one, to classify the brain tumors.

In Table 2, we have also included the result for a modified CNN adapted based on the proposed architecture introduced in Section 3. The modified CNN takes as input both brain images and bounding

Table 2. Comparison between the proposed approach and previous results. The bold number corresponds to the proposed approach, which outperforms its counterparts.

	Approach	Accuracy
1.	CapsNet given brain image as input [16].	78%
2.	CapsNet given segmented tumor as input [16].	86.56%
3.	Proposed CapsNet Architecture (Fig. 2).	90.89%
4.	CNN given brain image as input [27].	61.97%
5.	CNN given segmented tumor as input [27].	72.13%
6.	Modified CNN with brain image and tumor boundary box as inputs (Section 3).	88.33%

boxes, where the box coordinates are concatenated with the last fully connected layer of the CNN (the other components remain the same as described above). As it can be inferred from Table 2, the CapsNet architecture introduced in this paper outperforms CNN in all situations, and achieves the best performance when it is fed with brain images, and coarse tumor boundaries.

5. CONCLUSION

In this work, we have presented a CapsNet architecture that incorporates both the raw MRI brain images and the tumor coarse boundaries in order to classify the tumors. The proposed CapsNet architecture has two main advantageous: (i) First, the need for tumor exact annotation is eliminated, and; (ii) Second, it helps the CapsNet to focus on the main area, and at the same time, consider its relation with surrounding tissues. Our results show that the proposed approach is capable of increasing the classification accuracy, compared to the previous CapsNets and CNNs. Finally, it is worth mentioning that CapsNets are armed with properties that increase their interpretability, e.g., the output instantiation parameters of the true class can explain whether or not the network has captured correct features. As explainability is of high importance in medical decision making, in the future, we will investigate interpretability of CapsNets for the brain tumor classification.

6. REFERENCES

- [1] Rebecca L. Siegel, Kimberly D. Miller, *et al.*, “Cancer Statistics, 2016,” *A Cancer Journal for Clinicians*, 2016.
- [2] Jun Cheng, Wei Huang, *et al.*, “Enhanced Performance of Brain Tumor Classification via Tumor Region Augmentation and Partition,” *PloS one.*, 2015.
- [3] “Tumor Types: Understanding Brain Tumors,” *National Brain Tumor Society*, 2018.
- [4] Khalid Usman, Kashif Rajpoot, “Brain tumor classification from multi-modality MRI using wavelets and machine learning,” *Pattern Analysis and Applications*, vol. 20, no. 3, pp. 871-881, 2017.
- [5] Nidahl K. El Abbadi, Neamah E. Kadhim, “Brain Cancer classification Based on Features and Artificial Neural Network,” *International Journal of Advanced Research in Computer and Communication Engineering*, vol. 8, no. 1, Jan. 2017.
- [6] H. Mohsen, El-Sayed A.El-Dahshan, *et al.*, “Classification using deep learning neural networks for brain tumors,” *Future Computing and Informatics Journal*, vol. 3, no. 1, pp. 68-71, 2018.
- [7] Reema Mathew A, Babu Anto P, *et al.*, “Tumor detection and classification of MRI brain image using wavelet transform and SVM,” *2017 International Conference on Signal Processing and Communication (ICSPC)*, 2017.
- [8] P. Afshar, A. Mohammadi, K.N. Plataniotis, A. Oikonomou, H. Benali “From Hand-Crafted to Deep Learning-based Cancer Radiomics: Challenges and Opportunities,” Submitted to *IEEE Signal Processing Magazine*, 2018.
- [9] A. Oikonomou, F. Khalvati, *et al.*, “Radiomics Analysis at PET/CT Contributes to Prognosis of Recurrence and Survival in Lung Cancer Treated with Stereotactic Body Radiotherapy,” *Scientific Reports*, vol. 8, no. 1, 2018.
- [10] H.J. Aerts, E.R. Velazquez, *et al.*, “Decoding Tumour Phenotype by Noninvasive Imaging Using a Quantitative Radiomics Approach,” *Nature Communications*, vol. 5, 2014.
- [11] J. Griethuysen, A. Fedorov, *et al.* “Computational Radiomics System to Decode the Radiographic Phenotype,” *Cancer Research*, vol. 77, no. 21, pp. 104-107, 2017.
- [12] Alex Krizhevsky, Ilya Sutskever, Ilya Sutskever, “ImageNet Classification with Deep Convolutional Neural Networks,” *Neural Information Processing Systems (NIPS)* 2012.
- [13] Z. Li, Y. Wang, *et al.*, “Deep Learning Based Radiomics (DLR) and Its Usage in Noninvasive IDH1 Prediction for Low Grade Glioma,” *Scientific Reports*, vol. 7, no. 1, 2017.
- [14] Sara Sabour, Nicholas Frosst, Geoffrey E. Hinton, “Dynamic Routing Between Capsules,” *31st Conference on Neural Information Processing Systems (NIPS 2017)*, 2017.
- [15] G.E. Hinton, N.Frosst, and S.Sabour, “Matrix Capsules with EM Routing,” *Proc. ICLR*, Feb. 2018.
- [16] P. Afshar, A. Mohammadi, K.N. Plataniotis, “Brain Tumor Type Classification via Capsule Networks,” *IEEE International Conference on Image Processing (ICIP)*, 2018.
- [17] R. LaLonde and U. Bagci, “Capsules for Object Segmentation,” *arXiv preprint, arXiv:1804.04241*, 2018.
- [18] A. Mobiny and N.H. Van, “Fast CapsNet for Lung Cancer Screening,” *arXiv preprint, arXiv:1806.07416*, 2018.
- [19] C. Xiang, L. Zhang, W. Zou, Y. Tang and C. Xu, “MS-CapsNet: A Novel Multi-Scale Capsule Network,” *IEEE Signal Processing Letters*, 2018. In Press.
- [20] Z.H. Chen and D. Crandall, “Generalized Capsule Networks with Trainable Routing Procedure,” *arXiv preprint arXiv:1808.08692*, Aug. 2018.
- [21] A. Jaiswal, W. AbdAlmageed, and P. Natarajan, “CapsuleGAN: Generative Adversarial Capsule Network,” *arXiv preprint arXiv:1802.06167*, 2018.
- [22] M.T. Bahadori, “Spectral Capsule Networks,” *Proc. ICLR*, Feb. 2018.
- [23] J.O. Neill, “Siamese Capsule Networks,” *arXiv preprint, arXiv:1805.07242*, 2018.
- [24] Y. Lecun, L. Bottou, Y. Bengio, P. Haffner, “Gradient-based learning applied to document recognition,” *Proceedings of the IEEE*, 1998.
- [25] D. Ravi, C. Wong, *et al.*, “Deep Learning for Health Informatics,” *IEEE Journal of Biomedical and Health Informatics*, vol. 21, no. 1, pp. 4-21, 2017.
- [26] Jun Cheng, Wei Yang, *et al.*, “Retrieval of Brain Tumors by Adaptive Spatial Pooling and Fisher Vector Representation,” *PloS one.*, 2016.
- [27] Justin S. Paul, Andrew J. Plassard, *et al.*, “Deep Learning for Brain Tumor Classification,” *PROCEEDINGS OF SPIE*, 2017.
- [28] Diederik P. Kingma, Jimmy Ba, “Adam: A Method for Stochastic Optimization,” *3rd International Conference for Learning Representations, San Diego*, 2015.
- [29] Marcus D. Bloice, Andreas Holzinger, “A Tutorial on Machine Learning and Data Science Tools with Python,” *Machine Learning for Health Informatics*, 2016.