AI in healthcare: Activities of the University of Naples Federico II node of the CINI-AIIS Lab

Antonio Bosco[1], Salvatore Capuozzo[1], Biase Celano[2], Michela Gravina[1], Stefano Marrone[1], Maria Paola Maurelli[1], Vincenzo Moscato[1], Giuseppe Pontillo[1], Marco Postiglione[1], Antonio M. Rinaldi[1], Laura Rinaldi[1], Cristiano Russo[1], Giancarlo Sperlì[1], Cristian Tommasino[1], Giuseppe Cringoli[1] and Carlo Sansone[1,*]

[1]*University of Naples, Federico II, Naples, Italy*
[2]*Officina Elettronica SPA, Naples, Italy*

Abstract
Artificial Intelligence (AI) embraces techniques and algorithms that, over the years, have been applied to different fields and domains. The complexity and rise of data in healthcare have prompted research into exploring the application of AI techniques in the medical field, resulting in viable and promising approaches to sustain innovation, early diagnosis and support cures. In this work, we briefly report some of the projects we carried out in this domain at the University of Naples Federico II node of the CINI-AIIS Lab, focusing on the main aims and contributions.

Keywords
Artificial Intelligence, Healthcare, Deep Learning, Machine Learning

1. Introduction

Artificial intelligence (AI) is the simulation of human intelligence in machines to perform tasks like abstraction and problem-solving. Healthcare is one of the domains that has been impacted by AI, and it is considered one of the most promising applications of AI. Indeed, AI in healthcare can be used to support patients and physicians, as well as to transform patient care and administrative processes. Also, AI-powered systems can analyze patient data to identify potential health risks and help physicians make more accurate diagnoses, which is useful in situations where patients have complex medical conditions or multiple medical conditions.

AI applications are not limited to a specific technology, but rather a collection of them. Machine Learning (ML) is a subset of AI that includes the set of algorithms that provides the systems with the ability to automatically learn and improve from experience. In healthcare, the most common application of traditional ML is precision medicine- predicting what treatment protocols are likely to succeed on a patient considering his/her attributes and the treatment context. The great majority of ML and precision medicine applications require a training dataset for which the outcome variable (i.e. onset of disease) is known (this is called supervised learning.). Deep Learning (DL) is a family of ML algorithms that use Artificial Neural Networks (ANNs) to simulate the structure of the human brain. DL approaches have gained popularity in pattern recognition tasks, particularly in image processing, thanks to Convolutional Neural Networks (CNNs), a type of DL architecture that can learn the best set of features for a given task. As a consequence, DL is increasingly being used in healthcare, particularly in medical image analysis and radiomics to extract clinically relevant features from images beyond what can be perceived by the human eye. However, medical imaging is not the only domain in which DL is being proven useful. Making sense of human language has indeed been one of the main goals of AI researchers. Natural Language Processing (NLP), strongly based on ML and DL techniques, includes applications such as speech recognition and text analysis. In healthcare, the dominant applications of NLP involve the creation, understanding and classification of clinical documentation and published research. Moreover, NLP systems can analyse unstructured clinical notes on patients, prepare reports, transcribe patient interactions and conduct conversational AI.

It should be clear that thanks to its multitude of applications, AI plays a key role in healthcare. In this paper, we will illustrate some of the projects exploiting AI techniques in the medical field carried out at the University of

Ital-IA 2023: 3rd National Conference on Artificial Intelligence, organized by CINI, May 29–31, 2023, Pisa, Italy
*Corresponding author.
carlo.sansone@unina.it (C. Sansone)
0000-0002-1266-2154 (A. Bosco); 0000-0002-2578-9349 (S. Capuozzo); 0000-0001-5033-9617 (M. Gravina); 0000-0001-6852-0377 (S. Marrone); 0000-0002-7564-3356 (M. P. Maurelli); 0000-0002-0754-7696 (V. Moscato); 0000-0001-5425-1890 (G. Pontillo); 0000-0003-1470-8053 (M. Postiglione); 0000-0001-7003-4781 (A. M. Rinaldi); 0000-0002-0975-2052 (L. Rinaldi); 0000-0002-8732-1733 (C. Russo); 0000-0003-4033-3777 (G. Sperlì); 0000-0001-9763-8745 (C. Tommasino); 0000-0001-9828-0079 (G. Cringoli); 0000-0002-8176-6950 (C. Sansone)

Naples Federico II node of the CINI-AIIS Lab, highlighting their innovative aspects and contributions.

2. Multi-task BioNER

Biomedical Named Entity Recognition (BioNER) involves identifying mentions of biomedical entities (e.g. disorders, drugs, genetic information) from unstructured text data and is a fundamental step in many downstream tasks, such as the development of chatbots, Q/A systems or knowledge graphs. BioNER systems have evolved over time through deep learning techniques based on Bidirectional Long-Short Term Memory networks [1] and character-level features of words [2, 3]. Large-scale language models pre-trained on biomedical corpora [4, 5] have recently shown potential in enhancing the state-of-the-art. However, developing a BioNER system is still difficult due to the high frequency of synonyms, alternate spellings and polysemous words. Furthermore, due to the privacy concerns related to the healthcare domain, there is a lack of annotated data, and inference time and memory constraints arise when embedding models in clinical workflows. When building BioNER models, researchers are often forced to integrate datasets from different sources that include annotations for different entity types.

To handle these issues without having to design, train and deploy a single Transformer-based BioNER model for each available dataset — which is extremely impractical due to their constraining memory requirements and to the problems which would arise due to overlapping predictions, e.g. two models assigning two different entity types to the same mention —, we propose TaughtNet [6], a multi-task framework based on knowledge distillation that fine-tunes a single transformer architecture to recognize multiple entity types. TaughtNet produces a single, independent Student Transformer model that is capable of recognizing a variety of entity types. The proposed method is compared to other state-of-the-art approaches, and the results show that TaughtNet achieves higher performance across multiple benchmark datasets.

We show an overview of our framework in Figure 1. The methodological workflow assumes that a set of single-task Teachers represented by their parameters $\theta_T^1, \theta_T^2, \ldots, \theta_T^n$ (n being the number of teachers) has already been trained, and can be summarized with the steps listed below:

1. *Datasets aggregation.* The available single-entity datasets are merged together to build an aggregated multi-entity dataset $\mathcal{D}_S$.
2. *Retrieval of Teacher predictions.* Each sample in $\mathcal{D}_S$ is provided to each teacher as input and the resulting output distributions are stored.

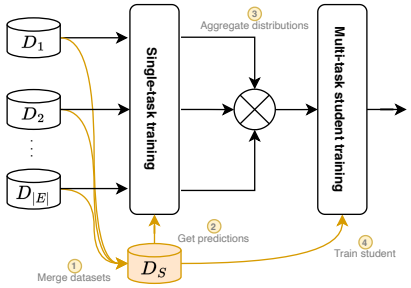

Figure 1: Overview of the training framework for *TaughtNet*. I) multiple datasets are combined to create a larger one called $\mathcal{D}_S$, to serve as a reference during the Student's training; II), the Teacher models make predictions for each sample in $\mathcal{D}_S$, and their output distributions are combined to create a new corpus; III), the Student is trained using two loss components, one based on the Teachers and the other based on the ground truth data from $\mathcal{D}_S$.

3. *Distributions aggregation.* A single output distribution is generated by integrating the output distributions from each teacher.
4. *Student Training.* A Student model with parameters θ_S is trained by taking both the ground truth and the knowledge of Teachers into consideration. Specifically, we define the loss function as follows:

$$\mathcal{L}(\mathcal{D}_S; \theta_S, \theta_T^1, \ldots, \theta_T^n) = \lambda \cdot \mathcal{L}_{KD}(\mathcal{D}_S; \theta_S, \theta_T^1, \ldots, \theta_T^n) + + (1 - \lambda) \cdot \mathcal{L}_{GT}(\mathcal{D}_S; \theta_S), \quad (1)$$

where $\mathcal{L}_{KD}$ and $\mathcal{L}_{GT}$ are the knowledge distillation and ground-truth loss, respectively, while λ is a hyperparameter controlling their weight on the overall loss $\mathcal{L}$.

In our experiments, we show that not only does TaughtNet allow to recognize multiple biomedical entity types while ensuring state-of-the-art performance, but it also can be applied to smaller and lighter students which can be more easily deployed on hardware with limited resources w.r.t. heavy transformer models.

3. The KFM

The scanning process of parasite eggs contained in faecal samples of farm animals is a common task in the veterinary medicine domain which can be automated thanks to Convolutional Neural Networks (CNNs) for Object Detection. The *Kubic FLOTAC Microscope* (KFM) is a compact, low-cost, versatile, and portable digital microscope

designed to autonomously analyze faecal specimens prepared with FLOTAC or Mini-FLOTAC, in both field and laboratory settings, for different parasites and hosts [7]. Having been proven to acquire images comparable to the view provided by traditional optical microscopes, the KFM can autonomously scan and acquire images in a few minutes, allowing the operator to focus on a different task. The KFM is composed of electro-mechanical components which enable both manual and automatic FLOTAC/Mini-FLOTAC reading discs 3D scans, a firmware that allows 3D movements of the camera, remote interactions, and scans retrieval, and external agents which enable users to connect with the KFM hardware (KFM web interface and KFM app) and to process scans for parasite eggs detection (KFM AI server). The device can be remotely controlled (the optical part with LED and digital camera and tray) by any user with external devices, like smartphones, tablets or PCs, through a dedicated web interface and app, making it possible to easily manage multiple KFMs with a single device, start several scans in parallel and switch from one KFM to another effortlessly. Images collected with the KFM can be stored in the KFM AI server and/or transmitted to diagnostic hubs in order to have a quick diagnosis or a parasitological consultation. The whole system flow is described in Figure 2.

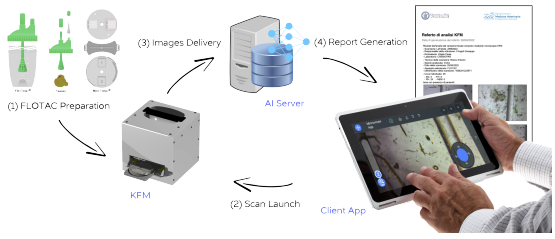

Figure 2: A schematic picture of the working flow of the KFM system, composed of the hardware, the AI server and the client app interface.

4. MGMT Promoter Methylation Identification

One of the most aggressive malignant tumours is the Glioblastoma Multiforme (GBM), which is known for its extremely low survival rate. Even if alkylating chemotherapy is typically adopted to fight this tumour, sometimes it may result inefficient since the O(6)-methylguanine-DNA methyltransferase (MGMT) enzyme repair abilities counter the cytotoxic effects of alkylating agents, preventing the tumours cells destruction. However, MGMT promoter regions may be subject to a phenomenon called methylation, a biological process preventing MGMT enzymes from destroying the alkyl agents [8]. Consequently, the presence of the methylation process in GBM patients' brains can be associated with a predictive biomarker of response to therapy and a prognosis factor [9]. Unfortunately, identifying methylation signs is not trivial, and often requires time-consuming, expensive and invasive procedures. In a recent work [10], we propose to face MGMT promoter methylation identification by analyzing Magnetic Resonance Imaging data with a Deep Learning based approach, consisting of a CNN operating on suspicious regions on the FLAIR series, which are pre-selected through an unsupervised Knowledge-Based filter leveraging both FLAIR and T1-weighted series.

Given the absence of segmentation masks, determining somehow the ROI representing the tumour region is a crucial and critical step. We proposed to select the area of interest in an unsupervised manner, leveraging past medical experience [11] for tumour recognition. In particular, we took advantage of simple characteristics of lesioned tissues for each considered series: in T1-W slices, tumour areas have pixels whose intensity is higher than cerebrospinal fluids (CSF) but lower than any other kind of tissue, while in FLAIR slices, pixels with the highest intensity belong to the tumour region. Leveraging these characteristics, it is possible to pre-select potentially lesioned tissues from volumes of these series by applying a threshold on the histogram of the signal intensities occurrences. The output of this process is a mask consisting of huge clusters, in case of tumour presence, corresponding to the ROIs, representing the potentially lesioned areas, with little outliers, or of sparse outliers in the opposite case, as shown in Figure 3. The resulting ROI is used to select from the FLAIR slices the portion of the image to be considered by the actual methylation detection module. We chose the FLAIR sequence for the good performance shown in the literature for related tasks [12]. The proposed module aims to compensate for the lack of segmentation masks with an early fusion technique, where information coming from multiple sources is merged to highlight different characteristics [13].

We adopted a CNN to face the task of MGMT promoter methylation identification. In particular, we built a sequential network from the ground called *MGMTClassifier*, which is composed of seven Convolutional Blocks and two fully connected layers separated by Rectified Linear Unit (ReLU) as an activation function. These Convolutional Blocks consist of a convolutional layer, followed by batch normalization and ReLU function. To reduce the number of training parameters and avoid overfitting at the same time, we adopt depth-wise separable convolution [14]. This simple but effective DL-based approach is able to outperform state-of-the-art solutions while consisting of less than 0.29% of their parameters (about 10 million of typical CNNs versus 40561 of the proposed approach).

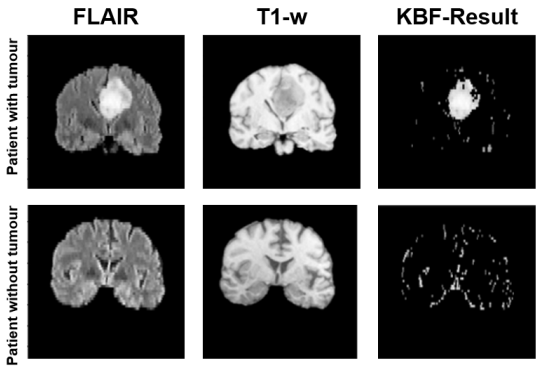

Figure 3: An illustrative example of the results produced by the ROI selection process on two patients: if the slice contains a tumour (**top row**) a huge cluster is generated in the pre-selection mask, while in the opposite case (**bottom row**) the mask contains sparse outliers.

5. Axillary Lymph Node Status Assessment in Breast Cancer

Among women, breast cancer (BC) is the most frequent form of tumour, and the axillary lymph node status (ALNS) is considered a crucial indicator, representing one of the most influencing and independent prognostic factors. Magnetic Resonance Imaging (MRI) is always performed for BC stage definition and plays a key role in primary tumour examination. The most important MRI sequence is the dynamic-contrast enhanced (DCE) that, thanks to the high contrast resolution, provides information about the tumour morphology, size, and perfusional behaviour allowing the distinction between benign and malignant lesions, the prediction of biological aggressiveness and the prognostic evaluation. The T2-weighted (T2) imaging is a standard component of breast MRI exams, most prominently utilized for the identification of cysts, allowing a better depiction of lesion morphology and perifocal or prepectoral edema within the breast. Diffusion-Weighted Imaging (DWI) is another sequence acquired during the MRI exam. It reflects the mobility of water molecules diffusing in tissues, revealing tissue organization at the microscopic level and providing complementary information for lesion assessment in comparison with the DCE scan.

Assessment of ALNS indicates inherent primary tumour properties, whose examination enables the discovery of minimally invasive solutions for the sentinel node biopsy currently being utilized. In [15], we only focused on the DCE sequence to analyze different tumour bounding options for the prediction of ALN status. In particular, the work described in [15] evaluates how the amount of the included non-tumour tissue impacts the assessment by exploiting CNNs in the solution. On the basis of the results reported in [15, 16], we aim to include multiple and complementary sequences acquired during the MRI exam, namely DCE, T2 and DWI, proposing an approach based on multimodal learning, where heterogeneous sources of data are fused to provide a shared representation. Moreover, we will focus on a methodology based on intermediate fusion, leveraging the ability of deep neural networks to provide an effective high-level representation of the input and taking into account both MR images and clinical information.

6. Assessing Brain Health with the Brain-Age Paradigm

In the search for objective imaging-derived markers of brain health and pathology, the brain-age paradigm has emerged as a promising approach. Briefly, machine learning methods are used to model chronological age as a function of structural brain MRI scans in healthy people, and the resulting model of 'normal' brain ageing is used for neuroimaging-based age prediction in unseen subjects [17]. The extent to which each subject deviates from healthy brain-ageing trajectories, expressed as the difference between predicted and chronological age (brain-predicted age difference, brain-PAD), has been proposed as an index of structural brain health, sensitive to brain pathology in a wide spectrum of neurological and psychiatric disorders [18]. We applied the brain-age paradigm to a target clinical population of patients with Fabry Disease (FD), a rare genetic multisystemic disorder that also involves the brain but lacks quantitative neuroimaging biomarkers. We selected MRI scans of FD patients and healthy controls from the same Institution. The Fabry stabilization index (FASTEX) was recorded as a measure of multi-organ FD severity. We trained and evaluated a model of healthy brain ageing on a large dataset (total N = 2160; male/female = 1293/867; mean age = 33 years, age range = 4-86) comprising 3D T1-weighted brain scans of healthy subjects from 8 publicly available sources. Our brain-age model was based on the DenseNet264 architecture [19] adapted from the implementation available at Project MONAI [1] by adding a linear regression layer for the prediction of a continuous variable and a 0.2 dropout rate after each dense layer to reduce the risk of overfitting. Mean absolute error (MAE) and coefficient of determination (R2) were used to quantify model performance. The final model was applied to the internal cohort of FD patients and healthy controls to generate brain-predicted ages and corresponding brain-PAD values (Figure 4). Lastly, within a linear modelling framework, brain-PAD was tested for

[1] https://docs.monai.io/en/stable/_modules/monai/networks/nets/densenet.html

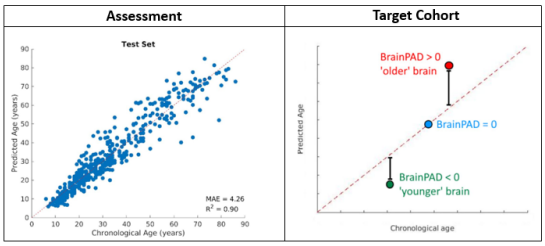

Figure 4: The model with the lowest validation loss is chosen, and performance is measured on the previously unseen cases of the test set (left). The final model is also applied to the target clinical population (right), composed of the internal cohort of FD patients and healthy controls, to generate brain-predicted ages and corresponding brain-PAD values.

demographics-adjusted associations with a diagnostic group (FD vs healthy controls) and FASTEX score.

We selected 52 FD patients (40.6 ± 12.6 years; 28 females) and 58 healthy controls (38.4 ± 13.4 years; 28 females). The brain-age model achieved accurate out-of-sample performance (MAE = 4.01y, R2 = 0.90). FD patients had significantly higher brain-PAD than healthy controls (estimated marginal means: 3.1 vs -0.1, p=0.01), and brain-PAD was also associated with multi-organ clinical severity as assessed with the FASTEX score (B = 0.10, p = 0.02). In conclusion, using deep learning and the brain-age paradigm, we found that patients with FD have older-appearing brains compared to healthy controls. The gap between brain-predicted and chronological age correlates with multi-organ disease severity, potentially representing a novel quantitative imaging biomarker of brain involvement in FD.

7. Histopathological images Deep Feature representation for CBIR in smart PACS

Pathological Anatomy is moving toward computerizing processes mainly due to the extensive digitization of histology slides that resulted in the availability of many Whole Slide Images (WSIs). Their use is essential, especially in cancer diagnosis and research, and raises the pressing need for increasingly influential information archiving and retrieval systems. Picture Archiving and Communication Systems (PACSs) represent an actual possibility to archive and organize this growing amount of data. The storage, retrieval, and analysis of biomedical images are essential tools of Picture Archiving and Communication Systems (PACSs). The same ones are also useful in related contexts such as computational pathology (CPATH), where their diffusion is still limited. A weakness of traditional PACSs concerns the capability to perform a query only employing metadata. The design and implementation of a robust and accurate methodology for querying them in the pathology domain using a novel approach are mandatory. In particular, the Content-Based Image Retrieval (CBIR) methodology can be involved in the PACSs using a query-by-example task. In this context, one of many crucial points of CBIR concerns the representation of images as feature vectors, and the accuracy of retrieval mainly depends on feature extraction. In order to allow intelligent multimodality query posing [20], Content-Based Image Retrieval (CBIR) techniques can be used [21]. The feature extraction methods and similarity functions are crucial aspects of a CBIR [22]. Over the years, different approaches for image retrieval have been proposed, also in the computational pathology field using different techniques such as deep learning [23]. Thus, our study explored different representations of WSI patches by features extracted from pretrained Convolution Neural Networks (CNNs). In order to perform a helpful comparison, we evaluated features extracted from different layers of state-of-the-art CNNs using different dimensionality reduction techniques. Furthermore, we provided a qualitative analysis of obtained results. To accomplish a more precise analysis, we analyzed in depth all the best CNN configurations using a confusion matrix. We intend to understand which dataset category is not correctly retrieved. According to Figure 5, it is clear that the queries using benign examples are not correctly retrieved. They are often misunderstood with in situ images and normal ones. Furthermore, normal and invasive queries are usually recognized with good precision value.

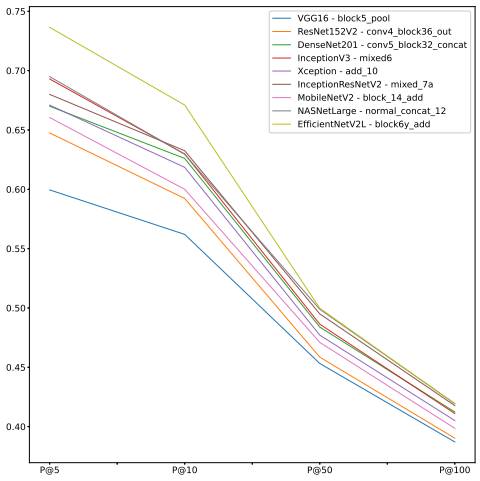

Figure 5: Precision comparison at 5, 10, 50 and 100 obtained from retrieval using CNN layers and global average pooling to reduce dimensionality.

References

[1] G. Lample, M. Ballesteros, S. Subramanian, K. Kawakami, C. Dyer, Neural architectures for named entity recognition, in: HLT-NAACL, 2016.

[2] X. Ma, E. Hovy, End-to-end sequence labeling via bi-directional LSTM-CNNs-CRF, in: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Berlin, Germany, 2016, pp. 1064–1074. URL: https://aclanthology.org/P16-1101. doi:10.18653/v1/P16-1101.

[3] M. Habibi, L. Weber, M. L. Neves, D. L. Wiegandt, U. Leser, Deep learning with word embeddings improves biomedical named entity recognition, Bioinformatics 33 (2017) i37 – i48.

[4] I. Beltagy, K. Lo, A. Cohan, Scibert: A pretrained language model for scientific text, in: EMNLP/IJCNLP, 2019.

[5] S. Gururangan, A. Marasović, S. Swayamdipta, K. Lo, I. Beltagy, D. Downey, N. A. Smith, Don't stop pretraining: Adapt language models to domains and tasks, in: ACL, 2020.

[6] V. Moscato, M. Postiglione, C. Sansone, G. Sperli, Taughtnet: Learning multi-task biomedical named entity recognition from single-task teachers, IEEE Journal of Biomedical and Health Informatics (2023) 1–12. doi:10.1109/JBHI.2023.3244044.

[7] G. Cringoli, A. Amadesi, M. P. Maurelli, B. Celano, G. Piantadosi, A. Bosco, L. Ciuca, M. Cesarelli, P. Bifulco, A. Montresor, et al., The kubic flotac microscope (kfm): a new compact digital microscope for helminth egg counts, Parasitology 148 (2021) 427–434.

[8] M. Christmann, B. Verbeek, W. P. Roos, B. Kaina, O6-methylguanine-dna methyltransferase (mgmt) in normal tissues and tumors: enzyme activity, promoter methylation and immunohistochemistry, Biochimica et Biophysica Acta (BBA)-Reviews on Cancer 1816 (2011) 179–190.

[9] W. Haque, E. Thong, S. Andrabi, V. Verma, E. B. Butler, B. S. Teh, Prognostic and predictive impact of mgmt promoter methylation in grade 3 gliomas, Journal of Clinical Neuroscience 85 (2021) 115–121.

[10] S. Capuozzo, M. Gravina, G. Gatta, S. Marrone, C. Sansone, A multimodal knowledge-based deep learning approach for mgmt promoter methylation identification, Journal of Imaging 8 (2022) 321.

[11] M. C. Clark, L. O. Hall, D. B. Goldgof, R. Velthuizen, R. Murtagh, M. S. Silbiger, Unsupervised brain tumor segmentation using knowledge-based fuzzy techniques, in: Fuzzy and neuro-fuzzy systems in medicine, CRC Press, 2017, pp. 137–170.

[12] S. Roozpeykar, M. Azizian, Z. Zamani, M. R. Farzan, H. A. Veshnavei, N. Tavoosi, A. Toghyani, A. Sadeghian, M. Afzali, Contrast-enhanced weighted-t1 and flair sequences in mri of meningeal lesions, American Journal of Nuclear Medicine and Molecular Imaging 12 (2022) 63.

[13] D. Ramachandram, G. W. Taylor, Deep multimodal learning: A survey on recent advances and trends, IEEE signal processing magazine 34 (2017) 96–108.

[14] F. Chollet, Xception: Deep learning with depthwise separable convolutions, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2017, pp. 1251–1258.

[15] D. Santucci, E. Faiella, M. Gravina, E. Cordelli, C. de Felice, B. Beomonte Zobel, G. Iannello, C. Sansone, P. Soda, Cnn-based approaches with different tumor bounding options for lymph node status prediction in breast dce-mri, Cancers 14 (2022) 4574.

[16] M. Gravina, E. Cordelli, D. Santucci, P. Soda, C. Sansone, Evaluating tumour bounding options for deep learning-based axillary lymph node metastasis prediction in breast cancer, in: 2022 26th International Conference on Pattern Recognition (ICPR), IEEE, 2022, pp. 4335–4342.

[17] J. H. Cole, K. Franke, Predicting age using neuroimaging: innovative brain ageing biomarkers, Trends in neurosciences 40 (2017) 681–690.

[18] T. Kaufmann, D. van der Meer, N. T. Doan, E. Schwarz, M. J. Lund, I. Agartz, D. Alnæs, D. M. Barch, R. Baur-Streubel, A. Bertolino, et al., Common brain disorders are associated with heritable patterns of apparent aging of the brain, Nature neuroscience 22 (2019) 1617–1623.

[19] G. Huang, Z. Liu, L. Van Der Maaten, K. Q. Weinberger, Densely connected convolutional networks, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2017, pp. 4700–4708.

[20] A. M. Rinaldi, C. Russo, C. Tommasino, Visual query posing in multimedia web document retrieval, in: 2021 IEEE 15th International Conference on Semantic Computing (ICSC), IEEE, 2021, pp. 415–420.

[21] V. Jeyakumar, B. Kanagaraj, A medical image retrieval system in pacs environment for clinical decision making, in: Intelligent Data Analysis for Biomedical Applications, Elsevier, 2019, pp. 121–146.

[22] A. M. Rinaldi, C. Russo, A content based image retrieval approach based on multiple multimedia features descriptors in e-health environment, in: 2020 IEEE International Symposium on Medical Measurements and Applications (MeMeA), IEEE, 2020, pp. 1–6.

[23] J. Ma, X. Jiang, A. Fan, J. Jiang, J. Yan, Image matching from handcrafted to deep features: A survey, International Journal of Computer Vision 129 (2021) 23–79.