




Contents lists available at ScienceDirect

Expert Systems With Applications

journal homepage: www.elsevier.com/locate/eswa

Exploiting causality signals in medical images: A pilot study with empirical results

Gianluca Carloni^{a,b,*}, Sara Colantonio^a^a National Research Council, Institute of Information Science and Technologies (ISTI) "Alessandro Faedo", Via Moruzzi, 1, Pisa, 56127, Italy^b Department of Information Engineering, University of Pisa, Via Caruso, 16, Pisa, 56122, Italy

ARTICLE INFO

Keywords:

Causality
Convolutional neural network
Deep learning
Medical imaging
Attention

ABSTRACT

We present a novel technique to discover and exploit weak causal signals directly from images via neural networks for classification purposes. This way, we model how the presence of a feature in one part of the image affects the appearance of another feature in a different part of the image. Our method consists of a convolutional neural network backbone and a causality-factors extractor module, which computes weights to enhance each feature map according to its causal influence in the scene. We develop different architecture variants and empirically evaluate all the models on two public datasets of prostate MRI images and breast histopathology slides for cancer diagnosis. We study the effectiveness of our module both in fully-supervised and few-shot learning, we assess its addition to existing attention-based solutions, we conduct ablation studies, and investigate the explainability of our models via class activation maps. Our findings show that our lightweight block extracts meaningful information and improves the overall classification, together with producing more robust predictions that focus on relevant parts of the image. That is crucial in medical imaging, where accurate and reliable classifications are essential for effective diagnosis and treatment planning.

1. Introduction

Automatic diagnosis models from medical data can potentially transform how patients are treated, especially in oncology. They could reduce the need for invasive tests and increase the likelihood of successful outcomes for the most severe cases. In this regard, some successful examples of machine learning (ML) systems for medical diagnosis exist, such as colon cancer diagnosis from gene expression profiling data (Su et al., 2022), diagnosis of neurological diseases using voice data (Wroge et al., 2018), and identification of patients with pulmonary hypertension using electronic health records (Kogan et al., 2023).

In recent years, the concepts of causal inference and causal reasoning have received increasing attention across the Artificial Intelligence (AI) community. This trend began with the very first work of the computer scientist Judea Pearl on Bayesian networks and the mathematical formalization of causality, which enabled the creation of computational systems that can automatically model causality (Pearl, 1985, 2009; Pearl & Mackenzie, 2018). Today, we have inspiring examples of the integration of causality into the ML community (Luo, Peng, & Ma, 2020; Schölkopf, 2022) and the deep learning research (Berrevoets, Kacprzyk,

Qian, & van der Schaar, 2023), with extensions to causal representation learning (Schölkopf, Locatello, Bauer, Ke, Kalchbrenner, Goyal, & Bengio, 2021), causal discovery under distribution shifts (Perry, Von Kügelgen, & Schölkopf, 2022) and with incomplete data (Wang, Menkovski, Wang, Du, & Pechenizkiy, 2020). Unfortunately, this line of research has always had in common the fact that the processed data are tabular, structured, not always real but simulated, and very often accompanied by a priori information about the process that generated them.

Unlike tabular data, when it comes to images, their representation does not include any explicit indications regarding objects or patterns. Instead, individual pixels are used to convey a particular scene visually, and image datasets do not usually provide labels describing the objects' dispositions. Additionally, unlike video frames, a single image cannot reveal the dynamics of the appearance and change of objects in a scene. These critical issues could explain why images have been neglected by research on the tabular causal discovery, where instead, there are established algorithms (Chickering, 2002; Spirtes & Glymour, 1991; Spirtes, Glymour, & Scheines, 2000). A particular case would be discovering hidden causalities among objects in an image dataset, as suggested by Terziyan and Vitko (2023), who conceive a way to

* Corresponding author at: National Research Council, Institute of Information Science and Technologies (ISTI) "Alessandro Faedo", Via Moruzzi, 1, Pisa, 56127, Italy.

E-mail addresses: gianluca.carloni@isti.cnr.it (G. Carloni), sara.colantonio@isti.cnr.it (S. Colantonio).

<https://doi.org/10.1016/j.eswa.2024.123433>

Received 17 August 2023; Received in revised form 2 January 2024; Accepted 7 February 2024

Available online 15 February 2024

0957-4174/© 2024 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

compute possible causal relationships within images. Although the idea is compelling, that work is preliminary, and a thorough investigation of the effectiveness of their method is lacking.

In our work, we intervene in this lack and propose a way to discover and exploit weak causal signals within images without requiring prior knowledge and use them to enhance convolutional neural network (CNN) classifiers. By combining a regular CNN with the proposed causality-factors extraction module, we present a new scheme based on feature map enhancement to enable “causality-driven” CNNs. This way, we weight each feature map according to its causal influence in the scene, in an attention-inspired fashion. We frame our system as an automatic diagnosis model from medical images since we study the efficacy of the proposed methods with extensive empirical evaluations on publicly available datasets of MRI images and histopathology slides. We study the effectiveness of our module both in fully-supervised and few-shot learning regimes, investigate its integration with existing attention-based solutions, and conduct ablation studies. Besides investigating the quantitative aspect, we also explored the concept of explainability of AI in our evaluation. The results on class activation maps demonstrate that our method improves classification and produces more robust predictions by focusing on the relevant parts of the image, thus enhancing reliability, trustworthiness, and user confidence.

Our paper is structured as follows. First, in Section 2, we provide the concepts behind the causal signals’ interpretation in images. Then, we start Section 3 by describing the novelty of our work, namely the methodological framework and the causality-factors extractor module we introduced. We also illustrate the datasets, the training scheme, and the evaluation details. Later, we present our main results in Section 4, explore the significance of our findings in the general discussion in Section 5, and pull the threads in Section 6.

2. Causality signals in images

Lopez-Paz, Nishihara, Chintala, Scholkopf, and Bottou (2017) propose the idea of “causal disposition” as a simple way to understand the hidden causes in images instead of using the methods of do-calculus and causal graphs from Pearl’s framework (Pearl, 2009; Pearl & Mackenzie, 2018). In their view, by counting the number $C(A, B)$ of images in which the causal dispositions of artifacts A and B are such that B disappears if one removes A , one can assume that artifact A causes the presence of artifact B when $C(A, B)$ is greater than the converse $C(B, A)$. For instance, they argue that the presence of a car causes the presence of a wheel, but not the other way around, because removing the car would make the wheel disappear, but removing the wheel would not make the car disappear. By studying such asymmetries, the authors find the causal direction between pairs of random variables representing features of objects and their contexts in images. Although the causal disposition concept is more primitive than the interventional approach, it could be the only way to proceed with limited a priori information. This concept leads to the intuition that any causal disposition induces a set of asymmetric causal relationships between the artifacts from an image (features, object categories, etc.) that represent (weak) causality signals regarding the real-world scene. A point of contact with machine vision systems would be to automatically infer such asymmetries from an observed image dataset.

Terziyan and Vitko (2023) suggest a way to compute estimates for possible causal relationships within images via CNNs. CNNs obtain the essential features required for classification not directly from the pixel representation of the input image but through a series of convolution and pooling operations designed to capture meaningful features from the image. Convolution layers are responsible for summarizing the presence of specific features in the image and generating a set of feature maps accordingly. Pooling consolidates the presence of particular features within groups of neighboring pixels in square-shaped sub-regions of the feature map. When a feature map F^i contains only non-negative numbers (e.g., thanks to ReLU functions) and is normalized in the

interval $[0, 1]$, we can interpret its values as probabilities of that feature to be present in a specific location. For instance, $F_{r,c}^i$ is the probability that the feature i is recognized at coordinates r, c . By assuming that the last convolutional layer outputs and localizes to some extent the object-like features, we may modify the architecture of a CNN such that the $n \times n$ feature maps (F^1, F^2, \dots, F^k) obtained from that layer got fed into a new module that computes pairwise conditional probabilities of the feature maps. The resulting $k \times k$ map would represent the causality estimates for the features and be called *causality map*. Given a pair of feature maps F^i and F^j and the formulation that connects conditional probability with joint probability, $P(F^i|F^j) = \frac{P(F^i, F^j)}{P(F^j)}$, Terziyan and Vitko (2023) suggest to heuristically estimate this quantity by adopting two possible methods, namely *Max* and *Lehmer*. The *Max* method considers the joint probability to be the maximal presence of both features in the image (each one in its location):

$$P(F^i|F^j) = \frac{(\max_{r,c} F_{r,c}^i) \cdot (\max_{r,c} F_{r,c}^j)}{\sum_{r,c} F_{r,c}^j} \quad (1)$$

On the other hand, the *Lehmer* method entails computing

$$P(F^i|F^j)_p = \frac{LM_p(F^i \times F^j)}{LM_p(F^j)} \quad (2)$$

where $F^i \times F^j$ is a vector of n^2 pairwise multiplications between each element of the two $n \times n$ feature maps, while LM_p is the generalized Lehmer mean function (Bullen, 2003) with parameter p , which is an alternative to power means for interpolating between minimum and maximum of a vector x via harmonic mean ($p = -1$), arithmetic mean ($p = 0$), and contraharmonic mean ($p = 1$): $LM_p(x) = \frac{\sum_{k=1}^n x_k^{p+1}}{\sum_{k=1}^n x_k^p}$. Eqs. (1) and (2) could be used to estimate asymmetric causal relationships between features F^i and F^j , since, in general, $P(F^i|F^j) \neq P(F^j|F^i)$. By computing these quantities for every pair i and j of the k feature maps, the $k \times k$ causality map is obtained. We interpret asymmetries in such probability estimates as weak causality signals between features, as they provide some information on the cause–effect of the appearance of a feature in one place of the image, given the presence of another feature within some other places of the image. Accordingly, a feature may be deemed to be the reason for another feature when $P(F^i|F^j) > P(F^j|F^i)$, that is $(F^i \rightarrow F^j)$, and vice versa. As an example, Fig. 1 depicts a causality map to give a visual interpretation of this concept.

In this work, we integrate a regular CNN with a new causality-extraction module to explore the features and causal relationships between them extracted during training. The previous work that inspired us (Terziyan & Vitko, 2023) is preliminary, and we introduce a novel attention-like scheme based on feature maps enhancement to enable “causality-driven” CNNs, providing an extensive empirical evaluation of the impact of this new introduction on real data. We hypothesize that it would be possible and reasonable to get some weak causality signals from the individual images of some medical datasets without adding primary expert knowledge and leverage them to better guide the learning phase. Ultimately, a model trained in such a manner would exploit weak causal dispositions of objects in the image scene to distinguish the tumor status of a medical image.

3. Material and methods

3.1. Embedding causality into CNNs

Usually, a CNN performs image classification based on the final set of (flattened) $n \times n \times k$ feature maps obtained just before the dense layers that constitute the classifier. In the following, we describe how the architecture of such a regular CNN (baseline) might be modified to make the classifier consider the information entailed in the estimated causality map.

Feature concatenation is a basic (yet popular) way to embed additional information in CNNs. Indeed, by concatenating the flattened

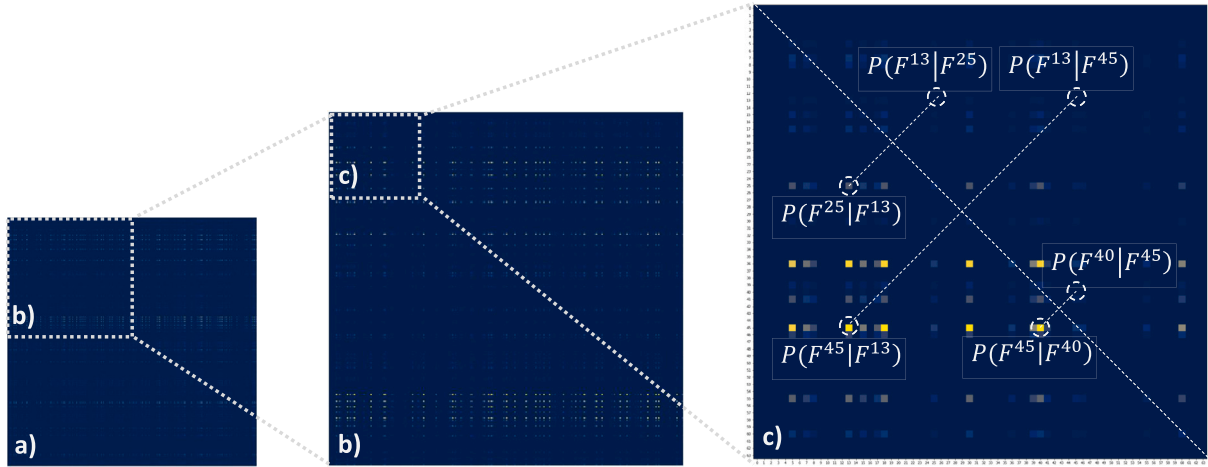


Fig. 1. Zoomed-in visualizations of a sample causality map computed with Eq. (1) on 512 feature maps extracted from an input image. (a) 512×512 original causality map; (b) 256×256 zoom-in of (a); (c) 64×64 zoom-in of (b), where dashed circles indicate exemplar elements and their corresponding elements opposite the main diagonal, representing conditional asymmetries of the type $P(F^i|F^j) \neq P(F^j|F^i)$. We can see that, for instance, $P(F^{25}|F^{13}) > P(F^{13}|F^{25})$, that is $F^{25} \rightarrow F^{13}$, and $P(F^{45}|F^{40}) > P(F^{40}|F^{45})$, that is $F^{45} \rightarrow F^{40}$.

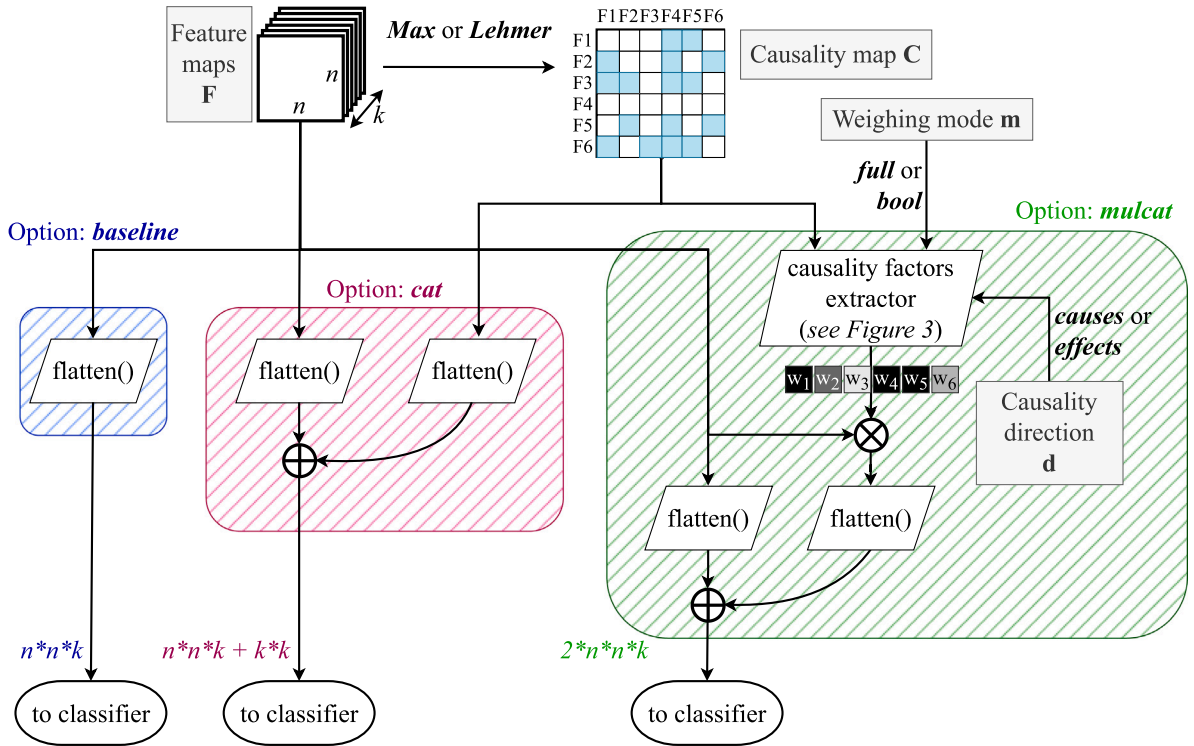


Fig. 2. Overview of the different settings investigated in this work: *Baseline*, *Cat*, and the proposed *Mulcat*. Assuming $k = 6$ feature maps as an example, the tensor F of feature maps that are obtained from a CNN just before the classifier can be either flattened and used as they are (Option *baseline*) or can be leveraged to compute the causality map C via the *Max* or *Lehmer* method. Once obtained, C can be flattened as well and concatenated to the feature maps (Option *Cat*) or fed to our proposed causality factors extractor (see Fig. 3) to implement the Option *Mulcat*. The latter produces a vector of causality factors that weighs the feature maps obtaining a causality-driven version of them, which is then concatenated to the original ones and fed to the classifier. Weighing mode m and causality direction d are two external signals used to tune the functioning of the system.

causality map to the flattened set of feature maps just before the classifier, Terziyan and Vitko (2023) let the CNN learn how these causality estimates influence image classification. That means that in addition to the $n \times n \times k$ features, the fully connected layers of the classifier will now have a $k \times k$ input, and the weights for the corresponding connections (i.e., actual causality influences) will be learned by back-propagation the same way as other neural network parameters. We will call this method the *Cat* (*concatenate*) option (see the magenta box in Fig. 2).

Alternatively, one could enhance or penalize parts of the existing information according to the newly gained one. Our proposition here is a new way to exploit the causality map: this time, it is used to

compute a vector of causality factors that multiply (i.e., weighs) the feature maps so that each feature map is strengthened according to its causal influence within the image's scene. After multiplication, the obtained causality-driven version of the feature maps is flattened and concatenated to the flattened original ones, producing a $2 \times n \times n \times k$ input to the classifier. We will call this method the *Mulcat* (*multiply and concatenate*) option (see the green box in Fig. 2).

At the core of the *Mulcat* option stands our *causality factors extractor* module, which yields the vector of weights needed to multiply the feature maps (see Fig. 3). The main idea here is to look for asymmetries between elements opposite the main diagonal of the causality map,

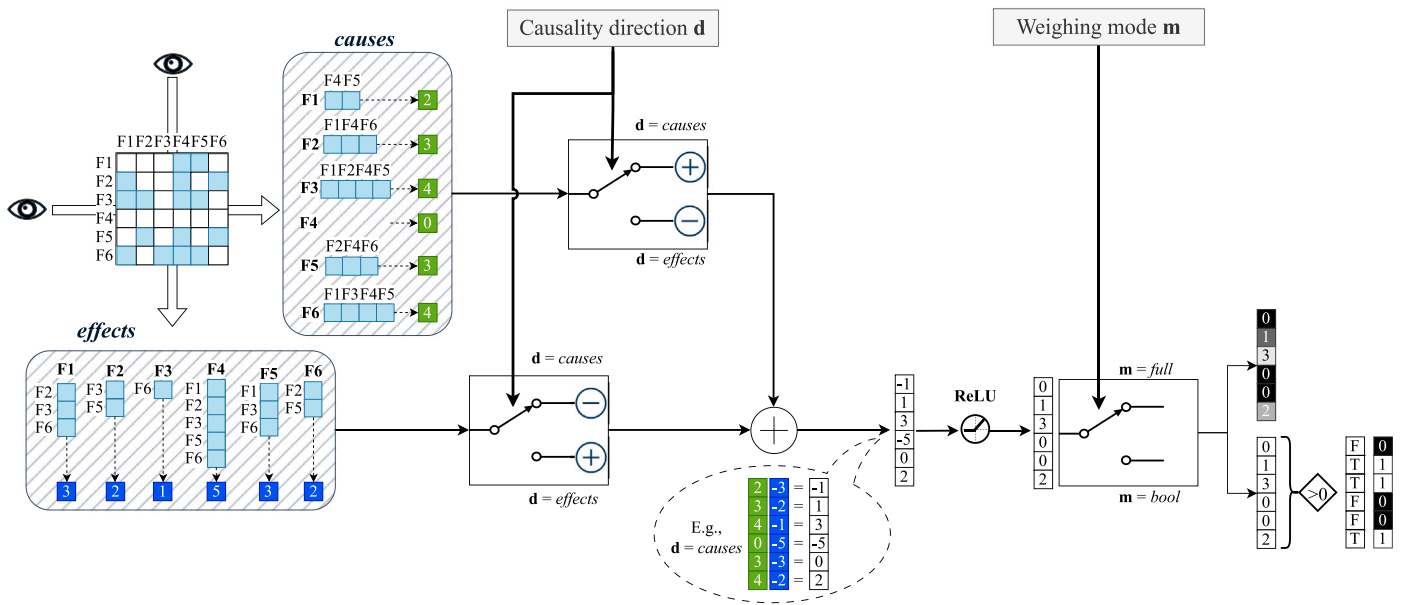


Fig. 3. The internals of the proposed *causality factors extractor* block of Fig. 2 given an example causality map. Cyan squares in the causality map indicate whether the probability value of one element is greater than its element opposite the main diagonal. The *causes* box shows how the causality map is processed row-wise for each feature map: the number of times that feature is a cause of another feature is registered. Similarly, the *effects* box shows how the causality map is processed column-wise for each feature map. Before being summed element-wise, those two vectors are either passed as they are or the sign of their elements is reversed according to the causality direction d . The obtained vector is rectified and then returned as it is or passed through boolean filtering depending on the weighing mode m .

as they represent conditional asymmetries entailing possible cause-effect relationships (e.g., Fig. 1). Indeed, some features may be more often found on the left side of the arrow (i.e., $F \rightarrow$) than on the right side (i.e., $\rightarrow F$). Accordingly, the 2D causality map is processed row-wise and column-wise. In the former case, we register the number of times each feature map F^i was found to cause another feature map F^j , that is, $P(F^i|F^j) > P(F^j|F^i)$. This way, we obtain a vector of values that quantify how much those feature maps can be called "causes". Conversely, in the column-wise processing, we register the number of times each feature map F^j was found to be caused by another feature map F^i , obtaining a vector of values that quantify how much the feature maps can be deemed "effects".

At this point, we propose two variants to the model's functioning. We allow an external signal d to represent the causality direction of analysis, which can be either *causes* or *effects*. When $d = \text{causes}$, the vector of *causes* (obtained row-wise) is not altered, while the sign is changed to the elements of the *effects* vector (obtained column-wise). Hence, as those two vectors enter a summation point, the difference between *causes* and *effects* is obtained as the weight vector. On the other hand, when $d = \text{effects}$, the vector of *effects* is not altered, while it is to the vector of *causes* that the sign is changed. Therefore, the difference between *effects* and *causes* is obtained at the summation point. As a result, the obtained weight vector is rectified to set any negative elements to zero.

In addition, we conceive two variants of the model controlled by another external signal m , that represents the weighing mode and can be one of:

- **full.** The vector of non-negative causality factors is left at its full count, being returned as it is. As a result of this choice, the model weighs features more according to their causal importance (a feature that is *cause* 10 times more than another receives 10 times more weight).
- **bool.** The factors undergo boolean thresholding where all the non-zero factors are assigned a new weight of 1 and 0 otherwise. As a result, this choice is more conservative and assigns all features that are most often *causes* the same weight.

In the following sections, we describe the data used for our empirical evaluations, the different types of model architectures we utilized, and the implementation details of the training process.

3.2. Datasets

To validate our proposed methods, we utilized multiple publicly available medical imaging datasets. To begin with, we exploited the Breast cancer Histopathological Image (BreakHis) dataset (Spanhol, Oliveira, Petitjean, & Heutte, 2015). On the other hand, we used the dataset from the PI-CAI challenge (Saha et al., 2023), comprising multi-parametric MRI (mpMRI) acquisitions of the prostate.

3.2.1. BreakHis dataset

The dataset has 7909 microscopic images of breast tumor tissues aggregated from 82 subjects at magnification levels of 40, 100, 200, and 400. There are eight classes in this dataset, namely adenosis, tubular adenoma, fibroadenoma, phyllodes tumor, papillary carcinoma, lobular carcinoma, mucinous carcinoma, and ductal carcinoma. In addition, a binary classification was provided, namely, *benign* and *malignant* lesions. In particular, the first four classes represent benign lesions, while the last four represent malignant lesions. We considered the images with a magnification level of 400 for a total of 1819 images. We split this dataset into training (1235 images), validation (218 images), and test (366 images) sets, ensuring class balancing according to the binary classification, i.e., benign and malignant.

In this study, we utilize the processed version of the BreakHis dataset, curated by Pereira (2023) to be used in ML tasks. Indeed, the original images were resized to 224×224 pixels and organized according to binary and multiclass classification tasks. We further resize the images to a consistent 128×128 pixel matrix. Some samples from the utilized dataset are presented in Fig. 4.

3.2.2. PI-CAI dataset

From the available 1500 acquisitions, we only selected T2-weighted (T2w) images. Within this cohort of patients and respective scans, some cases did not have any tumors (i.e., they had no biopsy examination),




Fig. 4. Some benign and malignant samples from the utilized BreakHis dataset.

while others had cancer lesions. For each of the latter, the dataset contained biopsy reports expressing the severity as Gleason Score (GS). In anatomopathology, a GS of 1 to 5 is assigned to the two most common patterns in the biopsy specimen based on the cancer severity. The two grades are then added together to determine the GS, which can assume all the combinations of scores from “1 + 1” to “5 + 5”. Additionally, the dataset included the assigned GS’s group affiliation, defined by the International Society of Urological Pathology (ISUP) (Egevad, Delahunt, Srigley, & Samaratunga, 2016), ranging from 1 to 5, which provides the tumor severity information at a higher granularity level. In this study, we included both cancerous and no-tumor patients. From the former case, we only considered lesions with $GS \geq 3 + 4$ (ISUP ≥ 2) and selected only the slices containing lesions by exploiting the expert annotations of the disease provided in the dataset. For the latter case, we considered all the available slices. In the end, we obtained a total number of 4159 images (from 545 patients), with a balanced distribution over the two classes: 2079 tumor images vs. 2080 no-tumor images. To constitute our subsets, we divided the available images into training (2830), validation (515), and testing (814) subsets. During the splitting process, we ensured patient stratification (i.e., images of the same patient were grouped to prevent data leakage) and class balancing.

We utilized the provided whole prostate segmentation to extract the mask centroid for each slice. We then standardized the field of view (FOV) at 100 mm in both x (FOV_x) and y (FOV_y) directions to ensure consistency across all acquisitions and subsequently cropped each image based on this value around its centroid. To determine the number of rows (N_{rows}) and columns (N_{cols}) corresponding to the fixed FOV, we utilized the pixel spacing in millimeters along the x -axis (px) and the y -axis (py). The relationships used to derive the number of columns and rows are $N_{cols} = \frac{FOV_x}{px}$ and $N_{rows} = \frac{FOV_y}{py}$, respectively. Furthermore, we resized all the images to a uniform matrix size of 96×96 pixels to maintain consistent pixel counts. Finally, we performed image normalization using an in-volume method. That involved calculating the mean and standard deviation of all pixels within the volume acquisition and normalizing each image based on these values using the z-score technique. Some samples from the utilized dataset are presented in Fig. 5.

3.3. Architecture and training

For each dataset, we built different CNN models to automatically classify input images in the two classes according to their diagnosis labels under full supervision. As for the architectures, we used the popular ResNet18 as the backbone for all the causality-driven models. To handle images of different sizes in image recognition, many common architectures use an adaptive average pooling layer that outputs a 1×1

shape before the classifier. It does this by adjusting its parameters (such as kernel size, stride, and padding) based on the input size. However, this reduces the dimensionality of the feature maps and ignores their 2D structure, which is needed for finding causalities. Therefore, we replaced the *AdaptiveAvgPool2D* layer of the ResNet18 with an identity layer in our experiments.

As described in Section 3.1, we could integrate the information of the causality map into the CNN classification in different manners. In this work, we developed six types of models for each dataset and trained them to test the efficacy of the newly proposed *Mulcat* architectures on medical image classification, namely:

- **ResNet18.** This model is a regular ResNet18 architecture to serve as a baseline, where we replaced its *AdaptiveAvgPool2D* layer with an identity layer. See Fig. 2 (blue box) for a visual representation.
- **ResNet18 + Cat.** This is a ResNet18 model we modified to embed the causal information via concatenation as in Terziyan and Vitko (2023). See Fig. 2 (magenta box) for a visual representation.
- **ResNet18 + Mulcat (full, causes).** This variant exploits our *causality factors extractor* to obtain weights for the feature maps. In this model, we set the causality direction $\mathbf{d} = \text{causes}$ and the weighing mode $\mathbf{m} = \text{full}$.
- **ResNet18 + Mulcat (bool, causes).** It is similar to the previous, but we set the weighing mode to $\mathbf{m} = \text{bool}$.
- **ResNet18 + Mulcat (full, effects).** This variant turns the way the set of causality factors is obtained within our *causality factors extractor* by setting the causality direction $\mathbf{d} = \text{causes}$. We use $\mathbf{m} = \text{full}$ in this model.
- **ResNet18 + Mulcat (bool, effects).** It is analogous to the previous, but setting the weighing mode to $\mathbf{m} = \text{bool}$.

As shown in Fig. 2, the different types of models we investigated expose the classifier to a different number of input features. Therefore, the classifier is modified for each type according to the number of new neurons entering the fully-connected layer.

We optimized the way we computed the causality maps (using either the *Max* option (Eq. (1)) or the *Lehmer* (Eq. (2))) and, for the *Lehmer* option, we tried six different values of its parameter p : $[-100, -2, -1, 0, 1, 100]$. Consequently, for each dataset, we trained seven models for each of the five types of causality-driven models, resulting in 35 causality-driven models plus one baseline model. We provide the pseudo-code for the algorithms utilized to compute the causality maps and the proposed causality factors in Algorithm 1 and Algorithm 2, respectively.

Regarding the training phase, we utilized the cross-entropy loss as the criterion and Adam as the optimizer, as well as performed data




Fig. 5. Some no-tumor and cancerous samples from the utilized PI-CAI dataset.

Algorithm 1 Causality map computation

Require: $\mathbf{x} \leftarrow$ feature maps $[k,n,n]$, $CM \leftarrow$ computation method, $lehmer_p \leftarrow$ Lehmer mean power.

- 1: $\mathbf{x} \leftarrow \mathbf{x}/\max(\mathbf{x})$ \triangleright Normalize in range $[0 - 1]$ by max activation
- 2: $\mathbf{cmap} \leftarrow 0$ \triangleright Size of cmap: $[k,k]$
- 3: **if** $CM == Max$ **then** \triangleright Eq. (1)
- 4: $\mathbf{sumValues} \leftarrow \text{sum}(\text{flatten}(\mathbf{x}), \text{dim} = 1)$ \triangleright Sum values by feature
- 5: $\mathbf{maxValues} \leftarrow \max(\text{flatten}(\mathbf{x}), \text{dim} = 1)$ \triangleright Max values by feature
- 6: $\mathbf{prod} \leftarrow \text{outerProduct}(\mathbf{maxValues}, \mathbf{maxValues})$ \triangleright Numerator of Eq. (1) for each i,j
- 7: $\mathbf{cmap} \leftarrow \mathbf{prod}/\mathbf{sumValues}$
- 8: **else**
- 9: **if** $CM == Lehmer$ **then** \triangleright Eq. (2)
- 10: $\mathbf{x} \leftarrow \text{flatten}(\mathbf{x})$
- 11: $\mathbf{crossMat} \leftarrow \text{outerProduct}(\mathbf{x}, \mathbf{x})$ \triangleright Multiply features pairwise
- 12: $\mathbf{num_a} \leftarrow \text{sum}(\exp(\mathbf{crossMat}, \text{lehmer_p} + 1))$
- 13: $\mathbf{num_b} \leftarrow \text{sum}(\exp(\mathbf{crossMat}, \text{lehmer_p}))$
- 14: $\mathbf{num} \leftarrow \mathbf{num_a}/\mathbf{num_b}$ \triangleright Numerator of Eq. (2)
- 15: $\mathbf{den_a} \leftarrow \text{sum}(\exp(\mathbf{x}, \text{lehmer_p} + 1))$
- 16: $\mathbf{den_b} \leftarrow \text{sum}(\exp(\mathbf{x}, \text{lehmer_p}))$
- 17: $\mathbf{den} \leftarrow \mathbf{den_a}/\mathbf{den_b}$ \triangleright Denominator of Eq. (2)
- 18: $\mathbf{cmap} \leftarrow \mathbf{num}/\mathbf{den}$
- 19: **end if**
- 20: **end if**
- 21: **return** \mathbf{cmap}

augmentation (random horizontal flip) at training time. We trained the models for 200 epochs and set up a learning rate (LR) scheduler to decrease the LR during training. Specifically, the scheduler starts by multiplying the LR by 1.0 after the first epoch, and then this factor linearly decreases to 0.1 at epoch 200. As for models' hyperparameters, we investigated different values of initial LR (0.01 and 0.001) and of weight decay (0.01, 0.001, and 0.0001). Accordingly, for each dataset, we trained the 36 models for each of the six combinations of hyperparameters and chose the best-performing model on the validation set. To prevent our results from being biased due to the random processes of the algorithms, we repeated the entire analysis (216 experiments) four times with different starting seeds that govern the random processes of the scripts.

3.4. Quantitative evaluation

During training, we utilized the loss and accuracy obtained by the models on the validation set to track their evolution during epochs, selecting the best-performing one once the training phase ended. Then, we evaluated such selected models on the external never-before-seen test set and reported their accuracy value. This way, we obtain a quantitative metric to compare the baseline architecture, the *Cat* model, and our proposed *Mulcat* architectures.

Algorithm 2 Causality factors extraction

Require: $\mathbf{x} \leftarrow$ feature maps $[k,n,n]$, $\mathbf{cmap} \leftarrow$ causality map $[k,k]$, $d \leftarrow$ causality direction, $m \leftarrow$ weighing mode.

- 1: $\mathbf{triu} \leftarrow \text{triu}(\mathbf{cmap}, \text{diag} = 1)$ \triangleright Upper triangular matrix
- 2: $\mathbf{tril} \leftarrow \text{tril}(\mathbf{cmap}, \text{diag} = -1).transpose()$ \triangleright Lower triangular matrix
- 3: $\mathbf{bool_ij} = (\mathbf{tril} > \mathbf{triu}).transpose()$
- 4: $\mathbf{bool_ji} = (\mathbf{triu} > \mathbf{tril})$
- 5: $\mathbf{bool_matrix} = \mathbf{bool_ij} + \mathbf{bool_ji}$ \triangleright Sum of booleans (OR logic)
- 6: $\mathbf{by_col} = \text{sum}(\mathbf{bool_matrix}, \text{dim} = 1)$ \triangleright Obtain the *causes* view
- 7: $\mathbf{by_row} = \text{sum}(\mathbf{bool_matrix}, \text{dim} = 0)$ \triangleright Obtain the *effects* view
- 8: **if** $d == \text{causes}$ **then**
- 9: $\mathbf{factors} = \text{ReLU}(\mathbf{by_col} - \mathbf{by_row})$ \triangleright *causes* minus *effects*
- 10: **if** $m == \text{full}$ **then**
- 11: **return** $\mathbf{factors}$
- 12: **else**
- 13: **if** $m == \text{bool}$ **then**
- 14: $\mathbf{factors} = 1 * (\mathbf{factors} > 0)$ \triangleright Boolean threshold
- 15: **return** $\mathbf{factors}$
- 16: **end if**
- 17: **end if**
- 18: **else**
- 19: **if** $d == \text{effects}$ **then**
- 20: $\mathbf{factors} = \text{ReLU}(\mathbf{by_row} - \mathbf{by_col})$ \triangleright *effects* minus *causes*
- 21: **if** $m == \text{full}$ **then**
- 22: **return** $\mathbf{factors}$
- 23: **else**
- 24: **if** $m == \text{bool}$ **then**
- 25: $\mathbf{factors} = 1 * (\mathbf{factors} > 0)$ \triangleright Boolean thresholding
- 26: **return** $\mathbf{factors}$
- 27: **end if**
- 28: **end if**
- 29: **end if**
- 30: **end if**

Ablation studies remove or damage specific components in a controlled setting to investigate all possible outcomes of system failure, thus understanding the contribution of a component to the overall system. The **ResNet18** (baseline) models already act as the ablation models for the remaining five types of models. Nevertheless, we wanted to do more than solely remove components. To gauge the significance of the values contained in causality maps and causality factors, we performed an additional test where we distort (i.e., damage) their information. We call these partially ablated versions of the networks **damaged**. Concerning the **ResNet18 + Cat** option, the only contribution of the causality map to the classification resides in the flattened elements that are concatenated to the actual (flattened) feature maps. Therefore, a natural *damaged* network for such a setting would be to create a fictitious causality map filled with random probability values. We called this model the **ResNet18 + Damaged-Cat**. On the

other hand, when it comes to the **ResNet18 + Mulcat** option, the key functionality is to extract a vector of meaningful causality factors that serve as weights to the feature maps. Hence, we created the **ResNet18 + Damaged-Mulcat** model, where we modify that vector to weigh features randomly rather than based on a principled way. This model comes in two variants according to the possible values of the causality factors mode, \mathbf{m} . Indeed, when $\mathbf{m} = \text{full}$, the $1 \times k$ vector of causality factors (i.e., weights) is replaced with a random vector of the same size with integer values ranging from 0 (a feature map is never *cause* of another feature) to $k - 1$ (it is *cause* of every other feature). Whereas, when $\mathbf{m} = \text{bool}$, the values of the weights are randomly assigned to either 0 or 1. Since, in this setting, weights are hand-crafted, there is no need to consider the causality direction used; therefore, the *damaged* study we performed is valid for both $\mathbf{d} = \text{causes}$ and $\mathbf{d} = \text{effects}$.

To observe how the different architectures differ in terms of memory requirements, we track the size of the trained models (in megabytes) and the number of corresponding parameters (in millions). To compute the former, we do not want to rely on the file size of the saved models (e.g., *.pth* files from PyTorch), as the file might be compressed. In fact, we calculate the number of parameters and buffers, multiply them by the element size, and accumulate these numbers.

3.5. Qualitative evaluation

To further investigate the possible benefits of integrating causality into CNNs for medical image classification, we performed explainable AI (XAI) experiments on the best-performing model for each type. Specifically, we aimed to obtain class activation maps (CAM) for the networks' decisions in all six types of models in our investigation: **Baseline** model, **Cat** model, **Mulcat-full-causes** model, **Mulcat-bool-causes** model, **Mulcat-full-effects** model, and **Mulcat-bool-effects** model. In all these models we assume ResNet18 as the backbone. Since investigating the variability of the visual output when changing the XAI method used is outside the scope of our work, we chose the popular Grad-CAM method (Selvaraju et al., 2017), implemented in the *pytorch-grad-cam* library (Gildenblat & Contributors, 2021). For the same reason, we selected the last convolutional layer of our architectures as the target layer for which we computed the CAM and performed the analysis with standard parameters. A more systematic analysis would require investigating the CAM output on all layers of the CNN and optimizing the smoothing parameters.

To evaluate the quality and robustness of the produced CAMs, we followed the following criteria for the two datasets:

- **BreakHis dataset.** To differentiate benign from malignant tumors, pathologists examine breast tissues at different magnification levels. Specifically, at $400\times$ magnification level, as the one used for our experiments, they analyze cytological features, such as shape and size of the nuclei, hyperchromatic nuclei, mitotic cells, and prominent nuclei (Young, O'Dowd, & Woodford, 2013). To highlight cell nuclei, they employ Hematoxylin and Eosin stains, which make the nuclei appear dark purple or blue, while the other structures appear in shades of pink, red, and orange (He, Long, Antani, & Thoma, 2012). For these reasons, we considered good explanations, the ones that focus on regions containing the dark purple/blue structures assumed as the nuclei of the cells.
- **PI-CAI dataset.** Based on the classification task, we considered explanations focusing on discriminative regions of the MRI (e.g., prostate gland area) to be better. In contrast, we considered explanations focusing on other structures, such as the rectum, bladder, or lateral muscle bundles, to be of lower quality and robustness.

3.6. Additional experiments

We conducted additional experiments on two very common application fronts to further test the effectiveness of our method. On the one hand, we proved that our module is easy to fit into existing convolutional models using other forms of visual attention, thus creating synergy. On the other, we verified its functioning in low-data scenarios, extending its applicability to Few-Shot Learning (FSL) (Fei-Fei, Fergus, & Perona, 2006; Fink, 2004).

3.6.1. Integrating bottleneck attention modules

The bottleneck attention module (BAM) (Park, Woo, Lee, & Kweon, 2020; Woo, Park, Lee, & Kweon, 2018) is a popular attention-based mechanism that, given a feature map, learns the attention map along two factorized axes, *channel* and *spatial*, to strengthen the representational power of CNNs. We thus investigated the addition of our module to models that already leveraged BAM. We used the same backbone as above (ResNet18) and placed multiple BAMs located after its layers 1, 2, and 3, to build hierarchical attention. After training with the same strategy as the main study, we compared the performance of BAM-based regular models (**ResNet18 + BAM**), BAM-based models integrating the *Cat* method (**ResNet18 + BAM + Cat**), and BAM-based models that integrate our *Mulcat* module (**ResNet18 + BAM + Mulcat**).

3.6.2. Few-shot experiments

In addition to fully-supervised studies, in this paper, we extended our recent investigation into causality-driven one-shot learning (OSL) (Carloni, Pachetti, & Colantonio, 2023) to the new BreakHis dataset, to understand how our *Mulcat* methods worked under the shortage of annotated data in the medical imaging domain. To make the analyses consistent, we only considered the *causes* direction in our *Mulcat* models. We adopted the meta-learning strategy and formulated each task (i.e., episode) of the training process as an *N*-way *1-shot* classification problem, that is, to classify *N* classes using only *1* support image per class.

Regarding the PI-CAI dataset, we utilized a subset of the data containing lesions and the clinical question was tumor grading (i.e., predict aggressiveness). From a higher-level perspective than that of GS scores and ISUP groups, prostate lesions with $GS \leq 3 + 3$ (ISUP = 1) and with $GS = 3 + 4$ (ISUP = 2) are considered low-grade (LG) tumors, while those with $GS > 3 + 4$ (ISUP > 2) are high-grade (HG) tumors. We considered only lesions whose GS was $\geq 3 + 4$ (ISUP ≥ 2). As a result, we had eight classes of GS and four classes of ISUP in our dataset. The total number of images was 2049 (from 382 patients), which we divided into training (1611), validation (200), and testing (238) subsets, and resized to 128×128 . We experimented with two classification scenarios on this dataset. In the first scenario (**2-way**), the meta-training data are labeled to the four ISUP classes, and the model is meta-trained by randomly picking *two* of the four classes in each task while distinguishing between LG and HG lesions during meta-testing. In the second scenario (**4-way**), we label meta-training data on the GS, and the model randomly picks *four* of the eight GS classes in each task while distinguishing between four ISUP classes in meta-testing.

Regarding the BreakHis dataset, we used the same subsets as for the main study and considered two scenarios: in the **2-way** scenario, the meta-training is performed by randomly picking two of the eight classes of aggressiveness in each task, and meta-testing is done on the two high-level classes *benign-vs-malignant*; in the **4-way** scenario, the meta-training is done on four out of eight random classes and the meta-testing is performed on four most prevalent classes (i.e., ductal carcinoma, fibroadenoma, phyllodes tumor, and tubular adenoma).

To increase the models' robustness to different data selections, we performed 600 meta-training tasks, 600 meta-validation tasks, and 600 meta-testing tasks for each experiment. To cope with the dataset unbalancing, we employed the AUC margin loss (AUCM) (Yang, 2022) and the proximal epoch stochastic method (PESG) (Guo, Yan, Yuan, & Yang,

Table 1

Results of the main study and the ablation study for the best-performing models w.r.t the causality setting, the mode of computing the causality factors, and the direction used to encode the causality factors. We report accuracy results on the test set as the mean and standard deviation (in lower script) over four repetitions of the experiments with different seeds. The top half of the table refers to the BreakHis dataset, while the bottom half of it refers to the PI-CAI dataset.

Architecture	Causality factors mode	Causality direction	Test set accuracy [%]
BreakHis dataset (main study)			
ResNet18	-	-	88.48 _{0,77}
ResNet18 + Cat (Terziyan & Vitko, 2023)	-	-	85.77 _{3,16}
ResNet18 + Mulcat (ours)	Full	Causes	91.32 _{1,63}
	Bool	Causes	91.06 _{1,32}
	Full	Effects	90.65 _{1,14}
	Bool	Effects	91.59 _{0,83}
BreakHis dataset (ablation study)			
ResNet18 + Damaged-Cat	-	-	50.66 _{1,93}
ResNet18 + Damaged-Mulcat	Full	Causes/Effects	81.57 _{0,50}
	Bool	Causes/Effects	82.24 _{2,15}
PI-CAI dataset (main study)			
ResNet18	-	-	68.38 _{1,50}
ResNet18 + Cat (Terziyan & Vitko, 2023)	-	-	70.07 _{1,86}
ResNet18 + Mulcat (ours)	Full	Causes	71.82 _{1,00}
	Bool	Causes	70.31 _{1,63}
	Full	Effects	69.96 _{0,55}
	Bool	Effects	71.13 _{0,36}
PI-CAI dataset (ablation study)			
ResNet18 + Damaged-Cat	-	-	52.08 _{0,93}
ResNet18 + Damaged-Mulcat	Full	Causes/Effects	49.63 _{0,88}
	Bool	Causes/Effects	67.18 _{0,99}

2020), maximizing the Area Under the ROC curve (AUROC), which we used as our training and evaluation metric. Specifically, in 2-way experiments, we computed the binary AUROC, while we calculated the AUROC using the *One-vs-rest* setting in 4-way experiments. Moreover, we evaluated the binary classification performance of the 4-way models by computing the AUROC of one significant class versus all the rest (i.e., 1-vs-3): *malignant* (ductal carcinoma) versus *benign* (fibroadenoma, phyllodes tumor, tubular adenoma) for the BreakHis dataset, and *LG* (ISUP=2) versus *HG* (ISUP=3, ISUP=4, ISUP=5). As with the main study, we performed ablation studies by repeating the OSL experiments with the *damaged* version of our *Mulcat* method.

3.7. Implementation details

All the experiments in this study ran on an NVIDIA A100 40 GB Tensor Core of the AI@Edge cluster of our Institute. We used Python 3.8.15 and back-end libraries of PyTorch (version 1.13.0, cuda 11.1), together with other libraries such as scikit-learn 1.2.0, grad-cam 1.4.8, pydicom 2.3.1, and pillow 9.4.0. Docker version 20.10.11 (build dea9396) was installed in the machine. To make results reproducible for each battery of experiments, we set a common seed for the random sequence generator of all the random processes and PyTorch functions. We release the codebase for our framework at https://github.com/gianlucaroni/causality_conv_nets.

4. Results

4.1. Main study

Table 1 shows the results of our main study for both datasets. We report the accuracy metric of the best-performing models on the external test set as the mean and standard deviation over four repetitions of the experiments with different seeds. Regarding the **BreakHis dataset**, the baseline models (ResNet18) achieved an accuracy of 88.48, while

the competing method *Cat* performed worse than the baseline, with an accuracy of 85.77. On the other hand, our proposed **Mulcat** models, where the causality factors are ultimately computed in different ways depending on the mode *m* and the direction *d*, demonstrate higher performance than both previous choices. For instance, the *full-causes* models achieved 91.32 accuracy, while their *bool* version achieved an accuracy of 91.06. On the other hand, the *full-effects* and *bool-effects* models reached accuracies of 90.65 and 91.59, respectively. Regarding the **PI-CAI dataset**, while the baseline models achieved an accuracy of 68.38, embedding causality in different forms improved performance. For instance, when the causality map was used with the *Cat* version, the models achieved an accuracy of 70.07. As for the **Mulcat** models, they all ranked above baseline, with the *full-causes* models that achieved 71.82 accuracy and the *bool* version achieving an accuracy of 70.31. On the other hand, the *full-effects* models achieved 69.96 accuracy, while using the *bool* version led the models to reach an accuracy of 71.13.

4.2. Ablation study

Table 1 also shows the results of purposely damaging the information contained in the causality maps and causality factors. These partial ablation studies for the two datasets reveal that, on BreakHis data, the **Damaged-Cat** models obtained an accuracy of 50.66, and the **Damaged-Mulcat** models obtained accuracy values of 81.57 and 82.24 when using *full* and *bool* mode, respectively. As for the PI-CAI dataset, the **Damaged-Cat** models achieved an accuracy of 52.08, and the **Damaged-Mulcat** models obtained accuracy values of 49.63 and 67.18 when using *full* and *bool* mode, respectively.

4.3. Memory requirements

The size of the trained models and the number of corresponding parameters for both experiments are given in Table 2. While our **Mulcat** models increase memory demand by a negligible amount compared to their baseline counterparts (+0.8%), using **Cat** models results in an overhead of up to approximately +4.7%.

Table 2

Memory requirements in terms of model size (in megabytes) and number of parameters (in millions) of the ResNet18 baseline models, the ResNet18 backbone with the addition of the *Cat* option, and the ResNet18 backbone with our *Mulcat* module. Lower is better.

Architecture	Model size (MB) [↓]	Number of parameters ($\times 10^6$) [↓]
BreakHis dataset (image size: $3 \times 128 \times 128$)		
ResNet18	42.73	11.19
ResNet18 + <i>Cat</i> (Terziyan & Vitko, 2023)	44.73	11.72
ResNet18 + Mulcat (ours)	42.80	11.21
PI-CAI dataset (image size: $1 \times 96 \times 96$)		
ResNet18	42.68	11.18
ResNet18 + <i>Cat</i> (Terziyan & Vitko, 2023)	44.68	11.70
ResNet18 + Mulcat (ours)	42.72	11.19




Fig. 6. Visual assessment of class activation maps for different malignant cases of the BreakHis dataset. Each row represents a different biopsy slide, and the columns represent the Grad-CAM outputs for the baseline model, the *Cat* option, and all the proposed *Mulcat* variants. ResNet18 is assumed as the backbone architecture in all models.

4.4. XAI evaluations

In addition to the quantitative experiments, we obtained qualitative results for the six models for each dataset by comparison of their CAMs given the same input test images. As an example, Fig. 6 shows the results for some BreakHis malignant cases for which all the models yielded the same correct prediction. Rows regard different bioptic slides, while columns represent from left to right the original input image, the CAM of the baseline (non-causality-driven) model, the CAM of the *Cat* models, and the CAMs of our proposed *Mulcat* models with their specific settings (i.e., direction d and mode m).

Similarly, Figs. 7 and 8 show results for the PI-CAI dataset on different cancerous and no-tumor cases, respectively. Again, rows represent

several scans, while columns represent, from left to right, the original T2w input image and the CAMs of each configuration.

4.5. Integration with BAMs

In Table 3, we show the result of integrating our *Mulcat* module to convolutional networks that utilize BAM attention. Regarding the BreakHis dataset, the regular BAM-based models achieved an accuracy of 89.63, and utilizing a *Cat* option worsened performance (85.77). Conversely, almost all the models that integrated our *Mulcat* modules performed better with a maximum accuracy of 92.07. We obtained similar results for the PI-CAI dataset (max accuracy: 71.93).




Fig. 7. Visual assessment of class activation maps for **cancerous** cases of the PI-CAI dataset. Each row represents a different scan, and the columns represent the Grad-CAM outputs for the baseline model, the *Cat* option, and all the proposed *Mulcat* variants. ResNet18 is assumed as the backbone architecture in all models.

Table 3

The effect of integrating our *Mulcat* method to models that utilize BAM modules, compared to baseline and competing models. We report accuracy results on the test set as the mean and standard deviation (in lower script) over four repetitions of the experiments with different seeds. The top half of the table refers to the BreakHis dataset, while the bottom half of it refers to the PI-CAI dataset.

Architecture	Causality factors mode	Causality direction	Test set accuracy [1]
BreakHis dataset			
ResNet18 + BAM (Park et al., 2020)	–	–	89.63 _{1,15}
ResNet18 + BAM + Cat (Terziyan & Vitko, 2023)	–	–	85.77 _{1,72}
ResNet18 + BAM + Mulcat (ours)	Full	Causes	84.75 _{2,01}
	Bool	Causes	90.65 _{2,01}
	Full	Effects	91.87 _{1,15}
	Bool	Effects	92.07 _{0,86}
PI-CAI dataset			
ResNet18 + BAM (Park et al., 2020)	–	–	68.00 _{0,61}
ResNet18 + BAM + Cat (Terziyan & Vitko, 2023)	–	–	66.09 _{0,41}
ResNet18 + BAM + Mulcat (ours)	Full	Causes	67.75 _{1,99}
	Bool	Causes	71.25 _{0,87}
	Full	Effects	71.93 _{0,37}
	Bool	Effects	66.52 _{1,67}




Fig. 8. Visual assessment of class activation maps for **no-tumor** cases of the PI-CAI dataset. Each row represents a different scan, and the columns represent the Grad-CAM outputs for the baseline model, the *Cat* option, and all the proposed *Mulcat* variants. ResNet18 is assumed as the backbone architecture in all models.

4.6. One-shot tasks

The main results of our OSL analysis are reported in Table 4 for both datasets. We report all values as mean and standard deviation AUROC across all the 600 meta-test tasks. Regarding the BreakHis dataset, the 2-way experiment was where our *Mulcat* module improved the models the most. Indeed, while the baseline achieved 0.51 AUROC, we achieved up to 0.69 AUROC with the ResNet18+ *Mulcat-Bool*. In contrast, we found this improvement to be lower in the case of 4-way experiments. Table 4 also shows the results for *damaged Mulcat* models, which consistently performed worse than their *Mulcat* counterparts. Concerning the PI-CAI dataset, embedding our *Mulcat* module improved the models in all scenarios, with a more pronounced improvement in the 4-way 1-shot*, where the models are trained to distinguish four classes (ISUP 2 – 5), but the AUROC is computed between ISUP 2 versus rest.

5. Discussion

In this work, we presented a new method for automatically classifying medical images that uses weak causal signals in the image to model how the presence of a feature in one part of the image affects the appearance of another feature in a different part of the image. Our plug-and-play *Mulcat* module leverages causality maps in a new way and extracts multiplicative factors that eventually weight feature maps according to their causal influence in the scene. Our results seem to indicate that this lightweight, attention-inspired mechanism makes it possible to exploit weak causality signals in medical images to improve neural classifiers without any additional supervision signal.

In our main study, we assessed the effectiveness of our method under a fully-supervised learning scheme. In general, all the models obtained with our *Mulcat* implementation achieved higher performance than the *baseline* (ResNet18) on the test set with both datasets (see Table 1). This superiority ranged from a minimum of +2.31% to a maximum of +5.03%. On the other hand, utilizing the *Cat* option from Terziyan and Vitko (2023) resulted in worse performance than most of our *Mulcat* model and, with BreakHis data, even of the *baseline*. We found that most best-performing models used the *Lehmer* method to get the causality map. Nevertheless, this choice comes with the drawback of necessitating more memory than the *Max* method. We experimented with six different integer values for the parameter p to sample the range of possible values. A possible improvement would be to let the network itself learn the parameter p by back-propagation instead of giving it a fixed value beforehand.

To further confirm the numerical results of our studies, we conducted partially ablating studies on the actual influence of the causality factors on generating useful causality-driven feature maps. As anticipated, when we damage the causal weights by replacing them with random vectors, the accuracy of the final model is lower than its *main study* counterpart (see Table 1).

The *Damaged Cat* performed worse than the *None (baseline)* because the network was likely to be confused by the large number of random values that were concatenated to the actual extracted features. We expected that concatenating a random vector, not trained in back-propagation, would be worse than concatenating nothing at all. In the *Damaged Mulcat*, the weights multiplying the features maps are random, and, being untrained, they are re-computed at each iteration without

Table 4

Results of the best-performing models under One-Shot Learning (OSL) settings in terms of AUROC values across all the 600 meta-test tasks as the mean and standard deviation (in lower script). The top half of the table refers to the BreakHis dataset, while the bottom half of it refers to the PI-CAI dataset. *: Trained to distinguish four classes, but the AUROC is computed with *One-vs-rest*.

Architecture	Causality factors mode	2-way 1-shot [↑]	4-way 1-shot [↑]	4-way 1-shot* [↑]
BreakHis dataset (OSL study)				
ResNet18	–	0.51 _{0,15}	0.59 _{0,06}	0.58 _{0,20}
ResNet18 + Mulcat	Full	0.66 _{0,24}	0.59 _{0,08}	0.59 _{0,15}
	Bool	0.69 _{0,25}	0.57 _{0,09}	0.57 _{0,15}
BreakHis dataset (OSL ablation study)				
ResNet18 + Damaged-Mulcat	Full	0.51 _{0,12}	0.52 _{0,06}	0.52 _{0,12}
	Bool	0.66 _{0,25}	0.50 _{0,05}	0.56 _{0,13}
PI-CAI dataset (OSL study)				
ResNet18	–	0.54 _{0,14}	0.58 _{0,07}	0.59 _{0,12}
ResNet18 + Mulcat	Full	0.55 _{0,14}	0.61 _{0,07}	0.71 _{0,12}
	Bool	0.56 _{0,14}	0.61 _{0,07}	0.71 _{0,12}
PI-CAI dataset (OSL ablation study)				
ResNet18 + Damaged-Mulcat	Full	0.53 _{0,14}	0.55 _{0,06}	0.55 _{0,11}
	Bool	0.54 _{0,14}	0.57 _{0,07}	0.61 _{0,12}

any optimization from previous iterations. This results in scenarios where depending on the multiplication factors, irrelevant features are amplified while important ones are suppressed. When *Damaged Mulcat* is used with the *full* option, this behavior is more pronounced (random weights can have very high values, up to $k-1$), and performance is low (even lower than baseline for PI-CAI data) because the network assigns a lot of importance to these potentially incorrect features. In contrast, when the *bool* option is used, this behavior is mitigated (random weights have a maximum value of 1), so the degree of confusion of the network is reduced, and performance is higher. Experiencing reduced performance when the causality maps and factors are completely ablated or partially corrupted suggests that our module is computing something significant. This observation indicates that, even if weak, the causality signals learned during training assist the network to perform better.

Although we notice the improvement of our **Mulcat** models over the baseline from a quantitative point of view, it seems that the different combinations of mode **m** and direction **d** behaved roughly the same way. Thus, we wanted to investigate the potential benefits of our proposition on a different level. We deepened the analysis and found that significant differences can emerge on the XAI side, supporting the role of causality in explainability. On BreakHis data, the **Mulcat-full** and **Mulcat-bool** models manage to focus on regions densely populated by nuclei, which is what pathologists do at this magnification level. Instead, *Baseline* models often pay attention to lateral zones or regions less critical for the malignancy classification. The latter behavior is also observed in the *Cat* models, which frequently focus on lateral, small, and/or irrelevant portions for classification purposes (see Fig. 6). As for the PI-CAI dataset, where the field of view is larger and comprises many different anatomic structures other than the prostate, we noticed a trend that **Mulcat-full-causes** and **Mulcat-bool-causes** are consistently more focused on the discriminative parts of the image (e.g., prostate gland area). Conversely, the other options led to models that often looked at the rectum, bladder, or lateral muscle bundles (see Figs. 7 and 8). This fact confirms our hypothesis that using causes and not effects allows the network to obtain more faithful results.

Among the methods that exploit the information of causality maps, **Cat** proves to be one of the worst. That is evident both quantitatively and qualitatively. The reason for this behavior could be the considerable complexity added to the model to account for all the combinations of feature maps. In fact, on the classifier, the number of input neurons goes from $n \times n \times k$ to $n \times n \times k + k \times k$, which with high k results in thousands of additional connections (e.g., $512 \times 512 = 262144$ new

neurons for a ResNet18). The overhead induced by the *Cat* method is quantitatively confirmed by the memory requirements summarized in Table 2. Instead, our *Mulcat* method increases the memory demand by a negligible amount compared to the baseline, promoting it as a low-cost improvement of regular architectures.

We conducted additional experiments to further add evidence of improved performance through our method. First, we showed that our *Mulcat* module can be easily integrated into existing architectures, such as attention-based BAM networks, and can create synergy in improving performance (see Table 3). Indeed, utilizing *Mulcat* led to an increase of up to +2.72% over the regular BAM-ResNet18 and of +7.34% over the BAM-ResNet18 that used *Cat*, for the BreakHis dataset. Similarly, BAM-ResNet18 which utilized our *Mulcat* option on the PI-CAI dataset achieved up to +5.78% and +8.83% w.r.t regular BAM-ResNet18 and *Cat* BAM-ResNet18, respectively.

Second, to tend towards a more generalized demonstration of classification problems, it was interesting to understand how our *Mulcat* method worked in practical application situations, such as the shortage of annotated data in the medical imaging domain. Thus, we performed One-Shot Learning (OSL) experiments, both in 2-way and 4-way settings (see Table 4). Our findings suggest that using *Mulcat* can be an effective choice even in low-data scenarios. Indeed, when performing binary classification over the BreakHis data (i.e., *benign-vs-malignant*), ResNet18+ *Mulcat* achieved an accuracy up to +35.2% compared to the baseline, while this improvement was broadly reduced when performing 4-way experiments. On the other hand, it is on the 4-way scenarios with PI-CAI data that ResNet18+ *Mulcat* outperformed the baseline the most, with an increase of 20.3% accuracy for the 4-way 1-shot* setting.

One of the limitations of our work is that we used only ResNet18 as the backbone architecture to extract latent representations for the different implementations, although this is consistent with the pilot nature of our study. Moreover, we acknowledge that our methods consider potential causal relationships in pairs rather than among more than two features. That, of course, can lead to suboptimal results, given the impossibility of excluding confounders. In future experiments, we would be interested in extending the operation to more variables and devising variations inspired by the classic PC algorithms of the literature on causal discovery in tabular data (Spirites & Glymour, 1991).

There could be other directions to explore from our work, both on the application and architectural level. It would be interesting to draw inspiration from the multi-depth, visual attention (Jetley, Lord, Lee, & Torr, 2018; Schlemper et al., 2019; Yan, Kawahara, & Hamarneh,

2019), which extracts information from the convolutional encoder at different depths (local and global features). In this regard, one could extract the causality map from the internal layers of the network (not only from the last one). Our work could also be expanded by proposing new methods to combine causal information besides concatenation and weighting of feature maps and by experimenting with ensemble methods. Additionally, visualizations such as that in Fig. 1 suggest one could potentially conceive a low-cost self-supervised feature pruning based on similarities of features across rows and columns of causality maps. That could help to disregard redundant features and consequently lower model complexity in a data-driven way. In the end, we foresee a possible integration of our methods within the convolutional block of generative models such as GANs (Goodfellow et al., 2020) and diffusion models (Ho, Jain, & Abbeel, 2020; Rombach, Blattmann, Lorenz, Esser, & Ommer, 2022), to guide the generation of more realistic images.

6. Conclusions

In this work, we introduced a novel technique to discover and exploit weak causal signals directly from medical images via neural networks for classification purposes. Our method consists of a CNN backbone and a causality-factors extractor module, which computes weights for the feature maps to enhance each feature map according to its causal influence in the image's scene in an attention-inspired fashion. We developed different architecture variants and empirically evaluated all of our models on two public datasets of medical images for cancer diagnosis. Moreover, we verified that our module can create synergies when introduced in existing attention-based architectures, and we verified its applicability to few-shot learning settings.

Our findings demonstrate how minor modifications to traditional models can enhance them. Indeed, our lightweight module can be easily integrated into regular CNN classification systems and produce better models without requiring additional trainable parameters. It enhances the overall classification results and makes the model focus more precisely on the critical regions of the image, leading to more accurate and robust predictions. This aspect is crucial in medical imaging, where accurate and reliable classification is essential for effective diagnosis and treatment planning. Nevertheless, what we propose in this paper may have a broader significance, such as non-medical tasks or application to other data types, such as videos. We believe that the new elements we introduce with our work are a way to connect machine vision and causal reasoning in a novel way, especially when no prior knowledge of the data is available, adding a unique dimension to the framework.

CRedit authorship contribution statement

Gianluca Carloni: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization. **Sara Colantonio:** Resources, Writing – review & editing, Supervision, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

I have shared the link to the codebase for our framework in the 'Implementation details' section of our manuscript.

Acknowledgments

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No 952159 (ProCAncer-I), and partially from the Regional Project PAR FAS Tuscany - NAVIGATOR. The funders had no role in the design of the study, collection, analysis, and interpretation of data, or writing the manuscript.

References

- Berrevoets, J., Kacprzyk, K., Qian, Z., & van der Schaar, M. (2023). Causal learning. arXiv preprint arXiv:2303.02186.
- Bullen, P. S. (2003). vol. 560, *Handbook of means and their inequalities*. Springer Science & Business Media.
- Carloni, G., Pachetti, E., & Colantonio, S. (2023). Causality-driven one-shot learning for prostate cancer grading from MRI. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 2616–2624).
- Chickering, D. M. (2002). Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 3(Nov), 507–554.
- Egevad, L., Delahunt, B., Srigley, J. R., & Samaratunga, H. (2016). International Society of Urological Pathology (ISUP) grading of prostate cancer - An ISUP consensus on contemporary grading. *APMIS: Acta Pathologica, Microbiologica, Et Immunologica Scandinavica*, 124(6), 433–435.
- Fei-Fei, L., Fergus, R., & Perona, P. (2006). One-shot learning of object categories. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(4), 594–611.
- Fink, M. (2004). Object classification from a single example utilizing class relevance metrics. In *Advances in neural information processing systems: Vol. 17*.
- Gildenblat, J., & Contributors (2021). PyTorch library for CAM methods. GitHub. <https://github.com/jacobgil/pytorch-grad-cam>.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., et al. (2020). Generative adversarial networks. *Communications of the ACM*, 63(11), 139–144.
- Guo, Z., Yan, Y., Yuan, Z., & Yang, T. (2020). Fast objective & duality gap convergence for nonconvex-strongly-concave min-max problems. arXiv preprint arXiv:2006.06889.
- He, L., Long, L. R., Antani, S., & Thoma, G. R. (2012). Histology image analysis for carcinoma detection and grading. *Computer Methods and Programs in Biomedicine*, 107(3), 538–556.
- Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. In *Advances in neural information processing systems: vol. 33*, (pp. 6840–6851).
- Jetley, S., Lord, N. A., Lee, N., & Torr, P. H. (2018). Learn to pay attention. arXiv preprint arXiv:1804.02391.
- Kogan, E., Didden, E.-M., Lee, E., Nnewiwe, A., Stamatiadis, D., Mataraso, S., et al. (2023). A machine learning approach to identifying patients with pulmonary hypertension using real-world electronic health records. *International Journal of Cardiology*, 374, 95–99.
- Lopez-Paz, D., Nishihara, R., Chintala, S., Scholkopf, B., & Bottou, L. (2017). Discovering causal signals in images. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 6979–6987).
- Luo, Y., Peng, J., & Ma, J. (2020). When causal inference meets deep learning. *Nature Machine Intelligence*, 2(8), 426–427.
- Park, J., Woo, S., Lee, J.-Y., & Kweon, I. S. (2020). A simple and light-weight attention module for convolutional neural networks. *International Journal of Computer Vision*, 128(4), 783–798.
- Pearl, J. (1985). Bayesian networks: A model of self-activated memory for evidential reasoning. In *Proceedings of the 7th conference of the cognitive science society* (pp. 15–17). Irvine, CA, USA: University of California.
- Pearl, J. (2009). *Causality*. Cambridge University Press.
- Pearl, J., & Mackenzie, D. (2018). *The book of why: the new science of cause and effect*. Basic books.
- Pereira, M. (2023). BreakHis - Breast cancer histopathological database.
- Perry, R., Von Kügelgen, J., & Schölkopf, B. (2022). Causal discovery in heterogeneous environments under the sparse mechanism shift hypothesis. *Advances in Neural Information Processing Systems*, 35, 10904–10917.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 10684–10695).
- Saha, A., Bosma, J., Twilt, J., van Ginneken, B., Yakar, D., Elschoot, M., et al. (2023). Artificial intelligence and radiologists at prostate cancer detection in MRI—The PI-CAI challenge. In *Medical imaging with deep learning, short paper track*.
- Schlemper, J., Oktay, O., Schaap, M., Heinrich, M., Kainz, B., Glocker, B., et al. (2019). Attention gated networks: Learning to leverage salient regions in medical images. *Medical Image Analysis*, 53, 197–207.
- Schölkopf, B. (2022). Causality for machine learning. In *Probabilistic and causal inference: the works of Judea Pearl* (pp. 765–804). Association for Computing Machinery.
- Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., et al. (2021). Toward causal representation learning. *Proceedings of the IEEE*, 109(5), 612–634.

- Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision* (pp. 618–626).
- Spanhol, F. A., Oliveira, L. S., Petitjean, C., & Heutte, L. (2015). A dataset for breast cancer histopathological image classification. *IEEE Transactions on Biomedical Engineering*, 63(7), 1455–1462.
- Spirtes, P., & Glymour, C. (1991). An algorithm for fast recovery of sparse causal graphs. *Social Science Computer Review*, 9(1), 62–72.
- Spirtes, P., Glymour, C. N., & Scheines, R. (2000). *Causation, prediction, and search*. MIT Press.
- Su, Y., Tian, X., Gao, R., Guo, W., Chen, C., Chen, C., et al. (2022). Colon cancer diagnosis and staging classification based on machine learning and bioinformatics analysis. *Computers in Biology and Medicine*, 145, Article 105409.
- Terziyan, V., & Vitko, O. (2023). Causality-aware convolutional neural networks for advanced image classification and generation. *Procedia Computer Science*, 217, 495–506.
- Wang, Y., Menkovski, V., Wang, H., Du, X., & Pechenizkiy, M. (2020). Causal discovery from incomplete data: a deep learning approach. arXiv preprint arXiv:2001.05343.
- Woo, S., Park, J., Lee, J.-Y., & Kweon, I. S. (2018). Cbam: Convolutional block attention module. In *Proceedings of the European conference on computer vision* (pp. 3–19).
- Wroge, T. J., Özkanca, Y., Demiroglu, C., Si, D., Atkins, D. C., & Ghomi, R. H. (2018). Parkinson's disease diagnosis using machine learning and voice. In *2018 IEEE signal processing in medicine and biology symposium* (pp. 1–7). IEEE.
- Yan, Y., Kawahara, J., & Hamarneh, G. (2019). Melanoma recognition via visual attention. In *Information processing in medical imaging: 26th international conference, IPMI 2019, Hong kong, China, June 2–7, 2019, proceedings. Vol. 26* (pp. 793–804). Springer.
- Yang, T. (2022). Algorithmic foundation of deep x-risk optimization. arXiv preprint arXiv:2206.00439.
- Young, B., O'Dowd, G., & Woodford, P. (2013). *Wheatler's functional histology e-book: a text and colour atlas*. Elsevier Health Sciences.