



A framework for falsifiable explanations of machine learning models with an application in computational pathology

David Schuhmacher^{a,b,1}, Stephanie Schörner^{a,c,1}, Claus Küpper^{a,c}, Frederik Großerueschkamp^{a,c}, Carlo Sternemann^{a,d}, Celine Lugnier^{a,e}, Anna-Lena Kraeft^{a,e}, Hendrik Jütte^{a,d}, Andrea Tannapfel^{a,d}, Anke Reinacher-Schick^{a,e}, Klaus Gerwert^{a,c}, Axel Mosig^{a,b,*}

^a Ruhr-University Bochum, Center for Protein Diagnostics, Bochum, 44801, Germany

^b Ruhr-University Bochum, Faculty of Biology and Biotechnology, Bioinformatics Group, 44801 Bochum, Germany

^c Ruhr-University Bochum, Faculty of Biology and Biotechnology, Department of Biophysics, 44801 Bochum, Germany

^d Institute of Pathology, Ruhr-University Bochum, 44789 Bochum, Germany

^e Department of Hematology, Oncology and Palliative Care, Ruhr-University Bochum, St. Josef-Hospital, 44791 Bochum, Germany

ARTICLE INFO

MSC:
62M45
68T27
62P10
68U10

Keywords:

Explainable artificial intelligence
U-Net
Tumor segmentation
Falsifiability

ABSTRACT

In recent years, deep learning has been the key driver of breakthrough developments in computational pathology and other image based approaches that support medical diagnosis and treatment. The underlying neural networks as inherent black boxes lack transparency and are often accompanied by approaches to explain their output. However, formally defining explainability has been a notorious unsolved riddle. Here, we introduce a hypothesis-based framework for falsifiable explanations of machine learning models. A falsifiable explanation is a hypothesis that connects an intermediate space induced by the model with the sample from which the data originate. We instantiate this framework in a computational pathology setting using hyperspectral infrared microscopy. The intermediate space is an activation map, which is trained with an inductive bias to localize tumor. An explanation is constituted by hypothesizing that activation corresponds to tumor and associated structures, which we validate by histological staining as an independent secondary experiment.

1. Introduction

Classifications obtained from machine learning models in medical applications are in general black-box decisions, and interpreting their outcome is often considered key for making machine learning based diagnoses or treatment decisions transparent and trustworthy (Abels et al., 2019; Kelly et al., 2019; Holzinger et al., 2019; Rudin, 2019). The power of deep learning based classification has become particularly evident in image-based diagnostic settings, most notably in radiology and pathology (Litjens et al., 2017; van der Laak et al., 2021), where the need for interpretation has commonly been addressed by approaches such as linear relevance propagation (LRP) (Bach et al., 2015; Lapuschkin et al., 2019), saliency maps (Simonyan et al., 2013), Grad-CAM (Selvaraju et al., 2017) or class activation maps (Zhou et al., 2016). These methods operate in a post-hoc manner in the sense that

they operate on the neural network after it has been trained. Using the trained network along with an input image, they yield segmentations or heatmap representations (Zech et al., 2018; Hägele et al., 2020; Ding et al., 2019; Korbar et al., 2017; Bai et al., 2020; Hosny et al., 2018) that localize relevant patterns and can be assessed by domain experts. However, while unveiling an interpretable output, existing approaches lack a clear definition what constitutes a valid interpretation or explanation.

Ever since Lipton diagnosed the notion of interpretability to be ill-defined (Lipton, 2018), several attempts have been made to define interpretability and explainability in the context of machine learning models, as surveyed in Roscher et al. (2020), Guidotti et al. (2018) and Holzinger et al. (2019) and recently in Samek et al. (2021). However, it has been stated explicitly that no consensus definition has emerged, neither in a general sense (Guidotti et al., 2018) nor specific

* Corresponding author at: .

E-mail addresses: david.schuhmacher@rub.de (D. Schuhmacher), stephanie.schoerner@rub.de (S. Schörner), claus.kuepper-2@rub.de (C. Küpper), frederik.grosserueschkamp@rub.de (F. Großerueschkamp), carlo.sternemann@pathologie-bochum.de (C. Sternemann), Celine.Lugnier@rub.de (C. Lugnier), al.kraeft@gmail.com (A.-L. Kraeft), hendrik.juette@rub.de (H. Jütte), andrea.tannapfel@pathologie-bochum.de (A. Tannapfel), anke.reinacher@rub.de (A. Reinacher-Schick), Klaus.Gerwert@ruhr-uni-bochum.de (K. Gerwert), axel.mosig@rub.de (A. Mosig).

¹ These authors contributed equally.

to biomedical image analysis (van der Laak et al., 2021). Montavon et al. (2018) define an interpretation as a *mapping of an abstract concept into a domain that the human can make sense of*, while an explanation is *the collection of features of the interpretable domain, that have contributed for a given example to produce a decision*. The lack of a clear definition what an interpretation is has also been noted by Murdoch et al. (2019), who embed an interpretation in a so-called data-science life cycle and introduce the *predictive, descriptive, relevant* framework as a basis for interpretable machine learning. Despite these attempts towards definitions, apparently none has emerged as a consensus.

The challenge to define interpretability and explainability in the context of experimental research stems from the seemingly dichotomous approaches behind experimental research and machine learning. Experimental research, on the one hand, is hypothesis-centric by following hypothetico-deductivism, where predictions derived from a general and falsifiable hypothesis are applied to specific cases that either corroborate or falsify the hypothesis. Machine learning, on the other hand, derives general models from specific training examples and thus follows inductive (Wolpert, 1996) (or sometimes transductive (Vapnik, 2006)) reasoning. In machine learning, human expert-centric studies are common and legitimate practice in the sense that models are trained inductively on ground truth labels obtained from a human expert or a panel of human experts (Bulten et al., 2020; Abels et al., 2019; Amgad et al., 2019; Sirinukunwattana et al., 2017). In the context of deductive experimental research, however, it appears reasonable if not inevitable to define explainability in a hypothesis-centric manner rather than an expert-centric manner. The early and remarkable work by Michalski (1983) and related work by Mitchell et al. (1986) on concept learning deserves credit for assigning the central role to hypotheses and explicitly addressing the problem of induction. Yet, Michalski's approach of learning structural descriptions from examples lacks an explicit role of validating hypotheses in a deductive manner. The same problem holds for more recently proposed frameworks by Murdoch et al. (2019) and Roscher et al. (2020): Explanations as feedback cycles in these frameworks only refer back to data, and thus inevitably underly the limitations of inductive learning (Wolpert, 1996).

It has been well-established that weakly supervised learning, in particular multiple instance learning in the context of computational pathology, bears inherent interpretability (Lu et al., 2021), for example obtained through the use of attention mechanisms (Ilse et al., 2018; Shi et al., 2020b,a). In these works, datasets are given as bags of items, and interpretability is obtained by inferring labels at the bag level while only item levels are provided. In our present work, we abstract this property of weakly supervised learning to a general concept of an internal or intermediate space that is inferred by a machine learning system.

Here, we introduce an abstract and generic framework for outcome interpretation that is built upon a falsifiable hypothesis as the central constituent. The framework involves a machine learning model whose output yields, besides a classification outcome, an interpretable domain which is explicitly inferred during training of the model. The underlying hypothesis is formulated as a connection between the interpretable domain and the sample from which input data to the classifier have been obtained. Since the hypothesis refers to the sample rather than the input data, specific predictions about the sample can be derived from the hypothesis, which can thus be validated, i.e., corroborated or falsified, by additional experiments.

We demonstrate the application of our framework in a computational pathology setting dealing with colorectal tissue samples. Specifically, we introduce a weakly supervised deep convolutional neural network approach for classifying cancer and localizing tumor in microscopic images of tissue thin sections. We primarily work on hyperspectral infrared (IR) microscopic images which at each pixel position provide a location-specific biochemical fingerprint of the sample with a spatial resolution of around $5 \mu\text{m}$ (Kallenbach-Thieltges et al., 2013; Raulf et al., 2020). Our *comparative segmentation network* approach (CompSegNet) is trained on sample labels only, but during training infers a segmentation of tumor which facilitates a hypothesis-based interpretation of classification outcome.

Summary of contributions

Our main contribution is a framework for outcome explanation of machine learning models which, to the best of our knowledge, is the first such framework that overcomes the limitations of inherently inductive supervised machine learning by involving deductive, hypothesis-based reasoning. We instantiate this framework by a weakly supervised learning approach which demonstrates how domain knowledge is translated into hypothesis-guided inductive bias. The resulting comparative segmentation network (CompSegNet) extends a previous version for small tissue microarray spots to facilitate tumor segmentation in whole-slide tissue images. In this context, the involvement of domain knowledge leads to a novel, tissue image specific way of dealing with pathology specific background.

2. Theory and calculations

2.1. A formal definition of model explanation

The computational approaches introduced here emerge from a formal definition of model interpretation in machine learning that is illustrated in Fig. 1. We assume a given machine learning model f that classifies input $x \in X$ derived from a specific sample as label $y = f(x)$. In this setting, interpretation requires an extension of the machine learning model to infer additional interpretable output, so that $f(x) = (y, h)$ with $h \in I$. Here, I is an *interpretable* space that is inductively inferred while (or subsequent to) training f . Formally, the extension is thus a mapping $f: X \rightarrow Y \times I$, where X is the space of measured data, Y the space of supervised training labels, and I is what will henceforth be referred to as *interpretable space*, or I -space for short. This allows us to define an *explanation* of $f(x) = (y, h)$ as a falsifiable hypothesis that establishes a connection between $h \in I$ and the sample from which data x have been obtained. We will refer to hypotheses satisfying this definition as *falsifiable explanations* of the underlying machine learning model. The definition can be restricted further by requiring the classification y to be obtained purely on the grounds of $h \in I$. Formally, this means that classification is performed in two steps, where first $h = f(x)$ and subsequently $y = g(h)$ is computed. In this case, we refer to the space I as an *intermediary I-space*, and correspondingly define a hypothesis connecting h with the original sample as a *falsifiable intermediary explanation* of the classifier $g \circ f$. Intermediary explanations are obviously much stronger, since the classification can be derived from the interpretable intermediate variable, while the less restricted general falsifiable explanations allow h to be computed fully independent from y . Also, post-hoc interpretations (Bach et al., 2015; Simonyan et al., 2013; Zhou et al., 2016; Selvaraju et al., 2017) can be phrased in this framework. If the machine learning model is described by M in the space of all trainable models, its output can be denoted by $f_M(x)$. The post-hoc interpretable space for input x is then computed through a function $h = g(M, x)$, and defining $f(x) := (f_M(x), g(M, x))$ matches the framework displayed in Fig. 1.

There are crucial differences between falsifiable explanations and earlier attempts (Montavon et al., 2018; Murdoch et al., 2019; Roscher et al., 2020) to define explainability and interpretability. First and foremost, our approach refers back to the physical reality of the sample from which the data originate rather than just referring back to the data $x \in X$. Second, we explicitly require an explanation to be a falsifiable hypothesis. Earlier definitions of explainability commonly involve understandability by a human expert. Our definition shifts understandability towards the hypothesis, which is implicitly assumed to be understandable by a human expert. Compared to Michalski (1983), the hypothesis is phrased ex-silico and outside the formalism of statistical learning, but within the rigor of hypothetico-deductive reasoning (Popper, 1959). By introducing a hypothesis, and by making this hypothesis falsifiable through referring back to the sample rather than the data obtained from it, our framework involves deductive reasoning.

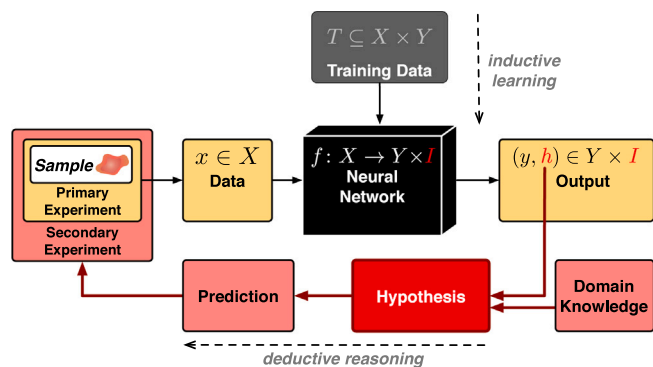


Fig. 1. Framework for a formal definition of machine learning model explanation. Conventionally, supervised training of a machine learning model f based on measurements obtained from data space X equipped with labels Y is an inductive step. Interpretation is facilitated by extending f to infer an additional output, so that $f(x) = (y, h)$, where $y \in Y$ is a label and $h \in I$ is additionally inferred output in an interpretable space I . A falsifiable explanation then is a hypothesis that connects h with the sample from which x originates.

To the best of our knowledge, our framework is the first approach that addresses the inherent problem of induction (Wolpert, 1996) behind supervised learning and, compared to related approaches (Michalski, 1983; Mitchell et al., 1986; Murdoch et al., 2019; Roscher et al., 2020) aims to overcome it by involving deductivism.

2.2. A neural network approach facilitating falsifiable explanation.

We present the *comparative segmentation network* (CompSegNet) as a weakly supervised neural network that yields an I-space through an activation map which is trained on image labels only. A preliminary version of the CompSegNet has been described earlier (Schuhmacher et al., 2020). Here, we extend the CompSegNet to a falsifiable explainable model in the context of whole-slide infrared microscopic images of histopathological tissue samples. In particular, we showcase how the CompSegNet yields an I-space and facilitates a falsifiable intermediary explanation of a classifier that distinguishes cancer samples from cancer-free samples.

As illustrated in Fig. 2, the CompSegNet is a U-Net (Ronneberger et al., 2015) extended by a pooling neuron which accumulates activation from the output layer of the U-Net. Applied to distinguishing cancerous samples from cancer-free samples, the loss function aims to maximize pooled activation in images labeled as cancerous, while minimizing pooled activation in images labeled as cancer-free. This loss function invokes an inductive bias for the U-Net to localize structures that are present in cancer samples only. In other words, we may *hypothesize* that the CompSegNet localizes tumor and associated structures. It must be emphasized that the I-space is the outcome of an inductive process and thus formally decoupled from epistemic concepts such as tumor. The concept of tumor is unmistakably characterizable on a molecular basis. The hypothesis connects the inductively obtained I-space with epistemic concepts from molecular reality, thus resolving the epistemological opacity (Durán and Jongasma, 2021) of the model.

The hypothesis that activation in I-space corresponds to tumor, henceforth referred to as *tumor-activation hypothesis* for short, is indeed falsifiable. In the given setting, images have been obtained from hyperspectral infrared microscopy, which is particularly amenable to further experimental characterization since the approach is staining-free, i.e., it does not involve chemical staining of the sample. Consequently, the sample is unaltered after infrared microscopic imaging and thus well accessible for secondary experiments. Specifically, as a secondary validation of our hypothesis, we performed conventional hematoxylin and eosin (H&E) staining, where tumor regions can be clearly identified, and the hypothesis can thus be validated.

2.3. Network architecture

Our network is essentially an extended U-Net architecture (Ronneberger et al., 2015) that, as displayed in Fig. 2, receives as input an image x_i and produces the activation map h_i as an I-space variable from which the output of the network function $f(x_i)$ is computed through a pooling neuron that performs a problem-specific piecewise linear pooling. The activation map h_i is represented by a layer that has the same spatial dimensions as the input image x_i , and each pixel neuron in h_i uses a standard sigmoid activation function. The pooling neuron performs average pooling, involving a problem-specific background mask and thresholds described in detail below. Beside sigmoid transfer function in the activation map and the extended optical dimensionality of the input, we employed the widely established topology of the U-Net. For the conventional H&E data, we use the U-net topology from Ronneberger et al. (2015) with an input dimensionality of $224 \times 224 \times 3$. In the experiments with the hyperspectral image data, the input size was adapted to 256×256 with 427 optical channels at each pixel, while the remaining layers of the network were unchanged.

We formally consider a microscopic image x involving d optical channels as a mapping $x: \Omega \times [1 : d] \rightarrow \mathbb{R}$ in a coordinate system Ω with $N \times N$ pixels. We consider conventional light microscopic H&E stained images involving $d = 3$ channels and hyperspectral infrared images where each pixel is represented by a spectrum involving $d = 427$ wavenumbers. In a complete dataset $\mathbf{X} = \{(x_1, y_1), \dots, (x_n, y_n)\}$, each image x_i is associated with a sample-level label y_i , with $y_i = 1$ indicating a diseased sample, and $y_i = 0$ indicating a healthy control sample. As a notational convention, we denote the optical spectrum at position $j \in \Omega$ by $x_i(j)$. The goal of training our network is to infer an activation map $h_i: \Omega \rightarrow [0, 1]$ that yields an I-space by localizing disease patterns in the diseased samples: If x_i is a cancer sample indicated by $y_i = 1$, the output map $h_i(j)$ should approach 1 precisely for those positions j where tumor or tumor-associated tissue is present, whereas at all tumor-free positions j in x_i , $x_i(j)$ should approach 0. Given a healthy control image x_i , all positions should receive an activation close to 0 in their activation map h_i .

The topology, parameter space and loss function of the network are explicitly tailored to accommodate modeling assumptions that adjust the inductive bias of the training process. One particular property of whole-slide-images to account for is sample background: On a microscopy slide with a tissue sample, the sample covers only a certain fraction of the complete slide (see Fig. 4), and all positions not covered by tissue should be treated as irrelevant in the loss function. To accommodate this nature of histopathological imaging data, positions in the sample that are not covered by tissue need to be identified and explicitly disregarded in certain parts of the loss function. To this end, we introduce a binary background mask B_i . We indicate $B_i(j) = 1$ if pixel j in image i is covered by tissue (regardless of whether the covering tissue is tumorous or tumor-free), and $B_i(j) = 0$ otherwise. For both infrared microscopic images and H&E stained images, B_i can be computed in a straightforward manner from x_i . For the infrared microscopic images, a threshold is calculated based on the histogram of the absorbance at 1656 cm^{-1} over the whole image using Otsu's binarization approach with the use of a gaussian filter of size 5. For H&E images, positions not covered by tissue appear bright white and can be similarly separated by an intensity threshold.

For the interpretation of network output, it is essential that the activation map $h_i(j)$ at each position j is bounded within the interval $[0, 1]$. To achieve this, a sigmoid activation function is used for each neuron in the neural network layer representing h . The activations of the neurons that constitute the activation map h_i are each weighted by $1/\#B_i^1$, with $\#B_i^1$ representing the number of foreground pixels in the mask B_i and then accumulated in a pooling neuron q ; positions masked as background are associated with zero activation. In this manner, the pooling neuron accumulates the ratio of tumor activation within tissue, rather than accumulating the ratio of tumor within pixels. A

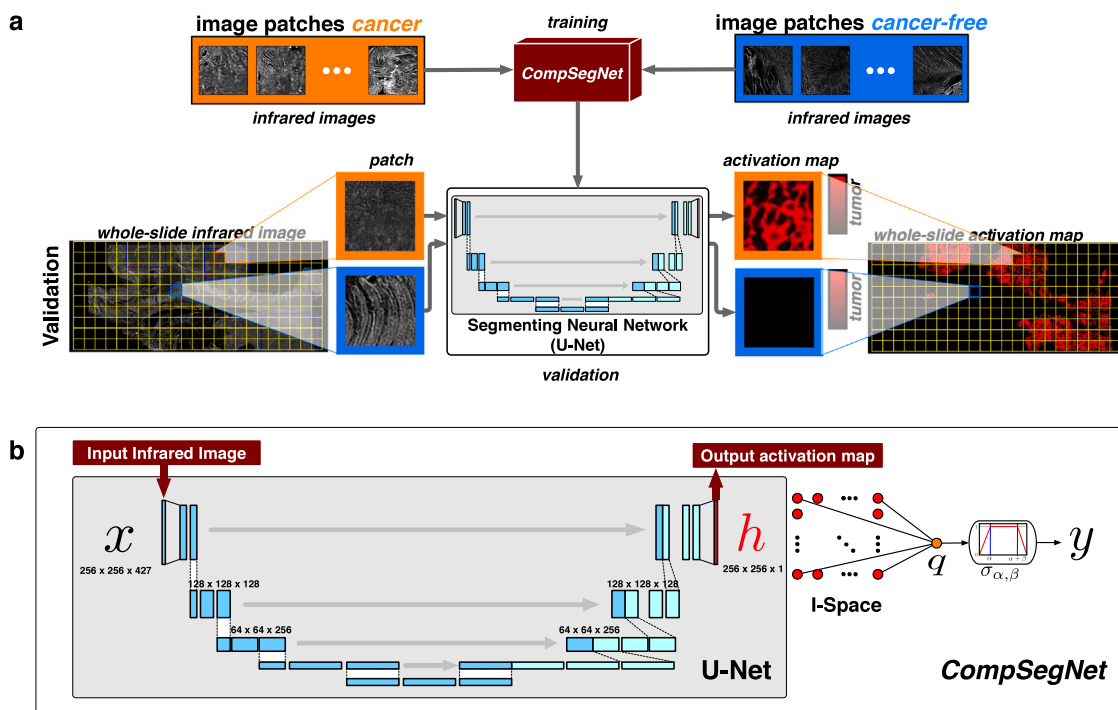


Fig. 2. Schematic overview of the Comparative Segmentation Network (CompSegNet). a The network is trained on sample labels only, but yields a U-Net (Ronneberger et al., 2015) whose output layer as an activation map facilitates segmentation. b The activation map is inferred with the help of a pooling neuron whose activation is maximized for cancer samples and minimized for control samples during training.

key modeling ingredient is the transfer function $\sigma_{\alpha,\beta}$ for the pooling neuron q . The transfer function is used to steer the network towards identifying a certain minimum and maximum fraction of tumor in input images labeled as cancer. At the same time, the transfer function needs to penalize the over-detection of cancer. Otherwise, without limiting over-detection, the network will escape towards the trivial solution and label each cancer sample completely as tumor. We translate these assumptions into a piecewise linear transfer function

$$\sigma_{\alpha,\beta}(x) = \begin{cases} x/\alpha & \text{if } x \leq \alpha \\ 1 & \text{if } \alpha < x \leq \alpha + \beta \\ \frac{(1-x)}{1-(\alpha+\beta)} & \text{if } x > \alpha + \beta, \end{cases} \quad (1)$$

which puts a maximum incentive for the network to identify a fraction of tumor that lies in between α and $\alpha + \beta$ where α and β are two parameters with an explicit modeling role. In our present networks, we used $\alpha = 0.05$ and $\beta = 0.9$ for the infrared image based CompSegNet and $\alpha = \beta = 0.4$ for the H&E image based CompSegNet. For training patches whose tumor ratio exceeds $\alpha + \beta$, this transfer function may lead to loss contributions that fluctuate between 0 and 1 during training. The loss function thus relies on the assumption that the number of cancer patches with close to 100% tumor is limited.

Through the pooling neuron q , a segmenting network becomes a binary classifier, whose classification performance is based on the segmentation resulting from the activation layer. Thus, in case of a successful classification by the pooling layer, the corresponding pixel-precise segmentation can be extracted.

The contribution L^i of image x_i to the loss function used is composed of two terms

$$L^i = L^i_{\text{class}} + \lambda L^i_{\text{bgr}} \quad (2)$$

coupled by a scaling parameter λ that was constantly set to 1 throughout all computations presented here. The first part L^i_{class} is a binary cross entropy with a class weighting. The class weights w_- and w_+ were determined by the relative prevalence of positive and negative samples in the dataset. For an input image x_i , we denote its given

disease label as y_i and the output of the network after the pooling neuron as $f(x_i)$. Then, we write the class-loss as

$$L^i_{\text{class}} = - (w_- \cdot y_i) \cdot \log(f(x_i)) - w_+ \cdot (1 - y_i) \cdot \log(1 - f(x_i)). \quad (3)$$

Activation of pixels that are either masked out as background through the mask B_i or belong to a negative cancer-free dataset is penalized by the term L^i_{bgr} through

$$L^i_{\text{bgr}} = - \frac{y_i}{\#B_i^0} \cdot \sum_{j \in \Omega} (1 - B_i(j)) \log(1 - h_i(j)) - \frac{1-y_i}{|\Omega|} \sum_{j \in \Omega} \log(1 - h_i(j)), \quad (4)$$

where $\#B_i^0$ denotes the number of background pixels in mask B_i . This background term also penalizes uninformative activation maps with global low activation. All networks were trained using a dropout rate of 20% in the hidden layers along with batch normalization and RMSprop as optimizer.

2.4. Activation map binarization

To binarize activation maps for computing Dice scores, we follow an approach that involves thresholds at two different levels. The primary threshold θ constitutes the actual activation threshold to binarize the activation map. A secondary threshold ρ is a tumor-fraction threshold such that a sample is classified as cancer whenever the ratio of tumor pixels exceeds ρ after binarization using θ . Now, each combination (θ, ρ) turns the activation layer of a CompSegNet-trained U-Net into a classifier whose performance can be measured on the validation set by computing the $F1_{\theta,\rho}$ -score. Since both θ and ρ are $[0, 1]$ -valued, we can determine $(\theta_{\text{opt}}, \rho_{\text{opt}}) = \arg \max_{\theta,\rho} F1_{\theta,\rho}$ using exhaustive search with a step size of 0.02 for θ and a step size of 0.01 for ρ .

Table 1

Patient characteristics for the infrared imaging dataset from the ColoPredict Plus 2.0 multicenter cohort. The total numbers are provided for the number of samples (N(s)), the number of measurements (N(m)) and the number of patches (N(patch)). The data provided for the characteristics are the numbers and the percentages corresponding to N(s). For the cancer-free samples only data about gender and age can be provided.

	Training Cohort		Validation Cohort		Test Cohort	
	Cancer	Cancer-free	Cancer	Cancer-free	Cancer	Cancer-free
Total number						
N(s)	50	50	25	25	25	25
N(m)	137	106	73	54	59	59
N(patch)	650	645	318	320	240	364
Gender						
female	32 (64%)	27 (54%)	15 (60%)	8 (32%)	14 (56%)	13 (52%)
male	18 (36%)	23 (46%)	10 (40%)	17 (68%)	11 (44%)	12 (48%)
Age in years						
mean ± SD	77 ± 11.9	68 ± 10.6	76 ± 13.2	72 ± 9.5	73 ± 10.4	71 ± 10.2
UICC Stage						
II	30 (60%)	n.a.	15 (60%)	n.a.	14 (56%)	n.a.
III	20 (40%)	n.a.	10 (40%)	n.a.	11 (44%)	n.a.
Localization						
left	18 (36%)	n.a.	7 (28%)	n.a.	11 (44%)	n.a.
right	32 (64%)	n.a.	18 (72%)	n.a.	14 (56%)	n.a.

3. Material and methods

3.1. Infrared imaging

Infrared imaging is a staining-free method for obtaining robust and observer-independent classifications of tissue thin sections on microscopic spectral images. Based on the molecular fingerprint in the spectral domain, supervised learning algorithms can be trained to identify disease-associated regions. It has been shown that this method allows a differentiation of tissue types with focus on cancer tissue sections (Kuepper et al., 2018). It also allows the biochemical differentiation of tumor tissue in terms of cancer subtyping and grading (Großerueschkamp et al., 2015; Kuepper et al., 2016). The original setup was based on Fourier transform infrared spectroscopy (FTIR), embedded in a microscope setup. Due to long acquisition times a transition to Quantum Cascade Laser (QCL) based imaging approaches has been made which reduced the measurement time from 90 h for the FTIR setup to approximately one hour for an average colon tissue thin section of 1.4 cm × 2 cm and yields high-quality spectral data (Kuepper et al., 2018). As a QCL based platform, the Spero QT IR microscope (Daylight Solution, San Diego, CA, USA) was utilized. The standard configuration covers the wavelength from 1800 cm⁻¹ to 950 cm⁻¹ with a spectral resolution of 2 cm⁻¹. The installed 4 × 0.3 NA objective was used with a coverage of 2 × 2 mm Field of View (FOV). The uncooled microbolometer focal plane array (FPA) detector consists of 480 × 480 pixels with a pixel size of 4.25 × 4.25 μm. During the measurement the sample chamber is constantly flooded with dry air.

3.2. Patients and datasets: IR imaging

The primary basis for our infrared microscopy based study is a collective of 100 samples of patients with colorectal carcinoma (UICC Stage II & III, older than 18 years) and 100 samples of cancer-free tissue samples, both randomly selected from the multicenter registry study ColoPredict Plus 2.0. The nation-wide study involves 188 clinical trial centers and collects samples retro- and prospectively. The 100 samples per group (cancer/cancer-free) were divided into 50% training data, 25% validation data and 25% independent test data, as shown in Table 1. The cohorts were balanced by gender, UICC Stage and localization with the whole cohort showing a tendency for a higher proportion of female patients, Stage II and a right-sided localization. For some samples, the dataset was extended by measurements on a second microscope of same type to include device variance in the

dataset, so that in total, the study is based on 200 samples and 488 IR measurements of 190 patients. For a detailed description of the cohort refer to Kallenbach-Thieltges et al. (2020). The subdivision into training, validation and test group was performed strictly at the patient level.

The samples were obtained during surgery, locally formalin-fixed and paraffin-embedded and subsequently treated according to the standard procedure at the Institute of Pathology. Samples were cut into 7 μm thick tissue thin sections following a previously established infrared microscopy workflow (Kuepper et al., 2018). Subsequent to infrared imaging, the tissue thin sections were stained by H&E following standard procedures in pathology. Based on the H&E stained tissue sections, from each staining-free IR measurement regions were manually selected that are slightly larger than the input patch size of the CompSegNet (≈1.1 mm). This was done on the IR overview image (integral of all wavenumbers) in comparison to the H&E staining by a pathologist. For the tumor-associated dataset, regions were chosen that contain a minimum of 20% tumor-related tissue in the 100 cancer samples. The tumor-free regions were selected from the measurements of tumor-free samples. Regions were chosen to cover a variety of different kinds of tumor-free tissue. The final dataset consists of cancer patches with patch-level label 1 and cancer-free patches with patch-level label 0 (class label). This class label is used for the class-loss according to Eq. (3).

For each patch, a binary mask is available representing pixels covered by tissue as 1 and pixels not covered by tissue (background) as 0. This binary mask is applied for the loss calculation according to Eq. (4). Examples for both cancer and cancer-free patches are shown in Supplementary Figure 1. As a data augmentation strategy, patches for training were obtained by randomly cropping patches matching the CompSegNet input size of 256 × 256 pixels. For further augmentation, the patches are flipped randomly after cropping.

To validate the tumor-activation hypothesis associated with the trained CompSegNet, the tissue sections from the test dataset were H&E stained as secondary experiment. In the resulting images, tumor-associated regions were annotated by a pathologist, yielding a binary mask with one class representing the tumor-associated tissue regions and the other class representing all tumor-free regions which includes tissue components such as crypts, lymphoid follicles, connective tissue, blood vessels and nerve tissue as well as muscle. The identification of tumor-associated regions was performed in a broad sense, including regions which may not contain tumor cells, but do not occur in cancer-free tissue sections, which is generally referred to as stroma. This includes, for example, the extracellular matrix, immune cells, infiltrated connective tissue and basement membranes.

3.3. Dataset: H&E images

For further validation of the CompSegNet approach, the H&E dataset from Kather et al. (2019a, 2018) was utilized for training and validation as a well-established dataset in computational pathology. It provides tissue samples from the National Center for Tumor diseases (NCT, Heidelberg, Germany) tissue bank and the University Medical Center Mannheim (UMM, Heidelberg University, Mannheim, Germany). The larger dataset provides 100,000 non-overlapping image patches from 86 H&E stained slides scanned by a conventional whole-slide scanning microscope and is referred to as *NCT-CRC-HE-100K*. Each image patch is 224 × 224 pixels in size at 0.5 μm per pixel and is assigned one out of nine different classes, with two classes associated with colorectal cancer (cancer-associated stroma and colorectal adenocarcinoma epithelium). The smaller dataset is referred to as *CRC-VAL-HE-7K* and consists of 7180 image patches, with the same size and resolution as the large dataset.

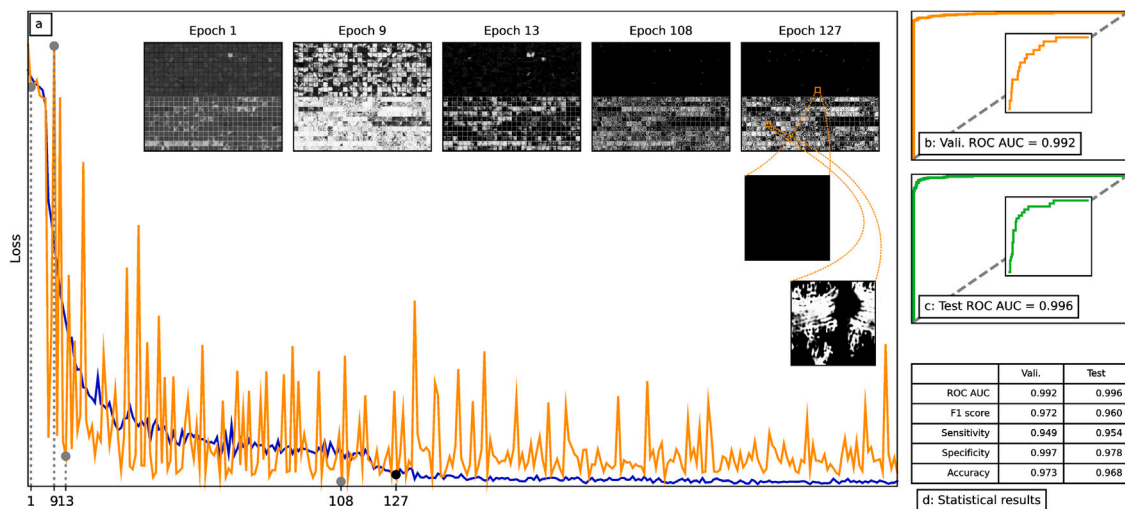


Fig. 3. Training curve, evolution in I-space and ROC curves of a CompSegNet trained on infrared images of colorectal histopathology tissue sections with patch size of 256×256 pixels. **a** Training progress for IR images with overview-images: The blue and orange curve show the training and validation loss over 300 epochs. Five epochs were chosen as examples to visually illustrate the course of training in the I-space for the validation data. The upper half of the five overview images displays activation within the tumor-free patches and the lower half within the tumor patches. **b-c** ROC curves and AUC values for validation ($N(\text{patch}) = 638$) and test dataset ($N(\text{patch}) = 604$) for the selected model from the 127th epoch on patch-level. **d** The table summarizes all statistical results on patch-level that were used for selecting the model.

4. Results

4.1. Choice of model-specific parameters

The CompSegNet architecture involves three hyper parameters, namely a lower bound α and an upper bound $\alpha + \beta$ representing minimum and maximum percentage of tumor within a single patch as well as the scaling parameter λ , introduced in Eq. (2), that weights the contribution of background loss to overall loss. To address the training stability against variations of those hyper-parameters, we evaluated training performance under varying combinations of α and β , as well as under varying λ .

For the tumor percentage thresholds, we tested all combinations of $\alpha \in \{.01, .05, .1\}$ and $\alpha + \beta \in \{.9, .95, .99\}$. With the given knowledge that cancer samples are selected to contain a minimum percentage of 20% and a maximum percentage or approx. 90%–95% tumor, these combinations represent a broad variation. Trained on the same cohort presented in Table 1, we did not observe major differences between different combinations when determining the maximal accuracy across 300 epochs of training. As shown in Supplementary Table 1, accuracies ranged between .9717 and .9921, where the highest accuracy was obtained for $\alpha = .01$ and $\alpha + \beta = .9$. We further assessed training stability for each combination by counting the number of epochs where accuracy exceeds .95. The results shown in Supplementary Table 1 show that for each combination of α and β , this was the case in 59 epochs on average, ranging between 23 epochs and 130 epochs. After inspection of segmentation results the model from epoch 127 with $\alpha = 0.05$ and $\alpha + \beta = .95$ was selected as best model for this analysis. Overall, our results show that, with a modest degree of relatively weak prior knowledge, a wide range of combinations of hyper-parameters α and β lead to a stable model.

For the background scaling parameter λ , we evaluated stability for $\lambda \in \{1., .66, .33, .0\}$. As shown in Supplementary Figure 2, we traced target loss, background loss and AUROC across 100 epochs of training. For $\lambda = 0$, training fails to exhibit any sign of convergence, even under a large range of learning rates. For any of the three non-zero values for λ , models converge stable with nearly equivalent training curves. These results show that the introduction of the background loss is essential, while the choice of the scaling parameter exhibits an essentially binary zero vs. non-zero characteristic.

4.2. Segmentation model classification performance

For the application of the CompSegNet to infrared data, the number of input channels was extended to 427 optical channels. The expected tumor fraction thresholds, as described in Section 2.3 and assessed in Section 4.1, were set to $\alpha = .05$ and $\beta = .9$, matching the expected rate of tumor content within an image patch and thus controls the inductive bias of the network based on modeling assumptions. The dropout rate was set to 20% along with a learning rate of $5 \cdot 10^{-4}$. The RMSprop parameters were set to $\rho = 10^{-6}$ with a momentum of 0. For the learning rate, a scheduler was implemented with a decay of 0.9 every 30 epochs along with a batch size of 4.

Fig. 3 illustrates the loss for the training and validation dataset over the 300 epochs of training. The first 6 epochs show a slowly dropping plateau phase. This is a consistent behavior of the CompSegNet and is explicitly guided by a background recognition term in the loss function. The length of the plateau phase is inversely proportional to the learning rate. In addition to the loss, a visual illustration has been added for five epochs over the course of training that shows an overview of the predictions of the validation patches in the I-space, where the evolution of the inferred I-space can be traced over training time. The evolution of the learned I-space can also be traced visually in the Supplementary Video. The final epoch (epoch 127) has been selected with regard to the area under receiver operating characteristic (AUC) values for the receiver operating characteristic (ROC), the F1 score, sensitivity, specificity and accuracy on the validation dataset. All statistical values are based on the patch-level classification. Validation on the independent test set exhibits similar performance as on the validation set as illustrated in Fig. 3, indicating a stable generalization performance.

4.3. Validating the tumor-activation hypothesis

In order to validate the tumor-activation hypothesis, we applied the trained CompSegNet to a set of three whole-slide images from the independent test set. The prediction for the whole-slides was performed by moving a 256×256 pixel window over each measurement with an overlap of 64 pixels while assigning maximum activation value in the overlapping areas. The resulting whole-slide activation map was binarized by a global activation threshold of $\theta = .628$ as described in Section 2.4. Using this activation threshold, the classification of the

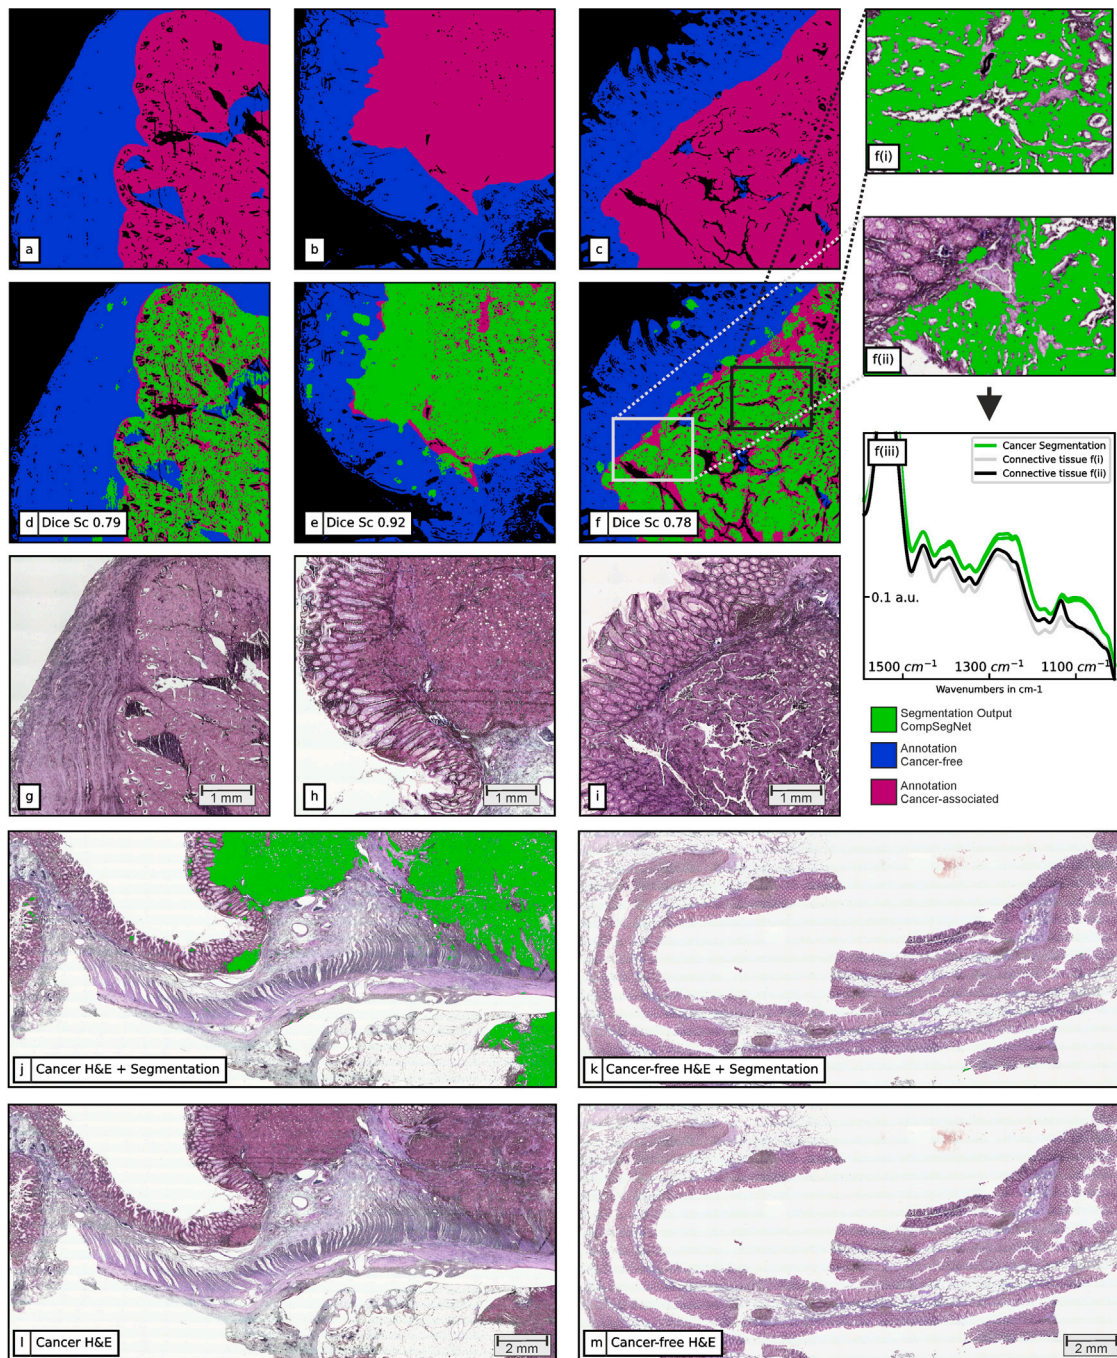


Fig. 4. Results for whole-slide segmentation on infrared measured samples from the test cohort of the ColoPredict Plus 2.0 study. a–c Ground truth annotation of H&E images with regard to tumor-free regions and tumor-related regions. Masked out background not covered by tissue is shown in black. d–f show the annotation from a–c overlaid with the segmentation output of the CompSegNet in green. f(i) and f(ii) show a magnification of f for two selected regions in which tumor-associated tissue was annotated, but no activation of the network appeared. Infrared spectra are plotted for cancer segmentation areas and highlighted non-activated areas. For interpretation of annotation and segmentation, g–h show the matching H&E stained sections for d–f. j–m illustrates two whole-slide H&E stained images, overlaid with the activation map of the infrared image (j: cancer-sample, k: cancer-free sample) and the corresponding H&E image without segmentation overlay (l: cancer-sample, m: cancer-free sample).

whole-slide images in the test set yields an ROC AUC value of 0.999. With a tumor-fraction threshold of $\rho = 5\%$, the sensitivity, specificity and accuracy are 100%, 98.2%, and 99.1%, respectively.

Agreement between the resulting binarized mask and tumor annotations was measured using Dice scores after transferring the H&E image and its annotation into the same coordinate system as the infrared image using the registration approach from Trukhan et al. (2020). Fig. 4 a–i illustrate the results on the whole sample for three thin tissue RoI from the test cohort. The Dice scores for all 25 whole slide images in

the test cohort is $.63 \pm .18$ indicating a good agreement between tumor-associated regions and the inferred I-space. The I-space is consistently a subset of the annotated regions. This may at least partially be explained by the inherent practical limitations of annotating the tumor-associated regions, which is complicated by the reduced H&E image quality that is inherent to the PET microscopy slides used for infrared microscopy. For the purpose of linking back the I-space to the sample, two subregions from Fig. 4f were selected that show examples of both agreement and disagreement between I-space activation and annotation. By looking at

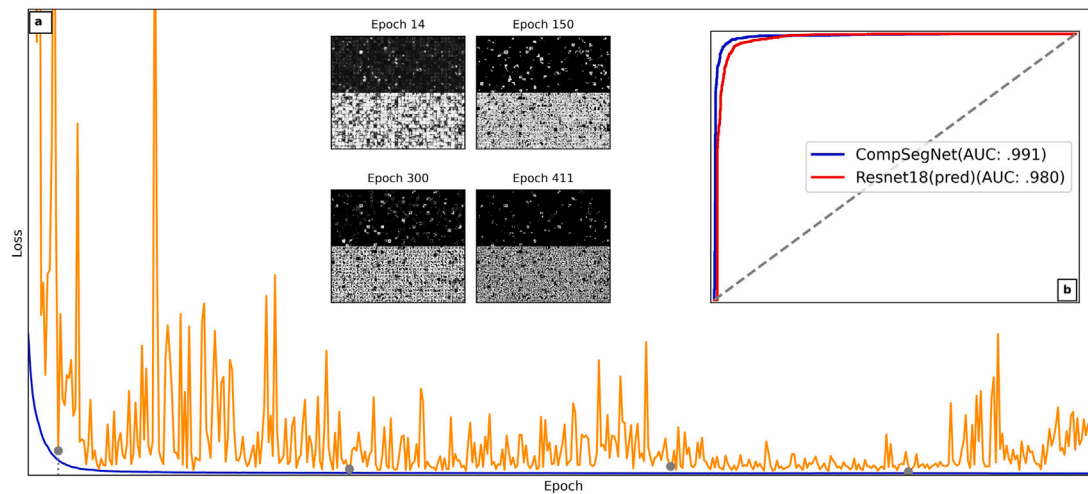


Fig. 5. Training curve, evolution in I-space and ROC curve of a CompSegNet trained on H&E stained images of colorectal histopathology image data. a Loss during training are shown in blue and orange for training set and validation set, respectively. Four epochs were chosen as examples to visually illustrate the course of training in the I-space for the validation data. The upper half of the four overview images displays activation within the tumor-free patches and the lower half within the tumor patches. b Comparison of ROC curves for the H&E CompSegNet and a ResNet18. The ROC curve was obtained based on the validation set from [Kather et al. \(2019a\)](#), using all 1200 tumor patches against 1200 patches randomly sampled from the remaining 8 classes as a class-balanced validation set. A visual snapshot displaying the activation maps of a random subset of the validation dataset is provided for four epochs. The ROC curve of the CompSegNet consistently exceeds the ROC curve of predictions through a ResNet-18.

the spectral property of the underlying tissue as shown in [Fig. 4f \(iii\)](#), a molecular interpretation becomes possible that closes the gap between I-space, hypothesis and the biochemical properties of the sample itself. For a high-level overview, two whole-slide H&E images are illustrated with the I-space activation in green in [Fig. 4j](#) and [k](#), showing one cancer and one cancer-free sample, respectively.

4.4. Tumor segmentation in conventional histopathology image data

In order to further validate the CompSegNet approach, we trained a network to segment colorectal carcinoma related structures using the H&E dataset from [Kather et al. \(2019b\)](#), as described in Section 3.3. As training data, the dataset provides 100,000 image patches from 86 H&E stained slides. Each image patch is 224×224 pixels in size and is assigned one out of nine different classes. In order to match the characteristics of the training data, we chose the expected tumor fraction parameters of the loss function as $\alpha = \beta = .4$, corresponding to an expected tumor content of 40%–80% in each patch labeled as tumor. The training dataset is complemented by a validation dataset with about 7000 images with the same parameters. For independent testing, we further supplemented the dataset with large subregions of ten further whole-slide images from the ColoPredict Plus 2.0 registry study summarized in [Table 1](#). Three examples are illustrated in [Fig. 6](#), results for all ten images are provided in Supplementary Figure 3. In the test samples, tumor and associated regions were annotated as ground truth. All H&E stained images were normalized using the approach by [Macenko et al. \(2009\)](#). Note that in the setting of applying the CompSegNet to H&E stained images, the H&E images themselves cannot be used as secondary experiment for validating the tumor activation hypothesis. Yet, in order to measure the agreement of inferred activation with tumor-associated tissue, we annotated tumor regions and measured the overlap with activation. Compared to the infrared-based activation maps, the H&E based activation maps illustrated in [Fig. 6](#) display over-detection in two out of three samples of tumor-related structure. In the third sample ([Fig. 6c](#)), one part of the tumor is not recognized, potentially due to tumor structures not represented in the training dataset.

We compared the I-space based classification of the CompSegNet with LRP as a reference for post-hoc explanation methods ([Hägele et al., 2020](#)). Conceptually, LRP exhibits a tendency towards satisfying the tumor-activation hypothesis because the relevance map is built on

the invariant that accumulated relevance equals network output ([Bach et al., 2015](#)). However, as a post-hoc approach, LRP lacks corresponding inductive bias during training. As detailed in Supplement A, we followed [Kather et al. \(2019b\)](#) and trained a ResNet-18 ([He et al., 2016](#)) on the same data as the CompSegNet, and subsequently performed LRP on the image patches from the validation set. The ResNet output was used to compute ROC curves on the validation set, which due to the relevance invariant is equivalent to a ROC curve over accumulated pixel relevance. The ResNet-18 exhibits slightly but consistently worse performance than the CompSegNet, indicating an effect of the inductive bias on optimizing I-space during training of the CompSegNet. The same trend is observed on the independent test set, where the LRP based segmentation achieves an average Dice score of .55, compared to an average of .66 obtained by the CompSegNet (Supplementary Figure 3).

As to be expected from the use of sigmoidal as well as piecewise linear transfer functions, convergence of the CompSegNet ([Fig. 5a](#)) is less smooth compared to the ResNet-18. With the rapid forward-calculations of the U-Net underlying the CompSegNet, the evolution of the learned I-space can be traced visually during training without much computational effort (Supplementary Video). For the backward propagation underlying LRP, this is impractical due to the much higher demand in computational resources (Supplement A).

5. Discussion

Several aspects need to be pointed out regarding our definition of a machine learning model explanation. First, it is crucial that the hypothesis refers back to the original ex-silico sample rather than the already observed in-silico data x . By referring back to the sample, the hypothesis can (and in general will) facilitate a prediction about the outcome of a secondary type of experiment performed on the same sample, rather than a prediction about the experiment that has been conducted to obtain x . Second, the definition introduces a new layer of validating a machine learning model. In addition to validating at the level of training labels, the explanation of the machine learning model can be validated by testing the stated hypothesis. In particular, the model may be invalidated by falsifying the hypothesis. Third, the definition unifies inductive machine learning with the hypotheto-deductivism of the scientific method, since the inductive task of supervised training is complemented by deductive reasoning:




Fig. 6. Segmentations of H&E stained images from the fully independent ColoPredict Plus 2.0 study test set. a–c Ground truth annotation of tumor and associated structures. Following Fig. 4, annotated tumor-associated regions are shown in purple, other tissue in blue. d–f Ground truth overlaid with binarized segmentation map (green). g–i H&E staining images underlying the annotation and the activation map.

Pure supervised training merely interpolates between data points, and the inductive bias of a machine learning method determines how to interpolate. Our approach introduces I-space, which is also subject to interpolation, but the arbitrariness of interpolation is blocked by a falsifiable hypothesis and predictions derived from it. In this sense, phrasing inductive bias through I-space can be thought of and used as a modeling tool that controls how the model interpolates. The deductive ex-silico validation level legitimates to train a model that optimizes towards certain features of I-space representations. In other words, I-space can be used to constrain and control the inductive bias of training f .

Our definition links to existing concepts and approaches of explainable machine learning. In the taxonomy introduced in Guidotti et al. (2018), interpretations matching our definition are outcome explanations since they explain the decision on a given object in contrast to model explanations which explain the internal logic of a machine learning model independent of a specific input. In general, any of the numerous and widely used model explanation approaches proposed in literature such as linear relevance propagation (Bach et al., 2015), saliency maps (Simonyan et al., 2013) or class activation maps (Zhou et al., 2016; Selvaraju et al., 2017) are potential sources of an I-space. These methods usually infer the computation of an input-specific interpretation in a post-hoc manner after a black box neural network has been trained. While such inferred interpretations may support a

hypothesis that withstands experimental validation, post-hoc methods of output explanation do not exploit the potential to optimize I-space inference towards a specific hypothesis by modeling inductive bias during training. It is important to notice that our CompSegNet approach yields a readily trained U-Net, which can process input using fast parallel forward calculations. This provides a practical and crucial advantage over most post-hoc methods, which require backpropagation steps that are more difficult to parallelize.

Some existing machine learning models implicitly match parts of our definition of falsifiable explanations. For example, applying multiple instance learning in histopathology (Campanella et al., 2019; Lu et al., 2021; Ilse et al., 2018; Li et al., 2019) tiles whole-slide samples into a large number of patches, and the key task is to identify the few disease-relevant patches among this large number. The few positive instances claim to localize disease, which yields a basis for a falsifiable hypothesis. In other words, multiple instance learning in histopathology yields an I-space at the level of patches. In this view of multiple instance learning, our CompSegNet approach is a weakly supervised learning approach that provides an I-space in the same coordinate system as the input image. Both the CompSegNet approach and the variants of multiple instance learning used in computational pathology (Campanella et al., 2019) involve different types of target functions that are tailored towards localizing disease under certain side assumptions. This optimization under side assumptions indeed

introduces the desired inductive bias while training the models. In more general terms, certain types of weakly supervised learning may provide a basis for falsifiable explanations.

A side-effect of directing inductive bias through modeling assumptions in the loss function is the introduction of hyperparameters. In case of the CompSegNet, the hyperparameters α and β possess a direct interpretation of a lower and upper bound of tumor fraction within a single patch and can thus be chosen by domain understanding and knowledge about the training data. Hyperparameters also occur in related weakly supervised learning methods, for example in the prominent and recently proposed CLAM system (Lu et al., 2021), where three hyperparameters balance certain contributions to the loss function. Similar to our results, the authors of CLAM observe robustness against changes in these hyperparameters. While introducing inductive bias through hyperparameters is inherently in danger of negatively affecting classification or generalization performance, our results as well as the results from CLAM indicate that this is not necessarily the case. As a general perspective, introducing explicit, and ideally hypothesis guided, inductive bias may be particularly useful when the amount of training data is inherently limited. This may be particularly fruitful in the context of clinical studies, where patient numbers are inevitably limited.

It is noteworthy that experimental validation is an inherent part of falsifiable explanations. It facilitates formal inclusion throughout all phases of clinical studies. This appears relevant since on the one hand, explainable models are often considered an important corner stone for trustworthy machine learning in medical applications (Abels et al., 2019; Kelly et al., 2019), while on the other hand no formal definition of what an explanation is has been established to date. The question arises whether explanations that are not connected to falsifiable hypotheses can or should formally be part of clinical studies. The necessity of interpretation has recently been argued by Durán and Jongsmá (2021), who identify epistemic opaqueness as the primary reason to distrust black box models. Since falsifiable explanations connect domain knowledge with machine learning models within the rigor of the scientific method, our proposed framework does establish epistemic transparency of the underlying machine learning model.

Neither the input space nor the I-space is limited in any way to image data, and the concept of falsifiable explanations carries beyond biomedical settings. While the input space is constituted by a given experimental setup, the only constraint on I-space is that it needs to connect to the hypothesis that constitutes the actual interpretation. It is important to realize that the hypothesis may and in general should involve further domain knowledge about the underlying sample. For example, the tumor-activation hypothesis involves knowledge about tumorigenesis. The same holds for the task of designing a secondary experiment. Examining the hypothesis as in our case by performing H&E staining on the sample involves background knowledge that tumor can be recognized by densely packed nuclei of certain morphology in particular regions of the sample. Other experiments may be imaginable for the same hypothesis, e.g. dissecting the areas with high activation using laser capture microdissection (Großerueschkamp et al., 2017; Selbach et al., 2021) and subsequently performing genome sequencing on the dissected sample material. A possible prediction derivable from the hypothesis would be to observe an accumulation of somatic mutations in the dissected areas compared to areas with no activation in I-space. In order to obtain robust molecular characterizations from laser microdissection, a highly precise localization of disease patterns is required. Inferring semantic segmentation at full input resolution based on coarse tile-level labels is indeed the major motivation behind our CompSegNet approach. In this sense, the approach differs from multiple instance learning approaches which infer tile-labels from whole-slide image labels.

Prospectively, a hypothesis-centric approach based on our framework could associate an inferred I-space with hypotheses regarding

molecular mechanisms behind specific tumor-subtypes. Such hypothesis could for instance claim the localization of either microsatellite instable or mismatch repair deficient regions within tissue samples (Li et al., 2020), which could be assessed deductively by combining laser microdissection with genomics or transcriptomics as a secondary experiment.

6. Conclusions

In summary, we have introduced a framework for falsifiable explanations of machine learning models which establishes a hypothesis rather than an expert as the major constituent of an explanation. We introduce the concept of I-space as the link between hypothesis and machine learning model. As we demonstrate, a suitable I-space can be obtained by either an explicit inductive bias during training or by post-hoc methods. We showcase these concepts by establishing a tumor-activation hypothesis for deep neural networks that identify cancer in histopathological imaging data.

Our comparison between the weakly supervised CompSegNet approach and LRP as a post-hoc approach demonstrates that in general both post-hoc and inductive bias based approaches are feasible approaches to infer a relevant I-space. Although our validations are limited to two types of histopathologic imaging data, the CompSegNet approach naturally translates to a variety of other image modalities, in particular if controlling the inductive bias through the design of dedicated loss functions is considered as a modeling step. Conceptually, the CompSegNet approach aims to overcome post-hoc machine learning models and takes a step towards models that are more, although certainly far from fully, intrinsically interpretable. Combined with the concept of falsifiable explanations, the method provides a framework for characterizing underlying or unidentified molecular changes or patterns.

Code availability

Source code will be made publicly available after publication.

CRedit authorship contribution statement

David Schuhmacher: Methodology, Software, Validation, Formal analysis, Writing – original draft, Visualization. **Stephanie Schörner:** Software, Validation, Formal analysis, Investigation, Data curation, Visualization. **Claus Küpper:** Investigation, Writing – review & editing. **Frederik Großerueschkamp:** Writing – review & editing, Supervision. **Carlo Sternemann:** Data curation. **Celine Lugnier:** Resources. **Anna-Lena Kraeft:** Resources. **Hendrik Jütte:** Data curation. **Andrea Tannapfel:** Resources, Data curation, Writing – review & editing, Project administration, Funding acquisition. **Anke Reinacher-Schick:** Resources, Writing – review & editing, Project administration, Funding acquisition. **Klaus Gerwert:** Writing – review & editing, Supervision, Funding acquisition. **Axel Mosig:** Conceptualization, Methodology, Writing – original draft, Visualization, Supervision, 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

The spectral and medical data that support the findings of this study are available from the corresponding author AM upon reasonable request.

Acknowledgments

The Center for Protein Diagnostics (PRODI) is funded by the Ministry of Culture and Science (MKW) of the State of North-Rhine Westphalia, Germany (grant number: 111.08.03.05-133974). Part of this research was conducted within the Slide2Mol project funded by the Computational Life Sciences program of the German Federal Ministry of Education and Research, grant number 031L0264.

Ethics statement

All methods and experimental protocols were approved by the relevant institutional review boards (registration numbers 4453-12, 20-6830 and 17-6151, Ethics Commission, Faculty of Medicine, Ruhr-University Bochum). Furthermore, all enrolled patients provided informed consent that the samples may be used for retrospective analysis. All procedures were conducted in accordance with the approved guidelines and regulations for human experimental research.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.media.2022.102594>.

References

- Abels, E., Pantanowitz, L., Aeffner, F., Zarella, M.D., van der Laak, J., Bui, M.M., Vemuri, V.N., Parwani, A.V., Gibbs, J., Agosto-Arroyo, E., et al., 2019. Computational pathology definitions, best practices, and recommendations for regulatory guidance: a white paper from the digital pathology association. *J. Pathol.* 249 (3), 286–294. <http://dx.doi.org/10.1002/path.5331>.
- Amgad, M., Elfandy, H., Hussein, H., Attveya, L.A., Elsebaie, M.A., Abo Elnasr, L.S., Sakr, R.A., Salem, H.S., Ismail, A.F., Saad, A.M., et al., 2019. Structured crowd-sourcing enables convolutional segmentation of histology images. *Bioinformatics* 35 (18), 3461–3467. <http://dx.doi.org/10.1093/bioinformatics/btz083>.
- Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.-R., Samek, W., 2015. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS One* 10 (7), e0130140. <http://dx.doi.org/10.1371/journal.pone.0130140>.
- Bai, H.X., Wang, R., Xiong, Z., Hsieh, B., Chang, K., Halsey, K., Tran, T.M.L., Choi, J.W., Wang, D.-C., Shi, L.-B., et al., 2020. Artificial intelligence augmentation of radiologist performance in distinguishing COVID-19 from pneumonia of other origin at chest CT. *Radiology* 296 (3), E156–E165. <http://dx.doi.org/10.1148/radiol.2020201491>.
- Bulten, W., Pinckaers, H., van Boven, H., Vink, R., de Bel, T., van Ginneken, B., van der Laak, J., Hulsbergen-van de Kaa, C., Litjens, G., 2020. Automated deep-learning system for Gleason grading of prostate cancer using biopsies: A diagnostic study. *Lancet Oncol.* 21 (2), 233–241. [http://dx.doi.org/10.1016/S1470-2045\(19\)30739-9](http://dx.doi.org/10.1016/S1470-2045(19)30739-9).
- Campanella, G., Hanna, M.G., Geneslaw, L., Mirafior, A., Silva, V.W.K., Busam, K.J., Brogi, E., Reuter, V.E., Klimstra, D.S., Fuchs, T.J., 2019. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nature Med.* 25 (8), 1301–1309. <http://dx.doi.org/10.1038/s41591-019-0508-1>.
- Ding, Y., Sohn, J.H., Kawczynski, M.G., Trivedi, H., Harnish, R., Jenkins, N.W., Lituiev, D., Copeland, T.P., Aboian, M.S., Mari Aparici, C., et al., 2019. A deep learning model to predict a diagnosis of Alzheimer disease by using 18F-FDG PET of the brain. *Radiology* 290 (2), 456–464. <http://dx.doi.org/10.1148/radiol.2018180958>.
- Durán, J.M., Jongasma, K.R., 2021. Who is afraid of black box algorithms? On the epistemological and ethical basis of trust in medical AI. *J. Med. Ethics* 47 (5), 329–335. <http://dx.doi.org/10.1136/medethics-2020-106820>.
- Großerueschkamp, F., Bracht, T., Diehl, H.C., Kuepper, C., Ahrens, M., Kallenbach-Thieltges, A., Mosig, A., Eisenacher, M., Marcus, K., Behrens, T., et al., 2017. Spatial and molecular resolution of diffuse malignant mesothelioma heterogeneity by integrating label-free FTIR imaging, laser capture microdissection and proteomics. *Sci. Rep.* 7 (1), 1–12. <http://dx.doi.org/10.1038/srep44829>.
- Großerueschkamp, F., Kallenbach-Thieltges, A., Behrens, T., Brüning, T., Altmeier, M., Stamatidis, G., Theegarten, D., Gerwert, K., 2015. Marker-free automated histopathological annotation of lung tumour subtypes by FTIR imaging. *Analyst* 2114–2120. <http://dx.doi.org/10.1039/C4AN01978D>.
- Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., Pedreschi, D., 2018. A survey of methods for explaining black box models. *ACM Comput. Surv.* 51 (5), 1–42. <http://dx.doi.org/10.1145/3236009>.
- Hägele, M., Seegerer, P., Lapuschkin, S., Bockmayr, M., Samek, W., Klauschen, F., Müller, K.-R., Binder, A., 2020. Resolving challenges in deep learning-based analyses of histopathological images using explanation methods. *Sci. Rep.* 10 (1), 1–12. <http://dx.doi.org/10.1038/s41598-020-62724-2>.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778.
- Holzinger, A., Lings, G., Denk, H., Zatloukal, K., Müller, H., 2019. Causability and explainability of artificial intelligence in medicine. *Wiley Interdiscip. Rev. Data Mining Knowl. Discov.* 9 (4), e1312. <http://dx.doi.org/10.1002/widm.1312>.
- Hosny, A., Parmar, C., Coroller, T.P., Grossmann, P., Zeleznik, R., Kumar, A., Bussink, J., Gillies, R.J., Mak, R.H., Aerts, H.J., 2018. Deep learning for lung cancer prognostication: A retrospective multi-cohort radiomics study. *PLoS Med.* 15 (11), e1002711. <http://dx.doi.org/10.1371/journal.pmed.1002711>.
- Ilse, M., Tomczak, J.M., Welling, M., 2018. Attention-based deep multiple instance learning. In: *35th International Conference on Machine Learning, ICML 2018. International Machine Learning Society (IMLS)*, pp. 3376–3391.
- Kallenbach-Thieltges, A., Großerueschkamp, F., Jütte, H., Kuepper, C., Reinacher-Schick, A., Tannapfel, A., Gerwert, K., 2020. Label-free, automated classification of microsatellite status in colorectal cancer by infrared imaging. *Sci. Rep.* 10, <http://dx.doi.org/10.1038/s41598-020-67052-z>.
- Kallenbach-Thieltges, A., Großerueschkamp, F., Mosig, A., Diem, M., Tannapfel, A., Gerwert, K., 2013. Immunohistochemistry, histopathology and infrared spectral histopathology of colon cancer tissue sections. *J. Biophotonics* 6 (1), 88–100. <http://dx.doi.org/10.1002/jbio.201200132>.
- Kather, J.N., Halama, N., Marx, A., 2018. 100,000 Histological images of human colorectal cancer and healthy tissue. <http://dx.doi.org/10.5281/zenodo.1214456>.
- Kather, J.N., Krisam, J., Charoentong, P., Luedde, T., Herpel, E., Weis, C.-A., Gaiser, T., Marx, A., Valous, N.A., Ferber, D., et al., 2019a. Predicting survival from colorectal cancer histology slides using deep learning: A retrospective multicenter study. *PLoS Med.* 16 (1), e1002730. <http://dx.doi.org/10.1371/journal.pmed.1002730>.
- Kather, J.N., Pearson, A.T., Halama, N., Jäger, D., Krause, J., Loosen, S.H., Marx, A., Boor, P., Tacke, F., Neumann, U.P., et al., 2019b. Deep learning can predict microsatellite instability directly from histology in gastrointestinal cancer. *Nature Med.* 25 (7), 1054–1056. <http://dx.doi.org/10.1038/s41591-019-0462-y>.
- Kelly, C.J., Karthikesalingam, A., Suleyman, M., Corrado, G., King, D., 2019. Key challenges for delivering clinical impact with artificial intelligence. *BMC Med.* 17 (1), 1–9. <http://dx.doi.org/10.1186/s12916-019-1426-2>.
- Korbar, B., Olofson, A.M., Mirafior, A.P., Nicka, C.M., Suriawinata, M.A., Torresani, L., Suriawinata, A.A., Hassanpour, S., 2017. Looking under the hood: Deep neural network visualization to interpret whole-slide image analysis outcomes for colorectal polyps. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 69–75.
- Kuepper, C., Großerueschkamp, F., Kallenbach-Thieltges, A., Mosig, A., Tannapfel, A., Gerwert, K., 2016. Label-free classification of colon cancer grading using infrared spectral histopathology. *Faraday Discuss.* 187, 105–118. <http://dx.doi.org/10.1039/C5FD00157A>.
- Kuepper, C., Kallenbach-Thieltges, A., Jütte, H., Tannapfel, A., Großerueschkamp, F., Gerwert, K., 2018. Quantum cascade laser-based infrared microscopy for label-free and automated cancer classification in tissue sections. *Sci. Rep.* 8, 1–10. <http://dx.doi.org/10.1038/s41598-018-26098-w>.
- van der Laak, J., Litjens, G., Ciompi, F., 2021. Deep learning in histopathology: The path to the clinic. *Nature Med.* 27 (5), 775–784. <http://dx.doi.org/10.1038/s41591-021-01343-4>.
- Lapuschkin, S., Wäldchen, S., Binder, A., Montavon, G., Samek, W., Müller, K.-R., 2019. Unmasking Clever Hans predictors and assessing what machines really learn. *Nature Commun.* 10 (1), 1–8. <http://dx.doi.org/10.1038/s41467-019-08987-4>.
- Li, K., Luo, H., Huang, L., Luo, H., Zhu, X., 2020. Microsatellite instability: a review of what the oncologist should know. *Cancer Cell Int.* 20 (1), 1–13. <http://dx.doi.org/10.1186/s12935-019-1091-8>.
- Li, M., Wu, L., Wiliem, A., Zhao, K., Zhang, T., Lovell, B., 2019. Deep instance-level hard negative mining model for histopathology images. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 514–522.
- Lipton, Z.C., 2018. The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue* 16 (3), 31–57.
- Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., Van Der Laak, J.A., Van Ginneken, B., Sánchez, C.I., 2017. A survey on deep learning in medical image analysis. *Med. Image Anal.* 42, 60–88. <http://dx.doi.org/10.1016/j.media.2017.07.005>.
- Lu, M.Y., Williamson, D.F., Chen, T.Y., Chen, R.J., Barbieri, M., Mahmood, F., 2021. Data-efficient and weakly supervised computational pathology on whole-slide images. *Nature Biomed. Eng.* 5 (6), 555–570. <http://dx.doi.org/10.1038/s41551-020-00682-w>.
- Macenko, M., Niethammer, M., Marron, J.S., Borland, D., Woosley, J.T., Guan, X., Schmitt, C., Thomas, N.E., 2009. A method for normalizing histology slides for quantitative analysis. In: *2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. IEEE, pp. 1107–1110. <http://dx.doi.org/10.1109/ISBI.2009.5193250>.

- Michalski, R.S., 1983. A theory and methodology of inductive learning. In: Machine Learning. Elsevier, pp. 83–134. <http://dx.doi.org/10.1016/B978-0-08-051054-5.50008-X>.
- Mitchell, T.M., Keller, R.M., Kedar-Cabelli, S.T., 1986. Explanation-based generalization: A unifying view. *Mach. Learn.* 1 (1), 47–80.
- Montavon, G., Samek, W., Müller, K.-R., 2018. Methods for interpreting and understanding deep neural networks. *Digit. Signal Process.* 73, 1–15. <http://dx.doi.org/10.1016/j.dsp.2017.10.011>.
- Murdoch, W.J., Singh, C., Kumbier, K., Abbasi-Asl, R., Yu, B., 2019. Definitions, methods, and applications in interpretable machine learning. *Proc. Natl. Acad. Sci.* 116 (44), 22071–22080. <http://dx.doi.org/10.1073/pnas.1900654116>.
- Popper, K.R., 1959. *The logic of scientific discovery*.
- Raulf, A.P., Butke, J., Küpper, C., Großesueschkamp, F., Gerwert, K., Mosig, A., 2020. Deep representation learning for domain adaptable classification of infrared spectral imaging data. *Bioinformatics* 36 (1), 287–294. <http://dx.doi.org/10.1093/bioinformatics/btz505>.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 234–241.
- Roscher, R., Bohn, B., Duarte, M.F., Garcke, J., 2020. Explainable machine learning for scientific insights and discoveries. *IEEE Access* 8, 42200–42216. <http://dx.doi.org/10.1109/ACCESS.2020.2976199>.
- Rudin, C., 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat. Mach. Intell.* 1 (5), 206–215. <http://dx.doi.org/10.1038/s42256-019-0048-x>.
- Samek, W., Montavon, G., Lapuschkin, S., Anders, C.J., Müller, K.-R., 2021. Explaining deep neural networks and beyond: A review of methods and applications. *Proc. IEEE* 109 (3), 247–278. <http://dx.doi.org/10.1109/JPROC.2021.3060483>.
- Schuhmacher, D., Gerwert, K., Mosig, A., 2020. A Generic Neural Network Approach to Infer Segmenting Classifiers for Disease-Associated Regions in Medical Images. *MedRxiv*, <http://dx.doi.org/10.1101/2020.02.27.20028845>.
- Selbach, L., Kowalski, T., Gerwert, K., Buchin, M., Mosig, A., 2021. Shape decomposition algorithms for laser capture microdissection. *Algorithms Mol. Biol.* 16 (1), 1–17. <http://dx.doi.org/10.1186/s13015-021-00193-6>.
- 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.
- Shi, X., Xing, F., Xie, Y., Zhang, Z., Cui, L., Yang, L., 2020a. Loss-based attention for deep multiple instance learning. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34, pp. 5742–5749.
- Shi, X., Xing, F., Xu, K., Chen, P., Liang, Y., Lu, Z., Guo, Z., 2020b. Loss-based attention for interpreting image-level prediction of convolutional neural networks. *IEEE Trans. Image Process.* 30, 1662–1675. <http://dx.doi.org/10.1109/TIP.2020.3046875>.
- Simonyan, K., Vedaldi, A., Zisserman, A., 2013. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*.
- Sirinukunwattana, K., Pluim, J.P., Chen, H., Qi, X., Heng, P.-A., Guo, Y.B., Wang, L.Y., Matuszewski, B.J., Bruni, E., Sanchez, U., et al., 2017. Gland segmentation in colon histology images: The GlaS challenge contest. *Med. Image Anal.* 35, 489–502. <http://dx.doi.org/10.1016/j.media.2016.08.008>.
- Trukhan, S., Tafintseva, V., Tøndel, K., Großesueschkamp, F., Mosig, A., Kovalev, V., Gerwert, K., Kohler, A., 2020. Grayscale representation of infrared microscopy images by extended multiplicative signal correction for registration with histological images. *J. Biophotonics* 13 (8), e201960223. <http://dx.doi.org/10.1002/jbio.201960223>.
- Vapnik, V., 2006. *Estimation of Dependences Based on Empirical Data*. Springer, New York.
- Wolpert, D.H., 1996. The lack of a priori distinctions between learning algorithms. *Neural Comput.* 8 (7), 1341–1390. <http://dx.doi.org/10.1162/neco.1996.8.7.1341>.
- Zech, J.R., Badgeley, M.A., Liu, M., Costa, A.B., Titano, J.J., Oermann, E.K., 2018. Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study. *PLoS Med.* 15 (11), e1002683. <http://dx.doi.org/10.1371/journal.pmed.1002683>.
- Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A., 2016. Learning deep features for discriminative localization. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2921–2929.