Aalborg Universitet



Self-supervised masked convolutional transformer block for anomaly detection

Madan, Neelu; Ristea, Nicolae Catalin; Ionescu, Radu Tudor; Nasrollahi, Kamal; Khan, Fahad Shahbaz; Moeslund, Thomas B.; Shah, Mubarak

Published in: IEEE Transactions on Pattern Analysis and Machine Intelligence

DOI (link to publication from Publisher): 10.1109/TPAMI.2023.3322604

Publication date: 2024

Document Version Accepted author manuscript, peer reviewed version

Link to publication from Aalborg University

Citation for published version (APA):

Madan, N., Ristea, N. C., Ionescu, R. T., Nasrollahi, K., Khan, F. S., Moeslund, T. B., & Shah, M. (2024). Selfsupervised masked convolutional transformer block for anomaly detection. IEEE Transactions on Pattern Analysis and Machine Intelligence, 46(1), 525-542. https://doi.org/10.1109/TPAMI.2023.3322604

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

Self-Supervised Masked Convolutional Transformer Block for Anomaly Detection

Neelu Madan, Nicolae-Cătălin Ristea, Radu Tudor Ionescu, *Member, IEEE,* Kamal Nasrollahi, Fahad Shahbaz Khan, *Senior Member, IEEE,* Thomas B. Moeslund, and Mubarak Shah, *Fellow, IEEE*

Abstract—Anomaly detection has recently gained increasing attention in the field of computer vision, likely due to its broad set of applications ranging from product fault detection on industrial production lines and impending event detection in video surveillance to finding lesions in medical scans. Regardless of the domain, anomaly detection is typically framed as a one-class classification task, where the learning is conducted on normal examples only. An entire family of successful anomaly detection methods is based on learning to reconstruct masked normal inputs (e.g. patches, future frames, etc.) and exerting the magnitude of the reconstruction error as an indicator for the abnormality level. Unlike other reconstruction-based methods, we present a novel self-supervised masked convolutional transformer block (SSMCTB) that comprises the reconstruction-based functionality at a core architectural level. The proposed self-supervised block is extremely flexible, enabling information masking at any layer of a neural network and being compatible with a wide range of neural architectures. In this work, we extend our previous self-supervised predictive convolutional attentive block (SSPCAB) with a 3D masked convolutional layer, a transformer for channel-wise attention, as well as a novel self-supervised objective based on Huber loss. Furthermore, we show that our block is applicable to a wider variety of tasks, adding anomaly detection in medical images and thermal videos to the previously considered tasks based on RGB images and surveillance videos. We exhibit the generality and flexibility of SSMCTB by integrating it into multiple state-of-the-art neural models for anomaly detection, bringing forth empirical results that confirm considerable performance improvements on five benchmarks: MVTec AD, BRATS, Avenue, ShanghaiTech, and Thermal Rare Event. We release our code and data as open source at: https://github.com/ristea/ssmctb.

Index Terms—anomaly detection, abnormal event detection, self-supervised learning, masked convolution, attention mechanism, transformer, self-attention.

1 INTRODUCTION

T HE applications of vision-based anomaly detection are very diverse, ranging from industrial settings, where the need is to detect faulty objects in the production line [1], [2], to video surveillance, where the need is to detect abnormal behavior [3] such as people fighting or shoplifting, and even medical imaging, where the need is to detect abnormal tissue [4] such as malignant lesions. One of the major challenges of the anomaly detection task is that the definition of what represents an anomaly implies a high dependence on context. For instance, a car driven in a pedes-

- N. Madan is with the Center for Research in Computer Vision (CRCV), Department of Computer Science, University of Central Florida, Orlando, FL, 32816 and the Department of Architecture, Design, and Media Technology, Aalborg University, Denmark.
- N.C. Ristea is with the Department of Telecommunications, University Politehnica of Bucharest, Romania, and the Department of Computer Science, University of Bucharest, Romania.
- R.T. Ionescu is with SecurifAI, the Romanian Young Academy, and the Department of Computer Science, University of Bucharest, Romania. R.T. Ionescu is the corresponding author. E-mail: raducu.ionescu@gmail.com
- K. Nasrollahi is with the Department of Architecture, Design, and Media Technology, Aalborg University, Denmark, and Milestone Systems, Denmark.
- F.S. Khan is with Mohamed bin Zayed University of Artificial Intelligence (MBZUAI), UAE, and Linköping University, Sweden.
- Thomas B. Moeslund is with the Department of Architecture, Design, and Media Technology, Aalborg University, Denmark.
- M. Shah is with the Center for Research in Computer Vision (CRCV), Department of Computer Science, University of Central Florida, Orlando, FL, 32816.

Manuscript received April 19, 2022; revised August 26, 2022.

trian area is labeled as anomalous, whereas the same action can be considered normal in a different context, *e.g.* when the car is driven on the road. Due to the reliance on context and the sheer diversity of possible anomalies, it is often very difficult to gather abnormal examples for training. As a result, anomaly detection is commonly devised as a one-class classification task, where the generic approach implicitly or explicitly learns the distribution of the normal training data. During inference, examples that do not belong to the normal training data distribution are labeled as abnormal. There are several categories of methods that are guided by this generic approach, such as dictionary-learning methods [5]–[10], change-detection frameworks [11]–[14], distancebased models [15]–[27], probabilistic frameworks [28]–[37], and reconstruction-based models [3], [38]–[48].

Our approach belongs to the category of reconstruction methods, which have recently become a prominent choice in anomaly detection [38], [39], [41], [43], [44], [46]–[48]. Reconstruction-based models implicitly learn the normal data distribution by minimizing the reconstruction error of the normal instances at training time. These models are based on the assumption that the learned latent manifold does not offer the means to reconstruct the abnormal samples robustly, due to the unavailability of such samples at training time. Hence, the reconstruction error is directly employed as the anomaly score. A particular subcategory of reconstruction-based models relies on learning to predict masked inputs [41], [42], [50], [51] as a self-supervised pretext task. In this case, the reconstruction error with



Fig. 1. An overview of our self-supervised masked convolutional transformer block (SSMCTB). At every location where the masked filters are applied, the proposed block has to rely on the visible regions (sub-kernels) to reconstruct the masked region (center area). A transformer module performs channel-wise self-attention to selectively promote or suppress reconstruction maps via a set of weights returned by a sigmoid (σ) layer. The block is self-supervised via the Huber loss (\mathcal{L}_{SSMCTB}) [49] between masked and returned activation maps. Best viewed in color.

respect to the masked information is used to assess the abnormality level of an input instance. Depending on the input type (image or video), methods in this subcategory mask various parts of the input, e.g. superpixels in images [41], future frames in video [42], or middle bounding boxes in object-centric temporal sequences [50], [51], and employ the whole model to reconstruct the masked input. We, on the other hand, propose to encapsulate the functionality of reconstructing the masked information into a novel neural block. There are two major benefits when wrapping the reconstruction task as a low-level architectural component: (i) it enables introducing the reconstruction of masked information as a self-supervised task at any layer of a neural network (not only at the input), and (ii) it eases integrating the self-supervised reconstruction task into a broad variety of neural architectures, regardless of whether the respective models are reconstruction-based or not. Due to its advantages, our block is very flexible and generic.

Our self-supervised reconstruction block consists of a dilated masked convolution followed by a channel-wise transformer module. The center area of our convolutional kernel is masked, hence hiding the center of the receptive field at every location where the filters are applied. In other words, each component of the input tensor is certainly masked at some point during the convolution operation, which means that the entire input tensor ends up being masked. Next, the convolutional activation maps are transformed into tokens using an average pooling layer. Then, the resulting tokens are passed through a transformer module [52], [53] that performs channel-wise self-attention. The proposed block is equipped with a transformer module to avoid the direct reconstruction of the masked area through linearly interpolating the visible regions of the convolutional kernels. The final activation maps are multiplied with the resulting attention tokens. Our block is designed in such a way that the output tensor has the same dimensions as the input tensor, which allows us to easily introduce a loss within our block to minimize the reconstruction error between the output tensor and

the masked input tensor. By integrating this loss, our block becomes a self-contained trainable component that learns to predict the masked information via self-supervision. As such, we coin the term *self-supervised masked convolutional transformer block* (SSMCTB) to designate our novel neural component for anomaly detection. As shown in Figure 1, SSMCTB learns to reconstruct the masked region based on the available context (visible regions of the receptive field), for each location where the dilated kernels are applied. Notably, we can graciously control the level (from local to global) of the contextual information by choosing the appropriate dilation rate for the masked kernels.

SSMCTB is an extension of the self-supervised predictive convolutional attentive block (SSPCAB) introduced in our recent CVPR 2022 paper [54]. In the current work, we modify SSPCAB in three different ways: (i) we replace the standard channel attention module in the original SSPCAB [54] with a multi-head self-attention module [52], [53] to increase the modeling capacity, (ii) we extend the masked convolution operation with 3D convolutional filters, enabling the integration of SSMCTB into networks based on 3D convolutional layers, and (iii) we replace the mean squared error (MSE) loss with the Huber loss [49], since the latter loss is less sensitive to outliers than the former loss. Aside from these architectural changes, we demonstrate the applicability of our block to more domains, adding anomaly detection in medical images and thermal videos to the previously considered tasks based on RGB images and surveillance videos. Moreover, we conduct a more extensive ablation study, thus providing a more comprehensive set of results. We also show that our module is suitable for both convolutional and transformer-based architectures.

We introduce SSMCTB into multiple state-of-the-art neural models [42], [44], [55]–[60] for anomaly detection and conduct experiments on five benchmarks: MVTec AD [1], BRATS [61], Avenue [9], ShanghaiTech [3], and Thermal Rare Event. The Thermal Rare Event data set is a novel benchmark for anomaly detection, which we constructed by manually labeling abnormal events from the Seasons in Drift data set [62]. The chosen benchmarks belong to various domains, ranging from industrial and medical images to RGB and thermal videos. This is to show that SSMCTB is applicable to multiple domains. When adding SSMCTB to the state-of-the-art models, our experiments show evidence of consistent improvements across all models and tasks, indicating that our block is generic and easily adaptable. When compared to SSPCAB, we observe performance gains in the majority of cases, showing that the multi-head selfattention and the Huber loss are beneficial in detriment of the standard channel attention [63] and the MSE loss, respectively.

In summary, with respect to our previous work [54], our current contribution is sixfold:

- We extend the 2D masked convolution to a 3D masked convolution that considers a 3D context, and we integrate the new 3D SSMCTB into two 3D networks for anomaly detection [55], [56].
- We replace the Squeeze-and-Excitation module [63] of SSPCAB with a transformer module that performs channel-wise attention.
- We substitute the MSE loss with the Huber loss, improving the sensitivity to outliers during selfsupervised learning.
- We conduct a more comprehensive set of experiments, including a new method and two new benchmarks from previously missing domains (medical images, thermal videos).
- We provide an extensive ablation study, including different variations of the proposed self-supervised block.
- We annotate a subset (one week of video) of the Seasons in Drift [62] data set with anomaly labels, obtaining a new benchmark for anomaly detection in thermal videos.

2 RELATED WORK

2.1 Transformers

Vaswani et al. [53] introduced the self-attention mechanism, sparking the research of neural architectures relying solely on attention, including research on vision transformers [52], [64]–[74]. These models are now embraced at a fast pace in the field of computer vision, certainly due to the imposing performance levels across a broad variety of tasks, ranging from object recognition [52], [69], [70] and object detection [64], [73], [74] to image generation [68], [71], [72] and anomaly detection [75]–[78]. Unlike approaches using only transformer-based attention [52], [64]-[70], [73], [74], [79], we propose a novel and flexible block that employs transformer-based attention along with masked convolution, which can be integrated into multiple architectures that are not necessarily transformer-based. To endorse this statement, we introduce SSMCTB into a variety of models and conduct a series of experiments showing that our block can bring significant performance gains. Another difference from vision transformers is that our block performs channelwise self-attention, while conventional vision transformers perform spatial attention [52]. We conduct an ablation study

to compare channel and spatial attention inside SSMCTB, showing that channel attention provides superior performance and faster processing.

2.2 Self-Supervision via Information Masking

The reconstruction of masked information has recently become an attractive area of interest [60], [80]-[83]. Models based on information masking are usually pre-trained on a self-supervised reconstruction task, being later employed for downstream visual tasks such as object detection and image segmentation. For instance, He et al. [60] proposed to reconstruct masked (erased) patches as a self-supervised pretext task for pre-training auto-encoders, subsequently using them for mainstream tasks, including object detection and object recognition. They reported optimal results when a majority (75%) of the patches is masked. Masked auto-encoders are directly applicable to anomaly detection. However, we show that SSMCTB can boost the performance of masked auto-encoders, suggesting that it can leverage information masking in a distinct way. Wei et al. [81] aimed at pre-training video models, proposing to mask spatiotemporal cubes from a video and predict the features of the masked regions. Chang et al. [82] introduced a bidirectional decoder that learns to predict masked tokens by attending them from all directions. The proposed method provides an efficient substitute for generative transformers. Yu et al. [83] used a masked point modeling task for pre-training a point cloud transformer. They showed that the representation learned by the model transfers well to new (downstream) tasks and domains. Distinct from such methods, we integrate information masking at a core operational level inside neural networks via our masked convolutional layer. We self-supervise our block (which incorporates masked convolution) through a reconstruction loss and show that modeling the context towards reconstructing the masked information results in an effective discriminative manifold for anomaly detection.

We underline that some recent approaches [42], [50], [84] utilize masking as a surrogate task for anomaly detection. We discuss these methods and explain how our approach is different in a separate subsection below.

2.3 Anomaly Detection

Anomaly detection frameworks are usually trained in a one-class setting, where only normal data is available at training time, whereas both normal and abnormal examples are present at test time. The anomaly detection methods operating in this setting can be classified into different categories, which are briefly presented below. Dictionary learning methods [5]-[10] construct a dictionary of atoms from normal instances, labeling examples that are not represented in the dictionary as abnormal. Change detection frameworks [11]-[14] are applied directly on test videos, measuring the degree of change between current and preceding video frames to detect anomalies. Probabilistic models [28]-[37] learn the probability density function of the normal data, flagging examples outside the distribution as abnormal. Distance-based approaches [15]-[27], [85] learn a distance function between samples, such that the distance between normal instances is lower than the distance

between normal and abnormal instances. Reconstructionbased methods [3], [38]–[48], [56], [86] learn to reconstruct normal examples, detecting anomalies based on the magnitude of the reconstruction error, as anomalies tend to have larger errors than normal instances.

Reconstruction-based methods. Since our block belongs to the category of reconstruction-based models, we discuss this category in more detail next. Reconstruction-based models are often chosen for both image and video anomaly detection [44], [50], [56], [59]. These approaches typically employ auto-encoders and generative adversarial networks (GANs) to learn a powerful latent manifold representing the normal data distribution. For the video domain, some anomaly detection approaches [17], [42], [59] incorporate additional cues by reconstructing the optical flow to capture motion information, enabling the detection of motion-based anomalies such as running and jumping. Doshi *et al.* [87] proposed a continual learning setup, which could be easily extended for future normal and abnormal patterns.

As the amount of normal training data is generally high, latent manifolds show a tendency to generalize too well, being capable of reconstructing abnormal instances with low error. In the context of anomaly detection, generalizing to out-of-distribution samples, e.g. anomalies, is not desired, although this would be mostly desirable in other application domains. To mitigate this issue, researchers employed various techniques, such as adding memory modules [39], [44], [59] or pseudo-anomalies during training [58], [84]. Memory-based auto-encoders [39], [59] generally employ an additional module to memorize the normal patterns observed in the training data. Consequently, memory modules increase the computational complexity of the model, and the faithful reconstruction of normal samples highly relies on the size of the memory module. Georgescu et al. [58] proposed to optimize the model on pseudo-anomalies with gradient ascent, while still using gradient descent to learn the normal data distribution. This results in a powerful discriminative subspace for the robust detection of the abnormal samples. The pseudo-abnormal instances are samples collected from different contexts, such as flowers, animals, cartoons, and textures, unrelated to the object distribution (comprising humans, cars, bicycles, etc.) observed in typical urban surveillance scenes. Similarly, Astrid et al. [84] generated pseudo-anomalies by skipping a few frames from the video and training an auto-encoder by maximizing the loss for pseudo-anomalies and minimizing it for normal samples. Introducing pseudo-anomalies increases the training time and may sometimes cause instability if the balance between gradient descent on normal data and gradient ascent on pseudo-abnormal data is not tuned. Different from related reconstruction-based methods, we increase the difficulty of the reconstruction task by masking information wherever SSMCTB is introduced into a neural model, thus making it harder for the model to generalize to abnormal data. As shown by our experimental results, our block adds a marginal computational overhead.

Masking for Anomaly Detection. Some approaches [38], [42], [50], [75], [78], [88], [89] are already using the prediction of masked inputs as a surrogate task for anomaly detection. These models form a distinctive subcategory of reconstruction-based methods. Liu *et al.* [42] proposed a

GAN for predicting a future frame based on a few past frames, where anomalies are classified according to the prediction error. Another GAN-based approach [90] performs joint detection and localization of anomalies via inpainting. The generator of this method learns to inpaint a patch from the input image, while the discriminator learns to identify if the inpainted patch is normal or abnormal. Interestingly, the inpainting task has also been studied in conjunction with vision transformers [78].

Generalizing over the method of Liu et al. [42], Yu et al. [89] employed the Cloze task [91], which is about learning to complete the video when certain frames are removed. Georgescu et al. [50] proposed the masking of the middle box of each temporal cube centered on an object. Anomalies are detected based on the assumption that motion reconstruction for an abnormal object is more difficult than for the normal ones. Fei et al. [38] proposed the Attribute Restoration Network (ARNet), where attributes such as color and orientation of the input are removed, and the network learns to restore those attributes. The idea is based on the assumption that the anomalous data can be distinguished based on the restoration error. Haselmann et al. [88] introduced an approach for surface anomaly detection by erasing a rectangular box from the center of the image and using the interpolation error for the classification of samples into normal or abnormal. Inspired by the success of masked auto-encoders [60], Jiang et al. [75] proposed a masked Swin Transformer [92] that is trained to inpaint masked regions. To cope with the lack of abnormal samples during training, the authors used simulated anomalies.

Unlike other models based on information masking, we propose a novel approach that incorporates the reconstruction-based functionality into a single neural block, which can be easily integrated into other state-ofthe-art anomaly detection models. Our experimental results confirm that our block is a valuable addition to various models, including both CNNs and transformers, which are applied to anomaly detection in a wide range of domains.

3 METHOD

3.1 Motivation and Overview

A wide set of computer vision tasks, including anomaly detection [44], [58], [59], [93], [94], are often addressed with convolutional neural networks (CNNs) [95], [96], due to the impressive performance levels reached by these models, sometimes even surpassing human-level accuracy. The defining component of a CNN architecture is the convolutional layer, which typically comprises multiple filters (kernels) that activate on discriminative local patterns captured within the receptive field of the respective filters. Each filter produces an activation map that is further given as input to the next convolutional layer. Since each filter in the subsequent layer processes all activation maps from the previous layer at once, the local features extracted by the previous layer are combined into more complex features. This sequential processing of features over multiple convolutional layers gives rise to a hierarchy of features during the learning process. Earlier convolutional layers activate on low-level features such as corners or edges, and later layers gradually shift to higher-level features such as car wheels or



Fig. 2. Our 2D masked convolutional kernel. The visible area of the receptive field is denoted by the regions K_i , $\forall i \in \{1, 2, 3, 4\}$, while the masked area is denoted by M. A dilation factor d controls the local or global nature of the visible information with respect to M. Best viewed in color.

human body parts, as shown by Zeiler *et al.* [97]. Although the learned hierarchy of features is very useful in solving discriminative tasks, CNNs do not have the direct means to model the global arrangement of local features [98], since they do not generalize well to novel viewpoints or affine transformations [99]. The inability of grasping the global arrangement of local features is mainly caused by the fact that convolutional filters operate on a limited (and typically small) receptive field, not making use of the context.

We hereby propose a self-supervised masked convolutional transformer block (SSMCTB), which is aimed at learning to reconstruct masked information based on contextual information. To accurately solve the reconstruction of its masked input, the proposed block is required to employ the context and learn the global structure of the local patterns. Hence, it inherently learns to cope with the problem stated by Sabour *et al.* [98], specifically the fact that CNNs lack the proper comprehension of the global arrangement of local features. To embed this learning capability into our block, we structure SSMCTB as a convolutional layer with dilated masked kernels, followed by a transformer module that performs channel attention. We attach a self-supervised loss function to our block in order to minimize the reconstruction error between the masked input and the predicted output.

We emphasize that SSMCTB is quite flexible, since it can be inserted at any level of almost any CNN or transformer model, generating powerful features that offer the capability of reconstructing masked information based on context. While the ability of learning and harnessing the global arrangement of local patterns is potentially useful in solving a broader set of computer vision tasks, we conjecture that anomaly detection is a natural and immediate application domain for SSMCTB, hence focusing our work in this direction. Indeed, since anomaly detection models are typically trained on normal data only, integrating SSMCTB into a neural model will lead to the learning of features that recover only masked normal data. Hence, when an anomalous sample is given as input during inference, SSM-CTB is likely less capable of reconstructing the masked information. This empowers the model to directly estimate the abnormality level of a data sample via the reconstruction error given by SSMCTB. Our claims are supported through the comprehensive set of experiments on image and video anomaly detection presented in Section 4.

3.2 Architecture

Our initial self-supervised block introduced in [54] was formed of a 2D masked convolution and a Squeeze-and-Excitation (SE) module [63]. To broaden the applicability of our block, we now introduce a 3D masked convolutional layer to replace the 2D masked convolution, whenever this is needed. Moreover, we replace the SE attention module with a modern transformer-based attention module [52], [53] to attend to the channels given as output by the masked convolution. We describe the individual components of our block below, while providing a graphical overview of SSMCTB in Figure 1.

2D Masked Convolution. Figure 2 shows our 2D masked convolutional kernel, where the corner regions of this kernel (in green color) are the learnable parameters (weights) defining the visible regions of the receptive field. The four learnable sub-kernels are denoted by $\mathbf{K}_i \in \mathbb{R}^{k' \times k' \times c}$, $\forall i \in \{1, 2, 3, 4\}$, where the spatial size $k' \in \mathbb{N}^+$ of each subkernel is a hyperparameter of our block, while the number of channels $c \in \mathbb{N}^+$ always matches the number of channels of the input tensor. Our masked region $M \in \mathbb{R}^{1 imes 1 imes c}$ (in pink color) is located at the center of the receptive field. Each sub-kernel K_i is located at a configurable distance $d \in \mathbb{N}^+$ (also referred to as *dilation rate*) from the masked region M. To keep the number of hyperparameters to a bare minimum, we fix the spatial size of the masked region to 1×1 . As a result, the spatial size k of the entire receptive field of our 2D masked convolution is k = 2k' + 2d + 1.

Let $oldsymbol{X} \in \mathbb{R}^{h imes w imes c}$ be the input tensor of the masked convolutional layer, where $c \in \mathbb{N}^+$ denotes the number of channels, and $h, w \in \mathbb{N}^+$ represent the height and width of the input tensor, respectively. When we apply our custom kernel at a given location (a, b) of the input tensor **X**, only the input values that overlap with the sub-kernels K_i are taken into consideration during the masked convolution operation, resulting in a single output value. We underline that our masked convolution is equivalent to convolving the input independently with the sub-kernels K_i , where each sub-kernel has a different spatial shift with respect to the current location (a, b), and the resulting values are summed up to produce a single output value. The output value at position (a, b) represents the reconstruction for only one value of the tensor M located at the same position (a,b). To reconstruct the entire tensor $M_{,}$ our layer requires the application of *c* masked convolutional filters, each reconstructing the masked value from a distinct channel at position (a, b). Convolving a single masked filter over the entire input generates a complete activation map. Since there are c masked convolutional filters, the output tensor Z is formed of c activation maps. Our aim is to apply the masked convolution such that every element in the input tensor is masked exactly once, *i.e.* we want to mask and predict the reconstruction for every spatial location of the input. As such, we set the stride to 1 and apply a zeropadding of k' + d in each direction. With this configuration in place, the output tensor \boldsymbol{Z} has $h \times w \times c$ components, exactly as the input tensor X. To obtain the final values, the output tensor Z is passed through Rectified Linear Units (ReLU) [100]. Finally, we emphasize that k' and d are the only tunable hyperparameters of our masked convolutional



Fig. 3. Our 3D masked convolutional kernel. The visible area of the receptive field is denoted by the regions $K_i, \forall i \in \{1, 2, ..., 8\}$, while the masked area is denoted by M. A dilation factor d controls the local or global nature of the visible information with respect to M. Best viewed in color.

layer.

3D Masked Convolution. Considering that anomaly detection is often applied on 3D inputs, *e.g.* video or medical scans, some researchers naturally resort to employing 3D CNNs. To this end, we extend our 2D masked convolution to the 3D domain, broadening the applicability of SSMCTB. We thus reformulate the 2D spatial reconstruction task into a more difficult one, which implies learning a global 3D structure of the discovered local patterns. Let $K_i \in \mathbb{R}^{k' \times k' \times k' \times c'}$, $\forall i \in \{1, 2, ..., 8\}$, be the learnable 3D sub-kernels depicted in Figure 3, where k' and c are defined above. The masked region M is located in the center of the 3D kernel, equally distant from the sub-kernels K_i . The size of the receptive field of our 3D masked convolution is $k \times k \times k$, where k = 2k' + 2d + 1.

To compute the feature response using the 3D masked convolutional layer, the input $X \in \mathbb{R}^{h \times w \times r \times c}$ is convolved with our custom masked kernel, where r represents the depth, and h, w and c are defined as before. The 3D filter is applied analogously to the 2D one, the only difference being that the input data and the kernel itself are 3D. The number of 3D convolutional filters is equal to the number of channels c, such that the spatial dimension of the output tensor $Z \in \mathbb{R}^{h \times w \times r \times c}$ is identical to that of the input X. The 3D masked convolution has the same number of configurable hyperparameters, these being k' and d.

Channel-wise transformer block. To better exploit the interdependencies between the different activation maps produced by the masked convolutional layer, we replace the Squeeze-and-Excitation module in SSPCAB [54] with a self-attention transformer-based module. The new attention module is able to capture more complex channel-wise interrelations through its higher modeling capacity, as it learns to assign attention weights to the reconstructed information corresponding to each masked convolutional filter in order to reduce the reconstruction error of SSMCTB.

Let $Z \in \mathbb{R}^{h \times w \times c}$ be the output tensor of a 2D masked convolutional layer with *c* filters. First, we apply a spatial average pooling, obtaining $\hat{Z} \in \mathbb{R}^{h' \times w' \times c}$, where $h' \leq h$ and $w' \leq w$. The average pooling layer is followed by a reshape operation, obtaining a matrix $A \in \mathbb{R}^{c \times n}$, which contains a vector of $n = h' \cdot w'$ components on each row to represent each masked filter. Next, A is fed into a linear projection layer to obtain the tokens $T \in \mathbb{R}^{c \times d_t}$, which are further summed up with the positional embeddings to obtain the final tokens $T^* \in \mathbb{R}^{c \times d_t}$.

Let *f* be a multi-head attention layer with $H \in \mathbb{N}^+$ heads, *g* a multi-layer perceptron, *norm* a normalization layer, and $P, R \in \mathbb{R}^{c \times d_t}$ some auxiliary tensors. The operations performed inside the transformer are formally described as follows:

$$\boldsymbol{P} = f(norm(\boldsymbol{R})) + \boldsymbol{R},\tag{1}$$

$$\boldsymbol{R} = g(norm(\boldsymbol{P})) + \boldsymbol{P}.$$
 (2)

As illustrated in Figure 1, the whole process described in Eq. (1) and Eq. (2) is repeated L times, where $L \in \mathbb{N}^+$ represents the number of transformer blocks inside the transformer module. For the first transformer block, R is initialized with T^* . In Eq. (1), the sequence of c tokens R is normalized, fed into the multi-head attention layer and added to itself, obtaining P. Further, P is normalized, fed into a multi-layer perceptron and also added to itself, according to Eq. (2).

The transformer is aimed at capturing the interaction among all c tokens by encoding each token in terms of the channel-wise contextual information. This is achieved via the multi-head attention layer f. Each head $j \in \{1, 2, ..., H\}$ comprises three learnable weight matrices denoted as $W^{Q_j} \in \mathbb{R}^{d_t \times d_q}, W^{K_j} \in \mathbb{R}^{d_t \times d_k}$ and $W^{V_j} \in \mathbb{R}^{d_t \times d_v}$, where $d_q = d_k$. The weight matrices are multiplied with the input tokens R, producing the queries Q_j , keys K_j and values V_j . In other words, the input sequence R is projected onto these weight matrices to get $Q_j = R \cdot W^{Q_j}, K_j = R \cdot W^{K_j}$ and $V_j = R \cdot W^{V_j}$, respectively. The output $Y_j \in \mathbb{R}^{c \times d_v}$ of each self-attention head is given by:

$$\boldsymbol{Y}_{j} = softmax \left(\frac{\boldsymbol{Q}_{j} \cdot \boldsymbol{K}_{j}^{\top}}{\sqrt{d_{q}}} \right) \cdot \boldsymbol{V}_{j}, \tag{3}$$

where K_j^{\dagger} is the transpose of K_j . The outputs returned by the self-attention heads are simply summed into Y, *i.e.*:

$$\boldsymbol{Y} = \sum_{j=1}^{H} \boldsymbol{Y}_j. \tag{4}$$

We can now rewrite Eq. (1) as follows:

$$P = Y + R. \tag{5}$$

The output sequence \boldsymbol{R} returned by the final transformer block is averaged along the token dimension, obtaining $\hat{\boldsymbol{R}} \in \mathbb{R}^{c \times 1}$, then fed into a sigmoid layer to generate the final attention weight assigned to each channel. Finally, the resulting attention weights are applied to the tensor \boldsymbol{Z} , obtaining the reconstructed output denoted by $\hat{\boldsymbol{X}} \in \mathbb{R}^{h \times w \times c}$, as follows:

$$\hat{\boldsymbol{X}} = \boldsymbol{Z} \otimes \sigma(\hat{\boldsymbol{R}}),$$
 (6)

where \otimes denotes the element-wise multiplication, and σ denotes the sigmoid layer. The entire processing performed by the transformer module is analogously applied when the preceding layer is a 3D masked convolution.

3.3 Self-Supervised Reconstruction Loss

We devise an integrated reconstruction loss to train the proposed SSMCTB in a self-supervised manner. To better cope with outlier values and reduce the sensitivity of the model to outliers, we define the self-supervised objective as the Huber loss between the reconstructed output \hat{X} and the input X, replacing the mean squared error (MSE) used by SSPCAB. The self-supervised objective enables our model to learn reconstructing the masked information at every location where the masked filters are applied. Let G denote the SSMCTB function. With this notation, the self-supervised reconstruction loss of our block can be computed as follows:

$$\mathcal{L}_{\text{SSMCTB}}(G, \mathbf{X}) = \begin{cases} \frac{1}{2} \cdot (G(\mathbf{X}) - \mathbf{X})^2, & \text{if } |G(\mathbf{X}) - \mathbf{X}| < \delta \\ \delta \cdot (|G(\mathbf{X}) - \mathbf{X}| - \frac{\delta}{2}), & \text{otherwise} \end{cases}$$
$$= \begin{cases} \frac{1}{2} \cdot (\hat{\mathbf{X}} - \mathbf{X})^2, & \text{if } |\hat{\mathbf{X}} - \mathbf{X}| < \delta \\ \delta \cdot (|\hat{\mathbf{X}} - \mathbf{X}| - \frac{\delta}{2}), & \text{otherwise} \end{cases}, \end{cases}$$
(7)

where $\delta \in \mathbb{R}^+$ is a hyperparameter representing the error threshold that determines when to switch from the squared loss (applied for errors below δ) to the absolute loss (applied for errors higher than or equal to δ).

When integrating SSMCTB into some neural network F, we can simply add our loss $\mathcal{L}_{\text{SSMCTB}}$ to the loss function \mathcal{L}_F of the respective neural model, resulting in a new loss function comprising both terms. Formally, the overall loss can be computed as follows:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_F + \lambda \cdot \mathcal{L}_{\text{SSMCTB}},\tag{8}$$

where $\lambda \in \mathbb{R}^+$ is a hyperparameter deciding the importance of $\mathcal{L}_{\text{SSMCTB}}$ with respect to \mathcal{L}_F . Naturally, the hyperparameter λ can vary across neural models or visual tasks.

4 EXPERIMENTS AND RESULTS

4.1 Data Sets

We carry out experiments on five benchmarks from various domains, considering the most popular data set choices, *e.g.* MVTec AD [1], BRATS [61], CUHK Avenue [9], ShanghaiTech [3], whenever such an option is available for a certain domain. For the thermal video domain, we build our own data set.

MVTec AD. MVTec AD [1] has become a standard data set for benchmarking anomaly detection methods applied in inspecting industrial defects. The data set contains over 5,000 images distributed over 15 different categories of textures (10) and objects (5). It comprises 3,629 defect-free training samples, as well as 1,725 test images with and without defects.

BRATS. BRATS [61] is a multimodal magnetic resonance imaging (MRI) data set for brain tumor segmentation. It is an intrinsically heterogeneous data set that contains brain tumors of different shape, appearance and histology. The data set comprises manually annotated MRI scans acquired by 19 institutions employing different clinical protocols. To evaluate anomaly detection models, we introduce a novel split of the data set, such that all training images are lesion-free, *i.e.* all images with lesions are kept for testing. The

TABLE 1

Rare events in our thermal anomaly detection data set along with the frequency of each event type.

Rare Event Type	Frequency
Activities in restricted zones	6
Jumping	4
Reverse driving	2
Unexpected activities	2
Unexpected interactions	14
Unexpected vehicle	1
Total	29

training set includes 11,280 slices (125 scans), which leaves 27,745 slices (180 scans) for the test set.

Avenue. CUHK Avenue [9] is one of the most widely-used data sets for video anomaly detection. It contains 16 videos for training and 21 videos for testing. The training videos comprise only normal events, whereas the test videos contain both normal and abnormal events. The data set contains videos from a single surveillance camera. Avenue contains people-related anomalies such as running, walking in the wrong direction, jumping, dancing, loitering and throwing objects.

ShanghaiTech. ShanghaiTech [3] is one of the largest benchmarks for video anomaly detection, comprising 330 training and 107 test videos. As in CUHK Avenue, abnormal instances appear only at test time. The data set includes videos from multiple scenes. Examples of anomalies are related to people, e.g. fighting, jumping and stealing, as well as vehicles, e.g. bikes and cars in pedestrian (forbidden) zones. Thermal Rare Event. To construct the Thermal Rare Event data set, we sampled one week of videos (330 clips) from the Seasons in Drift (SiD) data set [62]. The SiD data set [62] is an unlabeled thermal surveillance data set captured from a single view over a period of 8 months. The data set captures activities near a harbor front during day and night. Each clip is about 2 minutes long and contains 120 frames, being sampled at 1 frame per second (FPS). Out of the 330 clips, there are 29 clips containing rare (anomalous) events. We manually annotated these rare events at the frame level. In total, our Thermal Rare Event data set contains 36,120 frames for testing and 3,480 frames for training. The list of rare events in our data set along with their respective frequencies are summarized in Table 1. Examples of rare events from different categories are: activities in restricted zones (people sitting, standing, and running close to the pier), jumping (person jumping, group jumping), unexpected activities (doing yoga, smoking), unexpected interactions (running with stroller, embarking to a boat, debarking from a boat, chasing, dancing), unexpected vehicles (different types of trucks). We release the Thermal Rare Event data set along with our code at: https://github.com/ristea/ssmctb/.

4.2 Evaluation Measures

Image Anomaly Detection. Following Bergmann *et al.* [1], we carry out the evaluation on MVTec AD and BRATS considering the area under the receiver operating characteristics curve (AUROC) and the average precision (AP). To generate the ROC curve, the true positive rate (TPR) is plotted against the false positive rate (FPR). We evaluate both detection and

localization performance levels of anomaly detection methods. In anomaly detection, TPR is the proportion of images correctly classified as abnormal, and FPR is the proportion of normal images wrongly classified as abnormal. For the localization task, TPR denotes the proportion of correctly classified abnormal pixels, while FPR represents the proportion of normal pixels incorrectly classified as abnormal. For the localization task, we obtain anomaly segments by applying a threshold to produce a binary decision for each pixel, as described in [1]. The localization AP is obtained by taking the mean at different threshold levels.

Video Anomaly Detection. As the majority of previous works [101], we evaluate the detection performance of video anomaly detection methods using the frame-level area under the curve (AUC). To compute the AUC measure, a video frame is marked as abnormal if at least one pixel is abnormal. Inspired by Georgescu et al. [58], we employ both micro AUC and macro AUC. The micro AUC is computed by first concatenating all frames in all videos into a single video, while the macro AUC represents the average of the AUC scores which are independently computed for each single video in the test set. To evaluate the localization performance, we report the region-based detection criterion (RBDC) and the track-based detection criterion (TBDC) proposed by Ramachandra et al. [19]. RBDC considers each detected region, marking it as a *true positive* if the intersection over union (IOU) between the detected and the groundtruth anomalous region is greater than α . TBDC marks each tracked region as a *true positive* if the overlap with the ground-truth anomalous track is greater than β . We set the same values for α and β as previous works [19], [58], *i.e.* $\alpha = 0.1$ and $\beta = 0.1$.

4.3 Implementation Details

We choose eight state-of-the-art approaches [42], [44], [55]-[60] for image and video anomaly detection to serve as underlying models, on top of which we add SSPCAB [54] and SSMCTB (ours). We alternatively integrate SSPCAB and SSMCTB directly into the official implementations of the chosen baselines, while preserving all hyperparameter values, e.g. the number of epochs and the learning rate, as specified in the corresponding papers [42], [44], [55]-[60]. Even so, we are unable to exactly reproduce the original results for two baselines methods, *i.e.* those of Park *et al.* [44] and Liu et al. [42]. However, our reproduced quantitative results are still close to the originally reported results. For a fair comparison, we compare the models based on SSPCAB and SSMCTB with the reproduced baselines. Additionally, when we repurpose the approach of Park et al. [44] from the RGB domain to the thermal domain, we modify some hyperparameters, namely the number of epochs and the mini-batch size.

Following Ristea *et al.* [54], we replace the penultimate convolutional layer with SSMCTB in most underlying models. One exception is the architecture of Georgescu *et al.* [50], where SSPCAB and SSMCTB are integrated into the penultimate convolutional layer of the decoder instead of the final classification network. Another exception is the masked auto-enconder [60] based on the ViT backbone, where we place SSPCAB and SSMCTB before the first transformer block.

TABLE 2

Micro AUC scores (in %) obtained on the Avenue data set with different hyperparameter configurations, varying the kernel size (k'), the dilation rate (d), the loss type, and the attention type, while integrating SSMCTB into the method of Park *et al.* [44]. The top score is highlighted in bold.

Method	$\mathcal{L}_{\text{SSMCTB}}$	d	k'	Attention	Micro AUC
Park <i>et al</i> . [44]	-	-	-	-	82.8
		0	1		83.1
+SSMC	MAE	1	1	_	83.5
(no attention)	WIAL	2	1	_	84.2
		3 1			84.4
	0 1			83.7	
	MAE	1	1	CA	84.9
	IVIZIL	2	1		85.5
		3	1		85.9
		0	1		84.9
	MSE	1	1	CA	85.7
	INICE	2	1		85.4
		3	1		86.4
		0	1		83.3
	SSIM	1	1	CA	85.5
	00111	2	1		84.9
		3	1		83.0
		0	1		84.2
	Huber		1	CA	87.0
+SSMCTB		2			86.5
		3	1		86.1
			2		84.1
	Huber		2	CA	84.9
		2			84.8
		3	2		80.0
			3		84.5
	Huber	1	2	CA	05.0
		$\begin{vmatrix} 2 \\ 2 \end{vmatrix}$	2		80.4
		3	3		04.3 86.2
					84.7
	Huber	2	1	SA	85.8
		2			80.4
	Hubor	1	1	CA + SA	86.2
	Tuber		1	CA + 5A	00.2

In our previous work [54], we conducted a set of preliminary experiments to find an optimal value for the hyperparameter λ representing the contribution of our selfsupervised loss to the total loss defined in Eq. (7), taking values from 0.1 to 1 at an interval of 0.1. Following our previous work [54], we keep $\lambda = 0.1$ across all data sets. However, for two baselines [57], [59], we notice that the magnitude of our loss is too high with respect to the original losses of the respective models, dominating the optimization. Following our previous work [54], we decrease λ to 0.001 to reduce the dominant influence of our loss on these two particular models [57], [59].

For the channel-wise transformer, we fix the activation map size after the average pooling layer to 1×1 , the token size d_t to 64, the number of heads H to 4, as well as the number of successive transformer blocks L to 2. We discuss results for other transformer configurations in Section 4.7.

4.4 Preliminary Results

We conduct a series of preliminary experiments on Avenue to determine the hyperparameters of SSMCTB, namely the dilation rate d and the sub-kernel size k'. We perform experiments with $d \in \{0, 1, 2, 3\}$ and $k' \in \{1, 2, 3\}$. We



Fig. 4. Examples of image-level anomaly localization results from MVTec AD given by DRAEM [56], before (blue contour) and after (green contour) integrating SSMCTB. The ground-truth anomalies are shown in red. Best viewed in color.

also consider alternative attention types, namely channel attention (CA), spatial attention (SA) and both channel and spatial attention (CA+SA). Additionally, we alternate between multiple losses to self-supervise our block, such as the mean absolute error (MAE), the mean squared error (MSE), the Huber loss, and the Structured Similarity Index Measure (SSIM) loss. For the Huber loss, we set the hyperparameter δ to the default value, *i.e.* $\delta = 1$.

We employ the method of Park et al. [44] in our preliminary experiments, since this is the most lightweight method among the chosen ones [42], [44], [55]–[60]. The corresponding micro AUC scores are presented in Table 2. Except for a single SSMCTB configuration based on spatial attention (SA), all other SSMCTB configurations bring performance improvements over the approach of Park et al. [44] (first row). Our first set of preliminary experiments is aimed at evaluating the capacity of the standalone masked convolution. Even without the attention module, our masked convolution brings gains higher than 1% for d = 2 and d = 3. While adding the attention module is definitely useful, we conclude that it is clearly not the only factor responsible for the reported performance gains. To compare the losses on the one hand, and attention types on the other, we fix k' = 1. When alternating between MAE, MSE, SSIM and Huber as our self-supervised loss, we generally observe higher performance with Huber loss. We thus continue the experiments with Huber loss. Regarding the attention type, we note that channel attention (CA) generally leads to better results than spatial attention (SA). Hence, for the remaining experiments, we employ the transformer module based on channel attention. We continue by increasing the size of the sub-kernels, without obtaining further performance gains. We obtain the best micro AUC (86.7%) with d = 1 and k' = 1, while using channel attention. We make another attempt to further boost the performance by combining the channel and spatial attention (CA+SA), while fixing d = 1 and k' = 1. This attempt is also unsuccessful. Our final SSMCTB configuration, which we employ across all underlying models and data sets, is based on d = 1, k' = 1and channel attention.

We underline that the corresponding hyperparameters for SSPCAB were tuned in a similar manner, in our previous work [54]. Hence, we simply use the already tuned hyperparameters for SSPCAB. Importantly, we underline that our observations above are mostly consistent with those reported in our previous work [54], *i.e.* both SSPCAB and SSMCTB use channel attention, a dilation rate of d = 1and sub-kernels of size k' = 1. The only difference is that SSMCTB is based on the Huber loss instead of the MSE



Fig. 5. Examples of image-level anomaly localization results from BRATS given by DRAEM [56], before (blue contour) and after (green contour) integrating SSMCTB. The ground-truth anomalies are shown in red. Best viewed in color.

loss. We should also emphasize that it is not common for anomaly detection data sets to have validation splits. Since the training set contains normal instances only, keeping a representative training subset (with both normal and abnormal examples) for validation is not possible. This is the reason behind our decision to avoid hyperparameter tuning for each model and data set. We believe that this evaluation procedure is more fair because it avoids overfitting in hyperparameter space.

4.5 Anomaly Detection in Images

Baselines. We introduce SSMCTB into two state-of-theart baselines for image anomaly detection on MVTec AD, namely a self-supervised model based on natural synthetic anomalies (NSA) [57] and a discriminatively trained reconstruction anomaly embedding model (DRAEM) [56]. Both baselines are very recent, attaining strong results on MVTec AD. The NSA approach of Shülter et al. [57] generates synthetic anomalies using Poisson image editing, blending scaled patches of different sizes from separate images. In this way, it generates a wide range of synthetic anomalies that are similar to natural irregularities. DRAEM [56] comprises a reconstructive network and a discriminative network to detect and localize anomalies. The reconstructive network is based on a simple auto-encoder architecture which learns to reconstruct original images from artificially corrupted images. The discriminative network is a U-Net that learns to segment the introduced artifacts (corrupted regions).

Results on MVTec AD. We report the results on MVTec AD in Table 3. Considering the detection results, we observe that adding SSPCAB and SSMCTB leads to superior results for both DRAEM [56] and NSA [57]. Considering the localization results, the AUROC scores of DRAEM do not show any improvements when adding SSPCAB and SSMCTB. However, the localization AP of DRAEM exhibits gains of around 2% by adding SSPCAB and SSMCTB. In addition, the localization AUROC of NSA grows when SSPCAB and SSMCTB are introduced into the architecture.

In Figure 4, we present some examples of qualitative results from MVTec AD, obtained by DRAEM [56], before and after adding SSMCTB. In all shown cases, we observe that the anomaly localization results are better aligned with the ground-truth regions when SSMCTB is integrated into DRAEM.

Detection AUROC and localization AUROC/AP (in %) of two state-of-the-art methods [56], [57] on MVTec AD, before and after alternatively adding SSPCAB and SSMCTB. The best result for each model and each performance measure is highlighted in bold.

				Det	ection			Localization								
		DR	AEM	[56]	NSA	(logist	ic) [57]	DRAEM [56]				NSA (logistic) [57]				
		A	AURO	С		AURO	С	A	AURO	2		AP		AUROC		
	Class	Baseline	+SSPCAB	+SSMCTB	Baseline	+SSPCAB	+SSMCTB	Baseline	+SSPCAB	+SSMCTB	Baseline	+SSPCAB	+SSMCTB	Baseline	+SSPCAB	+SSMCTB
re	Carpet Grid	97.0 99.9	98.2 100	96.8 100	95.6 99.9	97.5 99.9	96.1 100	95.5 99.7	95.0 99.5	95.8 99.7	53.5 65.7	59.4 61.1	55.2 69.7	95.5 99.2	97.5 99.2	95.6 99.2
¥tu	Leather	100	100	100	99.9	99.9	100	98.6	99.5	97.6	75.3	76.0	65.5	99.5	99.5	99.6
Te	Tile	99.6	100	100	100	100	100	99.2	99.3	99.3	92.3	95.0	95.7	99.3	99.2	99.1
	Wood	99.1	99.5	100	97.5	97.7	97.8	96.4	96.8	94.8	77.7	77.1	75.6	90.7	90.4	93.5
	Bottle	99.2	98.4	99.4	97.7	97.7	97.7	99.1	98.8	99.2	86.5	87.9	89.9	98.3	98.3	98.4
	Cable	91.8	96.9	94.1	94.5	95.6	96.1	94.7	96.0	95.5	52.4	57.2	61.6	96.0	96.6	97.5
	Capsule	98.5	99.3	97.1	95.2	95.4	95.5	94.3	93.1	93.4	49.4	50.2	52.0	97.6	97.2	97.9
H	Hazelnut	100	100	100	94.7	94.2	97.1	99.7	99.8	99.5	92.9	92.6	89.1	97.6	97.9	97.9
jec	Metal Nut	98.7	100	100	98.7	99.0	99.5	99.5	98.9	99.3	96.3	98.1	94.7	98.4	98.6	98.3
ð	Pill	98.9	99.8	98.8	99.2	99.2	99.5	97.6	97.5	97.4	48.5	52.4	46.9	98.5	98.8	98.4
	Screw	93.9	97.9	99.0	90.2	91.1	90.4	97.6	99.8	99.5	58.2	72.0	70.1	96.5	96.2	96.4
	Toothbrush	100	100	100	100	100	100	98.1	98.1	99.0	44.7	51.0	69.0	94.9	95.3	95.4
	Transistor	93.1	92.9	96.0	95.1	95.6	96.2	90.9	87.0	89.1	50.7	48.0	45.8	88.0	87.1	88.3
	Zipper	100	100	100	99.8	99.8	99.9	98.8	99.0	99.0	81.5	77.1	76.5	94.2	94.5	94.7
	Overall	98.0	98.9	98.7	97.2	97.5	97.7	97.3	97.2	97.2	68.4	70.3	70.5	96.3	96.4	96.7

TABLE 4

Detection AUROC and localization AUROC/AP (in %) of two state-of-the-art methods [56], [57] on BRATS, before and after alternatively adding SSPCAB and SSMCTB. Additional results obtained by converting DRAEM to use 3D convolutions and integrating the 3D SSMCTB are also reported. The best result for each model and each performance measure is highlighted in bold.

	AUI	ROC	L
Method	Detection	Localization	Localization A
DRAEM [56]	41.06	42.40	45.41
DRAEM + SSPCAB [54]	44.19	46.66	46.89
DRAEM + SSMCTB (Ours)	50.27	53.98	50.75
NSA [57]	53.66	74.90	61.09
NSA + SSPCAB [54]	54.91	75.30	62.37
NSA + SSMCTB (Ours)	60.09	77.09	64.55
3D DRAEM [56]	43.74	44.12	45.97
3D DRAEM + 3D SSMCTB (Ours)	53.70	58.47	52.79

Results on BRATS. In Table 4, we present the brain lesion detection and localization results obtained by the anomaly detection models [56], [57] on BRATS, before and after adding SSPCAB and SSMCTB, respectively. Remarkably, we notice that the results of both DRAEM and NSA show significant performance improvements when integrating SSM-CTB. Moreover, the performance gain brought by SSMCTB is always higher than the gain brought by SSPCAB. When taking advantage of the 3D nature of the MRI scans by employing the 3D SSMCTB, we attain even higher performance with DRAEM.

In Figure 5, we present several examples of qualitative results from BRATS, given by DRAEM [56], before and after adding SSMCTB. In general, the localization results



Fig. 6. Frame-level anomaly scores of the method of Georgescu *et al.* [58], before (baseline) and after (ours) integrating SSMCTB, for test video 02 from the Avenue data set. Anomaly localization results correspond to the model based on SSMCTB. Best viewed in color.

based on SSMCTB exhibit a higher overlap with the groundtruth regions, explaining why SSMCTB leads to superior performance levels.

4.6 Anomaly Detection in Videos

Baselines. We select five recent methods [42], [44], [55], [58], [59] yielding state-of-the-art performance on Avenue and ShanghaiTech. Liu *et al.* [42] proposed a GAN-based framework to detect anomalies based on the future frame prediction error. Park *et al.* [44] presented a memory-based auto-encoder classifying anomalies based on the reconstruction error. The model comprises a memory module that memorizes prototypes of normal samples. Liu *et al.* [59] employed a hybrid framework based on flow reconstruc-

Micro-averaged frame-level AUC, macro-averaged frame-level AUC, RBDC, and TBDC scores (in %) of various state-of-the-art methods on Avenue and ShanghaiTech. Among the existing models, we select six models [42], [44], [55], [58]–[60] to show results before and after including SSPCAB and SSMCTB, respectively. The best result for each underlying model is highlighted in bold. The top score for each metric is shown in red.

	Avenue				ShanghaiTech			
Method	A	UC	RBDC	TBDC	A	UC	RBDC	TBDC
	Micro	Macro			Micro	Macro		
Liu <i>et al.</i> [13]	84.4	-	-	-	-	-	-	-
Sultani <i>et al</i> . [102]	-	-	-	-	-	76.5	-	-
Ionescu <i>et al</i> . [18]	88.9	-	-	-	-	-] -	-
Nguyen <i>et al.</i> [43]	86.9	-	-	-	-	-	-	-
Ionescu <i>et al</i> . [17]	87.4	90.4	15.77	27.01	78.7	84.9	20.65	44.54
Wu et al. [103]	80	5.6	-	-	-	-	-	-
Lee <i>et al</i> . [104]	90	0.0	-	-	-	-	-	-
Yu et al. [89]	89.6	-	-	-	74.8	-	-	-
Ramachandra <i>et al</i> . [19]	72	2.0	35.80	80.90	-	-	-	-
Ramachandra <i>et al</i> . [20]	82	7.2	41.20	78.60	-	-	-	-
Tang <i>et al</i> . [47]	8	5.1	-	-	73	3.0	-	-
Dong <i>et al.</i> [105]	84	4.9	-	-	73	3.7	-	-
Doshi <i>et al.</i> [106]	8	5.4	-	-	7	1.6	-	-
Sun <i>et al.</i> [107]	8	9.6	-	-	74	4.7	-	-
Wang <i>et al.</i> [108]	82	7.0	-	-	79	9.3	-	-
Astrid et al. [84]	84.7	-	-	-	73.7	-	-	-
Astrid et al. [109]	87.1	-	-	-	75.9	-	-	-
Georgescu <i>et al</i> . [50]	91.5	92.8	57.00	58.30	82.4	90.2	42.80	83.90
Liu et al. [42]	85.1	81.7	19.59	56.01	72.8	80.6	17.03	54.23
Liu et al. [42] + SSPCAB [54]	87.3	84.5	20.13	62.30	74.5	82.9	18.51	60.22
Liu et al. [42] + SSMCTB (Ours)	89.5	84.6	23.79	66.03	74.6	83.9	19.13	61.65
He et al. [60]	84.0	85.6	-	-	74.3	81.1	-	-
He et al. [60] + SSPCAB [54]	85.1	85.8	-	-	74.5	81.9	-	-
He et al. [60] + SSMCTB (Ours)	86.4	86.5	-	-	76.1	81.6	-	-
Park <i>et al.</i> [44]	82.8	86.8	-	-	68.3	79.7	-	-
Park <i>et al.</i> [44] + SSPCAB [54]	84.8	88.6	-	-	69.8	80.2	-	-
Park et al. [44] + SSMCTB (Ours)	87.0	87.7	-	-	70.6	80.3	-	-
Liu et al. [59]	89.9	93.5	41.05	86.18	74.2	83.2	44.41	83.86
Liu et al. [59] + SSPCAB [54]	90.9	92.2	62.27	89.28	75.5	83.7	45.45	84.50
Liu et al. [59] + SSMCTB (Ours)	89.6	93.9	46.49	86.43	75.2	83.8	45.86	84.69
Georgescu et al. [58]	92.3	90.4	65.05	66.85	82.7	89.3	41.34	78.79
Georgescu et al. [58] + SSPCAB [54]	92.9	91.9	65.99	64.91	83.6	89.5	40.55	83.46
Georgescu <i>et al.</i> [58] + SSMCTB (Ours)	93.2	91.8	66.04	65.12	83.3	89.5	40.52	81.93
Bărbălău <i>et al.</i> [55]	91.6	92.5	47.83	85.26	83.8	90.5	47.14	85.61
Bărbălău <i>et al.</i> [55] + 3D SSMCTB (Ours)	91.6	92.4	49.01	85.94	83.7	90.6	47.73	85.68

TABLE 6 Micro and macro AUC scores (in %) on Thermal Rare Event, obtained while alternatively including SSPCAB [54] and SSMCTB into the method of Park *et al.* [44].

Method	AUC				
Method	Micro	Macro			
Park <i>et al.</i> [44]	53.2	66.5			
Park et al. [44] + SSPCAB	53.6	66.6			
Park <i>et al.</i> [44] + SSMCTB (Ours)	58.9	66.6			

tion and frame prediction, using the accumulated error to detect anomalies. Georgescu *et al.* [58] introduced a training scheme where the latent subspaces of appearance and motion auto-encoders are improved by performing gradient ascent on pseudo-anomalies during training. Bărbălău *et al.* [55] extended the previous work of Georgescu *et al.* [50] with two 3D transformer-based self-supervised multi-task architectures trained on new sets of proxy tasks. Among the two versions proposed in [55], we opt for SSMTL++v2. We included this 3D model [55] because it serves as a good baseline for applying our 3D SSMCTB.

We also experiment with the recently proposed masked

auto-encoder framework [60], which is based on the ViT backbone [52]. We add this baseline model to demonstrate the applicability of SSMCTB to vision transformers.

Results on RGB videos. We present the results on Avenue and ShanghaiTech in Table 5. As for the image anomaly detection experiments, we compare the results of the underlying models before and after adding SSPCAB [54] and SSMCTB, respectively. For the method of Liu et al. [42], both SSPCAB and SSMCTB lead to performance improvements, but the gains brought by SSMCTB are always higher than those brought by SSPCAB. Since the methods of He et al. [60] and Park et al. [44] are only capable of detecting anomalies at the frame level, we only report their frame-level micro and macro AUC scores. The vanilla masked autoencoder obtains competitive results on both Avenue and ShanghaiTech. On Avenue, SSMCTB brings higher gains to the masked auto-encoder than SSPCAB. On ShanghaiTech, SSMCTB is better than SSPCAB in terms of the micro AUC, but SSPCAB exhibits higher macro AUC gains. In summary, both SSMCTB and SSPCAB improve the masked autoencoder, with SSMCTB having the upper hand. Considering the results of Park et al. [44] on Avenue, SSMCTB leads to



Fig. 7. Frame-level anomaly scores of the method of Georgescu *et al.* [58], before (baseline) and after (ours) integrating SSMCTB, for test video 10 from the Avenue data set. Anomaly localization results correspond to the model based on SSMCTB. Best viewed in color.



Fig. 8. Frame-level anomaly scores of the method of Liu *et al.* [59], before (baseline) and after (ours) integrating SSMCTB, for test video 02_0164 from the ShanghaiTech data set. Anomaly localization results correspond to the model based on SSMCTB. Best viewed in color.

higher gains in terms of the micro AUC (from 82.8% to 87.0%), while SSPCAB leads to a higher macro AUC (from 86.8% to 88.6%). On ShanghaiTech, we observe higher gains after adding SSMCTB rather than SSPCAB. Moving on to the object-centric models of Liu et al. [59] and Georgescu et al. [58], we observe that the top gains are mainly shared between SSPCAB and SSMCTB. When integrating our 3D SSMCTB into the 3D architecture presented in [55], we observe performance improvements according to most metrics. Overall, SSMCTB leads to the highest performance levels on Avenue for three metrics, namely the micro AUC (93.2%), the macro AUC (93.9%) and the RBDC (66.04%). At the same time, SSPCAB attains the highest TBDC score (89.28%) on Avenue. On ShanghaiTech, it appears that the best scores are obtained by adding the 3D SSMCTB into the underlying model of Bărbălău et al. [55], since our 3D SSMCTB brings performance gains for three metrics.

In Figure 6 and 7, we illustrate the anomaly detection performance on two test videos from Avenue, before and



Fig. 9. Frame-level anomaly scores of the method of Park *et al.* [44], before (baseline) and after (ours) integrating SSMCTB, for test video 39 from the Thermal Rare Event data set. Anomaly localization results correspond to the model based on SSMCTB. Best viewed in color.

after integrating SSMCTB into the model of Georgescu et al. [58]. Our approach produces superior frame-level anomaly scores, being able to detect the person running in the first video (Figure 6) and the person throwing an object in the second one (Figure 7). Moreover, in the second video, we also notice that SSMCTB increases the anomaly score for the penultimate abnormal event, resolving the false negative detection of the baseline. Similarly, in Figure 8, we show the effect of adding SSMCTB into the architecture of Liu et al. [59] applied on a test video from ShanghaiTech. Once again, SSMCTB improves the frame-level detection performance, being able to detect the person riding a bike in a pedestrian area, which is forbidden. SSMCTB correctly raises the anomaly scores for about 50 video frames, starting at around frame index 150, thus reducing the false negative rate.

Results on thermal videos. Since texture is not present in the thermal domain, there is no need to apply very deep architectures, as noticed by Nikolov et al. [62]. Moreover, object detectors pre-trained on natural images do not work equally well in the thermal domain due to the distribution shift. To this end, the object-centric [55], [58], [59] and very deep [42] baselines attain very poor results (micro AUC values under 50%). Hence, we resort to employing the architecture of Park et al. [44] as underlying model for SSPCAB and SSMCTB. As shown in Table 6, the chosen baseline attains a micro AUC of 53.2% and a macro AUC of 66.5%. Both SSPCAB and SSMCTB seem to have a positive influence on the micro AUC score, but the gains of the latter block are significantly higher (above 5%). In summary, the results reported on Thermal Rare Event demonstrate the utility of SSMCTB, further confirming the gains observed on RGB video data sets.

In Figure 9, we show the anomaly detection performance on a test video from Thermal Rare Event, before and after integrating SSMCTB into the model of Park *et al.* [44]. SSMCTB leads to important gains in terms of the frame-level scores, being able to detect the vehicle moving backwards. **Inference time.** Regardless of the underlying framework [42], [44], [55]–[60], similar to Ristea *et al.* [54], we add only

Inference time (in milliseconds) per example for three frameworks [42], [56], [58], before and after integrating SSPCAB and SSMCTB, respectively. The running times are measured on an Nvidia GeForce GTX 3090 GPU with 24 GB of VRAM.

		Time (ms)						
Method	Baseline	+SSPCAB	+SSMCTB	+3D SSMCTB				
Liu et al. [42]	2.1	2.4	2.5	-				
Georgescu et al. [58]	1.5	1.7	1.8	-				
Zavrtanik <i>et al</i> . [56]	26.4	-	-	26.6				

TABLE 8 Micro AUC (in %) on Avenue by incorporating SSMCTB into different conv blocks of the decoder proposed by Park *et al.* [44]. Along with the block placement, we also vary the dilation rate *d*.

Method	Decoder Conv Block	d	Micro AUC
Park et al. [44]	-	-	82.8
	early	0	83.6
	early	1	83.2
	early	2	84.6
	early	3	84.7
	early	4	84.9
	middle	0	84.3
	middle	1	83.2
SCMCTP	middle	2	84.5
+551VIC1D	middle	3	85.9
	middle	4	85.7
	late	0	86.2
	late	1	87.0
	late	2	85.9
	late	3	84.6
	late	4	85.8
	all	4,3,1	85.1

one instance of SSMCTB, usually replacing the penultimate convolutional layer. Considering that the channel attention from SSPCAB is replaced with a channel-wise transformer block in SSMCTB, we might expect a slightly higher processing time. To assess the amount of extra time added by SSMCTB, we present the running times before and after integrating SSPCAB and SSMCTB into two state-of-the-art frameworks [42], [58] in Table 7. For both baseline models, the time added by SSMCTB is at most 0.1 ms higher than the time taken by SSPCAB. Moreover, the computational time of SSMCTB does not exceed a difference of 0.4 ms with respect to the original baselines. Another important question is how does the 3D version of SSMCTB impact the running time. To answer this question, we take the DRAEM model [56] and measure the running time before and after adding the 3D SSMCTB. The reported time measurements show that the running time increase due to the 3D SSMCTB is still marginal, being around 0.2 ms. Hence, the processing delays caused by the introduction of the 2D or 3D SSMCTB versions are within the same range. In summary, we consider that the accuracy gains brought by SSMCTB outweigh the marginal running time expansions reported in Table 7.

TABLE 9

Micro AUC (in %) on Avenue by incorporating SSMCTB into the model of Park *et al.* [44], while varying the size of the masked region *M*.

Method	Size of M	Micro AUC
Park <i>et al</i> . [44]	-	82.8
	1×1	87.0
+SSMCTB	2×2	85.6
	3×3	84.9

4.7 Ablation Study

Block placement. Across all the experiments presented so far, recall that we introduce a single SSMCTB, which is usually placed near the end of the architecture (penultimate convolutional layer), as mentioned in Section 4.3. The number of blocks as well as their placement should be tuned on some validation set, which could lead to higher performance gains. However, anomaly detection data sets do not commonly contain a validation set and there is no way to keep a number of training samples for validation, as the training set comprises only normal examples. To this end, we employed a single configuration (one block, closer to the output) to fairly demonstrate the universality of SSMCTB. Certainly, this choice might not always be optimal. Hence, we perform ablation experiments by incorporating SSMCTB at different decoder levels of the network proposed by Park et al. [44], considering different dilation rates (d). We vary the dilation rate along with the block placement, because Duță et al. [110] observed that higher dilation rates are suitable for earlier dilated convolutional layers, and lower dilation rates are suitable for dilated convolutional layers closer to the output.

In Table 8, we show the corresponding results on the Avenue data set. We start by adding SSMCTB into the earliest stage of the decoder (first conv block), progressively moving the block to the layers closer to the output of the decoder, until we reach the very last one. For each decoder level (early, middle, late), we vary the dilation rate to find a suitable value. We attain the best micro AUC (87.0%)when integrating SSMCTB into the last conv block of the decoder, while using a dilation rate of d = 1. A dilation rate of d = 4 seems suitable when placing SSMCTB at an earlier stage, while, for the middle stage placement, the optimal dilation rate appears to be d = 3. Interestingly, these results are consistent with the observation made by Duță et al. [110], although their observation applies to dilated convolutions, while ours applies to masked convolutions. Nevertheless, all the results are consistently better than the baseline (82.8%), regardless of the block placement or the dilation rate. We do not observe major improvements when integrating multiple blocks, concluding that integrating a single SSMCTB is sufficient.

Size of masked region. Increasing the size of the masked region M can lead to a harder reconstruction task, at each location where our masked convolution is applied. However, it is unclear if making the task harder leads to better results. To this end, we vary the spatial size of M, considering three options: 1×1 , 2×2 and 3×3 . We present the corresponding results in Table 9. The empirical results indicate that increasing the size of M leads to lower anomaly detection scores. Hence, we conclude that a size of

Micro AUC (in %) on Avenue by incorporating SSMCTB into the model of Park *et al.* [44], while varying the hyperparameters of the channel-wise transformer, namely the activation map size $(h' \times w')$ after the average pooling layer, the token size (d_t) after the projection layer, the number of heads (H), as well as the number of successive transformer blocks (L).

Method	$h' \times w'$	d_t	H	L	Micro AUC
Park et al. [44]	-	-	-	-	82.8
	1×1	64	4	2	87.0
	2×2	64	4	2	85.3
	3×3	64	4	2	85.2
	4×4	64	4	2	85.6
	1×1	16	4	2	84.6
	1×1	32	4	2	85.6
	1×1	64	4	2	87.0
⊥ SSMCTB	1×1	128	4	2	85.1
+JJJVIC I D	1×1	64	3	2	85.6
	1×1	64	4	2	87.0
	1×1	64	5	2	87.0
	1×1	64	6	2	84.8
	1×1	64	4	1	85.1
	1×1	64	4	2	87.0
	1×1	64	4	3	84.0

1×1 for the masked region *M* is optimal.

Transformer architecture. In Table 10, we present further ablation experiments for the channel-wise transformer module. We keep the underlying model of Park *et al.* [44] and report the results on the Avenue data set. As variations for the transformer module, we consider the following hyperparameters: the activation map size $(h' \times w')$ after the average pooling layer, the token size (d_t) after the projection layer, the number of heads (H), as well as the number of successive transformer blocks (L).

First, we analyze how activation maps of different dimensions, given as output by the average pooling layer placed right before the transformer, influence the results. We observe that shrinking the maps to 1×1 gives the best micro AUC (87.0%). The optimal configuration of the average pooling layer (producing activation maps of 1×1) is equivalent to global average pooling. For the projection layer, we consider output dimensions in the set $d_t \in \{16, 32, 64, 128\}$. The optimal size for the projection layer is $d_t = 64$. We consider transformer modules having 3 to 6 heads. The empirical evidence indicates that using H = 4 or H = 5heads leads to equally good results. Finally, we experiment with transformer modules having 1 to 3 blocks. The best performance is achieved with L = 2 successive transformer blocks. We underline that all transformer configurations surpass the baseline model [44].

Huber loss hyperparameter. Huber loss is the combination of the L_1 (MAE) and L_2 (MSE) losses (see Eq. (7)), where δ is a hyperparameter representing the threshold that switches between the two loss functions. To study the effect of δ , we consider different values for the hyperparameter $\delta \in \{0.5, 1, 2\}$, reporting the results in Table 11. We find that the maximum improvement corresponds to $\delta = 1$, but the other values of δ also lead to superior results compared to the baseline.

Comparison with dilated convolution. In Table 12, we compare the dilated convolution against the proposed masked convolution, alternating between the two operations inside

Micro AUC (in %) on Avenue by incorporating SSMCTB into the model of Park *et al.* [44], while varying the hyperparameter δ of the Huber loss.

Method	δ	Micro AUC
Park <i>et al</i> . [44]	-	82.8
	0.5	84.1
+SSMCTB	1	87.0
	2	85.8

TABLE 12
Micro AUC (in %) on Avenue by incorporating SSMCTB into the model
of Park et al. [44], while switching between dilated and masked
convolution. Different values for the dilation rate d are tested for the two
operations.

Method	d	Micro AUC
Park <i>et al</i> . [44]	-	82.8
+SSDCTB (dilated conv)	1	85.1
	2	83.3
	3	85.0
+SSMCTB (masked conv)	1	87.0
	2	85.5
	3	85.9

SSMCTB. We denote the block based on dilated convolution through the acronym SSDCTB. When comparing the two convolutional operations, we consider multiple dilation rates between 1 and 3. The experiments show that the proposed masked convolution outperforms the dilated convolution, regardless of the dilation rate. This confirms that the two operations are not equivalent, essentially revealing the importance of the self-supervised task based on reconstructing the masked region M situated in the center of the receptive field.

5 CONCLUSION

In this paper, we extended our previous work [54] by introducing SSMCTB, a novel neural block composed of a masked convolutional layer and a channel-wise transformer module, which predicts a masked region in the center of the convolutional receptive field. Our neural block is trained in a self-supervised manner, via a reconstruction loss of its own. To show the benefits of using SSMCTB in anomaly detection, we integrated our block into a series of image and video anomaly detection methods [42], [44], [55]–[60]. In addition, we included two new benchmarks from domains that were not previously considered by Ristea et al. [54], namely medical images and thermal videos. Moreover, we extended the 2D masked convolution to a 3D masked convolution, broadening the applicability of the self-supervised block to 3D neural architectures. To showcase the utility of the new 3D SSMCTB, we integrated our 3D block into two 3D networks (3D DRAEM and SSMTL++v2) for anomaly detection in image and video, respectively. Our empirical results across multiple benchmarks and underlying models indicate that SSMCTB brings performance improvements in a vast majority of cases. Furthermore, with the help of SSMCTB, we are able to obtain new state-of-the-art levels on the widely-used Avenue and ShanghaiTech data sets. We consider this as a major achievement, which would not have been possible without SSMCTB.

In future work, we aim to apply our novel selfsupervised block on other tasks, aside from anomaly detection. For example, due to the self-supervised loss computed with respect to the masked region, our block could be integrated into various neural architectures to perform self-supervised pre-training, before applying the respective models to downstream tasks. Interestingly, the pretraining could be performed at multiple architectural levels, *i.e.* wherever the block is added into the model.

ACKNOWLEDGMENTS

This work was supported by a grant of the Romanian Ministry of Education and Research, CNCS - UEFISCDI, project no. PN-III-P2-2.1-PED-2021-0195, contract no. 690/2022, within PNCDI III. The research leading to these results has also received funding from the NO Grants 2014-2021, under project ELO-Hyp contract no. 24/2020. Additionally, this work has been funded by Milestone Systems through the Milestone Research Programme at AAU, and by SecurifAI.

REFERENCES

- P. Bergmann, M. Fauser, D. Sattlegger, and C. Steger, "MVTec AD – A Comprehensive Real-World Dataset for Unsupervised Anomaly Detection," in *Proceedings of CVPR*, pp. 9592–9600, 2019.
- [2] S. Lee, S. Lee, and B. C. Song, "CFA: Coupled-hyperspherebased feature adaptation for target-oriented anomaly localization," *IEEE Access*, vol. 10, pp. 78446–78454, 2022.
 [3] W. Luo, W. Liu, and S. Gao, "A Revisit of Sparse Coding Based
- [3] W. Luo, W. Liu, and S. Gao, "A Revisit of Sparse Coding Based Anomaly Detection in Stacked RNN Framework," in *Proceedings* of ICCV, pp. 341–349, 2017.
- [4] N. Shvetsova, B. Bakker, I. Fedulova, H. Schulz, and D. V. Dylov, "Anomaly Detection in Medical Imaging With Deep Perceptual Autoencoders," *IEEE Access*, vol. 9, pp. 118571–118583, 2021.
- [5] D. Carrera, F. Manganini, G. Boracchi, and E. Lanzarone, "Defect Detection in SEM Images of Nanofibrous Materials," *IEEE Transactions on Industrial Informatics*, vol. 13, no. 2, pp. 551–561, 2017.
- [6] K.-W. Cheng, Y.-T. Chen, and W.-H. Fang, "Video anomaly detection and localization using hierarchical feature representation and Gaussian process regression," in *Proceedings of CVPR*, pp. 2909–2917, 2015.
 [7] Y. Cong, J. Yuan, and J. Liu, "Sparse reconstruction cost for
- [7] Y. Cong, J. Yuan, and J. Liu, "Sparse reconstruction cost for abnormal event detection," in *Proceedings of CVPR*, pp. 3449– 3456, 2011.
- [8] J. K. Dutta and B. Banerjee, "Online Detection of Abnormal Events Using Incremental Coding Length," in *Proceedings of* AAAI, pp. 3755–3761, 2015.
- [9] C. Lu, J. Shi, and J. Jia, "Abnormal Event Detection at 150 FPS in MATLAB," in *Proceedings of ICCV*, pp. 2720–2727, 2013.
- [10] H. Ren, W. Liu, S. I. Olsen, S. Escalera, and T. B. Moeslund, "Unsupervised Behavior-Specific Dictionary Learning for Abnormal Event Detection," in *Proceedings of BMVC*, pp. 28.1–28.13, 2015.
- [11] A. Del Giorno, J. Bagnell, and M. Hebert, "A Discriminative Framework for Anomaly Detection in Large Videos," in *Proceed*ings of ECCV, pp. 334–349, 2016.
- [12] R. T. Ionescu, S. Smeureanu, B. Alexe, and M. Popescu, "Unmasking the abnormal events in video," in *Proceedings of ICCV*, pp. 2895–2903, 2017.
- [13] Y. Liu, C.-L. Li, and B. Póczos, "Classifier Two-Sample Test for Video Anomaly Detections," in *Proceedings of BMVC*, 2018.
- [14] G. Pang, C. Yan, C. Shen, A. v. d. Hengel, and X. Bai, "Selftrained Deep Ordinal Regression for End-to-End Video Anomaly Detection," in *Proceedings of CVPR*, pp. 12173–12182, 2020.
- [15] P. Bergmann, M. Fauser, D. Sattlegger, and C. Steger, "Uninformed Students: Student-Teacher Anomaly Detection With Discriminative Latent Embeddings," in *Proceedings of CVPR*, pp. 4183–4192, 2020.
 [16] T. Defard, A. Setkov, A. Loesch, and R. Audigier, "PaDiM: A
- [16] T. Defard, A. Setkov, A. Loesch, and R. Audigier, "PaDiM: A patch distribution modeling framework for anomaly detection and localization," in *Proceedings of ICPR*, pp. 475–489, 2021.

- [17] R. T. Ionescu, F. S. Khan, M.-I. Georgescu, and L. Shao, "Object-Centric Auto-Encoders and Dummy Anomalies for Abnormal Event Detection in Video," in *Proceedings of CVPR*, pp. 7842–7851, 2019.
- [18] R. T. Ionescu, S. Smeureanu, M. Popescu, and B. Alexe, "Detecting abnormal events in video using Narrowed Normality Clusters," in *Proceedings of WACV*, pp. 1951–1960, 2019.
- [19] B. Ramachandra and M. Jones, "Street Scene: A new dataset and evaluation protocol for video anomaly detection," in *Proceedings* of WACV, pp. 2569–2578, 2020.
- [20] B. Ramachandra, M. Jones, and R. Vatsavai, "Learning a distance function with a Siamese network to localize anomalies in videos," in *Proceedings of WACV*, pp. 2598–2607, 2020.
- [21] M. Ravanbakhsh, M. Nabi, H. Mousavi, E. Sangineto, and N. Sebe, "Plug-and-Play CNN for Crowd Motion Analysis: An Application in Abnormal Event Detection," in *Proceedings of* WACV, pp. 1689–1698, 2018.
- [22] M. Sabokrou, M. Fayyaz, M. Fathy, and R. Klette, "Deep-Cascade: Cascading 3D Deep Neural Networks for Fast Anomaly Detection and Localization in Crowded Scenes," *IEEE Transactions on Image Processing*, vol. 26, no. 4, pp. 1992–2004, 2017.
- [23] M. Sabokrou, M. Fayyaz, M. Fathy, Z. Moayed, and R. Klette, "Deep-Anomaly: Fully Convolutional Neural Network for Fast Anomaly Detection in Crowded Scenes," *Computer Vision and Image Understanding*, vol. 172, pp. 88–97, 2018.
- [24] V. Saligrama and Z. Chen, "Video anomaly detection based on local statistical aggregates," in *Proceedings of CVPR*, pp. 2112– 2119, 2012.
- [25] S. Smeureanu, R. T. Ionescu, M. Popescu, and B. Alexe, "Deep Appearance Features for Abnormal Behavior Detection in Video," in *Proceedings of ICIAP*, vol. 10485, pp. 779–789, 2017.
- [26] Q. Sun, H. Liu, and T. Harada, "Online growing neural gas for anomaly detection in changing surveillance scenes," *Pattern Recognition*, vol. 64, pp. 187–201, Apr. 2017.
- [27] H. T. Tran and D. Hogg, "Anomaly Detection using a Convolutional Winner-Take-All Autoencoder," in *Proceedings of BMVC*, 2017.
- [28] A. Adam, E. Rivlin, I. Shimshoni, and D. Reinitz, "Robust Real-Time Unusual Event Detection Using Multiple Fixed-Location Monitors," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 30, no. 3, pp. 555–560, 2008.
- *Intelligence*, vol. 30, no. 3, pp. 555–560, 2008.
 [29] B. Antic and B. Ommer, "Video parsing for abnormality detection," in *Proceedings of ICCV*, pp. 2415–2422, 2011.
- [30] Y. Feng, Y. Yuan, and X. Lu, "Learning deep event models for crowd anomaly detection," *Neurocomputing*, vol. 219, pp. 548– 556, 2017.
- [31] R. Hinami, T. Mei, and S. Satoh, "Joint Detection and Recounting of Abnormal Events by Learning Deep Generic Knowledge," in *Proceedings of ICCV*, pp. 3639–3647, 2017.
- [32] J. Kim and K. Grauman, "Observe locally, infer globally: A spacetime MRF for detecting abnormal activities with incremental updates," in *Proceedings of CVPR*, pp. 2921–2928, 2009.
- [33] V. Mahadevan, W.-X. LI, V. Bhalodia, and N. Vasconcelos, "Anomaly Detection in Crowded Scenes," in *Proceedings of CVPR*, pp. 1975–1981, 2010.
- [34] R. Mehran, A. Oyama, and M. Shah, "Abnormal crowd behavior detection using social force model," in *Proceedings of CVPR*, pp. 935–942, 2009.
- [35] M. Rudolph, B. Wandt, and B. Rosenhahn, "Same Same But DifferNet: Semi-Supervised Defect Detection with Normalizing Flows," in *Proceedings of WACV*, pp. 1907–1916, 2021.
- [36] B. Saleh, A. Farhadi, and A. Elgammal, "Object-Centric Anomaly Detection by Attribute-Based Reasoning," in *Proceedings of CVPR*, pp. 787–794, 2013.
- [37] S. Wu, B. E. Moore, and M. Shah, "Chaotic Invariants of Lagrangian Particle Trajectories for Anomaly Detection in Crowded Scenes," in *Proceedings of CVPR*, pp. 2054–2060, 2010.
- [38] Y. Fei, C. Huang, C. Jinkun, M. Li, Y. Zhang, and C. Lu, "Attribute Restoration Framework for Anomaly Detection," *IEEE Transactions on Multimedia*, vol. 24, pp. 116–127, 2022.
- [39] D. Gong, L. Liu, V. Le, B. Saha, M. R. Mansour, S. Venkatesh, and A. Van Den Hengel, "Memorizing Normality to Detect Anomaly: Memory-Augmented Deep Autoencoder for Unsupervised Anomaly Detection," in *Proceedings of ICCV*, pp. 1705–1714, 2019.

- [40] M. Hasan, J. Choi, J. Neumann, A. K. Roy-Chowdhury, and L. S. Davis, "Learning temporal regularity in video sequences," in Proceedings of CVPR, pp. 733-742, 2016.
- [41] Z. Li, N. Li, K. Jiang, Z. Ma, X. Wei, X. Hong, and Y. Gong, "Superpixel Masking and Inpainting for Self-Supervised Anomaly Detection," in Proceedings of BMVC, 2020.
- W. Liu, W. Luo, D. Lian, and S. Gao, "Future Frame Prediction for [42] Anomaly Detection - A New Baseline," in Proceedings of CVPR, pp. 6536-6545, 2018.
- T.-N. Nguyen and J. Meunier, "Anomaly Detection in Video [43] Sequence With Appearance-Motion Correspondence," in Proceedings of ICCV, pp. 1273-1283, 2019.
- H. Park, J. Noh, and B. Ham, "Learning Memory-guided Normal-ity for Anomaly Detection," in *Proceedings of CVPR*, pp. 14372-[44] 14381, 2020.
- [45] M. Ravanbakhsh, M. Nabi, E. Sangineto, L. Marcenaro, C. Regazzoni, and N. Sebe, "Abnormal Event Detection in Videos using Generative Adversarial Nets," in Proceedings of ICIP, pp. 1577-1581, 2017.
- M. Salehi, N. Sadjadi, S. Baselizadeh, M. H. Rohban, and H. R. [46] Rabiee, "Multiresolution Knowledge Distillation for Anomaly Detection," in Proceedings of CVPR, pp. 14902–14912, 2021.
- Y. Tang, L. Zhao, S. Zhang, C. Gong, G. Li, and J. Yang, "Integrat-[47] ing prediction and reconstruction for anomaly detection," Pattern Recognition Letters, vol. 129, pp. 123-130, 2020.
- [48] S. Venkataramanan, K.-C. Peng, R. V. Singh, and A. Mahalanobis, "Attention guided anomaly localization in images," in Proceed*ings of ECCV*, pp. 485–503, 2020. P. J. Huber, "Robust Estimation of a Location Parameter," *The*
- [49] Annals of Mathematical Statistics, vol. 35, no. 1, pp. 73–101, 1964.
- M.-I. Georgescu, A. Bărbălău, R. T. Ionescu, F. S. Khan, [50] M. Popescu, and M. Shah, "Anomaly Detection in Video via Self-Supervised and Multi-Task Learning," in Proceedings of CVPR, pp. 12742-12752, 2021.
- A. Acsintoae, A. Florescu, M. Georgescu, T. Mare, P. Sume-drea, R. T. Ionescu, F. S. Khan, and M. Shah, "Ubnormal: New [51] benchmark for supervised open-set video anomaly detection," in Proceedings of CVPR, pp. 20143-20153, 2022.
- [52] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al., "An image is worth 16x16 words: Transformers for image recognition at scale," in Proceedings of ICLR, 2021.
- A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. [53] Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in Proceedings of NIPS, vol. 30, 2017.
- [54] N.-C. Ristea, N. Madan, R. T. Ionescu, K. Nasrollahi, F. S. Khan, T. B. Moeslund, and M. Shah, "Self-Supervised Predictive Convolutional Attentive Block for Anomaly Detection," in Proceedings of CVPR, pp. 13576-13586, 2022.
- A. Bărbălău, R. T. Ionescu, M.-I. Georgescu, J. Dueholm, B. Ra-[55] machandra, K. Nasrollahi, F. S. Khan, T. B. Moeslund, and M. Shah, "SSMTL++: Revisiting Self-Supervised Multi-Task Learning for Video Anomaly Detection," Computer Vision and Image Understanding, vol. 229, p. 103656, 2023.
- [56] V. Zavrtanik, M. Kristan, and D. Skocaj, "DRAEM A Discriminatively Trained Reconstruction Embedding for Surface
- Anomaly Detection," in *Proceedings of ICCV*, pp. 8330–8339, 2021. H. M. Schlüter, J. Tan, B. Hou, and B. Kainz, "Natural synthetic [57] anomalies for self-supervised anomaly detection and localization," in Proceedings of ECCV, 2022.
- M. I. Georgescu, R. Ionescu, F. S. Khan, M. Popescu, and M. Shah, [58] "A Background-Agnostic Framework with Adversarial Training for Abnormal Event Detection in Video," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 44, no. 9, pp. 4505-4523, 2022.
- [59] Z. Liu, Y. Nie, C. Long, Q. Zhang, and G. Li, "A Hybrid Video Anomaly Detection Framework via Memory-Augmented Flow Reconstruction and Flow-Guided Frame Prediction," in Proceedings of ICCV, pp. 13588-13597, 2021.
- [60] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked Autoencoders Are Scalable Vision Learners," in Proceedings of CVPR, pp. 16000-16009, 2022.
- [61] B. H. Menze, A. Jakab, S. Bauer, J. Kalpathy-Cramer, K. Farahani, J. Kirby, Y. Burren, N. Porz, J. Slotboom, R. Wiest, L. Lanczi, E. Gerstner, M.-A. Weber, T. Arbel, B. B. Avants, N. Ayache, P. Buendia, D. L. Collins, N. Cordier, J. J. Corso, A. Criminisi, T. Das, H. Delingette, C. Demiralp, C. R. Durst, M. Dojat,

- S. Doyle, J. Festa, F. Forbes, E. Geremia, B. Glocker, P. Golland, X. Guo, A. Hamamci, K. M. Iftekharuddin, R. Jena, N. M. John, E. Konukoglu, D. Lashkari, J. A. Mariz, R. Meier, S. Pereira, D. Precup, S. J. Price, T. R. Raviv, S. M. S. Reza, M. Ryan, D. Sarikaya, L. Schwartz, H.-C. Shin, J. Shotton, C. A. Silva, N. Sousa, N. K. Subbanna, G. Szekely, T. J. Taylor, O. M. Thomas, N. J. Tustison, G. Unal, F. Vasseur, M. Wintermark, D. H. Ye, L. Zhao, B. Zhao, D. Zikic, M. Prastawa, M. Reyes, and K. Van Leemput, "The Multimodal Brain Tumor Image Segmentation Benchmark (BraTS)," IEEE Transactions on Medical Imaging, vol. 34, no. 10, pp. 1993-2024, 2015.
- I. Nikolov, M. Philipsen, J. Liu, J. Dueholm, A. Johansen, K. Nas-rollahi, and T. Moeslund, "Seasons in Drift: A Long-Term Ther-[62] mal Imaging Dataset for Studying Concept Drift," in Proceedings of NeurIPS, 2021.
- J. Hu, L. Shen, and G. Sun, "Squeeze-and-Excitation Networks," [63] in Proceedings of CVPR, pp. 7132–7141, 2018.
- [64] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, "End-to-end object detection with transformers," in Proceedings of ECCV, pp. 213–229, 2020.
- J. Chen, Y. Lu, Q. Yu, X. Luo, E. Adeli, Y. Wang, L. Lu, A. L. Yuille, and Y. Zhou, "TransUNet: Transformers Make [65] Strong Encoders for Medical Image Segmentation," arXiv preprint arXiv:2102.04306, 2021.
- S. Khan, M. Naseer, M. Hayat, S. W. Zamir, F. S. Khan, and [66] M. Shah, "Transformers in Vision: A Survey," ACM Computing Surveys, 2021.
- [67] N. Parmar, A. Vaswani, J. Uszkoreit, L. Kaiser, N. Shazeer, A. Ku, and D. Tran, "Image transformer," in Proceedings of ICML, pp. 4055-4064, 2018.
- [68] N.-C. Ristea, A.-I. Miron, O. Savencu, M.-I. Georgescu, N. Verga, F. S. Khan, and R. T. Ionescu, "CyTran: Cycle-Consistent Transformers for Non-Contrast to Contrast CT Translation," arXiv preprint arXiv:2110.06400, 2021.
- H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, and [69] H. Jégou, "Training data-efficient image transformers & distillation through attention," in Proceedings of ICML, pp. 10347–10357, 2021.
- H. Wu, B. Xiao, N. Codella, M. Liu, X. Dai, L. Yuan, and L. Zhang, [70] "CvT: Introducing Convolutions to Vision Transformers," in Proceedings of ICCV, pp. 22-31, 2021.
- X. Xu and N. Xu, "Hierarchical Image Generation via [71] Transformer-Based Sequential Patch Selection," in Proceedings of AAAI, pp. 2938–2945, 2022.
- B. Zhang, S. Gu, B. Zhang, J. Bao, D. Chen, F. Wen, Y. Wang, and [72] B. Guo, "StyleSwin: Transformer-based GAN for High-resolution Image Generation," in Proceedings of CVPR, pp. 11304-11314, 2022.
- [73] M. Zheng, P. Gao, X. Wang, H. Li, and H. Dong, "End-toend object detection with adaptive clustering transformer," in Proceedings of BMVC, 2020.
- [74] X. Zhu, W. Su, L. Lu, B. Li, X. Wang, and J. Dai, "Deformable DETR: Deformable Transformers for End-to-End Object Detection," in Proceedings of ICLR, 2020.
- J. Jiang, J. Zhu, M. Bilal, Y. Cui, N. Kumar, R. Dou, F. Su, and [75] X. Xu, "Masked Swin Transformer Unet for Industrial Anomaly Detection," IEEE Transactions on Industrial Informatics, vol. 19, no. 2, pp. 2200-2209, 2023.
- Y. Lee and P. Kang, "AnoViT: Unsupervised anomaly detec-[76] tion and localization with vision transformer-based encoderdecoder," IEEE Access, vol. 10, pp. 46717-46724, 2022.
- P. Mishra, R. Verk, D. Fornasier, C. Piciarelli, and G. L. Foresti, [77] "VT-ADL: A vision transformer network for image anomaly detection and localization," in Proceedings of ISIE, pp. 1-6, IEEE, 2021.
- J. Pirnay and K. Chai, "Inpainting transformer for anomaly [78] detection," in Proceedings of ICIAP, pp. 394-406, 2022.
- J. Xu, H. Wu, J. Wang, and M. Long, "Anomaly transformer: [79] Time series anomaly detection with association discrepancy," in Proceedings of ICLR, 2022.
- [80] D. Pathak, P. Krähenbühl, J. Donahue, T. Darrell, and A. Efros, "Context Encoders: Feature Learning by Inpainting," in Proceedings of CVPR, pp. 2536-2544, 2016.
- [81] C. Wei, H. Fan, S. Xie, C.-Y. Wu, A. Yuille, and C. Feichtenhofer, "Masked Feature Prediction for Self-Supervised Visual Pre-Training," in Proceedings of CVPR, pp. 14668–14678, 2022.

- [82] H. Chang, H. Zhang, L. Jiang, C. Liu, and W. T. Freeman, "MaskGIT: Masked Generative Image Transformer," in *Proceed*ings of CVPR, pp. 11315–11325, 2022.
- [83] X. Yu, L. Tang, Y. Rao, T. Huang, J. Zhou, and J. Lu, "Pointbert: Pre-training 3d point cloud transformers with masked point modeling," in *Proceedings of CVPR*, pp. 19313–19322, 2022.
- [84] M. Astrid, M. Z. Zaheer, and S.-I. Lee, "Synthetic Temporal Anomaly Guided End-to-End Video Anomaly Detection," in *Proceedings of ICCVW*, pp. 207–214, 2021.
- [85] K. Roth, L. Pemula, J. Zepeda, B. Schölkopf, T. Brox, and P. V. Gehler, "Towards Total Recall in Industrial Anomaly Detection," in *Proceedings of CVPR*, pp. 14298–14308, 2022.
- [86] S. Yamada, S. Kamiya, and K. Hotta, "Reconstructed Student-Teacher and Discriminative Networks for Anomaly Detection," in *Proceedings of IROS*, pp. 2725–2732, 2022.
- [87] K. Doshi and Y. Yilmaz, "Rethinking video anomaly detection a continual learning approach," in *Proceedings of WACV*, pp. 3961– 3970, 2022.
- [88] M. Haselmann, D. P. Gruber, and P. Tabatabai, "Anomaly detection using deep learning based image completion," *Proceedings of ICMLA*, pp. 1237–1242, 2018.
- [89] G. Yu, S. Wang, Z. Cai, E. Zhu, C. Xu, J. Yin, and M. Kloft, "Cloze Test Helps: Effective Video Anomaly Detection via Learning to Complete Video Events," in *Proceedings of ACMMM*, pp. 583–591, 2020.
- [90] M. Sabokrou, M. PourReza, M. Fayyaz, R. Entezari, M. Fathy, J. Gall, and E. Adeli, "AVID: Adversarial Visual Irregularity Detection," in *Proceedings of ACCV*, pp. 488–505, 2018.
- Detection," in *Proceedings of ACCV*, pp. 488–505, 2018.
 [91] D. Luo, C. Liu, Y. Zhou, D. Yang, C. Ma, Q. Ye, and W. Wang, "Video Cloze Procedure for Self-Supervised Spatio-Temporal Learning," in *Proceedings of AAAI*, pp. 11701–11708, 2020.
 [92] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and
- [92] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin Transformer: Hierarchical Vision Transformer using Shifted Windows," in *Proceedings ICCV*, pp. 10012–10022, 2021.
- [93] X. Guo, Z. Jin, C. Chen, H. Nie, J. Huang, D. Cai, X. He, and X. Hua, "Discriminative-Generative Dual Memory Video Anomaly Detection," arXiv preprint arXiv:2104.14430, 2021.
- [94] C. Li, K. Sohn, J. Yoon, and T. Pfister, "CutPaste: Self-Supervised Learning for Anomaly Detection and Localization," in *Proceedings* of CVPR, pp. 9664–9674, 2021.
- [95] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [96] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," in *Proceedings of NIPS*, pp. 1106–1114, 2012.
- [97] M. D. Zeiler and R. Fergus, "Visualizing and Understanding Convolutional Networks," in *Proceedings of ECCV*, pp. 818–833, 2014.
- [98] S. Sabour, N. Frosst, and G. E. Hinton, "Dynamic Routing Between Capsules," in *Proceedings of NIPS*, pp. 3859–3869, 2017.
- [99] A. Şandru, M.-I. Georgescu, and R. T. Ionescu, "Feature-level augmentation to improve robustness of deep neural networks to affine transformations," in *Proceedings of ECCVW*, 2022.
- [100] V. Nair and G. E. Hinton, "Rectified Linear Units Improve Restricted Boltzmann Machines," in *Proceedings of ICML*, pp. 807– 814, 2010.
- [101] B. Ramachandra, M. J. Jones, and R. R. Vatsavai, "A Survey of Single-Scene Video Anomaly Detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [102] W. Sultani, C. Chen, and M. Shah, "Real-World Anomaly Detection in Surveillance Videos," in *Proceedings of CVPR*, pp. 6479– 6488, 2018.
- [103] P. Wu, J. Liu, and F. Shen, "A Deep One-Class Neural Network for Anomalous Event Detection in Complex Scenes," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 31, no. 7, pp. 2609–2622, 2019.
- [104] S. Lee, H. G. Kim, and Y. M. Ro, "BMAN: Bidirectional Multi-Scale Aggregation Networks for Abnormal Event Detection," *IEEE Transactions on Image Processing*, vol. 29, pp. 2395–2408, 2019.
- [105] F. Dong, Y. Zhang, and X. Nie, "Dual Discriminator Generative Adversarial Network for Video Anomaly Detection," *IEEE Access*, vol. 8, pp. 88170–88176, 2020.
- [106] K. Doshi and Y. Yilmaz, "Any-Shot Sequential Anomaly Detection in Surveillance Videos," in *Proceedings of CVPRW*, pp. 934– 935, 2020.

- [107] C. Sun, Y. Jia, Y. Hu, and Y. Wu, "Scene-Aware Context Reasoning for Unsupervised Abnormal Event Detection in Videos," in *Proceedings of ACMMM*, pp. 184–192, 2020.
- [108] Z. Wang, Y. Zou, and Z. Zhang, "Cluster Attention Contrast for Video Anomaly Detection," in *Proceedings of ACMMM*, pp. 2463– 2471, 2020.
- [109] M. Astrid, M. Z. Zaheer, J.-Y. Lee, and S.-I. Lee, "Learning not to reconstruct anomalies," in *Proceedings of BMVC*, 2021.
- [110] I. C. Duță, M. I. Georgescu, and R. T. Ionescu, "Contextual convolutional neural networks," in *Proceedings of ICCVW*, pp. 403–412, 2021.



Neelu Madan is a Ph.D. student at the Department of Media Technology at Aalborg University, Denmark. She received her M.Sc. degree in Computer Science with a major as Interactive Systems and Visualization from the University of Duisburg-Essen, Germany. Her research interests include artificial intelligence, computer vision, machine learning, and deep learning. She is the author of a paper accepted for oral presentation at CVPR 2022.



Nicolae-Cătălin Ristea graduated as valedictorian from the Faculty of Electronics, Telecommunications and Information Technology, University Politehnica of Bucharest, in 2019. He received the M.Sc. degree in the image processing field and started the Ph.D. at the same university. Nicolae is first author of multiple papers accepted at top-tier conferences and journals, such as CVPR and INTERSPEECH. His research interests include AI, computer vision, machine learning, signal processing and deep learning.



Radu lonescu is professor at the University of Bucharest, Romania. He completed his Ph.D. at the University of Bucharest in 2013, receiving the 2014 Award for Outstanding Doctoral Research from the Romanian Ad Astra Association. His research interests include machine learning, computer vision, image processing, computational linguistics and medical imaging. He published over 100 articles at international venues (including CVPR, NeurIPS, ICCV, ACL, SIGIR, EMNLP, NAACL, TPAMI, IJCV, CVIU), and a

research monograph with Springer. Radu received the "Caianiello Best Young Paper Award" at ICIAP 2013.



Kamal Nasrollahi is working in a dual position, Professor of Computer Vision and Machine Learning at Aalborg University and Head of Machine Learning at Milestone Systems. He is interested in fair, ethical, and responsible use of technology, specifically machine learning applied to computer vision for topics like object detection, tracking, anomaly detection, and superresolution.



Fahad Khan is a faculty member at MBZ University of AI (MBZUAI), UAE and Linköping University, Sweden. Prior to joining MBZUAI, he worked as a Lead Scientist at the Inception Institute of Artificial Intelligence (IIAI), UAE. He received the M.Sc. degree in Intelligent Systems Design from Chalmers University of Technology, Sweden and a Ph.D. degree in Computer Vision from Autonomous University of Barcelona, Spain. His research interests include a wide range of topics within computer vision, such as object recogni-

tion, object detection, action recognition and visual tracking. He has published articles in high-impact computer vision journals and conferences in these areas.



Thomas B. Moeslund leads the Visual Analysis and Perception lab at Aalborg University, the Media Technology section at Aalborg University and the AI for the People Center at Aalborg University. His research covers all aspects of software systems for automatic analysis of visual data, especially including people.



Mubarak Shah, the UCF Trustee chair professor, is the founding director of the Center for Research in Computer Vision at the University of Central Florida (UCF). He is a fellow of the NAI, IEEE, AAAS, IAPR and SPIE. His research interests include video surveillance, visual tracking, human activity recognition, visual analysis of crowded scenes, video registration, UAV video analysis, among others. He has served as an ACM distinguished speaker and IEEE distinguished visitor speaker. He is a recipient of ACM

SIGMM Technical Achievement award; IEEE Outstanding Engineering Educator Award; Harris Corporation Engineering Achievement Award; an honorable mention for the ICCV 2005 "Where Am I?" Challenge Problem; 2013 NGA Best Research Poster Presentation; 2nd place in Grand Challenge at the ACM Multimedia 2013 conference; and runner up for the best paper award in ACM Multimedia Conference in 2005 and 2010. At UCF he has received the Pegasus Professor Award; University Distinguished Research Award; Faculty Excellence in Mentoring Doctoral Students; Scholarship of Teaching and Learning award; Teaching Incentive Program award; Research Incentive Award.