



HAL
open science

MixedTeacher: Knowledge Distillation for fast inference textural anomaly detection

Simon Thomine, Hichem Snoussi, Mahmoud Soua

► **To cite this version:**

Simon Thomine, Hichem Snoussi, Mahmoud Soua. MixedTeacher: Knowledge Distillation for fast inference textural anomaly detection. 2023 International Conference on Computer Vision Theory and Applications (VISAPP 2023), Feb 2023, Lisbonne, Portugal. hal-04131472

HAL Id: hal-04131472

<https://hal.science/hal-04131472>

Submitted on 16 Jun 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

MixedTeacher : Knowledge Distillation for fast inference textural anomaly detection

Simon Thomine^{1,2}, Hichem Snoussi¹, Mahmoud Soua²

¹University of technology Troyes, Troyes, France

² AQUILAE, Troyes, France

e-mail: simon.thomine@utt.fr, hichem.snoussi@utt.fr, m.soua@aquilae.tech

Keywords: Anomaly Detection - Texture - Knowledge Distillation - Layer Selection - Unsupervised

Abstract: For a very long time, unsupervised learning for anomaly detection has been at the heart of image processing research and a stepping stone for high performance industrial automation process. With the emergence of CNN, several methods have been proposed such as Autoencoders, GAN, deep feature extraction, etc. In this paper, we propose a new method based on the promising concept of knowledge distillation which consists of training a network (the student) on normal samples while considering the output of a larger pretrained network (the teacher). The main contributions of this paper are twofold: First, a reduced student architecture with optimal layer selection is proposed, then a new Student-Teacher architecture with network bias reduction combining two teachers is proposed in order to jointly enhance the performance of anomaly detection and its localization accuracy. The proposed texture anomaly detector has an outstanding capability to detect defects in any texture and a fast inference time compared to the SOTA methods.

1 INTRODUCTION

Anomaly detection in industry is a vast topic since there is a lot of possible applications. For instance, defect detection aims at identifying specific anomaly classes and locations in industrial manufacturing processes (Kähler, Schmedemann, and Schüppstuhl, 2022). This detection is crucial for ensuring the high quality of final products (Minhas and Zelek, 2019). A common property of defects is that their visual texture is inherently different from the defect-free surface. The specificity of textures is the pattern structure which, if known, allows the detection and the extraction of anomalies. However, the texture anomaly generally appears in a small region in few samples, which makes it difficult to build consistent normal and abnormal datasets to be used in supervised learning methods. Hence, unsupervised anomaly detection networks are very suitable for industrial scenarios as they represent the strong basis for building a detection model without any annotated samples (Huang et al., 2022). Several unsupervised anomaly detection methods have been introduced for texture anomaly detection. These methods could achieve high performance up to 99.6 AUROC. However, they suffer from complex networks and high latency.

In another context, knowledge distillation has been introduced with the purpose of reducing the network size while increasing performance. Knowledge distillation aims to train a smaller network (student) to imitate pretrained one or several larger ones (teachers) on normal samples. As the teacher is pretrained, it has the ability to generalize even if the sample contains an anomaly, whereas the student won't be able. Hence, by comparing the extracted features between the teacher and the student networks, an abnormal sample could be detected. According to some studies (Iglesias and Zseby, 2015), using too many features can significantly reduce the accuracy of anomaly detection. Recently, a Student-Teacher Feature Pyramid Matching (STPM) method has been proposed in (Wang et al., 2021), where the first three network layers are used in order to focus on edges, colors and shapes instead of context information. Even if using layer selection technique is an interesting approach, there is still a lack of explanation concerning the layer choice and the relevance of the relative information. Looking at the same layers for an object and for a texture reduces the relevance of the extracted information. For example, looking at context information in a texture is pointless and for an object, pure edge/color/texture information may not be the most interesting information.

Another recurrent problem is the classifier bias. The

best current methods use a pretrained classifier network on imageNet which is biased by the classes of imageNet and can have an impact on the localization and the detection of defects.

The main contributions of the paper are as follows:

- A new reduced student architecture for texture-specific object category.
- In order to reduce the classification bias, we propose a new architecture combining two teachers pretrained on imageNet but with different architectures (respectively ResNet-18 (He et al., 2015) and EfficientNet-b0 (Tan and Le, 2020)) and a single student network. This new mixed Teacher network structure outperforms competitive state-of-the-art methods both in inference time and SOTA scores, on anomaly datasets such as MVTEC AD textures and BTAD textures (Mishra et al., 2021). The proposed MixedTeacher model uses a score and anomaly localisation function based on each complementary teacher features with a careful feature selection.

The paper is organized as follows. In section 2, we review the related work, especially on MVTEC dataset and present the different approaches proposed in literature. In section 3, we compare the results of training with different architectures and different layer selection schemes and introduce our proposed texture-specific reduced student architecture. Section 4 is dedicated to describing a novel mixed Student-Teacher network. In section 5, we compare our results to the SOTA methods for both the reduced student architecture and the MixedTeacher in terms of AUROC, pixel-AUROC and inference time.

2 RELATED WORK

Anomaly detection is a problem that pops up in many areas and is often very difficult to deal with. Indeed, detecting the “abnormal” is a rather vague concept and is difficult to define according to the use cases, which makes research on this subject very specific.

For several years, the rise of deep learning has never ceased to impress with high quality results and interesting methods. Most of these methods are based on an unsupervised representation approach to discriminate outliers. Some specific work has been done for fabrics defect detection such as the multi-scale Convolutional denoising autoencoder (Mei, Wang, and Wen, 2018). For unsupervised anomaly detection in general, we can also cite the GEE, a gradient based VAE (Nguyen et al., 2019) or the Gaussian

mixture model VAE (Nguyen et al., 2019). Another common way to detect anomaly is to use generative adversarial networks (Goodfellow et al., 2014). AnoGAN (Schlegl et al., 2019) was one of the first utilization of GAN for anomaly detection but since then a lot of approaches emerged such as G2D (Pourreza et al., 2021) and OCR-GAN (Liang et al., 2022). Other interesting approaches rely on pretrained models especially on imageNet, using the feature extraction of pretrained network to extract useful information about a given sample. The idea is to extract features with a pretrained model and then train a normalizing flow model on good samples, so that the model is ready to find out if a given sample is an anomaly by looking at the reconstruction error. An advantage of normalizing flow is the reversible aspect which is useful to locate the anomaly pixel-wise. Many techniques based on this concept have been proposed such as differNet (Rudolph, Wandt, and Rosenhahn, 2021) and CSFLOW (Rudolph et al., 2021) which consider multi-scale normalizing flow and FastFlow (Yu et al., 2021) based on a 2D normalizing flow.

Recently, the concept of knowledge distillation has also been used for unsupervised anomaly detection. The student-teacher method consists of training a student teacher based on the output of a larger teacher model which is pretrained on ImageNet. The student network will learn to imitate the teacher on good samples only. Then, when an abnormal sample is tested, the teacher will be able to generalize and the student won't be, the difference between the output of the teacher and the output of the student will allow the detection of the anomaly. On the MVTEC dataset, four methods have been implemented, STPM (Wang et al., 2021) which trained the student on the 3 first layers of ResNet-18, RSTPM (Yamada and Hotta, 2022) which is basically the same method but with an attention layer, reverse distillation (Deng and Li, 2022) and CFA (Lee, Lee, and Song, 2022).

3 LAYER SELECTION AND REDUCED STUDENT

In this section, after a comparative study of layer selection methods for optimal texture anomaly detection, we present a new student architecture that both increases performance and reduces the inference time.

3.1 Layer selection

In deep neural networks, a common observation is that deep layer features contain context information and shallow layer features contain color, texture and

contour information. In a case of detection of defects on the fabric or on a generic texture, the context information is less important than the texture information, therefore, we will turn to shallow layer features. As reported in table 1, different combinations of shallow layers have been tried in order to select the optimal architecture with respect to detection performance evaluated by the AUC.

Table 1: Layers selection results

| Measures | Layer 1 and 2 AUC | Layer 2 and 3 AUC |
|---------------|-------------------|-------------------|
| Mean objects | 0.876 | 0.910 |
| Mean textures | 0.990 | 0.971 |

3.2 Reduced student

ResNet-18 architecture has been retained for the teacher network. As texture specific anomaly detection is the main objective of this work, we propose to add the ResNet-18 first layer after the first convolution to extract even more textural information. The second objective was to alleviate the student architecture to decrease inference time and possibly performance. As ResNet-18 presents several residual blocks with two identical convolutional layers, we first decided to take only one layer for each block in our student architecture. The classifier bias is another known problem while dealing with pretrained classifier and we tackled this problem by reducing features size with an adaptive average pooling layer in each Resnet residual block's output as presented in figure 1.

Given a training dataset of images without anomaly $D = [I_1, I_2, \dots, I_n]$, our goal is to extract the information of the L bottom layers. For an image $I_k \in R^{w \times h \times c}$ where w is the width, h the height and c the number of channels, the teacher outputs features $F_t^l(I_k) \in R^{w_l \times h_l \times c_l}$ and $F_s^l(I_k) \in R^{w_l/2 \times h_l/2 \times c_l/2}$ with $l > 1$ and $F_s^l(I_k) \in R^{w_l \times h_l \times c_l}$ if $l = 1$. For the loss function, we took the l_2 distance of normalized feature vectors like in the STPM original paper (Wang et al., 2021) while using an adaptive average pooling on teacher features where $l > 1$ and just sum all feature maps of all layers to obtain our loss with the same ratio for all layers (Eq.1).

$$F_t^{l>1}(I_k) = AAP(F_{Resnet18}^{l>1}(I_k)) \quad (1)$$

where AAP refers to the Adaptive Average Pooling. Pixel loss is defined in the following Eq.2:

$$loss^l(I_k)_{ij} = \frac{1}{2} \|norm(F_t^l(I_k)_{ij}) - norm(F_s^l(I_k)_{ij})\| \quad (2)$$

and for the layer 1, the loss is defined as:

$$loss^1(I_k) = \frac{1}{w_1 h_1} \sum_{i=1}^{w_1} \sum_{j=1}^{h_1} loss^1_{resNet}(I_k)_{ij} \quad (3)$$

and finally for the total loss is written as:

$$loss(I_k) = \sum^l loss^l(I_k) \quad (4)$$

Performance and inference speed are later reported in section 5 with comparison with SOTA networks on anomaly detection.

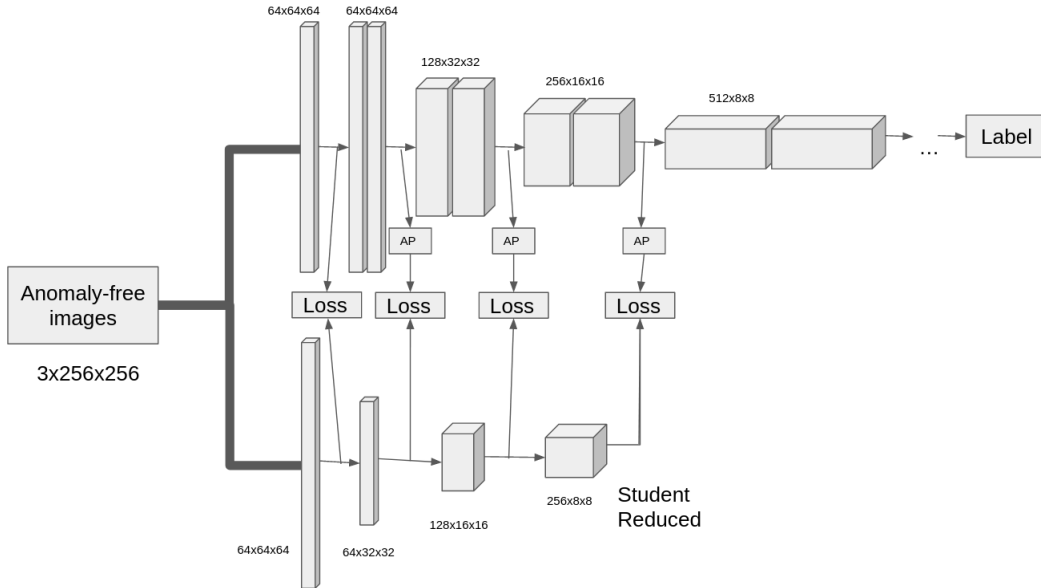


Figure 1: Reduced student architecture with AP for adaptive average pooling.

4 MIXED TEACHER

In this section, we introduce our new student teacher network structure that combines two teachers with the purpose of reducing the classifier bias, taking benefits from the two networks and exploiting the different layers in an optimal way.

4.1 Observation and main ideas

While testing our new student reduced architecture on the MVTEC AD textures, we obtained good results, but some noise still degrade results in terms of default localisation on specific images or texture-specific normal variation. Different teacher network architectures have been tested to conclude that ResNet-18 remains the best in terms of average precision and speed. However, interesting behaviors have been observed on the noise localisation for each architecture. In fact, every classifier had the capacity to locate the anomaly, but with output noise and anomaly detection mistakes.

The combination of two pretrained classifier networks has therefore been proposed with the purpose of interpolating their defect localisation to cancel noise and false detection/segmentation.

EfficientNet-b0 has been proposed as the second teacher when considering its performance in terms of precision and speed. For this network, it has been observed that for the bottom layers, one has good localisation but with a huge noise and with top layers, a coarse defect localisation but with minimal noise has been obtained, as illustrated in figure 2.

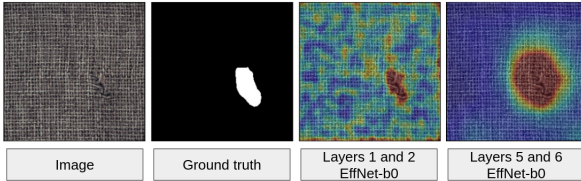


Figure 2: Difference between top layers and bottom layers for EfficientNet-b0 architecture.

4.2 Method description

The learning architecture is composed of two teachers: the ResNet-18 as main teacher and EfficientNet-b0 as a localisation confirmation teacher. For the ResNet-18 part, the reduced student proposed in section 3 is used in order to ensure a good inference speed and precision on texture samples. For EfficientNet-b0 student, we used one convolution for each efficientnet block without pooling because we

used deepest layers. In the student architecture, there is no communication between the networks except for the loss function as shown in figure 3.

For the training loss function, we used basically the same loss function as the one for the reduced teacher and we add an α factor to smooth the layer activation difference from the two teacher networks. As feature difference in efficientNet was about 10 times bigger than in ResNet-18, α has been set to 0.1.

$$loss_{effNet}^{l=5,6}(I_k)_{ij} = \frac{1}{2} \|norm(F_t^l(I_k)_{ij}) - norm(F_s^l(I_k)_{ij})\| \quad (5)$$

and

$$loss_{effNet}^{l=5,6}(I_k) = \frac{1}{w_l h_l} \sum_{i=1}^{w_l} \sum_{j=1}^{h_l} loss_{effNet}^l(I_k)_{ij} \quad (6)$$

and for Resnet-18 part :

$$loss_{resNet}^{l=1,2,3}(I_k) = \frac{1}{w_l h_l} \sum_{i=1}^{w_l} \sum_{j=1}^{h_l} loss_{resNet}^l(I_k)_{ij} \quad (7)$$

with $loss_{resNet}^l(I_k)_{ij}$ defined as in section 3. For the total loss with the α factor :

$$loss_{tot}(I_k) = \sum_{l=1}^3 loss_{resNet}^l(I_k) + \alpha \sum_{l=5}^6 loss_{effNet}^l(I_k) \quad (8)$$

As in every knowledge distillation method, the loss only impacts the student.

4.3 Anomaly score and localisation

In the test phase (inference), we want an anomaly map M of the original image size where every pixel at position (i, j) has an anomaly score M_{ij} . With a test image I and $F_{tResnet}^l, F_{tEffNet}^l$ the two teachers features of l th layer and $F_{sResnet}^l, F_{sEffNet}^l$ their corresponding l th layer student features, we perform an upsampling on the difference between the corresponding layers. The coarse localisation output of the efficientNet layers is obtained by summing each layer's anomaly map.

The anomaly map is obtained in the same way for the resnet part. Respectively :

$$A_{mapEffNet} = \sum_{l=5}^6 Upsample(F_{tEffNet}^l - F_{sEffNet}^l) \quad (9)$$

and :

$$A_{mapResnet} = \sum_{l=1}^3 Upsample(F_{tResnet}^l - F_{sResnet}^l) \quad (10)$$

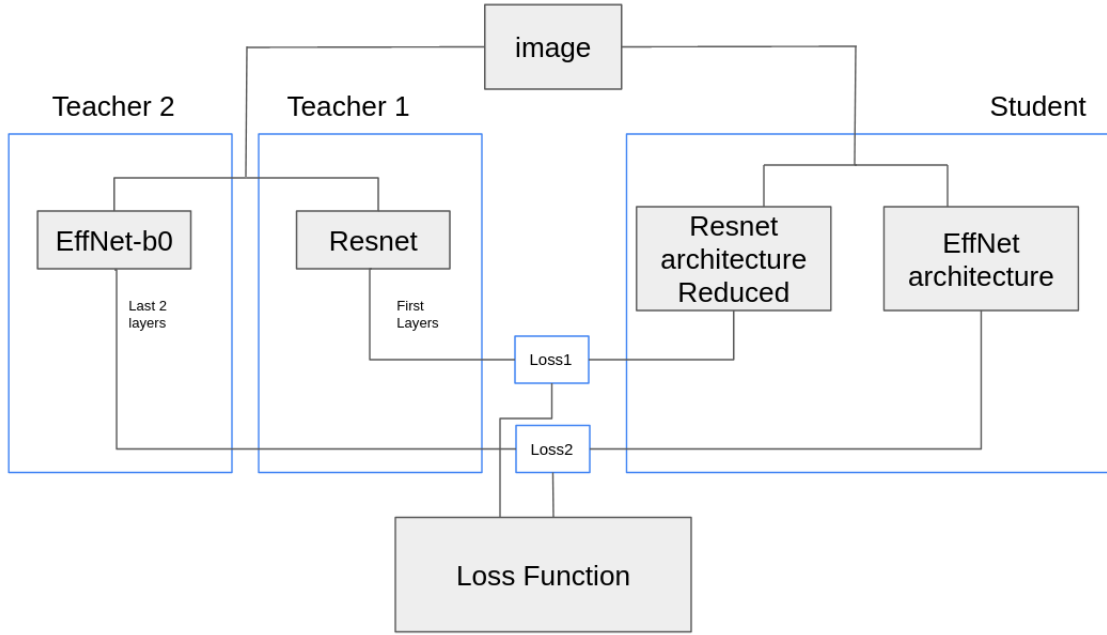


Figure 3: MixedTeacher architecture.

We then multiply the resnet anomaly map by the normalization of the effnet anomaly map multiplied by its mathematical extent. With $A_{mapEffnet}$, the anomaly map of efficientNet layers and $A_{mapResnet}$ the anomaly map of resnet layers, the final anomaly map is then defined as :

$$M = A_{mapResnet} * (\max(A_{mapEffnet}) - \min(A_{mapEffnet})) A_{mapEffnet} \quad (11)$$

The anomaly score is defined as :

$$score = \sum_{i=1}^w \sum_{j=1}^h M_{i,j} \quad (12)$$

with w and h are respectively the width and height of the anomaly map.

5 EXPERIMENTS

5.1 Datasets

We experiment our methods on the textures from the **MVTEC AD** (Bergmann et al., 2019) dataset which consists of 15 categories : 5 textures and 10 objects with a total of more than 5000 high resolution images. This dataset is used for unsupervised anomaly

detection therefore it contains only anomaly free images for the training. For the test part, it shows a good variety of defects with ground truth masks for anomaly localisation. We also used the texture of the **BTAD** (Mishra et al., 2021) dataset which is an unsupervised anomaly dataset with three different categories including one texture figure 4.




Figure 4: Overview of textures from MVTEC AD and BTAD dataset, samples with anomaly and ground truth. These images are only used for testing and unseen during the training.

The performance is evaluated with AUROC met-

ric image-level and pixel-level to compare our results with other methods.

5.2 Implementation and training metrics

Training and inference were done on an rtx 2080ti. To test the student reduced, we used the features of the first three blocks and the layer before the first block of ResNet-18. The Resnet network was pretrained on imageNet. We used stochastic gradient descent with a learning rate of 0.4 for 100 epochs with a batch size of 16. To test the MixedTeacher, we used the output features of the first two blocks and the layer before the first block of ResNet-18 and the output features of block 5 and 6 of EfficientNet-b0. We used stochastic gradient descent with a learning rate of 0.4 for 200 epochs with a batch size of 16. Both networks are pretrained on imageNet. We resized all the images to a size of 256x256 keeping 80% for training and 20% for validation. We kept the checkpoint with the lowest validation loss.

5.3 Reduced Student

5.3.1 Performance results

In this paragraph, we will compare reduced student AUROC results to SOTA methods. In 2, we present AUROC performance results of CFA (Lee, Lee, and Song, 2022), PatchCore (Roth et al., 2021), FastFlow (Yu et al., 2021), STPM (Wang et al., 2021), CutPaste (Li et al., 2021) and our reduced student on MVTEC AD textures.

Table 2: **Image-AUROC comparison on MVTEC AD : Reduced Student**

| Category | CutPaste | CFA | PatchCore | STPM | FastFlow | Ours |
|----------|----------|------|-----------|------|----------|------|
| carpet | 100 | 97.3 | 98.7 | 95.4 | 99.4 | 100 |
| tile | 100 | 99.4 | 98.7 | 94.9 | 100 | 98.7 |
| grid | 96.2 | 99.2 | 98.2 | 98.2 | 100 | 99.7 |
| wood | 99.1 | 99.7 | 99.2 | 96.1 | 99.2 | 99.6 |
| leather | 95.4 | 100 | 100 | 98.9 | 99.9 | 99.7 |
| Mean | 98.1 | 99.1 | 99.0 | 96.7 | 99.7 | 99.5 |

For FastFlow, we choosed to take the results from Anomalib as we were not able to reproduce their paper results (99.9 AUROC in paper). As seen in table 2, our reduced student is better than CFA for texture anomaly detection, which is the best actual knowledge distillation unsupervised anomaly detection method and is close to the SOTA results. We manage to gain 2.8 points against classic STPM with a network reduction and a wise layer selection aiming

for texture specific anomaly detection.

5.3.2 Inference time results

In table 3, we compare the reduced student inference time to other SOTA methods. The main purpose of reduced student was to propose a high processing speed to manage real time for several high resolution images. To get inference time results, we employ Anomalib. All the additional results come from this library to make sure the tests were carried out under the same conditions.

Table 3: **Inference time results**

| Category | PatchCore | FastFlow | STPM | Ours |
|--------------|-----------|----------|------|--------------|
| FPS | 5.8 | 21.8 | 83.2 | 108.1 |
| Latency (ms) | 172 | 45.9 | 12 | 9.2 |

The presented results are based on Anomalib inference time. In a self made code, we were able to obtain a 10x better inference time for STPM and reduced student. The most important thing to consider is that the STPM is by far the fastest anomaly detector and reduced student reduced its inference time by 30%.

5.4 MixedTeacher

5.4.1 Performance results

Unlike the reduced student, the MixedTeacher main purpose is performance and not inference time. In table 4 we compared AUROC of several SOTA methods in texture anomaly detection.

Table 4: **image-AUROC comparison on MVTEC AD : MixedTeacher**

| Category | CutPaste | CFA | PatchCore | FastFlow | ReducedStudent | Ours |
|----------|----------|-------------|-----------|------------|----------------|-------------|
| carpet | 100 | 97.3 | 98.7 | 99.4 | 100 | 99.8 |
| tile | 100 | 99.4 | 98.7 | 100 | 98.7 | 100 |
| grid | 96.2 | 99.2 | 98.2 | 100 | 99.7 | 99.7 |
| wood | 99.1 | 99.7 | 99.2 | 99.2 | 99.6 | 99.6 |
| leather | 95.4 | 100 | 100 | 99.9 | 99.7 | 100 |
| Mean | 98.1 | 99.1 | 99.0 | 99.7 | 99.5 | 99.8 |

Our method is the new state of the art texture anomaly detection on the MVTEC AD dataset.

5.4.2 Anomaly localisation

Even though anomaly localisation was not our main purpose, our approach uses EfficientNet-b0 with the objective of making the location more precise. To this end, we present in table 5 and table 6, our anomaly location results on textures from MVTEC AD dataset and BTAD respectively and we compare these results to the SOTA methods.

Table 5: Pixel-AUROC comparison on MVTEC AD : MixedTeacher

| Category | CutPaste | PatchCore | FastFlow | Ours |
|----------|----------|-----------|----------|------|
| carpet | 98.3 | 98.9 | 99.1 | 99.0 |
| tile | 90.5 | 95.6 | 96.6 | 95.9 |
| grid | 97.5 | 98.7 | 99.2 | 97.5 |
| wood | 95.5 | 95 | 94.1 | 94.9 |
| leather | 99.5 | 99.3 | 99.6 | 99.4 |
| Mean | 96.2 | 97.5 | 97.7 | 97.3 |

Table 6: Image-AUROC comparison on BTAD: MixedTeacher

| Category | FastFlow | Ours |
|--------------------|----------|-------------|
| 1 (wood from btad) | 96.0 | 97.0 |

5.4.3 Inference time results

In terms of inference speed, our MixedTeacher is 3x slower than the reduced student since it used two teacher networks and a more complex student architecture.

6 CONCLUSION

In this paper, we proposed two methods for efficient unsupervised anomaly detection using the principle of knowledge distillation applied to unsupervised anomaly training. Both methods offer different benefits. The reduced student proposes a high speed texture anomaly detector with an AUROC performance close to the state of the art, this method is to be used in situations where inference time is the most important priority (mobile device, low computational power, cost efficiency). The MixedTeacher propose the highest actual performance on anomaly detection with a localisation close to the state of the art on the MVTEC AD textures with still a fast inference. This method is to be used in situations where performance is the priority and the computational power is big enough (monitoring server etc ...)

References

Kähler, F., O. Schmedemann, and T. Schüppstuhl (2022). “Anomaly detection for industrial surface inspection: application in maintenance of aircraft components”. In: *Procedia CIRP* 107, pp. 246–251.

Minhas, M. S. and J. Zelek (Nov. 24, 2019). *AnoNet: Weakly Supervised Anomaly Detection in Textured Surfaces*.

Huang, J., C. Li, Y. Lin, S. Lian, and A. Innovation (2022). “Unsupervised Industrial Anomaly Detection via Pattern Generative and Contrastive Networks”. In: *arXiv:2207.09792 [cs]*.

Iglesias, F. and T. Zseby (Oct. 2015). “Analysis of network traffic features for anomaly detection”. In: *Machine Learning* 101.1, pp. 59–84.

Wang, G., S. Han, E. Ding, and D. Huang (Oct. 28, 2021). “Student-Teacher Feature Pyramid Matching for Anomaly Detection”. In: *arXiv:2103.04257 [cs]*.

He, K., X. Zhang, S. Ren, and J. Sun (Dec. 10, 2015). *Deep Residual Learning for Image Recognition*.

Tan, M. and Q. V. Le (Sept. 11, 2020). *EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks*.

Mishra, P., R. Verk, D. Fornasier, C. Piciarelli, and G. L. Foresti (June 20, 2021). “VT-ADL: A Vision Transformer Network for Image Anomaly Detection and Localization”. In: *2021 IEEE 30th International Symposium on Industrial Electronics (ISIE)*, pp. 01–06.

Mei, S., Y. Wang, and G. Wen (Apr. 2018). “Automatic Fabric Defect Detection with a Multi-Scale Convolutional Denoising Autoencoder Network Model”. In: *Sensors* 18.4, p. 1064.

Nguyen, Q. P., K. W. Lim, D. M. Divakaran, K. H. Low, and M. C. Chan (Mar. 15, 2019). “GEE: A Gradient-based Explainable Variational Autoencoder for Network Anomaly Detection”. In: *arXiv:1903.06661 [cs, stat]*.

Goodfellow, I. J., J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio (June 10, 2014). “Generative Adversarial Networks”. In: *arXiv:1406.2661 [cs, stat]*.

Schlegl, T., P. Seeböck, S. M. Waldstein, G. Langs, and U. Schmidt-Erfurth (May 2019). “f-AnoGAN: Fast unsupervised anomaly detection with generative adversarial networks”. In: *Medical Image Analysis* 54, pp. 30–44.

Pourreza, M., B. Mohammadi, M. Khaki, S. Bouindour, H. Snoussi, and M. Sabokrou (Jan. 2021). “G2D: Generate to Detect Anomaly”. In: *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*. 2021 IEEE Winter Conference on Applications of Computer Vision (WACV). event-place: Waikoloa, HI, USA. IEEE, pp. 2002–2011.

Liang, Y., J. Zhang, S. Zhao, R. Wu, Y. Liu, and S. Pan (Mar. 1, 2022). “Omni-frequency Channel-selection Representations for Unsupervised Anomaly Detection”. In: *arXiv:2203.00259 [cs]*.

- Rudolph, M., B. Wandt, and B. Rosenhahn (Jan. 2021). “Same Same But DifferNet: Semi-Supervised Defect Detection with Normalizing Flows”. In: *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*. 2021 IEEE Winter Conference on Applications of Computer Vision (WACV). event-place: Waikoloa, HI, USA. IEEE, pp. 1906–1915.
- Rudolph, M., T. Wehrbein, B. Rosenhahn, and B. Wandt (Oct. 6, 2021). “Fully Convolutional Cross-Scale-Flows for Image-based Defect Detection”. In: *arXiv:2110.02855 [cs]*.
- Yu, J., Y. Zheng, X. Wang, W. Li, Y. Wu, R. Zhao, and L. Wu (Nov. 16, 2021). “FastFlow: Unsupervised Anomaly Detection and Localization via 2D Normalizing Flows”. In: *arXiv:2111.07677 [cs]*.
- Yamada, S. and K. Hotta (2022). “Reconstruction Student with Attention for Student-Teacher Pyramid Matching”. In: *arXiv:2111.15376 [cs]*.
- Deng, H. and X. Li (2022). “Anomaly Detection via Reverse Distillation from One-Class Embedding”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2022*, pp. 9737–9746.
- Lee, S., S. Lee, and B. C. Song (June 9, 2022). *CFA: Coupled-hypersphere-based Feature Adaptation for Target-Oriented Anomaly Localization*.
- Bergmann, P., M. Fauser, D. Sattlegger, and C. Steger (June 2019). “MVTec AD — A Comprehensive Real-World Dataset for Unsupervised Anomaly Detection”. In: *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). Long Beach, CA, USA: IEEE, pp. 9584–9592.
- Roth, K., L. Pemula, J. Zepeda, B. Schölkopf, T. Brox, and P. Gehler (June 15, 2021). “Towards Total Recall in Industrial Anomaly Detection”. In: *arXiv:2106.08265 [cs]*.
- Li, C.-L., K. Sohn, J. Yoon, and T. Pfister (Apr. 8, 2021). “CutPaste: Self-Supervised Learning for Anomaly Detection and Localization”. In: *arXiv:2104.04015 [cs]*.