

Revisiting Multimodal Fusion for 3D Anomaly Detection from an Architectural Perspective

Kaifang Long^{1*}, Guoyang Xie^{2*}, Lianbo Ma^{1†}, Jiaqi Liu³, Zhichao Lu³

¹Software College, Northeastern University, Shenyang, China

²The Department of Intelligent Manufacturing, CATL, Ningde, China

³The Department of Computer Science, City University of Hong Kong, Hong Kong, China

longkf@stumail.neu.edu.cn, guoyang.xie@ieec.org, malb@swc.neu.edu.cn, liu_jiaqi_@outlook.com, luzhichao.cn@gmail.com

Abstract

Existing efforts to boost multimodal fusion of 3D anomaly detection (3D-AD) primarily concentrate on devising more effective multimodal fusion strategies. However, little attention was devoted to analyzing the role of multimodal fusion architecture (topology) design in contributing to 3D-AD. In this paper, we aim to bridge this gap and present a systematic study on the impact of multimodal fusion architecture design on 3D-AD. This work considers the multimodal fusion architecture design at the intra-module fusion level, i.e., independent modality-specific modules, involving early, middle or late multimodal features with specific fusion operations, and also at the inter-module fusion level, i.e., the strategies to fuse those modules. In both cases, we first derive insights through theoretically and experimentally exploring how architectural designs influence 3D-AD. Then, we extend SOTA neural architecture search (NAS) paradigm and propose 3D-ADNAS to simultaneously search across multimodal fusion strategies and modality-specific modules for the first time. Extensive experiments show that 3D-ADNAS obtains consistent improvements in 3D-AD across various model capacities in terms of accuracy, frame rate, and memory usage, and it exhibits great potential in dealing with few-shot 3D-AD tasks.

Introduction

Industrial anomaly detection is expected to accurately find out the difference between normal samples and anomalies like human inspectors. To achieve this, an emerging way is to exploit both image color (RGB) and depth information (rather than only RGB) for quality inspection, termed as 3D anomaly detection (3D-AD) (Qin et al. 2023; Zavrtnik, Kristan, and Skočaj 2024a). It works well since 3D depth information has shown an essential role in improving industrial detection accuracy (Cao, Xu, and Shen 2024). The typical way of existing 3D-AD methods is to train a reconstruction-based model via restoring synthetic abnormal samples to normal ones, and then utilize reconstruction features for defect identification of discriminator (Zhou et al. 2024). One key challenge of this way is how to realize correct fusion of the two modalities (i.e., RGB images and 3D point cloud) through multimodal

*Contribute equally.

†Corresponding author.

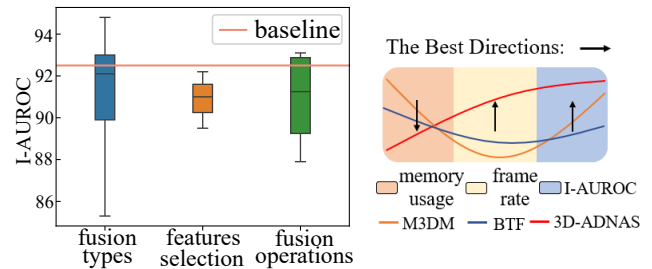


Figure 1: **(Left)** The impact of multimodal fusion architecture design on 3D-AD performance. This shows the distribution of 3D-AD performance with variations at the intra- and inter-module fusion levels. **(Right)** 3D-ADNAS vs. SOTA methods in terms of accuracy, FPS, and memory usage.

fusion network such that the integration of depth information will not interfere with color information (Gu et al. 2024a).

Regarding to the multimodal fusion network, it is natural that consistently evolved architectures outperform the original ones. In this sense, we have to figure out a question that whether the existing multimodal fusion architectures are ideal for 3D-AD. As previewed in Figure 1, we can observe that the architectural designs (in terms of modality-specific modules and combination of these modules) have a significant impact on 3D-AD’s accuracy. In fact, this observation can be validated via theoretical analyses (as provided in **Revisiting 3D-AD Fusion Architecture Section**). This motivates us to design multimodal fusion architectures/topologies tailored for 3D-AD.

The primary goal of this work is to bridge the gap between multimodal fusion architectures and 3D-AD via (i) *systematically studying the impact of architectural components on 3D-AD*, (ii) *recognizing crucial design schemes that can enhance 3D-AD*, and (iii) *proposing a simple yet effective 3D-ADNAS with a novel two-level search space tailored for 3D-AD*. We adopt empirical and theoretical approaches and conduct extensive experiments to realize this goal.

Given that EasyNet (Chen et al. 2023) serves as foundation framework for 3D-AD, we initiate this work from it to revisit multimodal fusion network and formulate the design of target multimodal fusion network around a two-level multimodal fusion problem, i.e., intra-module fusion and inter-

module fusion. Then we systematically assess the two main aspects of architecture design, intra-module fusion structure and inter-module fusion structure. Regarding to the former, we center around the investigation of modality-specific modules. Each module involves highly-correlated modality features, e.g., data-level or early features (Xu, So, and Dai 2021); middle-level or hybrid features (Bergmann and Sattlegger 2023); and classifier-level or late features (Wang et al. 2024a). Typical fusion operations include addition (weighted summation) and concatenation (Li et al. 2024; Zhang and Li 2023). Regarding to the latter, we seek to optimize fusion strategies, i.e., searching for optimal combination of those modules. To avoid randomness of the empirical observations, we repeat each ablation test multiple times with different random seeds. Then, we can obtain the following new observations:

❶ The single use of middle feature fusion is more favorable than early or late feature fusion, while the single use of late features degrades 3D-AD performance (**intra-module impact**). But, when combined with other fusion strategies, the late feature fusion can consistently improve performance across most 3D-AD tasks (**inter-module impact**). (Fig. 4-a)

❷ Selecting the first two layers of middle-level features for fusion is, in general, more beneficial with 3D-AD training, as opposed to selecting all middle features used in standard 3D-AD (**intra-module impact**). (Fig. 4-b)

❸ Employing multiple fusion operations in each fusion module (i.e., early, middle, or late feature fusion module) is more effective than only one operation. Particularly, the guided attention operation and weighted summation operation are good at improving 3D-AD (**intra-module impact**). (Fig. 4-c and Table 1)

❹ The multimodal fusion performance is highly dependent on the combination of early, middle and late features with relevant fusion operations. Especially, the automatic choice of the final fusion strategy according to the characteristics of features and operations is the core factor to realize 3D-AD (**inter-module impact**). (Fig. 4-a and Table 4)

❺ In summary, the multimodal fusion architectures in terms of both intra- and inter-module aspects contribute enormously to 3D-AD.

Based on these insights above, we can recognize key architectural designs that can boost the multimodal fusion of 3D-AD. Then, another question is arisen: **How to seek an efficient and 3D-AD-friendly multimodal fusion architecture?** A natural way is the Neural Architecture Search (NAS) technique (Liu, Simonyan, and Yang 2018), which can perform an automated search over the architecture for 3D-AD. In this paper, we extend NAS paradigm for 3D-AD and boost multimodal fusion from an architectural perspective. Note that, the performance of NAS is sensitive to the design of multimodal fusion search space (Pérez-Rúa et al. 2019; Yin, Huang, and Zhang 2022). Based on the investigation about the impact of architectural designs on 3D-AD’s performance, this work suggests a specialized two-level search space for 3D-AD, within which each architecture candidate can be concisely represented, facilitating an efficient search process.

The contributions of this work are as follows:

- We thoroughly analyze the impact of multimodal fusion architectures on 3D-AD from both empirical and theoretical

views, and recognize the crucial intra- and inter-module fusion designs that can boost 3D-AD performance.

- To the best of our knowledge, this work is the first attempt to utilize the NAS technique to shape a 3D-AD-friendly multimodal fusion architecture. To realize this, we design a 3D-ADNAS method to seek a promising fusion architecture across intra- and inter-module fusion strategies.
- Extensive experiments validate the effectiveness of our design in improving anomaly detection of multimodal fusion network in terms of accuracy, speed and efficiency. Furthermore, our proposed method shows great potential in dealing with few-shot 3D-AD tasks.

Preliminaries

In this section, we present the experimental setup for evaluating multimodal fusion architecture towards 3D-AD.

Architecture Skeleton. Fig. 2 shows the skeleton of 3D-ADNAS, which consists of an Anomaly Generator (AG), a Multimodal Reconstruction Network (MRN), a Multimodal Fusion Network (MFN), and an Anomaly Discrimination Network (ADN). AG aims to simulate abnormal image generation according to the mask. MRN restores abnormal images (from AG) to normal ones and then extracts multi-scale features, which can be divided as early features, middle features and late features (Xu, So, and Dai 2021). MFN consists of three types of independent modality-specific modules (MSMs), and it aims to fuse these multimodal features from MRN in a two-level manner. Then, ADN detects the products whether defective. Particularly, the two-level multimodal feature fusion of MFN is the target component of our design: (i) at the **intra-module fusion level**, we aim to optimize the inner structure of the three MSMs (i.e., early MSM, middle MSM and late MSM); (ii) at the **inter-module fusion level**, we optimize fusion strategies to combine those MSMs.

MSM Structure. Fig. 3 shows the basic structure of an MSM, which consists of a specific candidate feature pool \mathbb{F} (i.e., with early features, middle features or late features) and a fusion cell (i.e., EFAC, MIFC, or L AFC). Each cell is regarded as a directed acyclic graph containing two input nodes, \mathbb{K} intermediate nodes and an output node. It has three trainable architectural parameters that need to be optimized: α^{ex} denotes the weights of candidate features in the pool, α^{in} represents the weights of connections between input nodes and intermediate nodes, and β^{op} denotes the weights of candidate operations (e.g., addition, concatenation, GLU, and guided attention) in the cell.

Evaluation Metrics. We consider several popular metrics (Dai et al. 2024; Sui et al. 2024) to evaluate the 3D-AD performance, including I-AUROC (image-level area under the receiver operating characteristic curve), P-AUROC (pixel-wise area under the receiver operating characteristic curve), and AUPRO (area under the per-region overlap curve).

Implementation Details. For the tests about impact of intra/inter-module fusion architectures, we train the overall model of 3D-ADNAS by 600 epochs. For the tests about evaluation of searched architectures, we train MFN model by 80 epochs and follow the setting of DARTS to obtain best

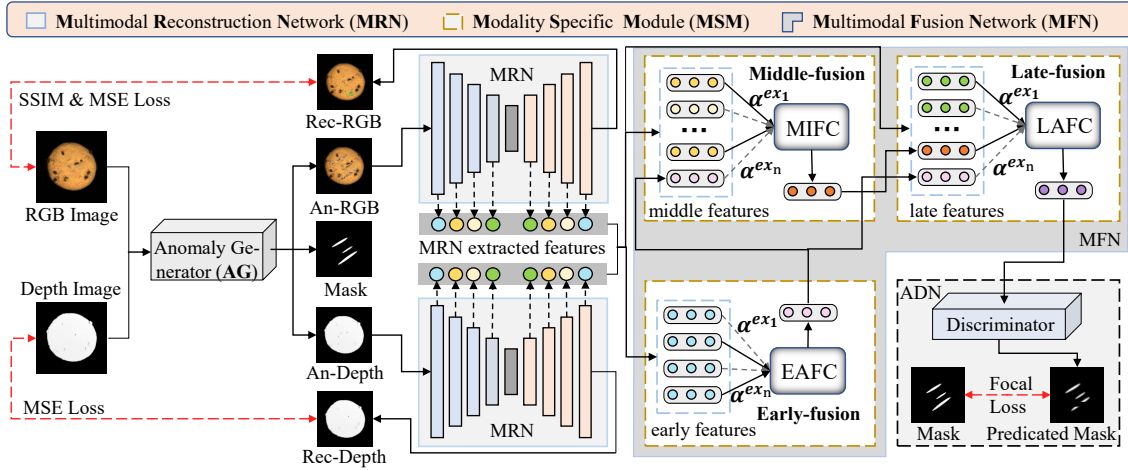


Figure 2: The overall framework of 3D-ADNAS, where the MFN architecture design is the core of this work, which is specified as a two-level search space: at the inter-module fusion level, the early fusion cell (EAFc), middle fusion cell (MIFc), and late fusion cell (LAFc) are configured to determine optimal combination of involved features and operations; at intra-module fusion level, it aims to seek best fusion strategy to combine those modules (MSMs).

multimodal fusion model, and then train the overall model of 3D-ADNAS with obtained MFN by 600 epochs.

Revisiting 3D-AD Fusion Architecture

Then, we conduct experimentation and theoretical analysis to verify the impact of fusion architectures on 3D-AD.

Impact of Inter-module Fusion

The three MSMs all show effectiveness in multimodal learning (Tu et al. 2024; Wang et al. 2024b). However, the impact of these fusion modules on 3D-AD remains unexplored. To bridge this gap, we conduct tests to scrutinize optimal combination of these MSMs on 3D-AD. For tests, we use EasyNet as the backbone for multimodal fusion of 3D-ADNAS. As shown in Fig. 4-a, we observe that these inter-module fusion strategies (i.e., the strategies about how to combine early, middle and late MSMs) indeed have significant impact on the performance of 3D-AD. For instance, when combined with middle MSM, the late MSM obtains obvious improvement of 3D-AD performance across most test tasks, but it shows harmfulness to 3D-AD if only using it itself. In most test cases, once the middle module is involved, the fusion module combination tends to have positive effect on 3D-AD, leading to significant performance improvement (see Table 4). However, this observation does not necessarily mean that the middle MSM is definitely superior to the others. So, we need to investigate the role of the intra-module fusion on 3D-AD.

Impact of Intra-module Fusion

For each MSM, the fusion of its involved multimodal features (e.g., early, middle, or late features) with fusion operations (e.g., addition (Gu et al. 2024b)) is the key to realize intra-module fusion of 3D-AD. Thus, following the settings of above tests, we perform a series of experiments by selecting different multimodal features and fusion operations within a specific MSM to explore their impact on 3D-AD. Here, the

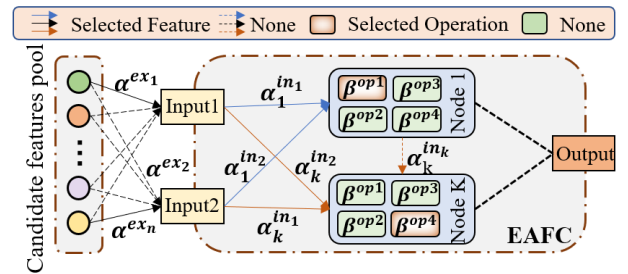


Figure 3: The inner structure of an MSM, where the early MSM is used as example with $\mathbb{K} = 2$. Note that the three types of MSMs share a similar structure.

middle MSM is used for test, which involves different stages of middle-level features. As shown in Fig. 4-b, we find that selecting partial layers of middle features is more effective for 3D-AD than selection all the features. From Fig. 4-c, we can see that the selection of fusion operations has important effect on 3D-AD, and employing multiple fusion operations is more beneficial with the fusion performance than only a single operation. Among these operations, the guided attention and summation operations (Lei et al. 2023) show best effectiveness in 3D-AD tasks. These results validate the actual impact of intra-module fusion design on 3D-AD.

Theoretical Analysis

Further, we theoretically analyze the impact of our target intra- and inter-module fusion design on 3D-AD using the Dempster-Shafer's evidence theory (DST) (Liu et al. 2017; Han et al. 2021). In the following, we first consider the impact of inter-module fusion design, where the DST's concepts including belief mass b and uncertainty u are utilized to evaluate the trustworthiness of target fusion model's outputs.

For an N -class classification task of MSM, given the

opinion of original MSM (e.g., the late MSM l): $L = \{\{b_l^n\}_{n=1}^N, u_l\}$, we aim to theoretically analyze whether the fusion of additional opinion (e.g., the middle MSM m) $M = \{\{b_m^n\}_{n=1}^N, u_m\}$ will influence the model's classification accuracy. Then, following the combination rule of DST, we fuse L into M and form a new opinion $F = \{\{b_f^n\}_{n=1}^N, u_f\}$. Here, b_f^n and u_f are new belief mass and uncertainty, respectively, which are given by

$$\begin{aligned} b_f^n &= (b_l^n b_m^n + b_l^n u_m + b_m^n u_l) / (1 - z), \\ u_f &= (u_l u_m) / (1 - z), \end{aligned} \quad (1)$$

where $z = \sum_{i \neq j} b_l^i b_m^j$ ($i, j \in [1, 2, \dots, N]$) is the measure of the conflict quantity between two belief mass sets of L and M , and $\frac{1}{1-z}$ is used as the normalization factor.

Then, we can give the following propositions (more detailed proofs are provided in the Supplementary Material¹).

Proposition 1. Under the conditions $b_m^g \geq b_l^{max}$, where $g \in N$ is the index of the ground-truth label, and b_l^{max} is the largest in $\{b_l^n\}_{n=1}^N$, fusing another opinion M makes the new opinion F satisfy $b_f^g \geq b_l^g$.

Proof.

$$\begin{aligned} b_f^g &= \frac{b_l^g b_m^g + b_l^g u_m + b_m^g u_l}{\sum_{n=1}^N b_m^n b_l^n + u_m + u_l - u_m u_l} \\ &\geq \frac{b_l^g b_m^g + b_l^g u_m + b_l^{max} u_l}{\sum_{n=1}^N b_m^n b_l^{max} + u_m + u_l - u_m u_l} \\ &\geq b_l^g \frac{(b_l^{max} + u_m + u_l)}{b_l^{max} + u_m + u_l} \geq b_l^g. \end{aligned} \quad (2)$$

Proposition 2. When u_m is large, $b_l^g - b_f^g$ will be limited, and it will have a negative correlation with u_m . As a special case, when u_m is large enough (i.e., $u_m = 1$), fusing another opinion will not reduce the performance (i.e., $b_f^g = b_l^g$).

Proof.

$$\begin{aligned} b_l^g - b_f^g &= b_l^g - \frac{b_l^g b_m^g + b_l^g u_m + b_m^g u_l}{\sum_{n=1}^N b_m^n b_l^n + u_m + u_l - u_m u_l} \\ &\leq b_l^g - \frac{b_l^g u_m}{b_m^{max} \cdot 1 + u_m + u_l - u_m u_l} \\ &\leq b_l^g - \frac{b_l^g u_m}{1 + u_l - u_m u_l} = b_l^g \frac{1 + u_l}{\frac{1}{(1-u_m)} + u_l}, \end{aligned} \quad (3)$$

To this end, we have the following conclusions: (i) According to **Proposition 1**, fusing an additional opinion (e.g., M) into the original opinion (e.g., L) has great potential to boost the model's accuracy. (ii) According to **Proposition 2**, the above fusion may lead to accuracy deterioration, but even this is limited under mild conditions.

For the impact of intra-module design, we can utilize the same analysis approach as mentioned above, and obtain the same conclusions as the case of the above inter-module design, which are provided in the Supplementary Material.

¹Please refer to the following for additional material: <https://github.com/longkaifang/3D-ADNAS>.

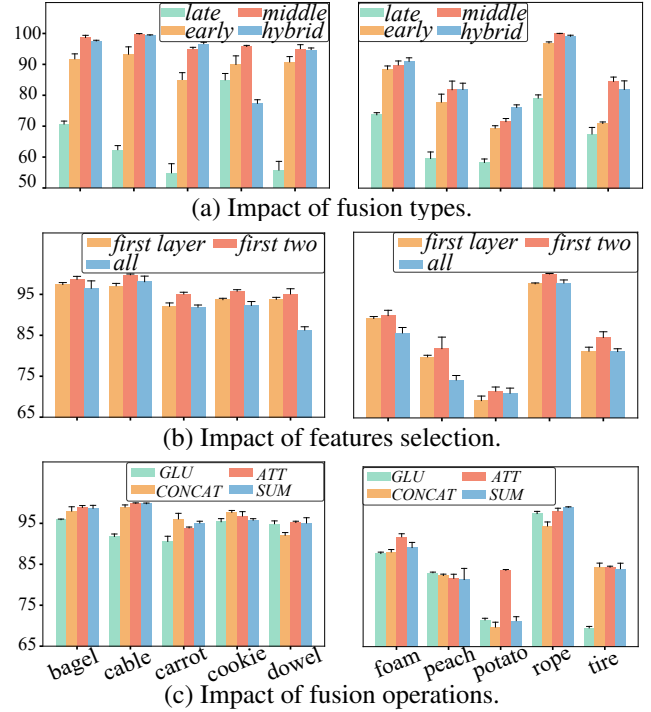


Figure 4: The impact of multimodal fusion architecture design on 3D-AD performance. Zoom in for details.

3D-ADNAS Method

To seek optimal architectural designs towards 3D-AD, we propose a simple yet powerful NAS method (3D-ADNAS) with a two-level search space and a simple search strategy grounded in gradient-based algorithms.

Search Space towards 3D-ADNAS

The above experimentation and theoretical analysis show the important role of the intra-module and inter-module fusion architectures in the 3D-AD skeleton. Therefore, to achieve a better 3D-AD performance, further exploration of the above two-level architectural design is necessary, i.e., the two-level design of 3D-AD-friendly search space.

Intra-module fusion level. The impact of various components, i.e., module-specific features (e.g., early, middle or late features), and fusion operations of the cell in each MSM, has not been thoroughly examined in prior experiments. Then, the optimal designs of the combination of these components, are necessary to search for, which needs to consider the following aspects.

What features are selected as cell inputs? As shown in Fig. 3, for each MSM, we need to select two features from the candidate feature pool (\mathbb{F}) as the inputs of the cell according to architectural parameters α^{ex} . That is, it is required to optimize α^{ex} and select the features with largest α^{ex} values as optimal solutions. To solve this problem, we utilize the continuous relaxation strategy of DARTS to convert the discrete feature selection problem into the continuous search problem, and then use gradient method to solve it. To achieve this, we

	Method	Year	Candy Cane	Chocolate Cookie	Chocolate Praline	Confetto	Gummy Bear	Hazelnut Truffle	Licorice Sandwich	Lollipop	Marsh-mallow	Peppermint Candy	Mean
RGB	EasyNet	MM23	72.3	92.5	84.9	96.6	70.5	81.5	80.6	85.1	97.5	96.0	85.8
	M3DM	CVPR23	64.8	94.9	94.1	100.0	87.8	63.2	93.3	81.1	99.8	100.0	87.9
	3D-ADNAS	-	79.8	99.8	92.6	100.0	83.5	86.2	91.2	96.0	97.8	100.0	92.7
3D	EasyNet	MM23	62.9	71.6	76.8	73.1	66.0	71.0	71.2	71.1	68.8	73.1	70.6
	M3DM	CVPR23	48.2	58.9	80.5	84.5	78.0	53.8	76.6	82.7	80.0	82.2	72.5
	3D-ADNAS	-	53.3	65.6	85.9	79.5	78.0	62.9	84.3	78.1	82.4	87.8	75.8
RGB+3D	EasyNet	MM23	73.7	93.4	86.6	96.6	71.7	82.2	84.7	86.3	97.7	96.0	86.9
	AST	WACV23	57.4	74.7	74.7	88.9	59.6	61.7	81.6	84.1	98.7	98.7	78.0
	M3DM	CVPR23	62.4	95.8	95.8	100.0	88.6	78.5	94.9	83.6	100.0	100.0	89.7
CFM	CVPR24	68.0	93.1	95.2	88.0	86.5	78.2	91.7	91.7	84.0	99.8	96.2	88.1
	3D-ADNAS	-	89.6	100.0	97.0	100.0	82.7	88.2	93.1	95.0	100.0	100.0	94.6

Table 1: I-AUROC scores on Eyecandies dataset. The red indicates the best results and the blue indicates the second best.

	Method	Year	Candy Cane	Chocolate Cookie	Chocolate Praline	Confetto	Gummy Bear	Hazelnut Truffle	Licorice Sandwich	Lollipop	Marsh-mallow	Peppermint Candy	Mean
RGB	M3DM	CVPR23	86.7	90.4	80.5	98.2	87.1	66.2	88.2	89.5	97.0	96.2	88.0
	3D-ADNAS	-	89.7	87.6	84.7	97.1	79.6	75.4	90.1	85.3	95.1	96.9	88.2
3D	M3DM	CVPR23	91.1	64.5	58.1	74.8	74.8	48.4	60.8	90.4	64.6	75.0	70.2
	3D-ADNAS	-	88.2	66.7	54.2	55.8	63.9	45.4	53.8	74.6	70.1	83.6	65.7
RGB+3D	AST	WACV23	51.4	83.5	71.4	90.5	58.7	59.0	73.6	76.9	91.8	87.8	74.4
	M3DM	CVPR23	90.6	92.3	80.3	98.3	85.5	68.8	88.0	90.6	96.6	95.5	88.2
	CFM	CVPR24	94.2	90.2	83.1	96.5	87.5	76.2	79.1	91.3	93.9	94.9	88.7
	3D-ADNAS	-	94.5	89.1	82.7	95.8	85.7	74.8	91.1	90.7	96.4	97.2	89.8

Table 2: AUPRO scores on Eyecandies dataset. The red indicates the best results and the blue indicates the second best.

reformulate the cell inputs through weighted summation of all candidate multimodal features in \mathbb{F} as:

$$\tilde{X}_i = \sum_{f_s \in \mathbb{F}} \frac{\exp(\alpha_i^{ex_i})}{\sum_{j=1}^{\mathbb{F}} \exp(\alpha_i^{ex_j})} \cdot f_s, \quad (4)$$

where \tilde{X}_i denotes the input features of each cell, $i \in [1, 2]$, f_s denotes the feature in \mathbb{F} . Accordingly, the two features serving as inputs of each cell’s can be determined by:

$$(f_s^j, f_s^h) = \arg \max_{i \in [1, \mathbb{F}]} (\alpha_1^{ex_i}, \alpha_2^{ex_i}), \quad (5)$$

where (f_s^j, f_s^h) denotes the choice of the j -th and h -th features from \mathbb{F} as cell inputs.

Similarly, following the above approach, we can also obtain the optimal connections (identified by α^{in}) between input nodes and intermediate nodes in each cell. More details are provided in the Supplementary Material.

What operations are selected for features fusion? To achieve best fusion of the above input features, we need to select appropriate fusion operations from the candidate operation pool \mathbb{O} (including addition, concatenation, GLU, and guided attention) for each intermediate node according to architectural parameters β^{op} . Similar to the case of feature selection, we relax the operations in \mathbb{O} to obtain the output of k -th intermediate node during the search, as follows:

$$\tilde{op}(t^y, t^z) = \sum_{op \in \mathbb{O}} \frac{\exp(\beta^{op})}{\sum_{op' \in \mathbb{O}} \exp(\beta^{op'})} op(t^y, t^z), \quad (6)$$

where op denotes the primitive operation in \mathbb{O} , t^y and t^z represent the input features of k -th intermediate node. Then, after gradient-based search, the best fusion operations for

each intermediate node can be obtained by discretizing β^{op} as follows:

$$op(\cdot) = \arg \max \beta^{op}. \quad (7)$$

Inter-module fusion level. After the optimization of the MSMs at the intra-module fusion level, a natural question is arisen: **how to combine those MSMs to achieve best final fusion performance?** For this issue, we formulate a simple yet effective inter-module fusion-level search space, where the output of early MSM is considered as a candidate input feature in the feature pool of the other two MSMs, the middle MSM output is considered as a candidate input feature of late MSM, and it is assigned with a weight α^{ex} , which can be optimized by gradient method. Finally, the output features of late MSM together with reconstructed image feature are fed into ADN for anomaly detection. In this way, we can search for optimal fusion strategy to integrate these MSNs in an efficient manner.

Search Strategy towards 3D-ADNAS

Since our relaxed continuous search space is differentiable, we can employ the popular gradient-based method to alternately optimize the architectural parameters (α^{ex} , α^{in} and β^{op}) and network weights (w) of 3D-ADNAS until the training of model converges. When the search is completed, we derive the optimal multimodal fusion architecture according to the values of the learned architectural parameters. In this work, for simplicity, we adopt the basic gradient descent method used in DARTS. Note that any other more effective gradient-based search paradigms can be applied in our target scenario for composite benefit. The detailed algorithm of 3D-ADNAS is provided in the Supplementary Material.

	Method	Year	Bagel	Cable Gland	Carrot	Cookie	Dowel	Foam	Peach	Potato	Rope	Tire	Mean
RGB	EasyNet	MM23	98.2	99.2	91.7	95.3	91.9	92.3	84.0	78.5	98.6	74.2	90.4
	M3DM	CVPR23	94.4	91.8	89.6	74.9	95.9	76.7	91.9	64.8	93.8	76.7	85.0
	3D-ADNAS	-	98.1	98.8	92.7	95.6	94.2	92.8	85.3	79.1	97.7	85.8	92.0
3D	EasyNet	MM23	73.5	67.8	74.7	86.4	71.9	71.6	71.3	72.5	88.5	68.7	74.7
	M3DM	CVPR23	94.1	65.1	96.5	96.9	90.5	76.0	88.0	97.4	92.6	76.5	87.4
	3D-ADNAS	-	79.4	85.7	69.9	94.6	69.5	68.6	70.5	87.3	95.3	66.7	78.8
RGB + 3D	BTF	CVPR23	93.8	76.5	97.2	88.8	96.0	66.4	90.4	92.9	98.2	72.6	87.3
	EasyNet	MM23	99.1	99.8	91.8	96.8	94.5	94.5	90.5	80.7	99.4	79.3	92.6
	AST	WACV23	98.3	87.3	97.6	97.1	93.2	88.5	97.4	98.1	100.0	79.7	93.7
	M3DM	CVPR23	99.4	90.9	97.2	97.6	96.0	94.2	97.3	89.9	97.2	85.0	94.5
	Shape_Guided	ICML23	98.6	89.4	98.3	99.1	97.6	85.7	99.0	96.5	96.0	86.9	94.7
3D-ADNAS	-	99.7	100.0	97.1	98.6	96.6	94.8	89.7	87.3	100.0	86.7	95.1	

Table 3: I-AUROC scores on MVTec 3D-AD dataset. The red indicates the best results and the blue indicates the second best.

MSM fusion components			Bagel	Cable Gland	Carrot	Cookie	Dowel	Foam	Peach	Potato	Rope	Tire	Mean
Early	Middle	Late											
✓	✗	✗	89.6	98.1	83.5	87.1	85.8	89.7	74.3	78.1	96.5	70.3	85.3
✗	✓	✗	99.5	99.9	94.5	95.6	96.2	91.6	85.7	74.0	100.0	83.7	92.1
✗	✗	✓	69.9	60.9	54.5	82.6	53.3	73.9	57.0	67.9	77.9	66.3	66.4
✓	✓	✗	99.4	100.0	92.9	95.4	93.9	93.3	84.7	83.8	99.8	87.1	93.0
✓	✗	✓	83.6	79.5	70.1	85.9	66.7	77.5	58.9	54.5	76.3	57.8	71.1
✗	✓	✓	97.4	100.0	96.2	78.5	94.3	89.8	82.9	76.9	98.8	84.6	89.9
✓	✓	✓	99.7	100.0	97.1	98.6	96.6	94.8	89.7	87.3	100.0	86.7	95.1

Table 4: Ablation study on different fusion components. The red indicates the best results and the blue shows the second best.

Method	MVTec 3D-AD		
	Frame Rate	Memory	I-AUROC
BTF	3.197	381.06	86.5
AST	4.966	463.94	93.7
M3DM	0.514	6526.12	94.5
3D-ADNAS (Ours)	24.693	269.33	95.1

Table 5: Comparison results in terms of frame rate, memory usage and accuracy metrics.

Metrics	Eyecandies			Full
	5-shot	10-shot	50-shot	
I-AUROC	77.5	80.7	86.8	94.6
P-AUROC	87.5	86.9	91.2	97.0
AUPRO	70.4	76.7	83.5	89.8

Table 6: Few-shot test results on Eyecandies.

Performance Evaluation towards 3D-ADNAS

We employ the I-AUROC, P-AUROC, and AUPRO (Tien et al. 2023) metrics to evaluate the overall performance of the searched 3D-ADNAS model with the optimal multimodal fusion topology. Specifically, our proposed model is trained and evaluated according to the settings in **Preliminaries**.

Evaluating the Improved Fusion Architecture

We then conduct empirical comparison experiments to validate the effectiveness of 3D-ADNAS.

Better 3D-AD performance

Benchmark Dataset. We evaluate 3D-ADNAS on Eyecandies (Bonfiglioli et al. 2022) and MVTec 3D-AD (Bergmann

Method	MVTec 3D-AD (I-AUROC)			
	5-shot	10-shot	50-shot	Full
BTF	67.1	69.5	80.6	86.5
AST	68.0	68.9	79.4	93.7
M3DM	82.2	84.5	90.7	94.5
3D-ADNAS (Ours)	82.6	84.8	89.0	95.1

Table 7: Few-shot test results on MVTec 3D-AD.

et al. 2022) datasets. Both contain 10 data categories, where MVTec 3D-AD is the real data collected from realistic scenes and Eyecandies is the synthetic virtual data.

Setup. We compare 3D-ADNAS with several baselines (AST (Rudolph et al. 2023), M3DM (Wang et al. 2023), EasyNet (Chen et al. 2023), CFM (Costanzino et al. 2024), BTF (Horwitz and Hoshen 2023), and Shape_Guide (Chu et al. 2023)) on both dataset. To make this comparison, we first reset the input image size to 256×256 adhering to the settings of (Chen et al. 2023) and then train the 3D-ADNAS with the Adam optimizer.

Evaluations on Eyecandies. Tables 1 and 2 show the comparison results between 3D-ADNAS and state-of-the-art (SOTA) methods on Eyecandies dataset. It is clear that 3D-ADNAS consistently outperforms the baselines in most test cases. Specifically, when both RGB and 3D depth images are used for training, 3D-ADNAS achieves 4.9% and 1.6% higher than M3DM in terms of I-AUROC and AUPRO metrics, respectively (16.6% and 15.4% higher than AST, and 6.5% and 1.1% higher than CFM, respectively). In fact, the above improvement surpasses the SOTA approaches. These experimental results validate the critical role of multimodal




Figure 5: Visualizations results on MVTEC 3D-AD.

fusion architectures in improving the 3D-AD performance.

Evaluations on MVTEC 3D-AD. Table 3 reports the I-AUROC scores of 3D-ADNAS on MVTEC 3D-AD dataset. As shown, we can see that 3D-ADNAS obtains best detection results in terms of both RGB-only and multimodal images in most test instances. In particular, 3D-ADNAS improves the I-AUROC value by 0.6% with 25 times less memory usage than M3DM, and by 1.4% with approximately 5 times faster frame rate than AST (while using comparable memory usage) (see Table 5). The above results and the visualization of Fig. 5 again demonstrate the effectiveness of our proposed multimodal fusion architecture design in enhancing 3D-AD. Due to the limitation of pages, more experimental results are provided in the Supplementary Material.

Impact of MSM Fusion Components

In this subsection, we again scrutinize the impact of the components of the multimodal fusion architecture design on 3D-AD tasks. As shown in Table 4, First, we can observe that in all test cases, 3D-ADNAS obtains better results when combining three MSM modules than only exploiting one MSM module, e.g., a 9.8% improvement over the early MSM, and a 3.0% improvement over the middle MSM. Second, the performance of 3D-AD can be improved when combining the middle MSM with other MSMs, e.g., when early MSM is combined with middle MSM, the performance is improved by 7.7% in terms of I-AUROC. Then, it is evident that our proposed two-level search space is able to reach a competitive multimodal fusion architecture for 3D-AD.

Higher Frame Rate and Lower Memory Usage

To further show the advantage of the proposed method, we report the results in terms of memory usage, frame rate, and I-AUROC on MVTEC 3D-AD. As shown in Table 5, we can see that 3D-ADNAS gets fastest frame rate, highest I-AUROC scores, and lowest memory usage when tested on single NVIDIA RTX 4090. **What contributes to the competitive performance of 3D-ADNAS?** Intuitively, it benefits from the fact that our method uses neither memory-bank-based strategies that increase memory nor large pre-trained language models that affect inference speed. In this sense, revisiting multimodal fusion from an architectural perspective is indeed an effective scheme for advancing 3D-AD.

Study on Few-shot Anomaly Detection

In resource-constrained scenarios, collecting a large number of samples is extremely expensive and infeasible. Thus, few-

shot anomaly detection (Kim et al. 2024; Duan et al. 2023) becomes a promising solution. To evaluate the effectiveness of 3D-ADNAS in few-shot scenarios, we randomly select 5, 10, and 50 images on Eyecandies and MVTEC 3D-AD datasets as the training set and perform inference on the full test set. As shown in Tables 6 and 7, 3D-ADNAS still exhibits promising performance.

Related Work

3D Anomaly Detection. Existing studies mainly focus on design of 2D-AD methods, i.e., detecting flaws in RGB images (Bergmann et al. 2019; Jiang et al. 2022; Hu et al. 2024), including feature-embedding-based methods (Liu et al. 2023; Li et al. 2021; Rudolph et al. 2022; Rudolph, Wandt, and Rosenhahn 2021; Deng and Li 2022; Bergmann et al. 2020; Cohen and Hoshen 2020; Defard et al. 2021) and reconstruction-based methods (Schlüter et al. 2022; Zavrtanik, Kristan, and Skočaj 2021; You et al. 2022; Zavrtanik, Kristan, and Skočaj 2022). Recently, 3D-AD has emerged as an improvement of 2D-3D, and received a surge of attention (Liu et al. 2024; Xie et al. 2024), since it exhibits a powerful detection ability via exploiting both RGB and depth images rather than only RGB (Reiss et al. 2022; Zavrtanik, Kristan, and Skočaj 2024b; Zhao et al. 2024). A series of 3D-AD methods have been proposed and developed, e.g., BTF (Horwitz and Hoshen 2023), M3DM (Wang et al. 2023), and CFM (Costanzino et al. 2024). However, existing works rarely focus on the impact of architectural design on 3D-AD. To bridge a gap, we systematically study the impact of multimodal fusion architecture design via theoretical analysis and experiments, and then propose a simple yet effective method.

Neural Architecture Search. NAS aims to automate the design of task-specific deep neural network architectures, which can be formulated as an optimization problem (Baker et al. 2017; Zoph et al. 2018). In this process, NAS uses a search strategy to traverse a specified search space comprising candidate architectures. Then, the architecture with the best performance is selected as the final design. It has been empirically demonstrated that NAS can shape architectures that surpass those manually designed (Yu et al. 2020; Lv et al. 2024). Note that, the key of applying NAS to 3D-AD lies in defining a suitable search space, which depicts a searchable subset of candidate architectures from the vast architecture space. Given the limited prior knowledge regarding the impact of architectural designs on 3D-AD tasks, defining a suitable search space remains challenging.

Conclusion

This paper investigates the impact of two-level multimodal fusion architecture design (including intra-module and inter-module fusion levels) on 3D-AD tasks. The proposed 3D-ADNAS bridges the gap between multimodal fusion architecture design and 3D-AD, achieving comprehensive improvement of 3D-AD performance in terms of detection accuracy, frame rate and memory usage. We believe that devising a friendly multimodal fusion architecture is practically meaningful for 3D-AD, and hope this work inspires further research on 3D-AD from the architectural perspective.

Acknowledgments

This work is supported by the National Natural Science Foundation of China under Grant 62472079.

References

- Baker, B.; Gupta, O.; Naik, N.; and Raskar, R. 2017. Designing Neural Network Architectures using Reinforcement Learning. In *International Conference on Learning Representations*.
- Bergmann, P.; Fauser, M.; Sattlegger, D.; and Steger, C. 2019. MVTEC AD - A Comprehensive Real-World Dataset for Unsupervised Anomaly Detection. In *Conference on CVPR*, 9592–9600.
- Bergmann, P.; Fauser, M.; Sattlegger, D.; and Steger, C. 2020. Uninformed students: Student-teacher anomaly detection with discriminative latent embeddings. In *Proceedings of the IEEE/CVF conference on CVPR*.
- Bergmann, P.; Jin, X.; Sattlegger, D.; and Steger, C. 2022. The MVTEC 3D-AD Dataset for Unsupervised 3D Anomaly Detection and Localization. In *Proceedings of the 17th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*, 202–213.
- Bergmann, P.; and Sattlegger, D. 2023. Anomaly detection in 3d point clouds using deep geometric descriptors. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2613–2623.
- Bonfiglioli, L.; Toschi, M.; Silvestri, D.; Fioraio, N.; and De Gregorio, D. 2022. The eyecandies dataset for unsupervised multimodal anomaly detection and localization. In *Proceedings of the Asian Conference on Computer Vision*, 3586–3602.
- Cao, Y.; Xu, X.; and Shen, W. 2024. Complementary pseudo multimodal feature for point cloud anomaly detection. *Pattern Recognition*, 110761.
- Chen, R.; Xie, G.; Liu, J.; Wang, J.; Luo, Z.; Wang, J.; and Zheng, F. 2023. EasyNet: An easy network for 3d industrial anomaly detection. In *Proceedings of the 31st ACM International Conference on Multimedia*, 7038–7046.
- Chu, Y.-M.; Liu, C.; Hsieh, T.-I.; Chen, H.-T.; and Liu, T.-L. 2023. Shape-guided dual-memory learning for 3D anomaly detection. In *Proceedings of the 40th International Conference on Machine Learning*, 6185–6194.
- Cohen, N.; and Hoshen, Y. 2020. Sub-image anomaly detection with deep pyramid correspondences. *arXiv preprint arXiv:2005.02357*.
- Costanzino, A.; Ramirez, P. Z.; Lisanti, G.; and Di Stefano, L. 2024. Multimodal industrial anomaly detection by cross-modal feature mapping. In *Proceedings of the IEEE/CVF Conference on CVPR*, 17234–17243.
- Dai, S.; Wu, Y.; Li, X.; and Xue, X. 2024. Generating and reweighting dense contrastive patterns for unsupervised anomaly detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1454–1462.
- Defard, T.; Setkov, A.; Loesch, A.; and Audigier, R. 2021. PaDiM: a patch distribution modeling framework for anomaly detection and localization. In *International Conference on Pattern Recognition*, 475–489.
- Deng, H.; and Li, X. 2022. Anomaly detection via reverse distillation from one-class embedding. In *Proceedings of the IEEE/CVF conference on CVPR*, 9737–9746.
- Duan, Y.; Hong, Y.; Niu, L.; and Zhang, L. 2023. Few-shot defect image generation via defect-aware feature manipulation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 571–578.
- Gu, Z.; Zhang, J.; Liu, L.; Chen, X.; Peng, J.; Gan, Z.; Jiang, G.; Shu, A.; Wang, Y.; and Ma, L. 2024a. Rethinking Reverse Distillation for Multi-Modal Anomaly Detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 8445–8453.
- Gu, Z.; Zhu, B.; Zhu, G.; Chen, Y.; Tang, M.; and Wang, J. 2024b. Anomalygpt: Detecting industrial anomalies using large vision-language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1932–1940.
- Han, Z.; Zhang, C.; Fu, H.; and Zhou, J. T. 2021. Trusted Multi-View Classification. In *International Conference on Learning Representations*.
- Horwitz, E.; and Hoshen, Y. 2023. Back to the feature: classical 3d features are (almost) all you need for 3d anomaly detection. In *Proceedings of the IEEE/CVF Conference on CVPR*, 2968–2977.
- Hu, T.; Zhang, J.; Yi, R.; Du, Y.; Chen, X.; Liu, L.; Wang, Y.; and Wang, C. 2024. Anomalydiffusion: Few-shot anomaly image generation with diffusion model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 8526–8534.
- Jiang, X.; Liu, J.; Wang, J.; Nie, Q.; Wu, K.; Liu, Y.; Wang, C.; and Zheng, F. 2022. Softpatch: Unsupervised anomaly detection with noisy data. *Advances in Neural Information Processing Systems*, 35: 15433–15445.
- Kim, S.; An, S.; Chikontwe, P.; Kang, M.; Adeli, E.; Pohl, K. M.; and Park, S. H. 2024. Few Shot Part Segmentation Reveals Compositional Logic for Industrial Anomaly Detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 8591–8599.
- Lei, J.; Hu, X.; Wang, Y.; and Liu, D. 2023. Pyramidflow: High-resolution defect contrastive localization using pyramid normalizing flow. In *Proceedings of the IEEE/CVF conference on CVPR*, 14143–14152.
- Li, C.-L.; Sohn, K.; Yoon, J.; and Pfister, T. 2021. Cutpaste: Self-supervised learning for anomaly detection and localization. In *Proceedings of the IEEE/CVF conference on CVPR*, 9664–9674.
- Li, W.; Xu, X.; Gu, Y.; Zheng, B.; Gao, S.; and Wu, Y. 2024. Towards Scalable 3D Anomaly Detection and Localization: A Benchmark via 3D Anomaly Synthesis and A Self-Supervised Learning Network. In *Proceedings of the IEEE/CVF Conference on CVPR*, 22207–22216.
- Liu, H.; Simonyan, K.; and Yang, Y. 2018. DARTS: Differentiable Architecture Search. In *International Conference on Learning Representations*.
- Liu, J.; Xie, G.; Wang, J.; Li, S.; Wang, C.; Zheng, F.; and Jin, Y. 2024. Deep industrial image anomaly detection: A survey. *Machine Intelligence Research*, 21(1): 104–135.

- Liu, Y.-T.; Pal, N. R.; Marathe, A. R.; and Lin, C.-T. 2017. Weighted fuzzy Dempster–Shafer framework for multimodal information integration. *IEEE Transactions on Fuzzy Systems*, 26(1): 338–352.
- Liu, Z.; Zhou, Y.; Xu, Y.; and Wang, Z. 2023. Simplenet: A simple network for image anomaly detection and localization. In *Proceedings of the IEEE/CVF conference on CVPR*, 20402–20411.
- Lv, J.; Sun, Y.; Ye, Q.; Feng, W.; and Lv, J. 2024. A multiscale neural architecture search framework for multimodal fusion. *Information Sciences*, 121005.
- Pérez-Rúa, J.-M.; Vielzeuf, V.; Pateux, S.; Baccouche, M.; and Jurie, F. 2019. Mfas: Multimodal fusion architecture search. In *Proceedings of the IEEE/CVF Conference on CVPR*, 6966–6975.
- Qin, J.; Gu, C.; Yu, J.; and Zhang, C. 2023. Image-Pointcloud Fusion based Anomaly Detection using PD-REAL Dataset. *arXiv preprint arXiv:2311.04095*.
- Reiss, T.; Cohen, N.; Horwitz, E.; Abutbul, R.; and Hoshen, Y. 2022. Anomaly detection requires better representations. In *European Conference on Computer Vision*, 56–68.
- Rudolph, M.; Wandt, B.; and Rosenhahn, B. 2021. Same same but different: Semi-supervised defect detection with normalizing flows. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1907–1916.
- Rudolph, M.; Wehrbein, T.; Rosenhahn, B.; and Wandt, B. 2022. Fully convolutional cross-scale-flows for image-based defect detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*.
- Rudolph, M.; Wehrbein, T.; Rosenhahn, B.; and Wandt, B. 2023. Asymmetric student-teacher networks for industrial anomaly detection. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, 2592–2602.
- Schlüter, H. M.; Tan, J.; Hou, B.; and Kainz, B. 2022. Natural synthetic anomalies for self-supervised anomaly detection and localization. In *European Conference on Computer Vision*, 474–489.
- Sui, W.; Lichau, D.; Lefèvre, J.; and Phelippeau, H. 2024. Cross-Modal Distillation in Industrial Anomaly Detection: Exploring Efficient Multi-Modal IAD. *arXiv preprint arXiv:2405.13571*.
- Tien, T. D.; Nguyen, A. T.; Tran, N. H.; Huy, T. D.; Duong, S.; Nguyen, C. D. T.; and Truong, S. Q. 2023. Revisiting reverse distillation for anomaly detection. In *Proceedings of the IEEE/CVF conference on CVPR*, 24511–24520.
- Tu, Y.; Zhang, B.; Liu, L.; Li, Y.; Xu, C.; Zhang, J.; Wang, Y.; Wang, C.; and Zhao, C. R. 2024. Self-supervised Feature Adaptation for 3D Industrial Anomaly Detection. *arXiv preprint arXiv:2401.03145*.
- Wang, C.; Zhu, H.; Peng, J.; Wang, Y.; Yi, R.; Wu, Y.; Ma, L.; and Zhang, J. 2024a. M3DM-NR: RGB-3D Noisy-Resistant Industrial Anomaly Detection via Multimodal Denoising. *arXiv preprint arXiv:2406.02263*.
- Wang, J.; Wang, X.; Hao, R.; Yin, H.; Huang, B.; Xu, X.; and Liu, J. 2024b. Incremental Template Neighborhood Matching for 3D anomaly detection. *Neurocomputing*, 581: 127483.
- Wang, Y.; Peng, J.; Zhang, J.; Yi, R.; Wang, Y.; and Wang, C. 2023. Multimodal industrial anomaly detection via hybrid fusion. In *Proceedings of the IEEE/CVF Conference on CVPR*, 8032–8041.
- Xie, G.; Wang, J.; Liu, J.; Lyu, J.; Liu, Y.; Wang, C.; Zheng, F.; and Jin, Y. 2024. Im-iad: Industrial image anomaly detection benchmark in manufacturing. *IEEE Transactions on Cybernetics*.
- Xu, Z.; So, D. R.; and Dai, A. M. 2021. Mufasa: Multimodal fusion architecture search for electronic health records. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 10532–10540.
- Yin, Y.; Huang, S.; and Zhang, X. 2022. Bm-nas: Bilevel multimodal neural architecture search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 8901–8909.
- You, Z.; Cui, L.; Shen, Y.; Yang, K.; Lu, X.; Zheng, Y.; and Le, X. 2022. A unified model for multi-class anomaly detection. *Advances in Neural Information Processing Systems*, 35: 4571–4584.
- Yu, Z.; Cui, Y.; Yu, J.; Wang, M.; Tao, D.; and Tian, Q. 2020. Deep multimodal neural architecture search. In *Proceedings of the 28th ACM International Conference on Multimedia*, 3743–3752.
- Zavrtnik, V.; Kristan, M.; and Skočaj, D. 2021. Draem-a discriminatively trained reconstruction embedding for surface anomaly detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, 8330–8339.
- Zavrtnik, V.; Kristan, M.; and Skočaj, D. 2022. DSR–A Dual Subspace Re-Projection Network for Surface Anomaly Detection. In *European Conference on Computer Vision*, 539–554.
- Zavrtnik, V.; Kristan, M.; and Skočaj, D. 2024a. Cheating depth: Enhancing 3d surface anomaly detection via depth simulation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2164–2172.
- Zavrtnik, V.; Kristan, M.; and Skočaj, D. 2024b. Keep DRÆMing: Discriminative 3D anomaly detection through anomaly simulation. *Pattern Recognition Letters*, 181: 113–119.
- Zhang, J.; and Li, W. 2023. Multi-Modal and Multi-Scale Temporal Fusion Architecture Search for Audio-Visual Video Parsing. In *Proceedings of the 31st ACM International Conference on Multimedia*, 3328–3336.
- Zhao, B.; Xiong, Q.; Zhang, X.; Guo, J.; Liu, Q.; Xing, X.; and Xu, X. 2024. PointCore: Efficient Unsupervised Point Cloud Anomaly Detector Using Local-Global Features. *arXiv preprint arXiv:2403.01804*.
- Zhou, Z.; Wang, L.; Fang, N.; Wang, Z.; Qiu, L.; and Zhang, S. 2024. R3D-AD: Reconstruction via Diffusion for 3D Anomaly Detection. *arXiv preprint arXiv:2407.10862*.
- Zoph, B.; Vasudevan, V.; Shlens, J.; and Le, Q. V. 2018. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE conference on CVPR*, 8697–8710.