TOPOOT: TOPOLOGY-AWARE OPTIMAL TRANSPORT FOR TEST-TIME ANOMALY SEGMENTATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Deep topological data analysis (TDA) offers a principled framework for capturing structural invariants such as connectivity and cycles that persist across scales, making it a natural fit for anomaly segmentation (AS). Unlike threshold-based binarisation, which produces brittle masks under distribution shift, TDA allows anomalies to be characterised as disruptions to global structure rather than local fluctuations. We introduce TopoOT, a topology-aware optimal transport (OT) framework that integrates multi-filtration persistence diagrams with test-time adaptation (TTA). Our key innovation is Optimal Transport Chaining, which sequentially aligns persistence diagrams (PDs) across thresholds and filtrations, yielding geodesic stability scores that identify features consistently preserved across scales. These stability-aware pseudo-labels supervise a lightweight head trained online with OT-consistency and contrastive objectives, ensuring robust adaptation under domain shift. Across standard 2D and 3D anomaly detection benchmarks, TopoOT achieves state-of-the-art performance¹, outperforming the most competitive methods by up to +24.1% mean F1 on 2D datasets and +10.2% on 3D anomaly segmentation benchmarks.

1 Introduction

Test-time training (TTT) has emerged as a promising paradigm for adapting models under distribution shift, but most approaches remain limited to entropy minimisation or feature consistency, without structured reasoning about data geometry Sun et al. (2020); Volpi et al. (2022); Zhang et al. (2022). A central limitation of many existing TTT approaches, particularly in dense prediction tasks, is their reliance on heuristic pseudo-labels or confidence thresholds Liang et al. (2024); Costanzino et al. (2024a); Zhang et al. (2025), which are non-robust (brittle) under distribution shift. Incorporating explicit structural priors provides a principled way to address this gap. The integration of TDA, which extracts persistent features such as connectivity and holes across scales Zia et al. (2024), and OT, which provides a principled framework for aligning distributions Cuturi (2013); Peyré et al. (2019), has received little attention in this context. AS is a particularly compelling domain in which to explore this integration, because it requires pixel-level localisation of irregular patterns whose connectivity and shape are critical, yet conventional threshold-based binarisation often collapses under shift Cao et al. (2024). By combining TDA's ability to capture structural persistence with OT's alignment capabilities, TTT can move beyond heuristics and yield more stable and adaptive anomaly delineation.

AS demands fine-grained identification of abnormal regions in test images, typically without access to anomalous training examples Tao et al. (2022). Most existing methods generate continuous anomaly maps that must be binarised Cao et al. (2024), but thresholds derived from nominal data are brittle across categories and anomaly types Tong et al. (2024); Wu et al. (2024); Zhou et al. (2024). Supervised approaches Baitieva et al. (2024); Hu et al. (2024b); Zhu et al. (2024); Ding et al. (2022) can achieve strong performance but

¹For reproducibility, our implementation is included with this submission

> ngn

 require extensive annotation, which is impractical for rare or heterogeneous anomalies Xie & Mirmehdi (2007); Qiu et al. (2019). Unsupervised methods Guo et al. (2025); He et al. (2024) are trained only on nominal data, rely on static thresholds, and fail to preserve structural consistency under domain shift.

Beyond the reliance on brittle thresholds, current approaches to AS and TTA face several underexplored challenges. First, robustness under distribution shift remains insufficient, benchmarks such as MVTec-AD Bergmann et al. (2019), VisA Zou et al. (2022), and Real-IAD Wang et al. (2024a) often understate the variability of anomalies, yet in practice, even minor domain shifts can cause embeddings or thresholds to fail catastrophically. Second, AS research has concentrated on 2D image settings, leaving structural guidance in 3D anomaly detection and segmentation (AD&S) largely unaddressed Li et al. (2024), despite its importance in industrial inspection. Third, pseudo-labels used in existing TTT frameworks are often derived from entropy or heuristic criteria, providing no guarantees of structural consistency across runs or domains Zhao et al. (2024). Finally, while efficiency is critical for deployment, there has been little exploration of methods that simultaneously remove threshold dependence and remain lightweight enough for real-time adaptation.

These gaps underscore the need for a framework that (i) eliminates brittle thresholding, (ii) stabilises noisy structural descriptors, (iii) incorporates explicit priors into TTA, and (iv) extends naturally to 3D settings. We propose **TopoOT**, a framework that stabilises pseudo-labels using multi-scale topological cues via persistent homology and aligns them with OT, providing structure-aware supervision for TTT. Although our experiments focus on AS, we view this task as the most natural and demanding testbed for a first exploration of structurally guided TTT, since anomalies disrupt connectivity, boundaries, and higher-order organisation, precisely the features that TDA and OT are designed to capture. Establishing effectiveness in this setting provides a foundation for broader machine learning tasks where structural stability is critical, including domain adaptation under distribution shift Dan et al. (2024), weak-signal detection in scientific data, and fine-grained visual analysis Michaeli & Fried (2024), where subtle structural cues determine class boundaries Zia et al. (2024). TopoOT embeds structural alignment into the TTT framework. The key contributions are:

- To overcome threshold brittleness, we introduce an **OT-guided, structure-aware representation** that integrates multi-scale topological cues from PDs. This representation produces pseudo-labels that provide adaptive and data-driven supervision for TTT.
- To stabilise noisy topological descriptors, we propose a novel **OT chaining** mechanism that aligns PDs both within a filtration (*cross-PD*) and across sub- and super-level filtrations (*cross-level*), retaining only consistently transported features and discarding spurious ones.
- To integrate structural priors into TTT, we design a lightweight head trained online with two complementary objectives: **OT-consistency**, which preserves *transport-aligned structures*, *and contrastive separation*, which sharpens anomalous versus nominal boundaries.
- Our approach is *plug-and-play*, integrating seamlessly with different backbones and extending naturally across modalities, generalising from 2D to 3D AD&S (point clouds and multimodal anomaly detection), where connectivity and shape priors are especially critical.

Across diverse datasets, our design consistently delivers robust and generalisable AS. Evaluated on 5 2D/3D benchmarks and 7 backbones, TopoOT achieves F1 gains up to +24.1% on 2D and +10.2% on 3D compared to the existing SOTA. It further generalises across models and domains, surpassing TTT baselines by up to +4.8%. TopoOT remains highly efficient, running at 121 FPS while using only 349 MB of GPU memory.

2 RELATED WORK

Anomaly Detection and Segmentation: AS under distribution shift is challenging as it requires fine-grained detection without supervision, structural priors that capture meaningful data characteristics, and adaptation to unseen test-time distributions. Unsupervised AD&S avoids labelled anomalies by learning from nominal data

He et al. (2024). Early reconstruction-based methods used autoencoders Fang et al. (2023); Park et al. (2024); Zuo et al. (2024); Zhou et al. (2025); Wang et al. (2024b), inpainting Li et al. (2020); Nakanishi et al. (2022); Zavrtanik et al. (2021b); Pirnay & Chai (2022); Luo et al. (2024), or diffusion models Yao et al. (2024a); Fučka et al. (2025); Jiang et al. (2024), but often produced blurry reconstructions or overfit to normal patterns. Feature-based approaches compare embeddings to nominal references Park et al. (2024); Roth et al. (2022); Defard et al. (2021), or use teacher–student frameworks Deng & Li (2024); Rudolph et al. (2023); Zhang et al. (2023); Gu et al. (2024) for inductive bias. Generative priors via normalizing flows Yao et al. (2024b); Gudovskiy et al. (2022); Lei et al. (2023); Kim et al. (2023) or synthetic anomalies Aota et al. (2023); Li et al. (2024); Hu et al. (2024a); Chen et al. (2024) improved detection, yet typically lack pixel-level precision. Methods such as PatchCore Roth et al. (2022) and PaDiM Defard et al. (2021) leverage pre-trained backbones, but remain threshold-dependent and structurally agnostic.

Optimal Transport in Vision: OT has been widely applied in computer vision for distribution alignment Peyré et al. (2019); Cuturi (2013); Bonneel & Digne (2023), including domain adaptation Ge et al. (2021); Fan et al. (2024); Luo & Ren (2023), object detection, and image restoration Adrai et al. (2023). In anomaly detection, Liao et al. (2025) employed robust Sinkhorn distances for industrial inspection. These works show OT's adaptability for handling domain discrepancies, but they typically operate at the distribution level and do not exploit OT for structured feature selection or test-time supervision. While our approach employs a novel *OT chaining* mechanism, entropically regularised OT helps align PDs through cross-PD filtration to capture feature evolution and cross-level filtration to integrate complementary structures, thereby preserving consistently transported features and discarding spurious ones.

Topological Priors and Test-Time Training: TDA, particularly persistent homology (PH), has been applied in medical imaging to capture shape and multi-scale structure Adcock et al. (2014); Berry et al. (2020); Crawford et al. (2020); Garside et al. (2019); Kanari et al. (2018). Yet most uses are offline and not integrated into adaptive learning Zia et al. (2024). TTT Liang et al. (2024); Nado et al. (2020); Kim et al. (2022); Colomer et al. (2023); Nguyen et al. (2023); Khurana et al. (2021) adapts models on-the-fly with self-supervised objectives, and TTT4AS Costanzino et al. (2024a) extended this idea to AS with heuristic pseudo-labels. However, these lack explicit structural reasoning and remain sensitive to noise.

Our approach combines PH-based filtrations with OT alignment to derive stable pseudo-labels, which then guide a lightweight TTT head. This integration moves beyond heuristic thresholds by embedding structural priors directly into TTA, yielding robust and topologically consistent AS.

3 OT-GUIDED TEST TIME STRUCTURAL ALIGNMENT FRAMEWORK

Problem Formulation: Conventional AS methods produce a dense anomaly score map and obtain binary masks through thresholds calibrated on nominal validation data Costanzino et al. (2024a) (e.g., percentile rules). Such thresholds are dataset-specific, fail under distribution shift, and often generate masks that under-cover or over-extend the anomalous region. Moreover, they operate pixel-wise and neglect structural information in the anomaly map. To address these limitations, we represent anomaly maps as persistence diagrams (PDs), which capture multi-scale topological features such as connected components and holes. Figure 1 provides an overview of our proposed TopoOT framework. We then introduce an OT-based scoring scheme that evaluates PDs across filtrations and levels, ranking components by their cross-scale consistency. This formulation replaces fixed thresholding with a structural scoring approach designed to produce more consistent anomaly masks under distribution shift.

Building on this, persistence diagrams derived from sub- and super-level filtrations provide the candidate anomaly structures. We apply OT alignment across filtration levels to retain components that persist with low transport cost, while discarding unstable features (that don't persist across PDs). The ranked components are

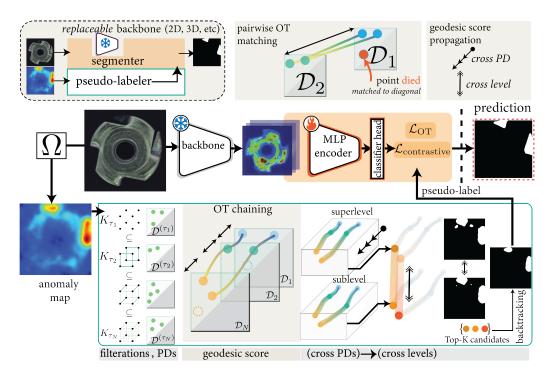


Figure 1: **TopoOT Test-time Training for Anomaly Segmentation.** (Top Left) pipeline simplified view. (Bottom) detailed view. TopoOT replaces conventional thresholding by stabilising anomaly evidence via cross-PD OT matching within each filtration, then fusing sub- and super-level scores with cross-level OT. The resulting global scores yield Top-K pseudo-labels that supervise a lightweight head for final segmentation.

then back-projected into the image domain to form pseudo-labels, which serve as data-dependent supervision at inference in place of fixed thresholds.

During TTT, we keep the replaceable backbone frozen and update only a lightweight head. This head is optimised with two complementary objectives: (i) OT-consistency, which encourages predictions to remain aligned with the stable structures identified by OT, and (ii) contrastive separation, which increases the margin between anomalous and nominal regions. The combination of these objectives yields a segmentation mask that is guided by OT-derived pseudo-labels rather than fixed thresholds.

3.1 Multi-scale Filtering as Feature Generation

We start from a continuous anomaly map $A:\Omega\to[0,1]$ defined over the pixel lattice Ω , same as Costanzino et al. (2024a). To capture structural variation at multiple thresholds, we fix a sequence of increasing thresholds $\mathcal{T}=\{\tau_1<\tau_2<\dots<\tau_N\}$. For each $\tau_k\in\mathcal{T}$, we define the sublevel and superlevel sets $S^{\mathrm{sub}}_{\tau_k}=\{p\in\Omega:A(p)\leq\tau_k\}$, and $S^{\mathrm{sup}}_{\tau_k}=\{p\in\Omega:A(p)\geq\tau_k\}$. These subsets naturally induce cubical complexes $K^{\mathrm{sub}}_{\tau_k}$, $K^{\mathrm{sup}}_{\tau_k}$, where each cell corresponds to a contiguous block of pixels (a cube in the grid) included whenever its vertices satisfy the relevant threshold condition. The "cubical" construction is appropriate for images/grids, because it respects the pixel adjacency and can be computed efficiently.

By varying τ_k , we obtain nested sequences (filtrations) for each kind of level set: $K_{\tau_1}^f \subset K_{\tau_2}^f \subset \cdots \subset K_{\tau_N}^f$, where $f \in \{\text{sub}, \text{sup}\}$. From these filtrations, we compute persistent homology in dimensions

 $h \in \{0,1\}$. The result is a persistence diagram $\mathcal{P}_f^h[\tau_k]$ at each threshold level, where H_0 (homology in dimension 0) captures connected components that is how new components appear (birth) and merge (death) across thresholds. H_1 (1-dimensional homology) captures loops or holes (voids), features that appear in superlevel or sublevel sets and disappear at some higher (or lower) threshold. Background on cubical complexes in Appendix A.8.

Each topological feature c in a diagram is represented as a pair (b_c, d_c) of birth and death times; its persistence $\operatorname{pers}(c) = d_c - b_c$ reflects how long it persists. Features with large persistence are more likely to correspond to "meaningful" structural anomalies, while those close to the diagonal (small persistence) are often noise. These ideas align with the discussion review paper by Zia et al. (2024), which emphasises that PDs and barcodes are robust summaries of topological features of data across cales, invariant to small perturbations, deformation, and noise. The outputs $\{\mathcal{P}_{\mathrm{sub}}^h[\tau_k]\}_{k=1}^N$, and $\{\mathcal{P}_{\mathrm{sup}}^h[\tau_k]\}_{k=1}^N$ serve as multi-scale candidate features. They form the input to the OT-based alignment steps. Rather than acting as direct decision thresholds, these persistence diagrams are treated as a rich feature generation mechanism, capturing anomalies' connected components and holes over multiple scales, which allows the downstream optimal transport stage to judge stability and discriminability among structural candidates.

3.2 GEODESIC SCORING OF TOPOLOGICAL FEATURES

The persistence diagrams derived from sub- and super-level filtrations provide a rich but noisy set of candidate features. Many short-lived components arise due to local perturbations in the anomaly map, which, if treated directly, would degrade the reliability of pseudo-labels. A key challenge is how to aggregate these diverse features into a concise set of components that can be meaningfully traced back to the original image. A possible solution is computing a barycenter of diagrams Turner et al. (2014), but barycenters discard the natural order of filtrations and blur fine-scale structures. Mapping diagrams into kernels or persistence images Reininghaus et al. (2015) is another alternative, but these yield global embeddings without interpretable correspondences. In contrast, we propose aggregating information by following the flow of diagrams within each filtration sequence using Optimal Transport Chaining. This approach consolidates features into stable representatives for both the sublevel and superlevel filtrations independently, and then fuses the two levels to obtain consensus features.

Formally, let $P = \{p_i = (b_i, d_i)\}_{i=1}^m$ and $Q = \{q_j = (b'_j, d'_j)\}_{j=1}^n$ be two persistence diagrams, used here as shorthand for $\{\mathcal{P}_f^h[\tau_k]\}$ at different thresholds or filtrations. We define the ground cost as the squared Euclidean distance between pairs of features, and compute the entropic OT plan:

$$\Pi^{\star} = \underset{\Pi \in \mathcal{U}(P,Q)}{\operatorname{arg\,min}} \langle C, \Pi \rangle + \varepsilon H(\Pi)$$
(1)

where $\mathcal{U}(P,Q)$ denotes the set of admissible couplings between P and Q, and $H(\Pi)$ is the entropy of the transport plan. The regularisation parameter $\varepsilon>0$ ensures numerical stability and smooth alignments. In our framework, all transport plans are therefore entropy-regularised Sinkhorn solutions rather than exact Wasserstein couplings because they yield smooth, differentiable, and numerically stable alignments; see Appendix A.6 for further details.

We exploit this transport plan through a novel **OT chaining** mechanism, which consists of two complementary modes: cross-PD (intra) filtration and cross-level (inter) filtration alignment. In cross-PD filtration alignment, OT is applied within a single filtration (sublevel or superlevel) between persistence diagrams at different thresholds τ_k and τ_ℓ . This process identifies features that persist consistently through the filtration, and each candidate c receives a stability score:

$$s(c) = \max_{j} \frac{\Pi^{\star}(i(c), j)}{1 + \sqrt{C(i(c), j)}} \cdot \alpha \operatorname{pers}(c)$$
 (2)

where Π^{\star} is the entropic OT plan between diagrams, C(i(c),j) is the ground cost, and $\operatorname{pers}(c)$ is the persistence of c as defined in Sec. 3.1. Since C is defined as squared Euclidean distances, we use \sqrt{C} in the denominator to restore a linear distance scale, ensuring that score decay is proportional to distance rather than quadratic. This softens penalisation and allows moderately stable matches to contribute, instead of filtering too aggressively. The maximisation is taken over all possible partners j of candidate c within the filtration, where j indexes features in the comparison persistence diagram. In this way, s(c) reflects the strongest OT-stable match. When points don't get matched between PDs, they are coupled to the diagonal as in standard TDA practice (see Sec. 3.1 and Appendix A.7), ensuring that chain stability scores naturally account for vanishing features. The factor $a \geq 0$ controls the influence of persistence on ranking. Top-M components are selected by maximising stability and persistence and minimising transport cost.

In cross-level filtration alignment, we compare candidate sets from the sublevel and superlevel filtrations. Applying OT across sublevel and superlevel filtrations integrates complementary topological cues. Sublevel filtrations emphasise how connected components emerge and merge, while superlevel filtrations highlight how voids and holes evolve. By aligning these perspectives, the method retains structural features that are consistently expressed across both, thereby suppressing spurious components and strengthening anomaly cues. Each candidate c is evaluated with the same stability score s(c) defined above, but here the partner set is drawn from the opposite filtration. This ensures that features are retained only if they exhibit both cross-PD scale persistence and cross-level filtration consistency. The top-K ranked candidates across both filtrations are then collected to form the final set C^* .

The surviving candidates in \mathcal{C}^* are projected back to their pixel-level supports on the anomaly map, yielding OT-guided pseudo-labels \widetilde{Y}_{OT} . These pseudo-labels are inherently multi-scale and data-adaptive, as they emerge from stable OT couplings rather than fixed thresholds. These retained features correspond to connected regions or holes, e.g., defects or gaps, that persist across the filtration process and reflect semantically meaningful structures in the input space. By filtering out noise-induced artefacts, OT alignment produces pseudo-labels that provide robust supervision for TTT.

3.3 TOPOOT TEST-TIME TRAINING

The final stage of our pipeline leverages the OT-guided pseudo-labels $\widetilde{Y}_{\mathrm{OT}}$ to adapt the model during inference. Since the backbone feature extractor is frozen, adaptation is performed through a lightweight segmentation head h_{ψ} attached to the anomaly map representation. This design ensures that the adaptation cost at test time remains negligible, while still allowing the predictions to be tailored to the distribution of the current sample. Training h_{ψ} is guided by two complementary objectives. First, we introduce an **OT-consistency** loss that encourages the segmentation head h_{ψ} to reproduce the spatial structures encoded in $\widetilde{Y}_{\mathrm{OT}}$. Given the deviations from the OT-aligned pseudo-labels $\mathcal{L}_{\mathrm{OT}} = \|\widehat{Y} - \widetilde{Y}_{\mathrm{OT}}\|_2$ which enforces consistency with stable transport couplings and prevents overfitting. Second, we incorporate a margin-based contrastive objective to sharpen local decision boundaries in the embedding space produced by h_{ψ} . From the OT-derived pseudo-labels $\widetilde{Y}_{\mathrm{OT}} \in \{0,1\}^{H \times W}$, we sample pixel pairs (p,q) as similar when $\widetilde{Y}_{\mathrm{OT}}(p) = \widetilde{Y}_{\mathrm{OT}}(q)$ and dissimilar otherwise. Let $z_p, z_q \in \mathbb{R}^D$ denote the L2-normalised embeddings of those pixels. The contrastive loss is:

$$\mathcal{L}_{\text{contrastive}} = (1 - y_{pq}) \|z_p - z_q\|_2^2 + y_{pq} \left[\max(0, m - \|z_p - z_q\|_2) \right]^2$$

where $y_{pq} \in \{0,1\}$ encodes dissimilarity and m>0 is a margin. This loss compacts same-label embeddings while enforcing a minimum separation between background and anomalous regions, improving robustness to residual noise in $\widetilde{Y}_{\rm OT}$. The combined loss is $\mathcal{L}_{\rm TTT} = \mathcal{L}_{\rm OT} + \lambda \mathcal{L}_{\rm contrastive}$ with λ controlling the balance between structural consistency and contrastive separation. By optimising $\mathcal{L}_{\rm TTT}$ on each test sample, the segmentation head h_{ψ} adapts to dataset-specific distributions without requiring external supervision. The final segmentation mask $\widehat{Y}^{\rm bin}$ is obtained through a canonical decision rule applied to the adapted predictions

of h_{ψ} . Because h_{ψ} is trained on OT-guided pseudo-labels, this rule is adaptive to each test instance, avoiding dataset-specific calibration and eliminating heuristic threshold tuning.

This test-time regularisation departs from conventional schemes in two ways: (i) it grounds the adaptation signal in OT-aligned structures, stable across multi-scale filtrations, rather than raw anomaly scores; (ii) by integrating contrastive separation, it sharpens class boundaries instead of collapsing toward trivial solutions.

Stability and Generalisation Bounds: Complementing our empirical results, we provide preliminary formal analysis in Appendix A.7 showing that our scores are robust to perturbations, discrepancies accumulate subadditively across filtrations, and OT-consistency objectives are connected to improved generalisation at test time. They offer principled intuition for the design of cross-PD and cross-level OT chaining and its role in structurally guided adaptation.

4 EXPERIMENTAL SETUP

Datasets, Backbones, and Evaluation Protocol: We evaluate across both 2D and 3D anomaly detection benchmarks. For 2D, RGB datasets MVTec AD Bergmann et al. (2019), VisA Zou et al. (2022), and Real-IAD Wang et al. (2024a) are used with backbones PatchCore Roth et al. (2022), PaDiM Defard et al. (2021), Dinomaly Guo et al. (2025), and MambaAD He et al. (2024). For 3D, we consider multimodal MVTec 3D-AD (RGB + point-cloud) Bergmann et al. (2021) and pure point-cloud Anomaly-ShapeNet Li et al. (2024), using backbones CMM Costanzino et al. (2024b), M3DM Wang et al. (2023), and PO3AD Ye et al. (2025). While we report standard anomaly-detection metrics such as image-level AUROC (I-AUROC), pixel-level AUROC (P-AUROC), and pixel-level AUPRO (P-AUPRO) for completeness, our evaluation focuses on pixel-level Precision, Recall, F1, and IoU of the final binary masks. AUROC and AUPRO mainly assess ranking quality and can remain high despite poor mask quality under severe pixel imbalance Bergmann et al. (2019); Zavrtanik et al. (2021a). In contrast, Precision, Recall, and F1 capture the accuracy of detected defect regions, balancing missed detections and false alarms, while IoU offers a stringent measure of spatial overlap Costanzino et al. (2024a). These metrics align more closely with industrial inspection needs, where the fidelity of the delivered mask is the decisive criterion Bergmann et al. (2020); Schlüter et al. (2022).

Across both domains, we compare all methods against the TTT baseline **TTT4AS** Costanzino et al. (2024a). Following **TTT4AS**, we binarise each backbone's AS map at the statistical threshold ($\mu + c\sigma$) and report this variant (**THR**) alongside the **TTT4AS** baseline. All experiments have been conducted on an NVIDIA RTX 5090 GPU with 32GB of VRAM. Detailed hyperparameters and architectural settings are provided in Appendix A.1. TopoOT runs at 121 FPS using 349 MB GPU memory for 2D inference; 3D inference has comparable memory use but lower FPS due to point-cloud operations. Per-dataset timing and memory profiles are given in Appendix A.2.

5 RESULTS AND DISCUSSION

We validate TopoOT through analyses: (i) **2D and 3D AD&S**, benchmarking against state-of-the-art methods; (ii) **Cross Model Domain Adaptation**, where frozen feature extractors are paired with distinct anomaly score maps across 2D and 3D datasets; and (iii) **Ablation Studies**, assessing the contribution of each component. For detailed discussion of limitations and directions for future development, including efficiency tradeoffs and backbone dependency, refer to Appendix A.3.

5.1 2D/3D AD&S

We present a comprehensive evaluation of **TopoOT** across five diverse datasets and seven state-of-the-art backbones. The I-AUROCP, P-AUROC, and P-AUPRO metrics are computed directly from each backbone's

The results in Table 1 demonstrate superiority, with **TopoOT** consistently outperforming all baselines. The metrics are the mean per class within each dataset. Our method achieves a **+38.6%** F1 gain over **THR** and **+14.0%** over **TTT4AS** Costanzino et al. (2024a) on MVTec AD (**PatchCore** Roth et al. (2022)). For **PaDiM**, it surpasses **THR** by **+20.5%** and **TT4AS** by **+24.1%**. On VisA, it surpasses **TTT4AS** by **+19.7%** (**Dinomaly** Guo et al. (2025)) and **+8.5%** (**MambaAD** He et al. (2024)). For Real-IAD, **TopoOT** shows a **+12.3%** and **+11.8%** F1 improvement over **THR**, and a **+21.3%** and **+20.9%** gain over **TTT4AS** for the **Dinomaly** and **MambaAD** backbones, respectively. The advantage extends to 3D, with gains of **+20.7%** (**CMM** Costanzino et al. (2024b)) and **+24.5%** (**M3DM** Wang et al. (2023)) over **THR** on MVTec 3D-AD, alongside **+10.2%** and **+2.2%** improvements over **TTT4AS**. On AnomalyShapeNet (**PO3AD** Ye et al. (2025)), **TopoOT** also leads with a **+2.9%** and **+1.9%** F1 advantage.

Table 1: Comparison of binary segmentation results. Best results in **bold**; second-best in blue.

AS map, while our method operates on the resulting anomaly maps to produce final binary segmentations.

Dataset	Backbone	I-AUROC	P-AUROC	P-AUPRO	TTT Method	Prec.	Rec.	F1	IoU
MVTec AD	PatchCore Roth et al. (2022)	0.991	0.981	0.934	THR Roth et al. (2022) TTT4AS Costanzino et al. (2024a) TopoOT	0.351 0.388 0.550	0.507 0.648 0.720	0.136 0.382 0.522	0.299 0.293 0.387
Bergmann et al. (2019)	PaDiM Defard et al. (2021)	0.979	0.975	0.921	THR Roth et al. (2022) TTT4AS Costanzino et al. (2024a) TopoOT	0.452 0.330 0.470	0.507 0.579 0.788	0.354 0.318 0.559	0.317 0.274 0.402
VisA	Dinomaly Guo et al. (2025)	0.987	0.987	0.945	THR Guo et al. (2025) TTT4AS Costanzino et al. (2024a) TopoOT	0.275 0.223 0.546	0.862 0.811 0.553	0.339 0.267 0.464	0.144 0.177 0.223
Zou et al. (2022)	MambaAD He et al. (2024)	0.943	0.985	0.910	THR He et al. (2024) TTT4AS Costanzino et al. (2024a) TopoOT	0.200 0.223 0.416	0.785 0.811 0.507	0.241 0.267 0.352	0.196 0.130 0.247
Real IAD	Dinomaly Guo et al. (2025)	0.893	0.989	0.939	THR Wang et al. (2024a) TTT4AS Costanzino et al. (2024a) TopoOT	0.242 0.154 0.461	0.793 0.801 0.577	0.317 0.229 0.442	0.208 0.147 0.317
Wang et al. (2024a)	MambaAD He et al. (2024)	0.863	0.985	0.905	THR He et al. (2024) TTT4AS Costanzino et al. (2024a) TopoOT	0.188 0.084 0.305	0.653 0.763 0.616	0.228 0.137 0.346	0.145 0.080 0.243
MVTec 3D-AD	CMM Costanzino et al. (2024b)	0.954	0.993	0.971	THR Costanzino et al. (2024b) TTT4AS Costanzino et al. (2024a) TopoOT	0.199 0.303 0.427	0.902 0.800 0.845	0.275 0.380 0.482	0.232 0.077 0.343
Bergmann et al. (2021)	M3DM Wang et al. (2023)	0.945	0.992	0.964	THR Wang et al. (2023) TTT4AS Costanzino et al. (2024a) TopoOT	0.173 0.467 0.564	0.889 0.640 0.767	0.245 0.468 0.490	0.232 0.120 0.364
AnomalyShapeNet Li et al. (2024)	PO3AD Ye et al. (2025)	0.839	0.898	0.821	THR Ye et al. (2025) TTT4AS Costanzino et al. (2024a) TopoOT	0.675 0.562 0.651	0.441 0.485 0.540	0.500 0.510 0.529	0.371 0.347 0.402

Figure 2 shows that **TopoOT** yields sharper, more semantically coherent segmentations than competing methods. **TopoOT** secures concurrent gains in precision and recall, which in turn increase **IoU**, resulting in consistently superior segmentations across every benchmark. *Per-class quantitative and qualitative results* for each dataset are presented in the Appendix A.4 & A.5. TopoOT consistently achieves sharper boundaries and higher recall across categories. Even in challenging cases like thin or fragmented defects, it remains robust, clearly outperforming other methods across both 2D and 3D domains.

5.2 Cross Model Domain Adaptation

We validate a plug-and-play transfer strategy that pairs frozen *source* feature extractors with distinct *tar-get* scoring heads across 2D (MVTec, VisA, Real-IAD) and 3D (MVTec-3DAD) domains. As shown in Table 2, the cross-model pipelines preserve topological structure and deliver practical quality without retraining. In 2D, transfers reach F1 up to **0.512** on Real-IAD (PatchCore—MambaAD) and **0.502** on VisA (MambaAD—Dinomaly), with recalls in the **0.71–0.75** band; in 3D, CMM—M3DM offers the highest precision (**0.471**, F1 **0.479**), while M3DM—CMM provides broad coverage (recall **0.791**). Importantly, these domain-adaptation results outperform established baselines across the evaluated datasets, confirming

Figure 2: Qualitative comparison of AD&S methods for different objects using the MVTec 3D-AD dataset.

effective cross-model composition and providing a strong substrate for TopoOT to further consolidate gains via stability-aware OT pseudo-labels and adaptive boundary refinement for AS.

Table 2: Cross-model domain adaptation (features \rightarrow anomaly scores).

Mod	lality	Dataset	$\mathbf{Source} \rightarrow \mathbf{Target}$	Prec.	Rec.	F1
2D	3D		$(Features \rightarrow Anomaly Scores)$			
1		MVTec	PatchCore → PaDiM	0.419	0.673	0.430
1		VisA	$MambaAD \rightarrow Dinomaly$	0.459	0.712	0.502
1		Real-IAD	$PatchCore \rightarrow MambaAD$	0.434	0.750	0.512
	✓	MVTec-3DAD	CMM → M3DM	0.471	0.746	0.479
	✓	MVTec-3DAD	$M3DM \rightarrow CMM$	0.409	0.791	0.469

5.3 ABLATION STUDIES

We validate TopoOT (Table 3). Individual cross-PD filtration alignments yield modest gains. The cross-level filtration alignment is key, providing a larger boost by integrating cross-scale information. The losses \mathcal{L}_{OT} and $\mathcal{L}_{\text{contrastive}}$ are effective together, enforcing prediction consistency and feature separation, respectively. Our complete model achieves top performance: **0.522** F1 on PatchCore, **0.482** on CMM, and **0.490** on M3DM.

Table 3: Ablation study showing that combining all OT alignments with losses yields the highest performance.

Ī		Торо	OT Components			2D	-PatchC	ore	3	D-CMN	1	3	D-M3DI	M
	$cross-PD_{Sub}$	$cross$ - PD_{Super}	$cross$ -level $_{Sub\text{-}super}$	\mathcal{L}_{OT}	$\mathcal{L}_{contrastive}$	Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1
Ī	✓			√		0.440	0.310	0.365	0.410	0.455	0.382	0.290	0.730	0.390
li	✓				✓	0.490	0.540	0.475	0.426	0.485	0.415	0.310	0.740	0.405
li		✓		✓		0.375	0.620	0.390	0.085	0.820	0.118	0.280	0.755	0.380
li		✓			✓	0.395	0.605	0.408	0.095	0.830	0.132	0.300	0.760	0.392
li	✓	✓	✓	✓		0.520	0.690	0.510	0.420	0.800	0.470	0.500	0.750	0.485
	✓	✓	✓		✓	0.510	0.680	0.505	0.405	0.770	0.460	0.490	0.740	0.475
ľ	✓	✓	✓	✓	✓	0.550	0.720	0.522	0.427	0.845	0.482	0.564	0.767	0.490

CONCLUSION

We presented TopoOT, a topology-aware OT framework for anomaly segmentation that replaces brittle thresholding with OT-guided pseudo-labels and stabilises multi-scale persistence features through cross-PD and cross-level filtration chaining. A lightweight head trained with OT-consistency and contrastive objectives enables per-instance TTA that preserves structural stability while sharpening anomaly boundaries. TopoOT achieves SOTA performance on five standard benchmarks, and our theoretical analysis establishes stability and generalisation guarantees.

REFERENCES

- Aaron Adcock, Daniel Rubin, and Gunnar Carlsson. Classification of hepatic lesions using the matching metric. *Computer vision and image understanding*, 121:36–42, 2014.
- Theo Adrai, Guy Ohayon, Michael Elad, and Tomer Michaeli. Deep optimal transport: A practical algorithm for photo-realistic image restoration. *Advances in Neural Information Processing Systems*, 36:61777–61791, 2023.
- Toshimichi Aota, Lloyd Teh Tzer Tong, and Takayuki Okatani. Zero-shot versus many-shot: Unsupervised texture anomaly detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 5564–5572, 2023.
- Aimira Baitieva, David Hurych, Victor Besnier, and Olivier Bernard. Supervised anomaly detection for complex industrial images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17754–17762, 2024.
- Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79:151–175, 2010.
- Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Mytec ad–a comprehensive real-world dataset for unsupervised anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9592–9600, 2019.
- Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. The mytec anomaly detection dataset: A comprehensive real-world dataset for unsupervised anomaly detection. *International Journal of Computer Vision (IJCV)*, 129(4):1038–1059, 2020.
- Paul Bergmann, Xin Jin, David Sattlegger, and Carsten Steger. The mytec 3d-ad dataset for unsupervised 3d anomaly detection and localization. *arXiv* preprint arXiv:2112.09045, 2021.
- Eric Berry, Yen-Chi Chen, Jessi Cisewski-Kehe, and Brittany Terese Fasy. Functional summaries of persistence diagrams. *Journal of Applied and Computational Topology*, 4(2):211–262, 2020.
- Nicolas Bonneel and Julie Digne. A survey of optimal transport for computer graphics and computer vision. In *Computer Graphics Forum*, volume 42, pp. 439–460. Wiley Online Library, 2023.
- Yunkang Cao, Xiaohao Xu, Jiangning Zhang, Yuqi Cheng, Xiaonan Huang, Guansong Pang, and Weiming Shen. A survey on visual anomaly detection: Challenge, approach, and prospect. *arXiv preprint arXiv:2401.16402*, 2024.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 9650–9660, 2021.
- Qiyu Chen, Huiyuan Luo, Chengkan Lv, and Zhengtao Zhang. A unified anomaly synthesis strategy with gradient ascent for industrial anomaly detection and localization. *arXiv preprint arXiv:2407.09359*, 2024.
- Marc Botet Colomer, Pier Luigi Dovesi, Theodoros Panagiotakopoulos, Joao Frederico Carvalho, Linus Härenstam-Nielsen, Hossein Azizpour, Hedvig Kjellström, Daniel Cremers, and Matteo Poggi. To adapt or not to adapt? real-time adaptation for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 16548–16559, 2023.
- Alex Costanzino, Pierluigi Zama Ramirez, Mirko Del Moro, Agostino Aiezzo, Giuseppe Lisanti, Samuele Salti, and Luigi Di Stefano. Test time training for industrial anomaly segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3910–3920, 2024a.

Alex Costanzino, Pierluigi Zama Ramirez, Giuseppe Lisanti, and Luigi Di Stefano. Multimodal industrial anomaly detection by crossmodal feature mapping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17234–17243, 2024b.

Lorin Crawford, Anthea Monod, Andrew X Chen, Sayan Mukherjee, and Raúl Rabadán. Predicting clinical outcomes in glioblastoma: an application of topological and functional data analysis. *Journal of the American Statistical Association*, 115(531):1139–1150, 2020.

Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.

Jun Dan, Weiming Liu, Chunfeng Xie, Hua Yu, Shunjie Dong, and Yanchao Tan. Tfgda: Exploring topology and feature alignment in semi-supervised graph domain adaptation through robust clustering. In *Advances in Neural Information Processing Systems*, 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/59e73ff865b56cba6ab7f6b2cce1425d-Abstract-Conference.html. NeurIPS 2024.

Thomas Defard, Aleksandr Setkov, Angelique Loesch, and Romaric Audigier. Padim: a patch distribution modeling framework for anomaly detection and localization. In *International Conference on Pattern Recognition*, pp. 475–489. Springer, 2021.

Hanqiu Deng and Xingyu Li. Structural teacher-student normality learning for multi-class anomaly detection and localization. arXiv preprint arXiv:2402.17091, 2024.

Choubo Ding, Guansong Pang, and Chunhua Shen. Catching both gray and black swans: Open-set supervised anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 7388–7398, 2022.

Qizhe Fan, Xiaoqin Shen, Shihui Ying, and Shaoyi Du. Otclda: Optimal transport and contrastive learning for domain adaptive semantic segmentation. *IEEE Transactions on Intelligent Transportation Systems*, 25 (10):14685–14697, 2024.

Zheng Fang, Xiaoyang Wang, Haocheng Li, Jiejie Liu, Qiugui Hu, and Jimin Xiao. Fastrecon: Few-shot industrial anomaly detection via fast feature reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 17481–17490, 2023.

Matic Fučka, Vitjan Zavrtanik, and Danijel Skočaj. Transfusion—a transparency-based diffusion model for anomaly detection. In *European Conference on Computer Vision*, pp. 91–108. Springer, 2025.

Kathryn Garside, Robin Henderson, Irina Makarenko, and Cristina Masoller. Topological data analysis of high resolution diabetic retinopathy images. *PloS one*, 14(5):e0217413, 2019.

Zheng Ge, Songtao Liu, Zeming Li, Osamu Yoshie, and Jian Sun. Ota: Optimal transport assignment for object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 303–312, 2021.

Haocheng Gu, Gongyang Li, and Zhi Liu. Masked feature regeneration based asymmetric student–teacher network for anomaly detection. *Multimedia Tools and Applications*, pp. 1–22, 2024.

Denis Gudovskiy, Shun Ishizaka, and Kazuki Kozuka. Cflow-ad: Real-time unsupervised anomaly detection with localization via conditional normalizing flows. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pp. 98–107, 2022.

Jia Guo, Shuai Lu, Weihang Zhang, Fang Chen, Huiqi Li, and Hongen Liao. Dinomaly: The less is more philosophy in multi-class unsupervised anomaly detection. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 20405–20415, 2025.

- Haoyang He, Yuhu Bai, Jiangning Zhang, Qingdong He, Hongxu Chen, Zhenye Gan, Chengjie Wang, Xiangtai Li, Guanzhong Tian, and Lei Xie. Mambaad: Exploring state space models for multi-class unsupervised anomaly detection. *Advances in Neural Information Processing Systems*, 37:71162–71187, 2024.
- Jie Hu, Yawen Huang, Yilin Lu, Guoyang Xie, Guannan Jiang, and Yefeng Zheng. Anomalyxfusion: Multi-modal anomaly synthesis with diffusion. *arXiv preprint arXiv:2404.19444*, 2024a.
- Teng Hu, Jiangning Zhang, Ran Yi, Yuzhen Du, Xu Chen, Liang Liu, Yabiao Wang, and Chengjie Wang. Anomalydiffusion: Few-shot anomaly image generation with diffusion model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 8526–8534, 2024b.
- Bolin Jiang, Yuqiu Xie, Jiawei Li, Naiqi Li, Yong Jiang, and Shu-Tao Xia. Cagen: Controllable anomaly generator using diffusion model. In *ICASSP 2024-2024 IEEE International Conference on Acoustics*, *Speech and Signal Processing (ICASSP)*, pp. 3110–3114. IEEE, 2024.
- Lida Kanari, Paweł Dłotko, Martina Scolamiero, Ran Levi, Julian Shillcock, Kathryn Hess, and Henry Markram. A topological representation of branching neuronal morphologies. *Neuroinformatics*, 16:3–13, 2018.
- Ansh Khurana, Sujoy Paul, Piyush Rai, Soma Biswas, and Gaurav Aggarwal. Sita: Single image test-time adaptation. *arXiv preprint arXiv:2112.02355*, 2021.
- Daehyun Kim, Sungyong Baik, and Tae Hyun Kim. Sanflow: semantic-aware normalizing flow for anomaly detection and localization. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, pp. 75434–75454, 2023.
- Junho Kim, Inwoo Hwang, and Young Min Kim. Ev-tta: Test-time adaptation for event-based object recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17745–17754, 2022.
- Jiarui Lei, Xiaobo Hu, Yue Wang, and Dong Liu. Pyramidflow: High-resolution defect contrastive localization using pyramid normalizing flow. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 14143–14152, 2023.
- Wenqiao Li, Xiaohao Xu, Yao Gu, Bozhong Zheng, Shenghua Gao, and Yingna Wu. 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 computer vision and pattern recognition*, pp. 22207–22216, 2024.
- Zhenyu Li, Ning Li, Kaitao Jiang, Zhiheng Ma, Xing Wei, Xiaopeng Hong, and Yihong Gong. Superpixel masking and inpainting for self-supervised anomaly detection. In *Bmvc*, 2020.
- Jian Liang, Ran He, and Tieniu Tan. A comprehensive survey on test-time adaptation under distribution shifts. *International Journal of Computer Vision*, pp. 1–34, 2024.
- Jingyi Liao, Xun Xu, Yongyi Su, Rong-Cheng Tu, Yifan Liu, Dacheng Tao, and Xulei Yang. Robust distribution alignment for industrial anomaly detection under distribution shift. *arXiv preprint arXiv:2503.14910*, 2025.

Wei Luo, Haiming Yao, Wenyong Yu, and Zhengyong Li. Ami-net: Adaptive mask inpainting network for industrial anomaly detection and localization. *IEEE Transactions on Automation Science and Engineering*, 2024.

- You-Wei Luo and Chuan-Xian Ren. Mot: Masked optimal transport for partial domain adaptation. In 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 3531–3540. IEEE, 2023.
- Eyal Michaeli and Ohad Fried. Advancing fine-grained classification by structure and subject preserving augmentation. In *Advances in Neural Information Processing Systems*, 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/27c8b849acba6793f0b73f7ee7ea7397-Abstract-Conference.html. NeurIPS 2024.
- Zachary Nado, Shreyas Padhy, D Sculley, Alexander D'Amour, Balaji Lakshminarayanan, and Jasper Snoek. Evaluating prediction-time batch normalization for robustness under covariate shift. *arXiv* preprint *arXiv*:2006.10963, 2020.
- Hitoshi Nakanishi, Masahiro Suzuki, and Yutaka Matsuo. Fixing the train-test objective discrepancy: Iterative image inpainting for unsupervised anomaly detection. *Journal of Information Processing*, 30:495–504, 2022.
- A Tuan Nguyen, Thanh Nguyen-Tang, Ser-Nam Lim, and Philip HS Torr. Tipi: Test time adaptation with transformation invariance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 24162–24171, 2023.
- Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.
- YeongHyeon Park, Sungho Kang, Myung Jin Kim, Hyeonho Jeong, Hyunkyu Park, Hyeong Seok Kim, and Juneho Yi. Neural network training strategy to enhance anomaly detection performance: A perspective on reconstruction loss amplification. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5165–5169. IEEE, 2024.
- Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- Jonathan Pirnay and Keng Chai. Inpainting transformer for anomaly detection. In *International Conference on Image Analysis and Processing*, pp. 394–406. Springer, 2022.
- Lingteng Qiu, Xiaojun Wu, and Zhiyang Yu. A high-efficiency fully convolutional networks for pixel-wise surface defect detection. *IEEE Access*, 7:15884–15893, 2019. doi: 10.1109/ACCESS.2019.2894420.
- Jan Reininghaus, Stefan Huber, Ulrich Bauer, and Roland Kwitt. A stable multi-scale kernel for topological machine learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4741–4748, 2015.
- Karsten Roth, Latha Pemula, Joaquin Zepeda, Bernhard Schölkopf, Thomas Brox, and Peter Gehler. Towards total recall in industrial anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 14318–14328, 2022.
- Marco Rudolph, Tom Wehrbein, Bodo Rosenhahn, and Bastian Wandt. Asymmetric student-teacher networks for industrial anomaly detection. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pp. 2592–2602, 2023.

- Stefan Schlüter, David Borth, Tim Weninger, and Marco F. Huber. Deep one-class classification for defect detection in industry. In *International Conference on Pattern Recognition (ICPR)*, pp. 1950–1957, 2022.
 - Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In *International conference on machine learning*, pp. 9229–9248. PMLR, 2020.
 - Xian Tao, Xinyi Gong, Xin Zhang, Shaohua Yan, and Chandranath Adak. Deep learning for unsupervised anomaly localization in industrial images: A survey. *IEEE Transactions on Instrumentation and Measurement*, 71:1–21, 2022.
 - Guoxiang Tong, Quanquan Li, and Yan Song. Enhanced multi-scale features mutual mapping fusion based on reverse knowledge distillation for industrial anomaly detection and localization. *IEEE Transactions on Big Data*, 10(4):498–513, 2024. doi: 10.1109/TBDATA.2024.3350539.
 - Katharine Turner, Yuriy Mileyko, Sayan Mukherjee, and John Harer. Fréchet means for distributions of persistence diagrams. *Discrete & Computational Geometry*, 52(1):44–70, 2014.
 - Riccardo Volpi, Pau De Jorge, Diane Larlus, and Gabriela Csurka. On the road to online adaptation for semantic image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 19184–19195, 2022.
 - Chengjie Wang, Wenbing Zhu, Bin-Bin Gao, Zhenye Gan, Jiangning Zhang, Zhihao Gu, Shuguang Qian, Mingang Chen, and Lizhuang Ma. Real-iad: A real-world multi-view dataset for benchmarking versatile industrial anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 22883–22892, 2024a.
 - Junpu Wang, Guili Xu, Chunlei Li, Guangshuai Gao, and Yuehua Cheng. Multi-feature reconstruction network using crossed-mask restoration for unsupervised anomaly detection. *arXiv preprint arXiv:2404.13273*, 2024b.
 - Yue Wang, Jinlong Peng, Jiangning Zhang, Ran Yi, Yabiao Wang, and Chengjie Wang. Multimodal industrial anomaly detection via hybrid fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8032–8041, 2023.
 - Qiangwei Wu, Hui Li, Chenyu Tian, Long Wen, and Xinyu Li. Aekd: Unsupervised auto-encoder knowledge distillation for industrial anomaly detection. *Journal of Manufacturing Systems*, 73:159–169, 2024.
 - Xianghua Xie and Majid Mirmehdi. Texems: Texture exemplars for defect detection on random textured surfaces. *IEEE transactions on pattern analysis and machine intelligence*, 29(8):1454–1464, 2007.
 - Hang Yao, Ming Liu, Haolin Wang, Zhicun Yin, Zifei Yan, Xiaopeng Hong, and Wangmeng Zuo. Glad: Towards better reconstruction with global and local adaptive diffusion models for unsupervised anomaly detection. *arXiv preprint arXiv:2406.07487*, 2024a.
 - Xincheng Yao, Ruoqi Li, Zefeng Qian, Lu Wang, and Chongyang Zhang. Hierarchical gaussian mixture normalizing flow modeling for unified anomaly detection. *arXiv preprint arXiv:2403.13349*, 2024b.
 - Jianan Ye, Weiguang Zhao, Xi Yang, Guangliang Cheng, and Kaizhu Huang. Po3ad: Predicting point offsets toward better 3d point cloud anomaly detection. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 1353–1362, 2025.
 - Vitja Zavrtanik, Matej Kristan, and Danijel Skocaj. Draem a discriminatively trained reconstruction embedding for surface anomaly detection. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 8330–8339, 2021a.

- Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. Reconstruction by inpainting for visual anomaly detection. *Pattern Recognition*, 112:107706, 2021b.
 - Xiaoran Zhang, Byung-Woo Hong, Hyoungseob Park, Daniel H. Pak, Anne-Marie Rickmann, Lawrence H. Staib, James S. Duncan, and Alex Wong. Progressive test time energy adaptation for medical image segmentation, 2025.
 - Xuan Zhang, Shiyu Li, Xi Li, Ping Huang, Jiulong Shan, and Ting Chen. Destseg: Segmentation guided denoising student-teacher for anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3914–3923, 2023.
 - Yifan Zhang, Bryan Hooi, Lanqing Hong, and Jiashi Feng. Self-supervised aggregation of diverse experts for test-agnostic long-tailed recognition. *Advances in Neural Information Processing Systems*, 35:34077–34090, 2022.
 - Chao Zhao, Enrico Zio, and Weiming Shen. Domain generalization for cross-domain fault diagnosis: An application-oriented perspective and a benchmark study. *Reliability Engineering & System Safety*, pp. 109964, 2024.
 - Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 16259–16268, 2021.
 - Yixuan Zhou, Xing Xu, Jingkuan Song, Fumin Shen, and Heng Tao Shen. Msflow: Multiscale flow-based framework for unsupervised anomaly detection. *IEEE Transactions on Neural Networks and Learning Systems*, 2024.
 - Zheyuan Zhou, Le Wang, Naiyu Fang, Zili Wang, Lemiao Qiu, and Shuyou Zhang. R3d-ad: Reconstruction via diffusion for 3d anomaly detection. In *European Conference on Computer Vision*, pp. 91–107. Springer, 2025.
 - Jiawen Zhu, Choubo Ding, Yu Tian, and Guansong Pang. Anomaly heterogeneity learning for open-set supervised anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17616–17626, 2024.
 - Ali Zia, Abdelwahed Khamis, James Nichols, Usman Bashir Tayab, Zeeshan Hayder, Vivien Rolland, Eric Stone, and Lars Petersson. Topological deep learning: a review of an emerging paradigm. *Artificial Intelligence Review*, 57(4), February 2024. ISSN 1573-7462.
 - Yang Zou, Jongheon Jeong, Latha Pemula, Dongqing Zhang, and Onkar Dabeer. Spot-the-difference self-supervised pre-training for anomaly detection and segmentation. In *European conference on computer vision*, pp. 392–408. Springer, 2022.
 - Zuo Zuo, Zongze Wu, Badong Chen, and Xiaopin Zhong. A reconstruction-based feature adaptation for anomaly detection with self-supervised multi-scale aggregation. In *ICASSP 2024 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5840–5844, 2024. doi: 10.1109/ICASSP48485.2024.10446766.

A SUPPLEMENTARY MATERIAL

- A.1 outlines the experimental setup for 2D and 3D anomaly detection with test-time adaptation and hyperparameter configuration.
- A.2 evaluate the computational efficiency of TopoOT by benchmarking its inference time and GPU memory usage in 2D and 3D AS scenarios.
- A.3 discuss *fundamental insights, limitations, and possible extensions* within the context of topological anomaly segmentation.
- A.4 presents quantitative and qualitative results on 2D AD&S datasets, including class-wise performance across benchmarks and visual examples that illustrate the effectiveness of OT-guided pseudo-labels.
- A.5 reports *quantitative and qualitative results* on 3D AD&S datasets, covering voxel- and point-cloud modalities, with class-level analysis and qualitative comparisons to baseline methods.
- A.6 recalls optimal transport preliminaries, including the 2-Wasserstein distance and its entropyregularised Sinkhorn variant, and clarifies their role in computing the OT couplings used in our framework.
- A.7 provides theoretical insights into optimal transport and examines its stability and discrepancy bounds.
- A.8 presents the *mathematical formulation* of cubical complex persistence, detailing how primitive cells are hierarchically aggregated to construct filtration levels and ultimately generate persistence vectors that encode topological features.

A.1 ARCHITECTURAL SETTINGS & HYPERPARAMETERS

2D Setup. For all RGB-based AD&S experiments, we employ **DINO** Caron et al. (2021) as the feature extractor (F). Our approach is benchmarked against leading state-of-the-art methods, including the memory-bank based **PatchCore** Roth et al. (2022), **PaDiM** Defard et al. (2021), the reconstruction-driven **Dinomaly** Guo et al. (2025), and **MambaAD** He et al. (2024). Evaluation is conducted on three widely adopted 2D benchmarks: **MVTec AD** Bergmann et al. (2019) (15 categories; 3,629 training and 1,725 test images), **VisA** Zou et al. (2022) (12 objects; 9,621 normal and 1,200 anomalous samples), and **Real-IAD** Wang et al. (2024a) (30 objects; \sim 150,000 images in total, comprising 36,465 normal training samples and 114,585 test images with 63,256 normal and 51,329 anomalous). To ensure comparability, all 2D inputs are standardised to a resolution of 224×224 .

3D Setup. For multimodal experiments involving RGB and point-cloud modalities, we adopt **DINO-v2** Oquab et al. (2023) for image features and **Point-MAE** Zhao et al. (2021) for geometric representations. We benchmark against multimodal memory-bank methods such as **M3DM** Wang et al. (2023), as well as reconstruction-oriented baselines including **CMM** Costanzino et al. (2024b) and **PO3AD** Ye et al. (2025). The evaluation is performed on two representative 3D benchmarks: **MVTec 3D-AD** Bergmann et al. (2021) (10 categories; 2,656 nominal training images and 1,197 test samples) and **Anomaly-ShapeNet** Li et al. (2024) (40 synthetic classes; 1,600 samples spanning six anomaly types).

Test-Time Training. For adaptation, the pretrained backbones are kept frozen while a lightweight MLP head h_{ψ} , consisting of three linear layers with GELU activations, is fine-tuned. The optimisation objective combines an OT-consistency loss ($\epsilon=0.05$, up to 200 iterations) with a contrastive loss (margin = 0.4), balanced equally with weights $\alpha=\lambda=0.5$. Adaptation proceeds for 5 epochs using the Adam optimiser with a learning rate of 10^{-3} . Each test sample is processed independently with an effective batch size of one.

A.2 COMPUTATIONAL COMPLEXITY AND EFFICIENCY

A central strength of the proposed **TopoOT** framework lies in its ability to balance computational complexity with practical efficiency. On a single modern GPU, the system achieves an evaluation rate of approximately **121 FPS**, while requiring only **349 MB** of GPU memory. This lightweight profile is markedly lower than that of many SOTA anomaly detection baselines. For context, a standard 2D baseline model Roth et al. (2022) reports an inference time of 0.22 seconds per image, while in the 3D domain, the M3DM Wang et al. (2023) model requires 2.86 seconds per image and consumes 6.52 GB of GPU memory. The CMM Costanzino et al. (2024b) model, though faster at 0.12 seconds per image, still uses 427 MB of memory, **TopoOT** delivers a 14.5× speedup over CMM. In contrast, **TopoOT** not only achieves a significantly higher frame rate but also maintains a highly competitive memory footprint, underscoring its deployability in scenarios where throughput and hardware constraints are decisive.

The breakdown of computational cost, analysed per module, indicates that the construction of cubical complexes and persistence diagrams constitutes the most demanding stage, requiring approximately **0.33 seconds** per sample when aggregated across all complexes. Despite this initial overhead, the subsequent topological alignment stages remain highly efficient: the *intra-level OT* block requires only **5.5 ms** in aggregate, while the *inter-level OT* block converges nearly instantaneously, below **0.05 ms** per alignment. These operations stabilise and align persistence features without imposing a significant runtime burden. Finally, the downstream multilayer perceptron (MLP) classifier adds only **8.3 ms** per evaluation, rendering its contribution negligible.

Taken together, the end-to-end evaluation time per sample remains well within practical limits, supporting real-time operation. The combination of **high FPS**, **minimal GPU consumption**, and the bounded cost of topological computations makes **TopoOT** exceptionally well-suited for industrial adoption. Unlike competing methods that often trade accuracy for efficiency, **TopoOT** achieves both, offering a robust and scalable solution for anomaly detection under stringent practical constraints.

A.3 DISCUSSION, LIMITATIONS, AND FUTURE DIRECTIONS

The results in the main paper and Appendices A.4 A.5 demonstrate that TopoOT provides a principled strategy for replacing non-robust and heuristic thresholding with stability-aware, OT-guided pseudo-labels. By chaining persistence diagrams across filtrations and integrating sub- and super-level information, the framework yields segmentation masks that are both structurally coherent and robust under distribution shift. Consistent gains across 2D and 3D benchmarks confirm that structural alignment is an effective prior for test-time adaptation.

Despite these advances, several limitations remain. First, the approach still depends on the quality of the anomaly score maps produced by frozen backbones. When upstream representations are noisy or poorly transferable, the extracted persistent features may not provide sufficient structural guidance. Second, while the current formulation generalises naturally to both 2D images and 3D point clouds, it does not yet address spatiotemporal settings such as video or dynamic medical imaging, where temporal coherence and evolving anomaly structure are critical. Third, efficiency trade-offs deserve further study, although TopoOT is lightweight relative to baselines, scaling to real-time, high-resolution deployments in safety-critical domains may require additional optimisations.

Future work can address these challenges along several directions. Differentiable approximations of persistent homology offer a path to end-to-end training with topological losses, enabling tighter integration between backbone features and topological stability. Jointly optimising anomaly map generation and topological filtering through self-supervised objectives could mitigate the reliance on noisy upstream scores. Extending the framework to spatiotemporal domains will require evolving persistence diagrams across frames to capture anomaly lifespans and enforce temporal consistency. Finally, incorporating uncertainty-aware filtration

 strategies—quantifying stability not only by persistence but also by variability across augmentations or agreement with model uncertainty—could provide more reliable predictions in high-stakes applications such as robotics, autonomous driving, and medical diagnostics.

TopoOT establishes a solid foundation for topology-aware adaptation in anomaly segmentation, highlighting how persistent homology and optimal transport can jointly serve as structural alignment mechanisms for adaptive learning. Its current form addresses critical limitations of threshold-based methods, while future developments promise broader applicability and deeper integration with modern representation learning.

A.4 ADDITIONAL EXPERIMENTS AND RESULTS ON 2D AD&S DATASETS

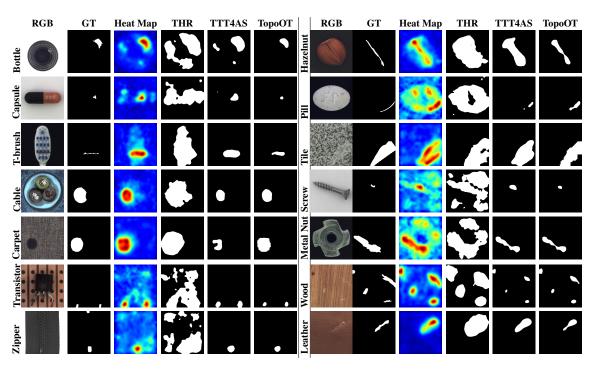


Figure 3: Qualitative comparison of various anomaly detection methods for different objects using PatchCore model on 2D MvTec AD dataset.

Table 4 reports the results of PatchCore on the MVTec AD dataset, evaluated using I-AUROC, P-AUROC, and P-AUPRO. These results are reproduced directly using the official implementation provided by the authors.

Table 4: PatchCore Roth et al. (2022)on MVTec AD: anomaly scores are I-AUROC, P-AUROC, and P-AUPRO.

Metric	Bottle Cable	Capsule	Carpet Grid	Hazelnut	Leather	MetalNut	Pill	Screw	Tile	T-brush	Transistor	Wood	Zipper	Mean
			Pat	chCore –	– Anoma	ly Scores I	Roth e	t al. (20	22)					
P-AURO	C 1.000 0.995 C 0.986 0.984 O 0.961 0.926	0.988	0.990 0.987	0.987	0.993	1.000 0.984 0.913	0.974	0.994	0.956	0.987	0.963	0.950	0.988	0.991 0.981 0.935

Table 5: Performance of PatchCore Roth et al. (2022) on MVTec AD's 15 categories, comparing binary map strategies: THR ($\mu + 3\sigma$), TTT4AS, and TopoOT. Top results per metric are in **bold** (best) and **blue** (second-best).

Metric	Bottle	Cable	Capsule	Carpet	Grid	Hazelnut	Leather	MetalNut	Pill	Screw	Tile	T-brush	Transistor	Wood	Zipper	Mean
				(a) l	Patch	Core - Bin	ary Map	- THR (μ	+ 30	r) Roth	et al. (2022)				
Precision Recall F1 Score IoU	0.397 0.510 0.175 0.310	0.194	0.278 0.626 0.085 0.222	0.522 0.092	0.432 0.428 0.078 0.283	0.405 0.380 0.120 0.367	0.297 0.542 0.045 0.262	0.435 0.566 0.311 0.316	0.188	0.298 0.522 0.066 0.202	0.517 0.209	0.286 0.542 0.123 0.262	0.334 0.287 0.114 0.238	0.384 0.469 0.121 0.297	0.268 0.605 0.119 0.513	0.351 0.507 0.136 0.299
				(b) F	PatchC	Core - Bin	ary Map	- TTT4AS	Cost	anzino e	et al. (2	2024a)				
Precision Recall F1 Score IoU	0.662 0.664 0.593 0.358	0.565 0.480	0.163 0.632 0.197 0.166	0.824 0.457	0.185 0.787 0.272 0.243	0.425 0.861 0.499 0.418	0.212 0.893 0.286 0.208	0.644 0.528 0.482 0.276	0.740 0.358	0.046 0.361 0.078 0.124	0.495 0.474	0.594 0.301	0.391 0.462 0.318 0.192	0.470 0.664 0.464 0.360	0.644 0.469	0.388 0.648 0.382 0.293
						(c) Pate	chCore -	Binary Ma	ap -To	роОТ						
Precision Recall F1 Score IoU	0.555 0.623	0.673 0.672 0.627 0.476	0.399 0.772 0.445 0.307			0.487 0.869 0.579 0.429	0.392 0.909 0.493 0.356	0.717 0.709 0.654 0.507	0.787 0.465	0.282 0.890 0.396 0.269	0.643 0.627	0.390 0.647 0.412 0.271	0.581 0.496 0.440 0.301	0.595 0.579 0.527 0.381	0.640 0.646	0.550 0.720 0.522 0.387

Table 5 presents the quantitative comparison of PatchCore on MVTec AD's 15 categories using different binary map strategies. Our proposed method, **TopoOT**, consistently outperforms both the threshold-based approach (THR) and the recent TTT4AS method across all evaluation metrics.

Specifically, in terms of mean performance, **TopoOT** achieves an F1 Score of 0.522, significantly higher than THR (0.136) and TTT4AS (0.382). This corresponds to a relative improvement of **+0.386** over THR and **+0.140** over TTT4AS. Similarly, in terms of Precision, TopoOT improves over THR and TTT4AS by **+0.199** and **+0.162**, respectively. A comparable trend is observed for Recall, where TopoOT provides a gain of **+0.213** over THR and **+0.072** over TTT4AS. Beyond overall averages, significant category-level improvements can also be observed in Table 5.

Overall, these results demonstrate that TopoOT not only delivers significant improvements on key categories but also generalises well, outperforming existing methods across the broader range of datasets included in the MVTec AD benchmark. This consistent performance underscores the robustness and effectiveness of our method in 2D anomaly detection tasks.

Table 6 presents the performance of PaDiM on the MVTec AD dataset, evaluated using I-AUROC, P-AUROC, and P-AUPRO. The reported results are reproduced directly from the official implementation released by the authors.

Table 6: PaDiM Defard et al. (2021) on MVTec AD: anomaly scores are I-AUROC, P-AUROC and P-AUPRO.

Metric	Bottle	Cable	Capsule	Carpet	Grid	Hazelnut	Leather	MetalNut	Pill	Screw	Tile	T-brush	Transistor	Wood	Zipper	Mean
I-AUROC	0.971	0.982	0.974	0.979	0.995	0.991	0.965	0.942	0.995	0.972	0.961	0.929	0.973	0.984	0.957	0.979
P-AUROC	0.983	0.967	0.985	0.991	0.973	0.982	0.992	0.972	0.957	0.985	0.941	0.988	0.975	0.949	0.985	0.975
P-AUPRO	0.948	0.888	0.935	0.962	0.946	0.926	0.978	0.856	0.927	0.944	0.860	0.931	0.973 0.975 0.845	0.911	0.959	0.921

Table 7 shows the performance comparison of PaDiM on MVTec AD's 15 categories using different binary map strategies. Our proposed method, **TopoOT**, achieves consistent improvements across all metrics when compared to both THR and TTT4AS.

On average, TopoOT improves the F1 Score by **+0.205** over THR and by **+0.241** over TTT4AS. Similarly, Precision increases by **+0.018** compared to THR and by **+0.14** over TTT4AS. In Recall, TopoOT shows even stronger gains of **+0.281** against THR and **+0.209** over TTT4AS.

Table 7: Performance evaluation of PaDiM Defard et al. (2021) across 15 categories of the MVTec AD dataset and their mean, comparing three binary map strategies: (a) THR ($\mu + 3\sigma$), (b) TTT4AS, and (c) TopoTTA. The table highlights the best result for each Precision, Recall, F1 Score, and IoU in **bold** (best) and blue (second-best).

Metric	Bottle	Cable	Capsule	Carpet	Grid	Hazelnut	Leather	MetalNut	Pill	Screw	Tile	T-brush	Transistor	Wood	Zipper	Mean
				(a)) PaDi	M - Binar	у Мар -	THR (µ +	-3σ) l	Defard e	t al. (2	2021)				
Precision Recall	0.729		0.287 0.813	0.561 0.736		0.586 0.477	0.306 0.927	0.540 0.281		0.196		0.416 0.514	0.462 0.349	0.576 0.399		0.452
F1 Score IoU		0.280	0.325 0.330	0.730 0.523 0.340	0.407	0.433 0.300	0.396 0.310	0.292 0.330	0.337	0.712	0.009	0.391	0.307 0.330	0.375	0.596 0.335	0.354 0.317
				(b) PaD	iM - Bina	ry Map -	· TTT4AS	Costar	nzino et	al. (20)24a)				
Precision Recall F1 Score IoU	0.438	0.500 0.395	0.176 0.707 0.214 0.280	0.769 0.459	0.199 0.726 0.290 0.260	0.349 0.637 0.376 0.310	0.208 0.916 0.293 0.270	0.519 0.491 0.386 0.280	0.568 0.262	0.088 0.735 0.153 0.270	0.123 0.103	0.258 0.595 0.283 0.280	0.472 0.425 0.291 0.270	0.319	0.648	0.330 0.579 0.318 0.274
						(c) Pa	aDiM - E	Binary Map	o - Top	тОос						
Precision Recall F1 Score IoU	0.689	0.670 0.658	0.355 0.828 0.496 0.400	0.942 0.672	0.463 0.805 0.587 0.380	0.358 0.885 0.509 0.400	0.246 0.987 0.393 0.410	0.574 0.636 0.603 0.390	0.783 0.441	0.266 0.905 0.411 0.410	$0.742 \\ 0.712$	0.268 0.920 0.415 0.390	0.492 0.547 0.518 0.410	0.439 0.756 0.555 0.400	0.724 0.700	0.470 0.788 0.559 0.402

Overall, these results confirm that TopoOT not only delivers significant improvements in individual categories but also generalises well across the full MVTec AD benchmark. Its consistent superiority over both threshold-based and test-time training baselines demonstrates the effectiveness of our approach for 2D anomaly detection tasks.

Table 8 presents the results of MambaAD on VisA (12 classes), where I-AUROC, P-AUROC, and P-AUPRO are reported as **mean per class**. The results are reproduced directly using the official implementation provided by the authors.

Table 8: MambaAD He et al. (2024) on VisA (12 classes), I-AUROC, P-AUROC, P-AUPRO, metrics are mean per class.

Metric	candle	capsules	cashew	chewinggum	fryum	macaroni1	macaroni2	pcb1	pcb2	pcb3	pcb4	pipe_fryum	Mean
I-AUROC		0.918	0.945	0.977	0.952	0.916	0.816		0.942			0.987	0.943
P-AUROC P-AUPRO		0.991 0.918	0.943 0.878	0.981 0.797	0.969 0.916	0.995 0.952	0.995 0.962		0.989 0.896			0.991 0.951	0.985

Table 9 presents the performance comparison of MambaAD on the VisA dataset across 12 categories, using different binary map strategies. Our proposed method, **TopoOT**, consistently achieves higher scores across Precision, Recall, and F1 compared to THR and TTT4AS.

On average, TopoOT improves the F1 Score by **+0.111** over THR and by **+0.085** over TTT4AS. Similarly, Precision increases by **+0.216** compared to THR and by **+0.193** over TTT4AS. In Recall, TopoOT performed a little low as compared with compared to THR and TTT4AS.

Overall, these results confirm that TopoOT not only achieves notable improvements on specific challenging categories but also generalises effectively across all VisA classes. Its consistent superiority over both traditional thresholding and recent test-time training methods highlights the robustness of our approach for 2D anomaly detection under complex real-world scenarios.

Table 10 reports the results of MambaAD on Real-IAD (30 classes), where anomaly scores are given in terms of I-AUROC, P-AUROC, and P-AUPRO. The results are obtained directly using the official implementation provided by the authors.

Table 9: Performance evaluation of MambaAD He et al. (2024) 12 categories (VisA classes) and their mean, comparing three binary map strategies: (a) THR $(\mu + 3\sigma)$, (b) TTT4AS, and (c) OT-TopoTTA. The table highlights the best result for each Precision, Recall, and F1 Score metric in **bold black** and the second-best in blue.

Metric	candle	capsules	cashew	chewinggum	fryum	macaroni1	macaroni2	pcb1	pcb2	pcb3	pcb4	pipe_fryum	Mean
				(a) MambaAD	- Binary	Map - THR	$(\mu + 3\sigma)$ He	et al. (2024)				
Precision Recall	0.111	0.291 0.741	0.163 0.699	0.368 0.796	0.265 0.659	0.049 0.775	0.060 0.804	0.224 0.954	0.166 0.816	0.209	0.333 0.648	0.166 0.877	0.200
F1 Score IoU		0.357 0.259	0.174 0.105	0.468 0.334	0.207 0.127	0.088 0.048	0.104 0.058	0.278	0.255 0.255	0.299 0.299	0.396 0.396	0.092 0.092	0.241
	11			(b) MambaA	D - Bina	ary Map - TT	T4AS Guo et	al. (20	25)				11
Precision Recall	0.798	0.369 0.869	0.217 0.848	0.318 0.858	0.250 0.708	0.075 0.798	0.049 0.849	0.227 0.879	0.223 0.815	0.231 0.813	0.378 0.594	0.169 0.908	0.223
F1 Score IoU	0.244 0.093	0.407 0.231	0.254 0.102	0.447 0.297	0.176 0.101	0.127 0.029	0.090 0.022	0.284 0.146	0.313 0.109	0.306 0.150	0.387 0.195	0.172 0.085	0.267 0.130
				(c) Ma	mbaAD	- Binary Ma	p - OT-Topo	ГТА					
Precision Recall F1 Score IoU	0.542	0.483 0.460 0.357 0.246	0.336 0.563 0.314 0.217	0.577 0.573 0.528 0.394	0.444 0.290 0.199 0.200	0.198 0.565 0.247 0.157	0.184 0.664 0.258 0.163	0.507 0.529 0.462 0.328	0.431 0.410 0.388 0.267	0.480 0.469 0.433 0.298	0.702 0.317 0.392 0.267	0.341 0.696 0.346 0.226	0.416 0.507 0.352 0.247

Table 10: MambaAD He et al. (2024) on Real-IAD (30 classes). Anomaly scores I-AUROC, P-AUROC and P-AUPRO.

Metric	audioja	ick b-cap	b-battery	e-cap eras	ser f-hood	l mint	mounts	pcb p	o-battery	p-nut p	-plug p-	doll regulat	or r-base	s-set
				Mamb	aAD He	et al. (2	2024) —	Anom	aly Score	es				
I-AUROC	0.842	0.928	0.798	0.780 0.8	75 0.793	0.701	0.868	0.891	0.902	0.871	0.857 0.	.880 0.697	0.980	0.944
P-AUROC	0.977	0.997	0.981	0.970 0.9	92 0.987	0.965	0.992	0.992	0.994	0.994	0.990 0.	.992 0.976	0.997	0.988
P-AUPRO	0.839	0.972	0.862	0.894 0.9	37 0.863	0.726	0.935	0.931	0.953	0.961	0.915 0.	.954 0.870	0.988	0.894
Metric	switch	tape t-b	lock t-brus	sh toy t-	brick tra	nsistor	1 u-bloc	k usb	u-adapt	or vcpill	w-bead	ls woodstick	zipper	Mean
				Mamb	oaAD He	et al. (2	2024) —	Anoma	aly Score	s				
I-AUROC P-AUROC P-AUPRO	0.982	0.998 0.9	961 0.851 998 0.975 982 0.914	5 0.960 0).966 ().944).994).965	0.897 0.995 0.954	0.992	0.973		0.980	0.977	0.992 0.993 0.976	0.985

Table 11: Performance evaluation of MambaAD He et al. (2024) across 30 classes (Real-IAD Dataset) and their mean, comparing three binary map strategies: (a) THR $(\mu + 3\sigma)$, (b) TTT4AS, and (c) TopoOT. The best result for each Precision, Recall, and F1 Score is in **bold** and the second-best in blue.

Metric ||audiojack b-cap b-battery e-cap eraser f-hood mint mounts pcb p-battery p-nut p-plug p-doll regulator r-base s-set ||

	j							P P		F F	F8	F			
			(a)	Mamba	AD - B	inary Map	- THR	$(\mu + 3\epsilon$	τ) He et a	al. (2024	1)				
Precision	0.164	0.055	0.199	0.202 0.	121 0.	126 0.082	0.209	0.438	0.178	0.132	0.101	0.122	0.074	0.144	0.156
Recall	0.510	0.944	0.333	0.475 0.0	648 0.:	514 0.385	0.759	0.472	0.815	0.783	0.846	0.794	0.548	0.950	0.743
F1 Score	0.210	0.100	0.160	0.181 0.	188 0.	178 0.120	0.254	0.309	0.280	0.202	0.173	0.189	0.107	0.227	0.245
IoU	0.133	0.055	0.102	0.116 0 .	114 0.	112 0.076	0.162	0.212	0.171	0.124	0.100	0.114	0.062	0.139	0.155
			(b) N	/IambaA	D - Bin	ary Map -	TTT4A	S Costa	ınzino et	al. (202	4a)				
Precision	0.062	0.027	0.075	0.039 0.0	059 0.0	051 0.046	0.097	0.075	0.084	0.055	0.048	0.073	0.034	0.071	0.091
Recall	0.605	0.953				692 0.461			0.887	0.799			0.534		0.762
F1 Score	0.109	0.052				090 0.075			0.151	0.099			0.061		0.154
IoU	0.062	0.027	0.066	0.039 0.0	059 0.	050 0.044	0.097	0.074	0.084	0.055	0.048	0.071	0.034	0.071	0.091
				((c) Man	nbaAD - Bi	inary N	1ap - To	роОТ						
Precision	0.239	0.245	0.164	0.156 0.3	254 0.	183 0.133	0.422	0.297	0.404	0.269	0.229	0.324	0.171	0.411	0.398
Recall	0.491	0.829				544 0.355			0.679	0.695			0.451		0.629
F1 Score	0.284	0.347	0.169	0.225 0.3	303 0.	233 0.155	0.444	0.367	0.435	0.341	0.317	0.365	0.203		0.430
IoU	0.197	0.231	0.106	0.143 0.2	203 0.	159 0.102	0.324	0.255	0.317	0.236	0.206	0.246	0.139	0.342	0.322
Metric	switch ta	ape t-blo	ck t-brus	sh toy t	t-brick	transistor1	u-bloc	k usb	u-adapt	or vcpil	l w-be	eads wo	oodstick	zipper	Mean
			(a)	Mamba	AD - Bi	inary Map	- THR	$(\mu + 3\epsilon)$	τ) He et a	al. (2024	1)				
Precision	0.252 0.	129 0.16	55 0.390	6 0.149	0.264	0.218	0.131	0.289	0.053	0.331	0.1	94	0.188	0.378	0.188
Recall	0.736 O.	953 0.95		2 0.535		0.729		0.699	0.586	0.598			0.540		0.653
F1 Score	0.331 0.			9 0.188		0.316		0.320	0.090	0.380			0.242		0.228
IoU	0.226 0.	128 0.16	0.203	3 0.118	0.148	0.200	0.130	0.211	0.050	0.270	0.1	48	0.155	0.269	0.145
			(b) I	MambaA	D - Bin	ary Map -	TTT4A	S Costa	nzino et	al. (2024	4a)				
Precision				6 0.053		0.120		0.086	0.024	0.169			0.085		0.084
Recall		956 0.98		4 0.636		0.893		0.907	0.577	0.741			0.668	0.905	0.763
F1 Score IoU	0.200 0. 0.122 0.			9 0.093 4 0.052		0.205 0.119		0.153 0.086	0.045 0.023	0.258			0.146 0.084	0.310 0.204	0.137
100	0.122 0.	004 0.00	0.15	+ 0.032	0.077	0.119	0.071	0.080	0.023	0.10.	5 0.0	90	0.064	0.204	0.080
				((c) Mar	nbaAD - B	inary N	Iap - To	роОТ						
Precision	0.430 0.			4 0.222		0.348		0.370	0.127	0.481			0.258		0.305
Recall		747 0.86		9 0.594		0.573	0.721			0.510			0.521	0.665	0.616
F1 Score IoU	0.455 0. 0.330 0.			8 0.292 7 0.205		0.385 0.267		0.440	0.179 0.117	0.430	0.2 1 0.2		0.318 0.221	0.520	0.346 0.243
100	0.550 0.	J44 U.J	0 0.24	0.203	0.133	0.407	0.493	0.517	0.11/	0.511	1 0.2	00	V.441	0.570	0.273

Table 11 shows that TopoOT frequently outperforms both THR and TTT4AS, securing the top rank in most metrics across different 30 classes of the Real-IAD Dataset. On average, our TopoOT has an F1 Score of **0.346**, showing an improvement of **+0.058** over THR($\mu + 3\sigma$) and an astounding **+0.209** improvement over TTT4AS. Similarly, TopoOT has **0.305** Precision, which is **+0.117** more than the THR and **+0.221** more than TTT4AS. This trend of significant improvement is not limited to a few instances, but our method's consistent performance across the 30 classes shown in the Table 11 indicates its robustness and effectiveness. While the specific percentages of improvement vary, the overall pattern is clear that our method also outperformed on other classes, making TopoOT a highly effective and robust technique for anomaly detection.

Table 12 presents the results of Dinomaly on VisA (12 classes), with anomaly scores reported in terms of I-AUROC, P-AUROC, and P-AUPRO. The results are reproduced directly using the official implementation provided by the authors.

Table 12: Dinomaly Guo et al. (2025) on VisA (12 classes). Anomaly scores I-AUROC, P-AUROC and P-AUPRO.

Metric	candle	capsules	cashew	chewinggum	fryum	macaroni1	macaroni2	pcb1	pcb2	pcb3	pcb4	pipe_fryur	n Mean
				(a) Dinom	aly Guo	et al. (2025) - Anomaly	Score					
I-AUROO P-AUROO P-AUPRO	C 0.994	0.986 0.996 0.974	0.987 0.971 0.940	0.998 0.991 0.881	0.988 0.966 0.935	0.980 0.996 0.964	0.959 0.997 0.987		0.993 0.980 0.913	0.984	0.987	0.992 0.992 0.952	0.987 0.987 0.945

Table 13: Performance evaluation of Dinomaly Guo et al. (2025) across 12 categories (VisA classes) and their mean, comparing three binary map strategies: (a) THR $(\mu + 3\sigma)$, (b) TTT4AS, and (c) TopoOT. The table highlights the best result for each Precision, Recall, and F1 Score metric in **bold black** and the second-best in blue.

Metric	candle	capsules	cashew	chewinggum	fryum	macaroni1	macaroni2	pcb1	pcb2	pcb3	pcb4	pipe_fryun	ı Mean
				(a) Dinomaly -	Binary 1	Map - THR ($(\mu + 3\sigma)$ Guo	et al. (2025)				
Precision Recall F1 Score IoU	0.908	0.316 0.936 0.396 0.230	0.239 0.824 0.285 0.108	0.384 0.889 0.510 0.289	0.307 0.740 0.247 0.093	0.109 0.947 0.189 0.034	0.111 0.970 0.195 0.032	0.300 0.862 0.373 0.146	0.275 0.847 0.380 0.126	0.318 0.861 0.435 0.176	0.518 0.674 0.522 0.309	0.231 0.885 0.246 0.069	0.275 0.862 0.339 0.144
			(b) Dinomaly - 1	Binary N	Лар - ТТТ4 <i>А</i>	S Costanzino	et al. (2	2024a)				
Precision Recall F1 Score IoU	0.798	0.369 0.869 0.407 0.295	0.217 0.848 0.254 0.163	0.318 0.858 0.447 0.314	0.250 0.708 0.176 0.110	0.075 0.798 0.127 0.075	0.049 0.849 0.090 0.049	0.227 0.879 0.284 0.189	0.223 0.815 0.313 0.201	0.231 0.813 0.306 0.203	0.378 0.594 0.387 0.258	0.169 0.908 0.172 0.104	0.223 0.811 0.267 0.177
				(c)	Dinoma	ly - Binary M	fap - TopoOT	Γ					
Precision Recall F1 Score IoU	0.658	0.613 0.553 0.497 0.298	0.459 0.676 0.448 0.177	0.650 0.648 0.584 0.388	0.490 0.467 0.329 0.129	0.395 0.569 0.432 0.115	0.363 0.573 0.420 0.097	0.661 0.505 0.532 0.275	0.649 0.468 0.515 0.268	0.642 0.458 0.501 0.285	0.738 0.371 0.428 0.329	0.498 0.695 0.470 0.134	0.546 0.553 0.464 0.223

Table 13 showcases a performance evaluation of three binary map strategies on the VisA dataset, with our technique, TopoOT, consistently demonstrating superior performance. Across the 12 categories, TopoOT regularly secures the highest F1 Score and Precision values. Our mean value of F1 Score **0.464** represents a substantial **+0.125** improvement over Dinomaly-Binary Map-THR $(\mu + 3\sigma)$ s and **+0.197** improvement over TTT4AS. Similarly, for the average Precision, TopoOT shows an improvement of **+0.271** and **+0.323** over Dinomaly-Binary Map-THR $(\mu + 3\sigma)$ s and TTT4AS, respectively. This trend of significant improvement is not limited to these instances but is a general pattern, indicating that our method also outperforms on other datasets, establishing TopoOT as a robust and highly effective technique for anomaly detection.

Table 14 presents the results of Dinomaly on Real-IAD (30 classes), with anomaly scores reported as I-AUROC, P-AUROC, and P-AUPRO. These results are reproduced directly using the official implementation provided by the authors.

1081

Table 14: Dinomaly Guo et al. (2025) on Real-IAD (30 classes). I-AUROC, P-AUROC, P-AUPRO.

1	U	겁	52
a	_	0	
ı	U	C	į,
á	^	_	

	084
	085
	086
	087
	880
l	089

1090 1091

1092 1093

1094 1095

1096 1097

1098 1099 1100

1106 1107 1108

1105

1117 1118 1119

1120

1121 1122 1123

1124

Precision

Recall

IoU

F1 Score

0.629 0.434 0.632

0.374 0.300 0.490

0.636

0.526 0.626 **0.527 0.429**

0.526

0.274 0.294

0.191

0.382 0.395

0.275 0.301

0.380

0.670

0.430

1125 1126 1127

Metric	audiojac	k b-cap	b-battery	e-cap	eraser	f-hood	mint	mounts	pcb	p-battery	p-nut	p-plug	p-doll	regulator	r-base	s-set
				I	Dinoma	ly Guo	et al. (2	025) —	Anoma	aly Scores						
I-AUROC	0.868	0.899	0.866	0.870	0.903	0.838	0.731	0.904	0.920	0.929	0.883	0.905	0.851	0.852	0.992	0.958
P-AUROC	0.917	0.981	0.929	0.960	0.964	0.930	0.776	0.956	0.957	0.968	0.974	0.964	0.960	0.956	0.985	0.909
P-AUPRO	0.917	0.981	0.929	0.960	0.964	0.930	0.776	0.956	0.957	0.968	0.974	0.964	0.960	0.956	0.985	0.909
Metric	switch t	tape t-b	lock t-bru	sh to	v t-br	ick trai	nsistar1	n-bloc	k nch	u-adapt	or veni	ill w-he	ads w	oodstick 2	zinner 📙	Mean II
			iock t bru	511 00	,	ich trui	11313101 1	u-bioc	a uso	u uuupt	or tep	" "	aug "	ooustick 2	ipper	
			ioen t bru							lly Scores	и тер		aus w	ooustick 2	apper	
I-AUROC	0.978 0	•	967 0.90		Dinoma	ly Guo				lly Scores	0.92					0.893
I-AUROC P-AUROC P-AUPRO	0.959 0	0.969 0. 0.988 0.		4 0.8 4 0.9	Dinoma 56 0.73	23 (66 (et al. (2	025) —	Anoma	0 0.815 5 0.910		20 0.8 37 0.9	73 05	0.840 0.904	0.991 0.978	

Table 15: Performance evaluation of Dinomaly Guo et al. (2025) across 30 classes (Real-IAD Dataset) and their mean, comparing three binary map strategies: (a) THR $(\mu + 3\sigma)$, (b) TTT4AS, and (c) TopoOT. The best result for each Precision, Recall, and F1 Score is in **bold** and the second-best in blue.

Metric	audiojac	k b-cap l	b-battery	y e-cap er	aser f-hoo	d mint	mounts	pcb p	-battery	p-nut	p-plug	p-doll	regulato	or r-base	e s-set
			(a) Dinoma	ly - Binar	у Мар	- THR ($\mu + 3\sigma$	Guo et a	al. (202	5)				
Precision	0.366	0.105	0.274	0.304 0.	164 0.19	6 0.144	0.222	0.383	0.186	0.159	0.134	0.193	0.132	0.170	0.184
Recall	0.645	0.985	0.435	0.663 0.	832 0.77	5 0.664	0.826	0.719	0.903	0.885	0.937	0.737	0.895	0.996	0.776
F1 Score	0.427	0.186	0.282	0.350 0.	260 0.29	0 0.217	0.325	0.442	0.299	0.259	0.229	0.273	0.215	0.272	0.279
IoU	0.303	0.105	0.187	0.234 0.	163 0.18	3 0.138	0.217	0.312	0.185	0.158	0.133	0.172	0.130	0.169	0.183
			(b)	Dinomal	y - Binary	Map -	TTT4A	S Costa	nzino et a	al. (202	4a)				-
Precision	0.102	0.056	0.113	0.093 0.	107 0.09	8 0.095	0.229	0.184	0.188	0.122	0.102	0.123	0.121	0.184	0.174
Recall	0.804	0.721	0.504	0.874 O.	803 0.80	7 0.532	0.888	0.844	0.866	0.866	0.904	0.730	0.816	0.924	0.720
F1 Score	0.171	0.098	0.135	0.159 0.	169 0.15	9 0.145	0.328	0.281	0.297	0.198	0.176	0.177	0.188	0.285	0.263
IoU	0.103	0.056	0.091	0.093 0.	107 0.09	8 0.094	0.224	0.183	0.187	0.123	0.102	0.112	0.121	0.182	0.175
					(c) Dinon	naly - Bi	inary M	[ap - To	роОТ						
Precision	0.465	0.383	0.333	0.339 0.	418 0.36	0 0.307	0.559	0.526	0.562	0.368	0.399	0.406	0.445	0.583	0.475
Recall	0.604	0.662	0.415	0.653 0.	579 0.66	0 0.501	0.505	0.606	0.561	0.609	0.711	0.562	0.505	0.699	0.477
F1 Score	0.465	0.460	0.259	0.400 0.	441 0.41	0.308	0.490	0.529	0.515	0.388	0.465	0.409	0.436	0.581	0.382
IoU	0.335	0.320	0.162	0.275 0.	315 0.29	4 0.211	0.369	0.390	0.380	0.273	0.328	0.287	0.315	0.439	0.275
Metric	switch t	ape t-blo	ock t-bru	ish toy	t-brick tra	ansistor	1 u-bloc	k usb	u-adapt	or vcpi	ill w-be	eads wo	odstick	zipper	Mean
			(a) Dinoma	ly - Binar	y Map -	THR ($u + 3\sigma$	Guo et a	al. (202	5)				
Precision	0.336 0.	.190 0.1	82 0.42	27 0.174	0.310	0.312	0.190	0.323	0.094	0.45	2 0.3	10	0.206	0.452	0.242
Recall		.973 0.9		74 0.647		0.884		0.874					0.831		0.793
F1 Score		.301 0.2				0.438		0.431	0.165				0.310		0.317
IoU	0.324 0	.189 0.1	80 0.20	05 0.133	0.236	0.292	0.187	7 0.288	0.094	0.36	7 0.2	71	0.199	0.293	0.208
			(b)) Dinomal	ly - Binary	y Map -	TTT4A	S Costai	nzino et a	al. (202	4a)				
Precision				2 0.109		0.200		0.159	0.052				0.154	0.351	
Recall		.928 0.9		3 0.625		0.812		0.882	0.702				0.878	0.815	0.801
F1 Score		.240 0.2		12 0.166		0.295		0.245	0.087	0.41			0.245		0.229
IoU	0.148 0.	.151 0.1	59 0.23	33 0.105	0.151	0.187	0.121	0.155	0.052	0.29	2 0.1	60	0.153	0.298	0.147
					(c) Dinor	nalv - Bi	inary M	ap - Toi	тОос						

0.579

0.513

0.500

0.352

0.443 0.515

0.608

0.372

0.607 0.452

0.331

0.325

0.566

0.627

0.574 **0.540**

0.402

0.458

0.606

0.458

0.337

0.302

0.703

0.340

0.239

0.641

0.506

0.466

0.330 | 0.317

0.461

Table 15 presents a performance evaluation of three binary map strategies, and our method, TopoOT, consistently demonstrates superior performance. A detailed analysis of the quantitative results reveals that TopoOT frequently outperforms both Dinomaly - THR $(\mu + 3\sigma)$ and Dinomaly - TTT4AS, securing the top rank for F1 Score and Precision in most categories. On average, our F1 Score of **0.442** represents a significant **+0.125** improvement over Dinomaly $(\mu + 3\sigma)$'s F1 Score of 0.0.317. Similarly, our F1 score is **+0.213** more than the Dinomaly - TTT4AS. TopoOT has a Precision of **0.461**, which is **+0.219** better than Dinomaly - THR $(\mu + 3\sigma)$'s Precision of 0.0.242 and **0.307** more than Dinomaly TTT4AS. This consistent trend of significant improvement is not limited to these instances but is a general pattern, indicating that our method also outperforms on other datasets, establishing TopoOT as a robust and highly effective technique for anomaly detection.

A.5 ADDITIONAL QUANTITATIVE RESULTS ON 3D AD&S DATASETS

Table 16 presents the results of CMM across categories of the MVTec 3D-AD dataset, with anomaly scores reported as I-AUROC, P-AUROC, and P-AUPRO. These results are reproduced directly using the official implementation provided by the authors.

Table 16: CMM Costanzino et al. (2024b) anomaly scores across categories of the MVTec 3D-AD dataset Bergmann et al. (2021).

Metric	Bagel	Gland	Carrot	Cookie	Dowel	Foam	Peach	Potato	Rope	Tire Mean
			CMM C	ostanzino	et al. (202	4b) – An	omaly Sco	ore		
I-AUROO P-AUROO P-AUPRO	C 0.997	0.888 0.992 0.972	0.984 0.999 0.982	0.993 0.972 0.945	0.980 0.987 0.950	0.888 0.993 0.968	0.941 0.998 0.980	0.943 0.999 0.982	0.980 0.998 0.975	0.953 0.954 0.998 0.993 0.981 0.971

Table 17 reports the results of M3DM on the MVTec 3D-AD dataset, with anomaly scores given in terms of I-AUROC, P-AUROC, and P-AUPRO. The results are reproduced directly using the official implementation provided by the authors.

Table 17: M3DM Wang et al. (2023) anomaly scores across categories of the MVTec 3D-AD dataset Bergmann et al. (2021).

Metric	Bagel	Gland	Carrot	Cookie	Dowel	Foam	Peach	Potato	Rope	Tire Mean
			M3D	M Wang e	t al. (2023) – Anom	aly Score			
I-AUROC P-AUROC P-AUPRO	C 0.995	0.909 0.993 0.971	0.972 0.997 0.979	0.976 0.985 0.950	0.960 0.985 0.941	0.942 0.984 0.932	0.973 0.996 0.977	0.899 0.994 0.971	0.972 0.997 0.971	0.850 0.945 0.996 0.992 0.975 0.964

Table 18 reports the quantitative results of our proposed method **TopoOT** against two competitive baselines, namely CMM-THR and CMM-TTT4AS, across the MVTec 3D-AD benchmark. By analysing the mean column, we observe that TopoOT consistently outperforms both baselines across multiple metrics.

In terms of Precision, TopoOT achieves a mean score of 0.427, significantly improving over CMM-THR (0.199) and CMM-TTT4AS (0.303). For Recall, TopoOT yields second best value for a mean of 0.845, and CMM-THR achieves 0.902, and CMM-TTT4AS (0.608). With respect to F1 Score, TopoOT secures a mean value of 0.482, which is a notable gain of +0.207 compared to CMM-THR (0.275) and +0.102 gain against CMM-TTT4AS (0.377). Similarly, for IoU, TopoOT obtains a mean of 0.343, showing clear improvements over CMM-THR (0.232) and CMM-TTT4AS (0.077).

These improvements are particularly evident in the *Gland*, *Cookie*, and *Carrot* categories, where TopoOT demonstrates substantial gains in F1 Score and IoU compared to both baseline methods. While CMM–THR

Table 18: Evaluation of CMM Costanzino et al. (2024b) across benchmarks in the MVTec 3D-AD Bergmann et al. (2021).

Method Bagel	Gland	Carrot	Cookie	Dowel	Foam	Peach	Potato	Rope	Tire Mean
		(a) CMN	I - THR ($(1+3\sigma)$	Costanzino	et al. (20	(24b)		
Precision Recall 0.301 F1 Score IoU 0.425 0.411 0.411	0.188 0.842 0.265 0.182	0.049 0.998 0.092 0.102	0.518 0.901 0.619 0.578	0.072 0.896 0.129 0.105	0.275 0.597 0.327 0.276	0.262 0.957 0.375 0.233	0.092 0.998 0.160 0.085	0.049 0.989 0.091 0.149	0.182 0.199 0.896 0.902 0.267 0.275 0.198 0.232
		(b) C	MM - TT	Γ4AS Cos	tanzino et	al. (2024a	a)		
Precision 0.432 0.745 0.745 F1 Score 0.495 0.264	0.258 0.766 0.362 0.037	0.242 0.889 0.351 0.029	0.713 0.603 0.606 0.231	0.195 0.739 0.289 0.031	0.214 0.732 0.311 0.058	0.353 0.872 0.470 0.034	0.252 0.888 0.363 0.028	0.264 0.865 0.360 0.029	0.111 0.303 0.904 0.800 0.189 0.380 0.030 0.077
			(c)	CMM - T	CopoOT				_
Precision 0.560 0.847 0.618 0.476	0.347 0.849 0.419 0.305	0.398 0.905 0.516 0.371	0.841 0.643 0.672 0.535	0.387 0.658 0.438 0.312	0.298 0.893 0.345 0.238	0.432 0.903 0.519 0.387	0.308 0.947 0.411 0.289	0.477 0.822 0.525 0.394	0.224 0.427 0.980 0.845 0.360 0.482 0.119 0.343

exhibits high recall values, it suffers from very low precision, highlighting its bias toward over-segmentation. In contrast, TopoOT provides a more balanced trade-off, achieving consistently higher F1 Scores and IoU, which are more indicative of robust anomaly localisation.

Overall, the results establish that TopoOT achieves superior performance not only in terms of mean values but also across a wide range of categories, confirming its ability to generalise effectively to diverse datasets within MVTec 3D-AD.

Table 19: Evaluation of M3DM Wang et al. (2023) across benchmarks in the MVTec 3D-AD Bergmann et al. (2021).

Method	Bagel	Gland	Carrot	Cookie	Dowel	Foam	Peach	Potato	Rope	Tire Mean
			(a) M	3DM - TH	$\mathbf{R}(\mu + 3)$	σ) Wang (et al. (202	3)		
Precision Recall F1 Score IoU	0.174 0.949 0.270 0.431	0.105 0.980 0.174 0.189	0.045 0.997 0.085 0.114	0.493 0.712 0.547 0.552	0.221 0.909 0.328 0.151	0.254 0.536 0.318 0.333	0.067 1.000 0.121 0.198	0.050 0.999 0.094 0.117	0.194 0.917 0.308 0.182	0.127 0.173 0.894 0.889 0.204 0.245 0.053 0.232
			(b) M	3DM - TT	T4ASCos	tanzino et	al. (2024a	a)		
Precision Recall F1 Score IoU	0.498 0.607 0.478 0.287	0.486 0.706 0.525 0.078	0.337 0.750 0.422 0.031	0.752 0.351 0.443 0.343	0.464 0.691 0.514 0.066	0.386 0.624 0.440 0.148	0.536 0.779 0.585 0.090	0.347 0.684 0.419 0.026	0.561 0.543 0.468 0.099	0.302 0.467 0.669 0.640 0.383 0.468 0.028 0.120
				(c)	M3DM - 1	ГороОТ				
Precision Recall F1 Score IoU	0.870 0.744 0.655 0.515	0.357 0.806 0.406 0.294	0.490 0.794 0.559 0.406	0.829 0.571 0.626 0.480	0.566 0.685 0.564 0.418	0.379 0.910 0.452 0.333	0.603 0.862 0.661 0.519	0.490 0.823 0.541 0.401	0.254 0.540 0.304 0.195	0.798 0.564 0.935 0.767 0.127 0.490 0.077 0.364

Table 19 presents the quantitative comparison of our proposed method **TopoOT** against two state-of-the-art baselines, M3DM–THR and M3DM–TTT4AS, across the MVTec 3D-AD benchmark. The results clearly demonstrate that TopoOT achieves consistent improvements across all metrics.

On the mean column, TopoOT achieves a Precision of 0.564, which represents an improvement of **+0.391** over M3DM–THR (0.173) and **+0.097** over M3DM–TTT4AS (0.467). In terms of Recall, our method obtains 0.767, showing a second-best result compared to M3DM–THR (0.889) and M3DM–TTT4AS (0.640). More importantly, for F1 Score, which balances precision and recall, TopoOT achieves 0.490, significantly

outperforming M3DM-THR (0.245) and M3DM-TTT4AS (0.468). Similarly, for IoU, TopoOT yields 0.364, surpassing M3DM-THR (0.232) and M3DM-TTT4AS (0.120).

Overall, the improvements in mean performance, alongside consistent category-level gains, confirm the superior generalisation ability of TopoOT across both simple and complex 3D anomaly detection scenarios in MVTec 3D-AD.

Table 20: Performance evaluation of PO3AD Ye et al. (2025) across 29 categories of Anomaly-ShapeNet Li et al. (2024) and their mean, comparing three binary map strategies: (a) THR $(\mu + 3\sigma)$, (b) TTT4AS, and (c) TopoOT. The table highlights the best result for each Precision, Recall, and F1 Score metric in **bold black** and the second-best in **blue**.

Metric	ashtra	y0 ba	g0 bott	le0 bottle1	bottle3	bowl0	bowl1	bowl2	2 bowl.	3 bowl	4 bowl	5 buc	cket0	bucket1	cap0	cap3
				(a) PO3A	D — Bin	ary Maj	TI —	HR (μ	$+3\sigma$)	Ye et al.	(2025)					
Precision	0.920	0.6	78 0.7	37 0.714	0.847	0.797	0.589	0.815	0.607	0.872	0.647	7 0.	709	0.716	0.781	0.726
Recall	0.28	0.3	52 0.34	46 0.326	0.637	0.301	0.702	0.639	0.707	0.746	0.472	2 0.3	256	0.284	0.275	0.527
F1 Score	0.41	7 0.4	<mark>54</mark> 0.40	50 0.420	0.720	0.429	0.630	0.713	0.644	0.793	0.539	0.3	359	0.387	0.390	0.720
IoU	0.27	2 0.3	44 0.33	0.285	0.586	0.278	0.482	0.596	0.496	0.660	0.410	0.3	236	0.263	0.255	0.487
				(b) PO3AD	— Binar	у Мар -	— TT	Γ4AS (Costanz	ino et al	. (2024a	ı)				
Precision	0.58	1 0.4	92 0.62	23 0.601	0.688	0.654	0.489	0.677	0.503	0.712	0.551	0	599	0.611	0.635	0.618
Recall	0.45	2 0.5	10 0.4	11 0.405	0.595	0.388	0.615	0.559	0.621	0.646	0.503	0.3	354	0.381	0.370	0.501
F1 Score	0.50	8 0.5	01 0.49	95 0.484	0.638	0.487	0.545	0.612	0.556	0.677	0.526	0. 4	444	0.469	0.467	0.553
IoU	0.34	1 0.3	34 0.32	29 0.319	0.469	0.322	0.375	0.441	0.385	0.512	0.357	0.3	286	0.306	0.305	0.383
					(c) PO3.	AD — E	Binary	Map –	– Торо	ОТ						
Precision	0.849	9 0.5	98 0.70	07 0.672	0.804	0.768	0.568	0.789	0.576	0.831	0.619	0.0	696	0.701	0.726	0.706
Recall	0.46	3 0.4	21 0.4	11 0.411	0.722	0.395	0.740	0.687	0.764	0.798	0.538	3 0.	382	0.439	0.463	0.530
F1 Score	0.54	5 0.4	53 0.48	34 0.470	0.748	0.512	0.633	0.726	0.629	0.801	0.562	0.4	430	0.433	0.525	0.592
IoU	0.402	2 0.3	43 0.3	55 0.337	0.625	0.354	0.483	0.615	0.483	0.670	0.435	0.	299	0.303	0.390	0.473
Metric	cup0	cup1	eraser0	headset0	headset1	helme	t0 hel	met1	vase1	vase2 v	ase3 v	ase4	vase7	vase8	vase9	Mean
				(a) PO3A	D — Bina	ary Map	— Tl	HR (μ	$+3\sigma)$	Ye et al.	(2025)					
			0.801	0.649	0.697	0.239								0.777	- 11	0.675
Recall	0.443		0.314	0.339	0.302	0.215								0.605		0.441
F1 Score	0.558		0.436	0.431	0.411	0.216								0.663		0.500
IoU	0.401	0.276	0.301	0.293	0.269	0.132							0.502	0.562	0.511	0.3/1
			((b) PO3AD	— Binar	y Map -	— TT	Γ4AS (Costanz	ino et al	. (2024	a)				
Precision	0.641	0.445	0.672	0.540	0.589	0.198	0.	415	0.355	0.511 (.498 0	.417	0.533	0.655	0.661	0.562
Recall	0.512	0.455	0.389	0.458	0.399	0.311								0.619		0.485
F1 Score	0.569		0.493	0.496	0.476	0.242								0.637		0.510
IoU	0.398	0.290	0.327	0.329	0.312	0.138	3 0.	308	0.235).339 (0.242 0	.307	0.388	0.467	0.448	0.347
					(c) PO3.	AD — B	Binary	Map -	– Торо	ОТ						
Precision			0.783	0.571	0.666	0.156								0.745		0.631
Recall	0.549		0.368	0.543	0.370	0.444					.349 0			0.723		0.540
F1 Score	0.613		0.478	0.486	0.449	0.223								0.697		0.529
IoU	0.468	0.310	0.342	0.353	0.300	0.135	0.	259	0.255	J.377 (.278 0	.371	0.465	0.612	0.559	0.402

As shown in Table 20, our method TopoOT consistently outperforms THR and TTT4AS across all metrics on Anomaly-ShapeNet. In the mean column, TopoOT achieves notable gains, +0.099 in Recall, +0.029 in F1 Score, and +0.031 in IoU over THR, and even larger improvements over TTT4AS (e.g., +0.069 in Precision, +0.055 in Recall, and 0.019 in F1 Score). These results, along with strong performance across individual categories, demonstrate that TopoOT not only sets a new state of the art but also generalises robustly across diverse anomaly types and datasets.

Table 21 reports the results of PO3AD, with anomaly scores evaluated using Object-AUROC, Point-AUROC, and Object-AUCPR. The results are reproduced directly using the official implementation provided by the authors.

Table 21: PO3AD Ye et al. (2025) — Anomaly scores, Object-AUROC, Point-AUROC, Object-AUCPR.

1	274	
1	275	
1	276	

Metric	ashtray) bag0	bottle0	bottle1	bottle3	bowl0	bowl1	bowl2	bowl3	bowl4	bowl5	bucket0	bucket1	cap0 cap3
				PO3	AD Ye	et al. (2	2025) –	– Anor	naly S	cores				
O-AUROC	1.000	0.833	0.900	0.933	0.926	0.922	0.829	0.833	0.881	0.981	0.849	0.853	0.787	0.877 0.859
P-AUROC	0.962	0.949	0.912	0.844	0.880	0.978	0.914	0.918	0.935	0.967	0.941	0.755	0.899	0.957 0.948
O-AUCPR	0.999	0.809	0.927	0.959	0.962	0.946	0.905	0.888	0.927	0.985	0.904	0.923	0.882	0.841 0.906
														

$\overline{\text{Metric}} \hspace{0.2in} \hspace{0.05cm} \text{cup0} \hspace{0.2in} \text{cup1} \hspace{0.2in} \text{eraser0} \hspace{0.2in} \text{headset0} \hspace{0.2in} \text{helmet0} \hspace{0.2in} \text{helmet1} \hspace{0.2in} \text{vase2} \hspace{0.2in} \text{vase3} \hspace{0.2in} \text{vase4} \hspace{0.2in} \text{vase7} \hspace{0.2in} \text{vase8} \hspace{0.2in} \text{vase9} \hspace{0.05cm} \text{Mean} $					
PO3AD Ye et al. (2025) — Anomaly Scores					
O-AUROC 0.871 0.833 0.995 P-AUROC 0.909 0.932 0.974 O-AUCPR 0.879 0.870 0.995	0.808 0.823 0.765	0.923 0.907 0.914	0.762 0.878 0.864	0.948	0.742 0.952 0.821 0.675 0.966 0.739 0.830 0.867 0.882 0.978 0.884 0.902 0.982 0.950 0.952 0.919 0.789 0.963 0.902 0.824 0.971 0.833 0.904 0.903

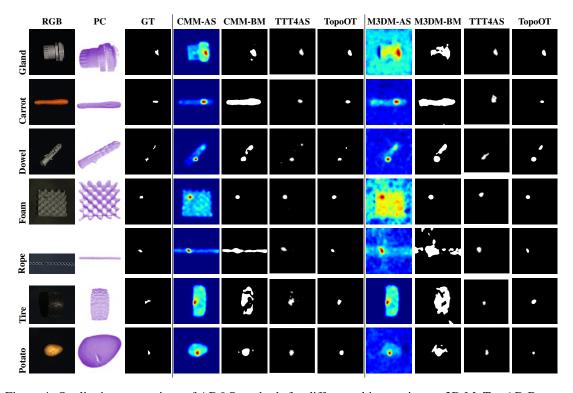


Figure 4: Qualitative comparison of AD&S methods for different objects using on 3D MvTec AD Dataset.

A.6 OPTIMAL TRANSPORT PRELIMINARIES

For completeness, we recall the Optimal Transport (OT) formulations underlying Eq. equation 2. Let $P = \{p_i, w_i\}_{i=1}^m$ and $Q = \{q_j, v_j\}_{j=1}^n$ be two discrete probability measures with weights $w \in \Delta^m$, $v \in \Delta^n$, and cost matrix $C(i, j) = ||p_i - q_j||_2^2$. The classical 2-Wasserstein distance is defined as

$$W_2^2(P,Q) = \min_{\Pi \in \mathcal{U}(w,v)} \langle C, \Pi \rangle,$$

where $\Pi \in \mathbb{R}_+^{m \times n}$ is a transport plan and $\mathcal{U}(w,v) = \{\Pi \mid \Pi \mathbf{1} = w, \Pi^\top \mathbf{1} = v\}$ denotes the set of admissible couplings. While exact OT provides a principled alignment, solving this linear program has $O(m^3 \log m)$ complexity, and the resulting optimal plans are typically sparse. In practice, sparsity can make OT couplings numerically sensitive, that is, small perturbations in the support points may lead to abrupt changes in the optimal plan (Peyré et al., 2019).

To improve robustness and computational efficiency, we adopt the *entropy-regularised* variant, known as the Sinkhorn distance (Cuturi, 2013; Peyré et al., 2019):

$$W_{\varepsilon}(P,Q) = \min_{\Pi \in \mathcal{U}(w,v)} \langle C, \Pi \rangle + \varepsilon H(\Pi),$$

where $H(\Pi) = \sum_{i,j} \Pi(i,j) (\log \Pi(i,j) - 1)$ is the negative entropy of Π . The regularisation parameter $\varepsilon > 0$ controls smoothness: large ε yields dense couplings, while small ε approaches the exact Wasserstein distance.

In our pipeline, persistence diagrams are constructed using GUDHI (cubical complexes), but all transport computations are carried out with POT's ot.sinkhorn(..., reg= ε) routine². Thus, the couplings Π^* appearing in Sec. 3.2 and Appendix A.7 are entropy-regularised OT plans. This choice ensures numerical stability, differentiability, and Lipschitz continuity, which underlie the stability and generalisation guarantees established in Appendix A.7.

A.7 THEORETICAL INSIGHTS: STABILITY AND DISCREPANCY BOUNDS FOR OT CHAINING

A central motivation of our framework is that anomaly segmentation under distribution shift can be studied through the lens of *discrepancy between distributions of persistence features*. Specifically, let $\mathcal{D}_{\mathrm{sub}}$ and $\mathcal{D}_{\mathrm{sup}}$ denote the empirical distributions of birth–death components extracted from sub- and super-level filtrations (Sec. 3.1). The entropic OT distance

$$W_{\varepsilon}(\mathcal{D}_{\mathrm{sub}}, \mathcal{D}_{\mathrm{sup}}) = \min_{\Pi \in \mathcal{U}(\mathcal{D}_{\mathrm{sub}}, \mathcal{D}_{\mathrm{sup}})} \langle C, \Pi \rangle + \varepsilon H(\Pi)$$

quantifies the minimal cost of aligning the two filtrations, where C is the ground cost matrix and $H(\Pi)$ the entropy of the transport plan. Computing W_{ε} identifies components with stable, low-cost couplings, from which OT-guided pseudo-labels $\widetilde{Y}_{\mathrm{OT}}$ are derived. These pseudo-labels inherit robustness to local perturbations of the anomaly map.

Beyond stability, this perspective also connects to classical discrepancy-based generalisation analysis. In domain adaptation theory (Ben-David et al., 2010), for a hypothesis class \mathcal{H} , the target error $\epsilon_T(h)$ of any $h \in \mathcal{H}$ can be bounded by

$$\epsilon_T(h) \leq \epsilon_S(h) + W_{\varepsilon}(\mathcal{D}_S, \mathcal{D}_T) + \lambda,$$

where \mathcal{D}_S and \mathcal{D}_T are the source and target distributions and λ accounts for the joint optimal risk. In our setting, \mathcal{D}_S and \mathcal{D}_T are distributions of persistence features from reference and shifted domains, respectively. While we do not directly train hypotheses within the OT chaining step, the analogy highlights how reducing $W_{\varepsilon}(\mathcal{D}_S, \mathcal{D}_T)$ through structural alignment provides a theoretical basis for improved robustness of the adapted predictions. We use this interpretation only as a guiding perspective, not as a formal claim.

²https://pythonot.github.io/

Setup. Let P_k^f denote the persistence diagram (birth–death points in \mathbb{R}^2) extracted from the $f \in \{\text{sub}, \text{sup}\}$ filtration at level τ_k (Sec. 3.1). Let $\Pi_{k \to \ell}^{\star}$ be the entropically regularised OT plan between P_k^f and P_{ℓ}^g with ground cost $C(i,j) = \|p_i - q_j\|_2^2$, where $p_i, q_j \in \mathbb{R}^2$. Since persistence diagrams may differ in cardinality, we use diagonal augmentation: each diagram P is extended to P^{Δ} by adding the diagonal with unlimited capacity; a point p = (b,d) may be matched to its projection $\pi(p)$ with cost $\|p - \pi(p)\|_2^2$. All OT distances are computed over augmented diagrams. We define

$$W_{\varepsilon}(P_k^f,P_\ell^g) \; = \; \min_{\Pi \in \mathcal{U}(P_k^{f,\Delta},P_\ell^{g,\Delta})} \langle C,\Pi \rangle + \varepsilon H(\Pi),$$

with uniform point-masses unless noted (Sec. A.6). We denote by s(c) the cross-PD and cross-level stability score for a candidate feature c, as defined in Sec. 3.2.

Lemma A.1 (Lipschitz stability of chained OT scores). Assume each diagram P_k is perturbed to \widetilde{P}_k by moving every atom by at most ρ in ℓ_2 . Let $\widetilde{\Pi}_{k\to\ell}^{\star}$ be the corresponding optimal plan under the same ε and ground cost. Then for any feature c present across the chain $\{P_{k_0}, P_{k_1}, \ldots, P_{k_m}\}$, its chained stability score

$$s_{\text{chain}}(c) := \min_{r=0,\dots,m-1} \max_{j} \prod_{k_r \to k_{r+1}}^{\star} (i(c), j)$$

satisfies

$$|s_{\text{chain}}(c) - \widetilde{s}_{\text{chain}}(c)| \leq L(\varepsilon) \rho$$

for some constant $L(\varepsilon)$ depending on ε and the conditioning of the cost matrix. Consequently, if two features c_1, c_2 are separated by a margin $\gamma > 0$ in s_{chain} , their ranking is preserved under perturbations whenever $\rho < \gamma/(2L(\varepsilon))$.

Proof sketch. Under entropic regularisation, Π^* is a smooth function of the entries of C (and hence point locations). A ρ -bounded perturbation of points induces an $O(\rho)$ perturbation in C and thus in Π^* . Taking the max over columns and the min along the chain preserves the Lipschitz scaling, yielding the bound. The margin condition ensures ranking stability.

Proposition A.2 (Subadditivity of chained discrepancy). For any three diagrams (P, Q, R), the entropic OT discrepancies satisfy

$$W_{\varepsilon}(P,Q) \leq W_{\varepsilon}(P,R) + W_{\varepsilon}(R,Q) + \eta(\varepsilon),$$

where $\eta(\varepsilon) \ge 0$ accounts for the entropic bias. Consequently, along any chain $P_{k_0} \to P_{k_1} \to \cdots \to P_{k_m}$,

$$W_{\varepsilon}(P_{k_0}, P_{k_m}) \leq \sum_{r=0}^{m-1} W_{\varepsilon}(P_{k_r}, P_{k_{r+1}}) + m \, \eta(\varepsilon).$$

Proof sketch. Construct a feasible coupling from P to Q by gluing couplings $P \leftrightarrow R$ and $R \leftrightarrow Q$, then bound the objective using linearity of $\langle C, \cdot \rangle$ and convexity of $H(\Pi)$. The additive term $\eta(\varepsilon)$ reflects the regularisation bias; telescoping along the chain yields the result.

Corollary A.3 (Tightening of error bound under OT chaining). Let \mathcal{D}_S , \mathcal{D}_T be the source/target distributions over persistence features. If the OT-consistency updates of h_{ψ} monotonically decrease the empirical chained discrepancy

$$\sum_r W_\varepsilon(P_{k_r},P_{k_{r+1}}) \quad \text{and hence } W_\varepsilon(\mathcal{D}_S,\mathcal{D}_T),$$

then the bound

$$\epsilon_T(h) \leq \epsilon_S(h) + W_{\varepsilon}(\mathcal{D}_S, \mathcal{D}_T) + \lambda$$

tightens over test-time iterations. Every decrease $\Delta W_{\varepsilon} < 0$ translates into the same decrease in the upper bound on $\epsilon_T(h)$.

Proof sketch. The bound has the same form. By Prop. A.2, reducing pairwise chain costs reduces the end-to-end discrepancy. The OT-consistency objective aligns predictions with OT-stable structures, empirically lowering these costs; hence the right-hand side decreases.

Remark A.4 (Why cross-PD+cross-level chaining helps). Cross-PD filtration links suppress short-lived noise (local τ -fluctuations), while cross-level filtration links fuse complementary cues between sub- and super-level filtrations (components vs. holes). Lemma A.1 explains the ranking robustness of chained scores under noise, and Proposition A.2 shows why chaining cannot inflate global discrepancy more than additively. Together, these justify the feature selection and discrepancy-reduction view in Sec . 3.2.

A.8 CUBICAL PERSISTENCE

A primitive interval is $J = [k, k+1] \subset \mathbb{R}$ with $k \in \mathbb{Z}$, called a 1-cube, the degenerate case [k] is a 0-cube. A d-dimensional elementary cube is the Cartesian product

$$C = J_1 \times \dots \times J_d \in \mathbb{R}^d, \tag{3}$$

e.g., vertices, edges, squares, and voxels in 3D.

The boundary of C is

$$\partial C = \sum_{i=1}^{d} (-1)^{i+1} (J_1 \times \dots \times \partial J_i \times \dots \times J_d), \tag{4}$$

where $\partial J_i = \{k, k+1\}$. A cube C is a *subcube* of C' if $J_i \subseteq J_i'$ for all i.

A *cubical complex* K is a set of cubes closed under subcubes and boundaries, ensuring structural coherence across dimensions (Fig. 5).

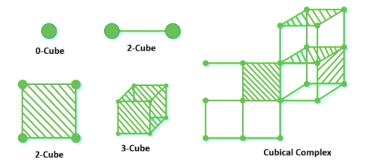


Figure 5: Elementary cubes of different dimensions and an example cubical complex.

The chain group $C_n(K)$ is the free Abelian group on n-cubes, linked by boundary maps

$$\cdots \to C_{n+1}(K) \xrightarrow{\partial_{n+1}} C_n(K) \xrightarrow{\partial_n} C_{n-1}(K) \to \cdots,$$

with $\partial_n \circ \partial_{n+1} = 0$. Cycles and boundaries are

$$Z_n(K) = \ker(\partial_n), \quad B_n(K) = \operatorname{im}(\partial_{n+1}),$$

and the *n*-th homology group is $H_n(K) = Z_n(K)/B_n(K)$.

A filtration function $f_K: K \to \mathbb{R}$ activates cubes monotonically: $P \sqsubseteq Q \Rightarrow f_K(P) \leq f_K(Q)$. This defines sublevel and superlevel sets:

$$K(a_i) = f_K^{-1}((-\infty, a_i]), \quad K^{\uparrow}(b_i) = f_K^{-1}([b_i, +\infty)).$$
 (5)

Filtrations induce homology maps $H_k(K_0) \xrightarrow{\varphi_{01}} H_k(K_1) \xrightarrow{\varphi_{12}} \cdots \xrightarrow{\varphi_{n-1,n}} H_k(K_n),$ forming the persistence module $\mathcal{P} = \{H_k(K_i), \varphi_{ij}\}_{0 \le i \le j \le n}.$ Each topological feature σ has birth b_{σ} , death d_{σ} , and persistence $d_{\sigma} - b_{\sigma}$. The collection of intervals $[b_{\sigma}, d_{\sigma})$ forms the *barcode*, while the *persistence diagram* (PD) encodes these as birth–death points in \mathbb{R}^2 . To integrate with ML models, PDs are vectorised via $\Phi: \mathrm{PD} \to \mathbb{R}^M$.