Multivariate Time Series Anomaly Detection with Idempotent Reconstruction

Xin Sun, Heng Zhou, Chao Li

Zhejiang University

NeurIPS 2025, San Diego

1. Problems that reconstruction-based methods are suffering

Reconstruction-based methods are a competitive choice for MTS Anomaly Detection, but they face two critical challenges:

(1) Over Generalization

(2) Balance between Robustness and Sensitivity

Challenge 1: Over Generalization

Models with excessive decoding power reconstruct abnormal inputs just as well as normal ones, making them indistinguishable.

Challenge 2: Robustness vs. Sensitivity

Models must be robust to real-world noise but also sensitive to subtle, early-stage anomalies. These goals are often in conflict.

1. Problems that reconstruction-based methods are suffering

Reconstruction-based methods are a competitive choice for MTS Anomaly Detection, but they face two critical challenges:

(1) Over Generalization

(2) Balance between Robustness and Sensitivity

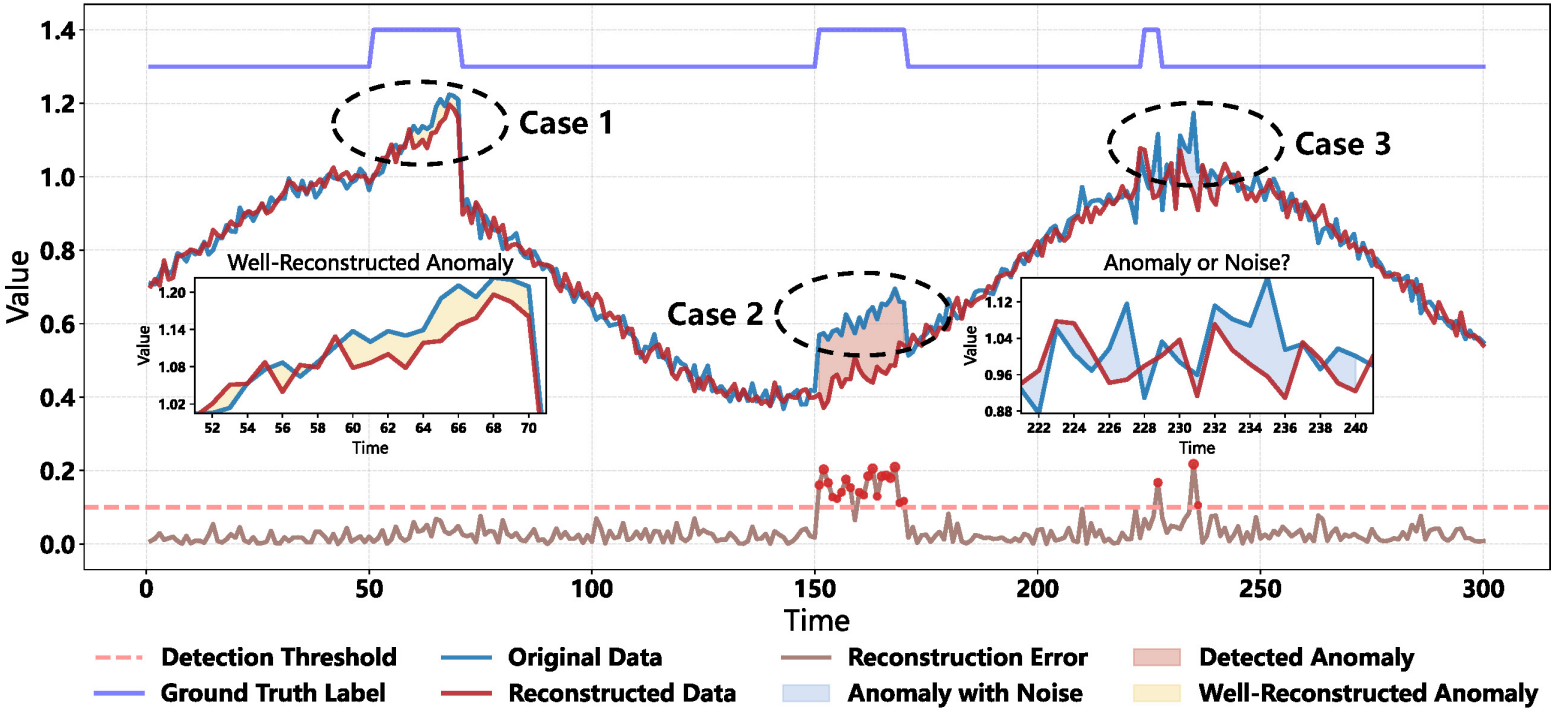

Case 1: Over Generalization

Case 2: Correct Detection

Case 3: Robustness vs. Sensitivity

2. Proposed solution from a Manifold Perspective

We propose Idempotent Generation for Anomaly Detection (IGAD), a flexible, parameter-free module inspired by idempotent generative network.

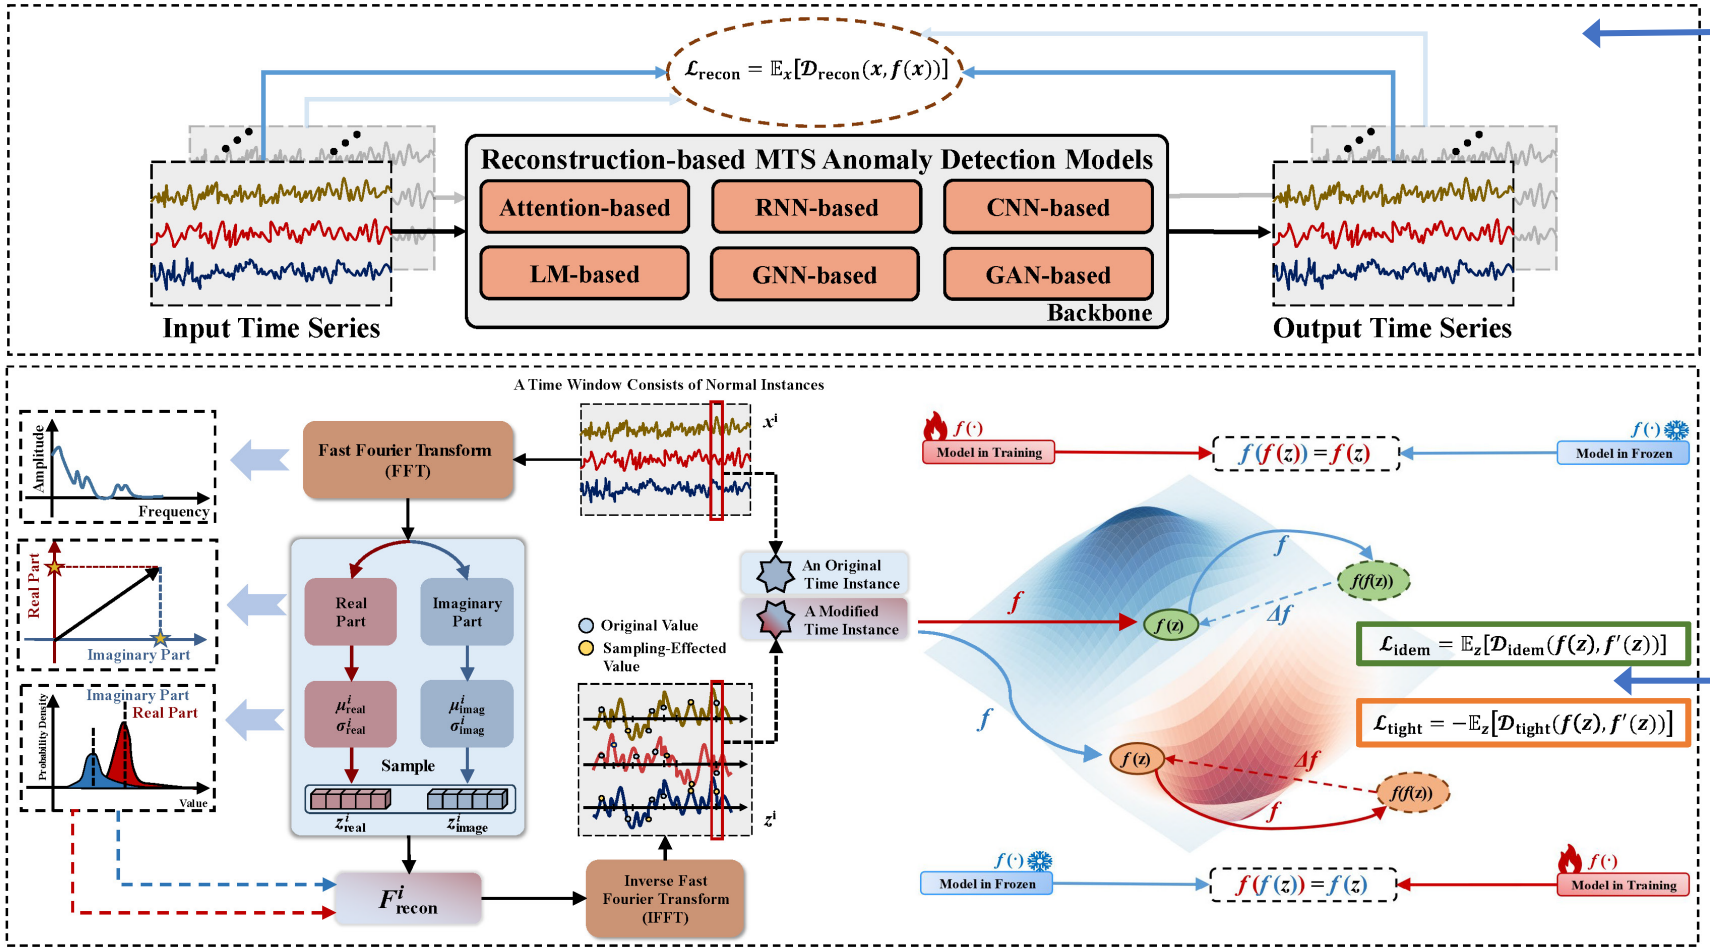

Reconstruction Loss

The Concept

1. Define a target manifold $\mathcal{M}_{\text{target}}$ that only contains normal instances.
2. Modify the manifold to ensure normal points, even noisy ones, map onto it.

Idempotent Loss & Tightness Loss

3. Optimization Objectives

Three objectives are included and optimized jointly.

(1) Reconstruction Objective

$$\mathcal{D}_{\text{recon}}\big(x, f(x)\big) = \frac{1}{m}\sum_{i=1}^{m} \| x^i - x^i_{\text{recon}} \|_2^2 \qquad \mathcal{L}_{\text{recon}}\big(x, f(x)\big) = \mathbb{E}_x\big[\mathcal{D}_{\text{recon}}\big(x, f(x)\big)\big]$$

(2) Idempotent Objective

$$\mathcal{D}_{\text{idem}}\Big(f(z), f'\big(f(z)\big)\Big) = \frac{1}{m}\sum_{i=1}^{m} \Big|f(z^i) - f'\Big(f\big(z^i\big)\Big)\Big| \qquad \mathcal{L}_{\text{idem}}\big(f(z), f'(z)\big) = \mathbb{E}_z\Big[\mathcal{D}_{\text{idem}}\Big(f(z), f'\big(f(z)\big)\Big)\Big]$$

(3) Tightness Objective

$$\mathcal{D}_{\text{tight}}\Big(f'(z), f\big(f'(z)\big)\Big) = \frac{1}{m}\sum_{i=1}^{m} \Big|f'(z^i) - f\Big(f'\big(z^i\big)\Big)\Big| \qquad \mathcal{L}_{\text{tight}}\big(f(z), f'(z)\big) = \mathbb{E}_z\Big[-\mathcal{D}_{\text{tight}}\Big(f'(z), f\big(f'(z)\big)\Big)\Big]$$

$$\mathcal{L}_{\text{idem}}\big(f(z), f'(z)\big) = -\mathcal{L}_{\text{tight}}\big(f'(z), f(z)\big)$$

Adversarial Training

4. Experiments

(1) VUS-PR on different models before and after applying IGAD.

Model	Venue	Dataset							
		SMD				MSL			
		w / o IGAD	w / IGAD	Δ_{model} (%)	p-value	w / o IGAD	w / IGAD	Δ_{model} (%)	p-value
CATCH	ICLR, 2025	0.1904 ± 0.0034	0.1970 ± 0.0022	+3.47	***	0.0331 ± 0.0010	0.0334 ± 0.0016	+0.91	***
M2N2	AAAI, 2024	0.0211 ± 0.0002	0.0362 ± 0.0066	+71.56	***	0.2024 ± 0.1532	0.8250 ± 0.2050	+307.61	***
FITS	ICLR, 2024	0.0409 ± 0.0027	0.0551 ± 0.0027	+34.72	***	0.0087 ± 0.0017	0.0107 ± 0.0038	+22.99	**
ModernTCN	ICLR, 2024	0.1521 ± 0.0019	0.1954 ± 0.0006	+28.47	***	0.0486 ± 0.0054	0.0620 ± 0.0419	+27.57	***
Peri-midFormer	NeurIPS, 2024	0.1545 ± 0.0073	0.1573 ± 0.0078	+1.81	***	0.0330 ± 0.0017	0.0356 ± 0.0011	+7.88	*
SARAD	NeurIPS, 2024	0.2266 ± 0.0047	0.2256 ± 0.0041	-0.44	***	0.0277 ± 0.0134	0.0504 ± 0.0549	+81.95	***
TimesNet	ICLR, 2023	0.0741 ± 0.0153	0.0759 ± 0.0265	+2.43	***	0.0072 ± 0.0028	0.0073 ± 0.0030	+1.39	***
OFA	NeurIPS, 2023	0.0559 ± 0.0027	0.1064 ± 0.0183	+90.34	***	0.0083 ± 0.0050	0.0105 ± 0.0091	+26.51	***
A.T.	ICLR, 2022	0.0259 ± 0.0043	0.0421 ± 0.0402	+62.55	***	0.0063 ± 0.0005	0.0064 ± 0.0009	+1.59	***
FGANomaly	TKDE, 2021	0.3611 ± 0.0041	0.3615 ± 0.0117	+0.11	***	0.0514 ± 0.0128	0.0578 ± 0.0087	+12.45	*
CAE-M	TKDE, 2021	0.0888 ± 0.0909	0.0863 ± 0.0985	-2.82	***	0.0043 ± 0.0001	0.0045 ± 0.0001	+4.65	***
MTAD-GAT	ICDM, 2020	0.3764 ± 0.0016	0.4169 ± 0.0250	+10.76	***	0.1731 ± 0.0022	0.2504 ± 0.1418	+44.66	*
OmniAnomaly	KDD, 2019	0.2096 ± 0.0022	0.2139 ± 0.0021	+2.05	***	0.0052 ± 0.0002	0.0087 ± 0.0028	+67.31	***
MSCRED	AAAI, 2019	0.3220 ± 0.0277	0.3245 ± 0.0206	+0.78	***	0.0111 ± 0.0007	0.0102 ± 0.0008	-8.11	*
DAGMM	ICLR, 2018	0.0322 ± 0.0031	0.1338 ± 0.0476	+315.53	***	0.0042 ± 0.0009	0.0068 ± 0.0021	+61.90	***
Δ_{data} (%)		Mean: 0.1554	Mean: 0.1752	+12.71		Mean: 0.0416	Mean: 0.0920	+120.89	

Model	Venue	Dataset							
		PSM				SMAP			
		w / o IGAD	w / IGAD	Δ_{model} (%)	p-value	w / o IGAD	w / IGAD	Δ_{model} (%)	p-value
CATCH	ICLR, 2025	0.1284 ± 0.0031	0.1326 ± 0.0028	+3.27	***	0.2882 ± 0.0012	0.2931 ± 0.0009	+1.70	***
M2N2	AAAI, 2024	0.2989 ± 0.0055	0.3010 ± 0.0021	+0.70	***	0.1934 ± 0.0046	0.1973 ± 0.0425	+2.02	***
FITS	ICLR, 2024	0.1163 ± 0.0003	0.1173 ± 0.0006	+0.86	***	0.2704 ± 0.0116	0.2851 ± 0.0128	+5.44	***
ModernTCN	ICLR, 2024	0.1383 ± 0.0002	0.1337 ± 0.0002	-3.33	***	0.4561 ± 0.0094	0.4144 ± 0.0077	-9.14	***
Peri-midFormer	NeurIPS, 2024	0.1310 ± 0.0005	0.1311 ± 0.0004	+0.08	***	0.5064 ± 0.0285	0.5067 ± 0.0254	+0.06	***
SARAD	NeurIPS, 2024	0.1499 ± 0.0078	0.1550 ± 0.0049	+3.40	***	0.8469 ± 0.0189	0.8477 ± 0.0176	+0.09	***
TimesNet	ICLR, 2023	0.1174 ± 0.0015	0.1297 ± 0.0055	+10.48	***	0.2678 ± 0.0758	0.2734 ± 0.0688	+2.09	***
OFA	NeurIPS, 2023	0.1261 ± 0.0002	0.1405 ± 0.0082	+11.42	***	0.2929 ± 0.0255	0.2969 ± 0.0254	+1.37	***
A.T.	ICLR, 2022	0.1158 ± 0.0083	0.1517 ± 0.0352	+31.00	***	0.2397 ± 0.0855	0.2578 ± 0.1130	+7.55	***
FGANomaly	TKDE, 2021	0.1970 ± 0.0046	0.1838 ± 0.0054	-6.70	***	0.9192 ± 0.0274	0.9835 ± 0.0012	+7.00	***
CAE-M	TKDE, 2021	0.1503 ± 0.0006	0.1504 ± 0.0129	+0.07	***	0.0736 ± 0.0001	0.0736 ± 0.0001	0.00	***
MTAD-GAT	ICDM, 2020	0.1433 ± 0.0018	0.1785 ± 0.0383	+24.56	***	0.2320 ± 0.0277	0.5131 ± 0.2348	+121.16	***
OmniAnomaly	KDD, 2019	0.1427 ± 0.0008	0.1431 ± 0.0009	+0.28	***	0.0780 ± 0.0002	0.9060 ± 0.0272	+1061.54	***
MSCRED	AAAI, 2019	0.1902 ± 0.0118	0.1727 ± 0.0171	-9.20	***	0.0958 ± 0.0009	0.1274 ± 0.0208	+32.99	***
DAGMM	ICLR, 2018	0.1672 ± 0.0015	0.1675 ± 0.0025	+0.18	***	0.0748 ± 0.0013	0.1083 ± 0.0140	+44.79	***
Δ_{data} (%)		Mean: 0.1542	Mean: 0.1592	+3.28		Mean: 0.3223	Mean: 0.4056	+25.83	

4. Experiments

(2) Performance under different weighted noise. The weights of noise vary from 1% to 20%.

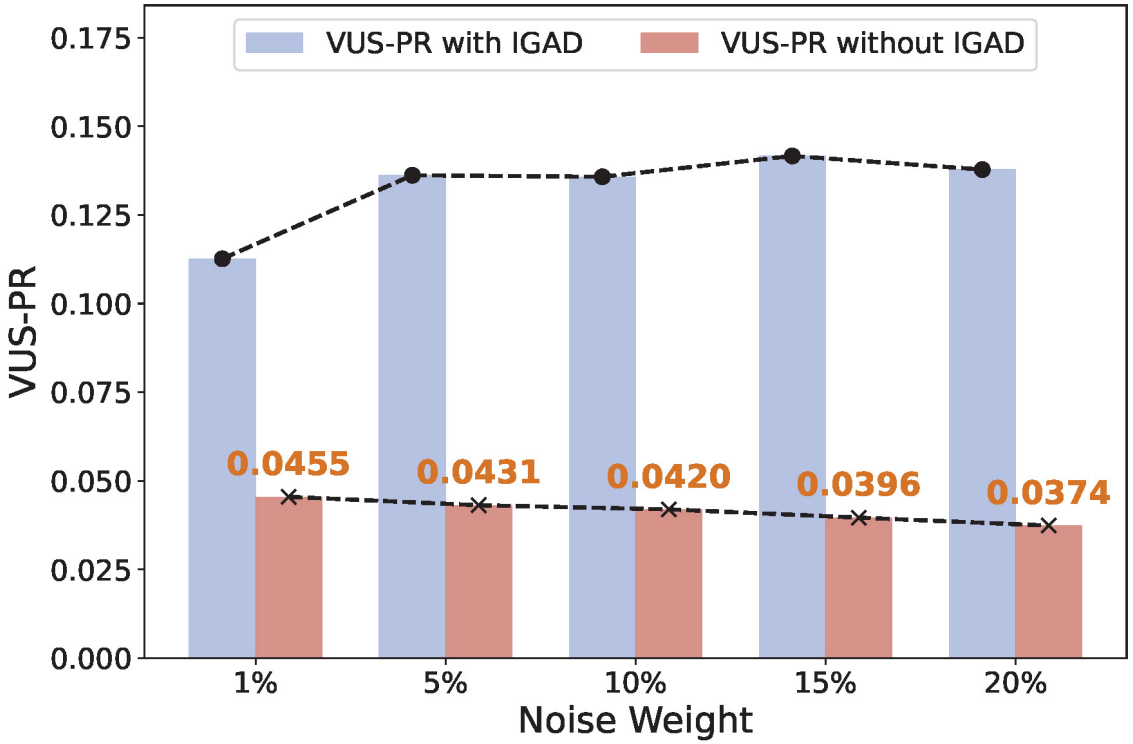

VUS-PR under noise for SARAD on MSL.

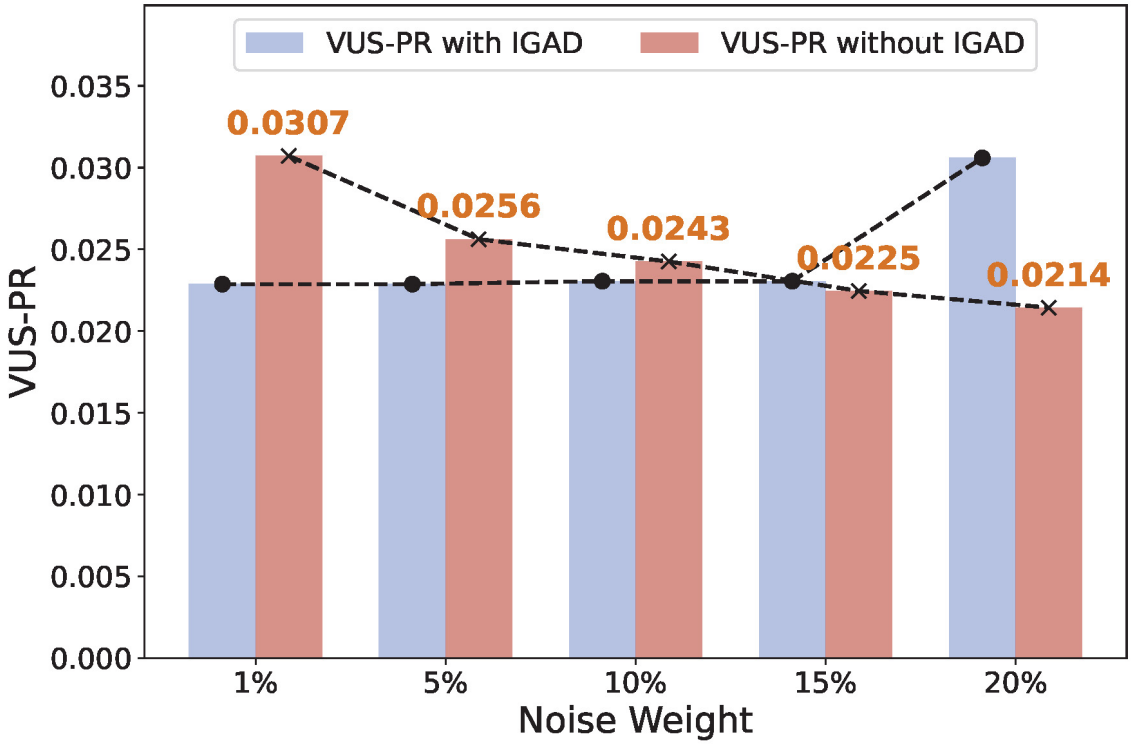

VUS-PR under noise for A.T. on SMD.

4. Experiments

(3) Anomaly score distributions. In each sub-figure, the left one shows the distribution without IGAD, and the right one with IGAD.

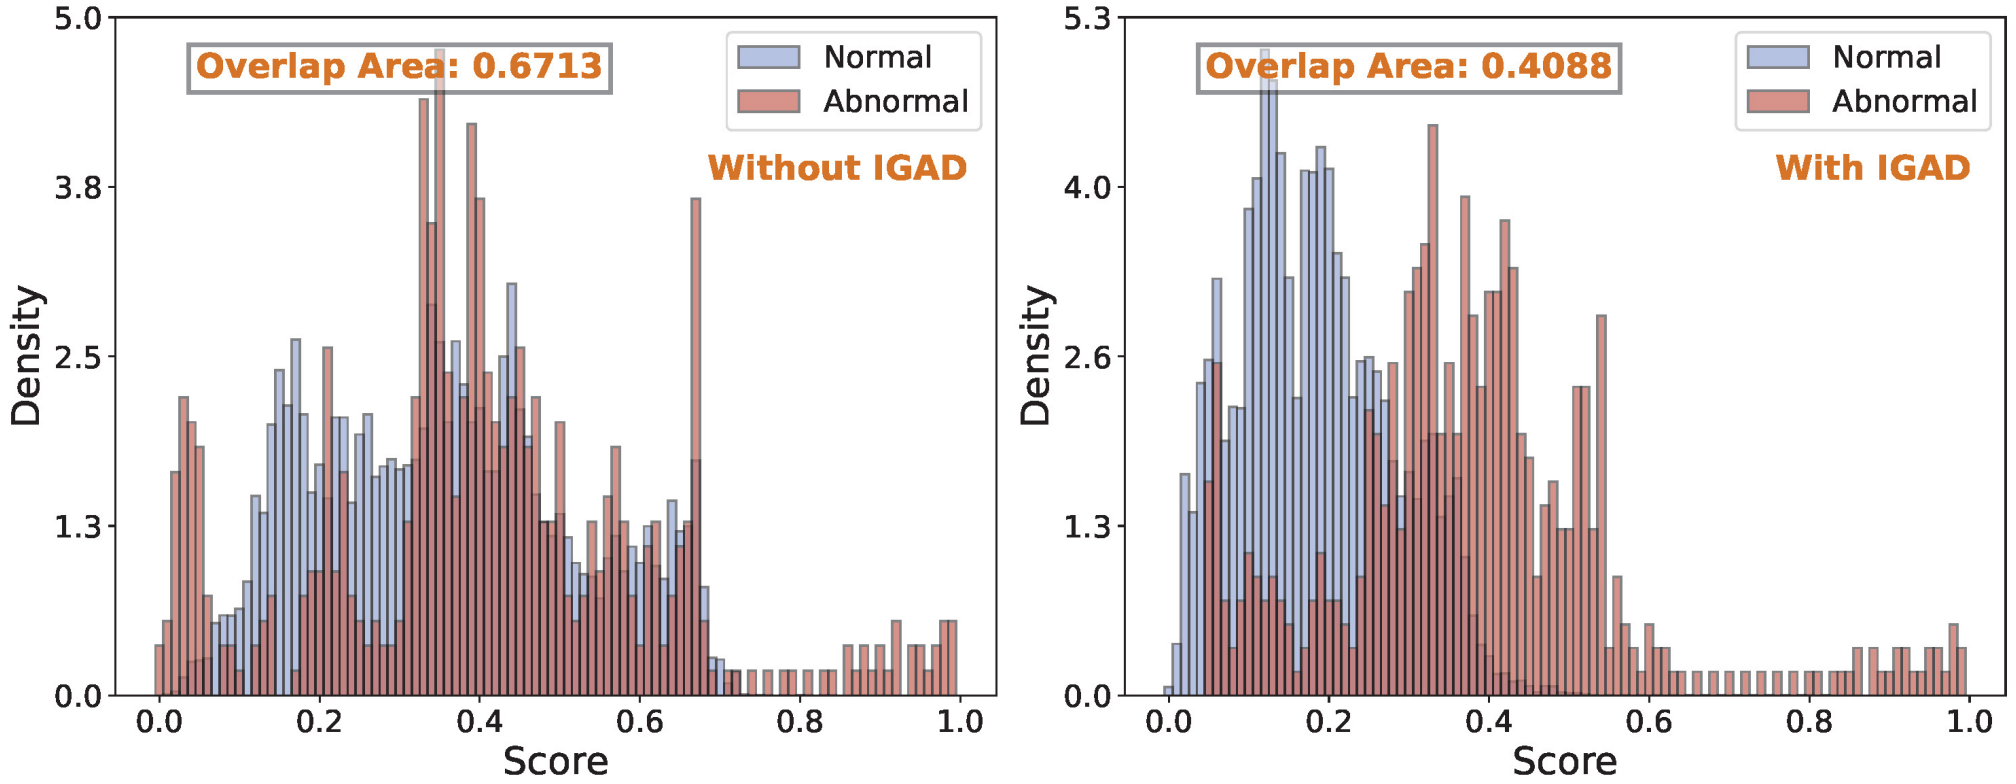

Anomaly scores for CAE-M on SMD.

4. Experiments

(3) Anomaly score distributions. In each sub-figure, the left one shows the distribution without IGAD, and the right one with IGAD.

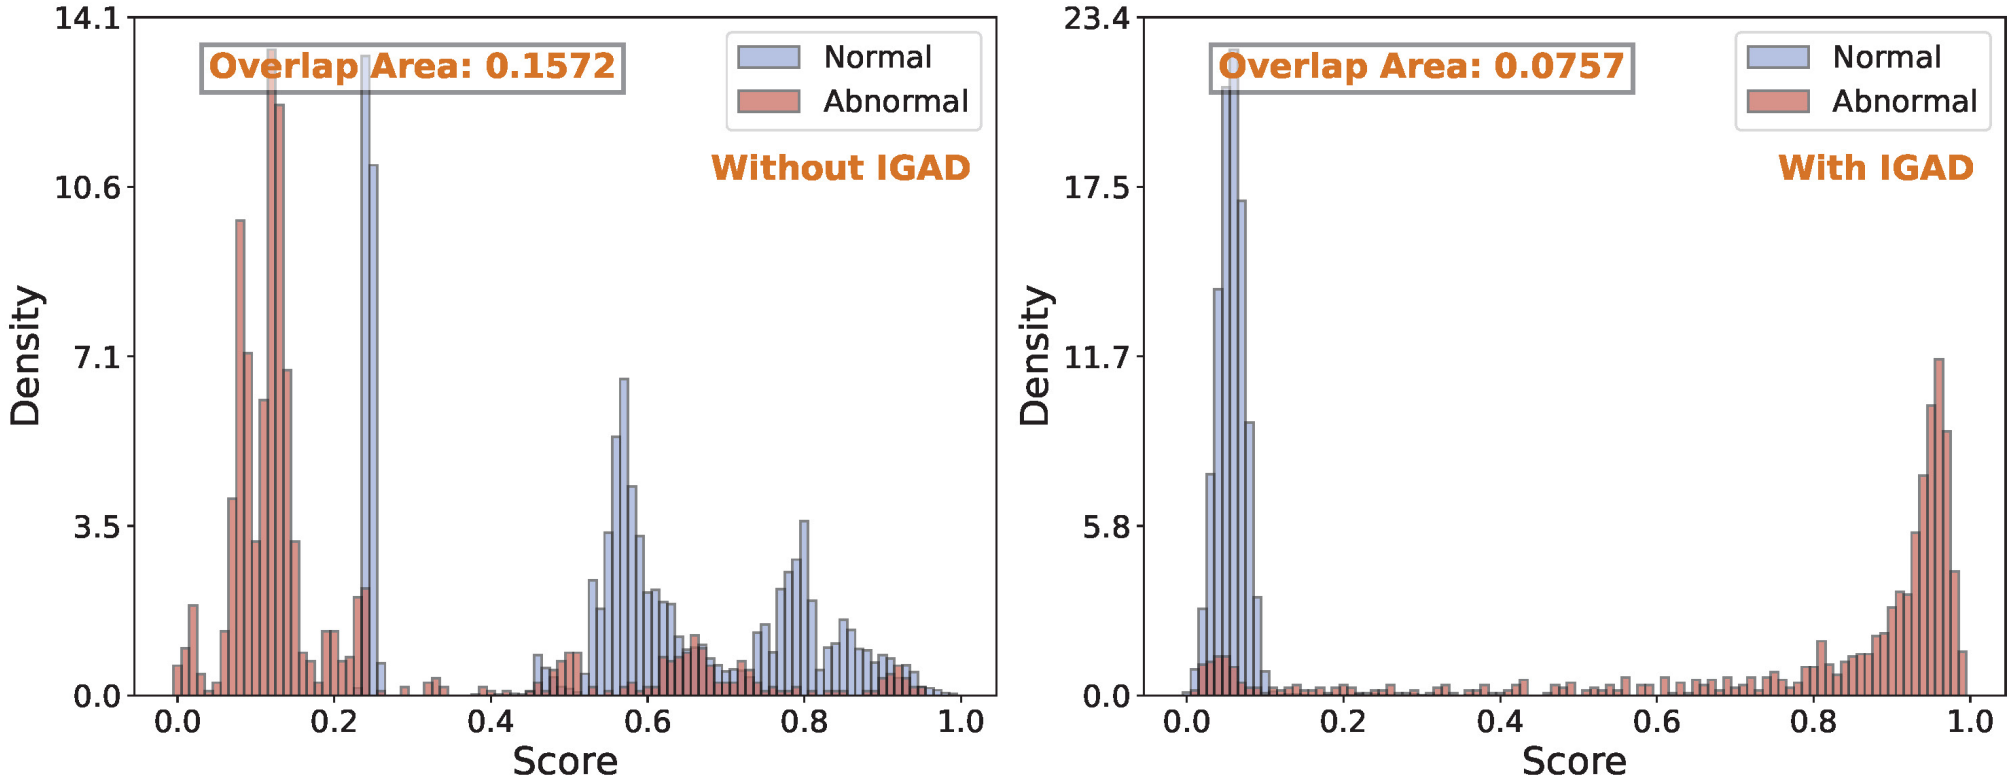

Anomaly scores for OmniAnomaly on SMAP.

4. Experiments

(4) Ablation study: loss function types and core components.

Table 2: Results on different types of loss functions. † denotes the average value of all the mean values reported in Tab.1.

Type	$\mathcal{L}_{\text{idem}}$	$\mathcal{L}_{\text{tight}}$	$\overline{\text{VUS-PR}}$	Δ (%)
only $\mathcal{L}_1$	$\mathcal{L}_1$	$\mathcal{L}_1$	**0.2080**†	N.A.
$\mathcal{L}_1 + \mathcal{L}_2$	$\mathcal{L}_1$	$\mathcal{L}_2$	0.1619	-22.16
$\mathcal{L}_2 + \mathcal{L}_1$	$\mathcal{L}_2$	$\mathcal{L}_1$	0.1857	-10.72
only $\mathcal{L}_2$	$\mathcal{L}_2$	$\mathcal{L}_2$	0.1684	-19.04

$\mathcal{L}_1$ is the MAE Loss and $\mathcal{L}_2$ is the MSE Loss. We calculate $\overline{\text{VUS-PR}}$ as the average values of all experiments shown in Tab.1 (15 basic models and 4 datasets) for each combination of $\mathcal{L}_{\text{idem}}$ and $\mathcal{L}_{\text{tight}}$. This experimental strategy is also performed in Tab.3.

Table 3: Results on different operation subset. † denotes the average value of all the mean values reported in Tab.1.

Component			$\overline{\text{VUS-PR}}$	Δ (%)
$\mathcal{L}_{\text{idem}}$	$\mathcal{L}_{\text{tight}}$	$\mathcal{L}^*_{\text{tight}}$		
✓	✓	✓	**0.2080**†	N.A.
×	✓	✓	0.1914	-7.98
✓	×	✓	Invalid	N.A.
✓	✓	×	0.1822	-12.40
✓	×	×	0.1704	-18.08
×	✓	×	0.1803	-13.32
×	×	✓	Invalid	N.A.
×	×	×	0.1684†	-19.04

4. Experiments

(4) Ablation study: VUS-PR on foundational models AE and VAE with different datasets.

Model	Dataset							
	SMD				MSL			
	w / o IGAD	w / IGAD	Δ_{model} (%)	p-value	w / o IGAD	w / IGAD	Δ_{model} (%)	p-value
AE	0.3113±0.0029	0.3823±0.0053	+22.81	***	0.0083±0.0003	0.0370±0.0208	+345.78	***
VAE	0.3556±0.0041	0.3556±0.0037	0.00	***	0.0086±0.0004	0.0087±0.0002	+1.16	***

Model	Dataset							
	PSM				SMAP			
	w / o IGAD	w / IGAD	Δ_{model} (%)	p-value	w / o IGAD	w / IGAD	Δ_{model} (%)	p-value
AE	0.1437±0.0001	0.1640±0.0098	+14.13	***	0.0969±0.0007	0.9149±0.0784	+844.17	***
VAE	0.1464±0.0004	0.1465±0.0003	+0.07	***	0.0930±0.0014	0.0947±0.0014	+1.83	***

4. Experiments

(4) Ablation study: A comparison of different signal augmentation strategies. Left: Frequency Resampling Strategy in IGAD; Middle: Time-Domain PCA with k = 5; Right: Time-Domain PCA with k = 10.

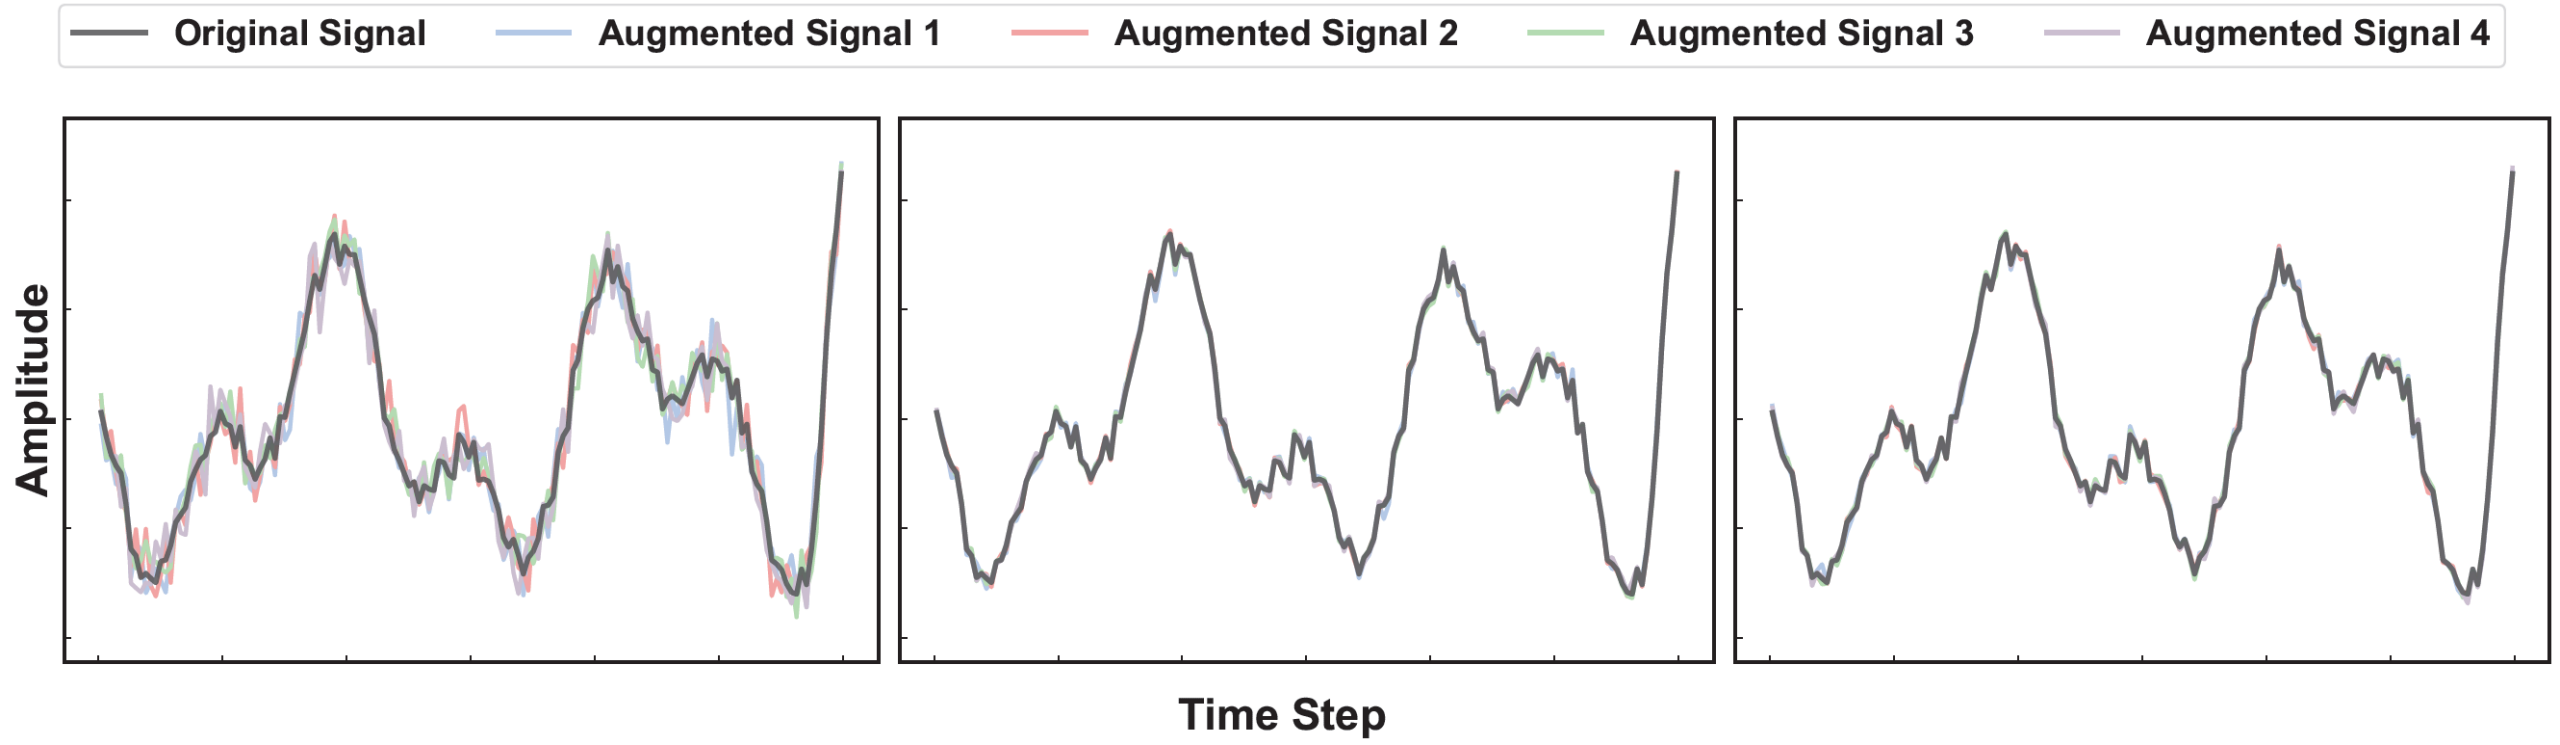

4. Experiments

(5) Visualize the latent space of different models before and after applying IGAD.

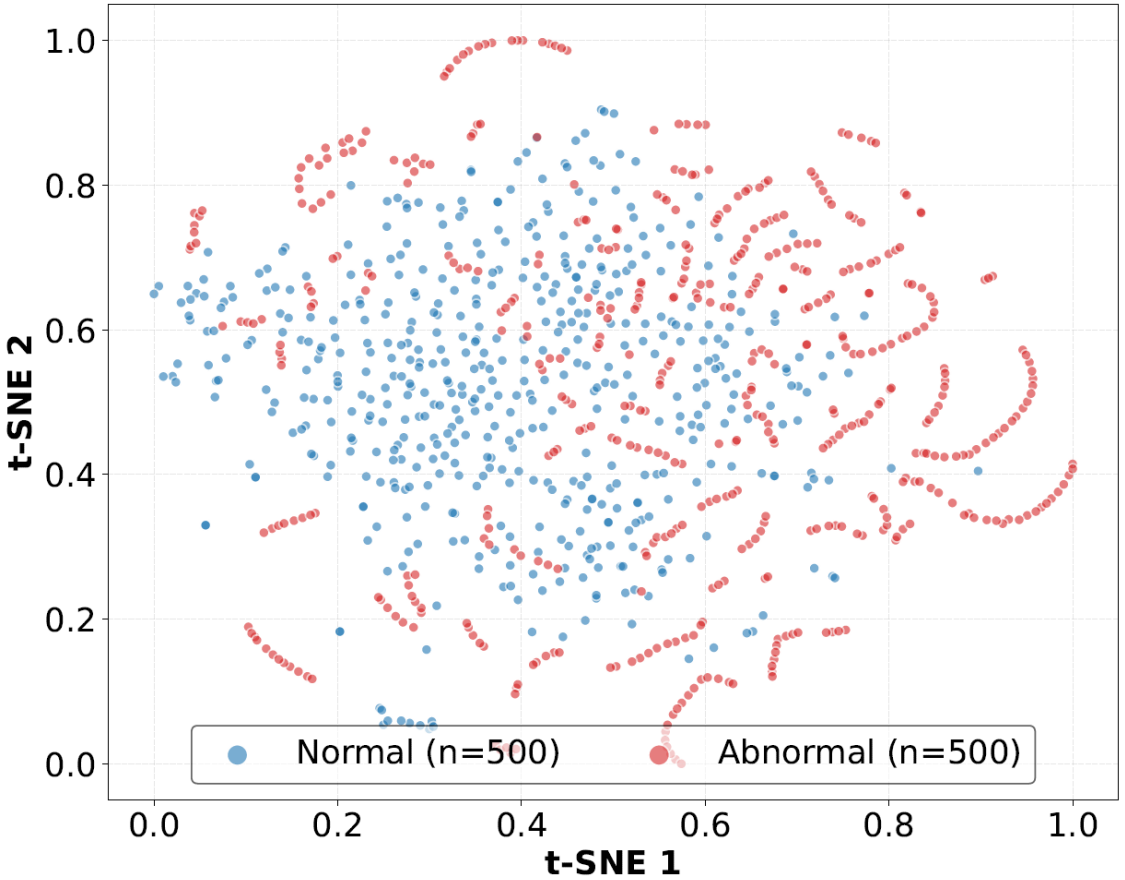

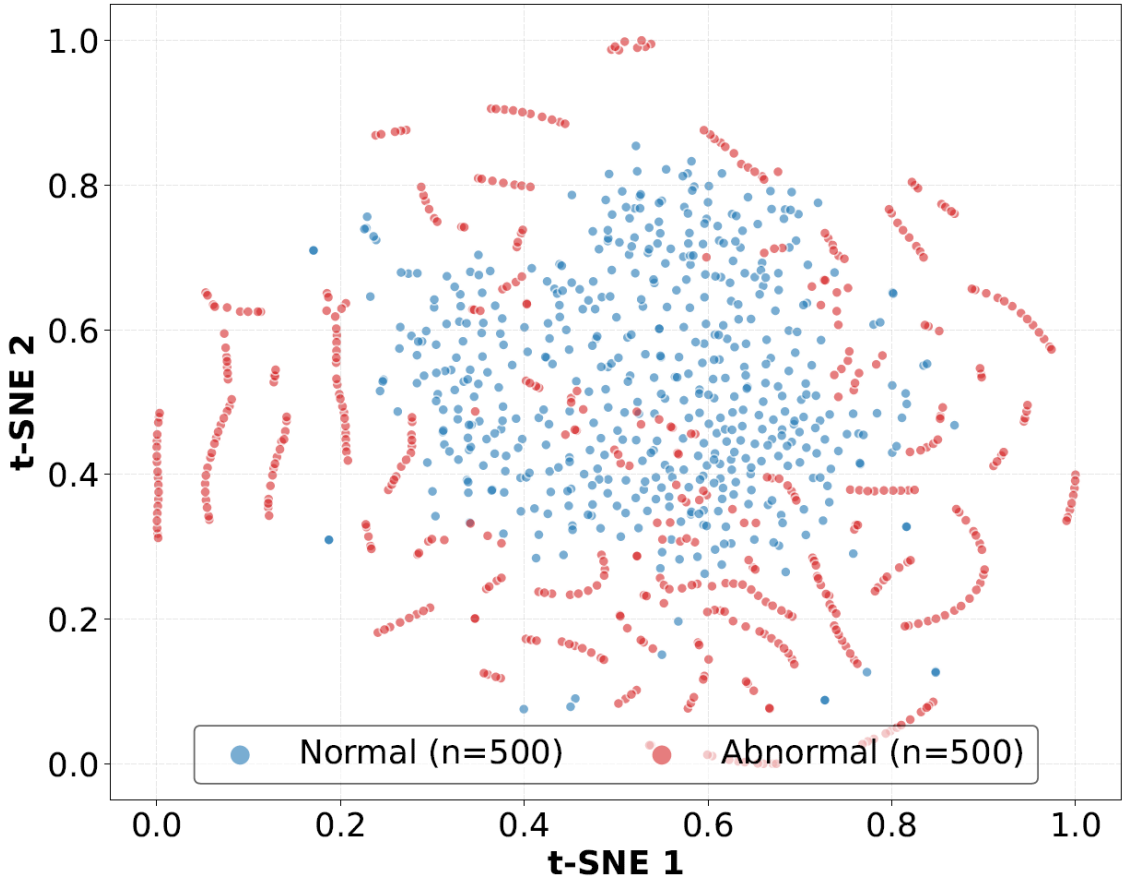

Under the effect of IGAD, the model gains a clearer boundary to distinguish normal instances from abnormal instances. This aligns with our design principles to modify and tighten the target manifold $\mathcal{M}_{\text{target}}$, with the aim of containing enough normal instances and dropping out potential abnormal instances.

4. Experiments

(5) Visualize the latent space of different models before and after applying IGAD.

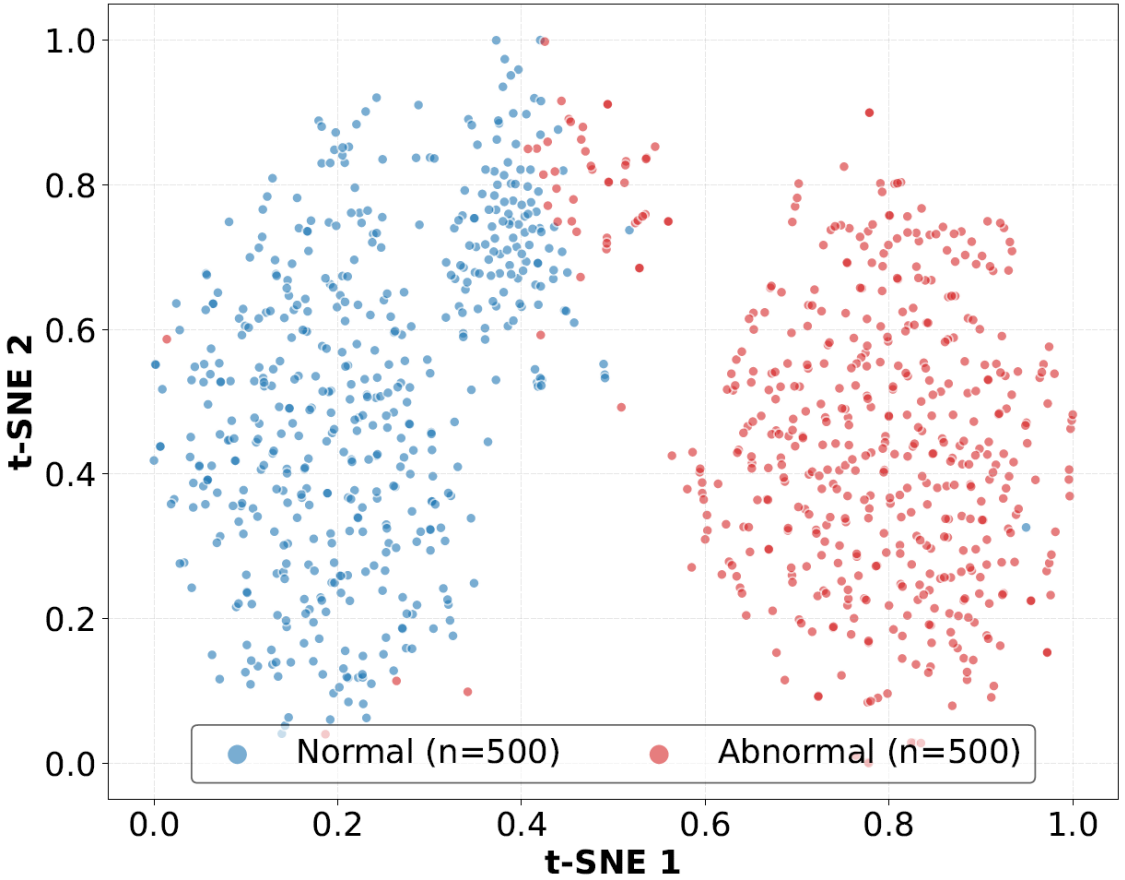

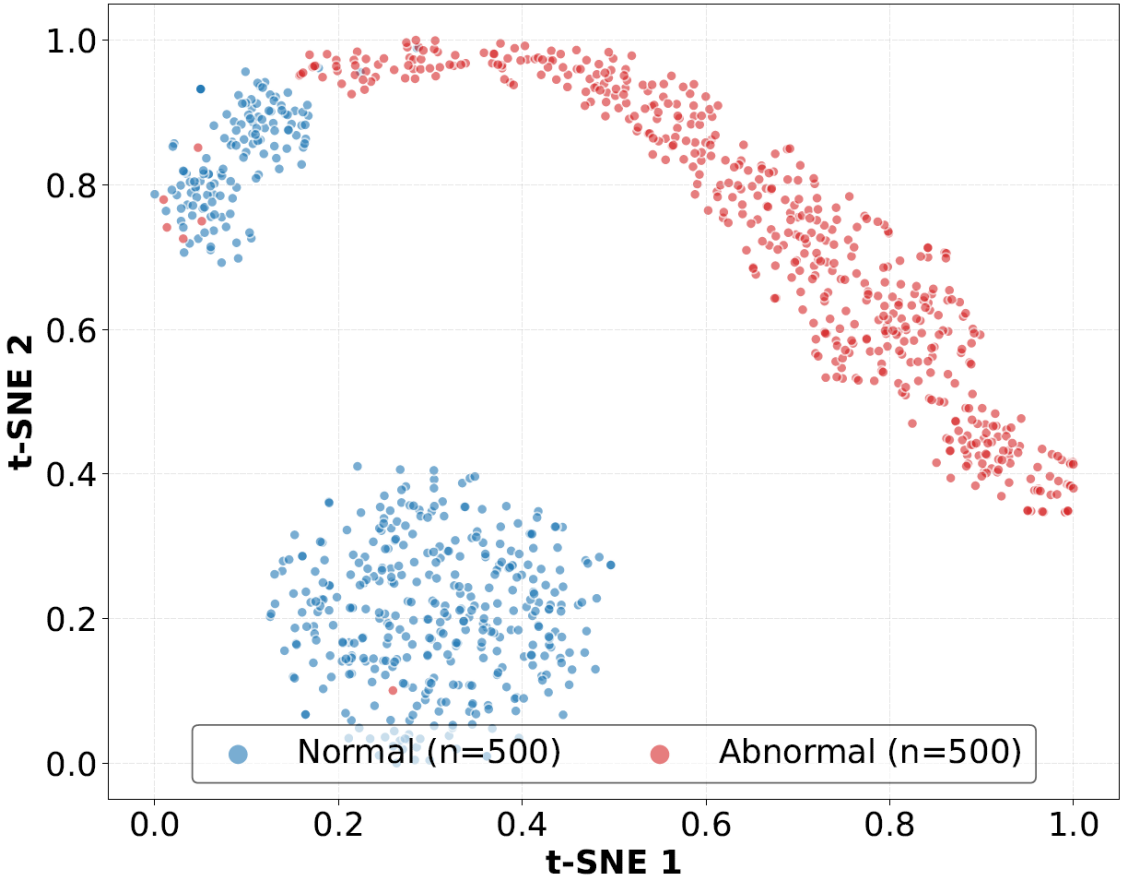

Under the effect of IGAD, the model gains a clearer boundary to distinguish normal instances from abnormal instances. This aligns with our design principles to modify and tighten the target manifold $\mathcal{M}_{\text{target}}$, with the aim of containing enough normal instances and dropping out potential abnormal instances.

5. Conclusion

This paper introduces IGAD, a novel module which can be easily integrated into reconstruction-based methods to enhance their effectiveness in detecting anomalies in MTS. Meanwhile, we conduct further experiments and explorations, addressing the issues of over generalization and overall performance balance of these models from the perspective of manifold. With defined optimization objectives, we aim to not only modify the target manifold, balancing the robustness and sensitivity to inherent normal patterns of time series, but also tighten the manifold to exclude potential abnormal time instances. The experimental results demonstrate the effectiveness of IGAD and show its great potential for further applications in MTS anomaly detection.

Thank you!

Multivariate Time Series Anomaly Detection with Idempotent Reconstruction

Xin Sun, Heng Zhou, Chao Li

Zhejiang University

NeurIPS 2025, San Diego