NeurIPS 2026

CAPS

Supervision Recovery for Time Series Anomaly Detection via Context-Anchored Pairing

Presented by Yifei Gao

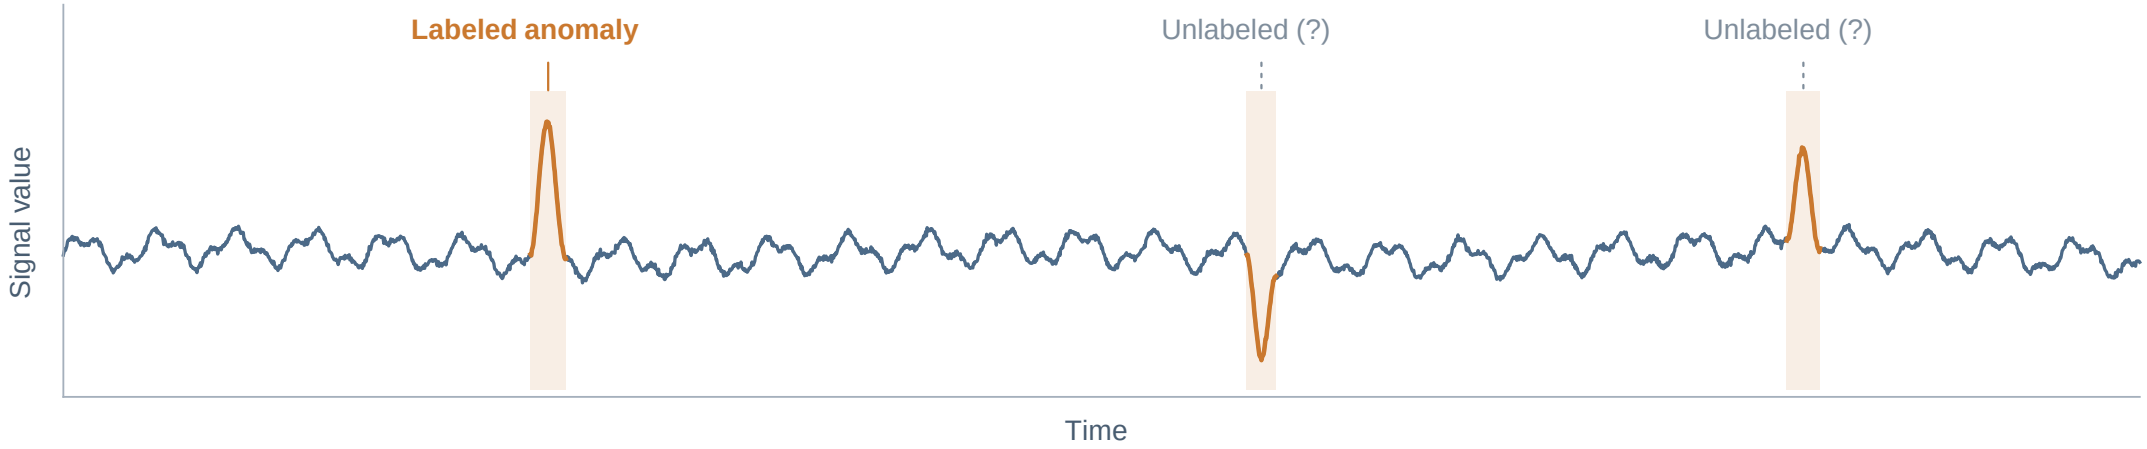
Yifei Gao, Tian Lan, Yimeng Lu, Xuming An, Meng Wang, Wenjun He,
Yijie Li, and Chen Zhang

github.com/gyf-0623/caps

Challenge 1: Target domains lack effective anomaly supervision

- Abnormal observations are rare compared with normal data
- Anomaly labels are limited or unavailable
- Detectors receive little direct supervision about abnormal behavior

Rare anomalies and sparse labels

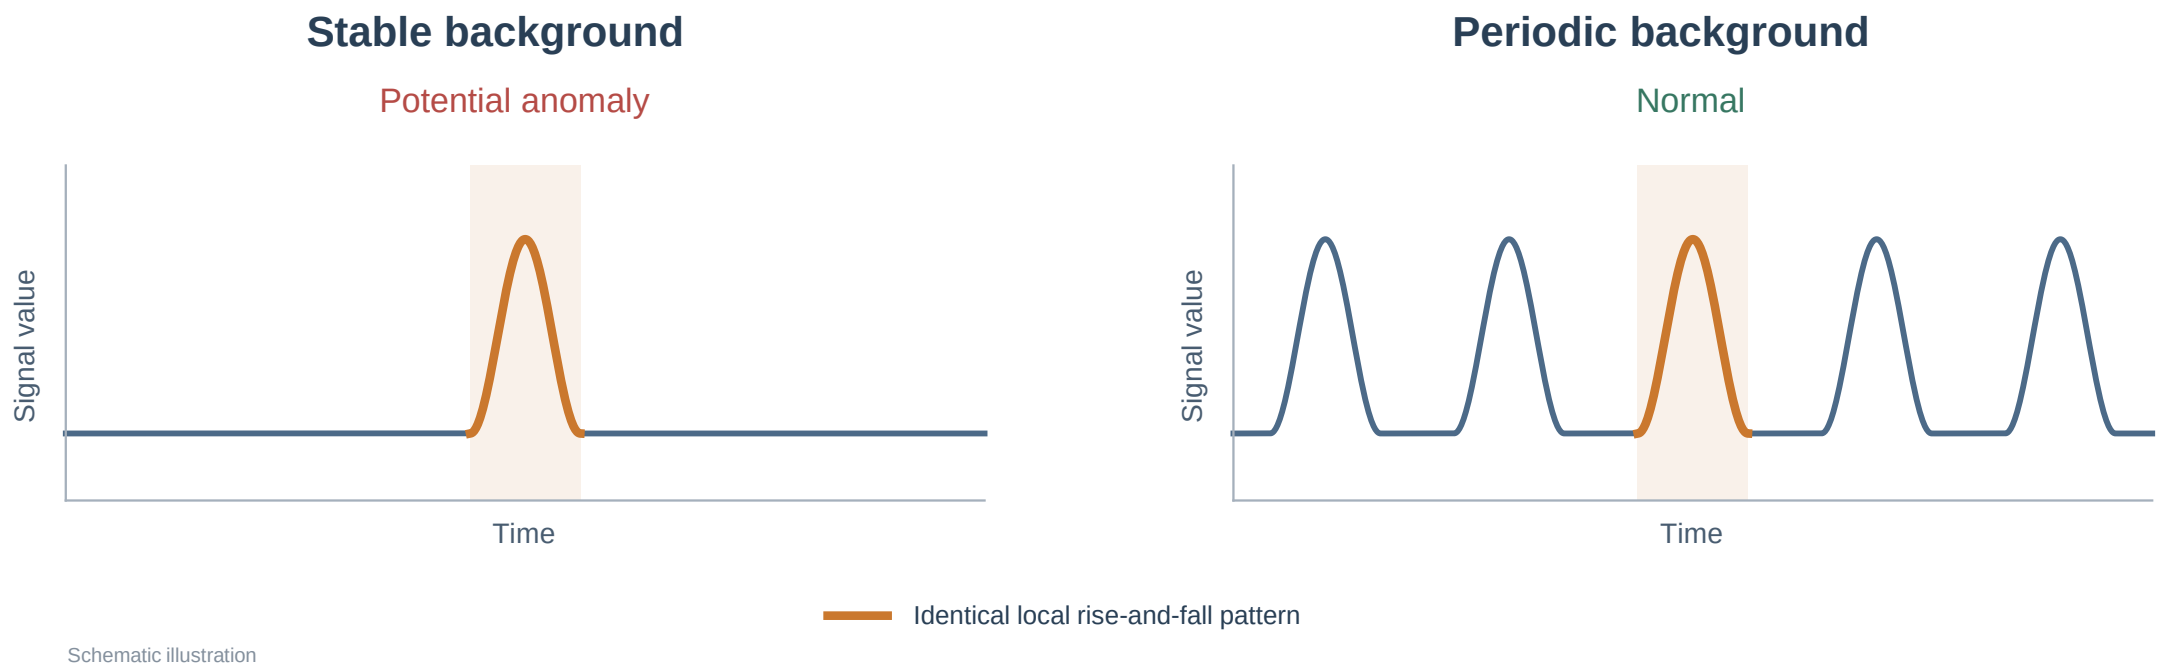


Only a small part of the sequence is anomalous, and fewer events have labels

Schematic illustration

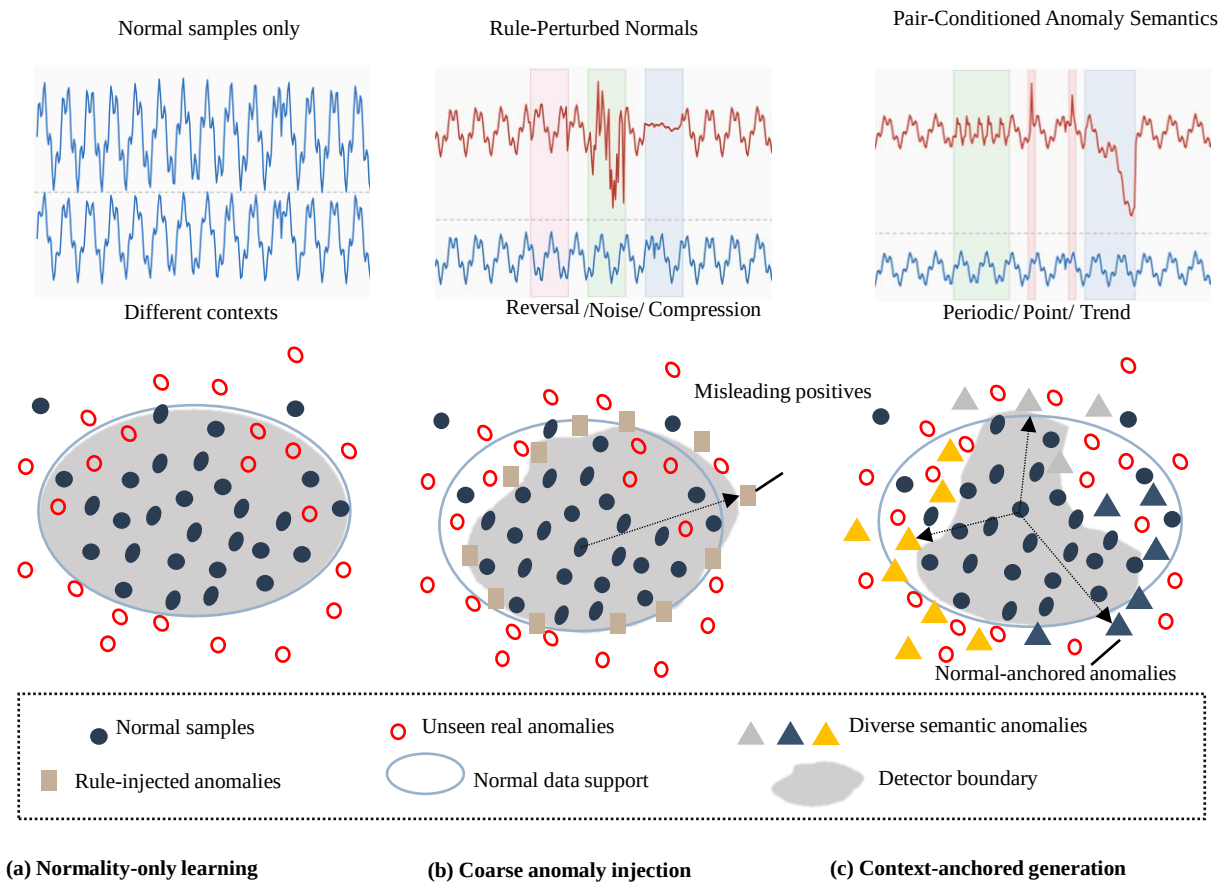
Challenge 2: Temporal anomalies are highly context-dependent

- The same local pattern can be normal or anomalous
- (An isolated rise and fall may indicate an anomaly in a stable background. The same pattern can be normal when it follows the surrounding periodic behavior)



Schematic illustration

Why Recover Anomaly Supervision?



Normality-only learning

Limited explicit evidence about abnormal behavior

Coarse anomaly injection

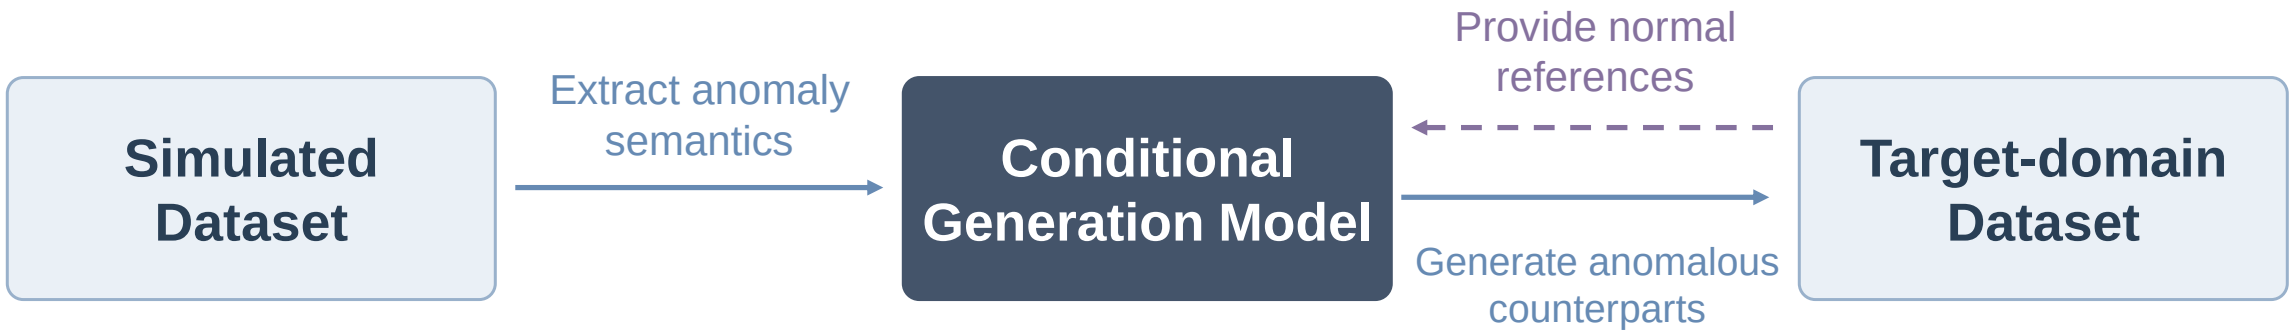
Predefined perturbations may overlook contextually meaningful deviations

Context-anchored supervision

Learn anomaly semantics from simulated normal–anomalous pairs

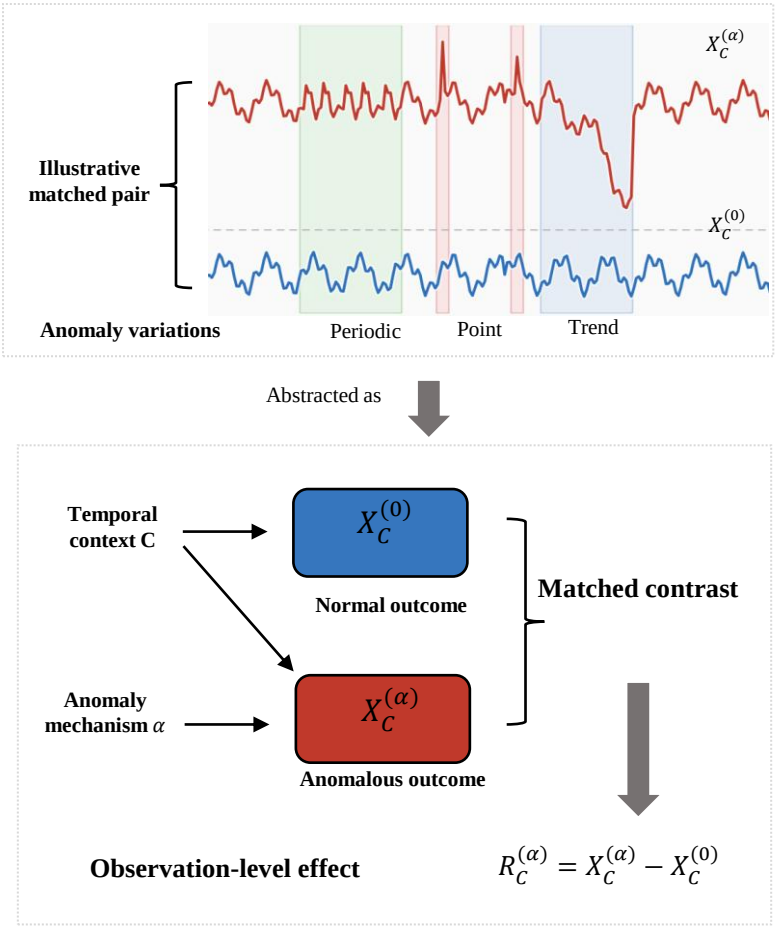
Realize diverse effects on target references while preserving their temporal backgrounds

How to Recover Anomaly Supervision?



1. Why use simulated data?
2. How do we construct the simulated dataset?
3. How do we extract anomaly semantics?
4. How do we learn conditional generation?
5. How do we instantiate anomalies on target data?

Simulated normal–anomalous pairs



(a) Matched-supervision view of temporal anomalies

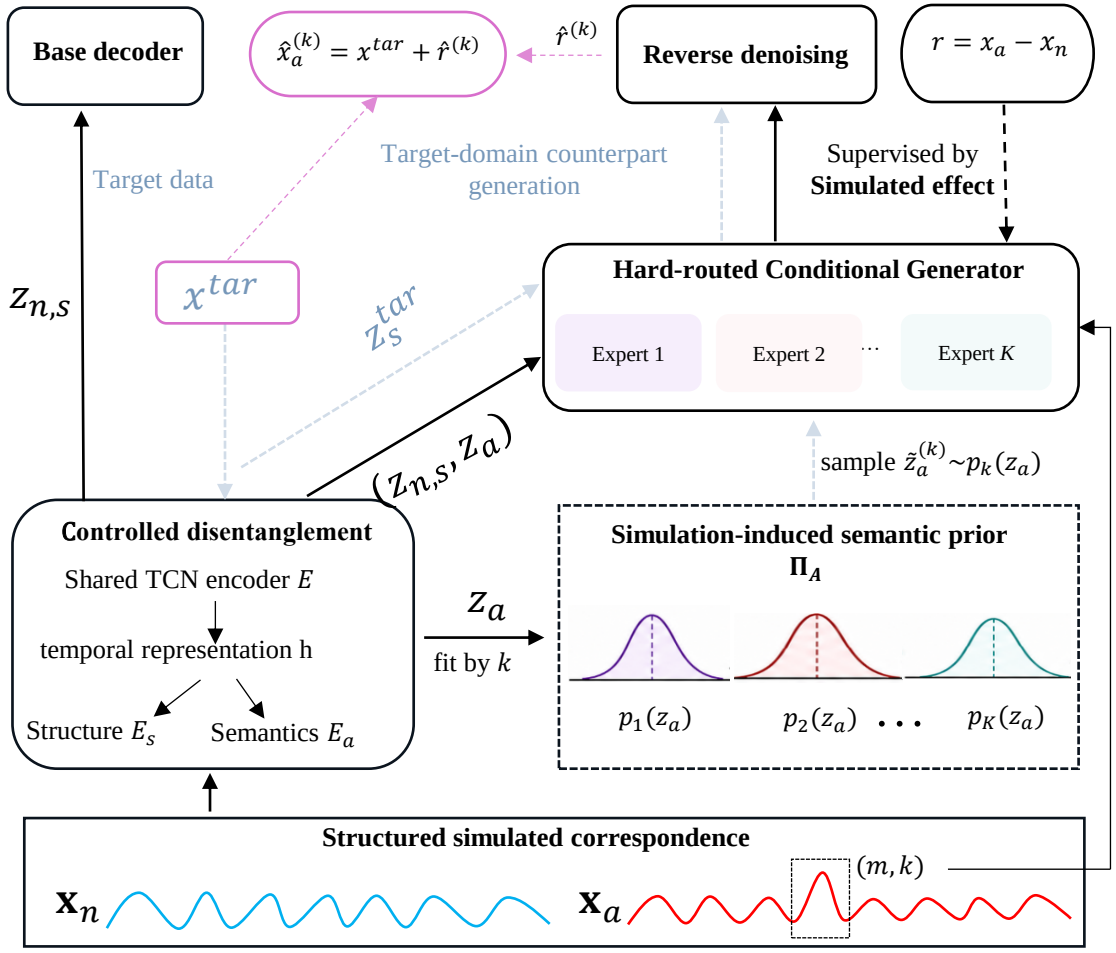
Why use simulated data ?

- Matched normal–anomalous outcomes are generally unavailable in real data;
- Simulation provides controlled anomaly variations under a shared temporal context;
- Paired comparisons reveal the changes induced by anomalies.

How do we construct the simulated dataset?

- Generate normal sequences across diverse temporal contexts;
- Introduce controlled periodic, point, or trend anomalies;
- Retain matched pairs with known anomaly supports and family labels.

The CAPS generation framework



(b) CAPS framework

How do we extract anomaly semantics?

- Learn separate representations of temporal structure and anomaly effects from matched pairs;
- Fit family-specific semantic priors for diverse effect sampling.

How do we generate anomalies?

- Train conditional diffusion experts to generate residual effects;
- Condition generation on anomaly semantics, reference structure, and anomaly support

How do we instantiate anomalies on target data?

- Combine sampled semantics with the target reference's temporal structure;
- Generate multiple anomalous counterparts to train an independent detector.

Learning Anomaly Effects Beyond Their Original Context

1. Align the shared structure

- Matched outcomes share the same background.

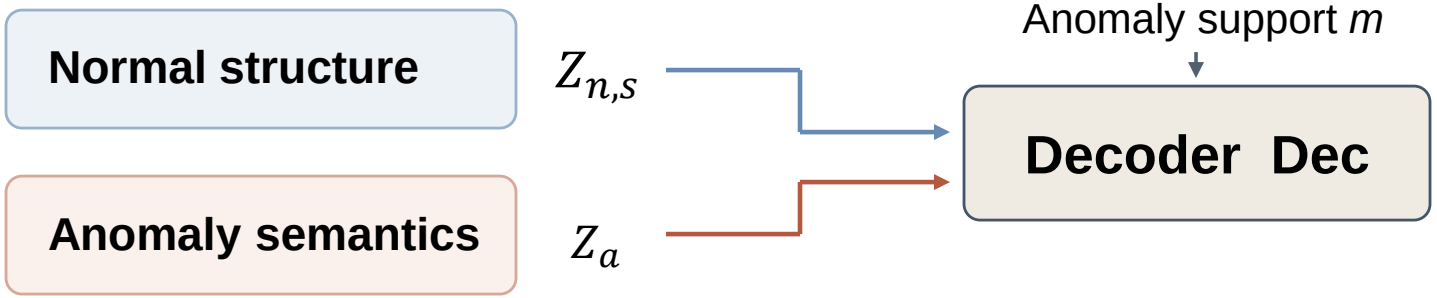
$$Z_{n,s} \approx Z_{a,s}$$

2. Anchor anomaly semantics

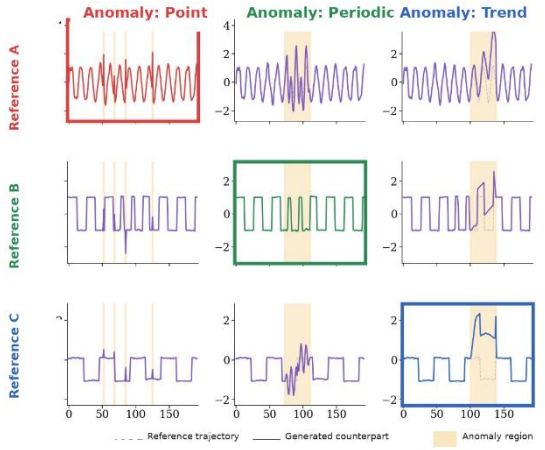
- Suppress anomaly codes for normal samples.

$$Z_{n,a} \approx 0$$

3. Reconstruct the matched anomalous outcome



$$L_{\text{cf}} = \text{MSE}(\text{DEC}(Z_{n,s}, Z_a, m), x_a)$$



Matched anomaly

These learned effects condition generation on new target references.

Performance on nine datasets under four metrics

Metric	Method	IOPS	MGAB	NAB	Power	SED	UCR	NEK	TODS	YAHOO
Affiliation-F	LOF	81.06	68.44	75.75	66.76	63.85	73.53	84.74	60.58	75.63
	IForest	52.81	68.82	39.84	0.00	70.09	50.56	71.15	44.17	33.30
	OmniAnomaly	80.32	67.35	92.35	78.16	61.26	73.53	86.30	50.73	71.31
	TranAD	83.19	67.28	90.28	71.56	61.03	73.31	85.02	52.76	76.08
	USAD	71.08	67.81	91.54	76.48	55.60	76.00	71.13	47.90	53.05
	AnomTrans.	70.79	67.65	79.03	71.57	68.21	80.03	75.32	44.57	64.75
	TimesNet	85.04	66.93	89.24	69.08	66.98	80.55	**89.79**	70.36	83.13
	FITS	86.36	67.77	89.03	67.96	66.99	74.36	87.76	71.77	81.19
	TCN-AE	83.91	67.61	85.74	83.34	67.75	74.54	85.74	72.61	82.47
	TCN-Injection	87.58	69.13	88.89	85.04	70.94	84.32	81.52	72.05	83.22
	TCN-Supervised	84.35	67.59	89.62	85.96	68.47	72.40	85.24	71.75	83.11
	TCN-Forecast	76.05	67.27	84.06	84.80	67.68	72.81	85.83	72.48	82.64
	CAPS	**89.24**	**69.22**	**93.25**	**85.98**	**74.09**	**88.41**	86.30	**86.85***	**94.63***
	CAPS(Diagnosis)	89.83*	69.61*	93.81*	86.35*	75.82*	89.72*	86.52*	81.32	94.57
$F1_T$	LOF	27.97	1.15	35.76	19.80	9.60	8.31	63.57	31.63	55.93
	IForest	7.64	0.84	21.44	0.00	9.54	6.36	65.56	11.06	4.90
	OmniAnomaly	51.17	1.61	40.09	23.48	9.68	8.47	**82.20**	14.33	24.16
	TranAD	22.63	1.65	37.28	22.36	9.57	7.75	69.97	13.51	8.41
	USAD	20.99	4.07	**61.46**	28.23	9.54	14.63	70.64	20.85	9.35
	AnomTrans.	2.56	0.56	8.13	7.81	15.07	4.09	22.62	6.18	1.39
	TimesNet	25.12	0.50	28.80	15.77	7.90	4.69	42.68	17.97	7.52
	FITS	15.38	0.54	31.98	15.77	7.95	4.95	42.90	16.39	4.86
	TCN-AE	39.38	3.25	36.63	19.92	10.23	13.87	68.58	24.85	58.94
	TCN-Injection	48.37	**23.89**	46.91	27.04	21.59	**44.79**	64.45	17.26	35.68
	TCN-Supervised	41.92	1.94	42.50	20.13	11.01	9.56	72.93	22.54	59.52
	TCN-Forecast	30.24	1.56	37.01	19.82	10.15	8.50	68.89	23.42	61.61
	CAPS	**60.62**	7.58	57.90	**29.20**	**24.27***	41.53	80.15	**46.17***	**71.35**
	CAPS(Diagnosis)	61.90*	7.99*	59.21*	30.36*	15.76	45.26*	83.33*	40.77	75.23*
Standard-F1	LOF	30.28	1.05	24.04	12.18	4.11	4.70	56.92	25.77	48.95
	IForest	8.37	0.73	29.41	19.77	3.81	4.09	58.10	13.35	3.20
	OmniAnomaly	47.05	1.44	28.81	23.50	0.43	5.11	74.03	12.65	21.40
	TranAD	34.85	1.46	27.33	22.36	2.63	4.40	60.36	11.94	5.70
	USAD	30.66	3.89	**56.15**	28.24	3.41	10.74	62.91	23.87	7.21
	AnomTrans.	3.74	0.58	19.35	15.77	16.87	3.44	30.23	13.58	2.31
	TimesNet	22.52	0.45	24.98	15.77	7.90	3.84	43.90	17.58	6.97
	FITS	15.81	0.49	29.00	15.77	7.95	3.86	44.27	18.08	4.58
	TCN-AE	27.15	2.26	28.53	20.08	10.22	9.06	58.61	20.72	49.27
	TCN-Injection	46.37	**14.61**	38.00	27.00	21.42	35.65	60.35	16.60	31.21
	TCN-Supervised	42.38	1.56	34.35	20.03	10.92	5.24	62.58	19.23	51.24
	TCN-Forecast	18.65	1.30	28.56	19.77	10.16	4.94	61.32	20.56	54.63
	CAPS	**51.63**	4.54	51.78	**29.37**	**24.11***	**37.70**	**76.23**	**41.24***	**69.58**
	CAPS(Diagnosis)	52.14*	6.84*	51.82*	30.30*	15.78	41.24*	76.78*	37.22	73.55*
VUS-PR	LOF	19.43	0.57	21.18	9.31	6.81	2.39	58.52	49.14	41.37
	IForest	8.59	0.62	23.57	11.56	7.71	2.88	56.50	46.62	10.47
	OmniAnomaly	25.35	0.64	27.17	14.32	6.20	2.40	74.51	45.55	29.26
	TranAD	21.61	0.64	24.82	13.04	5.75	2.25	61.63	47.33	25.78
	USAD	16.58	0.75	**55.03**	18.68	4.37	8.85	58.53	56.36	14.15
	AnomTrans.	4.60	0.43	13.11	8.84	8.29	2.12	22.67	44.74	8.00
	TimesNet	17.94	0.42	20.45	7.73	4.70	2.37	44.73	57.06	25.08
	FITS	13.94	0.40	25.43	7.46	5.08	2.20	38.77	56.00	18.81
	TCN-AE	17.27	0.87	21.56	14.95	8.99	3.63	61.04	46.30	41.03
	TCN-Injection	48.69	**3.27**	34.97	**21.87**	17.50	26.07	60.76	49.35	42.35
	TCN-Supervised	33.62	0.61	30.66	11.85	8.50	2.87	67.52	76.61	44.28
	TCN-Forecast	11.19	0.57	20.85	11.25	8.92	2.37	63.35	45.72	43.27
	CAPS	**49.00**	1.55	52.35	19.42	**20.09***	**28.47**	**80.52**	**86.26***	**83.70**
	CAPS(Diagnosis)	50.45*	1.68*	52.73*	21.00*	13.69	31.40*	81.91*	81.92	84.32*

Ablations and controlled comparisons across nine datasets

Variant	Affiliation-F	$F1_T$	Standard-F1	VUS-PR
Shuffled residuals	79.27	41.21	37.78	41.30
w/o residual parameterization	82.76	40.96	38.35	42.27
Direct source-residual injection	81.58	43.16	39.96	42.58
w/o counterfactual loss	83.65	42.76	39.11	43.11
CAPS	**85.33**	**46.53**	**42.91**	**46.82**

- Matched pairs preserve meaningful correspondence between context and anomaly effects.
- Context-conditioned generation outperforms direct source-residual injection.
- Recombination consistency strengthens semantic learning and improves detection.

Thanks!

Yifei Gao

Industrial Statistics and Data Analytics Lab

github.com/gyf-0623/caps