

GRAPH-GUIDED RECONSTRUCTION DIFFUSION FOR MULTIVARIATE TIME SERIES ANOMALY DETECTION

000
001
002
003
004
005 **Anonymous authors**
006 Paper under double-blind review
007
008
009
010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053
054
055
056
057
058
059
060
061
062
063
064
065
066
067
068
069
070
071
072
073
074
075
076
077
078
079
080
081
082
083
084
085
086
087
088
089
090
091
092
093
094
095
096
097
098
099
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
559
560
561
562
563
564
565
566
567
568
569
569
570
571
572
573
574
575
576
577
578
579
579
580
581
582
583
584
585
586
587
588
589
589
590
591
592
593
594
595
596
597
598
599
599
600
601
602
603
604
605
606
607
608
609
609
610
611
612
613
614
615
616
617
618
619
619
620
621
622
623
624
625
626
627
628
629
629
630
631
632
633
634
635
636
637
638
639
639
640
641
642
643
644
645
646
647
648
649
649
650
651
652
653
654
655
656
657
658
659
659
660
661
662
663
664
665
666
667
668
669
669
670
671
672
673
674
675
676
677
678
679
679
680
681
682
683
684
685
686
687
688
689
689
690
691
692
693
694
695
696
697
698
699
699
700
701
702
703
704
705
706
707
708
709
709
710
711
712
713
714
715
716
717
718
719
719
720
721
722
723
724
725
726
727
728
729
729
730
731
732
733
734
735
736
737
738
739
739
740
741
742
743
744
745
746
747
748
749
749
750
751
752
753
754
755
756
757
758
759
759
760
761
762
763
764
765
766
767
768
769
769
770
771
772
773
774
775
776
777
778
779
779
780
781
782
783
784
785
786
787
788
789
789
790
791
792
793
794
795
796
797
798
799
799
800
801
802
803
804
805
806
807
808
809
809
810
811
812
813
814
815
816
817
818
819
819
820
821
822
823
824
825
826
827
828
829
829
830
831
832
833
834
835
836
837
838
839
839
840
841
842
843
844
845
846
847
848
849
849
850
851
852
853
854
855
856
857
858
859
859
860
861
862
863
864
865
866
867
868
869
869
870
871
872
873
874
875
876
877
878
879
879
880
881
882
883
884
885
886
887
888
889
889
890
891
892
893
894
895
896
897
898
899
899
900
901
902
903
904
905
906
907
908
909
909
910
911
912
913
914
915
916
917
918
919
919
920
921
922
923
924
925
926
927
928
929
929
930
931
932
933
934
935
936
937
938
939
939
940
941
942
943
944
945
946
947
948
949
949
950
951
952
953
954
955
956
957
958
959
959
960
961
962
963
964
965
966
967
968
969
969
970
971
972
973
974
975
976
977
978
979
979
980
981
982
983
984
985
986
987
988
989
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1098
1099
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1198
1199
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1298
1299
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1398
1399
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1498
1499
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1598
1599
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1698
1699
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1798
1799
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1898
1899
1899
1900
1901
1902
1903
1904
19

054 which injects similarity-based priors into the attention mechanism to capture dynamic cross-feature
 055 dependencies.
 056

057 In summary, the contributions of this paper are:

058

- 059 • Proposes **GGRD**, a diffusion-based framework that integrates smooth multi-resolution de-
 060 composition with efficient one-step reconstruction.
- 061 • Introduces **GGN** with Graph-Guided Attention to explicitly model dynamic cross-feature
 062 dependencies.
- 063 • Demonstrates state-of-the-art performance and efficiency on multiple real-world bench-
 064 marks.

065

066 2 RELATED WORK

067

068 Research on time series data mainly focuses on classification, prediction, imputation, and anomaly
 069 detection(Jin et al., 2024). Among them, anomaly detection in time series has attracted considerable
 070 attention in both industry and academia(Zamanzadeh Darban et al., 2024), especially for multivariate
 071 time series data. The methods for multivariate time series anomaly detection can be roughly divided
 072 into three categories(Zhang et al., 2025): early traditional statistical methods, traditional machine
 073 learning methods, and deep learning methods.

074 Statistical methods mainly utilize statistical knowledge and some statistical indicators to achieve
 075 anomaly detection tasks, such as ARIMA(Yaacob et al., 2010) and COPOD(Li et al., 2020). Traditional
 076 machine learning methods such as PCA(Shyu et al., 2003), kNN(Ramaswamy et al., 2000),
 077 and IForest(Liu et al., 2008) use various linear transformations, proximity measures, or outlier de-
 078tection algorithms to effectively identify abnormal behaviors. However, these methods cannot effec-
 079tively capture the complex patterns of time series data.

080 In recent years, with the rapid development of deep learning, a large number of deep learning-
 081 based time series anomaly detection models have been proposed and achieved remarkable results
 082 in different scenarios, such as VAE(Kingma & Welling, 2013), LSTM-AD(Malhotra et al., 2015),
 083 GDN(Deng & Hooi, 2021), Anomaly Transformer(Xu et al., 2021), etc. It is worth noting that
 084 Diffusion Models have demonstrated outstanding performance in image generation(Xu & Chi, 2024;
 085 Luo et al., 2024; Epstein et al., 2023), prompting researchers to introduce them into time series
 086 anomaly detection tasks. ImDiffusion(Chen et al., 2023) uses diffusion models to mask and fill time
 087 series and combines ensemble strategies to enhance the robustness of anomaly detection, being one
 088 of the earliest works to apply diffusion models to time series anomaly detection. DiffAD(Xiao et al.,
 089 2023) proposes a new denoising diffusion-based imputation method and uses a density ratio-based
 090 strategy to flexibly select normal observations, thereby reducing the interference of dense anomaly
 091 regions on the model. D3R(Wang et al., 2023) proposes a dynamic decomposition and diffusion
 092 reconstruction framework for non-stationary time series, which significantly reduces the impact of
 093 drift on detection accuracy by achieving dynamic decomposition of stable and trend components
 094 and using noise diffusion to control the information bottleneck externally.

095 To fully utilize the information in the data, some multi-resolution methods have been applied to time
 096 series tasks. For example, MG-TSD(Fan et al., 2024) uses multi-granularity guided loss to enhance
 097 prediction performance to address the instability challenge caused by randomness. MR-diff(Shen
 098 et al., 2024) uses seasonal trend decomposition and a coarse-to-fine non-autoregressive method to
 099 solve prediction tasks. MODEM(Zhong et al., 2025) designs a multi-resolution decomposable dif-
 100 fusion model for the anomaly detection task of non-stationary time series, with the core being a
 101 coarse-to-fine diffusion process and a frequency domain enhanced decomposition network, which
 102 can capture long-term trends and short-term fluctuations at different time scales, thereby effectively
 103 distinguishing anomalies from non-stationary patterns.

104 The above-mentioned methods have promoted the development of the field of anomaly detection
 105 from the perspectives of statistics, generative models, and multi-resolution modeling. In this work,
 106 we for the first time explicitly introduce the correlation of different features of time series in different
 107 temporal granularity into the modeling. We consider multi-resolution data modeling and propose a
 108 novel reconstruction diffusion model to achieve the modeling ability of complex time series rela-
 109tionships.

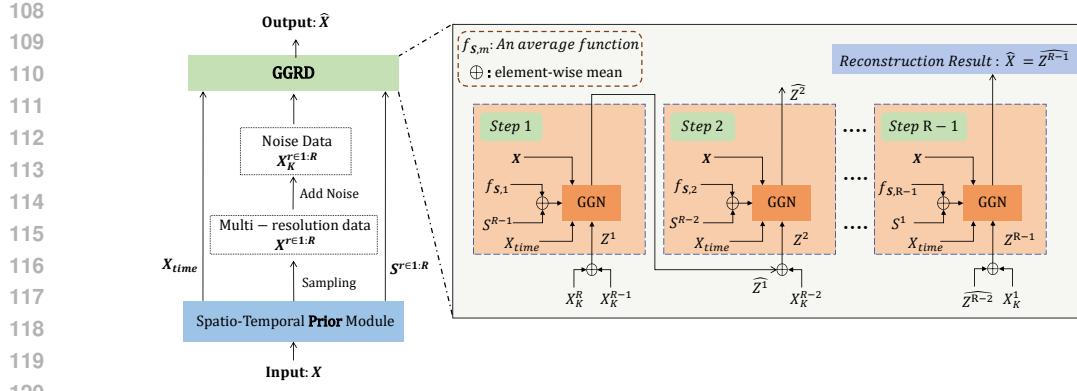


Figure 1: The overall structure of our model is summarized, mainly including STPM and GGRD.

3 PRELIMINARY

The goal of multivariate time series anomaly detection is to identify time steps at which the observed values deviate from normal behavior. Formally, let $\mathbf{X} \in \mathbb{R}^{T \times D}$ denote the input time series, where T is the number of time steps and D is the number of features. The anomaly labels are represented as $\mathbf{Y} \in \{0, 1\}^T$, where $y_t = 1$ indicates that the observation at time step t is anomalous and $y_t = 0$ otherwise. In the unsupervised setting considered here, labels \mathbf{Y} are unavailable during training and are used only for evaluation.

A common approach to unsupervised anomaly detection is to learn a generative model that captures the distribution of normal time series. The pipeline typically consists of three stages: (i) **Training**: fit a generative model $p_\theta(\mathbf{X})$ (or its conditional variant) using unlabeled historical data assumed to be mostly normal; (ii) **Reconstruction or prediction**: given a test input \mathbf{X} , obtain a reconstructed (or predicted) version $\hat{\mathbf{X}}$ using the generative model; (iii) **Scoring and thresholding**: compute an anomaly score s_t for each time step, e.g., via reconstruction error $s_t = \|\hat{\mathbf{x}}_t - \mathbf{x}_t\|_2^2$, and flag t as anomalous if s_t exceeds a learned or adaptive threshold.

In this work, we focus on diffusion-based models, which corrupt the input with noise and then learn to reconstruct it, using the reconstruction error as the anomaly indicator.

4 METHODOLOGY

4.1 OVERALL FRAMEWORK

Figure 1 illustrates the architecture of our proposed framework. The model consists of two major components: (i) a *Spatio-Temporal Prior Module (STPM)* that generates smooth multi-resolution representations and similarity-guided graph priors, and (ii) a *Graph-Guided Reconstruction Diffusion model (GGRD)* that performs one-step reconstruction guided by these priors.

Given an input time series \mathbf{X} , the STPM first generates a set of multi-resolution series $\{\mathbf{X}^{(r)}\}_{r=1}^R$ via sliding-window averaging, where R is the number of resolutions. For each resolution, we compute cosine similarity between feature dimensions to construct a similarity-guided graph tensor $\mathbf{S}^{(r)} \in \mathbb{R}^{P \times D \times D}$, which encodes the pairwise dependencies among features and time segments. The set $\{\mathbf{X}^{(r)}, \mathbf{S}^{(r)}\}$ is then corrupted by Gaussian noise following the forward diffusion process to obtain $\{\mathbf{X}_K^{(r)}\}$.

The GGN takes the noisy multi-resolution series as input and performs a *single-step* reconstruction to produce $\hat{\mathbf{X}}$. Unlike conventional diffusion models that require iterative denoising, GGN directly restores the clean signal in one step, significantly improving inference efficiency and mitigating error accumulation. The reconstruction is guided by graph priors through the *Graph-Guided Attention (GGA)* mechanism, enabling explicit modeling of dynamic cross-feature dependencies.

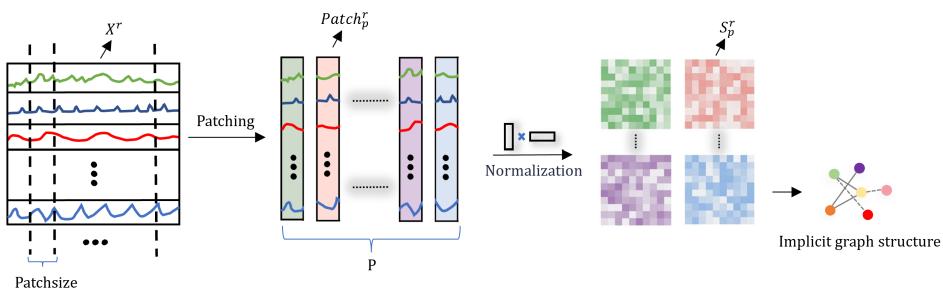


Figure 2: The process of obtaining similarity-guided Graph tensors.

4.2 SPATIO-TEMPORAL PRIOR MODULE

Smooth multi-resolution decomposition To obtain multi-resolution representations, we apply a sliding window of length rL with stride 1 over \mathbf{X} at each resolution $r \in [1, R]$. For each window, we take the mean of the included points to form the smoothed sequence $\mathbf{X}^{(r)} \in \mathbb{R}^{T \times D}$. This overlapping-window design preserves temporal smoothness and continuity, mitigating the staircase artifacts introduced by non-overlapping pooling.

Timestamp hard embedding Following D3R (Wang et al., 2023), we extract calendar-based features from absolute timestamps and build a fixed multi-granularity embedding $\mathbf{X}_{\text{time}} \in \mathbb{R}^{T \times d_{\text{time}}}$. We use $d_{\text{time}} = 5$ fields encoding *minute*, *hour*, *day*, *week*, and *month*. This hard-coded timestamp representation injects priors about periodic and seasonal patterns while preserving temporal ordering, and is subsequently consumed by the Time-Augment Encoder to enhance multi-level temporal modeling.

Similarity-guided graph construction As in Figure 2, for each resolution r , we divide $\mathbf{X}^{(r)}$ into P consecutive patches(time segments) along the temporal dimension:

$$P = \left\lceil \frac{T}{\text{patch size}} \right\rceil \quad (1)$$

For each patch p , we compute a cosine-similarity matrix:

$$\mathbf{S}_{ij}^{(r,p)} = \frac{\langle \mathbf{x}_{:,i}^{(r,p)}, \mathbf{x}_{:,j}^{(r,p)} \rangle}{\|\mathbf{x}_{:,i}^{(r,p)}\|_2 \|\mathbf{x}_{:,j}^{(r,p)}\|_2}, \quad (2)$$

where $\mathbf{x}_{:,i}^{(r,p)}$ denotes the series of feature i in patch p . The collection $\{\mathbf{S}^{(r,p)}\}$ forms a dynamic graph prior that captures time-varying, multi-resolution feature dependencies.

4.3 DIFFUSION FORWARD PROCESS

Following denoising diffusion probabilistic models (DDPM) (Ho et al., 2020), we gradually inject Gaussian noise into each $\mathbf{X}^{(r)}$ over K steps via a forward Markov chain:

$$q(\mathbf{X}_k^{(r)} | \mathbf{X}_{k-1}^{(r)}) = \mathcal{N}(\sqrt{1 - \beta_k} \mathbf{X}_{k-1}^{(r)}, \beta_k \mathbf{I}), \quad k = 1, \dots, K, \quad (3)$$

where $\beta_k \in (0, 1)$ is a variance schedule. In practice, $\mathbf{X}_k^{(r)}$ can be sampled in closed form at any k , allowing efficient generation of the fully-noised sample $\mathbf{X}_K^{(r)}$ in one step.

4.4 GRAPH-GUIDED NETWORK (GGN)

The GGN serves as the backbone for one-step reconstruction. It operates in a coarse-to-fine manner over R resolution levels and $R - 1$ steps. At step m , we first fuse the noisy input $\mathbf{X}_K^{(R-m)}$ with the previous reconstruction $\hat{\mathbf{Z}}^{(m-1)}$ through a channel-wise concatenation and a 1D projection layer:

$$\mathbf{Z}^{(m)} = \text{Proj}([\mathbf{X}_K^{(R-m)}; \hat{\mathbf{Z}}^{(m-1)}]), \quad \hat{\mathbf{Z}}^{(0)} \equiv \mathbf{X}_K^{(R)}, \quad m = 1, \dots, R - 1. \quad (4)$$

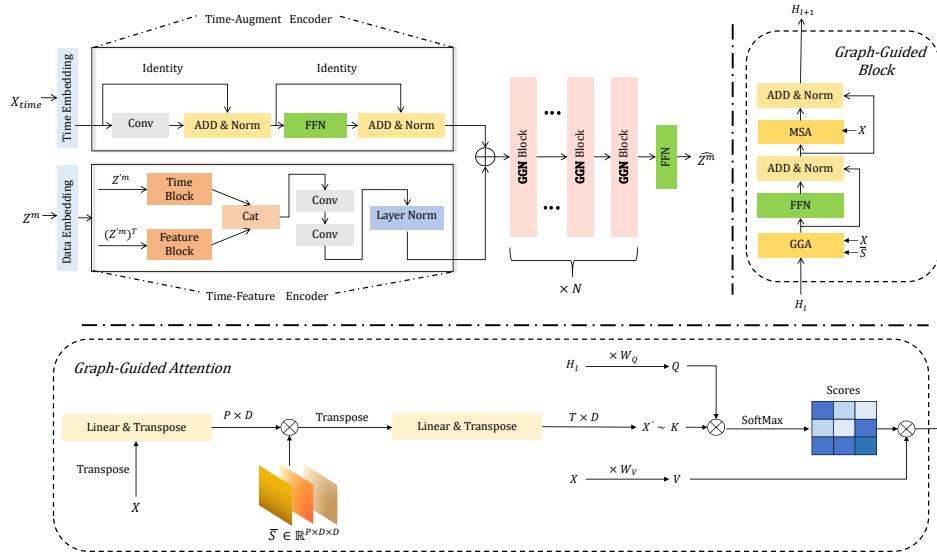


Figure 3: The overall structure of GGNetwork mainly includes Time-Feature Encoder, Time-Augment Encoder and some GGN Blocks. The structures of GGN Block and GGA are respectively to the right and below the dotted line.

This fusion passes information across resolutions, enabling progressively refined reconstruction.

Dual encoders The structure of GGN is presented in Figure 3. The fused sequence $\mathbf{Z}^{(m)}$ is embedded into a d_{model} -dimensional space and processed by two parallel encoders (the detailed structure is in the Appendix C.1): (i) the *Time-Feature Encoder* \mathcal{E}_{tf} , which extracts temporal dependencies and feature-wise interactions using self-attention and local convolutions, and (ii) the *Time-Augment Encoder* \mathcal{E}_{ta} , which projects the timestamp embedding \mathbf{X}_{time} into the same space. Their outputs are added elementwise:

$$\mathbf{H}_0^{(m)} = \mathcal{E}_{\text{tf}}(\mathbf{Z}^{(m)}) + \mathcal{E}_{\text{ta}}(\mathbf{X}_{\text{time}}). \quad (5)$$

Stacked GGN blocks with GGA The hidden state is refined through N stacked GGN blocks. Each block contains (as shown in Figure 3): (a) a **Graph-Guided Attention (GGA)** module that incorporates the similarity-guided graph prior $\bar{\mathbf{S}}^{(m)}$, which is obtained by a specially defined fusion function f (see Appendix C.2 for details), representing the information fusion between \mathbf{S} of the current step and \mathbf{S} of the previous step, and (b) a temporal self-attention layer that maintains global temporal context.

Given queries \mathbf{Q} , values \mathbf{V} projected from $\mathbf{H}_{\ell}^{(m)}$, $\text{Linear\&T}(\cdot)$ represents passing through the linear layer and immediately transposing. GGA modifies the attention computation by applying the graph prior to the keys:

$$\mathbf{K} = \text{Linear\&T}_2(\text{Linear\&T}_1(\mathbf{X}^T) \odot \bar{\mathbf{S}}^{(m)}) \quad (6)$$

$$\mathbf{A}^{(m)} = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}}\right), \quad \text{GGA}(\mathbf{H}) = \mathbf{A}^{(m)}\mathbf{V}. \quad (7)$$

This biases attention weights toward feature pairs with higher similarity in $\bar{\mathbf{S}}^{(m)}$, explicitly capturing dynamic cross-feature dependencies. The GGA output is added to the temporal attention output and passed through a lightweight feed-forward layer with residual connections and layer normalization.

Stage output After N blocks, the hidden representation is projected back to the input dimension to yield the reconstruction:

$$\hat{\mathbf{Z}}^{(m)} = \text{Head}(\mathbf{H}_N^{(m)}) \in \mathbb{R}^{T \times D}, \quad (8)$$

which is then passed to the next resolution level as input. The final reconstruction $\hat{\mathbf{X}} = \hat{\mathbf{Z}}^{(R-1)}$ is used for anomaly scoring.

270 4.5 TRAINING OBJECTIVE AND ANOMALY SCORING
271272 We train GGRD to minimize the mean squared error (MSE) between the original and reconstructed
273 series:

274
$$\mathcal{L} = \frac{1}{TD} \sum_{t=1}^T \sum_{d=1}^D (x_{t,d} - \hat{x}_{t,d})^2. \quad (9)$$

275

276 During inference, the anomaly score at time step t is defined as the reconstruction error $s_t = \|\hat{x}_t -$
277 $x_t\|_2^2$. Anomalies are identified by applying an adaptive thresholding method SPOT (Siffer et al.,
278 2017) to $\{s_t\}_{t=1}^T$.
279280 5 EXPERIMENTS
281282 5.1 EXPERIMENTAL SETTINGS
283284 **Datasets** This paper mainly uses five public time series datasets, including PSM (Pooled Server
285 Metrics) (Abdulaal et al., 2021), SMD (Server Machine Dataset) (Su et al., 2019), and SWaT (Secure
286 Water Treatment) (Mathur & Tippenhauer, 2016). The training set and validation do not contain la-
287 bels, only the test set data has labels. More descriptions of the datasets are in Appendix A.1.
288289 **Metrics** The experiments adopt Precision, Recall and F1 score as the main evaluation metrics.
290 Unlike the point adjustment method commonly used in most existing studies (Chen et al., 2023; Xiao
291 et al., 2023; Wen et al., 2025), we use the Affiliation based (Huet et al., 2022) strategy to calculate the
292 indicators. In a continuous anomaly interval, as long as any point is predicted, the point adjustment
293 strategy will be regarded as the entire interval being detected, thereby significantly improving TP
294 and masking the deficiency of the model in anomaly localization, which is prone to cause false
295 performance improvement (Wang et al., 2023). In contrast, the method based on Affiliation measures
296 the matching relationship between the predicted anomalies and the true anomaly intervals, which
297 more objectively reflects the performance of the model in interval-level anomaly detection. It can
298 avoid excessive bias towards long interval anomalies and thus obtain more reasonable results.
299300 **Experiment setup** GGRD uses Adam as the optimizer, with the learning rate set to $1e - 4$ and
301 weight decay set to $1e - 4$. For the unlabeled data in each dataset, we select 80% as the training set,
302 20% as the validation set, and the labeled data as the test set. For all datasets, the size of batchsize
303 is set to 8 and the training epoch is 10. Go through 1,000 steps from 0.0001 to 0.02. The number of
304 GGN blocks N is set to 4 in the SMD dataset and 2 in the rest of the datasets. The number of noise
305 additions K to the original time series data is 500. All the experiments of GGRD were carried out
306 under the Linux system, Pytorch, and a total of 4 NVidia A100 GPUs were used.
307308 **Baseline** The baseline methods we selected cover multiple paradigms, including probabilistic
309 modeling, linear transformation, deep neural networks, and Transformer methods, etc. For specific
310 descriptions, please refer to the Appendix A.2.
311312 5.2 DETECTION RESULTS
313314 We conducted experiments on the GGRD model and baselines on multiple real-world datasets, and
315 the experimental results are shown in Table 1. The GGRD model achieved the best performance on
316 the non-stationary datasets PSM and SMD, which were 3.28% higher than the second-best results
317 (from 0.8100 to 0.8428) and 1.44% higher (from 0.9238 to 0.9382), respectively. The mean of the
318 overall F1 on all datasets, Avg-F1, was 0.2% higher than the second-best result (from 0.8420 to
319 0.8440). However, on the SWaT dataset, GGRD performs slightly worse than D3R and MODEM.
320 This is mainly because SWaT data usually contains short-term burst exceptions, while the modeling
321 mechanisms of D3R and MODEM are more sensitive to such instantaneous changes. In contrast,
322 the advantage of GGRD lies in its ability to effectively capture the dependencies between non-
323 stationary features and complex features. Therefore, on time series datasets such as PSM and SMD,
324 which have complex dynamic behaviors and cross-feature correlations, GGRD demonstrates more
325 robust and superior anomaly detection performance, but it is not the case in SWaT (see the details
326 in Appendix A.4). These results indicate that although there is a slight gap in specific short-term
327 sudden abnormal scenarios, GGRD still has significant advantages when dealing with real complex
328 environments and high-dimensional multi-variable data.
329

324
 325 Table 1: The comparison test results with other models on three real-world datasets show that our
 326 GGRD leads on most datasets. The best F1 score is marked in bold, and the second best is marked
 327 with an underline. Avg-F1 represents the average F1 score.

328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343	Method	PSM			SMD			SWaT			Avg-F1
		P	R	F1	P	R	F1	P	R	F1	
COPOD	0.7602	0.3175	0.4479	0.6676	0.1366	0.2268	0.9876	0.1180	0.2108	0.2952	
ECOD	0.7460	0.3384	0.4656	0.7398	0.1615	0.2651	0.9761	0.1151	0.2059	0.3122	
OCSVM	0.8761	0.4744	0.6155	0.0000	0.0000	0.0000	0.6196	0.7558	0.6810	0.4322	
CBLOF	0.5990	0.9845	0.7449	0.8667	0.3352	0.4834	0.6308	0.7091	0.6677	0.6320	
HBOS	1.0000	0.0654	0.1228	0.5628	0.8007	0.6610	0.5771	0.8049	0.6722	0.4853	
IForest	1.0000	0.0335	0.0648	1.0000	0.0937	0.1713	0.6127	0.6280	0.6203	0.2855	
LODA	0.9266	0.4017	0.5605	0.5902	0.6618	0.6240	0.6117	0.7014	0.6535	0.6127	
VAE	0.6221	0.8772	0.7280	0.8209	0.4349	0.5686	0.6355	0.7218	0.6759	0.6575	
DeepSVDD	0.7405	0.5064	0.6015	0.6498	0.6477	0.6488	0.5911	0.9353	0.7244	0.6582	
LSTM-AE	0.7511	0.7586	0.7548	0.8496	0.4349	0.5753	0.6018	0.7219	0.6564	0.6622	
MTAD-GAT	0.7990	0.6014	0.6863	0.8590	0.6769	0.7571	0.6590	0.7751	0.7123	0.7186	
TFAD	0.7914	0.7163	0.7520	0.5632	0.9783	0.7149	0.6038	0.8196	0.6953	0.7207	
Anomaly Transformer	0.5201	0.8504	0.6455	1.0000	0.0319	0.0619	0.5541	0.5994	0.5759	0.4278	
Diff-AD	0.5564	0.7674	0.6450	0.5014	0.9093	0.6464	0.5183	0.7979	0.6284	0.6399	
D3R	0.6294	0.9619	0.7609	0.7715	0.9926	0.8682	0.7206	0.8529	0.7812	0.8034	
Imdiffusion	0.7556	0.8784	0.8100	0.9605	0.5271	0.6741	0.8387	0.2058	0.3297	0.6046	
MODEM	0.7348	0.8755	0.7990	0.8918	0.9582	0.9238	0.7436	0.8732	0.8032	0.8420	
ours	0.8827	0.8064	0.8428	0.9812	0.8988	0.9382	0.6529	0.8840	0.7511	0.8440	

346 5.3 ABLATION STUDIES

347 To verify the role of each module in the model, we conducted systematic ablation experiments on
 348 multiple datasets and compared the results with those of the complete model, and the ablation results
 349 are shown in Table 2.

351
 352 Table 2: The ablation experiment results reported the best F1.

354 355 356 357 358 359 360 361 362 363 364 365 366 367 368 369 370 371 372 373 374 375 376 377	354 355 356 357 358 359 360 361 362 363 364 365 366 367 368 369 370 371 372 373 374 375 376 377	w/o time-feature	w/o time-augment	w/o gga	w/o sliding window	w/o timestamp
SMD	0.9382	0.8922	0.9271	0.8515	0.9181	0.8931
PSM	0.8428	0.8078	0.8168	0.7721	0.8089	0.8231
SWaT	0.7511	0.7267	0.7331	0.7117	0.7377	0.7404
average	0.8440	0.8089	0.8257	0.7784	0.8216	0.8189

Spatio-Temporal Prior Module This module is mainly used to generate hard-coded timestamps, multi-resolution data and graph structures. In the ablation experiment, we removed the timestamps, denoted as **w/o timestamp**, and replaced the generation method of multi-resolution data from the proposed sliding window to non-overlapping average pooling(Zhong et al., 2025), denoted as **w/o sliding window**. The experimental results in Table 2 show that after removing the timestamp on the SMD dataset, the F1 value decreased from 0.9382 to 0.8931, and the performance decreased by 4.51%. Similarly, after replacing the sliding window with non-overlapping average pooling, the F1 of all three datasets decreased, indicating that timestamps and efficient multi-resolution modeling methods are of great significance for capturing the dynamic patterns of time series.

Reconstruction of Multivariate Time Series Module This module is the core part of the GGRD proposed in this paper, including Time-Feature Encoder, Time-Augment Encoder and GGA. We remove one of the components respectively, denoted as **w/o time-feature**, **w/o time-augment** and **w/o gga**. The experimental results show that on multiple datasets, the F1 value of the complete model is always superior to that of the ablated version. For example, in PSM, the F1 value of the original model was 0.8428, but it decreased to 0.8078 after removing the time-feature, a reduction of 3.5%. In SMD, the best F1 value decreased from 0.9382 to 0.8515 after removing GGA, with the largest performance decline and other datasets also have a significant impact, further verifying the key role of GGA in capturing cross-dimensional dependencies and dynamic features.

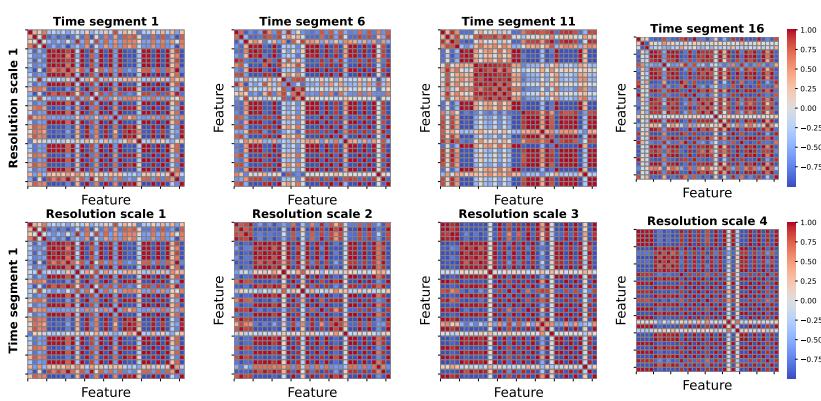


Figure 4: Visualization of Similarity-guided Graph Tensor \mathbf{S} on SMD. The four matrices in the first row represent the changes at different time segments when the resolution scale is 1. The second row represents the changes in different resolution scales within the same time segment (time segment 1).

5.4 EFFECTIVENESS ANALYSIS

Similarity-Guided Graph Tensor To further illustrate the effectiveness of GGA, in the previous paper, Similarity-guided Graph Tensor \mathbf{S} was constructed based on cosine Similarity, and the \mathbf{S} was visualized at four resolution scales and any four time segments of the SMD dataset. As shown in the Figure 4, the similarity between features shows significant differences at different time segments and resolutions, indicating that the feature dependency relationship is dynamic and multi-scale. For instance, when the resolution scale is 1(the first row), the similarity between features at different time segments varies the most. Similarly, when the time segmen is fixed (the second row), the feature similarity also shows significant differences at different resolutions.Compared with the methods that do not explicitly consider this information, \mathbf{S} can provide additional correlation information for multi-head self-attention, thereby enabling GGA to better capture cross-dimensional dependencies and improve the expressive ability of the model in anomaly detection.

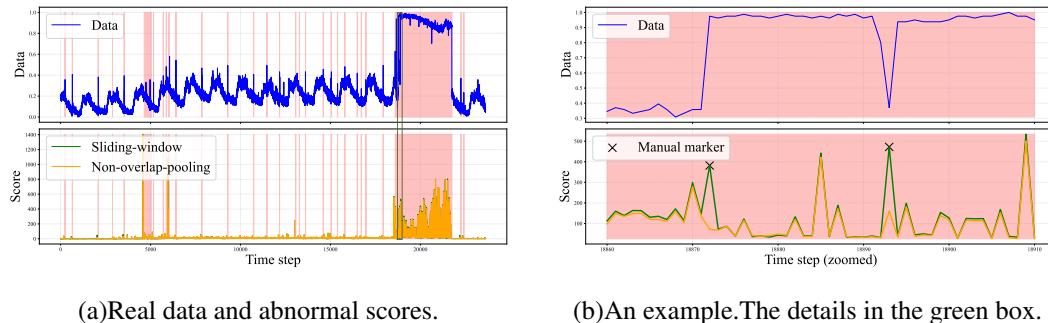


Figure 5: Real data and the abnormal scores corresponding to the two methods respectively. Both methods obtain abnormal scores under their own optimal models. The higher the abnormal score, the easier it is to be detected as an anomaly. The two types of data show differences at non-stationary points.

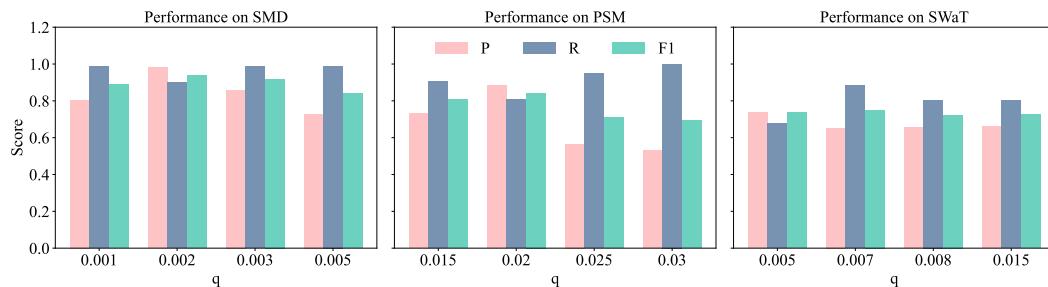
Obtain multi-resolution data using sliding windows When initially constructing the sliding window, we reduce the information loss between adjacent times. We use overlapping sliding Windows instead of non-overlapping average pooling. In Figure 5(a), the upper part is the original data of a feature in the SMD dataset, and the red transparent background represents the anomaly. The lower part shows the magnitudes of the anomaly scores corresponding to the two data generation methods. And the Figure 5(b) is an example. The black manually marked areas show obvious differences

432 because the real data of the adjacent timestamps corresponding to the two places are changing drastically.
 433 The sliding window method, being smoother compared to the latter method and less likely to
 434 lose more information, can obtain higher outliers, while lower outliers may lead to missed detections
 435 of anomalies. Similarly, when the label is normal, non-overlapping methods that are not stationary
 436 may also receive higher scores and be misjudged as abnormal.
 437

438 5.5 HYPERPARAMETER ANALYSIS

440 **Hyperparameters of GGRD** Here we mainly investigated the influence of the number of resolution
 441 categories R , the number of GGN modules N , *patch size*, and the initial window length L on
 442 the model performance. The specific experimental results can be found in the Appendix A.3.

443 **Anomaly detection threshold** Given a probability q , the SPOT algorithm can automatically obtain
 444 the detection threshold by using the abnormal scores of the training data and the test set. The
 445 following figure presents the metrics of different datasets under different q . In this experiment, we
 446 evaluated using the SPOT algorithm on three datasets (SWaT, PSM, and SMD), and observed the
 447 changes in Precision, Recall, and F1-score by adjusting the threshold q . The results are presented in
 448 Figure 6. The results show that as the q value increases, it is often accompanied by an increase in
 449 Recall, and Precision usually decreases, and q represents the proportion of outliers among extreme
 450 points (Siffer et al., 2017). This trend is particularly evident in the SWAT and PSM datasets, with
 451 the optimal F1-scores appearing at $q = 0.007$ and $q = 0.02$, respectively. On the SMD dataset,
 452 the model as a whole demonstrates high stability and robustness, with relatively stable Precision and
 453 Recall. Therefore, it is necessary to seek a trade-off between Precision and Recall to achieve the best
 454 F1. Overall, the SPOT algorithm can achieve high anomaly recognition results on different types of
 455 time series data and strike a good balance between accuracy and recall by reasonably selecting q ,
 456 demonstrating its applicability and reliability in multiple scenarios.
 457



466 Figure 6: The influence of different q on Model performance
 467

468 6 CONCLUSION AND LIMITATION

469 **Conclusion.** This paper proposes a Graph Guided Reconstruction Model (GGRD) to address
 470 the deficiencies of existing methods in anomaly detection of multivariate time series. Specifically,
 471 we generate multi-resolution data through a sliding window, reducing the resolution while
 472 retaining the original information features as much as possible. Furthermore, in this study, the
 473 reconstruction structure of GGRD was carefully designed and the correlation between features was
 474 considered. Similarity-Guided Graph Tensors were introduced into GGN to effectively guide
 475 feature interaction and improve the accuracy and robustness of data modeling. A large number of
 476 experiments have shown that GGRD outperforms existing anomaly detection methods on various
 477 datasets.
 478

479 **Limitations.** The length of the time segment in GGRD is fixed (depending on the *patch size* and
 480 T), but in fact, the correlation between time series features does not remain constant over a fixed-
 481 length time segment. That is to say, the duration of a certain correlation situation varies and is often
 482 highly uncertain. Therefore, a method for dynamically obtaining the graph structure prior is needed
 483 to enhance the detection ability of irregular fluctuations and short-term sudden anomalies. At the
 484 same time, in addition to cosine similarity (even though it is simple and effective), other forms can
 485 be considered for the graph structure tensor to further improve the modeling ability of time series.
 486

486 REFERENCES
487

488 Ahmed Abdulaal, Zhuanghua Liu, and Tomer Lancewicki. Practical approach to asynchronous
489 multivariate time series anomaly detection and localization. In *Proceedings of the 27th ACM
490 SIGKDD conference on knowledge discovery & data mining*, pp. 2485–2494, 2021.

491 Yuhang Chen, Chaoyun Zhang, Minghua Ma, Yudong Liu, Ruomeng Ding, Bowen Li, Shilin He,
492 Saravan Rajmohan, Qingwei Lin, and Dongmei Zhang. Imdiffusion: Imputed diffusion models
493 for multivariate time series anomaly detection. *arXiv preprint arXiv:2307.00754*, 2023.

494 Ailin Deng and Bryan Hooi. Graph neural network-based anomaly detection in multivariate time
495 series. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pp. 4027–
496 4035, 2021.

497 Dave Epstein, Allan Jabri, Ben Poole, Alexei Efros, and Aleksander Holynski. Diffusion self-
498 guidance for controllable image generation. *Advances in Neural Information Processing Systems*,
499 36:16222–16239, 2023.

500 Xinyao Fan, Yueying Wu, Chang Xu, Yuhao Huang, Weiqing Liu, and Jiang Bian. Mg-tsd:
501 Multi-granularity time series diffusion models with guided learning process. *arXiv preprint
502 arXiv:2403.05751*, 2024.

503 Markus Goldstein and Andreas Dengel. Histogram-based outlier score (hbos): A fast unsupervised
504 anomaly detection algorithm. *KI-2012: poster and demo track*, 1:59–63, 2012.

505 Zengyou He, Xiaofei Xu, and Shengchun Deng. Discovering cluster-based local outliers. *Pattern
506 recognition letters*, 24(9-10):1641–1650, 2003.

507 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in
508 neural information processing systems*, 33:6840–6851, 2020.

509 Alexis Huet, Jose Manuel Navarro, and Dario Rossi. Local evaluation of time series anomaly detec-
510 tion algorithms. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery
511 and Data Mining*, pp. 635–645, 2022.

512 Ming Jin, Huan Yee Koh, Qingsong Wen, Daniele Zambon, Cesare Alippi, Geoffrey I Webb, Irwin
513 King, and Shirui Pan. A survey on graph neural networks for time series: Forecasting, classifi-
514 cation, imputation, and anomaly detection. *IEEE Transactions on Pattern Analysis and Machine
515 Intelligence*, 2024.

516 Tung Kieu, Bin Yang, and Christian S Jensen. Outlier detection for multidimensional time series
517 using deep neural networks. In *2018 19th IEEE international conference on mobile data man-
518 agement (MDM)*, pp. 125–134. IEEE, 2018.

519 Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint
520 arXiv:1312.6114*, 2013.

521 Disen Lan, Guibin Zhang, and Rongjin Guo. Diffusion graph model for time series anomaly detec-
522 tion via anomaly-aware graph sparsification and augmentation. In *Companion Proceedings of the
523 ACM on Web Conference 2025*, pp. 2207–2214, 2025.

524 Zheng Li, Yue Zhao, Nicola Botta, Cezar Ionescu, and Xiyang Hu. Copod: copula-based outlier
525 detection. In *2020 IEEE international conference on data mining (ICDM)*, pp. 1118–1123. IEEE,
526 2020.

527 Zheng Li, Yue Zhao, Xiyang Hu, Nicola Botta, Cezar Ionescu, and George H Chen. Ecod: Unsu-
528 pervised outlier detection using empirical cumulative distribution functions. *IEEE Transactions
529 on Knowledge and Data Engineering*, 35(12):12181–12193, 2022.

530 Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation forest. In *2008 eighth ieee interna-
531 tional conference on data mining*, pp. 413–422. IEEE, 2008.

532 Guanxiong Luo, Shoujin Huang, and Martin Uecker. Autoregressive image diffusion: Generation
533 of image sequence and application in mri. *Advances in Neural Information Processing Systems*,
534 37:129094–129119, 2024.

540 Pankaj Malhotra, Lovekesh Vig, Gautam Shroff, Puneet Agarwal, et al. Long short term memory
 541 networks for anomaly detection in time series. In *Proceedings*, volume 89, pp. 94, 2015.
 542

543 Aditya P Mathur and Nils Ole Tippenhauer. Swat: A water treatment testbed for research and
 544 training on ics security. In *2016 international workshop on cyber-physical systems for smart*
 545 *water networks (CySWater)*, pp. 31–36. IEEE, 2016.

546 Tomáš Pevný. Loda: Lightweight on-line detector of anomalies. *Machine Learning*, 102(2):275–
 547 304, 2016.

548 Sridhar Ramaswamy, Rajeev Rastogi, and Kyuseok Shim. Efficient algorithms for mining outliers
 549 from large data sets. In *Proceedings of the 2000 ACM SIGMOD international conference on*
 550 *Management of data*, pp. 427–438, 2000.

551 Lukas Ruff, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexan-
 552 der Binder, Emmanuel Müller, and Marius Kloft. Deep one-class classification. In *International*
 553 *conference on machine learning*, pp. 4393–4402. PMLR, 2018.

554 Bernhard Schölkopf, John C Platt, John Shawe-Taylor, Alex J Smola, and Robert C Williamson.
 555 Estimating the support of a high-dimensional distribution. *Neural computation*, 13(7):1443–1471,
 556 2001.

557 Lifeng Shen, Weiyu Chen, and James Kwok. Multi-resolution diffusion models for time series
 558 forecasting. In *The Twelfth International Conference on Learning Representations*, 2024.

559 Mei-Ling Shyu, Shu-Ching Chen, Kanoksri Sarinnapakorn, and LiWu Chang. A novel anomaly
 560 detection scheme based on principal component classifier. 2003.

561 Alban Siffer, Pierre-Alain Fouque, Alexandre Termier, and Christine Largouet. Anomaly detection
 562 in streams with extreme value theory. In *Proceedings of the 23rd ACM SIGKDD international*
 563 *conference on knowledge discovery and data mining*, pp. 1067–1075, 2017.

564 Ya Su, Youjian Zhao, Chenhao Niu, Rong Liu, Wei Sun, and Dan Pei. Robust anomaly detection for
 565 multivariate time series through stochastic recurrent neural network. In *Proceedings of the 25th*
 566 *ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 2828–2837,
 567 2019.

568 Chengsen Wang, Zirui Zhuang, Qi Qi, Jingyu Wang, Xingyu Wang, Haifeng Sun, and Jianxin Liao.
 569 Drift doesn't matter: Dynamic decomposition with diffusion reconstruction for unstable multi-
 570 variate time series anomaly detection. *Advances in neural information processing systems*, 36:
 571 10758–10774, 2023.

572 Mi Wen, Zehui Chen, Yun Xiong, and Yichuan Zhang. Lgat: A novel model for multivariate
 573 time series anomaly detection with improved anomaly transformer and learning graph structures.
 574 *Neurocomputing*, 617:129024, 2025.

575 Qingsong Wen, Tian Zhou, Chaoli Zhang, Weiqi Chen, Ziqing Ma, Junchi Yan, and Liang Sun.
 576 Transformers in time series: A survey. *arXiv preprint arXiv:2202.07125*, 2022.

577 Chunjing Xiao, Zehua Gou, Wenxin Tai, Kunpeng Zhang, and Fan Zhou. Imputation-based time-
 578 series anomaly detection with conditional weight-incremental diffusion models. In *Proceedings*
 579 *of the 29th ACM SIGKDD conference on knowledge discovery and data mining*, pp. 2742–2751,
 580 2023.

581 Jiehui Xu, Haixu Wu, Jianmin Wang, and Mingsheng Long. Anomaly transformer: Time series
 582 anomaly detection with association discrepancy. *arXiv preprint arXiv:2110.02642*, 2021.

583 Xingyu Xu and Yuejie Chi. Provably robust score-based diffusion posterior sampling for plug-
 584 and-play image reconstruction. *Advances in Neural Information Processing Systems*, 37:36148–
 585 36184, 2024.

586 Asrul H Yaacob, Ian KT Tan, Su Fong Chien, and Hon Khi Tan. Arima based network anomaly
 587 detection. In *2010 Second international conference on communication software and networks*,
 588 pp. 205–209. IEEE, 2010.

594 Zahra Zamanzadeh Darban, Geoffrey I Webb, Shirui Pan, Charu Aggarwal, and Mahsa Salehi. Deep
595 learning for time series anomaly detection: A survey. *ACM Computing Surveys*, 57(1):1–42, 2024.
596

597 Chaoli Zhang, Tian Zhou, Qingsong Wen, and Liang Sun. Tfad: A decomposition time series
598 anomaly detection architecture with time-frequency analysis. In *Proceedings of the 31st ACM*
599 *international conference on information & knowledge management*, pp. 2497–2507, 2022.

600 Liang Zhang, Jianping Zhu, Guangjie Han, Bo Jin, Pengfei Wang, and Xiaopeng Wei. Self-
601 supervised disentangled representation learning for time series anomaly detection. *IEEE Internet*
602 *of Things Journal*, 2025.

603

604 Hang Zhao, Yujing Wang, Juanyong Duan, Congrui Huang, Defu Cao, Yunhai Tong, Bixiong Xu,
605 Jing Bai, Jie Tong, and Qi Zhang. Multivariate time-series anomaly detection via graph attention
606 network. In *2020 IEEE international conference on data mining (ICDM)*, pp. 841–850. IEEE,
2020.

607

608 Guojin Zhong, Jin Yuan, Zhiyong Li, Long Chen, et al. Multi-resolution decomposable diffusion
609 model for non-stationary time series anomaly detection. In *The Thirteenth International Confer-
610 ence on Learning Representations*, 2025.

611

612

613

614

615

616

617

618

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648 **A EXPERIMENTAL DETAILS**
649650 **A.1 DATASETS**
651652 Just like D3R(Wang et al., 2023),we only retained the continuous variables in the data for the exper-
653 iment and provided the statistical data of these datasets in the Table 3.
654655 Table 3: Dataset statistical description(AR represents the Rate of Anomalies in the test set).
656657

658 Dataset	659 Training size	660 Testing size	661 Dimensions	662 Frequency	663 AR(%)
SMD	23688	23689	33	1 minute	15.7
PSM	132481	87841	25	1 minute	27.8
SWAT	6840	7500	25	1 minute	12.6

664 **A.2 BASELINE**
665666 Some of the baselines selected for the experiment adopted the results of D3R and were implemented
667 following the configuration recommended in the original paper. Apart from the first one, the other
668 baselines are based on our operation and run with the recommended configuration in the original
669 paper.
670671 **COPOD**(Li et al., 2020) COPOD innovatively utilizes copula to construct an empirical distribution
672 and calculate the tail probability, thereby achieving a parameterless, interpretable and efficient
673 anomaly detection method.
674675 **ECOD**(Li et al., 2022) Based on the empirical cumulative distribution function to estimate the tail
676 probability, a parameter-free, easily interpretable and highly efficient and scalable anomaly detection
677 method has been implemented, which significantly outperforms existing methods on large-scale and
678 high-dimensional data.
679680 **OCSVM**(Schölkopf et al., 2001) This method extends the support vector machine to unsupervised
681 scenarios, constructs a discriminant function through the kernel function, divides the input space into
682 high-probability subsets and their complements to ensure that the probability of new samples falling
683 into this subset is controlled, and solves the extended coefficients through quadratic programming
684 to achieve efficient anomaly detection.
685686 **CBLOF**(He et al., 2003) It is an unsupervised anomaly detection method based on local clusters,
687 which identifies anomalies by evaluating the behavioral significance of data points in their respective
688 clusters and can effectively discover outliers with physical or statistical significance.
689690 **HBOS**(Goldstein & Dengel, 2012) It is an unsupervised anomaly detection method based on his-
691 tograms. By assuming feature independence, it achieves linear time scoring and can efficiently
692 identify global anomalies, but it performs weakly in local anomaly detection.
693694 **IForest**(Liu et al., 2008) Unsupervised anomaly detection methods based on the idea of isolation
695 achieve linear time complexity and low memory consumption by explicitly isolating outliers and
696 using sub-sampling, and exhibit excellent performance on large-scale, high-dimensional or datasets
697 with irrelevant features.
698699 **LODA**(Pevný, 2016) LODA is an unsupervised anomaly detection method based on weak detector
700 integration, which can efficiently handle large-scale or streaming data, deal with missing variables
701 and concept drift, and simultaneously identify the characteristics of anomaly occurrence. It outper-
702 forms many existing methods in terms of speed and accuracy.
703704 **VAE**(Kingma & Welling, 2013) By reparameterizing the variational lower bound and approxi-
705 mating the model, efficient learning and inference of directed probabilistic models with continuous
706 latent variables have been achieved, and effective optimization can be carried out even in the case of
707 posteriorly unresolvable and large-scale datasets.
708

702 **DeepSVDD**(Ruff et al., 2018) It is a deep method trained with anomaly detection as the goal.
 703 By directly optimizing the anomaly detection target in the neural network, it shows good results in
 704 image benchmark datasets and adversarial sample detection.

705 **LSTM-AE**(Kieu et al., 2018) This model generates statistical features for time series and recon-
 706 structs them using autoencoders to capture representative patterns, thereby identifying outliers that
 707 deviate from the reconstruction. At the same time, it combines convolutional and LSTM networks
 708 as well as context information to improve the accuracy of anomaly detection.

710 **MTAD-GAT**(Zhao et al., 2020) By capturing the dependencies of multivariate time series in the
 711 time and feature dimensions through the parallel graph attention layer, and combining prediction
 712 and reconstruction optimization, efficient anomaly detection is achieved, while also having good
 713 interpretability and anomaly diagnosis capabilities.

714 **TFAD**(Zhang et al., 2022) TFAD utilizes time-frequency joint analysis and enhances anomaly
 715 detection performance and interpretability through time series decomposition and data augmentation
 716 mechanisms.

717 **Anomaly Transformer**(Xu et al., 2021) By using the self-attention mechanism to calculate the
 718 correlation differences of points in the time series, and amplifying the distinguishability between
 719 normal and abnormal through abnormal attention and the minim-to-maximum strategy, advanced
 720 performance has been achieved in various unsupervised time series anomaly detection tasks.

721 **Diff-AD**(Xiao et al., 2023) By using the density ratio to select normal observations and com-
 722 bining with denoising diffusion interpolation with increasing conditional weights and multi-scale
 723 state space modeling, abnormal concentration scenarios can be effectively handled, achieving stable
 724 multi-variable time series anomaly detection.

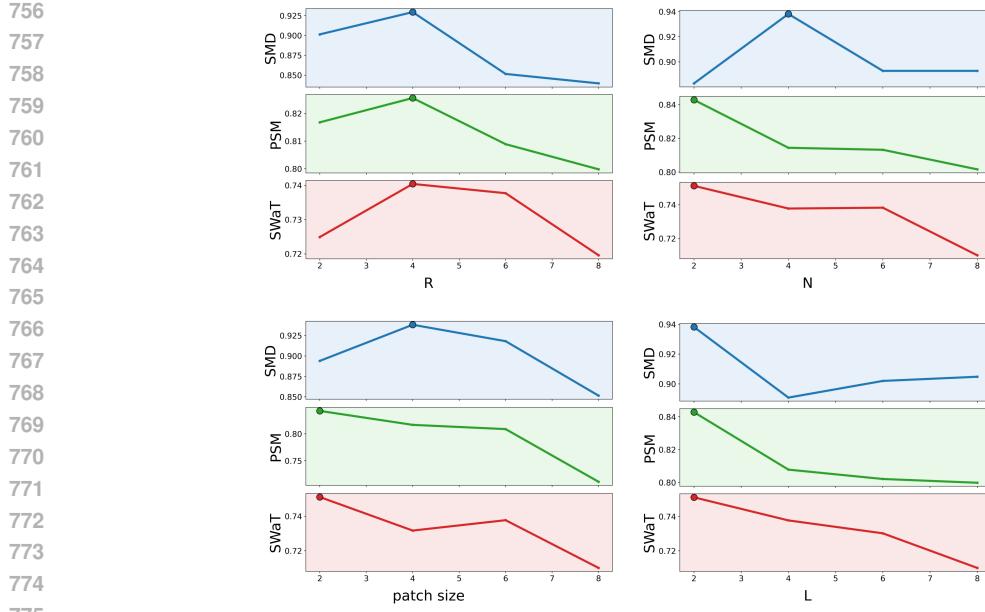
725 **D3R**(Wang et al., 2023) By combining the dynamic decomposition of data-time hybrid atten-
 726 tion with noise diffusion reconstruction, the time series is split into stable components and trend
 727 components, and the non-stationary multivariable time series is processed end-to-end, significantly
 728 reducing the false alarm rate caused by drift and improving the detection performance.

729 **Imdiffusion**(Chen et al., 2023) Combining time series interpolation with diffusion models,
 730 anomaly prediction is carried out through stepwise denoising to generate signals, accurately model-
 731 ing time series and cross-variable dependencies.

733 **MODEM**(Zhong et al., 2025) By jointly modeling non-stationary time series through multi-
 734 resolution diffusion and frequency-domain augmented networks, cross-resolution correlations are
 735 captured during the coarse-to-fine generation process to achieve precise anomaly detection of com-
 736 plex time series patterns.

737 738 A.3 HYPERPARAMETER OF GGRD ANALYSIS

741 The results are shown in Figure 7, and the following conclusions can be drawn: (1) As R in-
 742 creases, performance shows certain differences on different datasets. Overall, better results are
 743 usually achieved when $R = 4$. For example, on the SMD dataset, the F1-score reached 0.9382
 744 when $R = 4$, and the performance declined when it increased further. The lower resolution ($R = 2$)
 745 leads to insufficient expressive power of the model and a decline in discrimination. However, ex-
 746 cessive resolution division can lead to overly fine division, which affects the model’s judgment of
 747 the overall trend of the time series and instead impacts the model’s discriminative ability. (2) The
 748 effect of increasing the number of GGN blocks on different datasets is not consistent. On the SMD
 749 dataset, the performance is the best when $N = 4$, while on PSM and SWaT, too many GGN blocks
 750 ($N = 6, 8$) instead lead to a decline in performance. This indicates that a balance needs to be struck
 751 between complexity and expressive power, and overly deep stacking is not always beneficial. (3)
 752 Appropriate patch partitioning affects the shape of S and alters the number of time segmens a time
 753 series is divided into. The results show that on each dataset, a smaller *patch size* (such as 4 or 6)
 754 can achieve a higher F1-score, while an overly large size (such as 8) will lead to information loss. It
 755 is indicated that overly large patches will weaken the ability to capture local patterns. (4) The initial
 length L , in coordination with R , determines the size of the sliding window. It can be found that in
 any dataset, the smaller L can achieve better results.

Figure 7: Parameter sensitivity analysis on R , N , patch size and L

Based on these analyses, the optimal parameters of the SMD, PSM and SWaT datasets on $R, N, \text{patch size}$ and L are respectively $\{4, 4, 4, 2\}, \{4, 2, 2, 2\}$, and $\{4, 2, 2, 2\}$.

A.4 VISUALIZATION OF THE PROPERTIES OF SWAT DATASETS

The main advantage of GGRD lies in its ability to capture the similarity between displayed features. As shown in Figure 9, the values of \mathbf{S} in the SWaT dataset do not show significant differences at different time periods (Time segment 1, 11, 21, 32) and different resolution scales (Resolution 1, 2, 3, 4). This also limits the modeling upper limit of GGRD to a certain extent. Compared with large and cumbersome models like MODEM, the performance of our GGRD on the SWaT dataset is slightly lower.

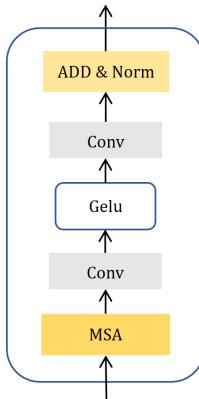
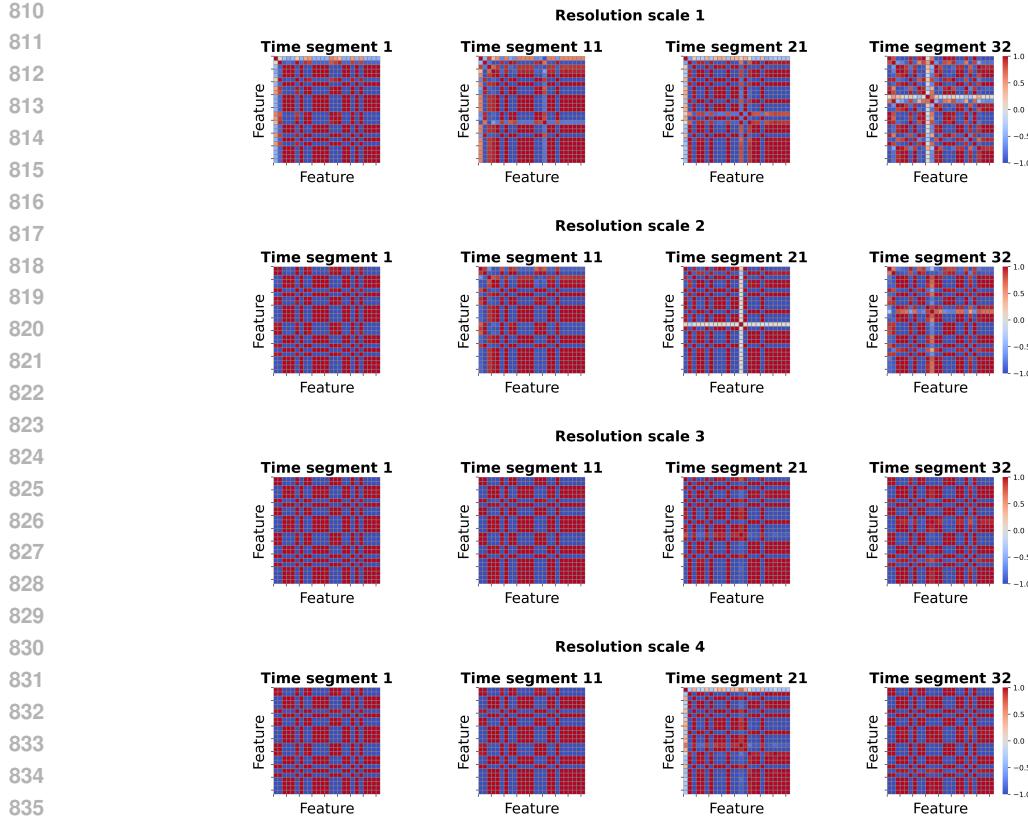


Figure 8: Time Block and Feature Block.

Figure 9: Visualization of Similarity-guided Graph Tensor \mathbf{S} on SWaT.

B ADDITIONAL EXPERIMENTS

B.1 ONE STEP VS MULTI-STEP

The core advantage of multi-step iterative denoising (such as DDPM) lies in generating high-quality and high-fidelity samples, which is crucial in tasks like image generation as it requires creating details from scratch.

However, in anomaly detection based on refactoring, our goal is not to 'create', but to 'test'. The ultimate goal is to calculate the reconstruction error. To verify the efficiency advantage of one-step denoising, we designed a multi-step denoising experiment. In multi-step denoising training, the model learns to recover from any noise level \mathbf{X}_k to \mathbf{X}_{k-1} , and in the inference stage, it iterates step by step from high noise to the final output in a completely consistent manner. We evaluated the impact of different denoising steps on model performance and computational efficiency, and the experiment was independently repeated five times.

The results in Table. 4 show that a more detailed multi-step denoising setting does not lead to an improvement in performance. Instead, it results in a significant increase in training and inference time. This further demonstrates the superiority of one-step denoising Settings in reconstruction-based anomaly detection tasks.

B.2 THE INFLUENCE OF NOISE ADDITION FREQUENCY

To illustrate the necessity and effectiveness of the denoising process of diffusion in this time series anomaly detection task, we conducted comparative experiments on K under different datasets and compared the classification metrics and MSE on the test sets. It can be found that the impact of the denoising level on the anomaly detection performance is significant, which proves the necessity of denoising in the model.

864 Table 4: The performance and efficiency of the model under different denoising steps on the SWaT
 865 dataset, with the best results indicated in bold.

Denoising Steps	F1-score (mean \pm std)	Training Time (min)	Inference Time (min)
1-step	0.7411 \pm 0.0105	27.40	5.84
5-steps	0.6855 \pm 0.0185	48.33	28.16
10-steps	0.6748 \pm 0.0159	54.01	39.23

872
 873 The experimental results Table. 5 indicate that either no noise ($K = 0$) or insufficient noise will
 874 affect the discriminative ability of anomaly detection. When the noise is moderate ($K = 500$), all
 875 the indicators of the model reach the optimum. Meanwhile, the results of MSE also show that in the
 876 task of time series anomaly detection based on reconstruction, a smaller reconstruction error does
 877 not necessarily lead to a better detection effect. In the reconstruction model, the model reconstructs
 878 by learning normal patterns to identify outliers. If no noise is added or the noise is insufficient during
 879 training or inference, the model may overfit the training data. It reconstructs the training set well but
 880 has poor generalization ability, resulting in significant reconstruction errors for normal samples with
 881 slight changes in the test set or actual scenarios, which affects abnormal judgment. Adding noise can
 882 force the model to learn denoising ability, enabling it to correctly reconstruct the normal mode even
 883 when the input contains disturbances, thereby enhancing the robustness of the reconstruction and
 884 the discrimination ability of anomaly detection. Appropriately increasing the number of noisy steps
 885 K can make the model pay more attention to the overall pattern of the sequence rather than point-
 886 by-point fitting, causing more obvious reconstruction differences in outliers and thereby improving
 887 the performance of anomaly detection.

888 Table 5: The influence of different noise levels K on the performance of anomaly detection, with the
 889 minimum value is bolded.

Datasets	K	P	R	F1	MSE
SMD	0	0.8159	0.9537	0.8794	15.64
	100	0.8088	0.8620	0.8346	15.85
	300	0.8491	0.9621	0.9021	16.98
	500	0.9812	0.8988	0.9382	14.74
PSM	0	0.7497	0.8368	0.7909	0.2540
	100	0.7160	0.8999	0.7975	0.2397
	300	0.7991	0.7943	0.7967	0.2563
	500	0.8827	0.8064	0.8428	0.2912
SWaT	0	0.5657	0.8889	0.6914	318.84
	100	0.6836	0.7665	0.7227	318.80
	300	0.5748	0.8887	0.6981	318.93
	500	0.6529	0.8840	0.7511	319.02

B.3 REPEATED EXPERIMENT STATISTICS

907 To verify the stability of the model’s performance. We independently conducted five experiments to
 908 obtain the standard deviation of F1 and the 95% confidence interval. As shown in Table. 6, the stan-
 909 dard deviation is controlled within 0.01, and the model demonstrates high stability and robustness
 910 on each dataset.

B.4 COMPARED WITH THE OVERHEAD OF OTHER ADVANCED MODELS

913 We compared the model overhead with the existing advanced models MODEM and D3R. The results
 914 of Table. 7 show that the proposed GGRD model has relatively reasonable training and inference
 915 overhead while ensuring detection performance. Compared with D3R, the total number of parame-
 916

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

Datasets	F1-score (mean \pm std)	95% CI (t-method)
PSM	0.8409 ± 0.0171	[0.8137, 0.8682]
SMD	0.9264 ± 0.0137	[0.9045, 0.9483]
SWaT	0.7411 ± 0.0105	[0.7281, 0.7542]

ters of GGRD has decreased by half, and the inference time is approximately 10 minutes, which is acceptable for actual deployment and maintenance.

930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

Table 7: The comparison with D3R regarding training time, inference time and parameters, with the best ones in bold

Model	Training Time (min)	Inference Time (min)	Total Params (MB)
D3R	42.33	18.29	199.18
MODEM	38.48	14.12	34.28
GGRD (ours)	28.70	10.31	84.92

B.5 REGARDING THE RELATIONSHIP BETWEEN PATCH SIZE AND EFFICIENCY

(i) The specific meaning of the 'graph structure' we proposed is the correlation between the features of different sensors within different time periods. Nodes represent different features, while edges represent the relationships between features. Given a $\mathbf{X}^{(r)} \in \mathbb{R}^{T \times D}$, we obtain the graph structure S by partitioning *patch*. Therefore, the size of *patch size* will affect the size of P , but will not influence the inherent number of nodes D . P represents the number of time periods that divide the time series of a sliding window T into.

(ii) Regarding the efficiency issue in handling long time series. Our model does not handle the entire infinitely long time series at one time. In practical applications, we perform the calculation within a fixed-length sliding window ($T = 64, 128, 512$). This means that no matter how long the total length of the input time series L is (weeks, months or even years), the computational complexity and memory consumption of our model at each time step are only related to the fixed window size T , and not to the total length L . The time complexity of constructing S is $O(T \times D^2)$, which is independent of *patch size*. When GGA subsequently utilizes S , it only involves an additional matrix multiplication and does not change the order of complexity.

(iii) We verified the model training and inference efficiency under different *patch size* and T (based on the SWaT dataset) in Table. 8 and Table. 9, it can be found that as the *patch size* decreases, the training and inference time does indeed increase. But the growth rate shows a clear nonlinear trend and does not increase exponentially. Therefore, a smaller *patch size* does not significantly reduce time efficiency, which is determined by the complexity of our algorithm.

962
963
964
965
966
967
968
969
970
971
Table 8: Training and inference time under different patch sizes ($T = 64$).

patch size	Training Time (s)	Inference Time (s)
2	1644.25	350.62
4	1367.52	273.94
8	1221.85	232.58
16	1136.84	215.79
32	703.78	124.81

972
973
974 Table 9: Training and inference time under different patch sizes (**T = 256**).
975
976
977
978
979
980
981

patch size	Training Time (s)	Inference Time (s)
2	4582.70	1128.08
4	2668.35	700.57
8	2430.75	550.15
16	1987.99	526.45
32	2153.20	505.33

982 C SOME SPECIFIC MODEL COMPONENTS
983984 C.1 TIME BLOCK AND FEATURE BLOCK STRUCTURE
985

986 The structures of the Time Block and the Feature Block are shown in the Figure 8, except that the
987 shapes of the input data are different (non-transposed and transposed). Firstly, the input data is ex-
988 changed for information through the multi-head self-attention (MSA) module and processed by the
989 GELU activation function. Subsequently, local features are extracted through the convolutional layer
990 and the final output is obtained by connecting the residual and normalizing the layer. This Block is
991 designed to extract the time information and feature information from the data respectively, thereby
992 providing an effective representation for subsequent time series modeling and feature interaction.
993

994 C.2 FUSION FUNCTION
995

996 In the process of coarse-to-fine reconstruction of time series data, for step m , in order to obtain the
997 similarity matrix $\bar{\mathbf{S}}^{(m)}$ and integrate the previous similarity information and also take into account
998 the current one, a function f is defined:
999

$$\bar{\mathbf{S}}^{(m)} = f(m), \quad (10)$$

$$1000 \quad f(m) = \frac{1}{2} \left(\frac{1}{m} \sum_{i=R-m+1}^R \mathbf{S}^{(i)} + \mathbf{S}^{(R-m)} \right) (m > 1), \quad f(1) = \frac{\mathbf{S}^{(R)} + \mathbf{S}^{(R-1)}}{2}. \quad (11)$$

1004 D LLM USAGE
1005

1006 We only use LLM for aid or polish writing in a very small part of the article.
1007

1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025