

ADAPTIVE TEST-TIME REASONING VIA REWARD-GUIDED DUAL-PHASE SEARCH

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
1905<br

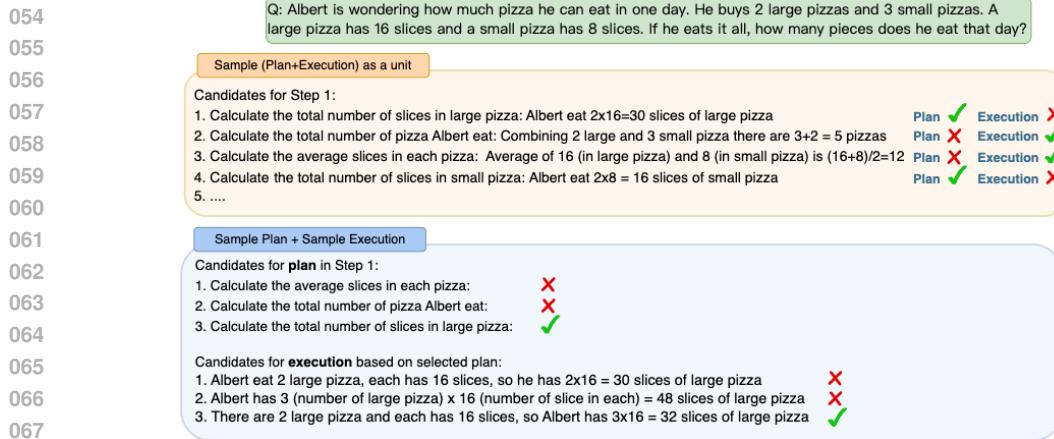


Figure 1: An example of reasoning with plan and execution as a single unit versus searched separately.

and execution, which involves carrying out precise computations or implementations (e.g., arithmetic calculations or code writing) (Zhou et al., 2022; Wang et al., 2024b; Hao et al., 2023; Wang et al., 2023a). Although the benefits of explicitly writing out plans during reasoning have been discussed (Zhou et al., 2022; Wang et al., 2024b; Hao et al., 2023; Wang et al., 2023a), most literature in test-time scaling treats planning and execution as a unified pipeline: the model generates a plan immediately followed by its execution, and both are evaluated together. The consequence is that, if a step has a correct plan but an incorrect execution, the entire candidate is discarded, wasting the useful partial structure. Conversely, if a step is already flawed at the planning stage, the search still wastes budget generating its executions. For example, consider the problem in Figure 1. Since both the plan and the execution may contain errors, sampling them as a unit often requires many trials before obtaining a step in which both are correct simultaneously.

The second limitation of existing tree-based test-time scaling methods is that they mainly adopt a fixed sampling budget per step (e.g., sampling k candidates at every reasoning step), ignoring the varying difficulty across different steps within the same question. This rigid allocation can lead to inefficient computation, especially when simple steps receive excessive attention while more challenging parts remain underexplored. While there are some studies that explore sample-wise budget allocation, i.e., dynamically distributing the overall budget across different questions or different candidate trajectories (Zuo & Zhu, 2025; Lin et al., 2025), these approaches do not address the problem of step-wise allocation within a single reasoning trajectory.

To address these limitations, we propose *DREAM*, a *D*ual-phase *R*Eward-guided *A*daptive *r*easoning *f*ramework at *t*est *t*ime. Unlike prior methods that treat each plan–execution pair as a single unit, DREAM explicitly conducts search in two stages: it first searches over multiple planning candidates and uses a reward model to select promising subgoals, and then, conditioned on the selected plans, it searches over execution candidates and applies a second reward evaluation to retain the most reliable solutions. For example, for the question in Figure 1, we first sample candidate plans and select the promising ones. Then, conditioned on these plans, we generate multiple execution candidates. This two-stage procedure ensures that poor plans are eliminated early, while promising plans can be paired with different execution attempts until the correct result is found, which ensure that computation is allocated more efficiently across the two phases. In addition to the above, we further incorporates DREAM with a *dynamic budget allocation* mechanism that adaptively adjusts the number of samples at both phases based on real-time reward feedback, enabling early stopping on easy steps and reallocating resources to harder ones.

To verify the effectiveness of the proposed algorithm, we conduct comprehensive evaluations across two domains: math reasoning and code generation. Experimental results show that our approach not only improves answer accuracy but also enhances test-time efficiency.

2 RELATED WORKS

Test-time Scaling. Test-time scaling methods improve reasoning quality without parameter updates by expending more computation at inference. According to Snell et al. (2024), two primary mechanisms for scaling test-time include (1) Distribution-Based Sampling and (2) Reward-Based Search-

108 ing. Methods of (1) include s1 (Muennighoff et al., 2025) and Reflexion (Shinn et al., 2023), which
 109 introduces sequential self-revision to iteratively refine candidate solutions, and Best-of-N (Wang
 110 et al., 2022), which samples multiple candidate reasoning chains in parallel and aggregates via
 111 majority voting (self-consistency).

112 Methods of (2) work by treating intermediate reasoning states as nodes in a search tree and expand
 113 continuations via the base LLM. They include MCTS-based methods, such as RAP (Hao et al.,
 114 2023), LiteSearch (Wang et al., 2024a), rStar (Qi et al., 2024) and rStar-Math (Guan et al., 2025),
 115 which apply Monte Carlo Tree Search to explore reasoning paths, and verifier-based methods, which
 116 rely on outcome-level judges (Cobbe et al., 2021; Snell et al., 2024) or process-supervised reward
 117 models (PRMs) (Lightman et al., 2023; Wu et al., 2024; Hooper et al., 2025) to score and prune
 118 candidates. Moreover, Setlur et al. (2025) indicate that verifier-based methods combined with
 119 search-based strategy are provably better than verifier-free approaches. While effective, most current
 120 approaches (even those that adopt a plan-execution format) still treat planning and execution as a
 121 single unified process, without performing separate search or adaptive budget allocation across the
 122 two phases. Moreover, we would like to highlight that, although a variety of test-time methods have
 123 been proposed, in our experiments we mainly consider reward-model-based methods as baselines to
 124 ensure a fair comparison, as reward models provide additional information beyond the base LLM.
 125

126 **Code Generation with LLMs.** Recent work has explored diverse strategies (including test-time
 127 scaling) to enhance LLMs for code generation. For example, S* (Li et al., 2025) employs parallel
 128 sampling with sequential scaling and adaptive input synthesis to improve code generation. Yu
 129 et al. (2025) introduce Z1, which trains the LLM on both short and long reasoning trajectories and
 130 leverages a shifted thinking window to enable the model to adaptively control the length of its ‘thinking’
 131 process according to problem complexity. In addition, PlanSearch (Wang et al., 2024b) boosts
 132 code generation by exploring diverse natural-language plans before translating them into code. In
 133 addition, tree-structured or agent-based searching frameworks like CodeTree (Li et al., 2024), Tree-
 134 of-Code (Ni et al., 2024) and Funcoder Chen et al. (2024) design stepwise generation or refinement
 135 algorithms for code generation, where candidate programs are expanded or revised through struc-
 136 tured search, demonstrating the benefits of structured exploration.

137 Although many of the above code generation methods also introduce a planning stage before the
 138 generation of code, they typically focus only on improving or selecting a good plan to guide execu-
 139 tion rather than performing separate search processes for planning and execution. In contrast, our
 140 work performs dual-phase scaling and selection, and assigns budget across phases, which has the
 141 potential of having a more effective use of computation.

142 3 PROPOSED METHOD

143 In this section, we present the general idea and design details of our proposed method. **In Sec-
 144 tion 3.1, we introduce the hierarchical structure of reasoning steps and the decomposition method.**
In Section 3.2, we introduce the dual-phase search over planning and execution, including both the
145 motivation and detailed implementation. We utilize math reasoning task and code generation task as
 146 examples to implement the main idea of the dual-phase search, and introduce the main algorithms
 147 for both tasks. In Section 3.3, we describe how to develop the reward models, including both the
 148 construction of training data and the design of the reward function.

149 3.1 HIERARCHICAL DECOMPOSITION OF REASONING STEPS

150 While chain-of-thought prompting and test-time scaling have been shown to significantly improve
 151 reasoning performance, existing approaches generally treat reasoning as a flat token-generation pro-
 152 cess, overlooking the inherent structural characteristics of reasoning.

153 To address this limitation, following recent latent-variable formulations of reasoning (He et al.,
 154 2024; Xie et al., 2021; Tutunov et al., 2023), we consider LLM reasoning considered as a hier-
 155 archical process which interleaves two structured reasoning roles: (i) global structural decisions
 156 (planning) and (ii) local step derivation (execution). This is similar to the hierarchical structure with
 157 the hidden/observed states in a Hidden Markov Model. Our “planning” phase represents the latent
 158 structural pattern, and the “execution” phase represents the observable realization of that structure.
 159 Specifically, the following are the formal definition grounded in their functional roles in generation:
 160

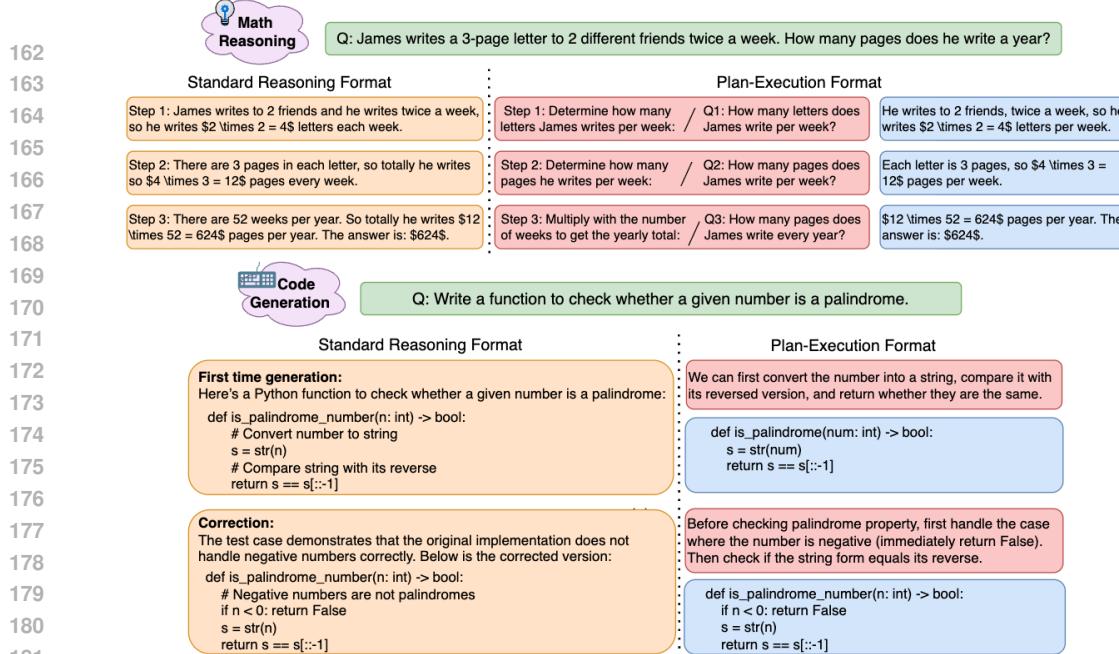


Figure 2: Plan–execution Format for Reasoning in Math and Code Tasks.(typo fixed)

Definition 3.1 (Planning Step) At reasoning step i , the planning step produces a structural decision z_i in the form of a token-level description (e.g., subgoal, decomposition strategy, or high-level direction). Formally, it is generated from the model via $z_i \in p_\theta(\cdot|h_i)$, where p_θ denotes the LLM and h_i is the dialogue history up to step i . The sequence $\{z_i\}$ forms a **hierarchical structure** of reasoning intent that organizes and guides the global trajectory of the reasoning process.

Definition 3.2 (Execution Step) Given the planning z_i , the execution step generates the corresponding concrete derivation x_i . Formally, it follows $x_i \in p_\theta(\cdot|h_i, z_i)$, where the generation is conditioned on both the prior context and the chosen structural decision. Execution operates at a lower hierarchical level, translating the structural directives in z_i into an explicit local derivations.

Although reasoning inherently includes structural (planning) and derivational (execution) components as we defined above, current LLMs do not explicitly separate them and these behaviors are often expressed in a mixed and entangled manner within regular generation. In order to make this structure explicit and observable, we utilize few-shot prompts to induce the model to generate reasoning in a separated plan–execution format. As shown in Figure 2, the prompt expresses the *plan* as a sub-question or subgoal, and the *execution* as a direct response that completes that subgoal through concrete derivations. This prompting strategy allows us to disentangle the two reasoning roles during inference, even though the underlying model was not trained with this structure. The whole few-shot prompt can be found in Appendix E.

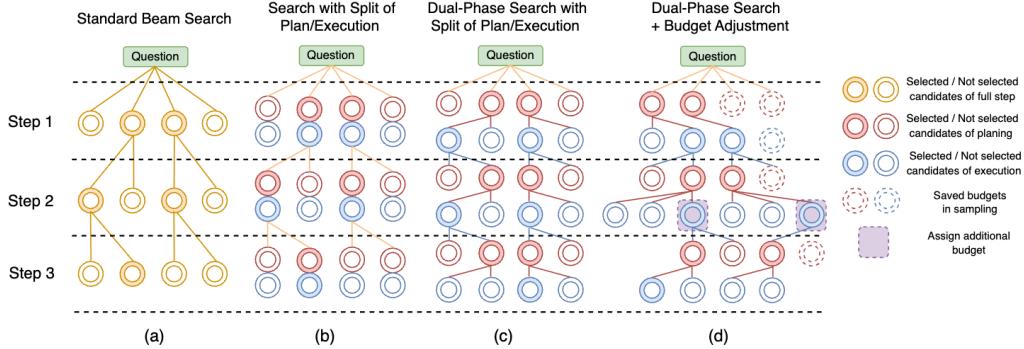
3.2 DUAL-PHASE SEARCH OVER PLANNING AND EXECUTION

Building on the hierarchical decomposition established above, in this subsection we first discuss the empirical observations that reveal distinct behaviors of planning and execution, which motivate the need for a dual-phase search strategy. We then introduce the dual-phase search algorithm in detail. Since the detailed implementation of test-time scaling varies among different type of tasks, following other literature in test-time scaling, e.g., (Hao et al., 2023; Li et al., 2024), we use two representative tasks, math reasoning and code generation, to present our algorithm.

Motivation for the Dual-phase Search. The motivation for dual-phase search stems from our observation that planning and execution exhibit fundamentally different behaviors and different error patterns. In short, planning steps are subjected to higher uncertainty, but their errors are usually less fatal. On the other hand, execution is simpler, but the calculation error in the math problem can be carried over in the later steps, being more harmful to the reasoning process. To justify the above, we provide the following two experiment results. **First**, we use the perplexity as a metric to evaluate

216
217
218
219
220
221
222
Table 1: Distribution of Perplexity and Reward

		mean	median	std	min	max
perplexity	plan	5.4646	3.8493	6.456	1.0679	88.4054
	exec	1.2614	1.2297	0.1831	1.0024	2.3179
reward	plan	0.804	0.957	0.274	0.008	1.000
	exec	0.768	0.957	0.319	0.002	1.000

233
234
Figure 3: Workflow of Standard Beam Search and Dual-Phase Search (with budget allocation).

235
236 the confidence of the LLM towards the planning and the execution step. The results are summarized
237 in Table 1. From the table, it is clear that the perplexity of the planning and execution are different,
238 demonstrating that the LLM is in general less confident in the planning step. **Second**, we also
239 check the distribution of the reward value of the planning and execution steps. The results are also
240 in Table 1. Based on the numbers, the reward for the execution steps are indeed lower than the
241 planning steps. Our conjecture is that, if the planning makes a mistake, the later planning steps can
242 still search and possibly reach the correct result. However, for the execution steps, once a step gives
243 a wrong number, it is very likely that the later steps will use this number for further calculation.

244 Given the above two differences on the planning and execution steps, we hypothesize that separating
245 them during the inference time and dynamically assigning a computation budget could enhance
246 the efficiency of test-time scaling. In particular, searching over planning steps enables the model
247 to quickly identify promising high-level directions, thereby improving the efficiency of exploration.
248 Meanwhile, searching over execution steps helps the model avoid committing to incorrect deriva-
249 tions and reducing error propagation.

250 **Dual-phase Search for Math Reasoning.** Our dual-phase search builds on the standard beam
251 search framework, whose workflow is shown in Figure 3 (a). In standard beam search, we utilize
252 the standard reasoning format, where planning and execution are not explicitly expressed, and use
253 the reasoning model to sample a fixed number of candidates. Reward models are applied to score
254 all candidates, and the top-ranked ones are retained for expansion in the next step. Motivated by
255 prior work highlighting the benefits of explicitly separating planning and execution (Zhou et al.,
256 2022; Wang et al., 2024b; 2023a), we further extend standard beam search into a variant that outputs
257 plan–execution pairs in a single step, as illustrated in Figure 3(b).

258 However, as mentioned in Section 1, the design of treating planning and execution as a single unit is
259 inefficient, since it fails to allocate computation appropriately across the two phases. To overcome
260 this drawback, our dual-phase search (Figure 3(c)) explicitly separates each step into a planning
261 phase (red nodes) and an execution phase (blue nodes). In the planning phase, N_1 candidate subgoals
262 are sampled and scored by a planning reward model (PRM_{plan}), where the top n_1 candidates are
263 selected. In the execution phase, N_2 continuations are generated conditional on the chosen plans
264 and scored by an execution reward model (PRM_{exec}). Then the top n_2 candidates will be selected for
265 the expansion of next step. This separation ensures that the weaker plans can be prune early while
266 promising plans can be given multiple execution attempts, which reduces the risk of discarding good
267 strategies due to execution errors.

268 In addition, we further extend our method to a *budget-adjusted dual-phase search* that incorporates
269 adaptive allocation. We design this mechanism based on the observation that reasoning difficulty
exhibits substantial variance: not only across different problems within a dataset, but also across
different steps within the same problem (see examples in Appendix A.2). As a result, allocating

270 the same sampling budget to every step of every example is inefficient: easy steps waste resources,
 271 while difficult steps remain underexplored. To address this limitation, we introduce an adaptive
 272 allocation strategy that stops early when confident candidates are already found and reallocates
 273 additional computation to more challenging steps, improving the overall accuracy–efficiency trade-
 274 off. The specific workflow is shown in Figure 3 (d) and the detailed algorithm can be found in
 275 Algorithm 1 in Appendix A.1. At each step, sampling in both the planning and execution phases
 276 follows a two-threshold rule. Specifically, as candidates are sampled and scored, if at least n_1 (for
 277 planning) or n_2 (for execution) candidates exceed a specific threshold τ_{p1} (for planning) or τ_{e1} (for
 278 execution), sampling is terminated early without consuming the full budget. Conversely, if after
 279 exhausting the full budget there is no candidate whose reward value higher than a lower threshold
 280 (τ_{p2}/τ_{e2}),¹ we allow at most an additional m_1 (planning) or m_2 (execution) samples to be generated
 281 in order to search harder steps more thoroughly. This mechanism prevents wasted computation on
 282 easy steps with confident high-reward candidates, while allocating extra exploration to uncertain or
 283 challenging steps. By combining dual-phase scoring with this adaptive budget policy, our method
 284 aligns computation with step-level difficulty and improves efficiency over standard beam search
 285 (as will be demonstrated in the experiments in Section 4.2). Furthermore, as will be discussed in
 286 Appendix C.2, the combination of dual-phase search and dynamic budget allocation has a synergy
 287 effect, as the two components mutually reinforce each other by reducing wasted computation and
 288 reallocating resources to harder steps.

289 **Dual-phase Search for Code Generation.** For code generation tasks, we follow the framework
 290 of CodeTree (Li et al., 2024) and extend it with our dual-phase search. We adopt this framework
 291 because its tree-based structure provides a natural backbone for implementing our dual-phase design.
 292 The key difference from math reasoning is that, in code generation task, a visible test set is available
 293 for debugging. Instead of treating each step as a partial components of the solution, in CodeTree,
 294 each step produces a complete program, and subsequent steps perform iterative debugging based on
 execution results from failed test cases of earlier solutions.

295 In the original CodeTree algorithm, both the initial generation and subsequent debugging involve
 296 sampling multiple planning candidates per step, which are then attempted sequentially. Whether to
 297 expand the current step or backtrack to alternative candidates depends on whether the current node’s
 298 reward exceeds that of the previous one. The node’s reward value is computed by combining two
 299 factors: (i) the percentage of passed test cases, and (ii) a score given by an LLM-based critic agent.

300 We modify the above framework in several ways to implement the dual-phase search. First, for each
 301 step, we apply a dedicated reward model to the N_1 sampled planning candidates, rank them, and
 302 prioritize higher-scoring plans for execution. Second, in the execution phase, we scale generation
 303 by producing N_2 candidate solutions conditioned on the chosen plan, and select the one with the
 304 highest reward. Third, as evidenced by experimental results shown in Appendix C.1, the critic agent
 305 provided only marginal benefit. As a result, we remove this component and rely solely on the per-
 306 centage of passed test cases as the execution reward. Finally, we incorporate a budget-adjustment
 307 criterion similar to that in mathematical reasoning: if the reward of a generated candidate exceeds
 308 a threshold, we stop further sampling and save the unused budget; if the rewards of all sampled
 309 candidates fall below another threshold, additional budget is allocated to generate more candidates.
 310 Through these modifications, we extend CodeTree into a dual-phase search framework that con-
 311 ducts separate searches for planning and execution, integrates reward methods for each phase, and
 312 allocates computation adaptively based on step-level difficulty.

313 **Remark.** We note that the general idea of dual-phase search is not limited to math or code reasoning,
 314 but can generalize to a wide range of reasoning tasks, as long as the solutions can be expressed in a
 315 plan–execution format and the reasoning process can be organized within a tree-based framework.

316 3.3 CONSTRUCTION OF THE REWARD MODELS

317 **Training data.** To develop the reward models, we follow the general idea of Wang et al. (2023b) to
 318 construct datasets that evaluates the quality of both planning and execution at each intermediate rea-
 319 soning step. We begin by generating complete multi-step reasoning trajectories. For each question
 320 in the training set, we sample multiple trajectories at a higher decoding temperature to ensure diver-
 321 sity. Each trajectory is expressed as a sequence of step-wise plan–execution pairs that progressively
 322 lead toward the final solution.

323 ¹In practice, for simplicity, we set $\tau_{p1}=\tau_{e1}$ and $\tau_{p2}=\tau_{e2}$

To annotate the steps, we adopt a rollout-based labeling strategy. For each intermediate plan or execution, we generate five independent continuations beginning from that step using the same LLM that produced the trajectory. If at least one of these rollouts leads to a correct final answer, the current step is labeled as positive (“+”); otherwise negative (“-”). This approach assesses the utility of a plan/execution based on its downstream impact on solving the problem.

Reward function: We implement the reward model by fine-tuning an instruction-tuned LLM. The input to the reward model, denoted as x , consists of the original question, all preceding reasoning steps (including both plans and executions), and the current plan or execution to be evaluated. In terms of the output, instead of adding a separate classification head, we follow Dong et al. (2024) to reformulate the prediction as a next-token prediction task: the final position of the input sequence is reserved for a binary label, and the model is trained to output either “+” or “-” at that position. The reward function is then given as:

$$\text{Reward}(x) = \text{softmax}(\ell(x))_+ = \frac{\exp(\ell_+(x))}{\exp(\ell_+(x)) + \exp(\ell_-(x))}$$

where $\ell_+(x)/\ell_-(x)$ are the logits output by the model when predicting the special tokens “+”/“-”.

4 EXPERIMENTS

4.1 SETUP

We evaluate our method in both math-reasoning and code generation tasks. In this subsection, we introduce the experiment setups for both tasks.

Maths Reasoning. We evaluate our method on two widely used math reasoning benchmarks: GSM8K (Cobbe et al., 2021) and MATH (Hendrycks et al., 2021). For GSM8K, the full training set of approximately 7.5k problems is used to construct training data for the reward model, while evaluation is performed on the 1.3k test set. For MATH, we leverage the 12.5k training problems to build reward-model training data and conduct evaluation on the standard MATH500 benchmark, a representative subset of 500 problems from the MATH test set. We generate large-scale synthetic trajectories to build up the training dataset: about 400k samples for GSM8K and 400k samples for MATH. The training trajectories and the rollout process for label assignment are produced using LLaMA-3-8B-Instruct (for GSM8K) and Qwen-2.5-3B-Instruct (for MATH). We then combine data from both datasets to train the reward models, fine-tuning Qwen-2.5-32B-Instruct (Yang et al., 2024b). The resulting reward model is applied in experiments on both benchmarks.

We compare our method with three baselines: majority vote, standard beam search, and REBASE (Wu et al., 2024), a tree-based search method that does not implement dual-phase search. To ensure a fair comparison, we format all reasoning in the plan–execution style and use the same reward model (trained with the rollout-based annotation strategy) across all methods which involve using reward models. For evaluation, we use three different LLMs on each benchmark: Qwen-2.5-MATH-1.5B-Instruct (Yang et al., 2024a), DeepSeekMath-7B-Instruct Shao et al. (2024), and LLaMA-3-8B-Instruct for GSM8K / LLaMA-3.1-8B-Instruct (Grattafiori & et al., 2024) for MATH. We use different versions of LLaMA for the GSM8K and MATH experiments to ensure that the reasoning models have moderate ability relative to the difficulty of each dataset. This choice allows us to better demonstrate the effectiveness of test-time scaling, since improvements are more evident when the base model is neither too strong nor too weak. [More details about the training and inference configurations and the settings of budgets and thresholds can be found in Appendix B.](#)

Code Generation Reasoning. For code generation, we conduct experiments on HumanEval (Chen et al., 2021) and MBPP (Austin et al., 2021), including their extended versions (HumanEval+ and MBPP+) (Liu et al., 2023), which include more challenging test cases. The reward model training data is drawn from the MBPP training set (approximately 600 examples), augmented with around 3,000 examples from CodeAlpaca (Chaudhary, 2023). The generation of training trajectories and the labeling process are carried out using a group of Qwen2.5-Coder models ranging in size from 1.5B to 32B. We use models with different capacities to produce a wider range of trajectories (both correct and incorrect), which improves the diversity of supervision for training the reward model. The final reward model is obtained by fine-tuning Qwen-2.5-Coder-7B-Instruct (Hui et al., 2024). We compare our method against the standard CodeTree (Li et al., 2024) and Reflexion Shinn et al. (2023), which prompts the LLM to repeatedly reflect on its previously generated code based on

378 test case results. The evaluations are conducted with two LLMs including LLaMA-3.1-8B-Instruct,
 379 Qwen-2.5-Coder-1.5B-Instruct (Hui et al., 2024).
 380

381 **4.2 MAIN RESULTS.**

382 In this subsection, we present the main experimental results to demonstrate the effectiveness of
 383 DREAM in achieving a better accuracy-efficiency trade-off in reasoning. The results for math rea-
 384 soning tasks and code generation tasks are present in Section 4.2.1 and 4.2.2, respectively.
 385

386 **4.2.1 MATH REASONING**

387 In this subsection, we present the main results of DREAM in math reasoning tasks. We show
 388 the accuracy–tokens frontier of each method in Figure 4. For DREAM, we consider the variants
 389 with/without budget allocation, which are labeled “DREAM” and “DREAM(+)”, respectively.
 390

391 Based on the results shown in Figure 4, we have the following observations. **First**, all tree-based
 392 search methods consistently achieve a significantly better accuracy–efficiency trade-off than major-
 393 ity vote. For example, on the MATH dataset with LLaMA-3.1-8B-Instruct, the performance gap
 394 can reach up to 20%. This confirms the effectiveness of tree-structured search mechanisms in im-
 395 proving the accuracy–efficiency trade-off. In addition, the consistently strong performance across
 396 models also highlights the effectiveness of the reward model trained with the rollout-based anno-
 397 tation strategy. **Second**, in general, DREAM outperforms standard beam search which does not
 398 explicitly separate planning and execution in searching. For example, on the MATH dataset with
 399 Qwen2.5-MATH-1.5B, DREAM continues to outperform beam search, and the advantage becomes
 400 more pronounced as the token budget increases. This suggests sampling and selecting planning and
 401 execution independently provides a more effective search of reasoning steps and leads to higher-
 402 quality candidate trajectories. **Third**, in many of our experiments, DREAM(+) with dynamic budget
 403 allocation provides additional gains in accuracy–efficiency trade-off, demonstrating the effective-
 404 ness of adaptively allocating computation based on the reward value. For example, on GSM8K with
 405 LLaMA-3-8B-Instruct, DREAM(+) consistently achieves about a 2% improvement in accuracy over
 406 DREAM at comparable token budgets. While in some cases (often on the MATH dataset), where
 407 the problems are more challenging, the improvement over standard dual-phase search is marginal,
 408 this can be explained by the fact that the adaptive budget mechanism is more effective when step
 409 difficulty is highly variable. When every step in a trajectory is uniformly hard, reallocating budget
 410 provides little benefit, and performance is mainly constrained by the inherent capacity of the reason-
 411 ing model and the reward model. This explanation is supported by the observation that in GSM8K, a
 412 large percentage (around 80%) of reasoning steps trigger early stopping, whereas in MATH this per-
 413 centage is much smaller (around 5%). This suggests that GSM8K contains more diverse step-level
 414 difficulty, while MATH problems are more uniformly hard.
 415

416 **Finally**, we note that in many of our settings, the models used to develop the training data for the
 417 reward model are different from the models used in the final evaluation. The only in-distribution
 418 setting is LLaMA-3-8B-Instruct with GSM8K, while all other evaluation settings are out of distri-
 419 bution. Nevertheless, in the out-of-distribution settings, DREAM consistently demonstrates strong
 420 performance, indicating that our reward model generalizes well across different backbone LLMs
 421 rather than overfitting to the one used in training.
 422

423 **4.2.2 CODE GENERATION**

424 In this subsection, we present the main results of DREAM combined with CodeTree on code gen-
 425 eration benchmarks, comparing with Reflexion and the standard CodeTree method. Figure 5 shows
 426 the accuracy–token curve of each approach. Consistent with the math reasoning experiments in Sec-
 427 tion 4.2.1, we report two variants of DREAM: the version with/without budget allocation (denoted
 428 as “DREAM/DREAM(+”).
 429

430 There are several observations from the experiment. **First**, according to the results in Figure 5, we
 431 observe that across all settings and datasets, the performance of CodeTree is significantly improved
 432 when combined with dual-phase search using reward models. For example, when the computa-
 433 tion budget is around 10^3 tokens, CodeTree+DREAM achieves an accuracy about 10% higher than
 434 CodeTree on both MBPP and HumanEval with Qwen2.5-Coder-1.5B-Instruct. This demonstrates
 435 the effectiveness of separated searching of planning and execution and leveraging dedicated reward
 436 signals for each phase. In addition, applying budget allocation consistently provides further gains
 437 in the accuracy–efficiency trade-off. This suggests that adaptively allocating computation not only
 438

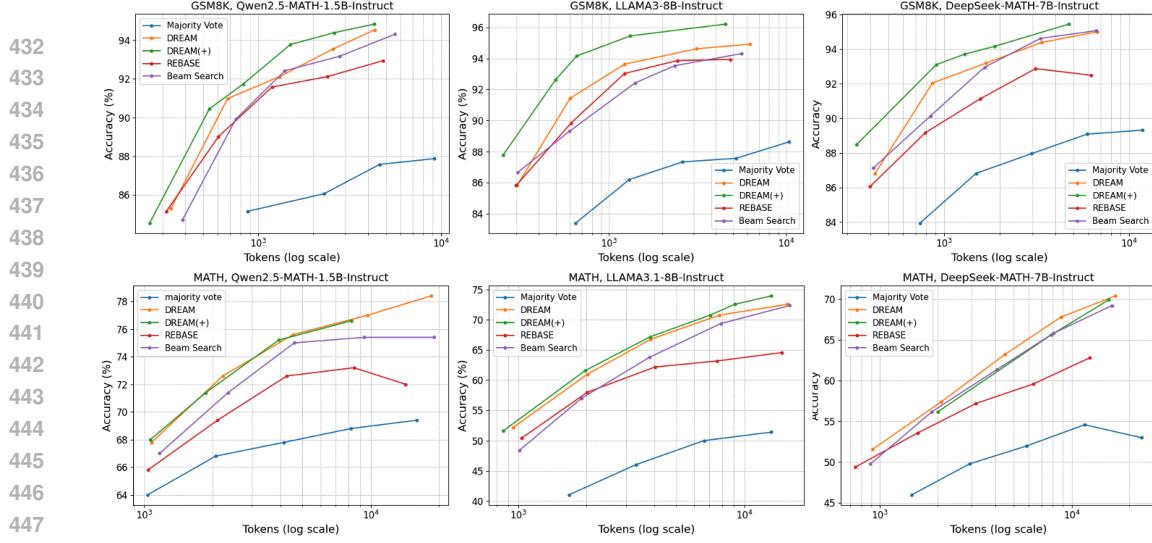


Figure 4: Accuracy vs Tokens (log) on GSM8K/MATH datasets

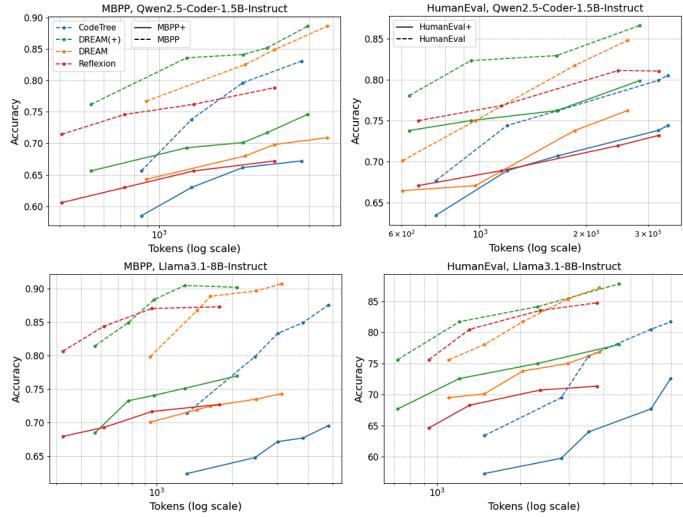


Figure 5: Accuracy vs Tokens (log) on MBPP/HumanEval datasets

improves accuracy but also reduces unnecessary generation cost. **Second**, in some cases, when the token budget is small, Reflexion also achieves a strong accuracy–efficiency trade-off. For instance, with LLAMA-3.1-8B-Instruct, Reflexion performs better when the token budget (log scale) is below 10^3 . However, as the budget increases, the accuracy of Reflexion grows more slowly compared to CodeTree+DREAM, indicating that Reflexion saturates earlier, while dual-phase search continues to benefit from additional computation. **Finally**, it is also worth noting that HumanEval does not appear in the training data of the reward model. The advantage of CodeTree+DREAM in Figure 5 also demonstrates the strong generalization ability of our reward model to unseen datasets.

In short, these findings highlight the advantages of integrating dual-phase search and adaptive budget allocation into tree-based code generation framework.

4.3 GENERALIZATION OF THE MATH REWARD MODEL TO UNSEEN DATASETS

In the main results, we demonstrated the transferability of the reward model in the code generation task on unseen datasets, as the HumanEval dataset does not appear in the training data of the reward model. In this section, we further examine the reward model’s transferability in the math reasoning domain. Specifically, we consider two out-of-distribution datasets: AMC23 (Math-AI, 2023), which consists of 40 competition-style problems from the American Mathematics Competitions 2023, and the test set of ASDiv (Miao et al., 2021), which contains 301 diverse grade-school math word problems. We evaluate reasoning with DREAM/DREAM(+) using LLAMA3/3.1-8B-Instruct, and compare it with a simple majority vote baseline. The results are presented in Table 2.

486
487 **Table 2: Performance in out-of-distribution datasets**
488

AMC23 (LLaMA3.1)								ADSIV (LLaMA3)									
DREAM(+)		DREAM		Majority Vote		DREAM(+)		DREAM		Majority Vote		DREAM(+)		DREAM		Majority Vote	
acc	# tokens	acc	# tokens	acc	# tokens	acc	# tokens	acc	# tokens	acc	# tokens	acc	# tokens	acc	# tokens	acc	# tokens
37.50%	3186.38	30.00%	3149.6	22.50%	2851.45	95.35%	160.55	93.02%	162.65	93.02%	193.93						
47.50%	6068.18	42.50%	6256.20	22.59%	5469.45	96.35%	218.76	95.35%	332.60	94.02%	388.29						
50.00%	11210.10	47.50%	12430.3	27.50%	11708.9	97.67%	543.02	96.01%	650.10	95.35%	782.80						
60.00%	23515.20	50.00%	22971.6	30.00%	22586.3	98.01%	1112.1	97.34%	1314.7	95.02%	1558.1						

493
494 From Table 2, we observe that across both datasets, DREAM with our reward model consistently
495 achieves higher accuracy at the same level of tokens compared to majority vote, and DREAM(+) provides
496 additional benefits in terms of accuracy-efficiency trade-off. Notably, on AMC23, which is
497 significantly more challenging than the datasets used to train the reward model, our approach still
498 provides strong guidance for intermediate reasoning, showing up to a 30% improvement over majority
499 voting. These findings suggest that our reward model generalizes beyond its training distribution
500 and demonstrates strong transferability to out-of-distribution math reasoning tasks.

501

4.4 ABLATION STUDIES

502 In this section, we conduct ablation studies regarding the reward model size for math reasoning
503 (Section 4.4.1), the effect of using the critic agent in code generation task, the synergy effect of
504 DREAM and empirical and theoretical analysis on the selection of the thresholds. Due to space
505 limitations, we defer the latter three studies to Appendix C.1, C.2 and C.3.

506

4.4.1 REWARD MODEL SIZE

507 In this subsection, we study the impact of reward model size. In our main experiments for math
508 reasoning, the reward model is fine-tuned from Qwen2.5-32B-Instruct, which is relatively large.
509 To assess whether smaller models can serve as effective alternatives, we additionally fine-tune a
510 reward model based on Qwen2.5-7B-Instruct. We then compare the performance of standard dual-
511 phase search (without budget allocation) under two configurations: using Qwen2.5-7B-Instruct or
512 Qwen2.5-32B-Instruct as the reward model. For reference, we also report results with simple
513 majority voting. We conduct the experiments using LLaMA3/3.1-8B-Instruct as the reasoning models
514 and the results are summarized in Table 3.

515 **Table 3: Performance comparison across different reward model size.**

GSM8K						MATH					
Qwen2.5-32B		Qwen2.5-7B		Majority vote		Qwen2.5-32B		Qwen2.5-7B		Majority vote	
acc	# tokens	acc	# tokens	acc	# tokens	acc	# tokens	acc	# tokens	acc	# tokens
91.43%	601.13	87.34%	606.191	83.40%	646.18	61.00%	2021.16	55.20%	1829.73	41.00%	1674.98
93.63%	1221.21	93.63%	1191.49	86.20%	1297.45	66.80%	3854.31	58.40%	3576.18	46.00%	3314.41
94.62%	2463.82	94.62%	2385.57	87.34%	2598.09	70.80%	7776.58	62.80%	7477.92	50.00%	6647.12
94.92%	4916.89	94.92%	4444.57	87.57%	5202.77	72.60%	15553.2	65.60%	14557.8	51.40%	13157.1

524 According to the results, the 32B reward model consistently outperforms the 7B version, achieving
525 better accuracy-efficiency trade-offs across datasets. This suggests that, as larger instruction-tuned
526 LLMs possess stronger intrinsic reasoning and representation capabilities, fine-tuning them with
527 rollout-based supervision enables the reward function to better capture nuanced signals of correctness.
528 Nevertheless, the 7B reward model also demonstrates strong effectiveness: although its step
529 selection is not as precise as the 32B model, it still achieves significantly better accuracy-efficiency
530 trade-offs compared to majority voting. This suggests that, with properly labeled training data, even
531 moderately sized reward models can provide substantial benefits while reducing computations.

532

5 CONCLUSION

533 In this work, we proposed DREAM, a dual-phase test-time scaling framework that explicitly separates
534 reasoning into planning and execution and conducts dedicated search in each phase. By
535 equipping each phase with a reward model and introducing an adaptive budget allocation mechanism,
536 our method enables finer-grained control over reasoning search, reduces wasted computation,
537 and improves accuracy on both math reasoning and code generation tasks. Empirical results show
538 that DREAM consistently outperforms standard beam search and prior PRM-based methods, high-
539 lighting the benefit of searching planning and execution separately and adaptively managing budget.

540 ETHICS STATEMENT
541542 We acknowledge the ICLR Code of Ethics and ensure that no concerns regarding the Code of Ethics
543 arise from our work. Our study does not involve human subjects, personal or sensitive data, or
544 experiments that could cause harm. All data used are either synthetic or publicly available under
545 appropriate licenses, and we have adhered to principles of fairness, transparency, and reproducibility
546 throughout.548 REPRODUCIBILITY STATEMENT
549550 To ensure reproducibility, we provide implementation details and experimental settings in Sec-
551 tion 4.1, and include additional information about hyperparameters and training configurations in
552 Appendix B.554 REFERENCES
555556 Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan,
557 Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language
558 models. *arXiv preprint arXiv:2108.07732*, 2021.560 Sahil Chaudhary. Code alpaca: An instruction-following llama model for code generation. <https://github.com/sahil280114/codealpaca>, 2023.562 Jingchang Chen, Hongxuan Tang, Zheng Chu, Qianglong Chen, Zekun Wang, Ming Liu, and Bing
563 Qin. Divide-and-conquer meets consensus: Unleashing the power of functions in code generation.
564 *Advances in Neural Information Processing Systems*, 37:67061–67105, 2024.566 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared
567 Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large
568 language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.569 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
570 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to
571 solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.573 Shizhe Diao, Pengcheng Wang, Yong Lin, Rui Pan, Xiang Liu, and Tong Zhang. Active prompting
574 with chain-of-thought for large language models. *arXiv preprint arXiv:2302.12246*, 2023.576 Hanze Dong, Wei Xiong, Bo Pang, Haoxiang Wang, Han Zhao, Yingbo Zhou, Nan Jiang, Doyen
577 Sahoo, Caiming Xiong, and Tong Zhang. Rlhf workflow: From reward modeling to online rlhf.
578 *arXiv preprint arXiv:2405.07863*, 2024.579 Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and
580 Graham Neubig. Pal: Program-aided language models. In *International Conference on Machine
581 Learning*, pp. 10764–10799. PMLR, 2023.583 Aaron Grattafiori and et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
584 URL <https://arxiv.org/abs/2407.21783>.585 Xinyu Guan, Li Lyra Zhang, Yifei Liu, Ning Shang, Youran Sun, Yi Zhu, Fan Yang, and Mao Yang.
586 rstar-math: Small llms can master math reasoning with self-evolved deep thinking. *arXiv preprint
587 arXiv:2501.04519*, 2025.589 Shibo Hao, Yi Gu, Haodi Ma, Joshua Jiahua Hong, Zhen Wang, Daisy Zhe Wang, and Zhiting Hu.
590 Reasoning with language model is planning with world model. *arXiv preprint arXiv:2305.14992*,
591 2023.592 Pengfei He, Zitao Li, Yue Xing, Yaling Li, Jiliang Tang, and Bolin Ding. Make llms better zero-shot
593 reasoners: Structure-orientated autonomous reasoning. *arXiv preprint arXiv:2410.19000*, 2024.

594 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
 595 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv*
 596 *preprint arXiv:2103.03874*, 2021.

597 Coleman Hooper, Sehoon Kim, Suhong Moon, Kerem Dilmen, Monishwaran Maheswaran,
 598 Nicholas Lee, Michael W Mahoney, Sophia Shao, Kurt Keutzer, and Amir Gholami. Ets: Ef-
 599 ficient tree search for inference-time scaling. *arXiv preprint arXiv:2502.13575*, 2025.

600 Binyuan Hui, Jian Yang, Zeyu Cui, Jiaxi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang,
 601 Bowen Yu, Keming Lu, et al. Qwen2. 5-coder technical report. *arXiv preprint arXiv:2409.12186*,
 602 2024.

603 Dacheng Li, Shiyi Cao, Chengkun Cao, Xiuyu Li, Shangyin Tan, Kurt Keutzer, Jiarong Xing,
 604 Joseph E Gonzalez, and Ion Stoica. S*: Test time scaling for code generation. *arXiv preprint*
 605 *arXiv:2502.14382*, 2025.

606 Jierui Li, Hung Le, Yingbo Zhou, Caiming Xiong, Silvio Savarese, and Doyen Sahoo. Code-
 607 tree: Agent-guided tree search for code generation with large language models. *arXiv preprint*
 608 *arXiv:2411.04329*, 2024.

609 Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan
 610 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *The Twelfth*
 611 *International Conference on Learning Representations*, 2023.

612 Junhong Lin, Xinyue Zeng, Jie Zhu, Song Wang, Julian Shun, Jun Wu, and Dawei Zhou. Plan
 613 and budget: Effective and efficient test-time scaling on large language model reasoning. *arXiv*
 614 *preprint arXiv:2505.16122*, 2025.

615 Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. Is your code generated by
 616 chatGPT really correct? rigorous evaluation of large language models for code generation. In
 617 *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=1qvx610Cu7>.

618 Math-AI. Amc23: Competition-level math problems. <https://huggingface.co/datasets/math-ai/amc23>, 2023. Hugging Face Dataset, accessed 2025-09-17.

619 Shen-Yun Miao, Chao-Chun Liang, and Keh-Yih Su. A diverse corpus for evaluating and developing
 620 english math word problem solvers. *arXiv preprint arXiv:2106.15772*, 2021.

621 Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke
 622 Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time
 623 scaling. *arXiv preprint arXiv:2501.19393*, 2025.

624 Ziyi Ni, Yifan Li, Ning Yang, Dou Shen, Pin Lv, and Daxiang Dong. Tree-of-code: A tree-structured
 625 exploring framework for end-to-end code generation and execution in complex task handling.
 626 *arXiv preprint arXiv:2412.15305*, 2024.

627 Zhenting Qi, Mingyuan Ma, Jiahang Xu, Li Lyra Zhang, Fan Yang, and Mao Yang. Mutual reason-
 628 ing makes smaller llms stronger problem-solvers. *arXiv preprint arXiv:2408.06195*, 2024.

629 Amrith Setlur, Nived Rajaraman, Sergey Levine, and Aviral Kumar. Scaling test-time compute
 630 without verification or rl is suboptimal. *arXiv preprint arXiv:2502.12118*, 2025.

631 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
 632 Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathemati-
 633 cal reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

634 Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion:
 635 Language agents with verbal reinforcement learning. *Advances in Neural Information Processing*
 636 *Systems*, 36:8634–8652, 2023.

637 Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally
 638 can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*, 2024.

648 Rasul Tutunov, Antoine Grosnit, Juliusz Ziomek, Jun Wang, and Haitham Bou-Ammar. Why can
 649 large language models generate correct chain-of-thoughts? *arXiv preprint arXiv:2310.13571*,
 650 2023.

651

652 Ante Wang, Linfeng Song, Ye Tian, Baolin Peng, Dian Yu, Haitao Mi, Jinsong Su, and Dong Yu.
 653 Litesearch: Efficacious tree search for llm. *arXiv preprint arXiv:2407.00320*, 2024a.

654

655 Evan Wang, Federico Cassano, Catherine Wu, Yunfeng Bai, Will Song, Vaskar Nath, Ziwen Han,
 656 Sean Hendryx, Summer Yue, and Hugh Zhang. Planning in natural language improves llm search
 657 for code generation. *arXiv preprint arXiv:2409.03733*, 2024b.

658

659 Lei Wang, Wanyu Xu, Yihuai Lan, Zhiqiang Hu, Yunshi Lan, Roy Ka-Wei Lee, and Ee-Peng Lim.
 660 Plan-and-solve prompting: Improving zero-shot chain-of-thought reasoning by large language
 661 models. *arXiv preprint arXiv:2305.04091*, 2023a.

662

663 Peiyi Wang, Lei Li, Zhihong Shao, RX Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang
 664 Sui. Math-shepherd: Verify and reinforce llms step-by-step without human annotations. *arXiv*
 665 *preprint arXiv:2312.08935*, 2023b.

666

667 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdh-
 668 ery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models.
 669 *arXiv preprint arXiv:2203.11171*, 2022.

670

671 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 672 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in*
 673 *neural information processing systems*, 35:24824–24837, 2022.

674

675 Yangzhen Wu, Zhiqing Sun, Shanda Li, Sean Welleck, and Yiming Yang. Inference scaling laws:
 676 An empirical analysis of compute-optimal inference for problem-solving with language models.
 677 *arXiv preprint arXiv:2408.00724*, 2024.

678

679 Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. An explanation of in-context
 680 learning as implicit bayesian inference. *arXiv preprint arXiv:2111.02080*, 2021.

681

682 An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jian-
 683 hong Tu, Jingren Zhou, Junyang Lin, et al. Qwen2. 5-math technical report: Toward mathematical
 684 expert model via self-improvement. *arXiv preprint arXiv:2409.12122*, 2024a.

685

686 An Yang et al. Qwen2.5: Technical report. *arXiv preprint arXiv:2412.15115*, 2024b. URL <https://arxiv.org/abs/2412.15115>.

687

688 Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao.
 689 React: Synergizing reasoning and acting in language models. In *International Conference on*
 690 *Learning Representations (ICLR)*, 2023.

691

692 Zhaojian Yu, Yinghao Wu, Yilun Zhao, Arman Cohan, and Xiao-Ping Zhang. Z1: Efficient test-time
 693 scaling with code. *arXiv preprint arXiv:2504.00810*, 2025.

694

695 Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuur-
 696 mans, Claire Cui, Olivier Bousquet, Quoc Le, et al. Least-to-most prompting enables complex
 697 reasoning in large language models. *arXiv preprint arXiv:2205.10625*, 2022.

698

699 Bowen Zuo and Yinglun Zhu. Strategic scaling of test-time compute: A bandit learning approach.
 700 *arXiv preprint arXiv:2506.12721*, 2025.

701

702 **A APPENDIX**
703704 **A.1 ALGORITHM OF DREAM(+)**
705706 We present the detailed algorithm of DREAM(+) in Algorithm 1.
707708 **Algorithm 1** Dual-Phase Search with budgets adjustment.
709

710 **Require:** Question Q , Planning/Execution Budget N_1/N_2 , Thresholds τ_{p_1}/τ_{p_2} and τ_{e_1}/τ_{e_2} , Beam
711 Width n_1/n_2 , Additional Budget limit m_1/m_2 .

712 1: Initialize finished paths $F \leftarrow \emptyset$
713 2: Initialize step counter $s \leftarrow 1$
714 3: Initialize beam set $B \leftarrow \emptyset$
715 4: **while** $s \leq \text{max_steps}$ **do**
716 5: **if** all $b \in B$ is finished **then**
717 **break**
718 6: **end if**
719 7: C_p (candidates storage) $\leftarrow \emptyset$
720 8: **if** $s = 1$ **then**
721 Generate up to N_1 candidates for planning p_1 with reward scores r
722 Early stop if there are n_1 planning all having reward $r > \tau_{p_1}$
723 $C_p \leftarrow C_p \cup \{(p_1^{(i)}, r_1^{(i)}) \mid i = 1, \dots, n\}$, n denotes the number of actual sampling
724 **if** all $r < \tau_{p_2}$ **then**
725 Generate up to additional m_1 candidates for planning p_1 with reward scores r
726 Early stop if there are n_1 planning having $r > \tau_{p_1}$
727 $C_p \leftarrow C_p \cup \{(p_1^{(i)}, r_1^{(i)}) \mid i = 1, \dots, n\}$
728 **end if**
729 18: **else**
730 **for** all $b \in B$ **do**
731 Generate up to N_1/n_1 candidates for planning p_s with reward scores r
732 Early stop if there are N_1/n_1 planning having $r > \tau_{p_1}$
733 $C_p \leftarrow C_p \cup \{(p_s^{(i)}, r_s^{(i)}) \mid i = 1, \dots, n\}$
734 **if** all $r < \tau_{p_2}$ **then**
735 Generate up to m_1 additional candidates for planning p_s with reward scores r
736 Early stop if there are N_1/n_1 planning having $r > \tau_{p_1}$
737 $C_p \leftarrow C_p \cup \{(p_s^{(i)}, r_s^{(i)}) \mid i = 1, \dots, n\}$
738 **end if**
739 **end for**
740 30: **end if**
741 Sort C_p and take the top- n_1 planning $\{(p_1^1, e_1^1, \dots, p_{s-1}^1, e_{s-1}^1, p_s^1, r_s^1), \dots, (p_1^{n_1}, e_1^{n_1}, \dots, p_{s-1}^{n_1}, e_{s-1}^{n_1}, p_s^{n_1}, r_s^{n_1})\}$
742 Update beam $B \leftarrow \{(Q, p_1^1, e_1^1, \dots, p_{s-1}^1, e_{s-1}^1, p_s^1), \dots, (Q, p_1^{n_1}, e_1^{n_1}, \dots, p_{s-1}^{n_1}, e_{s-1}^{n_1}, p_s^{n_1})\}$
743 C_e (candidates storage) $\leftarrow \emptyset$
744 33: **for** all $b \in B$ **do**
745 Generate up to N_2/n_2 candidates for executions e_s based on p_s with reward scores r'
746 Early stop if there are N_2/n_2 executions having reward $r' > \tau_{e_1}$
747 $C_e \leftarrow C_e \cup \{(p_s^{(i)}, e_s^{(i)}, r_s'^{(i)}) \mid i = 1, \dots, n\}$, n denotes the number of actual sampling
748 **if** all $r' < \tau_2$ **then**
749 Generate up to m_2 additional candidates for executions e_s with reward scores r'
750 Early stop if there are N_2/n_2 executions having reward $r' > \tau_{e_1}$
751 $C_e \leftarrow C_e \cup \{(p_s^{(i)}, e_s^{(i)}, r_s'^{(i)}) \mid i = 1, \dots, n\}$
752 **end if**
753 41: **end for**
754 42: Sort C_e and take the top- n_2 executions $\{(p_1^1, e_1^1, \dots, p_s^1, e_s^1, r_s'^1), \dots, (p_1^{n_2}, e_1^{n_2}, \dots, p_s^{n_2}, e_s^{n_2}, r_s'^{n_2})\}$
755 43: Update beam $B \leftarrow \{(Q, p_1^1, e_1^1, \dots, p_s^1, e_s^1), \dots, (Q, p_1^{n_2}, e_1^{n_2}, \dots, p_s^{n_2}, e_s^{n_2})\}$
756 45: $s \leftarrow s + 1$
757 46: **end while**
758 47: **return** Path with highest reward in B

756 A.2 EXAMPLES TO DEMONSTRATE DIFFICULTY VARIANTS IN REASONING STEPS.
757

758 In this subsection, we present real examples to illustrate difficulty diversity both across problems
759 in a dataset and across steps within a single problem, motivating the need for dynamic allocation.
760 Tables 4 and 5 provide examples from GSM8K and MATH: for each dataset, we include one easy
761 problem, where intermediate steps have high average reward values and low variance, and one hard
762 problem, where intermediate steps have lower average reward values and higher variance. For refer-
763 ence, we also provide the ground-truth solution for each example, offering a more intuitive sense of
764 the difficulty of each problem. These examples demonstrate that difficulty variance arises not only
765 across problems but also within individual reasoning trajectories.

766 Table 4: Examples of samples with different diversity in GSM8K datasets
767

769 Q: Josh decides to try flipping a house. He buys a house for \$80,000 and then puts in \$50,000 in repairs.
770 This increased the value of the house by 150%. How much profit did he make?

771 Ground truth solution:

772 The cost of the house and repairs came out to $80,000 + 50,000 = 130,000$
773 He increased the value of the house by $80,000 * 1.5 = 120,000$
774 So the new value of the house is $120,000 + 80,000 = 200,000$
775 So he made a profit of $200,000 - 130,000 = 70000$
776 ##### 70000

777 Average reward score for intermediate steps: 0.3613
778 Standard deviation of reward score: 0.184845

779 Q: A robe takes 2 bolts of blue fiber and half that much white fiber. How many bolts in total does it take?

780 Ground truth solution:
781 It takes $2/2 = 1$ bolt of white fiber
782 So the total amount of fabric is $2 + 1 = 3$ bolts of fabric
783 ##### 3

785 Average reward score for intermediate steps: 0.9990
786 Standard deviation of reward score: 0.001746

788 Table 5: Examples of samples with different diversity in MATH datasets
789

790 Q: What is the smallest positive integer n such that all the roots of $z^4 + z^2 + 1 = 0$ are n^{th} roots of unity?

791 Ground Truth:
792 Multiplying the equation $z^4 + z^2 + 1 = 0$ by $z^2 - 1 = (z - 1)(z + 1)$, we get $z^6 - 1 = 0$.
793 Therefore, every root of $z^4 + z^2 + 1 = 0$ is a sixth root of unity.
794 The sixth roots of unity are $e^0, e^{2\pi i/6}, e^{4\pi i/6}, e^{6\pi i/6}, e^{8\pi i/6}$, and $e^{10\pi i/6}$.
795 We see that $e^0 = 1$ and $e^{6\pi i/6} = e^{\pi i} = -1$, so the roots of $z^4 + z^2 + 1 = 0$
796 are the remaining sixth roots of unity, namely $e^{2\pi i/6}, e^{4\pi i/6}, e^{8\pi i/6}$, and $e^{10\pi i/6}$.
797 The complex number $e^{2\pi i/6}$ is a primitive sixth root of unity, so by definition,
798 the smallest positive integer n such that $(e^{2\pi i/6})^n = 1$ is 6.
799 Therefore, the smallest possible value of n is [6].

800 Average reward score for intermediate steps: 0.4206503
801 Standard deviation of the reward score: 0.3246094

803 Q: Simplify $\sqrt{242}$.

804 Ground truth:
805 Factor 242 as $11^2 \cdot 2$.
806 Then $\sqrt{242} = \sqrt{11^2 \cdot \sqrt{2}} = [11\sqrt{2}]$.

808 Average reward score for intermediate steps: 0.9947
809 Standard deviation of reward score: 0.0024

810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
Table 6: Candidate value of beam widths and sampling budgets

N_1	2	4	8	16	32
N_2	2	4	8	16	32
n_1	1	2	2	4	4
n_2	1	2	2	4	4

819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

B ADDITIONAL DETAILS OF EXPERIMENTS SETTING

In this subsection, we provide additional details of experiment settings to ensure reproducibility.

820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

B.1 TRAINING AND INFERENCE CONFIGURATIONS

- For math reasoning, we develop the reward model by fine-tuning Qwen2.5-32B-Instruct for 2 epochs with a learning rate of 2.0×10^{-6} , using the Adam optimizer and a cosine learning-rate scheduler. The fine-tuning is conducted on 8 H100 GPUs with a batch size of 32. For code generation, we fine-tune Qwen2.5-Coder-7B-Instruct for 3 epochs, while keeping the other hyperparameters the same.
- During inference on math reasoning, we set the sampling temperature to 1.0 for LLaMA3/3.1-8B-Instruct and DeepSeek-MATH-7B-Instruct, and 0.5 for Qwen2.5-MATH-1.5B-Instruct.
- For code generation, we set the temperature to 1.0 across all experiments in DREAM and DREAM(+), while using 0.0 for CodeTree and Reflexion.

837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

B.2 SETTINGS OF THRESHOLD AND BUDGETS.

For GSM8K, we use $\tau_{p1} = \tau_{e1} = 0.99$, $\tau_{p2} = \tau_{e2} = 0.9$, and $m_1 = m_2 = 2$ for all models. For MATH, we set $\tau_{p1} = \tau_{e1} = 0.9$, $\tau_{p2} = \tau_{e2} = 0.5$, and $m_1 = m_2 = 2$. To control test-time scaling, we vary the sampling budget N_1 , N_2 and the beam widths n_1 , n_2 . The candidate values used in our experiments are shown in Table 6; each column represents one configuration, corresponding to a single point in the accuracy–budget curves in the main results in Figure 4. Increasing these values proportionally increases inference-time computation.

Notably, we set $N_1 = N_2$ and $n_1 = n_2$ to keep the search space balanced between the planning and execution phases. Since both phases play important roles in producing a high-quality reasoning step, using symmetric widths prevents the search from being unintentionally biased toward one phase and also avoids introducing additional hyperparameter tuning.

849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

C ADDITIONAL ABLATION STUDIES

852
853
854
855
856
857
858
859
860
861
862
863

C.1 USAGE OF CRITIC AGENTS IN CODE GENERATION

In this subsection, we provide empirical evidence showing that when applying dual-phase search (DREAM) in the CodeTree method, the critic agent does not yield clear benefits in performance and, in fact, reduces efficiency. Table 7 presents results for DREAM+CodeTree under settings with and without critic agents. The accuracies on MBPP/HumanEval and MBPP+/HumanEval+ are reported as “weak acc” and “acc,” respectively.

From the results, we observe that to achieve the same level of reasoning accuracy, DREAM with a critic agent consistently requires substantially more generated tokens. This indicates that, although the critic agent proposed by Li et al. (2024) was originally shown to improve accuracy given sufficient computation budget, its efficiency is inferior to simply scaling generation multiple times and directly applying the percentage of passed test cases as the reward signal. These findings support our design choice of removing the critic agent from the code generation framework, simplifying the system while preserving performance and improving efficiency.

Table 7: Comparison of performance with/without critic agent

MBPP(+)												HumanEval(+)											
without critic						with critic						without critic						with critic					
weak acc	acc	# tokens	weak acc	acc	# tokens	weak acc	acc	# tokens	weak acc	acc	# tokens	weak acc	acc	# tokens	weak acc	acc	# tokens	weak acc	acc	# tokens	weak acc	acc	# tokens
81.49%	68.52%	569.59	81.49%	68.52%	1051.52	75.61%	67.69%	721.28	75.61%	67.69%	1235.92	86.51%	74.10%	771.66	87.58%	74.88%	1348.16	81.71%	72.57%	1205.10	78.66%	70.13%	1596.59
88.37%	74.09%	976.72	86.25%	73.02%	1627.35	84.14%	72.57%	1619.51	80.49%	71.34%	1827.49	90.48%	75.14%	1292.83	88.37%	76.19%	2074.12	84.14%	75.00%	2307.87	81.71%	71.95%	2311.83
90.21%	76.98%	2081.73	90.73%	75.94%	2620.38	87.80%	78.04%	4552.75	87.20%	76.83%	4833.02												

C.2 SYNERGY EFFECT OF DREAM

In this subsection, we demonstrate the synergy between dual-phase search and dynamic budget allocation. Specifically, we compare the benefits of applying budget allocation to dual-phase search versus standard beam search. Figure 6 plots the accuracy–tokens frontier of four methods: DREAM, DREAM(+), Beam Search, and Beam Search(+), where the “+” variants denote the incorporation of dynamic budget allocation. From the figure, we observe that budget allocation also improves the accuracy–efficiency trade-off when applied to standard beam search, indicating the general effectiveness of this design. However, the gains for Beam Search(+) are consistently smaller than those achieved by DREAM(+). This comparison highlights the complementary nature of the two components: dual-phase search and dynamic budget allocation. The synergy arises because dual-phase search reduces wasted computation by separating planning and execution, while dynamic budget allocation further reallocates saved resources to harder steps, making the two mechanisms mutually reinforcing.

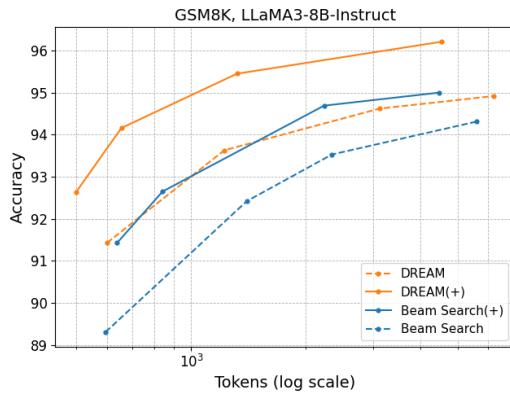


Figure 6: Accuracy-Tokens Frontier across methods.

C.3 ABLATION STUDIES ON THE THRESHOLDS.

In this subsection, we provide (1). Justification about why $\tau_{e1}=\tau_{p1}$; and (2). Theoretical and empirical analysis on how different choices of the thresholds $\tau_{e1}, \tau_{p1}, \tau_{e2}, \tau_{p2}$ affect the performance of DREAM+.

C.3.1 JUSTIFICATION ABOUT WHY $\tau_{e1}=\tau_{p1}$

To justify why $\tau_{e1}=\tau_{p1}$, empirically, as shown in Table 8, using the same threshold ($\tau_{e1}=\tau_{p1}$) yields slightly better performance.

The intuition behind this observation is twofold. Firstly, although planning and execution operate on different sampling spaces, their reward scores share a comparable numerical scale: both reward models evaluate step-level reasoning quality (whether a subgoal is sensible, or whether a subexecution is correct). Since the rewards are trained under the same supervision paradigm, applying the same threshold yields consistent pruning behavior in both phases.

918 Table 8: Performance of DREAM in gsm8k with Qwen2.5-MATH-1.5B-Instruct when τ s take dif-
 919 ferent values. (We set $N_1 = N_2 = 4$, $n_1 = n_2 = 2$)

$\tau_2 = 0.7$			$\tau_2 = 0.9$		
τ_{p_1}/τ_{e_1}	acc	tokens	τ_{p_1}/τ_{e_1}	acc	tokens
0.9/0.8	89.40%	501.276	0.99/0.9	90.00%	569.846
0.8/0.9	89.00%	492.466	0.9/0.99	91.20%	525.868
0.85/0.85	90.00%	492.182	0.95/0.95	91.60%	522.408

928 Table 9: Performance of DREAM+ with different values of thresholds (GSM8k, Qwen2.5-MATH-
 929 1.5B-Instruct)

$\tau_{e_1} = \tau_{p_1} = 0.99$			$\tau_{e_2} = \tau_{p_2} = 0.3$		
τ_{e_2}/τ_{p_2}	acc	tokens	τ_{e_1}/τ_{p_1}	acc	tokens
0.9/0.9	0.92	538.754	0.99/0.99	0.892	489.742
0.7/0.7	0.902	520.430	0.8/0.8	0.884	472.432
0.5/0.5	0.896	506.496	0.6/0.6	0.878	467.584
0.3/0.3	0.892	489.742	0.4/0.4	0.876	465.156

940 Secondly, Different thresholds can easily create imbalanced pruning: Both planning and execution
 941 are essential for producing a good reasoning step. If one phase is pruned more aggressively than
 942 the other, the search becomes unbalanced. In such cases, either valid plans are discarded too early,
 943 or promising plans cannot be properly expanded. Therefore, using the same threshold keeps the
 944 pruning strength aligned across phases and prevents either side from dominating the search.

946 C.4 THEORETICAL AND EMPIRICAL ANALYSIS ON THE INFLUENCE OF THE THRESHOLDS.

948 **Empirical Results.** In Table 9, we provide empirical results on how different values of
 949 $\tau_{e1}, \tau_{p1}, \tau_{e2}, \tau_{p2}$ influence the performance of DREAM+.

950 The left block reports results under a fixed high early step threshold ($\tau_{e1} = \tau_{p1} = 0.99$) while
 951 varying the extra-budget thresholds τ_{e2}/τ_{p2} . The right block reports results under a fixed low extra-
 952 budget thresholds ($\tau_{e2} = \tau_{p2} = 0.3$) while varying the early stop thresholds τ_{e1}/τ_{p1} .

953 From the results in Table 9, we observe two consistent trends.

955 First, when τ_1 is fixed, decreasing τ_2 leads to lower accuracy and fewer generated tokens. This
 956 result is expected, as a higher τ_2 imposes a stricter extra budget standard, causing the model to more
 957 frequently determine that the current candidates are not sufficient and therefore allocate more budget
 958 to explore further. As τ_2 decreases, fewer steps trigger budget expansion, reducing exploration and
 959 slightly lowering accuracy.

960 Second, when τ_2 is fixed, decreasing τ_1 also reduces accuracy and the number of generated tokens. This
 961 is because the thresholds τ_1 determine whether the remaining budget for a step should continue
 962 to be used. A higher τ_1 corresponds to a stricter stopping criterion and encourages generating more
 963 candidates within the step, thereby resulting in higher accuracy and higher computation costs.

964 Overall, τ_1 and τ_2 control different aspects of intra-step search, and their effects on accuracy and
 965 efficiency follow intuitive patterns.

966 Theoretical Analysis: Markov-based Analysis of the Threshold Behaviors.

968 To build a theoretical understanding of how the thresholds τ_1 and τ_2 influence DREAM+, we begin
 969 by modeling the within-step search as a simplified Markov process. To begin, we model the single
 970 step generation in DREAM+ as a simple Markov process with states $\{1, 2, 3, \dots, n, *\}$ where state
 971 * denotes the end state. For any states a and b , let $P(x_{i+1} = a \mid x_i = b) = P_{b \rightarrow a}$. We adopt the
 972 following assumptions for the transition matrix.

972 1. $P_{b \rightarrow b+1} = \frac{1}{2}$ for any $1 \leq b < n$
 973
 974 2. $P_{a \rightarrow *} = \frac{1}{2}$ for any $1 \leq a \leq n$
 975
 976 3. $P_{n \rightarrow n} = \frac{1}{2}$
 977

978 At step i , each beam samples multiple candidates for next states and selects the one with the highest
 979 reward value based on the reward model. This process can be modeled as: at state $x_i = b$, we
 980 draw several independent samples of x_{i+1} under a sampling budget k , and retain the candidate with
 981 the highest reward. We further assume that the reward assigned by the reward model is positively
 982 correlated with the transition likelihood toward the end state $*$, $P_{x_{i+1} \rightarrow *}^\infty$. Then we have the following
 983 theorem:

984
 985 **Theorem C.1** *Under the assumptions described above, suppose the maximum number of reasoning
 986 steps is M (M is sufficiently large). Then starting from state b , the probability of eventually reaching
 987 the terminal state $*$ within the remaining $M - i$ steps is: $P_{b \rightarrow *}^{(M-i)} = 1 - 2^{-(M-i-b+1)}$. Furthermore,
 988 if at step i we sample k candidates for x_{i+1} and keep the one with the highest reward, then the
 989 probability of successfully reaching the end state becomes: $P^* = 1 - (\frac{1}{2^{M-i-b+1}})^k$.*

990 Based on the above threshold, we have the following two propositions on the impact of τ_1 and τ_2
 991 towards P^* . For simplicity, each proposition considers τ_1 and τ_2 respectively.

992
 993 **Proposition C.1** *Let $p = 1 - \frac{1}{2^{(M-i-b)}}$. Assume the sampling budget is and $\tau_2 = 0$. Then:*

994 (1) *If $\tau_1 \in (0, p]$, the probability of successfully reaching the end states is: $P^* = 1 - (\frac{1}{2^{M-i-b+1}})^k$.*
 995 (2) *If $\tau_1 \in (p, 1]$, $P^* = 1 - (\frac{1}{2^{M-i-b+1}})^k$*
 996 *When $k > 1$, P^* is higher when $\tau_1 \in (p, 1]$.*

997
 998 **Proposition C.2** *Let $p = 1 - \frac{1}{2^{(M-i-b)}}$. Assume that the original sampling budget is k , the maximum
 999 additional budget triggered by τ_2 is k' , and $\tau_1 = 1$ is fixed. Then:*

1000 (1) *If $\tau_2 \in (0, p]$, the probability of successfully reaching the end states is: $P^* = 1 - (\frac{1}{2^{M-i-b+1}})^k$.*
 1001 (2) *If $\tau_2 \in (p, 1]$, $P^* = 1 - (\frac{1}{2^{M-i-b+1}})^{(k+k')}$*
 1002 *Thus, P^* is higher when $\tau_2 \in (p, 1]$.*

1003
 1004 The two proposition provide evidence that, when all the other conditions are fixed, increasing either
 1005 τ_1 or τ_2 increase the probability of successfully reaching the end state increase. This theoretical
 1006 analysis explains why larger thresholds generally lead to better performance in DREAM+, as they
 1007 effectively allocate more sampling budget to promising reasoning paths.

1008 For simplicity, our analysis considers the effect of each threshold independently by fixing one and
 1009 varying the other. Intuitively, increasing both τ_1 and τ_2 at the same time would compound their
 1010 effects, which tightening selection and increasing the effective sampling budget, thereby further
 1011 enhancing the probability of reaching the correct solution.

1015 D ADDITIONAL EVALUATION

1016 D.1 EXPERIMENTS IN MORE CHALLENGING DATASETS AND MORE CAPABLE MODELS

1017 To provide further empirical evidence for the effectiveness of our method, in this subsection we ad-
 1018 ditionally evaluate Qwen2.5-MATH-72B (Yang et al., 2024b), one of the strongest math-reasoning
 1019 models, on two harder benchmarks: AMC23 (Math-AI, 2023) and GSM8K-Hard (Gao et al., 2023).
 1020 In these experiments, we use the same reward model as the one used in the experiments of Sec-
 1021 tion 4.2.1. The results are presented in Table 12 and 11 and plot in Figure 8. From the results, we
 1022 can observe that, DREAM consistently achieves a better accuracy-efficiency trade-off across both
 1023 datasets compared with the baselines, demonstrating that the benefits of our method generalize to
 1024 stronger models and more difficult reasoning tasks.

1026 D.2 COMPARISON BETWEEN PLANNING/EXECUTION DECOMPOSITION AND INCREASING
1027 GRANULARITY OF CoT
10281029 In this subsection, we provide comparison between planning/execution decomposition and increasing
1030 granularity of CoT to further demonstrate that the effectiveness brought by our method does not
1031 come from simply making the reasoning steps more fine-grained, but from the structural decomposi-
1032 tion that separates high-level subgoal planning from low-level execution.1033 To begin with, we would like to highlight the fundamental difference between plan–execution de-
1034 composition and simply increasing step granularity. Specifically, increasing granularity means di-
1035 viding a single reasoning step into smaller sequential segments (e.g., splitting one step into two
1036 micro-steps). In contrast, plan–execution decomposition separates two different types of reasoning
1037 behaviors. The planning component focuses on forming high-level strategic decisions or subgoals,
1038 while the execution component focuses on concretely carrying out these decisions. Consequently,
1039 decomposing a step into plan and execution changes the structure of the reasoning process rather
1040 than merely adjusting the level of granularity. To better illustrate the difference, we provide real
1041 examples demonstrating both strategies and the combination of both strategies:1042 Q: Jack deposited \$4000 in a bank with a 1% annual interest rate. What will be the total amount
1043 after one year?1044 **1. Standard CoT:**1045 Step 1: With 1% interest rate, after the first year, the total value is $10000 \times (1+1\%) = 10100$ 1046 Step 2: With 2% interest rate, after the second year, the total value is $10100 \times (1+2\%) = 10302$ 1047 **2. CoT with increased granularity:**1048 Step 1: With 1% interest rate, the first year's interest is $10000 \times 1\% = 100$.1049 Step 2: The total value after the first year is $10000 + 100 = 10100$.1050 Step 3: With 2% interest rate, the second year's interest $10100 \times 2\% = 202$.1051 Step 4: The total value after the second year is $10100 + 202 = 10302$.1052 **3. CoT with plan/execution decomposition:**

1053 Step 1: Determine the total value after the first year:

1054 - Given that the interest rate of the 1st year is 1%. The total value becomes $10000 \times (1+1\%) = 10100$

1055 Step 2: Determine the total value after the second year:

1056 - Given that the interest rate of the 1st year is 2%. The total value becomes $10100 \times (1+2\%) = 10302$.1057 **4. CoT with plan/execution decomposition + increased granularity:**

1058 Step 1: Determine the first year's interest:

1059 - Given that the interest rate of the 1st year is 1%. The first year's interest is $1\% \times 10000 = 100$

1060 Step 2: Determine the total value after the first year:

1061 - After the first year, the interest is $10000 + 100 = 10100$

1062 Step 3: Determine the second year's interest:

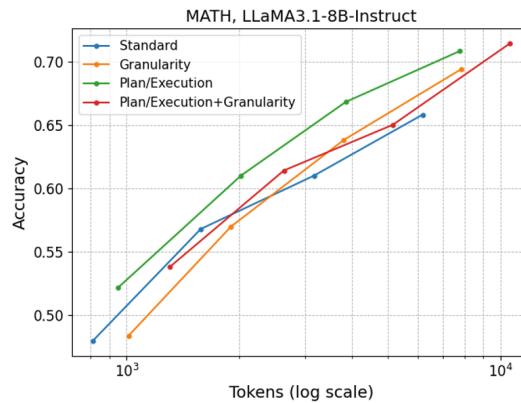
1063 - Given that the interest rate of the 2nd year is 2%. The second year's interest is $2\% \times 10100 = 202$

1064 Step 4: Determine the total value after the second year:

1065 - After the second year, the total value becomes $10100 + 202 = 10302$.1066 To better demonstrate that the improvement of DREAM comes from structural plan–execution de-
1067 composition rather than simply splitting steps into smaller units, we compare all the four CoT vari-
1068 ants shown in the above example using DREAM under different beam widths. The numerical results
1069 are provided in Table 10 and visualized in Figure 7.1070 Based on the results, we have the following observations. First, when comparing entries within
1071 each row of the table i.e., under the same searching width, we observe both increasing granularity
1072 and applying plan–execution decomposition can improve accuracy, but both approaches also require
1073 more tokens and therefore higher computation. However, the accuracy–efficiency trade-off shown
1074 in Figure 1 reveals a clear difference: there is no obvious gain from finer-grained steps in terms of
1075 accuracy–efficiency trade-off, whereas the gain from plan–execution decomposition is obvious.1076 Second, when the beam width is fixed, combining increased granularity with plan–execution decom-
1077 position leads to higher accuracy than using plan–execution alone. However, this combined strategy
1078 substantially increases the number of generated tokens and makes the reasoning trace significantly

1080
1081
1082 Table 10: Comparison of four CoT formats on the MATH dataset with LLaMA3.1-8B-Instruct
1083
1084

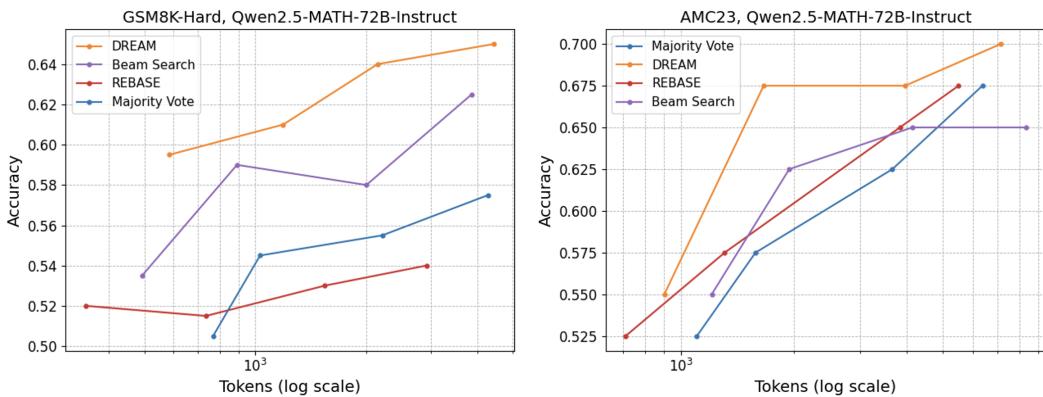
Standard		Granularity		Plan / Execution		Plan/Execution + Granularity	
acc	tokens	acc	tokens	acc	tokens	acc	tokens
0.480	811.938	0.484	1011.53	0.522	949.49	0.538	1302.44
0.568	1578.09	0.570	1902.53	0.610	2021.16	0.614	2631.79
0.610	3170.16	0.638	3805.05	0.668	3854.31	0.650	5149.69
0.658	6188.42	0.694	7872.97	0.708	7776.57	0.714	10587.93

1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103 Figure 7: Comparison of accuracy-efficiency trade-off of the four CoT formats
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133

more verbose. As a consequence, the computational cost grows faster than the improvement in accuracy. When examining the accuracy–efficiency trade-off, we do not observe advantages beyond using plan–execution decomposition alone.

In summary, the results confirm that the benefit of DREAM does not come from simply making the reasoning steps more fine-grained, but from the structural decomposition that separates high-level subgoal planning from low-level execution, enabling more effective exploration and pruning and ultimately yielding a better accuracy–efficiency trade-off.

E FEW-SHOT PROMPTS FOR GUIDING REASONING IN THE PLAN–EXECUTION FORMAT

1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267
2268
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2290
2291
2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321
2322
2323
2324
2325
2326
2327
2328
2329
2330
2331
2332
2333
2334
2335
2336
2337
2338
2339
2340
2341
2342
2343
2344
2345
2346
2347
2348
2349
2350
2351
2352
2353
2354
2355
2356
2357
2358
2359
2360
2361
2362
2363
2364
2365
2366
2367
2368
2369
2370
2371
2372
2373
2374
2375
2376
2377
2378
2379
2380
2381
2382
2383
2384
2385
2386
2387
2388
2389
2390
2391
2392
2393
2394
2395
2396
2397
2398
2399
2400
2401
2402
2403
2404
2405
2406
2407
2408
2409
2410
2411
2412
2413
2414
2415
2416
2417
2418
2419
2420
2421
2422
2423
2424
2425
2426
2427
2428
2429
2430
2431
2432
2433
2434
2435
2436
2437
2438
2439
2440
2441
2442
2443
2444
2445
2446
2447
2448
2449
2450
2451
2452
2453
2454
2455
2456
2457
2458
2459
2460
2461
2462
2463
2464
2465
2466
2467
2468
2469
2470
2471
2472
2473
2474
2475
2476
2477
2478
2479
2480
2481
2482
2483
2484
2485
2486
2487
2488
2489
2490
2491
2492
2493
2494
2495
2496
2497
2498
2499
2500
2501
2502
2503
2504
2505
2506
2507
2508
2509
2510
2511
2512
2513
2514
2515
2516
2517
2518
2519
2520
2521
2522
2523
2524
2525
2526
2527
2528
2529
2530
2531
2532
2533
2534
2535
2536
2537
2538
2539
2540
2541
2542
2543
2544
2545
2546
2547
2548
2549
2550
2551
2552
2553
2554
2555
2556
2557
2558
2559
2560
2561
2562
2563
2564
2565
2566
2567
2568
2569
2570
2571
2572
2573
2574
2575
2576
2577
2578
2579
2580
2581
2582
2583
2584
2585
2586
2587
2588
2589
2590
2591
2592
2593
2594
2595
2596
2597
2598
2599
2600
2601
2602
2603
2604
2605
2606
2607
2608
2609
2610
2611
2612
2613
2614
2615
2616
2617
2618
2619
2620
2621
2622
2623
2624
2625
2626
2627
2628
2629
2630
2631
2632
2633
2634
2635
2636
2637
2638
2639
2640
2641
2642
2643
2644
2645
2646
2647
2648
2649
2650
2651
2652
2653
2654
2655
2656
2657
2658
2659
2660
2661
2662
2663
2664
2665
2666
2667
2668
2669
2670
2671
2672
2673
2674
2675
2676
2677
2678
2679
2680
2681
2682
2683
2684
2685
2686
2687
2688
2689
2690
2691
2692
2693
2694
2695
2696
2697
2698
2699
2700
2701
2702
2703
2704
2705
2706
2

1134

1135

1136

1137

1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

1149

1150

1151

1152

1153

1154

1155

1156

1157

1158

1159

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187

Table 11: Qwen2.5-MATH-72B-Instruct on GSM8k-Hard

DREAM		Beam		Rebase		Majority Vote	
acc	tokens	acc	tokens	acc	tokens	acc	tokens
0.595	584.345	0.535	492.910	0.5200	347.075	0.505	768.035
0.610	1190.035	0.590	891.890	0.5150	735.220	0.545	1030.260
0.640	2146.050	0.580	2001.450	0.530	1541.035	0.555	2210.765
0.650	4436.815	0.625	3863.665	0.540	2911.175	0.575	4281.115

Table 12: Qwen2.5-MATH-72B-Instruct on AMC23

DREAM		Beam Search		REBASE		Majority Vote	
acc	tokens	acc	tokens	acc	tokens	acc	tokens
0.55	901.7	0.550	1206.625	0.525	710.425	0.525	1101.425
0.675	1660.2	0.625	1943.300	0.575	1306.475	0.575	1579.400
0.675	3960.83	0.650	4142.125	0.650	3837.675	0.625	3663.275
0.700	7128.45	0.650	8314.875	0.675	5503.650	0.675	6373.825

Table 13: Few-shot Prompt for the GSM8K dataset

1188	Learn from the following examples to answer the given question. For each step, begin by stating the
1189	general plan, followed by executing that plan (started with “-”), separated by a colon.
1190	When the answer is finished, end with “The final answer is:”.
1191	
1192	Q: A father is building a playset for his son and needs to purchase lumber, nails, and fabric. When he
1193	started planning the project, the necessary lumber cost \$450, the nails cost \$30, and the fabric cost \$80.
1194	However, recent economic inflation has caused the price of lumber to increase by 20%, the price of nails
1195	to increase by 10%, and the price of fabric to increase by 5%. In dollars, how much more money will it
1196	cost to complete the project now (after inflation) than it would have when the father began planning?
1197	A: Let’s think step by step.
1198	Step 1: Find the new cost of lumber:
1199	- The lumber originally cost \$450, and with a 20% increase it becomes $450 \times 1.20 = 540$.
1200	Step 2: Find the new cost of nails:
1201	- The nails originally cost \$30, and with a 10% increase they become $30 \times 1.10 = 33$.
1202	Step 3: Find the new cost of fabric:
1203	- The fabric originally cost \$80, and with a 5% increase it becomes $80 \times 1.05 = 84$.
1204	Step 4: Find the total original cost:
1205	- Before inflation, the total was $450 + 30 + 80 = 560$.
1206	Step 5: Find the total new cost:
1207	- After inflation, the total is $540 + 33 + 84 = 657$.
1208	Step 6: Compare the costs:
1209	- The difference between the new and old total is $657 - 560 = 97$. The final answer is: 97.
1210	
1211	Q: There are five times as many swordfish as pufferfish in an aquarium exhibit.
1212	If there are 90 fish total, how many pufferfish are there?
1213	A: Let’s think step by step.
1214	Step 1: Express swordfish in terms of pufferfish:
1215	- If the number of pufferfish is x , then the number of swordfish is $5x$.
1216	Step 2: Write the total number of fish:
1217	- The total is $x + 5x = 6x$.
1218	Step 3: Set up the equation:
1219	- Since the total is 90, we have $6x = 90$.
1220	Step 4: Solve for x :
1221	- Dividing both sides by 6 gives $x = 15$. The final answer is: 15.
1222	
1223	Q: Mrs. Johnson’s class raised \$2300,
1224	which is twice what Mrs. Sutton’s class raised. If Mrs. Sutton’s class raised 8 times less than Miss Rollin’s
1225	class, and Miss Rollin’s class raised one-third of the total,
1226	how much money did the school raise after deducting 2% fees?
1227	A: Let’s think step by step.
1228	Step 1: Find Mrs. Sutton’s amount:
1229	- Since Mrs. Johnson’s \$2300 is twice Sutton’s, Sutton’s amount is $2300 / 2 = 1150$.
1230	Step 2: Find Miss Rollin’s amount:
1231	- Rollin’s class raised 8 times more than Sutton’s, so $1150 \times 8 = 9200$.
1232	Step 3: Find the school’s total before fees:
1233	- Since Rollin’s amount is one-third of the total, the total is $9200 \times 3 = 27600$.
1234	Step 4: Calculate the deduction:
1235	- 2% of 27600 is $27600 \times 0.02 = 552$.
1236	Step 5: Find the net amount:
1237	- Subtracting the deduction gives $27600 - 552 = 27048$. The final answer is: 27048.
1238	
1239	Q: John buys 500 newspapers. Each newspaper sells for \$2. He sells 80% of them. He buys them all for 75%
1240	less than the price at which he sells them. How much profit does he make?
1241	A: Let’s think step by step.
1242	Step 1: Determine how many he sells:
1243	- He sells 80% of 500, which is $0.8 \times 500 = 400$.
1244	Step 2: Determine the selling price:
1245	- Each newspaper sells for \$2.
1246	Step 3: Compute revenue:
1247	- Selling 400 at \$2 each gives $400 \times 2 = 800$.
1248	Step 4: Determine purchase price per paper:
1249	- He bought them 75% cheaper, meaning at 25% of 2, which is 0.50.
1250	Step 5: Compute total cost:
1251	- Buying 500 at \$0.50 each costs $500 \times 0.50 = 250$.
1252	Step 6: Calculate profit:
1253	- Revenue minus cost is $800 - 250 = 550$. The final answer is: 550.

1242

1243

Continue:

1244

1245

1246

1247

Q: A house's living room, dining room and kitchen take up 1,000 sq ft. The house also has a guest bedroom and a master bedroom suite. If the guest bedroom is a quarter of the size of the master bedroom suite, and the house is 2,300 sq ft total, how large is the master bedroom suite?

A: Let's think step by step.

Step 1: Find how much space is left for the bedrooms:

- The house is 2,300 sq ft total, and 1,000 sq ft is already used by the living room, dining room, and kitchen. So the bedrooms must occupy $2300 - 1000 = 1300$ sq ft.

Step 2: Express the guest bedroom in terms of the master bedroom:

- If the master bedroom is x sq ft, the guest bedroom is one quarter of that, or $x/4$.

Step 3: Write the total space of both bedrooms:

- Together the master and guest bedrooms occupy $x + x/4 = (5/4)x$.

Step 4: Solve for the master bedroom size:

- Since the two bedrooms together are 1,300 sq ft, we have $(5/4)x = 1300$.

Solving gives $x = (1300 \times 4)/5 = 1040$. The final answer is: 1040.

1257

1258

1259

Q: Carter usually bakes 6 cheesecakes, 5 muffins, and 8 red velvet cakes regularly for a week. For this week he was able to bake triple the number of cheesecakes, muffins, chocolate moist cakes, and red velvet cakes. How much more cakes was Carter able to bake for this week?

A: Let's think step by step.

Step 1: Calculate Carter's usual total:

- Normally he bakes $6 + 5 + 8 = 19$ cakes per week.

Step 2: Find his cheesecakes this week:

- Tripling 6 cheesecakes gives $6 \times 3 = 18$.

Step 3: Find his muffins this week:

- Tripling 5 muffins gives $5 \times 3 = 15$.

Step 4: Find his red velvet cakes this week:

- Tripling 8 red velvet cakes gives $8 \times 3 = 24$.

Step 5: Add up the new totals:

- Altogether he baked $18 + 15 + 24 = 57$ cakes this week.

Step 6: Compare to his usual:

- Compared to 19 cakes normally, the increase is $57 - 19 = 38$. The final answer is: 38.

1271

1272

1273

Q: Four adults with 32 teeth went to the dentist. The first had $1/4$ removed, the second $3/8$ removed, the third half removed, and the fourth 4 removed. What's the total number of teeth removed?

A: Let's think step by step.

Step 1: Find the first person's removal:

- One-quarter of 32 teeth is $32 \times 1/4 = 8$.

Step 2: Find the second person's removal:

- Three-eighths of 32 teeth is $32 \times 3/8 = 12$.

Step 3: Find the third person's removal:

- Half of 32 teeth is $32 \times 1/2 = 16$.

Step 4: Note the fourth person's removal:

- This was directly stated as 4 teeth.

Step 5: Add all removals together:

- $8 + 12 + 16 + 4 = 40$ teeth removed in total. The final answer is: 40.

1283

1284

Q: Joy fosters dogs. The mom dog eats 1.5 cups of food three times a day. The puppies (5 total) each eat 0.5 cups twice a day. How much food will Joy need for the next 6 days?

A: Let's think step by step.

Step 1: Calculate the mom's daily food:

- She eats 1.5 cups at each of 3 meals, so $1.5 \times 3 = 4.5$ cups per day.

Step 2: Calculate one puppy's daily food:

- Each puppy eats 0.5 cups at 2 meals, so $0.5 \times 2 = 1$ cup per day.

Step 3: Calculate all puppies' daily food:

- With 5 puppies, their total is $1 \times 5 = 5$ cups per day.

Step 4: Find the combined daily food:

- Adding the mom and puppies gives $4.5 + 5 = 9.5$ cups per day.

Step 5: Find the total for 6 days: Over 6 days,

- Joy needs $9.5 \times 6 = 57$ cups. The final answer is: 57.

1294

Q: [New Question to be Solved]

1296

1297

1298

1299

1300 Learn from the following examples to answer the given question. For each step, begin by stating the
1301 general plan, followed by executing that plan (started with "-"), separated by a colon.
1302 When the answer is finished, end with "The final answer is:".1303 Q: Karl was attempting to calculate economic figures. He found the following equation to be
1304 true: $fp - w = 10000$. If $f = 5$ and $w = 5 + 125i$, what is p ?1305 A: Step 1: Write down the expression for fp in terms of w :1306 - $fp = 10000 + w$.1307 Step 2: Write down the expression for fp using the value of w :1308 - Substitute $w = 5 + 125i$ into the equation, we get $fp = 10000 + 5 + 125i = 10005 + 125i$.1309 Step 3: Calculate the value of p :1310 - Since $f = 5$, we compute $p = \frac{fp}{f} = \frac{10005+125i}{5} = 2001 + 25i$. The final answer is: $2001 + 25i$.1311 Q: The roots of the equation $2x^2 - mx + n = 0$ sum to 6 and multiply to 10.1312 What is the value of $m + n$?

1313 A: Step 1: Identify the relationships between the roots and coefficients of a quadratic:

1314 - For a quadratic $ax^2 + bx + c = 0$, the sum of the roots is $-\frac{b}{a}$ and the product is $\frac{c}{a}$.1315 Step 2: Use the sum of roots to find m :1316 - Here, $a = 2$ and $b = -m$, so the sum is $-\frac{-m}{2} = \frac{m}{2}$.1317 Setting this equal to 6 gives $\frac{m}{2} = 6 \Rightarrow m = 12$. Step 3: Use the product of roots to find n :1318 - The product is $\frac{n}{2}$, and it's given as 10. So $\frac{n}{2} = 10 \Rightarrow n = 20$.1319 Step 4: Add m and n :1320 - $m + n = 12 + 20 = 32$. The final answer is: 32.

1321

1322 Question: If the two roots of the quadratic $7x^2 + 3x + k$ are $\frac{-3 \pm i\sqrt{299}}{14}$, what is k ?

1323 A: Step 1: Recall the quadratic formula:

1324 - The roots of $ax^2 + bx + c = 0$ are given by $\frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$.

1325 Step 2: Identify coefficients:

1326 - From the equation $7x^2 + 3x + k$, we have $a = 7$, $b = 3$, $c = k$.1327 Step 3: Write the discriminant in terms of k :1328 - $b^2 - 4ac = 3^2 - 4 \cdot 7 \cdot k = 9 - 28k$.

1329 Step 4: Use the root form to determine the discriminant:

1330 - Since the roots are $\frac{-3 \pm i\sqrt{299}}{14}$, the discriminant is -299 . So $9 - 28k = -299 \Rightarrow 28k = 308 \Rightarrow k = 11$.

1331 The final answer is: 11.

1332

1333 Question: Rationalize the denominator of $\frac{1}{\sqrt{2} + \sqrt{3} + \sqrt{7}}$, and write your answer in the form1334 $\frac{A\sqrt{2} + B\sqrt{3} + C\sqrt{7} + D\sqrt{E}}{F}$. What is $A + B + C + D + E + F$?

1335 A: Step 1: Multiply by a conjugate to simplify:

1336 - Multiply numerator and denominator by $(\sqrt{2} + \sqrt{3}) - \sqrt{7}$.

1337 Step 2: Simplify the denominator:

1338 - $(\sqrt{2} + \sqrt{3})^2 = 2 + 3 + 2\sqrt{6} = 5 + 2\sqrt{6}$, and the full denominator becomes $(5 + 2\sqrt{6}) - 7 = -2 + 2\sqrt{6}$.

1339 Step 3: Multiply by the conjugate of the new denominator:

1340 - Multiply numerator and denominator by $-2 - 2\sqrt{6}$: $\frac{\sqrt{2} + \sqrt{3} - \sqrt{7}}{-2 + 2\sqrt{6}} \cdot \frac{-2 - 2\sqrt{6}}{-2 - 2\sqrt{6}} = \frac{4\sqrt{2} + 3\sqrt{3} - \sqrt{7} - \sqrt{42}}{10}$.

1341 Step 4: Identify constants and sum:

1342 - $A = 4$, $B = 3$, $C = -1$, $D = -1$, $E = 42$, $F = 10$,1343 so the total is: $A + B + C + D + E + F = 4 + 3 - 1 - 1 + 42 + 10 = 57$. The final answer is: 57.

1344

1345 Question: For what values of b is -2 not in the range of the function $f(x) = x^2 + bx + 2$?

1346 Express your answer in interval notation.

1347 A: Step 1: Understand when -2 is in the range:1348 - -2 is in the range if there exists some x such that $f(x) = -2$.

1349 Step 2: Set up the equation:

1340 - $f(x) = -2 \Rightarrow x^2 + bx + 2 = -2 \Rightarrow x^2 + bx + 4 = 0$.

1341 Step 3: Determine when the equation has no real solutions:

1342 - A quadratic has no real solutions when the discriminant is negative.

1343 Step 4: Solve for values of b :1344 - $b^2 - 16 < 0 \Rightarrow b^2 < 16 \Rightarrow -4 < b < 4$. The final answer is: $(-4, 4)$.

1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373

Continue:

1374 Question: What is the value of $\sqrt[3]{3^5 + 3^5 + 3^5}$?
1375 A: Step 1: Calculate 3^5 :
1376 - $3^5 = 243$.
1377 Step 2: Add three copies of 3^5 :
1378 - $3^5 + 3^5 + 3^5 = 243 + 243 + 243 = 729$.
1379 Step 3: Take the cube root:
1380 - $\sqrt[3]{729} = 9$ since $9^3 = 729$.
The final answer is: 9.

1381 Q: [New Question to be Solved]
1382

1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403