

# OOLONG: EVALUATING LONG CONTEXT REASONING AND AGGREGATION CAPABILITIES

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  
1040  
1041  
1042  
1043  
1044  
1045  
1046  
1047  
1048  
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  
1906  
1907  
1

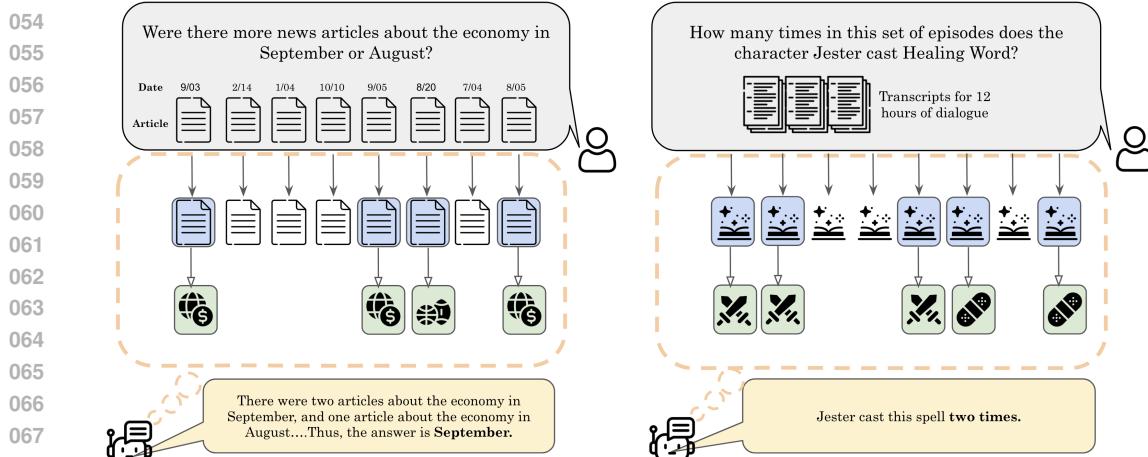


Figure 1: OOLONG poses questions that require performing a multi-step information aggregation process to determine the solution. OOLONG-synth uses ICL-based tasks, which could be easily decomposed and solved iteratively, as a proxy for real-world aggregation tasks over long inputs. OOLONG-real poses challenging information aggregation questions over transcripts from live-action Dungeons & Dragons shows, which can not be easily decomposed into component pieces.

questions about character states and campaign statistics from live-play Dungeons & Dragons role-playing transcripts, using human-annotated gold answers.

Each task requires identifying the relevant segments of the input; this ranges from only a handful of instances or lines of dialogue to questions that require the use of *every* line of the input. The relevant segments must be classified or categorized, and these individual decisions must be aggregated to produce a final answer. By framing the benchmark around problems that are *simple individually*, we ensure that we are measuring capability on long context reasoning and not accuracy on the underlying task. By requiring identification of relevant context, classification decisions over that context, and numeracy skills to produce statistics about the input, OOLONG requires models to perform multi-step reasoning over long inputs.

Although these component capabilities have been well studied and the frontier models demonstrate strong performance in each (Yen et al. (2024), Agarwal et al. (2024), *inter alia*), we find that models struggle with information aggregation as the input length grows. None of the models benchmarked score higher than 50% on OOLONG at 128K context. We study why this task is so challenging using the more flexible OOLONG-synth, ablating settings that decrease context length dramatically and remove the line-by-line classification task, and find that identification and aggregation of information is the bottleneck, not labeling (§4). We also study the role of reasoning behavior and identify consistent error cases in reasoning models that limit their performance in this type of task (§4). We believe OOLONG is a usefully challenging evaluation of long-context reasoning abilities.

## 2 OOLONG-SYNTH

We aim to construct a synthetic information aggregation task that allows fine-grained control over the types of information and number of steps necessary to solve the task. To do this, we need building blocks: simple, short-context documents to build into a longer collection. We construct OOLONG-synth by constructing challenging corpus-level questions over existing in-context learning (ICL) datasets.

### 2.1 DATA

We collect 10 common text classification datasets with between 2 and 10 labels. We select for tasks that are possible for the authors to perform without difficulty (and validate that these tasks are similarly simple for models during our dataset filtering). Table 1 describes the datasets in more

108	Dataset	Task	# Labels	Input Len
109	Spam (Almeida et al., 2011)	SMS spam classification	2	57
110	TREC-QC-coarse (Li & Roth, 2002; Hovy et al., 2001)	Question type classification	6	39
111	AGNews (Zhang et al., 2015)	Headline topic classification	4	90
112	App Reviews (Zur, 2017)	Review sentiment classification	2 <sup>†</sup>	49
113	Pavlick Formality (Lahiri, 2015; Pavlick & Tetreault, 2016)	Formality classification	2 <sup>†</sup>	51
114	IMDB reviews (Maas et al., 2011)	Sentiment analysis	2 <sup>†</sup>	376
115	HiTZ Negation (García-Ferrero et al., 2023)	Verify claims about definitions	2	45
116	Yahoo Topics (Zhang et al., 2015)	Question topic classification	10	74
117	MultiNLI (Williams et al., 2018)	Entailment	3	70
	Metaphors (Bizzoni & Lappin, 2018)	Metaphor meaning validation	2	51

118  
119 Table 1: Datasets used for OOLONG-synth. <sup>†</sup> indicates cases where we combined similar labels to  
120 produce a smaller label-set for OOLONG. The input length is the average per-instance length (in  
121 Llama 2 tokens), including our added date and user metadata.

122  
123 detail. We split the data into two validation tasks and eight test tasks. The validation tasks are  
124 selected to avoid underlying task overlap with test tasks (in the style of Min et al. (2022)), and  
125 so that there is both a 2-label and a multi-label (6-label) validation task. For the sentiment and  
126 formality datasets with more than two labels, we reduce the label space to a single positive and a  
127 single negative label, to reduce the difficulty of the task.

128  
129 **Data filtering** In a normal in-context learning task, if a few examples are unusually hard or even  
130 mislabeled, this has a small impact on the overall score. Because we intend to require aggregation  
131 across many ICL examples at once, one particularly difficult example can affect the score on many  
132 questions downstream. However, many ICL datasets are known to contain mislabelings (Ying &  
133 Thomas, 2022; Chong et al., 2022; Klie et al., 2023).

134  
135 We perform an additional screening step to remove mislabeled or unusually hard examples. We  
136 select two models that we do not expect to be substantially stronger than the models we are eval-  
137 uating: GPT-4.1 nano and Llama 4 Maverick. We perform zero-shot ICL with a minimal instruction  
138 that provides the label space and type of task. We then exclude all examples that *both* models get  
139 incorrect, for exclusion rates ranging from 0% to 0.635% by dataset. We report the percentage of  
140 examples screened out by dataset and provide examples of excluded and validated datapoints for  
141 each dataset in Appendix A. Many of the excluded instances are clearly mislabeled; the remainder  
142 generally represent challenging cases.

## 143 2.2 CONTEXT WINDOW CONSTRUCTION

144  
145 We determine the number of examples in the context window by using an estimate of the number  
146 of tokens required for each example, for the general task instructions, and for the specific question.  
147 Following Yen et al. (2024), we compute these averages using the Llama 2 tokenizer, and use ap-  
148 proximately 500K tokens of context for each estimate. We then estimate the number of examples  
149 to use by determining how many average-length examples would reach 95% of the target context  
150 length.

151  
152 We sample a distribution over the label classes so that the model cannot use any information about  
153 the true distribution over labels (e.g., by guessing that most sentiment classification datasets are  
154 balanced). Then, we sample examples to fill this distribution; if there are less total examples than  
155 required, we sample with replacement. For each instance, we sample a date and user ID. User IDs  
156 are drawn such that 80% of instances have an ID in the 20% of IDs that are most common. Dates  
157 are drawn uniformly with replacement from an approximately 40-month range.

158  
159 Once a context window is constructed, we use it for 25 questions; reusing the same context window  
160 allows for prompt caching, reducing the time to evaluate local models and the cost to evaluate on  
161 most APIs. At each context length, we sample two context windows per dataset, for a total of  
50 questions per dataset or 400 total questions in the test split per context length. We construct  
OOLONG-synth questions for every power of 2 from 1K to 4M.

162 **Questions** We construct three types of OOLONG-synth questions, in order of increasing complexity. *Counting* questions concern simple statistical properties of the label distribution, ranging from identifying the most frequent label to determining the number of examples with each label class. If the label distribution was provided, all tasks in this set should be trivial; these tasks can be seen implicitly as the task of labeling every example and reporting summary statistics. *User information* questions require additional cross-reference with the *user ID* field. *Timeline* questions ask about changes in distribution before or after a certain date, between years, or between months across years. This is more challenging than user information questions because it requires reasoning about *before* or *after* a date, rather than matching a list of IDs (see §B.1 in Appendix for questions).

171 **Instructions** For each dataset, we provide a brief (one-sentence) description of the task and label space, as well as the number of examples in-context. We provide instructions at the start and end of the input, as suggested by OpenAI’s long context prompting guide.<sup>1</sup> However, we provide the question only at the end, to enable prompt caching.

172 2.3 EVALUATION

173 **Baseline** A random baseline is non-trivial for these tasks, because the output space for each question varies. We construct an algorithm for a random baseline with the following rules: (1) in cases where there is an  $n$ -way choice (e.g. choosing between  $n$  labels or choosing a month where some criteria occurs), we select an answer from the set of valid answers at random; (2) if the question requires a numerical answer, we return  $N/|L|$ , where  $N$  is the number of data points in context and  $L$  is the labelset; (3) if the question requires a date or user ID, we sample this from the list of dates/IDs in context at random. We compute the random baseline performance as the expected value of this procedure over the dataset.

174 **Parsing answers** The task prompts specify an output format. We make a best-effort attempt to parse answers using this format; if this template is not present in the output, we take the candidate answer to be the last substring in the output that appears to match the desired answer type (e.g. a comparison or a number). Generally, this occurs if the model runs out of output token budget before providing an answer.

175 **Scoring** For questions that require a label, date, user ID, or comparison (e.g. “greater than” or “less than”), we score on exact match. For questions that require a numerical answer, we set the score to be

$$\text{score}(\hat{y}) = 0.75^{|y - \hat{y}|}$$

176 This allows for partial credit for answers close to the exact value.

177 **On iterative solutions** Although OOLONG-synth allows for a high degree of control over the types of input, these inputs are composed of independent examples from the source dataset. An iterative setup could complete these tasks by calling a model to label each individual example and then deterministically aggregating the resulting labels. We note that this is true for many capability evaluation tasks; for instance, Needle-in-a-Haystack tasks could be solved almost trivially by asking a model if each of the  $N$  input sentences individually contains the needle.

178 **Aside from its merits as a capability evaluation, this setting reflects the way many real-world users interact with language models.** Real user queries tend to be underspecified and rarely use prompt engineering strategies (Sarkar et al., 2025; Xue et al., 2025). When users do edit prompts, they are more likely to add additional context or more detailed instructions, rather than removing or subdividing tasks (Desmond & Brachman, 2024). This is further exacerbated by recent systems supporting uploads of arbitrary size files for analysis (Google, 2025; Anthropic, 2025). We argue posing questions over a large block of context represents a realistic use scenario. To further motivate the information aggregation problem, we turn to a setting where the input cannot be so trivially decomposed.

1<sup>https://cookbook.openai.com/examples/gpt4-1\_prompting\_guide#prompt-organization</sup>

216	<i>Counting</i>
217	Total number of rolls in this episode?
218	What is the count of Crits? (natural rolls of value 1 or 20)?
219	How many {spell type} spells were cast during this episode?
220	How many characters cast {spell name} spell all across episodes?
221	What is the second spell cast in the episode {episode index}?
222	
223	<i>Enumeration</i>
224	What are the first {count} spells cast in this episode? Return a comma separated list.
225	List the last spell cast in each episode? Return a comma separated list.
226	
227	<i>Indexing</i>
228	What is the cumulative total of rolls by the end of episode {episode index}? Count the number of rolls and not the values of the rolls.
229	What is the second spell cast in the episode {episode index}?
230	List the last spell cast in each episode? Return a comma separated list.
231	

Table 2: Question types covered in OOLONG-real dataset.

### 3 OOLONG-REAL

We complement OOLONG-synth with questions derived from real conversational data. OOLONG-real is compiled from the transcripts of a Dungeons and Dragons (D&D) role-playing game, where a group of players collaboratively build a story through in-character actions and success depends on rolling dice. Stories unfold over narratives (“campaigns”) that span dozens to hundreds of episodes, with each episode generally lasting 4-5 hours of play. These transcripts involve several levels of conversation, from out-of-character chitchat to rules discussion to in-character actions and speech. Though lightly edited for readability, they reflect naturalistic speech instead of carefully planned written text. Because the conversation is unscripted and can involve tangents or side channels, conversational turns require variable amounts of prior context to resolve. In some instances, the same event (e.g. casting a spell) is discussed for many turns, or brought up again after a long interlude; in other cases, a prior event is “retconned” or revised post hoc.

In Dungeons and Dragons, there are limitations to how frequently characters can take certain actions (e.g., cast certain types of spells). Additionally, fans are often interested in metadata of the play (e.g., whether a certain person is particularly unlucky with their dice rolls or if a character uses a signature spell more or less as the campaign progresses). Because these shows are extremely long and wide-ranging and improvised live, the creators do not plan for or provide this type of information. In lieu of an official source, this information is often annotated meticulously, and with multiple levels of verification, by dedicated fans of the work.

We take this as an example of an information aggregation task in the wild that is *not* simple to reframe as an iterative task. We consider the series Critical Role and the data compiled by CritRoleStats,<sup>2</sup> a fan project that tracked per-episode statistics for the first several campaigns of Critical Role. We devise questions related to characters, dice rolls, and spells cast during episodes and use human-labeled CritRoleStats to compute gold answers.

#### 3.1 DATASET COMPILATION

For OOLONG-real, we consider two campaigns from the Critical Role TV series. We used episode transcripts from the Critical Role Dungeons and Dragons Dataset (CRD3) (Rameshkumar & Bailey, 2020), which includes full episode transcripts from the first two campaigns of the Critical Role TV series. For our testbed, we used the first campaign, which consists of 115 episodes. Each line in the transcript includes an utterance with the player name labeled.

<sup>2</sup><https://www.critrolestats.com/>

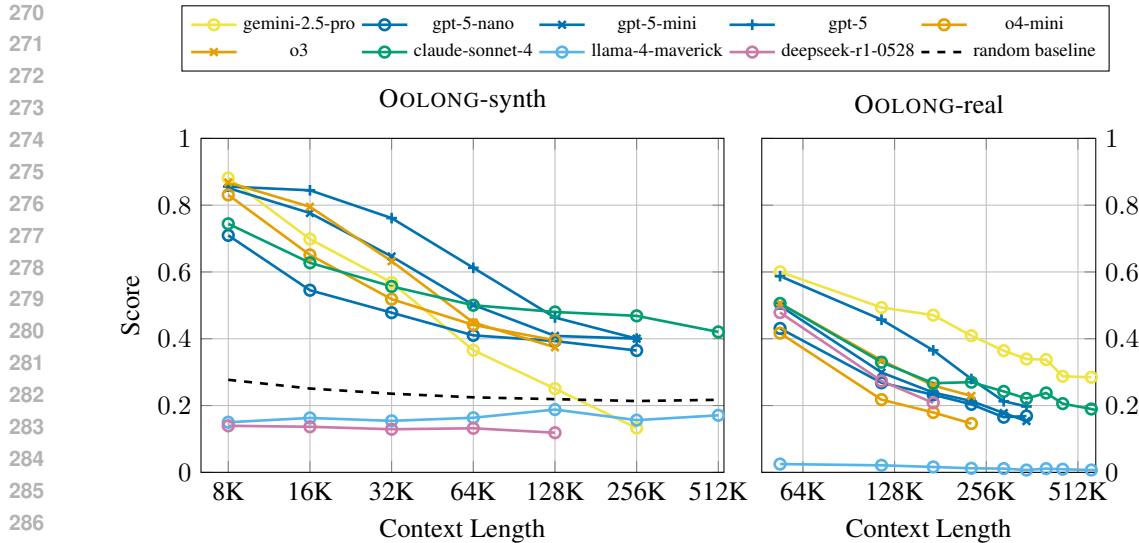


Figure 2: Scores by context window length for OOLONG-synth and OOLONG-real.

**QA pairs** We use the game statistics compiled by the authors of CritRoleStats. Specifically, we utilize statistics about dice rolls and spells cast in each episode. We design a set of questions that cover a variety of information aggregation tasks; see Table 2 for some examples. We include questions that require processing of single- or multiple-episode transcripts. For multi-episode questions, we concatenate transcripts and use delimiters to highlight the start and end of each transcript. See Table 19 and Table 20 of the Appendix for a full list of OOLONG-real questions. We include varying context windows to evaluate the model’s ability to use long-context reasoning and aggregation capabilities. Using a single episode transcript as a context unit, we include context windows ranging from one to 24 episode transcripts. This covers input lengths of 55K to 1.3M tokens.

### 3.2 EVALUATION

We closely follow the evaluation setup from OOLONG-synth.

**Parsing** Our task prompt requires the model to place the answer in `\boxed{}`. If the answer cannot be extracted successfully, we attempt to extract the answer for a given question using GPT-5-nano.

**Scoring** OOLONG-real contains three types of answers: numeric, string, and a list of strings. For numerical answers, we use the same scoring scheme as OOLONG-synth to allow partial credit. We use exact match for string answer types and set overlap for answers of a type list.

## 4 RESULTS AND ANALYSIS

On OOLONG, we benchmark a strong suite of frontier models to study their information aggregation capabilities over long context. We include a mix of models of varying sizes, levels of reasoning, and long-context capabilities.

In Figure 2, we visualize the model performance for context windows up to 512K tokens. As expected, we see a significant drop in performance at higher context windows. The two splits are of similar difficulty at the same context length, although direct comparison is challenging because real data do not necessarily align with fixed context window buckets.<sup>3</sup> A breakdown of results with an aggregate rating over the two datasets is available in the Appendix (Table 4).

<sup>3</sup>While performance on OOLONG-real starts lower, this is because the shortest inputs in this dataset are a single episode, with an average length of 55K tokens; models perform similarly on OOLONG-synth around the same context length.

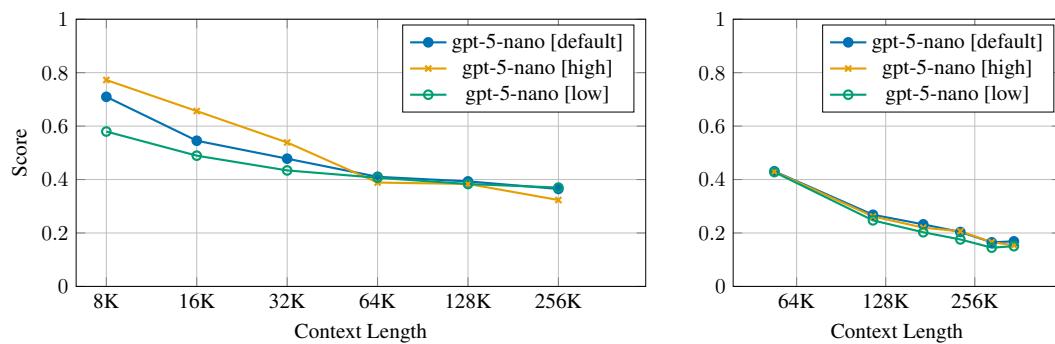


Figure 3: Comparison across reasoning levels.

#### 4.1 IMPACT OF REASONING LEVEL

We compare “high” and “low” reasoning effort for GPT-5-nano. Figure 3 shows the results for both the splits of OOLONG. Although OOLONG is a reasoning-intensive task, specifying a higher reasoning effort is only useful for short contexts; after 64k, there is little discernible difference between reasoning levels, with the “high” reasoning setting even slightly underperforming “low” reasoning at 256k on OOLONG-synth. Although it is difficult to draw firm conclusions without the ability to view the reasoning trace, we hypothesize that at context lengths where there is sufficient remaining room in the context window to enumerate labels for each example in-context, adding more reasoning effort may encourage the model to take this strategy. However, since the default routing for this model performs reasonably well at all lengths and especially for longer inputs, we do not explicitly specify a reasoning level for the remainder of the runs.

#### 4.2 WHAT TYPES OF QUESTIONS ARE CHALLENGING?

On OOLONG-synth, we further break down results by question type (counting, user, or temporal) and answer type (the expected format of the answer, e.g. as a label or number or date) in Figure 7. Consistently, temporal questions are the most challenging for models, highlighting the difficulty of temporal reasoning. Figure 7 shows that temporal questions are challenging for models; this is reflected in both the performance by task type and the distribution of score by answer type. Questions that require a date or month/year (e.g. “January 2021”) as an answer show generally lower performance for the same model, and show greater spread in model capabilities than the other answer categories. For instance, the gap between GPT-5 and GPT-5-nano performance is more than 4x larger for questions that require outputting a date than for questions that require outputting a label. Relative model performance is mostly stable across answer types, although Claude-Sonnet-4 is relatively much stronger on numerical reasoning and comparisons than the GPT series models.

On OOLONG-real, we divide the questions into ten categories: counting, enumeration, frequency, indexing, player, character, roll type/value, spell level/type (Table 19, Table 20). Across models, we find ‘frequency’ to be the hardest and ‘spell type’ the easiest (Figure 8). Except for Gemini-2.5-Pro and GPT-5, all models struggle with enumeration and indexing questions. Gemini-2.5 Pro and GPT-5 also make effective use of the player-character mapping provided in the system prompt.

#### 4.3 DIFFERENCES BETWEEN OOLONG-SYNTH AND OOLONG-REAL

For most models, the performance trends are consistent across the two datasets. We examine two notable exceptions in more detail.

**Gemini 2.5 Pro** Gemini 2.5 Pro is the strongest model on OOLONG-real, but exhibits performance dropoff at longer contexts on OOLONG-synth. We observe that the model increasingly exceeds its maximum output length during reasoning for long OOLONG-synth inputs. Unlike the other APIs used in benchmarking, the Gemini API does not return any tokens if the max token count is exceeded during reasoning; this results in an automatic score of 0 for any cases of overlong reasoning, whereas

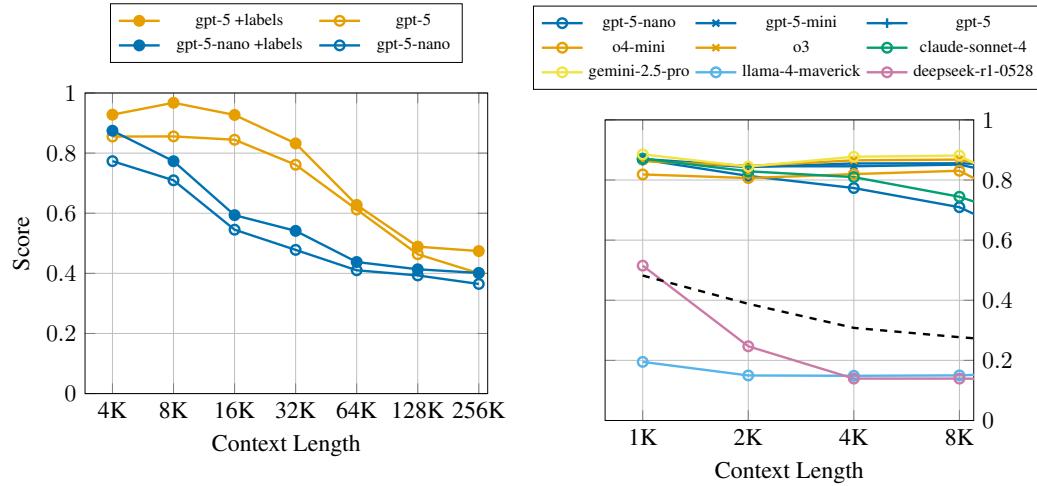


Figure 4: Comparison on OOLONG-synth: (a) we provide the gold labels in the input. This leads to a consistent but small improvement, (b) short context performance; while the top models have similar short-context performance, differences emerge as the context length grows.

a truncated output returned from another model might still score partial credit. Gemini also will return an empty output if the “recitation” filter (indicating substantial regurgitation of pretraining data) is triggered. Because copies of some parts of these inputs (e.g. the IMDB movie reviews) are prevalent on the web, this filter is occasionally triggered when running OOLONG-synth, particularly at longer contexts. We do not observe either of these behaviors on OOLONG-real, and as a result Gemini maintains its strong performance at longer context lengths on this split.

**Deepseek R1** Deepseek R1 is a strong reasoning model and outperforms GPT-5 nano and o4-mini on OOLONG-real. However, on OOLONG-synth, it achieves performance below the random baseline. What causes this discrepancy?

In manual inspection of the traces, we observe pathologies in the reasoning traces for OOLONG-synth that are not present in traces for OOLONG-real and appear only occasionally in traces from other models. We identify these behaviors through discussion between the authors, then prompt GPT-5 nano to label 2,400 randomly selected traces from Deepseek-R1 with a yes/no classification about whether an identified behavior is present in this trace. Many traces (60%) do not provide an answer at all. We hypothesize that, because OOLONG-synth is such an information-dense task, the model’s apparent strategy of labeling each example before deciding which are relevant results in running out of context tokens. 64% of traces end in an incomplete sentence.<sup>4</sup> In 17% of cases, the model spends at least some time debating whether the task is impossible or intractable given the length of the input, and in 4% of cases it refuses to respond completely. Appendix C shows example traces for Deepseek-R1 from both OOLONG-synth and OOLONG-real.

#### 4.4 SIMPLER SETTINGS FOR OOLONG-SYNTH

**Shorter context** Most models show declining performance with context length. We consider a short-context version of OOLONG-synth, with inputs between 1k and 4k tokens, in Figure 4. While models perform better on this task, several still struggle, showing that even short-context aggregation remains challenging. Performance differences between the top models are difficult to distinguish in the short-context regime, suggesting that these models have the ability to perform the task at some context length. However, no model exceeds 85% performance at any context length.

<sup>4</sup>Note that some of these incomplete traces do contain an answer: in a small number of cases, the model provides a candidate answer, then begins to double-check the answer and then exceeds the maximum output length.

432 **Aggregation without classification** For OOLONG-synth, we can construct an easier version of  
 433 the task by providing the label for each ICL example in-context. This reduces the task for most  
 434 questions to simply identifying the relevant instances and summing the occurrence of each label  
 435 type. Figure 4 shows the behavior of GPT-5 and GPT-5-nano with and without labels provided in-  
 436 context. As expected, adding labels improves accuracy; however, this improvement ranges from  
 437 10.9 points to only 0.79 points. We do not see consistently higher gains from providing labels in  
 438 longer inputs, which suggests that the lower performance at longer context lengths is not primarily  
 439 due to an accumulation of mislabeling errors.

440 Additionally, the improvement is not larger for GPT-5-nano than it is for the more powerful GPT-5,  
 441 suggesting that the performance difference observed between these models is not due to differing  
 442 ability to perform the classification task. This is by design; our classification task validation was  
 443 designed to eliminate overly challenging or misleading examples, as the ability to aggregate infor-  
 444 mation is the main capability we aim to measure.

## 445 5 RELATED WORK

446 **Long-context benchmarks** A variety of long-context benchmarks have included some type of  
 447 aggregation or reasoning-focused task. RULER (Hsieh et al., 2024) benchmark includes synthetic  
 448 retrieval tasks at varying context lengths. It includes multi-hop tracing and aggregation tasks. HEL-  
 449 MET (Yen et al., 2024) expands the tasks in RULER to include tasks related to the downstream use  
 450 of LMs (reranking, ICL, LongQA, summarization). This includes a set of long-context question-  
 451 answering tasks from Zhang et al. (2024). LongMemEval (Wu et al., 2025) and multi-round core-  
 452 ference resolution tasks (MRCR) evaluate long-context capabilities with conversational data. MRCR  
 453 was first introduced in Gemini (Vodrahalli et al., 2024) and was further extended in OpenAI’s MRCR  
 454 (OpenAI, 2025). Unlike the standard NIAH task, the needles and the distractors in MRCR are sam-  
 455 pled from the same distributions. OpenAI’s MRCR includes variants of 2, 4, and 8 needles, and the  
 456 task involves retrieving the  $i$ th instance of one of the needles. The documents in MRCR are syn-  
 457 synthetic conversations generated using GPT-4o. BABILong (Kuratov et al., 2024) tests long-context  
 458 reasoning by placing BABI (Weston et al., 2015) reasoning tasks within a long context and eval-  
 459 uating performance as the amount of distractor text increases; however, this cannot vary the amount  
 460 of information needed to perform the reasoning task, only the amount of distraction present, and  
 461 thus could be solved by a filter-then-reason approach. TLDM (Hamilton et al., 2025) includes two  
 462 related tasks about character location and time passed in a narrative setting, but only considers  
 463 public-domain novels, which likely suffer from data contamination from online discussion of the  
 464 texts Palavalli et al. (2024). ZeroSCROLLS (Shaham et al., 2023) introduces a task of identifying  
 465 the percentage of reviews on Amazon for a product that are positive, which is the closest conceptual  
 466 ancestor of our work; OOLONG-synth encompasses this type of task along with other aggregation  
 467 tasks. GSM-infinite (Zhou et al., 2025) stress tests long-context reasoning through an adaptable  
 468 framework for varying task difficulty and length; our work is complementary, as their focus is pri-  
 469 marily on increasingly difficult mathematical reasoning through synthetically generated problems,  
 470 while we focus on a more noisy, naturalistic text setting with less challenging mathematical rea-  
 471 soning required. MoNaCo (Wolfson et al., 2025) considers challenging information seeking queries  
 472 that require retrieval and aggregation; their emphasis is on retrieval of the necessary documents  
 473 for intermediate reasoning steps in an agentic setting, a setting completely disjoint from OOLONG.  
 474 Finally, procedural generation benchmarks (Ye et al., 2025) measure a different type of reasoning  
 475 task by evaluating the ability to plan and produce coherent long outputs. **Table 3 summarizes the**  
 476 **distinctions between OOLONG and the most similar prior long-context reasoning benchmarks across**  
 477 **a range of desiderata.**

478 **Dungeons & Dragons data** Several prior works consider D&D as a potential testbed for model ca-  
 479 pabilities; for instance, we use a version of the Critical Role transcripts preprocessed by Rameshku-  
 480 mar & Bailey (2020), who proposed an episode summarization task using fan-written summaries.  
 481 D&D data has also been used to study theory of mind (Zhou et al., 2023), user assistant development  
 482 Zhu et al. (2023b), and dialogue generation, with game state either inferred from forum-based games  
 483 (Callison-Burch et al., 2022) or captured through an external tool (Zhu et al., 2023a); to the best of  
 484 our knowledge, we are the first to use fan annotations of gold labels and to consider the generation  
 485 of these statistics as a task in its own right.

Dataset	Real vs synthetic data			Aggregation capabilities		Measurability	
	Realistic inputs	Realistic questions	Relevant text matches distractors	Multi-step reasoning	Numeracy	Must use full input	Easy to eval
MRCR	✓	✗	✓	✓	✓	✗	✓
BABILong	✗	✓	✗	✓	✗	✗	✓
RULER	✗	✗	✗	†	†	†	✓
HELMET Summ	✓	✓	✓	✓	✗	✓	✗
HELMET LongQA	✓	✓	✓	✓	✗	†	✓
Oolong-synth	✗	✓	✓	✓	✓	✓	✓
Oolong-real	✓	✓	✓	✓	✓	†	✓

Legend: ✓ = yes, ✗ = no, † = in some cases

Table 3: Comparing long-context reasoning benchmarks. OOLONG measures aggregation capabilities in controllable, realistic, and simple-to-evaluate settings.

**Aggregation as an NLP concept** Goldman et al. (2024) argue for the definition of long context tasks in terms of the information dispersion and input scope required to find the answer; under this taxonomy, OOLONG is high dispersion (because relevant info is distribution over the full context length) and high scope (because most of the input is necessary for the task, particularly for OOLONG-synth). DeYoung et al. (2024) measure synthesis in multi-document summarization by asking if summaries convey the consensus opinion of the input documents; this is a related task to the counting tasks in OOLONG-synth, but framed as a measure of summarization capabilities instead of long context capabilities, and thus requires different affordances related to evaluating generated outputs.

## 6 CONCLUSION

We introduce OOLONG, a challenging long-context information aggregation benchmark in two parts. OOLONG-synth uses synthetic aggregation tasks over ICL data to enable finer-grained control of the benchmark settings, while OOLONG-real poses questions over real long-context conversational data and human-annotated labels. On both splits, models struggle, with performance dropping with increasing context length even when controlling for the potential compounding of mislabeling errors.

The low performance on OOLONG suggests more work is necessary in developing both reasoning abilities and more general long context abilities. Even in short context regimes (8K tokens or less), strong models cannot perfectly perform the multi-step reasoning tasks in OOLONG-synth. However, in longer context regimes, many models perform roughly equivalently on OOLONG-synth and OOLONG-real at the same context length, despite the much higher information density in OOLONG-synth. This suggests that better length generalization, regardless of reasoning ability, may also improve performance on this type of challenging long context aggregation task.

However, some models behave differently in extremely information-dense regimes. In particular, both Gemini-2.5-Pro and Deepseek-R1-0528 are strong models that perform well on OOLONG-real but fail in OOLONG-synth because of an over-reasoning phenomenon. This highlights a direction for future work. While prior work on reasoning chains has focused on over-reasoning for adversarial (Kumar et al., 2025) or overly simple problems (Chiang & yi Lee, 2024; Sui et al., 2025), OOLONG indicates that even strong models may struggle to plan the reasoning quantity for information-dense inputs, where it may be desirable to accept a less optimal or more error-prone reasoning strategy (e.g. double-checking less frequently) in order to avoid running into the maximum reasoning tokens. Future models that are aware of the maximum reasoning token budget may be able to more carefully plan reasoning strategies according to the allowed budget.

We see substantial headroom between strong open weights models and API-based models on this task. Overall, our results suggest that there is still a long way to go in designing robust long-context aggregation capabilities for LLMs.

540  
541 **ETHICS STATEMENT**542 This work is intended to further the evaluation of long-context LLMs. We do not foresee any par-  
543 ticular additional risks introduced by our evaluation, although information aggregation, like nearly  
544 any LM capability, could be useful for both beneficial and harmful actors.545 CritRoleStats, whose data we use as part of constructing gold labels for the OOLONG-real split, ex-  
546 plicitly allows the use of their data for statistical analyses in their FAQ, so long as they are credited.<sup>5</sup>  
547 We have also notified them of our use of their data.548 In the course of this work, we used LM assistance in the writing of some data preprocessing, analy-  
549 sis, and visualization scripts.551  
552 **REPRODUCIBILITY STATEMENT**  
553554 We release code for dataset construction and evaluation, in addition to the final dataset version. API-  
555 based models typically exhibit some small variability across API calls, and versions of each model  
556 may be phased out over time by the providers, so exact replication of outputs may be infeasible;  
557 however, we will release the model outputs for each model run reported in the paper.559  
560 **REFERENCES**

561 Software applications user reviews. 2017.

562 Rishabh Agarwal, Avi Singh, Lei M. Zhang, Bernd Bohnet, Luis Rosias, Stephanie Chan, Biao  
563 Zhang, Ankesh Anand, Zaheer Abbas, Azade Nova, John D. Co-Reyes, Eric Chu, Feryal Be-  
564 hbahani, Aleksandra Faust, and Hugo Larochelle. Many-shot in-context learning, 2024. URL  
565 <https://arxiv.org/abs/2404.11018>.566 Tiago A. Almeida, Jose Maria Gomez Hidalgo, and Akebo Yamakami. Contributions to the study  
567 of sms spam filtering: New collection and results. In *Proceedings of the 2011 ACM Symposium  
568 on Document Engineering (DOCENG'11)*, 2011.569 Anthropic. What kinds of documents can i upload to  
570 claude? <https://support.claude.com/en/articles/8241126-what-kinds-of-documents-can-i-upload-to-claude>, 2025.  
571 Accessed: 2025-11-04.572 Amanda Bertsch, Maor Ivgi, Emily Xiao, Uri Alon, Jonathan Berant, Matthew R Gormley, and  
573 Graham Neubig. In-context learning with long-context models: An in-depth exploration. *arXiv  
574 preprint arXiv:2405.00200*, 2024.575 Yuri Bizzoni and Shalom Lappin. Predicting human metaphor paraphrase judgments with deep  
576 neural networks. In Beata Beigman Klebanov, Ekaterina Shutova, Patricia Lichtenstein, Smaranda  
577 Muresan, and Chee Wee (eds.), *Proceedings of the Workshop on Figurative Language Processing*,  
578 pp. 45–55, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi:  
579 10.18653/v1/W18-0906. URL <https://aclanthology.org/W18-0906/>.580 Chris Callison-Burch, Gaurav Singh Tomar, Lara Martin, Daphne Ippolito, Suma Bailis, and David  
581 Reitter. Dungeons and dragons as a dialog challenge for artificial intelligence. In *Proceedings  
582 of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 9379–9393.  
583 Association for Computational Linguistics, 2022. doi: 10.18653/v1/2022.emnlp-main.637. URL  
584 <http://dx.doi.org/10.18653/v1/2022.emnlp-main.637>.585 Cheng-Han Chiang and Hung yi Lee. Over-reasoning and redundant calculation of large language  
586 models, 2024. URL <https://arxiv.org/abs/2401.11467>.587 Derek Chong, Jenny Hong, and Christopher D. Manning. Detecting label errors by using pre-trained  
588 language models, 2022. URL <https://arxiv.org/abs/2205.12702>.5<sup>5</sup><https://www.critrolestats.com/faqs>

594 Michael Desmond and Michelle Brachman. Exploring prompt engineering practices in the enter-  
 595 prise, 2024. URL <https://arxiv.org/abs/2403.08950>.

596

597 Jay DeYoung, Stephanie C. Martinez, Iain J. Marshall, and Byron C. Wallace. Do multi-document  
 598 summarization models synthesize? *Transactions of the Association for Computational Linguistics*,  
 599 12:1043–1062, 2024. doi: 10.1162/tacl\_a\_00687. URL <https://aclanthology.org/2024.tacl-1.58/>.

600

601 Iker García-Ferrero, Begoña Altuna, Javier Alvez, Itziar Gonzalez-Dios, and German Rigau. This  
 602 is not a dataset: A large negation benchmark to challenge large language models. In Houda  
 603 Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empir-  
 604 ical Methods in Natural Language Processing*, pp. 8596–8615, Singapore, December 2023.  
 605 Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.531. URL  
 606 <https://aclanthology.org/2023.emnlp-main.531>.

607

608 Omer Goldman, Alon Jacovi, Aviv Slobodkin, Aviya Maimon, Ido Dagan, and Reut Tsarfaty. Is it  
 609 really long context if all you need is retrieval? towards genuinely difficult long context NLP. In  
 610 Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference  
 611 on Empirical Methods in Natural Language Processing*, pp. 16576–16586, Miami, Florida, USA,  
 612 November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.  
 613 924. URL <https://aclanthology.org/2024.emnlp-main.924/>.

614

615 Google. Google notebooklm: Ai research tool & thinking partner. [https://notebooklm.  
 616 google/](https://notebooklm.google/), 2025. Accessed: 2025-11-04.

617

618 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu  
 619 Zhang, Shirong Ma, Xiao Bi, et al. Deepseek-r1 incentivizes reasoning in llms through reinforce-  
 620 ment learning. *Nature*, 645(8081):633–638, 2025.

621

622 Sil Hamilton, Rebecca M. M. Hicke, Matthew Wilkens, and David Mimno. Too long, didn't model:  
 623 Decomposing llm long-context understanding with novels, 2025. URL [https://arxiv.  
 624 org/abs/2505.14925](https://arxiv.org/abs/2505.14925).

625

626 Eduard Hovy, Laurie Gerber, Ulf Hermjakob, Chin-Yew Lin, and Deepak Ravichandran. To-  
 627 ward semantics-based answer pinpointing. In *Proceedings of the First International Confer-  
 628 ence on Human Language Technology Research*, 2001. URL <https://www.aclweb.org/anthology/H01-1069>.

629

630 Cheng-Ping Hsieh, Simeng Sun, Samuel Kriman, Shantanu Acharya, Dima Rekesh, Fei Jia, and  
 631 Boris Ginsburg. RULER: What's the real context size of your long-context language models? In  
 632 *First Conference on Language Modeling*, 2024. URL [https://openreview.net/forum?  
 633 id=kIoBbc76Sy](https://openreview.net/forum?id=kIoBbc76Sy).

634

635 Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik  
 636 Narasimhan. Swe-bench: Can language models resolve real-world github issues? *arXiv preprint  
 637 arXiv:2310.06770*, 2023.

638

639 Jan-Christoph Klie, Bonnie Webber, and Iryna Gurevych. Annotation error detection: Analyzing  
 640 the past and present for a more coherent future. *Computational Linguistics*, 49(1):157–198, 03  
 641 2023. ISSN 0891-2017. doi: 10.1162/coli\_a\_00464. URL [https://doi.org/10.1162/coli\\_a\\_00464](https://doi.org/10.1162/coli_a_00464).

642

643 Abhinav Kumar, Jaechul Roh, Ali Naseh, Marzena Karpinska, Mohit Iyyer, Amir Houmansadr, and  
 644 Eugene Bagdasarian. Overthink: Slowdown attacks on reasoning llms, 2025. URL <https://arxiv.org/abs/2502.02542>.

645

646 Yuri Kuratov, Aydar Bulatov, Petr Anokhin, Ivan Rodkin, Dmitry Sorokin, Artyom Sorokin, and  
 647 Mikhail Burtsev. Babilong: Testing the limits of llms with long context reasoning-in-a-haystack,  
 648 2024. URL <https://arxiv.org/abs/2406.10149>.

649

650 Shibamouli Lahiri. SQUINKY! A corpus of sentence-level formality, informativeness, and implica-  
 651 ture. *arXiv preprint arXiv:1506.02306*, 2015.

648 Xin Li and Dan Roth. Learning question classifiers. In *COLING 2002: The 19th International Conference on Computational Linguistics*, 2002. URL <https://www.aclweb.org/anthology/C02-1150>.

649

650

651

652 Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pp. 142–150, Portland, Oregon, USA, June 2011. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/P11-1015>.

653

654

655

656

657 Sewon Min, Mike Lewis, Luke Zettlemoyer, and Hannaneh Hajishirzi. Metaicl: Learning to learn in context, 2022. URL <https://arxiv.org/abs/2110.15943>.

658

659

660 OpenAI. OpenAI MRCR: Long context multiple needle in a haystack benchmark, 2025. URL <https://huggingface.co/datasets/openai/mrcr>.

661

662 Medha Palavalli, Amanda Bertsch, and Matthew R. Gormley. A taxonomy for data contamination in large language models, 2024. URL <https://arxiv.org/abs/2407.08716>.

663

664

665 Ellie Pavlick and Joel Tetreault. An empirical analysis of formality in online communication. *Transactions of the Association for Computational Linguistics*, 2016.

666

667 Revanth Rameshkumar and Peter Bailey. Storytelling with dialogue: A Critical Role Dungeons and Dragons Dataset. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 5121–5134, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.459. URL <https://aclanthology.org/2020.acl-main.459>.

668

669

670

671

672 Rupak Sarkar, Bahareh Sarrafzadeh, Nirupama Chandrasekaran, Nagu Rangan, Philip Resnik, Longqi Yang, and Sujay Kumar Jauhar. Conversational user-ai intervention: A study on prompt rewriting for improved llm response generation, 2025. URL <https://arxiv.org/abs/2503.16789>.

673

674

675

676

677 Uri Shaham, Maor Ivgi, Avia Efrat, Jonathan Berant, and Omer Levy. Zeroscrolls: A zero-shot benchmark for long text understanding, 2023. URL <https://arxiv.org/abs/2305.14196>.

678

679

680 Yang Sui, Yu-Neng Chuang, Guanchu Wang, Jiamu Zhang, Tianyi Zhang, Jiayi Yuan, Hongyi Liu, Andrew Wen, Shaochen Zhong, Na Zou, Hanjie Chen, and Xia Hu. Stop overthinking: A survey on efficient reasoning for large language models, 2025. URL <https://arxiv.org/abs/2503.16419>.

681

682

683

684 Kiran Vodrahalli, Santiago Ontanon, Nilesh Tripuraneni, Kelvin Xu, Sanil Jain, Rakesh Shivanna, Jeffrey Hui, Nishanth Dikkala, Mehran Kazemi, Bahare Fatemi, Rohan Anil, Ethan Dyer, Siamak Shakeri, Roopal Vij, Harsh Mehta, Vinay Ramasesh, Quoc Le, Ed Chi, Yifeng Lu, Orhan Firat, Angeliki Lazaridou, Jean-Baptiste Lespiau, Nithya Attaluri, and Kate Olszewska. Michelangelo: Long context evaluations beyond haystacks via latent structure queries, 2024. URL <https://arxiv.org/abs/2409.12640>.

685

686

687

688

689

690

691 Jason Weston, Antoine Bordes, Sumit Chopra, Alexander M. Rush, Bart van Merriënboer, Armand Joulin, and Tomas Mikolov. Towards ai-complete question answering: A set of prerequisite toy tasks, 2015. URL <https://arxiv.org/abs/1502.05698>.

692

693

694

695

696

697

698

699

700 Adina Williams, Nikita Nangia, and Samuel Bowman. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 1112–1122. Association for Computational Linguistics, 2018. URL <http://aclweb.org/anthology/N18-1101>.

701

702 Tomer Wolfson, Harsh Trivedi, Mor Geva, Yoav Goldberg, Dan Roth, Tushar Khot, Ashish Sabharwal, and Reut Tsarfaty. Monaco: More natural and complex questions for reasoning across dozens of documents. *Transactions of the Association for Computational Linguistics*, 2025.

702 Di Wu, Hongwei Wang, Wenhao Yu, Yuwei Zhang, Kai-Wei Chang, and Dong Yu. Longmemeval:  
 703 Benchmarking chat assistants on long-term interactive memory. In *The Thirteenth International*  
 704 *Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=pZiyCaVuti>.

705

706 Haoning Xue, Yoo Jung Oh, Xinyi Zhou, Xinyu Zhang, and Berit Oxley. User prompting strategies  
 707 and chatgpt contextual adaptation shape conversational information-seeking experiences, 2025.  
 708 URL <https://arxiv.org/abs/2509.25513>.

709

710 Xi Ye, Fangcong Yin, Yinghui He, Joie Zhang, Howard Yen, Tianyu Gao, Greg Durrett, and Danqi  
 711 Chen. Longproc: Benchmarking long-context language models on long procedural generation,  
 712 2025. URL <https://arxiv.org/abs/2501.05414>.

713

714 Howard Yen, Tianyu Gao, Minmin Hou, Ke Ding, Daniel Fleischer, Peter Izsak, Moshe Wasserblat,  
 715 and Danqi Chen. Helmet: How to evaluate long-context language models effectively and thor-  
 oughly. *arXiv preprint arXiv:2410.02694*, 2024.

716

717 Cecilia Ying and Stephen Thomas. Label errors in BANKING77. In Shabnam Tafreshi, João Sedoc,  
 718 Anna Rogers, Aleksandr Drozd, Anna Rumshisky, and Arjun Akula (eds.), *Proceedings of the*  
 719 *Third Workshop on Insights from Negative Results in NLP*, pp. 139–143, Dublin, Ireland, May  
 720 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.insights-1.19. URL  
 721 <https://aclanthology.org/2022.insights-1.19/>.

722

723 Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. Character-level convolutional networks for text  
 724 classification. In *NIPS*, 2015.

725

726 Xinrong Zhang, Yingfa Chen, Shengding Hu, Zihang Xu, Junhao Chen, Moo Khai Hao, Xu Han,  
 727 Zhen Leng Thai, Shuo Wang, Zhiyuan Liu, and Maosong Sun.  $\infty$ bench: Extending long context  
 728 evaluation beyond 100k tokens, 2024. URL <https://arxiv.org/abs/2402.13718>.

729

730 Pei Zhou, Andrew Zhu, Jennifer Hu, Jay Pujara, Xiang Ren, Chris Callison-Burch, Yejin Choi, and  
 731 Prithviraj Ammanabrolu. I cast detect thoughts: Learning to converse and guide with intents and  
 732 theory-of-mind in dungeons and dragons, 2023. URL <https://arxiv.org/abs/2212.10060>.

733

734 Yang Zhou, Hongyi Liu, Zhuoming Chen, Yuandong Tian, and Beidi Chen. Gsm-infinite: How do  
 735 your llms behave over infinitely increasing context length and reasoning complexity?, 2025. URL  
<https://arxiv.org/abs/2502.05252>.

736

737 Andrew Zhu, Karmanya Aggarwal, Alexander Feng, Lara J. Martin, and Chris Callison-Burch.  
 738 FIREBALL: A dataset of dungeons and dragons actual-play with structured game state infor-  
 739 mation. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Proceedings of the*  
 740 *61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Pa-  
 741 pers)*, pp. 4171–4193, Toronto, Canada, July 2023a. Association for Computational Linguis-  
 742 tics. doi: 10.18653/v1/2023.acl-long.229. URL [https://aclanthology.org/2023.acl-long.229/](https://aclanthology.org/2023.acl-long.229).

743

744 Andrew Zhu, Lara Martin, Andrew Head, and Chris Callison-Burch. Calypso: Llms as dungeon  
 745 master’s assistants. *Proceedings of the AAAI Conference on Artificial Intelligence and Interac-  
 746 tive Digital Entertainment*, 19(1):380–390, Oct. 2023b. doi: 10.1609/aiide.v19i1.27534. URL  
 747 <https://ojs.aaai.org/index.php/AIIDE/article/view/27534>.

748

749

750

751

752

753

754

755

Model	Avg.	OOLONG-synth					OOLONG-real			
		8K	16K	32K	64K	128K	Avg.	55K	118K	175K
GPT-5	<b>70.75</b>	85.56	<b>84.45</b>	<b>76.12</b>	<b>61.24</b>	46.36	47.00	58.74	45.72	36.53
Gemini-2.5-Pro	55.29	<b>88.13</b>	69.84	56.83	36.56	25.06	<b>52.16</b>	<b>60.04</b>	<b>49.36</b>	<b>47.07</b>
o3	62.37	86.80	79.52	63.23	44.86	37.45	36.71	50.57	33.57	25.99
GPT-5-mini	63.68	85.13	77.65	64.64	50.14	40.85	34.55	49.86	29.90	23.89
Claude-Sonnet-4	58.18	74.43	62.75	55.67	50.04	<b>48.02</b>	36.75	50.58	32.98	26.70
o4-mini	56.74	83.07	65.10	51.86	44.15	39.53	27.13	41.69	21.77	17.93
GPT-5-nano	50.73	70.96	54.53	47.81	41.02	39.31	31.05	43.09	26.82	23.23
Deepseek-R1	13.11	13.94	13.65	12.91	13.20	11.87	32.00	47.85	27.35	20.81
Llama-4-Maverick	16.37	15.00	16.29	15.42	16.35	18.80	2.07	2.48	2.11	1.62

Table 4: OOLONG results on a number of strong models. All models we test support at least 200K context; thus, we report OOLONG scores as an average over scores on 8K-175K inputs. OOLONG-synth is a more information-dense task than OOLONG-real; the ranking of some models shifts slightly between the two settings. Models are sorted by the average between the two benchmarks.

Dataset	% Removed
Spam	0.635%
TREC-Q-coarse	0.048%
AGNews	0.026%
App Reviews	0.051%
Formality	0.108%
IMDB	0.042%
Metaphors	0.000%
MultiNLI	0.109%
Negation	0.016%
Yahoo Topics	0.188%

Table 5: Only a small fraction of examples are removed during validation for each source dataset.

## A ICL LABEL VALIDATION

Table 1 shows the fraction of examples removed by the validation step for each dataset.

For each ICL dataset used to construct OOLONG-synth, we provide example instances that failed validation (and thus were discarded) and that passed validation and were used for constructing context windows. These examples are selected nearly-randomly; many of the instances that fail validation contain sexual content or offensive language, and we screen these out of the examples shown wherever possible. While the instances screened out vary by dataset, we note that the datasets that use labels scraped from web content (i.e. Yahoo Topics, AGNews, IMDB reviews, App Reviews) generally have higher rates of validation failures. Some of these failures appear to be genuine mislabelings, which is likely because of noise in the user behavior (e.g. a user asks a question about Business & Finance but mistakenly posts it in the Sports topic, or a user writes a review with mostly negative text but gives the product 4 stars).

Input	Label	Passed validation?
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	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	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
‘Virtual Girlfriend’ Demands Gifts HONG KONG - She needs to be coddled with sweet talk and pampered with gifts, but you’ll never see her in the flesh. A Hong Kong company has developed a “virtual girlfriend” for new cell phones with video capability...	World	✗
Court Hears Case of Brain Damaged Woman (AP) AP - The Florida Supreme Court questioned lawyers Tuesday about the extent of the power handed to Gov. Jeb Bush under a law that let him order the reinsertion of a brain-damaged woman’s feeding tube.	Sci/Tech	✗
Rooney backs fiercesome threesome (AFP) AFP - Wayne Rooney believes the three-pronged attack of himself, Michael Owen and Jermain Defoe can put England on the fast track to the World Cup finals.	World	✗
From mouths of babes: What’s hot, what’s not How hot are the “Hot Dozen” toys? To find out, The Boston Globe put six toys from the 2004 Toy Wishes magazine Hot Dozen list in front of 13 kids from the Charlestown Boys and Girls Club.	Business	✗
Crematory operator to get 12 years Ray Brent Marsh, who is to enter the plea Friday, had faced up to 8,000 years in a case that shocked the nation two years ago when investigators found hundreds of corpses at his rural northwest facility.	Business	✗
Vending Machines Making Room for Healthy Products WASHINGTON (AP) – The typical vending machine fare consists of chocolate bars and potato chips, leaving few options for people seeking low-calorie or low-salt snacks. That is changing now as companies develop markets for products they expect to satisfy both nutritionists and consumers...	Sci/Tech	✓
Venezuelans Line Up to Vote on Chavez CARACAS, Venezuela - Summoned by bugle calls and the detonations of huge firecrackers, Venezuelans turned out in unprecedented numbers Sunday to vote on whether to force leftist President Hugo Chavez from office. Some lines at polling places extended for 1.25 miles...	World	✓
Singh Snare PGA Title Vijay Singh outlasts Justin Leonard and Chris DiMarco in a three-way playoff to win the PGA Championship on Sunday at Whistling Straits in Haven, Wisconsin.	Sports	✓
China’s Panchen Lama visits Tibet The boy named by the Chinese authorities as the 11th Panchen Lama visits a temple in Tibet.	World	✓
The Region’s Highest-Paid Executives Pay for the Washington area’s top executives rose significantly last year, reversing the downward trend that set in with the recession in 2001.	Business	✓

Table 6: AGNews examples that failed and passed validation.

Input	Label	Passed validation?
Kill Me..	positive	✗
Used to be great app. But since the last 7 updates always get - Authentication failed for: x@gmail.com	positive	✗
App force closing on changing THEME PLZ FIX!	positive	✗
You can't uninstall & I don't like my space being used. But am glad for the blind & deaf	positive	✗
GOOGLE TALKBACK Read carefully & then give at least a good reviews this app is made only for persons who got disability such as a blind person & if you want to get rid this simply just ROOT your device dumb head!!	negative	✗
This thing is great. I am at the bottom of the learning curve but in the few minutes I have played with it it seems like it will be easy to learn and to use. Just what I hoped for.	positive	✓
Crash when set a demodulation! :(	negative	✓
doesn't work with rtl-2832 Says it works with rtl devices but doesn't. never displays anything even though the rtl 2832 works with other spectrum analyzers. no help files anywhere so good luck figuring out what's wrong!	negative	✓
Great app! Works really well with Galaxy note 5	positive	✓
Works great I am using HackRF One with this apps.	positive	✓

Table 7: App Reviews examples that failed and passed validation.

## B QUESTION TYPES

## B.1 OOLONG-SYNTH

Table 16, Table 17, and Table 18 list the questions used for each type of task in OOLONG-synth.

## B.2 OOLONG-REAL

Table 19 and Table 20 lists the question types used for single and multi-episode settings in OOLONG-real dataset.

Input	Label	Passed validation?
The Unit secret agent has signed on to appear in three episodes as a businessman being shown some houses - and, presumably, one bedroom in particular - by our randy realtor heroine.	formal	✗
Barely 12, with large brown eyes and stick-like arms, Fandi is 3 years older than his brother – in his eyes almost a man.	formal	✗
D-Lister Avril Lavigne appears on the cover of Z-List magazine Savvy this month.	formal	✗
ORHS is going through it now, but from all I can perceive the district is out for the quick fix, chop off the principal, instead of really trying to assist the students and dig to the bottom to find the truth.	formal	✗
Little Sally Draper (Kiernan Shipka) is a "Patty Duke"-era girl - indeed, "Mad Men" is currently set during the month "The Patty Duke Show" premiered in 1963 - and on a "Patty Duke"-like show, she'd just be a lispy Shirley Temple doll with a crush on Daddy.	formal	✗
I have tried everything possible to attract business.	formal	✓
Never heard that one before.	informal	✓
"Yup."	informal	✓
That's how we pass our traits to the next generation: through DNA	formal	✓
Wedding rituals differ in different regions and communities in India.	formal	✓

Table 8: Formality examples that failed and passed validation.

972 973 974	Input	Label	Passed validation?
975 976 977 978 979 980	Good movie, very 70s, you can not expect much from a film like this. Sirpa Lane is an actress of erotic films, a nice body but nothing exceptional. Not demand a lot from these films are light years away from the movies today, the world has changed incredibly. The plot is simple and the actors not extraordinary.	positive	✗
981 982 983 984 985	This film Evil Breed: The legend of samhain contains very little thought or effort. It is ridiculed with specs of ultra fast "slasher" style death and plain disgusting acts of death. The acting was rated a D as the actors show very little ability, and the stupidity of them in the film is too questionable.	positive	✗
986 987 988 989 990 991	The movie "Holly" may make the audience want to donate money towards organizations that improve the life for these poor youngsters, but the film's dramatic weaknesses may reduce its chances of being seen by enough people to make a difference. Overall, I think the concept is better as a documentary and it was not as touching as a movie.	positive	✗
992 993 994 995 996	Although this film put Davis on the map due to her brilliantly intense performance, this film is strangely unsatisfying to me as a whole. What I cannot fathom for the life of me is just how or why Phillip would take the constant abuse this tramp constantly dishes out towards him.	positive	✗
997 998 999 1000 1001	This film is just plain horrible. John Ritter doing pratt falls, 75% of the actors delivering their lines as if they were reading them from cue cards, poor editing, horrible sound mixing, and a plot that really goes nowhere. If I could sum this film up in one word, that word would be: Suckotrocity	negative	✓
1002 1003 1004 1005 1006 1007 1008 1009 1010 1011 1012	Zentropa has much in common with The Third Man, another noir-like film set among the rubble of postwar Europe. Like TTM, there is much inventive camera work. There is an innocent American who gets emotionally involved with a woman he doesn't really understand, and whose naivety is all the more striking in contrast with the natives. But I'd have to say that The Third Man has a more well-crafted storyline. Zentropa is a bit disjointed in this respect. Perhaps this is intentional: it is presented as a dream/nightmare, and making it too coherent would spoil the effect. This movie is unrelentingly grim—"noir" in more than one sense; one never sees the sun shine. Grim, but intriguing, and frightening.	positive	✓
1013 1014 1015 1016 1017	Never in my life have I come across a movie as bad as The Zombie Chronicles. Filmed on a budget of what looks to be about 20 bucks, TZC is a completely horrible horror movie that relies on lame, forgettable actors. Simply put, avoid TZC like a sexually-transmitted disease.	negative	✓
1018 1019 1020 1021 1022 1023	Without wishing to be a killjoy, Brad Sykes is responsible for at least two of the most dull and clichéd films I've ever seen. The acting is terrible, the print is shoddy, and everything about this film screams "seriously, you could do better yourself".	negative	✓

Table 9: IMDB examples that failed and passed validation. Only eight examples are shown because of the example length.

1026

1027

1028

1029

1030

1031

1032

1033

1034

1035

1036

1037

1038

1039

1040

1041

1042

1043

1044

1045

1046

1047

Input	Label	Passed validation?
Draupadi's eyes were diamonds. $\leftrightarrow$ Draupadi's eyes were beautiful.	correct	✓
The faculty meeting was an easy breeze $\leftrightarrow$ The faculty meeting was very easy and relaxing	correct	✓
The house was a tomb. $\leftrightarrow$ The house was big.	incorrect	✓
I had already planted the idea in her mind. $\leftrightarrow$ I had already scared her about the idea.	incorrect	✓
It is sad to observe the fruits of ignorance. $\leftrightarrow$ It is sad to observe the effects of ignorance.	correct	✓

Table 10: Metaphors examples that passed validation. No examples failed validation for this dataset.

1048

1049

1050

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

Input	Label	Passed validation?
You're safe. $\rightarrow$ You have nothing to worry about.	neutral	✗
asks Burton, cupping his ear. $\rightarrow$ Burton wanted to ask what was going on, but he couldn't do much of anything while bound and gagged.	contradiction	✗
I "poison my dearest Emily!" $\rightarrow$ How could you think I would poison Emily?	entailment	✗
When he finally succeeded, after a prolonged siege and heavy losses, he punished the local population by cutting off the noses and lips of all men except those who played wind instruments. $\rightarrow$ All men who played wind instruments were tasked with helping cut off people's noses.	neutral	✗
They died slowly, their eyes bulging and faces turning blue. $\rightarrow$ They died asphyxiated.	neutral	✗
well you see that on television also $\rightarrow$ You can see that on television, as well.	entailment	✓
Vrenna and I both fought him and he nearly took us. $\rightarrow$ Neither Vrenna nor myself have ever fought him.	contradiction	✓
This analysis pooled estimates from these two studies to develop a C-R function linking PM to chronic bronchitis. $\rightarrow$ The analysis proves that there is no link between PM and bronchitis.	contradiction	✓
He turned and smiled at Vrenna. $\rightarrow$ He smiled at Vrenna who was walking slowly behind him with her mother.	neutral	✓
We sought to identify practices that were commonly implemented by the agencies within the past 5 years. $\rightarrow$ We want to identify practices commonly used by agencies in the last 5 years	entailment	✓

Table 11: MultiNLI examples that failed and passed validation.

1080

1081

1082

1083

Input	Label	Passed validation?
No care is anything that serves as an enticement.	True	✗
No care may be anything that serves as an enticement.	True	✗
A loud utterance is an appropriate definition of produce in no context.	True	✗
Complete attention is an appropriate definition of candidate in no context.	True	✗
No stress is a message received and understood.	True	✗
Action refers to a military engagement.	True	✓
Action never stands for a distinguishing quality.	True	✓
Action does not stand for a military engagement.	False	✓
Action may stand for a military engagement.	True	✓
Not a single action is a military engagement.	False	✓

1099

1100

Table 12: Negation examples that failed and passed validation.

1101

1102

1103

1104

1105

1106

1107

Input	Label	Passed validation?
Go chase after her and run her over while she's crossing the street	ham	✗
i want to grasp your pretty booty :)	ham	✗
No da if you run that it activate the full version da.	ham	✗
i am seeking a lady in the street and a freak in the sheets. Is that you?	ham	✗
Should i send you naughty pix? :)	ham	✗
WINNER!! As a valued network customer you have been selected to receive a £900 prize reward! To claim call 09061701461. Claim code KL341. Valid 12 hours only.	spam	✓
Had your mobile 11 months or more? U R entitled to Update to the latest colour mobiles with camera for Free! Call The Mobile Update Co FREE on 08002986030	spam	✓
I'm gonna be home soon and i don't want to talk about this stuff anymore tonight, k? I've cried enough today.	ham	✓
SIX chances to win CASH! From 100 to 20,000 pounds txt CSH11 and send to 87575. Cost 150p/day, 6days, 16+ TsandCs apply Reply HL 4 info	spam	✓
I HAVE A DATE ON SUNDAY WITH WILL!!	ham	✓

1129

1130

1131

1132

1133

Table 13: Spam examples that failed and passed validation. The examples that failed validation were almost exclusively ham messages with sexual content; the selected messages are the least inappropriate of the examples that failed validation.

Input	Label	Passed validation?
What explosive do you get by mixing charcoal, sulfur and saltpeter?	entity	✗
What is the procedure called for drilling a hole in your skull to achieve a higher consciousness?	entity	✗
Name 11 famous martyrs.	human being	✓
What's the Olympic motto?	description and abstract concept	✓
What is the highest waterfall in the United States?	location	✓
What does the abbreviation AIDS stand for?	abbreviation	✓
How many points make up a perfect fivepin bowling score?	numeric value	✓

Table 14: TREC-coarse-Q examples that failed and passed validation. Only two examples failed validation.

Input	Label	Passed validation?
Is there a God? The question to end all questions, and begin them.	Business & Finance	✗
why do we need to lie?	Health	✗
try to type the word supercalifragilisticexpialidocious. 20 times fast you can only make 10 mistakes? This is for fun!	Education & Reference	✗
do you think it is okay to tell a lie?	Education & Reference	✗
where can you purchase cesium carbonate?	Science & Mathematics	✗
What are good sources to find out about new gospel artists? Is there a site that focuses primarily on gospel?	Entertainment & Music	✓
How a black hole is formed? I would like to know how a black hole can possibly be formed. Are there any experimental evidence of such creation?	Science & Mathematics	✓
Economics of running a restaurant? Running a restaurant looks like hard work and long hours. What percentage of restaurants are profitable?	Business & Finance	✓
Why doesn't the NBA implement a minor leagues? I don't want to see any more High School kids on the court.	Sports	✓
how does a helicopter fly this is a miracle- I always wanted to learn to fly one of those. Can someone explain how can it get airborne?	Science & Mathematics	✓

Table 15: Yahoo examples that failed and passed validation. One class of examples that failed validation, not demonstrated here, are sexually explicit questions labeled as seemingly random categories (likely from spam posters on the original Yahoo Answers forum).

1188

1189

1190

1191

1192

1193

---

Counting

---

1194

1195

In the above data, which of the labels is the most common? Give your final answer in the form ‘label: answer’ where answer is one of the labels: {label\_list}.

1196

1197

In the above data, which of the labels is the least common? Give your final answer in the form ‘label: answer’ where answer is one of the labels: {label\_list}.

1198

1199

In the above data, is label ‘{A}’ more common, less common, or the same frequency as label ‘{B}’? Give your final answer in the form ‘Answer: {A} is [X]{B}’, where [X] is ‘more common than’, ‘less common than’, or ‘same frequency as’.

1200

1201

In the above data, how many data points should be classified as label ‘{label}’? Give your final answer in the form ‘Answer: number’.

---

1202

1203

Table 16: Counting questions used in the OOLONG-synth dataset.

1204

1205

1206

1207

1208

1209

1210

1211

1212

1213

1214

---

User

---

1215

1216

In the above data, which user is represented most often? Give your final answer in the form ‘User: [X]’, where [X] is the user ID.

1217

1218

In the above data, which user is represented the second most often? Give your final answer in the form ‘User: [X]’, where [X] is the user ID.

1219

1220

For the following question, only consider the subset of users with IDs {user\_names}. Among these users, which user is represented most often? Give your final answer in the form ‘User: [X]’, where [X] is the user ID.

1221

1222

For the following question, only consider the subset of users with IDs {user\_names}. Among these users, which user is represented the second most often? Give your final answer in the form ‘User: [X]’, where [X] is the user ID.

1223

1224

For the following question, only consider the subset of instances that are associated with user IDs {user\_names}. Among instances associated with these users, {any of the Counting questions above}

1225

1226

For the following question, only consider the subset of users with IDs {user\_names}. Among these users, which user has the most instances with the label {label}? Give your final answer in the form ‘User: [X]’, where [X] is the user ID.

1227

1228

In the above data, which user has the most instances with the label {label}? Give your final answer in the form ‘User: [X]’, where [X] is the user ID.

1229

1230

In the above data, which user has more instances with the label {label}: User {A} or User {B}? Give your final answer in the form ‘User: [X]’, where [X] is the user ID.

---

1231

1232

1233

1234

1235

1236

1237

1238

1239

1240

1241

Table 17: User questions used in the OOLONG-synth dataset.

1242  
 1243  
 1244  
 1245  
 1246  
 1247  
 1248  
 1249  
 1250  
 1251  
 1252  
 1253

---

### Timeline

---

1254 In the above data, which date is represented most often? Give your final answer in the form  
 1255 ‘Date: [X]’, where [X]is the date in the format MM/DD/YYYY.  
 1256 In the above data, which date is represented most often? Give your final answer in the form  
 1257 ‘Date: [X]’, where [X]is the date in the format MM/DD/YYYY.  
 1258 In the above data, which date is represented the second most often? Give your final answer in  
 1259 the form ‘Date: [X]’, where [X]is the date in the format MM/DD/YYYY.  
 1260 In the above data, how many dates are represented exactly {n} times? Give your final answer  
 1261 in the form ‘Answer: [X]’, where [X]is the number of dates represented exactly {n} times.  
 1262 In the above data, was label ‘{key}’ more common, less common, or the same frequency before  
 1263 {time}, as compared to after {time}? Give your final answer in the form ‘Answer: {key} is  
 1264 [X]before {time}’, where [X]is ‘more common’, ‘less common’, or ‘the same frequency’.  
 1265 In the above data, was label ‘{key}’ more common, less common, or the same frequency before  
 1266 {time}, as compared to after {time}? Give your final answer in the form ‘Answer: {key} is  
 1267 [X]before {time}’, where [X]is ‘more common’, ‘less common’, or ‘the same frequency’.  
 1268 For the following question, only consider the subset of instances that occur in {month\_name} of  
 1269 any year. Among instances occurring in {month\_name},{any of the Counting questions above}  
 1270 For the following question, only consider the subset of instances that occur between  
 1271 {starting\_date} and {ending\_date}, inclusive. Among instances occurring in this date range,  
 1272 {any of the Counting questions above}  
 1273 In which month did the label ‘{label1}’ first occur more often than the label ‘{label2}’? Give  
 1274 your final answer in the form ‘Answer: [month][year]’, where [month]is the name of the month  
 1275 and [year]is the four-digit year where ‘{label1}’ first occurred more often than ‘{label2}’.  
 1276 For how many months does the label ‘{label1}’ occur more frequently than the label  
 1277 ‘{label2}’? Disregard months where there is a tie. Give your final answer in the form ‘Answer:  
 1278 [X]’, where [X]is the number of months where ‘{label1}’ occurs more often than ‘{label2}’.  
 1279 For how many months is the label ‘{label}’ the single most frequently occurring label? Disre-  
 1280 gard months where there is a tie for the most common label. Give your final answer in the form  
 1281 ‘Answer: [X]’, where [X]is the number of months where ‘{label}’ is the most common label.

---

1285  
 1286  
 1287  
 1288  
 1289  
 1290  
 1291  
 1292  
 1293  
 1294  
 1295

Table 18: Timeline questions used in the OOLONG-synth dataset.

1296	Rolls
1297	Total number of rolls in this episode? (counting)
1298	Total number of rolls by the character {character name} in this episode? (counting, character)
1299	Total number of rolls by the player {player name} in this episode? (counting, player)
1300	Total number of rolls of type {roll type} in this episode? (counting, roll type)
1301	Number of rolls of natural value {roll value} in this episode? (counting, roll value)
1302	In this episode, what percentage of rolls were of value {roll value}? round to the nearest integer. (counting, roll value)
1303	What is the most common roll type in this episode? Return a comma separated list. (frequency)
1304	What is the least common roll type in this episode? Only include types with more than one roll. Return a comma separated list. (frequency)
1305	What is the most common natural roll value in this episode? Return a comma separated list. (frequency)
1306	What is the least common natural roll value in this episode? Only include values with more than one roll. Return a comma separated list. (frequency)
1307	What is the count of Crits? (natural rolls of value 1 or 20)? (counting, roll value)
1308	What is the count of Nat20s (natural rolls of value 20)? (counting, roll value)
1309	What is the count of Nat1s (natural rolls of value 1)? (counting, roll value)
1310	Spells
1311	How many spells were cast during this episode? (counting)
1312	How many spells were cast by the character {character name} in this episode? (counting, character)
1313	How many spells were cast by the player {player name} in this episode? (counting, player)
1314	How many {spell type} spells were cast during this episode? (counting, spell type)
1315	What is the first spell cast in this episode? (enumeration)
1316	What are the first two spells cast in this episode? (enumeration)
1317	What are the first three spells cast in this episode? (enumeration)
1318	What is the last spell cast in this episode? (enumeration)
1319	What are the last two spells cast in this episode? (enumeration)
1320	What are the last three spells cast in this episode? (enumeration)
1321	What is the first spell cast by each character in this episode? Return a comma separated list and retain the order of spells as they appear in the episode. (enumeration)
1322	What is the last spell cast by each character in this episode? Return a comma separated list and retain the order of spells as they appear in the episode. (enumeration)
1323	How many characters cast {spell name} spell in this episode? (counting, spell type)
1324	What is the most common spell in this episode? Return a comma separated list. (frequency)
1325	What is the least common spell in this episode? Only include spells that were cast at least once. Return a comma separated list. (frequency)
1326	Which spells were cast by more than one character in this episode? Return a comma separated list. (counting, enumeration)
1327	What is the total number of cantrip spells cast in this episode? (counting, spell type)
1328	In this episode, how many times was a spell cast at a level higher than its base level? (counting)
1329	In this episode, which spells were cast at a level higher than their base level? Return a comma separated list of unique spells. (enumeration)
1330	Table 19: Question types used in the OOLONG-real dataset (single episode). <a href="#">For reference, we provide the question types in parentheses.</a>
1331	
1332	
1333	
1334	
1335	
1336	
1337	
1338	
1339	
1340	
1341	
1342	
1343	<b>C DEEPSEEK R1 TRACE EXAMPLES</b>
1344	
1345	Figure 5 and Figure 6 contrast the behavior of this model on representative examples from the OOLONG-synth and OOLONG-real splits. See the main text for more discussion.
1346	
1347	
1348	
1349	

1350	
1351	<b>Rolls</b>
1352	What is the cummulative total of rolls by the end of episode {episode index}? Count the
1353	number of rolls and not the values of the rolls. (counting)
1354	What is the cummulative total of rolls by the character {character name} at the end of episode
1355	{episode index}? Count the number of rolls and not the values of the rolls. (counting, character)
1356	Total number of rolls across all the episodes? (counting)
1357	Total number of rolls by the character {character name} across all episodes? (counting, char- acter)
1358	Total number of rolls by the player {player name} across all episodes? (counting, player)
1359	Total number of rolls of type {roll type} across all episodes? (counting, roll type)
1360	Number of rolls of natural value {roll value} across all episodes? (counting, roll value)
1361	Across all episodes, what percentage of rolls were of value {roll value}? round to the nearest integer. (counting, roll value)
1362	What is the most common roll type across all episodes? Return a comma separated list. (fre- quency)
1363	What is the least common roll type across all episodes? Only include types with more than one roll. Return a comma separated list. (frequency)
1364	What is the most common natural roll value across all episodes? Return a comma separated list. (frequency)
1365	What is the least common natural roll value across all episodes? Only include values with more than one roll. Return a comma separated list. (frequency)
1366	What is the total count of Crits across all episodes? (natural rolls of value 1 or 20)? (counting, roll value)
1367	What is the total count of Nat20s across all episodes? (natural rolls of value 20)? (counting, roll value)
1368	What is the total count of Nat1s across all episodes? (natural rolls of value 1)? (counting, roll value)
1369	
1370	
1371	
1372	
1373	
1374	
1375	
1376	<b>Spells</b>
1377	What is the cummulative total of spells cast by the end of episode {episode index}? (counting, indexing)
1378	What is the first spell cast in the episode {episode index}? (enumeration, indexing)
1379	What is the second spell cast in the episode {episode index}? (enumeration, indexing)
1380	What is the third spell cast in the episode {episode index}? (enumeration, indexing)
1381	List the first spell cast in each episode? Return a comma separated list. (enumeration)
1382	List the last spell cast in each episode? Return a comma separated list. (enumeration)
1383	List the first spell cast by the character {character name} in each episode? Return a comma separated list. (enumeration, character)
1384	List the last spell cast by the character {character name} in each episode? Return a comma separated list. (enumeration, character)
1385	How many spells were cast across all episodes? (counting)
1386	How many spells were cast by the character {character name} across all episodes? (counting, character)
1387	How many spells were cast by the player {player name} across all episodes? (counting, player)
1388	How many {spell type} spells were cast across all episodes? (counting, spell type)
1389	How many characters cast {spell name} spell across all episodes? (counting, spell type)
1390	What is the most common spell across all episodes? Return a comma separated list. (frequency, enumeration)
1391	What is the least common spell across all episodes? Only include spells that were cast at least once. Return a comma separated list. (frequency, enumeration)
1392	What is the total number of cantrip spells cast across all episodes? (counting, spell type)
1393	Across all episodes, how many times was a spell cast at a level higher than its base level? (counting, spell level)
1394	Across all episodes, which spells were cast at a level higher than their base level? Return a comma separated list of unique spells. (enumeration, spell level)
1395	
1396	
1397	
1398	
1399	
1400	
1401	
1402	
1403	

Table 20: Question types used in the OOLONG-real dataset (multi-episode). For reference, we provide the question types in parentheses.

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

<think> We are given 21 pairs. We need to classify each pair as one of ‘neutral’, ‘contradiction’, or ‘entailment’. Then, we will count the frequency of ‘contradiction’ and ‘entailment’ and compare them. Let’s go through each pair one by one: 1. “and then the class uh double A Dallas Fort Worth Spurs we later uh they went to the Texas league - Dallas Fort Worth Spurs got into the Texas League.” - The first sentence says they went to the Texas league. The second says they got into the Texas League. This seems to be the same meaning. So, entailment. 2. “Specifically, although most public companies are required to have a board that is comprised of a majority of independent directors and certain key committees are required to be comprised solely of outside directors, a closer look reveals that the independence of many boards may be more a matter of form than substance. - Public companies can put whoever they want on their boards.” - The first sentence says that public companies are required to have independent directors, but the independence might be in form only. The second sentence says they can put whoever they want. This contradicts the requirement mentioned. So, contradiction. **[....authors truncated classifications for space]** 19. “And the bureaucratic approach prevents the section as a whole from taking positions on—or creating—literary issues. - Bureaucracy makes groups not take positions or a stand on literary issues.” - The second is a restatement of the first. Entailment. 20. “They are linked not only by geographical location but also by here are the upmarket and fashionable neighborhoods, with affluent residential sections, trendy restaurants, and L.A.’s fabulous shopping areas. - They are linked because they both have very exclusive shoe stores.” - The first says they are linked by upmarket and fashionable neighborhoods, etc. The second says they are linked because of exclusive shoe stores. The first does not mention shoe stores. So, the second is a specific reason that is not stated. This could be neutral. However, note that the first lists shopping areas, which might include shoe stores, but it doesn’t specify. The second is too

Figure 5: Example reasoning trace for OOLONG-synth. The model fails to complete the task because it runs out of tokens during the reasoning trace. The majority of the trace (classifying items 3-18) is omitted for space.

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

Based on the provided transcript of Episode 2, the second spell cast is Disintegrate, cast by Lyra (Felicia Day's character).

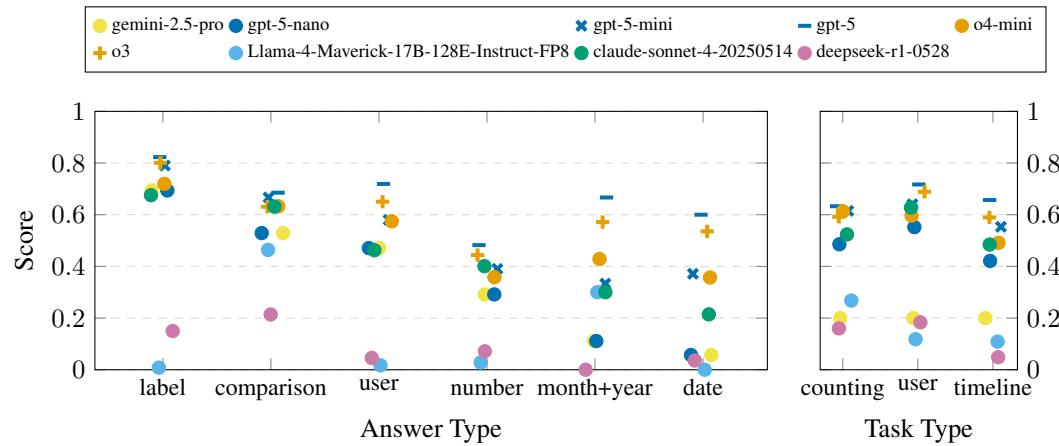
Sequence of Events: 1. First Spell: Vex'ahlia (Laura Bailey) casts Hunter's Mark on a frost giant (during her turn in combat). 2. Second Spell: Lyra (Felicia Day) casts Disintegrate on the same frost giant (during her subsequent turn).

Key Excerpt: Felicia: "I have a spell called Disintegrate. Can I cast that on its wing? [...] I'm doing Disintegrate. And it's 10d6 plus a d10, plus 40." Matt: "Disintegrate is nasty. [...] As you release the torrent of electrical energy, it courses through the body [...] the giant looks severely messed up."

Thus, the second spell in Episode 2 is Disintegrate.

Disintegrate

Figure 6: Example reasoning trace for OOLONG-real. The model completes the task within a relatively short reasoning trace, despite this task using substantially longer input.

1512 **D QUESTION TYPE ANALYSIS**  
15131514 **D.1 OOLONG-SYNTH**  
15151516 In Figure 7, we report the model performance by question and task types on OOLONG-synth.  
15171533 Figure 7: The performance trend for models by type of answer and type of task on OOLONG-synth.  
15341535 **D.2 OOLONG-REAL**  
15361537 In Figure 8, we report model performance by question type on OOLONG-real.  
15381539  
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

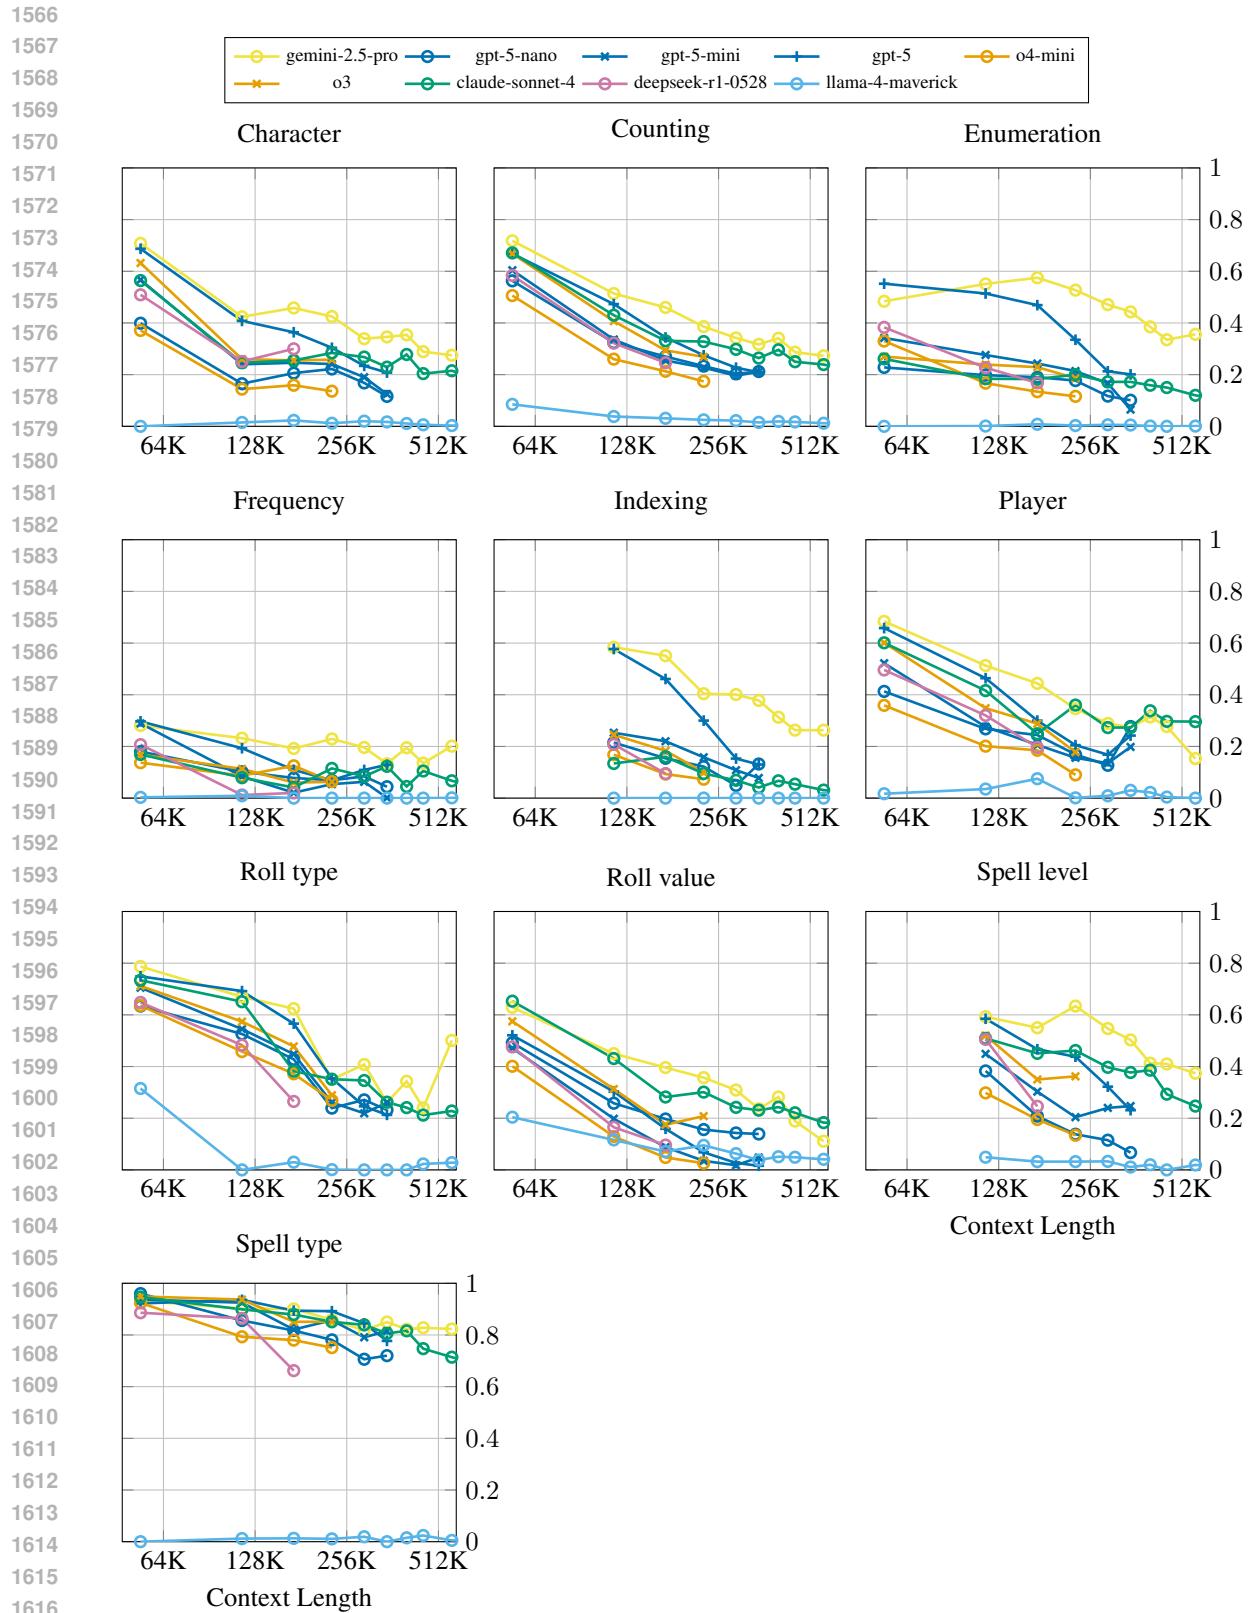


Figure 8: Scores by context window length for various question types in OOLONG-real. Some questions are only relevant for multi-episode settings.

1620 E OOLONG-REAL DATA

1621

1622

1623

1624

1625

1626

1627

1628

1629

## E.1 DATA QUALITY AND CURATION

1630

1631

1632

1633

1634

1635

1636

The episode transcripts were first hand-transcribed by members of the fan community, who implemented a multi-step editing process with editorial guidelines. These were then further cleaned by Rameshkumar & Bailey (2020) to resolve any remaining inconsistencies. We refer the reader to the original paper for more details on the curation process. The stats were manually compiled by volunteers and the team at CritRoleStats. While it's not possible to completely eliminate sources of potential error, these types of fan-documentation processes are fastidiously annotated and cross-checked by multiple people— a standard of care far higher than that of a paid annotator with no connection to the task or data domain.

1637

1638

1639

1640

1641

1642

1643

## E.2 CHALLENGES OF OOLONG-REAL

1644

1645

1646

1647

In the tables below, we highlight various unique challenges presented by OOLONG-real. Figure 9 covers text that includes in-character, out-of-character and narrative content. Additionally, we show complex discussions surrounding dice rolls (Figure 10, Figure 12) and damage calculation (Figure 11).

1648

1649

1650

1651

1652

1653

1654

1655

1656

1657

**Matt:** All right. As your smaller friend curls awake, Sam, would you like to describe your character?

**Sam:** Yeah. Um. I am a little goblin girl. (laughter)

**Sam:** I am a goblin. So, you know, the green skin, the green hair, the yellow eyes. And she wears not-great clothes, just like her traveling companion there. She hides in the shadows a lot, because she knows goblins aren't welcome in this part, and that's about it. I mean, she's a little skittish, and right now she's probably stirring awake as well, right?

**Matt:** Well. What's your name?

**Sam:** Oh. Nott the Brave.

**Matt:** So Caleb, as you come to consciousness, you glance over and can see, slowly snoring and rousing at about the same time, Nott's eyes blink open, her slowly groaning face looking over towards you.

**Sam:** (high-pitched Cockney accent) Oh! You're finally awake, I see. Oh yeah, motherfuckers. It's on. You were out for quite some time, there. Rough day, eh?

**Liam:** (light German accent) Not our best day, no.

**Sam:** No, I mean, usually you're so good at everything, but yesterday you were just— maybe you needed the sleep, is what you needed.

1672

1673

Figure 9: Example of mixed in-character, out-of-character, and narrative text.

1674  
 1675 **Matt:** The old man's body stops quaking, his flesh now grey and mangled like an ancient tree  
 1676 trunk. He turns, his eyes blood-red and bulging, his lips curled into a horrifying grimace.  
 1677 The two Crown's Guard begin to try and make their way through the panicked crowd, but  
 1678 the people, like a wave of chaos, are keeping them at bay. I need everyone to roll initiative.  
 1679 (yelling)  
 1680 **Liam:** The miniatures come out, guys!  
 1681 **Travis:** First map! (yelling)  
 1682 **Marisha:** I rolled a natural one.  
 1683 **Sam:** That bodes well.  
 1684 **Matt:** So you guys, this guy is actually over here. You guys are all up here in the front.  
 1685 **Laura:** Oh yeah, we were right in the fucking front row. Look at us! We're so cute!  
 1686 **Taliesin:** Oh my god, it's so pretty.  
 1687 **Sam:** We're all going to die.  
 1688 **Liam:** Probably. But then the third campaign begins.  
 1689 **Laura:** Yay.  
 1690 **Sam:** I like the tents and donkeys outside. Amazing.  
 1691 **Marisha:** Look, Taliesin, it's a bardo. Oh my god!  
 1692 **Taliesin:** No, go for it. Oh boy.  
 1693 **Matt:** All right, so. 20 to 15?  
 1694 **Laura:** 21.  
 1695 **Matt:** All right.  
 1696 **Laura:** Oh, wait, 22.  
 1697 **Matt:** Nice. 15 to ten?  
 1698 **Taliesin:** Ten.  
 1699 **Sam:** Ten.  
 1700 **Travis:** 12.  
 1701 **Ashley:** 13.  
 1702 **Liam:** 11.  
 1703 **Matt:** So 13 and then ten and ten? Sorry, 12. So Yasha got 13, then we have Fjord. And  
 1704 then 11. And then we have Nott at ten. And what'd you get?  
 1705 **Marisha:** Five. Rolled terribly.

Figure 10: A complex discussion of dice rolls. While many numbers are mentioned, there are actually seven rolls discussed in this segment: one initiative value for each player.

1708  
 1709  
 1710  
 1711  
 1712  
 1713  
 1714  
 1715 **Matt:** 16 hits. Go ahead and roll damage.  
 1716 **Laura:** Yay. 4d6.  
 1717 **Matt:** Yasha, you're almost on deck.  
 1718 **Laura:** Ten, 16, 18, and then any attacks in the future– the next attack has advantage.  
 1719 **Matt:** How much damage was that? 2d6, 18?  
 1720 **Sam:** 4d6, you said.  
 1721 **Laura:** Yeah, it was 4d6 radiant damage. So yeah, it was 18. And then I'm going to run the  
 1722 opposite direction away from him, but I'm going to keep my duplicate up there.  
 1723 **Matt:** Okay, which, for your duplicate, I should go ahead and grab something.

Figure 11: A complex discussion of a damage roll calculation. Many numbers are floated in this discussion; the model must recognize that the discussion refers to a single damage event and the final, resolved damage was 18.

1728

**Sam:** I thought you just said I had advantage!**Matt:** You had advantage, but you have disadvantage because they're beyond the-**Liam:** It cancelled out your advantage. It's just a straight roll and you lost the sneak attack.**Sam:** Could I run forward and still get the sneak attack?**Matt:** You could try it.**Sam:** That's what I will do. Undo.**Laura:** Roll again, because you might get natural 20!**Sam:** Now I have to roll for stealth, right?**Matt:** You've been stealthed as you crept up on this from the last time, so go for it.**Sam:** I'm rolling to hit?**Matt:** You already hit on that.**Laura:** Roll again for advantage, you might get a 20.**Sam:** Okay. Still 17 plus six.**Matt:** Now you get a sneak attack, go for it.

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

Figure 12: This discussion highlights an *undo* of a dice roll. To tackle, the models needs to understand the narrative around the original roll and correct its count because of the undo.