

000 001 002 003 004 005 CHILDEVAL: HOW LARGE LANGUAGE MODELS MEET 006 CHILDREN'S PERSONALITIES 007 008 009

010 **Anonymous authors**
011 Paper under double-blind review
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028

ABSTRACT

029
030
031 The remarkable success of Large Language Models (LLMs) has revolutionized
032 LLM-based chatbots for personalized tasks beyond generic dialogues. Personal-
033 alization involves customizing LLMs to generate text responses based on user
034 preferences. One promising endeavor is to enable personalized interactions for
035 children's caretakers while also promoting development and learning. However,
036 dedicated research is required to determine whether LLMs can effectively deliver
037 personalized responses based on children's preferences, given their interactions
038 differ from adults and there are no child-oriented preference assessment method-
039 ologies. Facing these challenges, we introduce ChildEval, a benchmark to eval-
040 uate LLMs' capacity to infer, interpret, and follow child-centered preferences in
041 a long-context conversational setting. Our benchmark comprises of 29K synthe-
042 sized children's (ages 3-6) persona profiles, which are related to their preferences
043 in both explicit and implicit manners. Implicit preferences are integrated inside
044 dialogues consisting of 6 to 10 turns. The preferences cover 5 top-level and 14
045 sub-level topics that involve children's daily lives and development. We further
046 propose fine-grained child-centric evaluation protocols to systematically evaluate
047 the performance of open-source LLMs. Experimental results demonstrate the im-
048 pact of various personalized representations on LLM responses and indicate that
049 fine-tuning on this dataset may enhance performance.¹
050
051

1 INTRODUCTION

052 Large Language Models (LLMs) (e.g., ChatGPT (OpenAI et al., 2024), Gemini (Gemini et al.,
053 2025), and Claude (Anthropic, 2024)) have achieved remarkable success in effectively under-
054 standing and generating human language, leading to a revolutionary era in LLMs. Beyond generic dia-
055 logues, LLMs have been utilized in a wide range of individual daily tasks (e.g., healthcare (Xu et al.,
056 2024) and finance (Easin et al., 2024)) to deliver personalized user experiences based on preferences
057 (Kumar et al., 2024). One promising endeavor is to enable personalized interactions for children's
058 caretakers while also promoting development and learning (Feng et al., 2024; Seo et al., 2024; Chen
059 et al., 2025a), instead of just giving the “correct” answers.

060 Previous research into the potential of LLMs for personalized interactions has been focused on adult
061 preferences and tasks. Qiu et al. (2025) advance LLMs toward effective personalization by intro-
062 ducing new methods for extracting user preferences from historical profiles. Other studies (Salemi
063 et al., 2023; Jiang et al., 2025) attempt to address the challenge of a lack of adequate benchmarks
064 for comprehensively evaluating the LLM's personalized capabilities. However, the proposed bench-
065 marks focus on general preferences (e.g., the number of dialogue turns) with generic tasks (e.g.,
066 ticket booking and restaurant recommendations) for adults. There are several benchmarks (Rath
067 et al., 2025; Liu & Fournassi, 2024) proposed for children. Rath et al. (2025) examine child safety
068 without considering a diversity of child-centered tasks. Liu & Fournassi (2024) explore how LLM
069 generate replies that mimic the child's style without considering the child's developmental and learn-
070 ing requirements. Dedicated research is required to determine whether LLMs can effectively deliver
071 personalized responses based on children's preferences, as their interactions differ from those of
072 adults.

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
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1

In particular, we identify the following primary gaps in current research on child personalization benchmarks in LLMs. 1) First, existing benchmarks fail to align with children’s interactions, which show both non-conventional patterns and behaviors, e.g., word omissions, semantic errors, and incoherence (Liu & Fourtassi, 2024), while still responding in an age-appropriate manner (Chen et al., 2025b). 2) Second, there is a lack of a comprehensive taxonomy of evaluation protocols specific to children’s personalization. Current personalized evaluation studies mainly cover standard adult-centered preferences (e.g., general preference following) and fail to address the specific needs of children (e.g., generating child-appropriate content while fostering children’s creativity).

Facing these challenges, we introduce ChildEval, a benchmark to evaluate LLMs’ capacity to infer, interpret, and follow child-centered preferences in a long-context conversational setting. We focus on the preschool children (aged 3-6) who have a greater need for companionship from LLM-based chatbots. ChildEval comprises of 29K synthesized children’s persona profiles, which are related to their preferences in both explicit and implicit manners. Implicit preferences are integrated inside dialogues consisting of 6 to 10 turns. The preferences cover 5 top-level and 14 sub-level topics that involve children’s daily lives and development according to the guidelines published by the ministry of education (Antle, 2008; Wang, 2013). We further propose fine-grained evaluation protocols for child-oriented preferences to systematically evaluate the performance of open-source LLMs. Experimental results demonstrate the impact of various personalized representations on LLM responses and indicate that fine-tuning on this dataset may enhance performance.

2 RELATED WORK

2.1 DATA AND EVALUATION FOR PERSONALIZATION

Recent advances in LLM personalization have emphasized building data resources and evaluation benchmarks. PersonaMem (Jiang et al., 2025) and HiCUPID (Mok et al., 2025), for instance, simulate multi-attribute profiles and multi-turn conversations to assess whether models maintain user-specific consistency over time. Evaluation frameworks further consider metrics such as style alignment, preference fidelity, and user satisfaction (Salemi et al., 2023). On the data side, researchers have explored synthetic dialogue generation (Braga et al., 2024), profile summarization (Zhang, 2024), and memory retrieval from past interactions (Qian et al., 2025) to address the scarcity of high-quality personalized corpora.

Child-centered personalization has begun to emerge as a distinct line of research. KidLM (Nayeem & Rafiei, 2024) introduces a curated child-friendly corpus and training strategies for age-appropriate responses. Other works highlight the need for style simplification (Valentini et al., 2023) and safety evaluation frameworks tailored to children (Rath et al., 2025). Yet, most benchmarks and datasets target general users, leaving the open question of how well LLMs can adapt to children’s preferences in multi-turn interactions.

2.2 METHODOLOGICAL ADVANCES IN LLM PERSONALIZATION

Methodological advances in personalization can be grouped into three main directions. First, *prompt-based methods* leverage explicit profile descriptions or inferred traits from past interactions. Structured attributes (e.g., demographics) improve personalization (Liu et al., 2025), while implicit profiles capture subtler user traits (Li et al., 2021), and profile placement affects model behavior (Wu et al., 2024). Second, *memory-based methods* incorporate past interactions as context, either directly or via retrieval systems such as MemPrompt (Madaan et al., 2022). Hierarchical memory structures (Pan et al., 2025; Magister et al., 2024) and integrated frameworks like PRIME (Zhang et al., 2025a) enable robust long-term personalization. Third, *preference modeling* seeks to capture user preferences. While LLMs still struggle with adherence (Zhao et al., 2025a), persona synthesis (Ryan et al., 2025) and representation editing (Zhang et al., 2025b) improve alignment.

Beyond non-parametric approaches, *parametric adaptation* embeds user traits into model parameters via supervised fine-tuning (full or parameter-efficient methods like LoRA or prefix-tuning). Representative works include OPPU (Tan et al., 2024), E2P (Huber et al., 2025), and controllable vectors such as Hydra (Zhuang et al., 2024). Reinforcement learning methods similarly adapt reward functions to individual preferences, e.g., PRLHF (Li et al., 2024) and RLPA (Zhao et al., 2025b).

108 While these methods benefit the general population, their effectiveness for children—characterized
 109 by implicit interests and markedly different communication patterns—requires new evaluation pro-
 110 tocols.

111

112 3 THE CHILDEVAL BENCHMARK

113 3.1 PROBLEM FORMULATION

116 To evaluate whether an LLM can perceive and adapt to a child’s preference ρ when it communicates
 117 with the child, the full prompt sent to the model could be formulated by:

118

$$119 \quad \mathcal{B} = H + u^* \quad (1)$$

120 where

121

- 123 • $+$ denotes the concatenation of texts.
- 124 • $H = \{S_1, S_2, \dots, S_t, \dots, S_T\}$ denotes a multi-session conversation history between a
 125 child and an LLM. Each session $S_t = \{(u_{t,1}, m_{t,1}), \dots, (u_{t,K_t}, m_{t,K_t})\}$ consists of K_t
 126 dialogue turns, where u_{t,k_t} is the child utterance and m_{t,k_t} the model response.
- 127 • u^* is a child utterance related to the child preference ρ , and would be used as a utterance
 128 that the LLM shall respond to.

129

130 Sessions in H are categorized as:

131

- 132 • **Relevant session:** Following the setting of Zhao et al. (2025a), the first session S_1 of H is a
 133 session with dialogues related to the user preference queried by u^* .
- 134 • **Irrelevant session:** The remaining sessions of H contain dialogue turns unrelated to u^* .

135

136 In each relevant session, the child preference ρ can be revealed explicitly or implicitly:

137

- 138 • **Explicit:** Such a session contains a single dialogue turn $S_1 = \{(u_{1,1}, m_{1,1})\}$, and $u_{1,1}$
 139 directly expresses the child preference.
- 140 • **Implicit:** Such a session contains multiple dialogue turns, and the user preference could be
 141 implicitly inferred by partial user dialogue turns in this session.

142

143 The task used for evaluating the LLM could then be formulated as:

144

$$145 \quad f(p, \mathcal{B}) \longrightarrow \hat{m} \quad (2)$$

146

147 where p denotes the child persona, i.e. the persistent attributes that embody the child’s consistent
 148 personality traits (e.g., age and gender) and long-standing interests. $f(\cdot)$ denotes the model to be
 149 evaluated, and \hat{m} is the response generated by the model given the prompt \mathcal{B} . A good response shall
 150 align with the child persona p and the child’s preference ρ revealed by the historical conversation
 151 displaced in \mathcal{B} .

152


153 Given such a problem formulation, a sample data of ChildEval is illustrated in Figure 1(b). The
 154 ChildEval benchmark comprises the child persona, the child preference statement, which is identical
 155 to the child utterance in the historical dialogue that explicitly reveals the preference, the historical
 156 multi-turn dialogues that implicitly reveal the preference, and a piece of preference-related child
 157 utterance used for LLM response evaluation. The composition of the benchmark enables us to
 158 assess not only a model’s capability to identify and adapt to user preferences in dialogue but also
 159 how user personas influence the model’s behavior patterns during interaction.

160

161 3.2 DATA CONSTRUCTION PIPELINE OF CHILDEVAL

162

163 Step 1: We generate 29K child personas using Qwen2.5-72B through an iterative generation-and-
 164 refinement process. Semantically similar candidates are removed using FAISS (Douze et al., 2024)



178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201

Figure 1: Overview of the ChildEval benchmark.(a) The data construction pipeline.(b) A data sample consists of the child persona, the child preference statement, which is identical to the child utterance in the historical dialogue that explicitly reveals the preference, the historical multi-turn dialogues that implicitly reveal the preference, and a piece of preference-related child utterance used for LLM response evaluation.

based on text embeddings (Xiao et al., 2023), ensuring diversity and comprehensive representation of personality traits. To mitigate privacy risks, all child names and identifiers are systematically removed via LLM processing and human review.

Step 2: Each child preference is generated from a child persona and a sampled topic, covering 5 top-level and 14 sub-level categories, following established guidelines on children’s daily lives and development (Antle, 2008; Wang, 2013). For each child preference, a corresponding child utterance is also generated, which serves as the child’s turn in a child-LLM dialogue, providing input for the LLM’s response. Each child preference is expressed as a single first-person sentence and is linked to one sub-level topic and its corresponding top-level topic (see Table 1). Two preferences are generated per persona, yielding 58K preferences, of which 46K are retained after FAISS-based filtering of semantically similar instances.

Step 3: Historical dialogues implicitly revealing preferences are generated using prompt-based generation with self-verification, producing 6-10-turn child-LLM conversations. For more prompts, see Section A.5.

200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215

Table 1: Distribution of the 14 preference topics within ChildEval, which are related to children’s daily life and development.

Topic	Subtopic			
Art enlightenment (21.64%)	Music (6.68%)	Dance (6.74%)	Painting & Crafts (8.22%)	
Cognitive development and exploration (29.20%)	Science (7.72%)	Nature & universe (7.19%)	Technology (7.09%)	Learning (7.20%)
Nutrition and physical activity (13.11%)	Outdoor activity (6.24%)	Health eating (6.87%)		
Language and literacy development (22.09%)	Story (7.39%)	Language (7.31%)	Reading (7.39%)	
Social and emotional development (13.97%)	Social interaction (6.94%)	Play (7.03%)		

3.3 FINE-GRAINED EVALUATION METRICS FOR CHILD PREFERENCES

As we stated previously, there is a lack of a comprehensive taxonomy of evaluation protocols specific to children’s personalization. Existing personalized evaluation studies mainly cover standard adult-centered preferences. For example, Zhao et al. (2025a) and Jiang et al. (2025) focus on the

preference following performance from different settings. They fail to address the specific needs of children (e.g., generating child-appropriate content while fostering children’s creativity). Therefore, we propose fine-grained child-oriented evaluation metrics to systematically evaluate the LLM response to child utterance in addition to preference consistency.

(1) Preference Consistency (PC). This dimension assesses the LLM’s capability to produce responses aligned with explicitly expressed or implicitly inferred child preferences. This evaluation principle is universally applicable and not confined to child-centric dialogues, as any personalized dialogue system must meet this fundamental requirement. Since prior studies have already investigated preference consistency, our benchmark adopts the established evaluation criteria (Zhao et al., 2025a) for the unique context of child-oriented conversations.

(2) Child-Oriented Evaluation. In addition to maintaining preference consistency, child-centered dialogues necessitate extra fine-grained evaluation dimensions different from typical adult-oriented communication. Hence, we newly propose a set of novel child-oriented evaluation metrics, which concentrate on distinctive linguistic and contextual characteristics inherent in child-centered conversations and cover four sub-dimensions:

Emotional Adaptation (EA). The LLMs should be sensitive to the emotions expressed by the children, providing empathetic, supportive, and age-appropriate responses that help to maintain a positive atmosphere of interaction.

Interaction Scaffolding (IS). The LLM should be able to scaffold effective child-centered conversation with prompts, clarifications, or playful follow-ups in a natural conversational flow.

Developmental Appropriateness (DA). The LLM’s responses should match the cognitive and linguistic abilities of 3-to-6-year-old children, avoiding overly complex vocabulary or reasoning while providing informative and stimulating content.

Engagement (EG). The LLM should be able to produce lively and appealing utterances, using child-specific markers such as playful particles, reduplication, or culturally grounded scenarios, to keep children actively interested in the dialogue.

4 EXPERIMENTS

In addition to assessing whether an LLM can effectively infer and leverage child preferences during interactions with children, the experiments are designed to validate the positive influence of persona information in personalizing LLMs and to explore how such information can be efficiently integrated into LLMs with minimal computational overhead.

4.1 EXPERIMENTAL SETUP

We evaluate multiple state-of-the-art open-source models—Qwen2.5-3B-Instruct, Qwen3-4B-Instruct, LLaMA3.1-8B-Instruct, DeepSeek-R1-671B, and Mistral-7B-Instruct-v0.3—across three strategies for adapting to child preferences: prompt-based method (PBM), LoRA fine-tuning (LoRA), and our proposed persona steer model (PSM). PSM integrates a pluggable Persona Steer Module (Section A.6) to test whether child persona information in ChildEval enhances personalization. All experiments are zero-shot on a bilingual (Chinese & English) dataset. To examine the effect of persona information under varying context lengths, We construct long-context multi-session data to test LLMs’ ability to capture and apply child preferences in long context. Details of long context settings are provided in the Appendix A.2. Qwen2.5-3B-Instruct serves as the SFT backbone, with fine-grained evaluations conducted using Qwen2.5-72B-Instruct.

4.2 EVALUATING LLMs ON PREFERENCE CONSISTENCY

SOTA LLMs struggle to maintain personalization across long-term interactions. As shown in Figure 2, all prompt-based LLMs exhibit a decrease in accuracy when generating personalized responses after inserting irrelevant dialogues, compared to directly expressing preferences without any intervening turns. However, as the number of irrelevant turns increases, the performance degradation gradually slows down. Interestingly, for some models (e.g. Qwen3-4B-instruct), additional irrelevant turns even lead to a slight recovery or improvement, suggesting a potential stabilizing.




Figure 2: Zero-shot consistency of LLMs with children’s explicit (left) and implicit (right) preferences across n-turn dialogues. Each n-turn dialogue uses a fixed token set (See Table 3)

LLMs face greater difficulty in deducing implicit preferences than in understanding explicit ones. Comparing the results in the left and right panels of Figure 2, it is evident that personalization consistency on implicit-preference datasets is lower than on explicit-preference datasets across almost all the LLMs evaluated. This suggests that inferring user preferences from dialogue context poses greater challenges for LLMs than directly leveraging explicitly stated preferences. The gap highlights the difficulty of capturing subtle cues embedded in conversation, underscoring the need for more robust mechanisms to enhance implicit personalization.

Incorporating persona enhances the model’s personalized outputs. As shown in Figure 3, adding persona information in the prompts consistently improves performance across all models. The improvement is most pronounced on Qwen3-4B, where accuracy on the explicit dataset increases from 78.7% to 89.1%, while the smallest gain is observed on Qwen2.5-3B for the implicit dataset, improving from 75.7% to 75.8%. The varying magnitude of improvement across models suggests that some LLMs are more effective at leveraging persona cues to generate personalized responses, whereas others exhibit smaller gains, reflecting differences in how models utilize persona information. This underscores that the ability to exploit persona is model-dependent.




Figure 3: Performance of models on preference consistency: response on dataset with 5 irrelevant turns inserted under 0-shot prompting (with vs. without persona).

4.3 EVALUATING LLMs ON CHILD-ORIENTED EVALUATION

Personality preference consistency does not align with child-oriented capabilities. Comparing data in Table 2 with data in Figure 3 across multiple models, while consistency accuracy may be similar (e.g., 74%–76%), performance on child-oriented evaluation varies widely across dimensions, particularly in IS and DA. This suggests that a high consistency score alone does not necessarily reflect strong child-oriented personalization.

LLMs show limited capability in Interaction Scaffolding (IS). Across all models, performance on the IS dimension is lower than on other evaluation dimensions. For example, on the explicit dataset, Qwen2.5-3B-Instruct achieves 35.8% accuracy on IS. This substantial gap highlights a key

324
 325 Table 2: Performance of models on child-oriented evaluation: response on dataset with 5 irrelevant
 326 turns inserted under 0-shot prompting (with vs. without persona).

Model	Without Persona				With Persona			
	EA	IS	DA	EG	EA	IS	DA	EG
Explicit Data (%)								
Qwen2.5-3B-instruct	77.23	35.8	97.31	75.99	94.50	70.13	96.51	93.82
Qwen3-4B-instruct	96.29	52.28	99.82	97.59	98.05	82.19	99.58	96.96
Llama3.1-8B-instruct	79.02	28.42	89.66	72.66	86.67	59.06	93.66	83.71
DeepSeek-R1	88.25	50.24	98.34	87.73	95.75	79.72	98.59	97.55
Implicit Data (%)								
Qwen2.5-3B-instruct	78.73	38.58	98.01	77.41	93.16	69.92	96.33	93.67
Qwen3-4B-instruct	96.45	59.85	99.84	97.64	97.39	83.22	99.85	97.09
Llama3.1-8B-instruct	79.89	28.67	91.67	74.4	84.55	57.02	94.74	81.72
DeepSeek-R1	88.34	55.96	98.92	88.97	94.61	80.07	99.02	96.75

341 limitation of current approaches, as sensitivity to subtle cues is critical for building engaging and
 342 personalized child interactions.

343
 344 **LLMs exhibit considerable variation across dimensions in child-oriented evaluation.** In par-
 345 ticular, models consistently achieve much stronger results on the DA dimension (e.g., Qwen3-4B
 346 achieves 99.82%) compared with other dimensions, underscoring a clear imbalance across subtasks.
 347 Such uneven distribution suggests that the evaluation of child-oriented dialogue systems must be
 348 multi-dimensional, as relying on aggregated or single metrics may conceal important deficiencies.

349
 350 **Incorporating child persona leads to improvements across all evaluation dimensions of the**
 351 **child-oriented evaluation.** The most substantial improvements are observed in EA, IS, and EG,
 352 where absolute and relative increases are notably larger. By contrast, DA dimension also improves,
 353 but with a smaller margin. This pattern suggests that child persona information primarily strengthens
 354 dimensions tied to individual child preferences and sensitivity to implicit cues, while its influence
 355 on group-level preferences, such as DA, which catches broader developmental norms, remains more
 356 modest.

357
 358 **LLMs consistently struggle to maintain child-oriented evaluation performance over long-term**
 359 **interactions.** As shown in Figure 4, although the overall trend under irrelevant dialogue insertion
 360 resembles that of Preference Consistency, a key distinction emerges: the performance difference be-
 361 tween explicit and implicit datasets is relatively small. This suggests that, in child-oriented settings,
 362 models are less reliant on whether preferences are directly or indirectly expressed, and can preserve
 363 comparable dialogue quality across both conditions.

364 4.4 FINE-TUNING ON CHILDEVAL TO ENHANCE CHILD PERSONALIZATION

365
 366 **Supervised Fine-Tuning on ChildEval leads to consistent improvements in children’s person-
 367 alization performance across open-source LLMs.** As illustrated in Figure 5, applying LoRA SFT,
 368 both with and without persona injection, leads to substantial gains in both preference consistency
 369 and child-oriented evaluation compared with the base models. Interestingly, LoRA SFT with per-
 370 sona shows slightly lower improvements in preference consistency than LoRA without persona.
 371 One possible reason is that adding persona signals may introduce additional constraints, and the
 372 persona itself may contain noise related to the explicit and implicit preference expressions in the
 373 ChildEval dataset, which could slightly limit the model’s ability to fully optimize for consistency
 374 and child-oriented performance.

375
 376 **The choice of persona utilization strategy significantly affects the performance of models fine-
 377 tuned with SFT.** As shown in Figure 5, LoRA, which incorporates persona directly into dialogue
 378 prompts, achieves higher preference consistency than the PSM approach that encodes persona as
 379 vectors. The advantage is more evident on the explicit dataset, while the difference on the implicit
 380 dataset is relatively small. On child-oriented benchmarks, both strategies show limited differences.
 381 A possible explanation for the stronger gains of LoRA on the explicit dataset is related to its injection




Figure 4: Accuracy of LLMs on different dimensions of child-oriented evaluation with varying numbers of inserted irrelevant turns (n-turn).




Figure 5: Fine-tuning results for children’s personalities on explicit and implicit datasets (both with 5-turn test dialogues). Explicit training added 5 unrelated utterances; implicit training used 6–10 consecutive turns. “Persona” denotes inclusion of child persona information during fine-tuning.

mechanism. The explicit dataset contains many irrelevant dialogues, and LoRA—by incorporating persona cues at the prompt level—may help the model maintain consistency under such noisy contexts. In contrast, PSM only adjusts persona information at the final vector layer, which can be considered a post-hoc modification. Since it does not directly influence the model’s intermediate reasoning or attention distribution, persona signals may be more easily diluted in the presence of irrelevant dialogues, potentially resulting in weaker consistency compared with LoRA.

LLMs exhibit the most marked improvement in Interaction Scaffolding (IS) after fine-tuning. One possible reason is that IS tasks require the model to generate coherent and contextually appropriate responses, which benefit directly from the additional supervision provided during fine-tuning. Fine-tuning helps the model better capture the underlying patterns of guidance and scaffolding strategies in child-oriented dialogues, enabling more effective interaction management.

5 ERROR TYPE ANALYSIS

Preference consistency errors include Unhelpful Response, Inconsistency Violation, Preference Hallucination Violation, and Preference-Unaware Violation (Zhao et al., 2025a). Figure 6 shows their distribution across 10-turn dialogues on explicit and implicit datasets under different methods. Initially, Preference-Unaware Violations dominate, reflecting LLMs’ limited awareness of user prefer-




Figure 6: Distribution of preference consistency errors across 10-turn dialogues. Base refers to Qwen2.5-3B-Instruct; Base+Persona applies prompting with persona. LoRA-5 and PSM-5 denote LoRA- and PSM-based methods trained with 5-turn inserted context, with or without persona. LoRA-implicit and PSM-implicit are trained with implicit context.

ences. With various methods, this error decreases while Inconsistency Violations appear, indicating ongoing challenges in generating preference-aligned responses. Fine-tuning methods amplify inconsistency errors compared to prompt-based approaches. Incorporating persona information has mixed effects: LoRA shows more Inconsistency Violations than PSM, while on the implicit dataset, PSM produces many Unhelpful Responses (79%), whereas LoRA and prompt-based methods exhibit more preference-related errors, reflecting a trade-off between proactive preference-following and reliability.

We further analyze the effect of inserted context length on preference consistency (details in Appendix A.7.1 and A.7.2). Under zero-shot prompting, Preference-Unaware Violations rise with longer irrelevant context. Fine-tuning methods reduce these violations but show trade-offs: LoRA tends toward Inconsistency Violations, while PSM shifts from Preference Hallucinations in short contexts to Unhelpful Responses in longer ones. On explicit datasets, LoRA remains proactive; PSM becomes conservative, especially on implicit datasets.

6 CONCLUSION

In this work, we introduced ChildEval, a benchmark designed to evaluate LLMs’ ability to infer, interpret, and follow child-centered preferences in long-context conversational settings. By constructing 29K synthetic persona profiles and multi-session dialogues grounded in preschool children’s daily lives, we provide a systematic evaluation framework that highlights the strengths and limitations of existing open-source LLMs in child-focused personalization. Our experiments reveal that LLMs struggle to consistently capture personal preferences, particularly implicit ones in long-interaction scenarios. Moreover, we find that different personalized representations affect their responses, and that tailored fine-tuning on this dataset can enhance performance, highlighting promising avenues for advancing child-centric AI systems.

As future work, we plan to extend ChildEval toward richer real-world scenarios and more diverse developmental contexts. An important next step is to ensure that personalization methods for children not only improve alignment with preferences but also rigorously safeguard privacy and safety, which are critical when designing trustworthy LLM-based companions for young users.

486 **7 ETHICS STATEMENT**

487

488 This research adheres to the ethical standards by confirming that no studies involving human par-
 489 ticipants or animals were conducted by any of the authors. The absence of human/animal subjects
 490 ensures full alignment with ethical principles of voluntary participation, informed consent, and non-
 491 harm. Data analysis relies exclusively on publicly available or synthetic datasets, with no privacy
 492 violations or welfare concerns. This approach reflects our commitment to integrity in scholarly
 493 work. All codes and data will be open-source.

494

495 **REFERENCES**

496

497 Anthropic. The claude 3 model family: Opus, sonnet, haiku. 2024. URL <https://api.semanticscholar.org/CorpusID:268232499>.

498

500 Alissa N Antle. Child-based personas: need, ability and experience. *Cognition, Technology & Work*,
 501 10(2):155–166, 2008.

502 Marco Braga, Pranav Kasela, Alessandro Raganato, and Gabriella Pasi. Synthetic data generation
 503 with large language models for personalized community question answering. In *2024 IEEE/WIC
 504 International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT)*, pp.
 505 360–366. IEEE, 2024.

506

507 Jiaju Chen, Minglong Tang, Yuxuan Lu, Bingsheng Yao, Elissa Fan, Xiaojuan Ma, Ying Xu, Dakuo
 508 Wang, Yuling Sun, and Liang He. Characterizing llm-empowered personalized story reading and
 509 interaction for children: Insights from multi-stakeholder perspectives. In *Proceedings of the 2025
 510 CHI Conference on Human Factors in Computing Systems*, pp. 1–24, 2025a.

511 Jiaju Chen, Minglong Tang, Yuxuan Lu, Bingsheng Yao, Elissa Fan, Xiaojuan Ma, Ying Xu, Dakuo
 512 Wang, Yuling Sun, and Liang He. Characterizing llm-empowered personalized story reading and
 513 interaction for children: Insights from multi-stakeholder perspectives. *Conference on Human
 514 Factors in Computing Systems*, 2025b.

515

516 Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvassy, Pierre-
 517 Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. The faiss library. 2024.

518

519 Arafat Md Easin, Saha Sourav, and Orosz Tamás. An intelligent llm-powered personalized assistant
 520 for digital banking using langgraph and chain of thoughts. In *2024 IEEE 22nd Jubilee Interna-
 521 tional Symposium on Intelligent Systems and Informatics (SISY)*, pp. 625–630. IEEE, 2024.

522

523 Tiantian Feng, Anfeng Xu, Rimita Lahiri, Helen Tager-Flusberg, So Hyun Kim, Somer Bishop,
 524 Catherine Lord, and Shrikanth Narayanan. Can generic llms help analyze child-adult interactions
 525 involving children with autism in clinical observation?, 2024. URL <https://arxiv.org/abs/2411.10761>.

526

527 Gemini, :, Rohan Anil, Sebastian Borgeaud, and et al. Gemini: A family of highly capable multi-
 528 modal models, 2025. URL <https://arxiv.org/abs/2312.11805>.

529

530 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
 531 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
 532 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

533

534 Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,
 535 and Weizhu Chen. Lora: Low-rank adaptation of large language models, 2021. URL <https://arxiv.org/abs/2106.09685>.

536

537 Bernd Huber, Ghazal Fazelnia, Andreas Damianou, Sebastian Peleato, Max Lefarov, Praveen
 538 Ravichandran, Marco De Nadai, Mounia Lalmas-Roellke, and Paul N Bennett. Embedding-to-
 539 prefix: Parameter-efficient personalization for pre-trained large language models. *arXiv preprint
 arXiv:2505.17051*, 2025.

540 Bowen Jiang, Zhuoqun Hao, Young-Min Cho, Bryan Li, Yuan Yuan, Sihao Chen, Lyle H. Ungar,
 541 Camillo J. Taylor, and Dan Roth. Know me, respond to me: Benchmarking llms for dynamic
 542 user profiling and personalized responses at scale. *CoRR*, abs/2504.14225, 2025. doi: 10.48550/
 543 ARXIV.2504.14225. URL <https://doi.org/10.48550/arXiv.2504.14225>.

544

545 Ishita Kumar, Snigdha Viswanathan, Sushrita Yerra, Alireza Salemi, Ryan A. Rossi, Franck Der-
 546 noncourt, Hanieh Deilamsalehy, Xiang Chen, Ruiyi Zhang, Shubham Agarwal, Nedim Lipka,
 547 Chien Van Nguyen, Thien Huu Nguyen, and Hamed Zamani. Longlamp: A benchmark for per-
 548 sonalized long-form text generation, 2024. URL <https://arxiv.org/abs/2407.11016>.

549

550 Juntao Li, Chang Liu, Chongyang Tao, Zhangming Chan, Dongyan Zhao, Min Zhang, and Rui Yan.
 551 Dialogue history matters! personalized response selection in multi-turn retrieval-based chatbots.
 552 *ACM Transactions on Information Systems (TOIS)*, 39(4):1–25, 2021.

553

554 Xinyu Li, Zachary C. Lipton, and Liu Leqi. Personalized language modeling from personalized
 555 human feedback. *CoRR*, abs/2402.05133, 2024. doi: 10.48550/ARXIV.2402.05133. URL
 556 <https://doi.org/10.48550/arXiv.2402.05133>.

557

558 Jiahong Liu, Zexuan Qiu, Zhongyang Li, Quanyu Dai, Jieming Zhu, Minda Hu, Menglin Yang, and
 559 Irwin King. A survey of personalized large language models: Progress and future directions.
 560 *arXiv preprint arXiv:2502.11528*, 2025.

561

562 Jing Liu and Abdellah Fourtassi. Benchmarking llms for mimicking child-caregiver language in
 563 interaction. *arXiv preprint arXiv:2412.09318*, 2024.

564

565 Aman Madaan, Niket Tandon, Peter Clark, and Yiming Yang. Memory-assisted prompt editing to
 566 improve GPT-3 after deployment. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (eds.),
 567 *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing,*
 568 *EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pp. 2833–2861. Asso-
 569 ciation for Computational Linguistics, 2022. doi: 10.18653/V1/2022.EMNLP-MAIN.183. URL
 570 <https://doi.org/10.18653/v1/2022.emnlp-main.183>.

571

572 Lucie Charlotte Magister, Katherine Metcalf, Yizhe Zhang, and Maartje ter Hoeve. On the way to
 573 llm personalization: Learning to remember user conversations. *arXiv preprint arXiv:2411.13405*,
 574 2024.

575

576 Jisoo Mok, Ik hwan Kim, Sangkwon Park, and Sungroh Yoon. Exploring the potential of llms as
 577 personalized assistants: Dataset, evaluation, and analysis, 2025. URL <https://arxiv.org/abs/2506.01262>.

578

579 Mir Tafseer Nayeem and Davood Rafiei. KidLM: Advancing language models for children – early
 580 insights and future directions. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.),
 581 *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*,
 582 pp. 4813–4836, Miami, Florida, USA, November 2024. Association for Computational Linguis-
 583 tics. doi: 10.18653/v1/2024.emnlp-main.277. URL <https://aclanthology.org/2024.emnlp-main.277/>.

584

585 OpenAI, :, Aaron Jaech, Adam Kalai, and et al. Openai o1 system card, 2024. URL <https://arxiv.org/abs/2412.16720>.

586

587 Zhuoshi Pan, Qianhui Wu, Huiqiang Jiang, Xufang Luo, Hao Cheng, Dongsheng Li, Yuqing Yang,
 588 Chin-Yew Lin, H. Vicky Zhao, Lili Qiu, and Jianfeng Gao. On memory construction and
 589 retrieval for personalized conversational agents, 2025. URL <https://arxiv.org/abs/2502.05589>.

590

591 Hongjin Qian, Zheng Liu, Peitian Zhang, Kelong Mao, Defu Lian, Zhicheng Dou, and Tiejun
 592 Huang. Memorag: Boosting long context processing with global memory-enhanced retrieval
 593 augmentation. In Guodong Long, Michale Blumestein, Yi Chang, Liane Lewin-Eytan, Zi He-
 594 llen Huang, and Elad Yom-Tov (eds.), *Proceedings of the ACM on Web Conference 2025, WWW*
 595 *2025, Sydney, NSW, Australia, 28 April 2025- 2 May 2025*, pp. 2366–2377. ACM, 2025. doi:
 596 10.1145/3696410.3714805. URL <https://doi.org/10.1145/3696410.3714805>.

594 Yilun Qiu, Xiaoyan Zhao, Yang Zhang, Yimeng Bai, Wenjie Wang, Hong Cheng, Fuli Feng, and Tat-
 595 Seng Chua. Measuring what makes you unique: Difference-aware user modeling for enhancing
 596 llm personalization. *arXiv preprint arXiv:2503.02450*, 2025.

597 Prasanjit Rath, Hari Shrawgi, Parag Agrawal, and Sandipan Dandapat. Llm safety for children,
 598 2025. URL <https://arxiv.org/abs/2502.12552>.

600 Michael J. Ryan, Omar Shaikh, Aditri Bhagirath, Daniel Frees, William Held, and Diyi Yang.
 601 SynthesizeMe! inducing persona-guided prompts for personalized reward models in LLMs. In
 602 Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Pro-
 ceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long
 Papers)*, pp. 8045–8078, Vienna, Austria, July 2025. Association for Computational
 603 Linguistics. ISBN 978-8-98176-251-0. doi: 10.18653/v1/2025.acl-long.397. URL
 604 <https://aclanthology.org/2025.acl-long.397/>.

605 Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. Lamp: When large
 606 language models meet personalization. *arXiv preprint arXiv:2304.11406*, 2023.

607 Woosuk Seo, Chanmo Yang, and Young-Ho Kim. Chacha: Leveraging large language models
 608 to prompt children to share their emotions about personal events. In *Proceedings of the CHI
 Conference on Human Factors in Computing Systems*, CHI ’24, pp. 1–20. ACM, May 2024. doi:
 609 10.1145/3613904.3642152. URL <http://dx.doi.org/10.1145/3613904.3642152>.

610 Zhaoxuan Tan, Qingkai Zeng, Yijun Tian, Zheyuan Liu, Bing Yin, and Meng Jiang. Democ-
 611 ratizing large language models via personalized parameter-efficient fine-tuning. In Yaser Al-
 612 Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on
 613 Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November
 12-16, 2024*, pp. 6476–6491. Association for Computational Linguistics, 2024. doi:
 614 10.18653/V1/2024.EMNLP-MAIN.372. URL <https://doi.org/10.18653/v1/2024.emnlp-main.372>.

615 Maria Valentini, Jennifer Weber, Jesus Salcido, Téa Wright, Eliana Colunga, and Katharina
 616 Kann. On the automatic generation and simplification of children’s stories. *arXiv preprint
 617 arXiv:2310.18502*, 2023.

618 Ping Wang. *Interpretation of the Learning and Development Guidelines for Children Aged 3-6*.
 619 Northeast Normal University Press, 2013.

620 Bin Wu, Zhengyan Shi, Hossein A Rahmani, Varsha Ramineni, and Emine Yilmaz. Under-
 621 standing the role of user profile in the personalization of large language models. *arXiv preprint
 622 arXiv:2406.17803*, 2024.

623 Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighoff. C-pack: Packaged resources to
 624 advance general chinese embedding, 2023.

625 Xuhai Xu, Bingsheng Yao, Ziqi Yang, Shao Zhang, Ethan Rogers, Stephen Intille, Nawar Shara,
 626 Guodong Gao, and Dakuo Wang. Talk2care: Facilitating asynchronous patient-provider commu-
 627 nication with large-language-model. In *Proceedings of the AAAI Symposium Series*, pp. 146–151,
 628 2024.

629 Jiarui Zhang. Guided profile generation improves personalization with llms. *arXiv preprint
 630 arXiv:2409.13093*, 2024.

631 Xinliang Frederick Zhang, Nick Beauchamp, and Lu Wang. PRIME: large language model person-
 632 alization with cognitive memory and thought processes. *CoRR*, abs/2507.04607, 2025a. doi: 10.
 633 48550/ARXIV.2507.04607. URL <https://doi.org/10.48550/arXiv.2507.04607>.

634 Yijing Zhang, Dyah Adila, Changho Shin, and Frederic Sala. Personalize your LLM: fake it
 635 then align it. In Luis Chiruzzo, Alan Ritter, and Lu Wang (eds.), *Findings of the Associa-
 636 tion for Computational Linguistics: NAACL 2025, Albuquerque, New Mexico, USA, April
 637 29 - May 4, 2025*, pp. 7287–7301. Association for Computational Linguistics, 2025b. doi:
 638 10.18653/V1/2025.FINDINGS-NAACL.407. URL <https://doi.org/10.18653/v1/2025.findings-naacl.407>.

648 Siyan Zhao, Mingyi Hong, Yang Liu, Devamanyu Hazarika, and Kaixiang Lin. Do llms recognize
 649 your preferences? evaluating personalized preference following in llms. In *The Thirteenth In-*
 650 *ternational Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025.*
 651 OpenReview.net, 2025a. URL <https://openreview.net/forum?id=QWunLKbBGF>.
 652

653 Weixiang Zhao, Xingyu Sui, Yulin Hu, Jiahe Guo, Haixiao Liu, Biye Li, Yanyan Zhao, Bing Qin,
 654 and Ting Liu. Teaching language models to evolve with users: Dynamic profile modeling for
 655 personalized alignment. *CoRR*, abs/2505.15456, 2025b. doi: 10.48550/ARXIV.2505.15456. URL
 656 <https://doi.org/10.48550/arXiv.2505.15456>.
 657

658 Wenting Zhao, Xiang Ren, Jack Hessel, Claire Cardie, Yejin Choi, and Yuntian Deng. Wildchat: 1m
 659 chatgpt interaction logs in the wild, 2024. URL <https://arxiv.org/abs/2405.01470>.
 660

661 Yuchen Zhuang, Haotian Sun, Yue Yu, Rushi Qiang, Qifan Wang, Chao Zhang, and Bo Dai. HY-
 662 DRA: model factorization framework for black-box LLM personalization. In Amir Globersons,
 663 Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng
 664 Zhang (eds.), *Advances in Neural Information Processing Systems 38: Annual Conference on*
 665 *Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December*
 666 *10 - 15, 2024*, 2024. URL http://papers.nips.cc/paper_files/paper/2024/hash/b6b4906c1334656e97cc9968ccfca073-Abstract-Conference.html.
 667

668 A APPENDIX

669 A.1 USAGE OF LARGE LANGUAGE MODEL

670 We appreciate the assistance provided by DeepSeek-R1 (Guo et al., 2025) , ChatGPT (OpenAI et al.,
 671 2024) in writing aid and sentence-level polishing.

672 A.2 LONG-CONTEXT SETTINGS

673 To simulate realistic conversational dynamics, we adopt a methodology similar to Zhao et al.
 674 (2025a). We incorporate multi-session dialogue turns from the WildChat-1M dataset (Zhao et al.,
 675 2024), which contains real user and LLM interactions across diverse topics. For SFT training, we
 676 randomly select 3, 5, and 10 round conversations to construct three training sets. For testing, we
 677 sample multi-session contexts up to 21K tokens, interleaving dialogue turns between the disclosure
 678 of children’s preferences and the related utterances. Although we initially considered extending
 679 dialogues to 50K tokens, the backbone model supports at most 30K, beyond which outputs be-
 680 came unstable. This setup creates a challenging evaluation for LLMs to infer, retrieve, and utilize
 681 children’s preferences in long dialogues, especially when interspersed with unrelated topics. For
 682 dialogues of varying turn counts, we randomly sample and fix their lengths, with token statistics
 683 reported in Table 3.
 684

685 Table 3: Token number in the long context.
 686

687 Number of Turns in the Long Context	688 5-turn	689 10-turn	690 20-turn	691 30-turn	692 50-turn
693 Chinese	694 2156	695 4369	696 10390	697 12389	698 15380
699 English	700 2754	701 4010	702 10522	703 12420	704 21817

705 A.3 MODEL VERSION

706 In our experiments, we employ the Bge-Large-Zh model as the text encoder. Table 4 provides an
 707 overview of the evaluated LLMs and their versions, together with the text encoder version.
 708

709 A.4 CHILDEVAL EXAMPLE

710 An example from ChildEval is presented in Table 5.
 711

702
 703 Table 4: Overview of the benchmarked LLMs, their versions, and the text encoder version used in
 704 the experiments.

705 Model Name	706 Version
706 Qwen2.5-3B-Instruct	707 qwen.qwen2.5-3B-instruct-v1:0
707 Qwen3-4B-Instruct	708 qwen.qwen3-4B-instruct-v1:0
708 LLaMA3.1-8B-Instruct	709 meta.llama3.1-8b-instruct-v1:0
709 Mistral-7B-Instruct	710 mistral.mistral-7b-instruct-v0:3
710 DeepSeek-R1-671B	711 deepseek-ai.deepseek-r1-v1:0
711 Bge-Large-Zh	baai.bge-large-zh-v1:5

712
 713 **A.5 PROMPT DESIGN AND EXAMPLES**

714
 715 **A.5.1 PROMPTS FOR DATA CONSTRUCTION**

716
 717 The prompts used within this work are listed in Figures 7–9. Some prompts are too long to fit on a
 718 single page, so we split them into two figures, as shown in Figure 8 and 9.

719
 720 **A.5.2 PROMPTS FOR THE PROMPTING-BASED APPROACH**

721
 722 We extensively evaluate a variety of state-of-the-art LLMs using zero-shot prompts, both with and
 723 without persona information. In the default zero-shot setting, the LLM answers the user’s query
 724 directly without any additional prompting. However, these models are not specifically designed for
 725 child-oriented dialogue. If used without modification, they tend to generate overly long responses
 726 that do not reflect the conversational style of young children. To ensure a fair evaluation, we accord-
 727 ingly augmented the original dialogue prompts as follows, corresponding to the with-persona and
 728 without-persona settings.

729
 730 **zero-shot-without-persona:** Provide clear, concise, and conversational responses in 1-3
 731 sentences, prioritizing accuracy and a friendly tone while avoiding unnecessary details.

732
 733 **zero-shot-with-persona:** Never use any names or personal identifiers from the profile
 734 “{persona}”. Always address the child directly as ‘you’ when it feels natural, or give sug-
 735 gestions without using a subject, based on the user information in the profile. Provide clear,
 736 concise, and conversational responses in 1-3 sentences, prioritizing accuracy and a friendly
 737 tone.

738
 739 **A.5.3 EVALUATION PROMPTS FOR CHILD-ORIENTED TASKS**

740
 741 The evaluation prompts for child-oriented tasks are shown in Figures 10–13, which correspond
 742 respectively to Emotional Adaptation, Interaction Scaffolding, Developmental Appropriateness and
 743 Engagement.

744
 745 **A.6 ARCHITECTURE OF THE PERSONA STEER MODEL**

746
 747 To assist in examining whether providing the child persona information in our benchmark would
 748 contribute to better LLM personalization, we propose a persona steer model that leverages persona
 749 information to guide the LLM’s outputs toward personalized behaviors. The architecture of our
 750 persona steer model is depicted in Figure 14, whose core is the Personalized Steer Module. While the
 751 pre-trained LLM provides robust general language comprehension and generation, the Personalized
 752 Steer Module enables effective user adaptation without huge computational burdens.

753
 754 Specifically, as shown in Figure 14, the Personalized Steer Module is designed to introduce user-
 755 specific information into the language model in a precisely controlled manner. A Persona Encoder
 processes textual descriptions of children’s personas, converting them into dense embeddings that
 effectively encapsulate the semantic essence of user characteristics. The persona embedding is then

756

757

758

759

760

761

762

763

764

765

766

767

768

769

770

771

772

773

774

775

776

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

Table 5: An example set from ChildEval.

Topic	Persona	Explicit Preference	Implicit Preference	Utterance
Cognitive development and exploration Nature&universe	5-year-old boy who shows a strong interest in technology and mechanics. He enjoys playing with LEGO bricks, especially the sets that can be assembled into vehicles and robot models. His room is filled with various completed builds, each brimming with creativity and imagination. He also likes to play video games, particularly those that require puzzle-solving and strategy. He has a certain sensitivity towards numbers and programming, enjoying the creation of his own simple games using basic programming software. He is not very fond of sports and outdoor activities, preferring to stay indoors exploring new things. He has a mild allergic constitution and needs to be particularly cautious with certain foods and pollen. At bedtime, he loves listening to stories about technology and inventions told by his father, dreaming of becoming an inventor one day.	I really hate pollen; as soon as I come into contact with it, I start sneezing and my eyes get itchy. It's so uncomfortable.	user: Dad said that spring has arrived, and we can go to the park to play. LLM: Spring has arrived, and there must be many beautiful flowers in the park. Which flower is your favorite? user: I don't like flowers; the pollen makes me sneeze and my eyes itch. LLM: That does sound uncomfortable. So, what do you like to do in the park? user: I like to play in clean places, such as a sandbox, where I can build sandcastles. LLM: Playing in the sandbox is indeed a lot of fun! What kind of sandcastles have you built? user: I built very tall sandcastles and decorated them with small stones. LLM: It sounds like you're really good at playing with sand! Can you make other things with sand as well? user: I also make little cars and animals with sand, which is a lot of fun. LLM: You're so creative! What's your favorite small animal to make? user: I love making dogs and cats; they are so cute. LLM: Puppies and kittens are indeed very cute! Can you also make these little animals with building blocks? user: I build small cars and robots with LEGO bricks, and sometimes I even make dogs and cats. LLM: The things you build with LEGO bricks must look great! What is your most impressive creation? user: My proudest creation is a moving robot that even has glowing eyes. LLM: Wow, that's amazing! You're really a little inventor! user: Dad said he would take me to the botanical garden in spring, where there are many beautiful flowers. Have you been there? LLM: Wow, there are indeed many beautiful flowers in the botanical garden, in all sorts of colors; they are really lovely! However, since you're allergic to pollen, remember to wear a mask and sunglasses, and put on long-sleeved clothes when you go. This can help reduce your contact with pollen. You can also bring a small bottle of water to clean your nose and eyes if you feel uncomfortable. This way, you can enjoy the beautiful flowers happily! What do you think about this?	Dad said he would take me to the botanical garden in spring, where there are many beautiful flowers. Have you been there?

810 aggregated with the output of a Steer Adapter module, which contains information extracted from
 811 the historical dialogue by the LLM. The aggregated embedding $e_{aggregate}$ is then processed by a
 812 Gating Network, whose gate value $g(\cdot)$ is dynamically calculated based on h_L . The final output
 813 embedding e'_v before applying the *Softmax* function could be obtained by:
 814

$$815 \quad e'_v = h_L W^T + g(h_L) e_{aggregate} W^T = h_L W^T + \Delta(h_L) W^T \quad (3)$$

816 where W is the weight matrix employed by the LLM. It should be noted that only the Steer Adapter
 817 and the Gating Network used to calculate $\Delta(h_L)$ contain parameters learnable during the training,
 818 while other parameters remain fixed.

819 To minimize the learnable parameters in the Steer Adapter, we implement a parameterization ap-
 820 proach inspired by low-rank adaptation (LoRA) (Hu et al., 2021). Specifically, rather than learning
 821 a full-rank transformation matrix, we decompose it into two low-rank matrices, as shown in Figure
 822 14. This decomposition maintains the transformation’s expressive power while allowing efficient
 823 integration of personalized information, seamlessly merging it into the LLM’s representations to
 824 facilitate effective user adaptation and stable generation. Additionally, it opens possibilities for in-
 825 corporating more sophisticated personalized models into LLM generation.

826 A.7 ADDITIONAL RESULTS

827 A.7.1 EFFECT OF INSERTED CONTEXT LENGTH ON PREFERENCE CONSISTENCY ERROR 828 TYPES

829 Figures 15 and 16 illustrate the changes in preference consistency error types across different num-
 830 bers of inserted irrelevant turns in the explicit and implicit datasets, respectively. Under zero-shot
 831 prompting without persona, Preference-Unaware Violations become increasingly prominent as the
 832 number of irrelevant turns increases, indicating that LLMs struggle more to maintain awareness
 833 of user preferences when exposed to longer irrelevant context. With the introduction of various
 834 methods, including fine-tuning approaches such as LoRA and PSM, the proportion of Preference-
 835 Unaware Violations decreases, while Hallucination Violations increase and Inconsistency Violations
 836 begin to appear, reflecting the challenges models face in generating responses that are both aligned
 837 with retrieved preferences and free from hallucinated information.

838 On the explicit dataset, LoRA is more prone to Inconsistency Violations across n-turn scenarios,
 839 whereas PSM exhibits higher rates of Preference Hallucination Violations in shorter contexts; how-
 840 ever, as the number of irrelevant turns increases beyond 30, the rate of Preference Hallucination
 841 Violations in PSM decreases, while Unhelpful Responses become increasingly dominant. On the
 842 implicit dataset, Unhelpful Responses constitute the primary error type for PSM, indicating a ten-
 843 dency to refuse or provide unhelpful answers rather than attempt alignment with user preferences.

844 Overall, these results highlight the trade-offs between proactive preference-following and robustness
 845 to irrelevant context. Notably, as the length of irrelevant context increases, PSM becomes increas-
 846 ingly conservative, producing unhelpful responses, whereas LoRA is more proactive, continuing to
 847 attempt responses aligned with user preferences, although alignment issues remain.

848 A.7.2 EFFECT OF INSERTED CONTEXT LENGTH ON FINE-TUNING RESULTS

849 Figure 17 illustrates how persona-informed finetuning methods (i.e., PSM and LoRA) evolve with
 850 increasing dialogue length on both explicit and implicit datasets (i.e., datasets with historical dia-
 851 logues that explicitly and implicitly reveal the child preference). On the explicit dataset, PSM-based
 852 models show a relatively sharp decline in preference consistency as the number of inserted irrelevant
 853 dialogue turns increases, while LoRA-based models exhibit a moderate decrease. Moreover, results
 854 on the PC (Preference Consistency) dimension indicate that training with longer irrelevant dialogues
 855 yields greater robustness on equally long test dialogues than training with shorter ones. Interestingly,
 856 within the child-oriented dimensions, most metrics remain relatively stable across dialogue lengths,
 857 whereas developmental appropriateness (DA) exhibits the largest fluctuations, indicating its height-
 858 ened sensitivity to contextual length.

864 On the implicit dataset, model trends largely mirror those observed on the explicit dataset. In the
 865 PC dimension, PSM-based models remain relatively stable compared to LoRA-based models on
 866 dialogues shorter than 30 turns (12K tokens) and benefit more from inserted irrelevant dialogues.
 867 However, in longer dialogues (e.g., 50 turns, PSM-based models show a sharper decline, falling
 868 below LoRA-based models. This pattern may be due to the fact that PSM relies on the aggregation
 869 of final-layer vectors to incorporate the persona information into an LLM, which works well when
 870 the inserted irrelevant dialogue is short, but may be negatively affected by accumulated noise when
 871 the irrelevant dialogue is long. In contrast, LoRA’s low-rank adaptation maintains greater stability
 872 in extended contexts.

873 **A.7.3 PERFORMANCE ON ENGLISH VERSION**

875 To gain a comprehensive understanding, we conduct more experiments on the English version of
 876 ChildEval. Figure 18 presents the evaluation results across different numbers of inserted irrelevant
 877 dialogue turns. These dynamics indicate that tasks and models exhibit varying levels of robustness
 878 and adaptability across different dialogue stages in both the explicit and implicit datasets. IS remains
 879 the most challenging dimension for all models. However, the overall performance on the English
 880 dataset is slightly lower, likely because it is a translated counterpart of the Chinese corpus and may
 881 not fully capture the natural distribution of native English dialogues.

882 Figure 18 presents the evaluation results after incorporating the persona information into the prompt,
 883 and different models exhibit divergent patterns. Notably, LLaMA3.1-8B-instruct shows substantial
 884 fluctuations on the EA, IS, and EG dimensions of the child-oriented evaluation. The performances
 885 of the other two models show a decreasing trend with small fluctuations as the number of irrelevant
 886 dialogue turns increases. Comparing Figure 18 and Figure 19, the inclusion of persona leads to
 887 significant improvements across PC, EA, IS, and EG for all models, with the only exception being a
 888 slight decrease on DA observed for Minstral-7B-instruct.

889
 890
 891
 892
 893
 894
 895
 896
 897
 898
 899
 900
 901
 902
 903
 904
 905
 906
 907
 908
 909
 910
 911
 912
 913
 914
 915
 916
 917

918

919

920

921

922

923

You need to generate the following content:

924

Preference: The user (3-6-year-old children) clearly expresses a specific and unique like, ability, or dislike in the first person (e.g., "I like xx more than xx," "I really hate xx," "I only care about xx," "I cannot xx," etc.). This preference or ability should be clear and distinctive enough for the intelligent assistant to remember. It should be concise and unique, summarized in 1-2 sentences.

925

The preference should consider diversity from different aspects of children.

926

Utterance: The user (3-6-year-old children) initiates the conversation or question using the first-person expressions "I" or "myself." The wording of the question or request should be careful to avoid contradicting or revealing the declared preference.

927

The dialogue should be naturally aligned with the child's personality and make it difficult for an intelligent assistant to give a satisfying answer if the preference is unknown, but it must not conflict with the preference.

928

Brief explanation (1-2 sentences): Explain why a conventional answer might violate the child's preference and how the intelligent assistant should respond or make suggestions based on the child's preference.

929

Scoring criteria:

930

Generate preference-utterance pairs with a high probability of violation:

931

High violation probability means:

932

$P(\text{answer} \mid \text{utterance}) >> P(\text{answer} \mid \text{preference}, \text{utterance})$ — i.e., without knowing the preference, conventional responses are very likely to violate the child's preference.

933

High violation probability example: <HighViolationExample>

934

Low violation probability example: <LowViolationExample>

935

Additional high violation probability examples: <Examples>

936

Do not generate:

937

Contradictory or too obvious combinations (utterance directly negates the preference, or perfectly matches it).

938

Utterance completely incompatible with the preference, or answers too simple/direct.

939

Preference or utterance lacking key information (like location or specific details).

940

Key points:

941

"Preference setting": starts with "I," written in 3-6-year-old style, short sentences with particles like "la," "ne," "ya," "ma," etc., avoiding complex words.

942

"Utterance": initiated by the child, natural and non-contradictory with the preference.

943

Utterance and preference must be strongly related and diverse.

944

Utterance and theme must be strongly correlated, not multi-theme ambiguous.

945

Child's dialogue style: oral style with particles, simple vocabulary, avoids adult-like wording.

946

Dialogue is strictly between child and assistant. Mentions of parents allowed as indirect statements only.

947

Answer strategy:

948

If the child's preference is unknown, the assistant's answers are likely to trigger the aversion objects; if the preference is known, the assistant should adjust responses to avoid violating it.

949

Based on the following child persona and topic, generate 2 different realistic scenarios with high violation probability (realistic, innovative, challenging):

950

Child persona: {persona} Topic: {topic}

951

Do not number; generate content directly using the following format:

<task>

<preference>...</preference>

<utterance>...</utterance>

<explanation>...</explanation>

</task>

952

953

954

955

956

957

958

959

960

961

962

963

964

965

966

967

968

969

970

971

Figure 7: Prompt used for generating explicit preference and utterance.

972

973

974

975

976 Please generate an $\{n\}$ -turn dialogue between a child and an intelligent assistant based on the child's persona and
 977 explicit preference.

978 Input:

979 Persona: Based on the basic information and long-term stable preference traits of children (3-6 years old).

980 Explicit Preference: For the given persona, the user clearly expresses a specific and unique like, ability, or dislike
 981 in the first person (e.g., "I like xx more than xx," "I really hate xx," "I only care about xx," "I cannot xx," etc.).

982 Topic: The theme around which the dialogue is built.

983 Output:

984 1. An analysis of the "forgetting-prevention self-check," following the required checking order (written inside
 985 <explain> tags).

986 2. An $\{n\}$ -turn dialogue between a 3-6-year-old child and the intelligent assistant (written inside
 987 <conversations> tags).

988 Forgetting-prevention self-check requirements (must be checked in this order and written in <explain> tags):

989 1. Whether names were mistakenly added: remove all specific personal names.

990 2. Whether the last turn includes: remove all closing phrases or polite endings.

991 3. Whether the dialogue addresses a child user: limit filler words appropriately.

992 4. Whether the intelligent assistant is described with human actions: the assistant can only provide suggestions.

993 5. Whether the dialogue is exactly $\{n\}$ turns: if fewer than $\{n\}$, extend the topic (through questions or additional
 994 information).

995 6. Whether the generation format tags are complete: check that all tags are correctly closed.

996 7. Whether the dialogue allows the explicit preference {preference} to be inferred naturally.

997 Multi-turn dialogue requirements (written inside <conversations> tags):

998 Strictly follow the rules below. Before each response, re-check compliance.

999 1. The dialogue must revolve around the theme, match the persona, and align with the speaking style of
 1000 3-6-year-old children:

1001 - Oral style, frequently using particles like "la," "ne," "ya," "ma," etc., to show a child's identity. For example:
 1002 "I don't like noisy ne" instead of the complex adult expression "I don't like noisy and chaotic environments."

1003 - Simple vocabulary (avoid complex words such as "recommend," "suggest"). Do not use adult-style
 1004 expressions like "Do you have any good food suggestions?" Instead, use child-style wording such as "What
 1005 yummy things are there? I want to eat yummy food!"

1006 2. Use concise, friendly, conversational expressions and avoid mechanical tone.

1007 3. The dialogue must not explicitly mention the input's explicit preference, but the child–assistant conversation
 1008 should make the preference inferable.

1009 4. The dialogue is strictly between the child and the intelligent assistant, following these rules:

1010 - Objective mentions are allowed: e.g., "Dad said..." "Mom said...", but the child cannot speak directly to
 1011 parents (e.g., "Dad, let's go play").

1012 - Interaction restriction: the child can only talk to the assistant (using "you" to refer to the assistant).

1013 - Direct conversation with parents or third parties is prohibited (e.g., "Mom, we...").

1014 - Scene restriction: if family activities are mentioned, they must be expressed indirectly (e.g., "Dad said we can
 1015 go to the park") instead of directly addressing parents.

1016 5. The dialogue must have exactly $\{n\}$ turns, where 1 turn = 1 <user> + 1 <assistant>. $\{n\}$ turns = $\{n\}$ <user> and
 1017 $\{n\}$ <assistant>.

1018 6. No specific personal names (like "Xiao An") or role names (like "little assistant," "smart helper") should appear.
 1019 <user> and <assistant> already indicate roles, no repetition needed.

1020

1021

1022

1023

1024 Figure 8: Prompt used for generating child–LLM dialogue to infer the implicit preference: Part 1 –
 1025 Inputs.

1026
 1027
 1028
 1029
 1030
 1031
 1032
 1033 7. The assistant's responses must not include human behaviors (e.g., attending activities, eating, walking). The
 1034 assistant must always remain non-embodied, only providing content.
 1035 8. The last turn of the assistant's reply must not contain a closing phrase (e.g., "Goodbye," "Ask me anytime").
 1036 The ending should feel naturally continuous.
 1037 Output must strictly follow the fixed format below, without modifying tag names, order, or nesting.
 1038 <explain>
 1039 [1] Name check: No personal names used, compliant.
 1040 [2] Closing phrase check: No closing phrase in the last turn, compliant.
 1041 [3] Tone check: Language is mild and natural, matching the style of 3-6-year-old children.
 1042 [4] Assistant behavior check: Assistant is not personified and contains no self-involvement in activities.
 1043 [5] Turn count check: Exactly {n} turns (i.e., {n} <user> and {n} <assistant>).
 1044 [6] Tag check: All tags spelled correctly and fully closed.
 1045 [7] Preference inference check: From the dialogue, the child's attitude toward "xxx" can naturally reveal the
 1046 explicit preference.
 1047 </explain>
 1048 <conversations>
 1049 <!-- Turn 1 -->
 1050 <user>...</user>
 1051 <assistant>...</assistant>
 1052 ...
 1053 <!-- Turn {n} -->
 1054 <user>...</user>
 1055 <assistant>...</assistant>
 1056 </conversations>
 1057 Given Input:
 1058 Persona: {persona}
 1059 Explicit Preference: {preference}
 1060 Topic: {topic}
 1061 Non-compliant examples:
 1062 <assistant>Xiao An can try chips</assistant> (contains personal name)
 1063 <assistant>Goodbye</assistant> (closing phrase in last turn)
 1064 <user>Hello, little helper! I recently learned lots of new words ne!</user> (contains role name)
 1065 <assistant>I've been on a picnic, ya, so fun! On a picnic you can eat yummy food and play on the
 1066 grass.</assistant> (assistant described as doing activities)
 1067 <assistant>Not a problem! I'm happy to help. If you have other questions or topics, feel free to tell me! (tags
 1068 incomplete, missing closing)</assistant>
 1069 Example Output:
 1070 <Examples>
 1071
 1072
 1073
 1074 Figure 9: Prompt used for generating child-LLM dialogue to infer the implicit preference: Part 2 –
 1075 Self-check and Outputs.
 1076
 1077
 1078
 1079

1080
 1081
 1082
 1083
 1084
 1085
 1086
 1087
 1088
 1089
 1090
 1091
 1092
 1093 You will analyze a conversation between a child and an assistant, focusing on whether the assistant shows
 1094 Emotional Adaptation in answering the child's message.
 1095 Evaluate the response based on these stringent criteria:
 1096 Check if the assistant explicitly or implicitly acknowledges the child's emotional state (happy, sad, frustrated,
 1097 excited, scared, etc.) in its answer.
 1098 The accuracy of the guessed emotion does not matter; only whether the assistant makes an effort to adapt
 1099 emotionally.
 1100 Answer "Yes" if:
 1101 1. The response explicitly refers to the child's emotion. Examples include: "I can see you are feeling sad, let
 1102 me cheer you up with a story.>"; "Since you are excited about dinosaurs, let's play a dinosaur game!"; "You seem
 1103 worried, don't worry, I will stay with you."
 1104 2. The response implicitly adapts to the child's emotion by mirroring or matching tone, even without naming
 1105 it. Example: Child says "I'm scared of the dark." Assistant replies: "It's okay, I'll be your flashlight friend so you
 1106 don't feel alone."
 1107 Answer "No" if the response does not show any emotional adaptation. For instance, if the child says "I'm sad" and
 1108 the assistant just gives a factual answer or irrelevant reply like "Okay, here's a list of animals."
 1109 Here is the assistant's response:
 1110 The user's question is: {question}
 1111 Assistant response: {generation}
 1112 Examine the response meticulously and answer. Please answer in this exact XML format without any additional
 1113 text:
 1114 <explanation>[1 short sentence explanation]</explanation>
 1115 <answer>[Yes/No]</answer>
 1116
 1117
 1118
 1119
 1120
 1121 Figure 10: Evaluation prompt used for checking Emotional Adaptation.
 1122
 1123
 1124
 1125
 1126
 1127
 1128
 1129
 1130
 1131
 1132
 1133

1134
 1135
 1136
 1137
 1138
 1139
 1140
 1141
 1142
 1143
 1144
 1145
 1146
 1147

You will analyze a conversation between a child and an assistant, focusing on whether the assistant shows Interaction Scaffolding in its response.

Evaluate the response based on these stringent criteria:

Check if the assistant makes an effort to guide the child's participation (e.g., asking follow-up questions, breaking down complex tasks, offering hints, or encouraging elaboration).

The correctness of the content is not important; only whether scaffolding is present.

Answer "Yes" if:

1. The assistant explicitly encourages the child to take part. Examples include: "Can you think of another animal that lives in the ocean?"; "Let's try this step by step: first, can you name the colors you see?"; "Do you want to hear a harder riddle or an easier one?"

2. The assistant implicitly scaffolds the interaction by providing structured choices or gradual hints instead of just giving a direct answer. Example: Child asks "Why is the sky blue?" Assistant replies: "That's a great question! Do you remember what happens when light passes through water? Something similar happens in the air."

Answer "No" if the assistant only provides an answer or information without offering any interactive guidance. For instance, if the child says "Tell me about dinosaurs" and the assistant only replies "Dinosaurs lived millions of years ago," without further prompting or scaffolding.

Here is the assistant's response:

The user's question is: {question}

Assistant response: {generation}

Examine the response meticulously and answer. Please answer in this exact XML format without any additional text:

<explanation>[1 short sentence explanation]</explanation>

<answer>[Yes/No]</answer>

1174
 1175
 1176
 1177
 1178
 1179
 1180
 1181
 1182
 1183
 1184
 1185
 1186
 1187

Figure 11: Evaluation prompt used for checking Interaction Scaffolding.

1188
 1189
 1190
 1191
 1192
 1193
 1194
 1195
 1196
 1197
 1198
 1199
 1200
 1201 You will analyze a conversation between a child and an assistant, focusing on whether the assistant's response is
 1202 Developmentally Appropriate for the child.
 1203
 1204 Evaluate the response based on these stringent criteria:
 1205 Check if the assistant's response matches the child's likely age level in terms of language, complexity, and content.
 1206 The scientific correctness of the content is not the main concern; only whether the response is suitable for a young
 1207 child's developmental stage.
 1208 Answer "Yes" if:
 1209 1. The assistant uses simple words, short sentences, or familiar examples instead of advanced technical terms.
 1210 Examples include: "The sun is like a big lamp in the sky that keeps us warm.;" "A volcano is like a mountain that
 1211 can burp hot lava.;" "Let's count together how many stars we can imagine."
 1212 2. The assistant adjusts explanations or provides analogies that fit a child's world. Example: Child asks:
 1213 "What is electricity?" Assistant replies: "It's like invisible energy that makes your toys and lights work when you
 1214 plug them in."
 1215 Answer "No" if the response uses adult-level, abstract, or overly complex explanations that a child is unlikely to
 1216 understand. For example, if the child asks "Why is the sky blue?" and the assistant replies with "Due to Rayleigh
 1217 scattering of shorter wavelengths of light in the atmosphere," this would be developmentally inappropriate.
 1218
 1219 Here is the assistant's response:
 1220
 1221 The user's question is: {question}
 1222 Assistant response: {generation}
 1223
 1224 Examine the response meticulously and answer. Please answer in this exact XML format without any additional
 1225 text:
 1226 <explanation>[1 short sentence explanation]</explanation>
 1227 <answer>[Yes/No]</answer>

Figure 12: Evaluation prompt used for checking Developmental Appropriateness.

1242
 1243
 1244
 1245
 1246
 1247
 1248
 1249
 1250
 1251
 1252
 1253
 1254
 1255
 1256
 1257
 1258
 1259
 1260
 1261
 1262
 1263
 1264
 1265
 1266
 1267
 1268
 1269
 1270
 1271
 1272

You will analyze a conversation between a child and an assistant, focusing on whether the assistant shows Engagement in its response.

Evaluate the response based on these stringent criteria:

Check if the assistant's response makes the conversation lively, fun, or interesting for the child.

Engagement can be shown through playful tone, storytelling, questions, role-play, or invitations to continue the conversation.

Answer "Yes" if:

1. The assistant explicitly uses playful or inviting language to keep the child engaged. Examples include: "Wow, that's a great question! Do you want to imagine we are astronauts and fly to space together?"; "Haha, dinosaurs are awesome! Which one do you like best?"; "Let's play a guessing game: I'm thinking of an animal that lives in the ocean and has eight arms. Can you guess what it is?"

2. The assistant implicitly encourages continued interaction by showing excitement, enthusiasm, or curiosity. Example: Child: "I like cats." Assistant: "Me too! Cats are so soft and playful. Do you have a favorite color for a cat?"

Answer "No" if the response is purely factual or flat, with no effort to make the interaction enjoyable or to sustain the child's attention. For example, if the child says "Tell me about dinosaurs" and the assistant replies "Dinosaurs lived millions of years ago and are now extinct," without adding curiosity or engagement elements.

Here is the assistant's response:

The user's question is: {question}

Assistant response: {generation}

Examine the response meticulously and answer. Please answer in this exact XML format without any additional text:

<explanation>[1 short sentence explanation]</explanation>

<answer>[Yes/No]</answer>

Figure 13: Evaluation prompt used for checking Engagement.




Figure 14: The architecture of the persona steer model.

1293
 1294
 1295

1296

Preference consistency errors on the explicit dataset

1297

1298

1299

1300

1301

1302

1303

1304

1305

1306

1307

1308

1309

1310

1311

1312

1313

1314

1315

1316

1317

1318

1319

1320

1321

1322

1323

1324

1325

1326

1327

1328

1329

1330

1331

1332

1333

1334

1335

1336

1337

1338

1339

1340

1341

1342

1343

1344

1345

1346

1347

1348

1349




Figure 15: Preference consistency error types under different numbers of inserted irrelevant turns (n-turn).






Figure 16: Preference consistency error types under different numbers of inserted irrelevant turns (n-turn).



1374 Figure 17: Accuracy of LLMs on preference consistency (PC) and child-oriented dimensions under
1375 different numbers of inserted irrelevant turns (n-turn).



1397 Figure 18: LLMs performances on preference consistency (PC) and the child-oriented evaluation
1398 under different numbers of inserted irrelevant dialogue turns on the English dataset.




Figure 19: LLMs performances on preference consistency (PC) and the child-oriented evaluation under different numbers of inserted irrelevant dialogue turns on the English dataset, after integrating persona information into the prompt.