

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 DECOUPLING TASK-SOLVING AND OUTPUT FORMATTING IN LLM GENERATION

Anonymous authors

Paper under double-blind review

ABSTRACT


Large language models (LLMs) are increasingly adept at following instructions containing task descriptions to solve complex problems, such as mathematical reasoning and automatic evaluation (LLM-as-a-Judge). However, as prompts grow more complex, models often struggle to adhere to all instructions. This difficulty is especially common when instructive prompts intertwine reasoning directives—specifying what the model should solve—with rigid formatting requirements that dictate how the solution must be presented. The entanglement creates competing goals for the model, suggesting that more explicit separation of these two aspects could lead to improved performance. To this front, we introduce DECO-G, a decoding framework that explicitly decouples format adherence from task solving. DECO-G handles format compliance with a separate tractable probabilistic model (TPM), while prompts LLMs with only task instructions. At each decoding step, DECO-G combines next token probabilities from the LLM with the TPM calculated format compliance likelihood to form the output probability. To make this approach both practical and scalable for modern instruction-tuned LLMs, we introduce three key innovations: instruction-aware distillation, a flexible trie-building algorithm, and HMM state pruning for computational efficiency. We demonstrate the effectiveness of DECO-G across a wide range of tasks with diverse format requirements, including mathematical reasoning, LLM-as-a-judge, and event argument extraction. Overall, our approach yields 1.0% to 6.0% relative gain over regular prompting practice with guaranteed format compliance.

Mathematical Reasoning (GSM8k)


Question: Lee runs the 400-meter hurdles in 38 seconds ... what is Gerald's new time?

Format Constraints: The final answer is ...

Prompting w/o Format Constraints



Prompting w/ Format Constraints



Deco-G




Figure 1: Example of GSM8k responses. LLM prompted without format constraints gets the correct answer, but the number is embedded in a sentence with mixed types, making it hard to capture. LLM prompted with format constraints gets the answer wrong. DECO-G prompts the model with task information and handles the format constraints by employing a Format Estimation Module (FEM). The framework generates the correct answer in the required format, making it easy to harvest.

054 **1 INTRODUCTION**

055

056 Instruction fine-tuning (Wei et al., 2021; Chung et al., 2024) enables large language models (LLMs)
 057 to follow user instructions and solve complex tasks. Given a task description and a desired output
 058 format, LLMs can perform tasks such as model evaluation and event extraction without additional
 059 training. Prompting strategies like Chain-of-Thought (Wei et al., 2022) and Tree-of-Thought (Yao
 060 et al., 2023) further enhance performance by encouraging structured reasoning. However, emerging
 061 evidence suggests that complex instructions—especially those with strict formatting—can nega-
 062 tively impact model performance (Tam et al., 2024; Long et al., 2025; He et al., 2024). For example,
 063 Long et al. (2025) reveal that LLM performance varies based on the required output format. On the
 064 MMLU (Hendrycks et al., 2020) benchmark, a model might fail to provide the correct answer when
 065 forced into one output structure, or provide the correct answer but fail a minor formatting instruction,
 066 which complicates automatic evaluation. In addition, Tam et al. (2024) point out that stricter format
 067 constraints generally lead to a greater degradation in performance on reasoning tasks. Therefore,
 068 the current paradigm of stacking task instructions and format instructions in the input prompt (as
 069 shown in Figure 1, the attachment of format instruction enclosed in the red hexagon) appears to be
 070 a limiting factor for harnessing LLM capabilities.

071 Attempts have been made to reduce format constraints’ impact on LLM generation. For in-
 072 stance, Tam et al. (2024) employ a less strict format to give LLM more flexibility. Long et al. (2025)
 073 and He et al. (2024) explore formats that are more intuitive for the LLM to follow. Yet, they still
 074 pose certain constraints to the LLM, impairing its reasoning skills. Other works (Beurer-Kellner
 075 et al., 2024; guidance-ai, 2024; Willard & Louf, 2023) perform non-neural inference-time control
 076 centered toward constraint satisfaction. They guarantee format compliance by enforcing the model
 077 to decode certain tokens. This mechanism fails to consider the interplay with LLM reasoning, often
 078 resulting in incoherent output. The situation thus highlights the need for a framework that seamlessly
 079 decouples format constraints from LLM task solving to unlock the full potential of LLMs.

080 In this paper, we propose a decoupled generation framework DECO-G that separates output
 081 formatting from task reasoning, thereby allowing the LLM to focus on the task without the burden
 082 of format adherence. We leverage the modularity of existing controllable text generation methods
 083 (e.g. GeLaTo (Zhang et al., 2023), CtrlG (Zhang et al., 2024)) and delegate the format adherence
 084 responsibility to an auxiliary Tractable Probabilistic Model (TPM), which estimates compliance
 085 rate and reweights token probability. While GeLaTo and CtrlG provide pathways for controllable
 086 generation with keyphrase and length constraints, they face significant challenges when applied to
 087 instruction-tuned LLMs with complex output templates. These challenges stem from a domain shift
 088 and computational bottlenecks that hinder scalability and efficiency. To make our framework prac-
 089 tical and effective for general instructive tasks, we introduce three key techniques. Firstly, we train
 090 the HMM on LLM’s instruction-response pairs to better captures task-oriented behaviors. Secondly,
 091 we employ a flexible trie-based algorithm to efficiently construct automata for complex, multi-part
 092 output templates. Thirdly, we implement HMM hidden state pruning to accelerate inference speed
 093 and ensure practical usability. To our knowledge, we are the **first** to propose the direct separation
 094 of task solving and format adherence in LLM generation to preserve its full potential.

095 To assess DECO-G’s effectiveness in handling tasks of different natures, including reasoning and
 096 multi-phrase templates adherence, we test the framework on three different tasks: mathematical rea-
 097 soning, LLM-as-a-judge evaluation, and generative event argument extraction. Experiment results
 098 show that DECO-G is able to improve overall task performance through multiple aspects: 1) impro-
 099 ving the format satisfaction rate, 2) encouraging more natural and flexible integration of format in the
 100 output, and 3) allowing LLM to concentrate on task solving without the burden of format following.
 101 Our contributions are as follow:¹

- 102 • We propose a framework to separate format compliance from task-solving to enhance overall
 103 performance of LLMs on various tasks through the use of a tractable probabilistic model.
- 104 • Our framework achieves high efficiency and effectiveness through technical innovations, includ-
 105 ing instruction-aware distillation, a flexible trie-building algorithm, and HMM state pruning.
- 106 • We secure improved task performance on multiple tasks compared to baseline methods, observ-
 107 ing relative gains ranging from 1.0% to 6.0% , and provide an analysis of DECO-G’s steering
 108 process with insights into its parameter setup from an entropy perspective.

109 ¹Code and model weights will be released upon paper acceptance.

108 **2 PRELIMINARIES**

109

110 In this section, we present our goal of task-format decoupling within a probabilistic formulation
 111 of language model generation. We further discuss how prior controllable generation methods align
 112 with this objective and provide a strong foundation for our approach.

113

114 **2.1 GENERATION WITH ATTRIBUTE CONTROL**

115

116 We frame the problem of controllable text generation using a probabilistic formulation. The auto-
 117 regressive generation of a token sequence $x_{1:n}$ given a desired attribute α can be expressed as:

$$P(x_{1:n}|\alpha) = \prod_t P(x_t|x_{<t}, \alpha) \quad (1)$$

120 The objective is to generate a sequence $x_{1:n}$ that exhibits the attribute α . At each generation step
 121 t , the target distribution for producing text with the desired attribute is $P(x_t|x_{<t}, \alpha)$. Using Bayes'
 122 rule, we can rewrite this as:

123

$$P(x_t|x_{<t}, \alpha) = P_{\text{LM}}(x_t|x_{<t}) \frac{P_{\text{LM}}(\alpha|x_t, x_{<t})}{P_{\text{LM}}(\alpha|x_{<t})} \quad (2)$$

126 Here, the first term, $P_{\text{LM}}(x_t|x_{<t})$, is the language model's next-token probability, which is respon-
 127 sible for generating fluent and coherent content. The second term, the ratio $\frac{P_{\text{LM}}(\alpha|x_t, x_{<t})}{P_{\text{LM}}(\alpha|x_{<t})}$, acts as a
 128 control signal. It quantifies how the choice of the current token x_t influences the probability that the
 129 final, complete sequence will satisfy the attribute α . However, directly calculating this ratio is in-
 130 tractable, as it requires marginalizing over all possible future sequences to compute the likelihoods.
 131 Thus, a key challenge in controllable generation is to find a tractable approximation for this term.

132

133 **2.2 ESTIMATING LIKELIHOOD OF ATTRIBUTE SATISFACTION**

134

135 Recent controllable generation frameworks such as GeLaTo (Zhang et al., 2023) and Ctrl-G (Zhang
 136 et al., 2024) leverage a tractable probabilistic model (TPM) to efficiently estimate the marginal
 137 probability $P(\alpha|x_t, x_{<t})$, serving as a signal to steer an LLM's generation, following

$$P(x_t|x_{<t}, \alpha) \propto P_{\text{LM}}(x_t|x_{<t}) P_{\text{TPM}}(\alpha|x_t, x_{<t}) \quad (3)$$

138 These approaches first distill a Hidden Markov Model (HMM) as a probabilistic approximation of
 139 the LLM and then encode logical constraints to formal structure that the HMM can reason over.

140

141 **Sequence modeling with Hidden Markov Models.** A Hidden Markov Model (HMM) is the spe-
 142 cific type of TPM used in these frameworks, chosen for its ability to model sequential data tractably.
 143 The joint probability distribution over a sequence of observed variables (tokens, $x_{1:n}$) and a corre-
 144 sponding sequence of hidden state variables $z_{1:n}$, is modeled as

145

$$P_{\text{HMM}}(x_{\leq t}, z_{\leq t}) = P(z_1) P(x_1|z_1) \prod_{t=2}^T P(z_t|z_{t-1}) P(x_t|z_t) \quad (4)$$

146 Critically, the Markov property of HMMs enables efficient probabilistic inference over all possible
 147 future sequences, a task that is intractable for language models. In frameworks like GeLaTo and Ctrl-
 148 G, the HMM is distilled from the LLM using samples drawn unconditionally from the LLM. This
 149 process involves training the HMM via maximum likelihood on the sampled completions, equivalent
 150 to minimizing the KL-divergence between the two models' distributions $D_{\text{KL}}(P_{\text{LM}}||P_{\text{HMM}})$.

151

152 **Formalizing Constraints with Deterministic Finite Automata.** To enforce a constraint using
 153 the HMM, the constraints must be expressed in a formal language. Zhang et al. (2024) propose
 154 representing logical constraints as Deterministic Finite Automata (DFA). A DFA is an abstract state
 155 machine that recognizes patterns in sequences. Formally, a DFA is a 5-tuple $\mathcal{D} = (Q, \Sigma, \delta, q_0, F)$,
 156 where Q is a finite set of states, Σ is the alphabet (the LLM's token vocabulary), $\delta : Q \times \Sigma \rightarrow Q$
 157 is the transition function, $q_0 \in Q$ is the initial state, and $F \subseteq Q$ is the set of accept states. A
 158 sequence is “accepted” if it drives the machine from its initial state to an accept state; otherwise, it
 159 is “rejected.” This formalism is capable of representing logical constraints including the presence of
 160 keyphrases and word counts by defining the appropriate states and transitions.

161




Figure 2: DECO-G decouples task and format—prompting LLM with task-only information and sending format constraints to FEM. DECO-G decodes from the posterior constructed by multiplying LLM token probabilities with FEM estimated satisfaction rate.

Probabilistic reasoning over logical constraints. The core idea of these prior frameworks is to use the TPM to perform a probabilistic lookahead—that is, to efficiently compute $P_{\text{TPM}}(\alpha|x_t, x_{<t})$, the probability that the full generated sequence will satisfy the constraint α . This is accomplished by marginalizing the joint HMM-DFA state space over all possible future sequences that reach an accepting state in the DFA. According to Zhang et al. (2024), this marginalization can be calculated efficiently using a backward recurrence relation. Refer to Section B for detailed derivation.

2.3 FROM PRIOR WORK TO DECO-G

Prior frameworks like GeLaTo (Zhang et al., 2023) and Ctrl-G (Zhang et al., 2024) successfully use Hidden Markov Models (HMM) as tractable generative models to guide LLM generation, ensuring outputs satisfy specific logical constraints in tasks such as keyphrase generation and text editing. While GeLaTo introduces this concept, its use of Conjunctive Normal Forms (CNF) is primarily limited to keyphrase constraints. Ctrl-G generalizes this approach by specifying logical constraints through Deterministic Finite Automata (DFA), which can represent constraints on bounded-length sequences. While this foundation is promising, significant challenges arise when adapting this framework to decouple format from task reasoning for modern, instruction-tuned LLMs.

- **Domain shift:** the paradigm shifts from logical-constrained generation to separating task-instructed generation into two sub-tasks: problem solving (LLM side) and format adherence (TPM side). This incurs domain mismatch, as prior methods train HMMs on random generation without context, which is a poor proxy for an LLM conditioned on specific task instructions.
- **Complexity of format templates:** The intricate nature of real-world format templates presents a major challenge, as the overhead from constructing complex constraint automata creates a computational bottleneck.
- **Inference-time inefficiency:** The large vocabulary size of modern LLMs introduces substantial computational overhead during the inference-time guidance step, which severely impedes the framework’s latency

3 DECO-G

In this section, we present DECO-G, a framework that realizes the decoupling of task reasoning from output formatting. As shown in Figure 2, our method separates the input prompt: the LLM receives only the task-specific information, while a dedicated Format Estimation Module (FEM) receives the format constraints. At each decoding step, the FEM estimates the likelihood of future compliance with the given format constraints α . This likelihood is then used to reweigh the LLM’s original token probabilities, steering the generation towards a format-compliant output. We now describe the key components that enable this framework.

3.1 INSTRUCTION-AWARE HMM DISTILLATION

An HMM can approximate a large language model’s (LLM) output distribution to guide controllable generation. The fidelity of this approximation is critical—ideally, an HMM that perfectly replicates the LLM’s probabilities would yield an exact posterior for format-decoupled generation, per Equa-

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
 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
 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
 864
 865
 866
 867
 868
 869
 870
 871
 872
 873
 874
 875
 876
 877
 878
 879
 880
 881
 882
 883
 884
 885
 886
 887
 888
 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
 924
 925
 926
 927
 928
 929
 930
 931
 932
 933
 934
 935
 936
 937
 938
 939
 940
 941
 942
 943
 944
 945
 946
 947
 948
 949
 950
 951
 952
 953
 954
 955
 956
 957
 958
 959
 960
 961
 962
 963
 964
 965
 966
 967
 968
 969
 970
 971
 972
 973
 974
 975
 976
 977
 978
 979
 980
 981
 982
 983
 984
 985
 986
 987
 988
 989
 990
 991
 992
 993
 994
 995
 996
 997
 998
 999
 1000
 1001
 1002
 1003
 1004
 1005
 1006
 1007
 1008
 1009
 1010
 1011
 1012
 1013
 1014
 1015
 1016
 1017
 1018
 1019
 1020
 1021
 1022
 1023
 1024
 1025
 1026
 1027
 1028
 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
 1060
 1061
 1062
 1063
 1064
 1065
 1066
 1067
 1068
 1069
 1070
 1071
 1072
 1073
 1074
 1075
 1076
 1077
 1078
 1079
 1080
 1081
 1082
 1083
 1084
 1085
 1086
 1087
 1088
 1089
 1090
 1091
 1092
 1093
 1094
 1095
 1096
 1097
 1098
 1099
 1100
 1101
 1102
 1103
 1104
 1105
 1106
 1107
 1108
 1109
 1110
 1111
 1112
 1113
 1114
 1115
 1116
 1117
 1118
 1119
 1120
 1121
 1122
 1123
 1124
 1125
 1126
 1127
 1128
 1129
 1130
 1131
 1132
 1133
 1134
 1135
 1136
 1137
 1138
 1139
 1140
 1141
 1142
 1143
 1144
 1145
 1146
 1147
 1148
 1149
 1150
 1151
 1152
 1153
 1154
 1155
 1156
 1157
 1158
 1159
 1160
 1161
 1162
 1163
 1164
 1165
 1166
 1167
 1168
 1169
 1170
 1171
 1172
 1173
 1174
 1175
 1176
 1177
 1178
 1179
 1180
 1181
 1182
 1183
 1184
 1185
 1186
 1187
 1188
 1189
 1190
 1191
 1192
 1193
 1194
 1195
 1196
 1197
 1198
 1199
 1200
 1201
 1202
 1203
 1204
 1205
 1206
 1207
 1208
 1209
 1210
 1211
 1212
 1213
 1214
 1215
 1216
 1217
 1218
 1219
 1220
 1221
 1222
 1223
 1224
 1225
 1226
 1227
 1228
 1229
 1230
 1231
 1232
 1233
 1234
 1235
 1236
 1237
 1238
 1239
 1240
 1241
 1242
 1243
 1244
 1245
 1246
 1247
 1248
 1249
 1250
 1251
 1252
 1253
 1254
 1255
 1256
 1257
 1258
 1259
 1260
 1261
 1262
 1263
 1264
 1265
 1266
 1267
 1268
 1269
 1270
 1271
 1272
 1273
 1274
 1275
 1276
 1277
 1278
 1279
 1280
 1281
 1282
 1283
 1284
 1285
 1286
 1287
 1288
 1289
 1290
 1291
 1292
 1293
 1294
 1295
 1296
 1297
 1298
 1299
 1300
 1301
 1302
 1303
 1304
 1305
 1306
 1307
 1308
 1309
 1310
 1311
 1312
 1313
 1314
 1315
 1316
 1317
 1318
 1319
 1320
 1321
 1322
 1323
 1324
 1325
 1326
 1327
 1328
 1329
 1330
 1331
 1332
 1333
 1334
 1335
 1336
 1337
 1338
 1339
 1340
 1341
 1342
 1343
 1344
 1345
 1346
 1347
 1348
 1349
 1350
 1351
 1352
 1353
 1354
 1355
 1356
 1357
 1358
 1359
 1360
 1361
 1362
 1363
 1364
 1365
 1366
 1367
 1368
 1369
 1370
 1371
 1372
 1373
 1374
 1375
 1376
 1377
 1378
 1379
 1380
 1381
 1382
 1383
 1384
 1385
 1386
 1387
 1388
 1389
 1390
 1391
 1392
 1393
 1394
 1395
 1396
 1397
 1398
 1399
 1400
 1401
 1402
 1403
 1404
 1405
 1406
 1407
 1408
 1409
 1410
 1411
 1412
 1413
 1414
 1415
 1416
 1417
 1418
 1419
 1420
 1421
 1422
 1423
 1424
 1425
 1426
 1427
 1428
 1429
 1430
 1431
 1432
 1433
 1434
 1435
 1436
 1437
 1438
 1439
 1440
 1441
 1442
 1443
 1444
 1445
 1446
 1447
 1448
 1449
 1450
 1451
 1452
 1453
 1454
 1455
 1456
 1457
 1458
 1459
 1460
 1461
 1462
 1463
 1464
 1465
 1466
 1467
 1468
 1469
 1470
 1471
 1472
 1473
 1474
 1475
 1476
 1477
 1478
 1479
 1480
 1481
 1482
 1483
 1484
 1485
 1486
 1487
 1488
 1489
 1490
 1491
 1492
 1493
 1494
 1495
 1496
 1497
 1498
 1499
 1500
 1501
 1502
 1503
 1504
 1505
 1506
 1507
 1508
 1509
 1510
 1511
 1512
 1513
 1514
 1515
 1516
 1517
 1518
 1519
 1520
 1521
 1522
 1523
 1524
 1525
 1526
 1527
 1528
 1529
 1530
 1531
 1532
 1533
 1534
 1535
 1536
 1537
 1538
 1539
 1540
 1541
 1542
 1543
 1544
 1545
 1546
 1547
 1548
 1549
 1550
 1551
 1552
 1553
 1554
 1555
 1556
 1557
 1558
 1559
 1560
 1561
 1562
 1563
 1564
 1565
 1566
 1567
 1568
 1569
 1570
 1571
 1572
 1573
 1574
 1575
 1576
 1577
 1578
 1579
 1580
 1581
 1582
 1583
 1584
 1585
 1586
 1587
 1588
 1589
 1590
 1591
 1592
 1593
 1594
 1595
 1596
 1597
 1598
 1599
 1600
 1601
 1602
 1603
 1604
 1605
 1606
 1607
 1608
 1609
 1610
 1611
 1612
 1613
 1614
 1615
 1616
 1617
 1618
 1619
 1620
 1621
 1622
 1623
 1624
 1625
 1626
 1627
 1628
 1629
 1630
 1631
 1632
 1633
 1634
 1635
 1636
 1637
 1638
 1639
 1640
 1641
 1642
 1643
 1644
 1645
 1646
 1647
 1648
 1649
 1650
 1651
 1652
 1653
 1654
 1655
 1656
 1657
 1658
 1659
 1660
 1661
 1662
 1663
 1664
 1665
 1666
 1667
 1668
 1669
 1670
 1671
 1672
 1673
 1674
 1675
 1676
 1677
 1678
 1679
 1680
 1681
 1682
 1683
 1684
 1685
 1686
 1687
 1688
 1689
 1690
 1691
 1692
 1693
 1694
 1695
 1696
 1697
 1698
 1699
 1700
 1701
 1702
 1703
 1704
 1705
 1706
 1707
 1708
 1709
 1710
 1711
 1712
 1713
 1714
 1715
 1716
 1717
 1718
 1719
 1720
 1721
 1722
 1723
 1724
 1725
 1726
 1727
 1728
 1

270

Table 1: GSM8k results.

271

Method	Format (%)	Acc. (%)
<i>Llama-3.1-8B-Instruct</i>		
NL	96.3	82.3
NL-S	100	81.3
JSON	64.7	51.8
JSON-S	100	75.7
DECO-G	100	85.2
<i>Qwen2.5-7B-Instruct</i>		
NL	98.0	83.6
NL-S	99.9	82.7
JSON	93.3	74.8
JSON-S	99.8	79.0
DECO-G ^{$\gamma=2$}	100	88.6
<i>Qwen3-8B</i>		
NL	97.4	90.5
NL-S	100	88.3
JSON	66.9	61.4
JSON-S	99.2	90.6
DECO-G ^{$\gamma=2$}	100	91.7

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

To mitigate this, we introduce **HMM hidden state pruning**, an optimization technique to reduce the computational load while preserving guidance quality. This technique is predicated on the observation that, at any given generation step, the probability mass of the hidden state distribution is concentrated within a small subset of states (see Section D). Consequently, rather than employing the full set of h states for the emission probability calculation, we prune the distribution by considering only the top- k most probable states. Our empirical validation demonstrates that selecting a minimal fraction of states—specifically, the top 5% ($k = 200$) based on their probability magnitudes—is sufficient to retain over 98% of the full model’s performance.

This pruning strategy drastically improves efficiency. The complexity of the emission step is reduced from $O(h|\mathcal{V}|)$ to $O(h \log h + k|\mathcal{V}|)$, where $k \ll h$. The $O(h \log h)$ term represents the cost associated with selecting the top- k states, while the dominant matrix multiplication is reduced to an $O(k|\mathcal{V}|)$ operation. This optimization achieves a considerable reduction in inference time for a negligible loss in performance, thereby enhancing the practical viability of DECO-G.

4 EXPERIMENT

Experimental Setup. We assess DECO-G’s overall performance over three tasks: (1) math problem solving with reasoning, (2) LLM-as-a-judge for summary evaluation, and (3) event argument extraction as a generative task. We apply DECO-G on performant instruction models *Llama-3.1-8B-Instruct* (Grattafiori et al., 2024), *Qwen2.5-7B-Instruct* (Yang et al., 2025b), and *Qwen3-8B* (Yang et al., 2025a) to verify its effectiveness. The baselines we include for comparison are as follows,

- **NL**: prompt LLM with task instruction and natural language output constraints, free generation
- **NL-S**: prompt LLM with task instruction and natural language output constraints, structured generation enforced through *Outlines* (Willard & Louf, 2023)
- **JSON**: prompt LLM with task instruction and JSON output constraints, free generation
- **JSON-S**: prompt LLM with task instruction and JSON output constraints, structured JSON generation enforced through *Outlines*

For DECO-G, the HMM for each LLM has hidden states of size $h=4096$, output space of $|\mathcal{V}|=128k$ for *Llama* and $|\mathcal{V}|=152k$ for *Qwen* models, and is trained for 100 epochs on one-million LLM generated responses (sampling takes 56 GPU hours and training takes 1 GPU hour on NVIDIA A100). For the following experiments, we adopt greedy decoding to ensure fair comparison with baseline methods and evaluate zero-shot performance.

Table 2: Generative EAE results on ACE05.

Method	AI	AC	AI+	AC+
<i>Llama-3.1-8B-Instruct</i>				
NL	36.8	27.3	34.8	25.5
NL-S	37.1	27.8	35.1	26.0
JSON	35.2	26.2	33.4	24.6
JSON-S	33.6	25.2	31.6	23.6
DECO-G	39.4	28.7	37.0	26.8
<i>Qwen2.5-7B-Instruct</i>				
NL	32.6	25.5	31.2	24.4
NL-S	33.2	24.9	31.1	23.7
JSON	31.9	24.1	30.5	22.9
JSON-S	34.1	26.1	32.5	24.7
DECO-G ^{$\gamma=2$}	35.2	25.9	33.4	24.5
<i>Qwen3-8B</i>				
NL	33.2	24.6	31.5	23.1
NL-S	33.3	24.2	31.5	22.6
JSON	31.7	23.4	3.01	21.8
JSON-S	32.5	23.0	30.8	21.5
DECO-G ^{$\gamma=2$}	34.0	24.6	32.1	23.1

324 4.1 MATHEMATICAL REASONING
325

326 In this task, we evaluate our framework on GSM8k (Cobbe et al., 2021), a collection of grade school
327 math problems that take two to eight steps to solve. Models are expected to carry out step-by-step
328 reasoning and arrive at the answer. Following Tam et al. (2024), a group of task instructions is
329 adopted to prompt the model to first reason about the math problem and then yield an integer as
330 its answer. For JSON format output, we prompt the model to output valid JSON blob with keys
331 “*reason*” and “*answer*.*”* For natural language output, format instructions are used to encourage the
332 model to generate the template phrase “*The final answer is ...*” Meanwhile, this phrase is specified
333 as a key phrase to appear in DECO-G’s generation.

334 **Evaluation Metrics.** We measure *Format Compliance* as the rate to which the generated answer
335 follows format requirement. In addition, we measure *Accuracy* as exact match of ground truth
336 answer.

337 **Results.** As shown in Table 1, unstructured NL generation offers decent performance, with *Llama*
338 scoring 82.3% and *Qwen* 83.6% on accuracy. However, together with unstructured JSON, free
339 generation methods completely rely on the LLM for following the format constraint and thus suffer
340 from low compliance rate. Structured generation, on the contrary, guarantees format compliance, but
341 its performance is negatively impacted by the invasive intervention that sometimes cut the generation
342 flow and alter course abruptly. DECO-G guarantees a 100% format compliance rate and achieves
343 the best performance over all three models. In practice, we find out that *Qwen* models have more
344 skewed token distribution. We thus raise the control factor λ to exert stronger control on the output.

345 4.2 LLM-AS-A-JUDGE EVALUATION
346

347 We then use LLMs as judges to evaluate the quality of summaries and assess how well it aligns
348 with human annotation. This evaluation is performed on the SummEval (Fabbri et al., 2021) dataset
349 which consists 1600 machine-generated summaries for 100 news articles, and human annotated
350 scores over four dimensions: Coherence, Consistency, Fluency, and Relevance. The models are
351 asked to analyze the summary and assign a score from 1 to 5 based on the given criteria suggested
352 by ChatGPT (OpenAI, 2025). We use the format “*The rating is ...*” for natural language output and
353 “*rating*” as the key for harnessing JSON output.

354 **Evaluation Metrics.** Following Liu et al. (2023), we adopt the summary-level Spearman and
355 Kendall-Tau correlation to gauge the performance of each method. Higher number indicates bet-
356 ter alignment with human annotated scores.

357 **Results.** For this task, *Qwen* models perform well in following the output format in unstructured
358 settings, securing over 99.7% compliance rate. This may attribute to a less intensive reasoning phase
359 compared to mathematical reasoning. As indicated in Table 3, DECO-G demonstrates the strongest
360 average correlation with human across models, enhancing over Consistency, Fluency, and Relevance
361 when applied to *Qwen* models. With *Llama*, DECO-G gains over Coherence and Relevance while
362 showing relative weakness in evaluating Consistency and Fluency. A close inspection of model
363 generated outputs suggests that DECO-G encourages a more flexible integration of the key phrase in
364 different places of the response: the beginning, middle, and end of response.

365 4.3 EVENT ARGUMENT EXTRACTION
366

367 The generative event argument extraction (EAE) task mainly assess a model’s ability in identifying
368 role-related arguments from source text. We evaluate on the ACE05-EN dataset (Doddington et al.,
369 2004), in which a model is be presented with an article, a trigger word, and a set of roles to determine
370 whether arguments associated with the roles are present in the article. This is naturally a templated
371 task as generative model has to specify which word is extracted for which role. Regarding JSON
372 output, we ask model to generate a JSON blob with roles as keys and extracted arguments as values.
373 For natural language output, we specify the template “*The <role_i> is ...*” for every relevant roles.
374 For DECO-G, we construct a flexible DFA that fuses the template phrases together with empty slots
375 allowing LLM predict arguments spanning from 1 to 5 tokens.

Table 3: SummEval results, measured over Coherence, Consistency, Fluency, and Relevance.

Method	Coherence		Consistency		Fluency		Relevance		Avg		
	ρ	τ	ρ	τ	ρ	τ	ρ	τ	Format	ρ	τ
<i>Llama-3.1-8B-Instruct</i>											
NL	0.381	0.311	0.383	0.351	0.321	0.291	0.405	0.337	95.8	0.372	0.322
NL-S	0.376	0.308	0.375	0.343	0.316	0.287	0.435	0.364	100	0.376	0.325
JSON	0.449	0.368	0.446	0.415	0.326	0.296	0.424	0.358	99.8	0.411	0.359
JSON-S	0.450	0.369	0.447	0.416	0.334	0.302	0.424	0.358	100	0.414	0.361
DECO-G	0.458	0.379	0.439	0.404	0.331	0.298	0.441	0.371	100	0.418	0.363
<i>Qwen2.5-7B-Instruct</i>											
NL	0.407	0.339	0.442	0.407	0.291	0.265	0.399	0.340	100	0.385	0.338
NL-S	0.403	0.337	0.448	0.412	0.279	0.254	0.408	0.347	100	0.384	0.338
JSON	0.411	0.334	0.488	0.455	0.305	0.280	0.383	0.326	99.7	0.396	0.349
JSON-S	0.412	0.335	0.489	0.457	0.309	0.284	0.387	0.330	100	0.399	0.351
DECO-G ^{$\gamma=2$}	0.327	0.271	0.506	0.470	0.348	0.311	0.452	0.380	100	0.408	0.358
<i>Qwen3-8B</i>											
NL	0.510	0.416	0.540	0.504	0.479	0.441	0.464	0.392	100	0.498	0.439
NL-S	0.507	0.413	0.544	0.509	0.477	0.439	0.468	0.396	100	0.499	0.439
JSON	0.504	0.409	0.491	0.459	0.444	0.410	0.450	0.382	99.8	0.472	0.415
JSON-S	0.486	0.393	0.486	0.454	0.406	0.376	0.441	0.374	100	0.455	0.399
DECO-G ^{$\gamma=2$}	0.490	0.395	0.546	0.516	0.499	0.456	0.494	0.414	100	0.507	0.445

Evaluation Metrics. We measure performance by calculating the *f1-score* comparing the extracted tuples and the ground truth tuples for the following categories:

- Argument Id (AI): argument span and event type.
- Argument Class (AC): argument span, event type, and role type.
- Argument-attached Id (AI+): argument span, event type, and event trigger.
- Argument-attached Class (AC+): argument span, event type, event trigger, and role type.

Results. The *f1-scores* reported in Table 2 suggest that EAE remains a challenging task for generative models. LLMs suffer from identifying correct relations in the article and presenting valid predictions that indeed exist in the original text—without modifying entity format or referring to exterior content. Baseline methods show inconsistent trends across models, indicating LLMs’ lack of robustness in event argument extraction. Employing DECO-G enhances overall extraction quality for *Llama* and *Qwen3*, while mainly improving over AI and AI+ for *Qwen2.5*. DECO-G’s gain on AI and AI+ is more evident than its improvement on AC and AC+, suggesting that DECO-G can further benefit from a tighter association between roles and extracted arguments—possible through designing more natural and intuitive control phrase for DECO-G.

5 ANALYSIS

5.1 THE STEERING PROCESS

DECO-G takes advantage of HMM to estimate the future format satisfaction rate and adjust token probabilities based on the estimation. To better understand this steering process, we examine the control signals produced by the FEM and visualize the control for a span of decoding step. We track the original LLM distribution, FEM distribution, and their composed distribution, which DECO-G decodes from. Figure 3 provides an illustration of DECO-G encouraging the generation of key phrase after step by step reasoning. While the LLM tends to conclude its response with “*The total number of ... is ...*” DECO-G assigns high probabilities to the token “*final*,” steering the LLM generation to conform with format constraints.

As LLMs are trained to provide clear and concise response, they tend to avoid repeating themselves when presenting the final answer. DECO-G captures this intricacy and replaces LLM's intended conclusive phrase with the format phrase "*The final answer is*" to reduce repetition. We consider

432 this format integration to be more natural than forcing LLM to generate certain phrases as in regex-
 433 structured generation.
 434

435 **5.2 TOKEN ENTROPY AND STEERING STRENGTH**
 436

437 In the previous section, we report DECO-G’s results with hyperparameter $\gamma = 2$ for controlling
 438 *Qwen* models, as $\gamma = 1$ doesn’t provide enough power to steer the model away from its own
 439 generation course. We hypothesize that *Qwen* models’ token distributions are more skewed than
 440 *Llama*’s, making it difficult for the control signal to actually make an impact on the distribution.
 441 To verify this, we draw 100 examples from GSM8k responses and measure the average step-wise
 442 entropy of LLM token distribution. As shown in Figure 4, *Llama*’s entropy is significantly higher
 443 than those of *Qwen2.5* and *Qwen3*, suggesting that *Llama*’s token probabilities are more spread out
 444 and diverse, whereas *Qwen* models’ token distributions are more peaky. This increased peakiness
 445 could be a consequence of the distribution squeezing induced by more intensive fine-tuning and
 446 preference optimization of the LLM (Ren & Sutherland, 2025). It is thus intuitive to amplify DECO-
 447 G’s control strength for LLM with more skewed distribution to guarantee format compliance.
 448

449 Within the same model, structured generation methods (NL-S and JSON-S) have slightly higher
 450 entropy than their unstructured counterparts (NL and JSON). This may attribute to imposed template
 451 tokens provoking more uncertainty in future token prediction. Meanwhile, DECO-G produces lowest
 452 LLM entropy, indicating that an absence of format constraint in task solving may lead to LLM
 453 providing the most confident response.
 454

455 **6 RELATED WORK**
 456

457 In the paper, we explore a controllable text generation (CTG) method to decouple task solving from
 458 format adherence. There are two branches in CTG that provide avenues for achieving this format-
 459 task decoupling—content-wise hard control and attribute-wise soft control.
 460

461 Content-wise structured generation aims to produce outputs that conform to predefined schemes or
 462 templates. The guaranteed adherence to specified format ensures high reliability when integrating
 463 LLM with external systems. This line of methods (Willard & Louf, 2023; guidance-ai, 2024), how-
 464 ever, exerts invasive control over the LLM generation which often produce abrupt cut-off, resulting
 465 in incomplete and incoherent responses.
 466

467 Attribute-wise soft control offers a more flexible paradigm, focusing on conditioning the generation
 468 based on a desired attribute. One line of works instills attribute information into the the LLM and
 469 updates model weights, through retraining (Keskar et al., 2019; Arora et al., 2022), fine-tuning (Wei
 470 et al., 2021; Zeldes et al., 2020; Li & Liang, 2021; Lester et al., 2021), or reinforcement learning
 471 (Ouyang et al., 2022; Stiennon et al., 2020; Zeng et al., 2024; Dai et al., 2024). This method
 472 benefits from no added computational load during inference, but the expense of training the LLM
 473 for updates can be significant. The other set of works (Dathathri et al., 2019; Yang & Klein, 2021;
 474 Krause et al., 2021; Schick et al., 2021; Liu et al., 2021; Khandelwal et al., 2021; Sridhak et al.,
 475 2022; Wen et al., 2023; Deng & Raffel, 2023) keeps the LLM as-is and instead modifies the genera-
 476 tion probabilities at inference time, also known as *weighted decoding*. These methods typically train
 477 a lightweight auxiliary model to guide the LLM’s generation at decoding time according to Bayes’
 478 rule. In light of these prior works, DECO-G takes the weighted decoding measure to compute the
 479 posterior given format constraints as an attribute.
 480

481 **7 CONCLUSION**
 482

483 In this paper, we present DECO-G, a novel decoding framework designed to decouple the respon-
 484 sibilities of task reasoning and format adherence. It achieves this responsibility separation by em-
 485 ploying an auxiliary Format Estimation Module to estimate future format satisfaction and modify
 486 token probabilities, thus allowing the LLM to concentrate solely on problem-solving. Experiments
 487 on mathematical reasoning, LLM-as-a-judge evaluation, and event argument extraction demon-
 488 strate this decoupling approach leads to overall performance gain, attributing to improved format com-
 489 pliance, more natural format integration, and more confident response from the LLM. Limitation of
 490 this work is covered in Section A.
 491

486 ETHICAL CONSIDERATIONS
487488 We conduct experiments on mathematical reasoning, LLM-as-a-judge evaluation, and event argu-
489 ment extraction. The score assigned by an LLM should not be considered an accurate reflection
490 of quality of the summary. In addition, the LLM responses to the GSM8k questions should not be
491 referenced for math instruction as they may include hallucination.492 We acknowledge the use of AI assistants for improving the manuscript’s prose, generating tables
493 in LaTeX format, figure design, and assisting with code implementation for the analysis of HMM
494 hidden state pruning. All generated content, particularly the data in tables, was manually verified
495 for accuracy against our experimental results.
496497 REPRODUCIBILITY STATEMENT
498500 A detailed description of our experimental setup, including the specific models used, HMM training
501 parameters, and decoding strategy, is provided in the introductory paragraph of Section 4. To allow
502 for replication of our experiments, the full prompts used for the mathematical reasoning (GSM8k),
503 LLM-as-a-judge (SummEval), and event argument extraction (ACE05) tasks are detailed in Sec-
504 tion F. Regarding computational overhead, a breakdown of the FLOPs required for the HMM for-
505 ward pass, both with and without pruning, is presented in Section E. Code and model weights will
506 be made publicly available upon paper acceptance.
507508 REFERENCES
509510 Kushal Arora, Kurt Shuster, Sainbayar Sukhbaatar, and Jason Weston. Director: Generator-
511 classifiers for supervised language modeling. In *Proceedings of the 2nd Conference of the Asia-
512 Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint
513 Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 512–526, 2022.514 Leonard E Baum et al. An inequality and associated maximization technique in statistical estimation
515 for probabilistic functions of markov processes. *Inequalities*, 3(1):1–8, 1972.516 Luca Beurer-Kellner, Marc Fischer, and Martin Vechev. Guiding LLMs the right way: Fast, non-
517 invasive constrained generation. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian
518 Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st
519 International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning
520 Research*, pp. 3658–3673. PMLR, 21–27 Jul 2024. URL [https://proceedings.mlr.
521 press/v235/beurer-kellner24a.html](https://proceedings.mlr.press/v235/beurer-kellner24a.html).
522523 Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li,
524 Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned lan-
525 guage models. *Journal of Machine Learning Research*, 25(70):1–53, 2024.526 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
527 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to
528 solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.529 Josef Dai, Xuehai Pan, Ruiyang Sun, Jiaming Ji, Xinbo Xu, Mickel Liu, Yizhou Wang, and Yaodong
530 Yang. Safe rlhf: Safe reinforcement learning from human feedback. In *The Twelfth International
531 Conference on Learning Representations*, 2024.532 Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosin-
533 ski, and Rosanne Liu. Plug and play language models: A simple approach to controlled text
534 generation. *arXiv preprint arXiv:1912.02164*, 2019.535 Haikang Deng and Colin Raffel. Reward-augmented decoding: Efficient controlled text generation
536 with a unidirectional reward model. In *Proceedings of the 2023 Conference on Empirical Methods
537 in Natural Language Processing*, pp. 11781–11791, 2023.

540 George R Doddington, Alexis Mitchell, Mark A Przybocki, Lance A Ramshaw, Stephanie M
 541 Strassel, and Ralph M Weischedel. The automatic content extraction (ace) program-tasks, data,
 542 and evaluation. In *Lrec*, volume 2, pp. 837–840. Lisbon, 2004.

543 Alexander R Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and
 544 Dragomir Radev. Summeval: Re-evaluating summarization evaluation. *Transactions of the Association
 545 for Computational Linguistics*, 9:391–409, 2021.

546 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad
 547 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd
 548 of models. *arXiv preprint arXiv:2407.21783*, 2024.

549 550 guidance-ai. Guidance: A guidance language for controlling large language models. <https://github.com/guidance-ai/guidance>, 2024.

551 552 Jia He, Mukund Rungta, David Koleczek, Arshdeep Sekhon, Franklin X Wang, and Sadid Hasan.
 553 Does prompt formatting have any impact on llm performance?, 2024. URL <https://arxiv.org/abs/2411.10541>.

554 555 Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and
 556 Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*, 2020.

557 558 Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child,
 559 Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language
 560 models. *arXiv preprint arXiv:2001.08361*, 2020.

561 562 Nitish Shirish Keskar, Bryan McCann, Lav R Varshney, Caiming Xiong, and Richard Socher.
 563 Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint arXiv:1909.05858*, 2019.

564 565 Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. Generalization
 566 through memorization: Nearest neighbor language models. In *International Conference on Learning Representations*, 2021.

567 568 Ben Krause, Akhilesh Deepak Gotmare, Bryan McCann, Nitish Shirish Keskar, Shafiq Joty, Richard
 569 Socher, and Nazneen Fatema Rajani. Gedi: Generative discriminator guided sequence generation.
 570 In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pp. 4929–4952,
 571 2021.

572 573 Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt
 574 tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*,
 575 pp. 3045–3059, 2021.

576 577 Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In
 578 *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*,
 579 pp. 4582–4597, 2021.

580 581 Alisa Liu, Maarten Sap, Ximing Lu, Swabha Swayamdipta, Chandra Bhagavatula, Noah A Smith,
 582 and Yejin Choi. Dexperts: Decoding-time controlled text generation with experts and anti-experts.
 583 In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 6691–6706, 2021.

584 585 Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. G-eval: Nlg
 586 evaluation using gpt-4 with better human alignment. *arXiv preprint arXiv:2303.16634*, 2023.

587 588 Do Xuan Long, Ngoc-Hai Nguyen, Tiviatis Sim, Hieu Dao, Shafiq Joty, Kenji Kawaguchi, Nancy F.
 589 Chen, and Min-Yen Kan. LLMs are biased towards output formats! systematically evaluating
 590 and mitigating output format bias of LLMs. In Luis Chiruzzo, Alan Ritter, and Lu Wang (eds.),
 591 *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for
 592 Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 299–
 593 330, Albuquerque, New Mexico, April 2025. Association for Computational Linguistics. ISBN
 979-8-89176-189-6. URL <https://aclanthology.org/2025.naacl-long.15/>.

594 Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. Cross-task generalization
 595 via natural language crowdsourcing instructions. In *ACL*, 2022.

596

597 OpenAI. ChatGPT, 2025. URL <https://chat.openai.com>. Large language model. Ac-
 598 cessed May 19, 2025.

599 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
 600 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to fol-
 601 low instructions with human feedback. *Advances in neural information processing systems*, 35:
 602 27730–27744, 2022.

603 Yi Ren and Danica J. Sutherland. Learning dynamics of llm finetuning, 2025. URL <https://arxiv.org/abs/2407.10490>.

604

605 Timo Schick, Sahana Udupa, and Hinrich Schütze. Self-diagnosis and self-debiasing: A proposal for
 606 reducing corpus-based bias in nlp. *Transactions of the Association for Computational Linguistics*,
 607 9:1408–1424, 2021.

608

609 Askhat Situdikov, Nikita Balagansky, Daniil Gavrilov, and Alexander Markov. Classifiers are better
 610 experts for controllable text generation. *arXiv preprint arXiv:2205.07276*, 2022.

611 Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford,
 612 Dario Amodei, and Paul F Christiano. Learning to summarize with human feedback. *Advances
 613 in Neural Information Processing Systems*, 33:3008–3021, 2020.

614

615 Zhi Rui Tam, Cheng-Kuang Wu, Yi-Lin Tsai, Chieh-Yen Lin, Hung yi Lee, and Yun-Nung Chen. Let
 616 me speak freely? a study on the impact of format restrictions on performance of large language
 617 models, 2024. URL <https://arxiv.org/abs/2408.02442>.

618 Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du,
 619 Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. *arXiv preprint
 620 arXiv:2109.01652*, 2021.

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

624 Zhihua Wen, Zhiliang Tian, Zhen Huang, Yuxin Yang, Zexin Jian, Changjian Wang, and Dongsheng
 625 Li. Grace: gradient-guided controllable retrieval for augmenting attribute-based text generation.
 626 In *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 8377–8398, 2023.

627

628 Brandon T Willard and Rémi Louf. Efficient guided generation for llms. *arXiv preprint
 629 arXiv:2307.09702*, 2023.

630 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang
 631 Gao, Chengan Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu,
 632 Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin
 633 Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang,
 634 Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui
 635 Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang
 636 Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger
 637 Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan
 638 Qiu. Qwen3 technical report, 2025a. URL <https://arxiv.org/abs/2505.09388>.

639 Kevin Yang and Dan Klein. Fudge: Controlled text generation with future discriminators. In *Pro-
 640 ceedings of the 2021 Conference of the North American Chapter of the Association for Compu-
 641 tational Linguistics: Human Language Technologies*, pp. 3511–3535, 2021.

642 Qwen: An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan
 643 Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang,
 644 Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin
 645 Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li,
 646 Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang,
 647 Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2025b.
 648 URL <https://arxiv.org/abs/2412.15115>.

648 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik
649 Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *Ad-*
650 *vances in neural information processing systems*, 36:11809–11822, 2023.

651

652 Yoel Zeldes, Dan Padnos, Or Sharir, and Barak Peleg. Technical report: Auxiliary tuning and its
653 application to conditional text generation. *arXiv preprint arXiv:2006.16823*, 2020.

654

655 Yongcheng Zeng, Guoqing Liu, Weiyu Ma, Ning Yang, Haifeng Zhang, and Jun Wang. Token-level
656 direct preference optimization. In *Proceedings of the 41st International Conference on Machine*
657 *Learning*, pp. 58348–58365, 2024.

658 Honghua Zhang, Meihua Dang, Nanyun Peng, and Guy Van den Broeck. Tractable control for
659 autoregressive language generation, 2023. URL <https://arxiv.org/abs/2304.07438>.

660 Honghua Zhang, Po-Nien Kung, Masahiro Yoshida, Guy Van den Broeck, and Nanyun Peng. Adapt-
661 able logical control for large language models, 2024. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2406.13892)
662 2406.13892.

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

APPENDIX

A LIMITATIONS

While we show DECO-G enhances LLM task performance in various tasks, a few limitations should be taken into consideration when using DECO-G. Firstly, the HMM used to estimate format satisfaction rate is specific to an LLM, meaning that one has to distill a new HMM when switching to a different LLM. Although, in practice, we find out that an HMM can be applied to larger LLMs in the same family, it is not an accurate representation of their token distributions. Secondly, similar to other CTG methods that includes additional module for attribute modeling, DECO-G introduces additional computation overhead during decoding. As suggested by (Zhang et al., 2024), the probabilistic traversal of future generation courses using HMM has a complexity that is linear to the number of edges in DFA and quadratic to the number of hidden states in HMM. Complex format constraints, converted to larger DFA, are thus likely to increase generation runtime. Finally, finding the optimal hyperparameter γ for LLMs with highly peaked token distributions may require empirical explorations. Such distributions require increased γ to ensure robust format compliance, yet an excessive value may adversely affect the output quality.

B PROBABILISTIC REASONING OVER LOGICAL CONSTRAINTS

Zhang et al. (2023) and Zhang et al. (2024) use a TPM to perform a probabilistic lookahead—that is, to efficiently compute $P_{\text{TPM}}(\alpha|x_{\leq t})$, the probability that the full generated sequence will satisfy the constraint α . The constraint is encoded as a DFA, and compliance is denoted by the event $\mathcal{D}(\alpha) = 1$. Then, the marginal probability over all sequences accepted by $\mathcal{D}(\alpha)$ is expressed as

$$P_{\text{TPM}}(\alpha|x_{\leq t}) = \frac{P_{\text{TPM}}(\mathcal{D}(\alpha) = 1, x_{\leq t})}{P_{\text{HMM}}(x_{\leq t})} \quad (7)$$

where the numerator—likelihood of satisfying the constraint given a prefix $x_{\leq t}$ —is found by marginalizing over the HMM hidden states z_t and the DFA states s_t

$$P_{\text{TPM}}(\mathcal{D}(\alpha) = 1, x_{\leq t}) = \sum_{z_t} P_{\text{TPM}}(\mathcal{D}(\alpha) = 1|z_t, s_t) P_{\text{HMM}}(z_t, x_{\leq t}) \quad (8)$$

The conditional compliance probability $P_{\text{TPM}}(\mathcal{D}(\alpha) = 1|z_t, s_t)$, as shown in Zhang et al. (2024), is calculated using a backward recurrence relation. This sums the probabilities of all valid transitions from step t to $t + 1$, weighted by the HMM’s transition and emission probabilities

$$P(\mathcal{D}(\alpha) = 1|z_t, s_t) = \sum_{z_{t+1}} P(z_{t+1}|z_t) \sum_{s_{t+1}} P(\mathcal{D}(\alpha) = 1|z_{t+1}, s_{t+1}) \sum_{\delta(s_t, x_{t+1}) = s_{t+1}} P(x_{t+1}|z_{t+1}) \quad (9)$$

The computed probability then serves as the tractable approximation of the format compliance likelihood, which is used to guide the LLM’s next-token generation as shown in Equation 3.

C SUBOPTIMAL CONTROL FROM UNCONDITIONED HMM DISTILLATION

When an instruction-tuned model is prompted with no specific user input for unconditional sampling, it often defaults to generic conversational phrases like, “*Is there something I can help you with?*” This behavior is a byproduct of its safety and helpfulness training. An HMM distilled from thousands of such non-substantive responses learns a token distribution that is unrepresentative of the model’s capabilities in actual problem-solving scenarios.

Consequently, when this poorly-approximated HMM is applied to a complex reasoning task, it provides a suboptimal control signal. The HMM, having not learned the patterns of reasoned thought, cannot accurately predict the LLM’s token distribution during task execution. This leads to improper guidance that can disrupt the generation process. For example, when we applied CtrlG to the GSM8k dataset, its control mechanism prematurely forced the model to generate the required format phrase (“*The final answer is ...*”), suppressing the step-by-step reasoning necessary to solve

756 the problem. This resulted in an accuracy of only 60.6%, a significant drop compared to the standard
 757 natural language baseline.
 758

759 The table below provides an example of this failure mode on a GSM8k problem, contrasting CtrlG’s
 760 flawed output with the coherent response from DECO-G, which uses an instruction-aware HMM.
 761

762 Table 4: Comparison of CtrlG and DECO-G outputs on a GSM8k reasoning task.
 763

764 Method	764 Output
765 CtrlG	765 The final answer is 0.36 (INCORRECT)
766 DECO-G	766 To find the probability that both tickets are winners, we need to multiply the probabilities of each ticket winning. 767 768 1. The probability of the first ticket winning is 20% or 0.2. 769 2. The probability of the second ticket winning is three times more likely, so it’s 770 $3 \times 0.2 = 0.6$. 771 3. The probability of both tickets winning is the product of their individual probabilities: $0.2 \times 0.6 = 0.12$. 772 4. To express this as a percentage, we multiply by 100: $0.12 \times 100 = 12\%$. 773 774 The final answer is 12. (CORRECT)

775 D HMM PRUNING AND EFFICIENCY

781 As established in Section 3.4, hidden state pruning is employed to mitigate the computational over-
 782 head of the HMM. This optimization is empirically justified by the highly concentrated nature of
 783 the hidden state probability distribution, as illustrated in Figure 5. Our analysis confirms that for
 784 *Llama*- and *Qwen*-distilled HMMs, the top 5% of hidden states ($k = 200$) retain over 97.8% of the
 785 total probability mass on average.

786 This high mass retention translates to a negligible impact on task performance. As shown in Table 6,
 787 the accuracy degradation on the GSM8k benchmark is minimal when pruning is applied: -0.2% for
 788 *Llama*, -1.2% for *Qwen2.5*, and +0.9% for *Qwen3*. This result validates that the pruned HMM
 789 provides sufficient guidance, confirming the efficacy of the optimization.

790 The primary benefit of this approach is a substantial improvement in computational efficiency. By
 791 reducing the computation of the HMM emission stage, pruning achieves a 13x reduction in the
 792 FLOPs required by the HMM forward function at each decoding step (from approx. 1.08 GFLOPs
 793 to 0.08 GFLOPs for *Llama*, see Section E for calculation). When compared to the LLM’s own for-
 794 ward pass, which requires approximately 16 GFLOPs per token (Kaplan et al., 2020), the pruned
 795 FEM’s computational cost constitutes only about 0.53% of the main inference workload. This opti-
 796 mization renders the guidance overhead practically insignificant, thereby enhancing the viability of
 797 the DECO-G framework.

798 E HMM FORWARD COMPUTATION COST

801 This section details the computational cost (in FLOPs) of the HMM’s forward pass. The calculation
 802 uses the HMM parameters for the *Llama* model: hidden states $h=4096$, vocabulary size $|\mathcal{V}|=128k$,
 803 and top-k states for pruning $k=200$.
 804

805 **Before Pruning** The total cost is the sum of the state transition cost ($2h^2$) and the emission cost
 806 ($2h|\mathcal{V}|$).
 807

$$808 \text{Total FLOPs} = (2 \times 4096^2) + (2 \times 4096 \times 128,000) \\ 809 = (3.36 \times 10^7) + (1.05 \times 10^9) \approx \mathbf{1.08 \text{ GFLOPs}}$$

810 **After Pruning** The cost is the sum of the state transition cost and the pruned emission cost ($2k|\mathcal{V}|$).

811

$$812 \text{Total FLOPs} = (2 \times 4096^2) + (2 \times 200 \times 128,000)$$

813

$$814 = (3.36 \times 10^7) + (5.12 \times 10^7) \approx \mathbf{0.08 \text{ GFLOPs}}$$

815 This optimization reduces the HMM’s computational overhead from 1.08 GFLOPs to 0.08 GFLOPs,
816 a **~13x reduction** per decoding step.

817

F PROMPT CONSTRUCTION

820 We present the set of prompts used in the experiments. For GSM8k (Table 7), we sample from a
821 set of task instructions and a set of format instructions to construct prompts for baseline methods.
822 For SummEval (Table 8), we include domain specific scoring criteria in the task instructions to
823 help LLM align better with human annotations for all methods. For ACE05 (Table 9), an event
824 description is appended to the task instructions which further explains the event of interest.

825

826 **Table 5: Full ACE05 Results.**

827

828 Model	829 Method	830 AI			831 AC			832 AI+			833 AC+		
		Precision	Recall	f1	Precision	Recall	f1	Precision	Recall	f1	Precision	Recall	f1
830 Llama-3.1-8B-Instruct	NL	33.7	40.5	36.8	24.2	31.4	27.3	30.5	40.5	34.8	21.7	31.0	25.5
	NL-S	33.7	41.3	37.1	24.4	32.3	27.8	30.5	41.3	35.1	21.8	32.1	26.0
	JSON	30.5	41.8	35.2	21.5	33.3	26.2	27.8	41.7	33.4	19.6	33.0	24.6
	JSON-S	28.0	42.0	33.6	19.8	34.4	25.2	25.4	41.6	31.6	18.0	34.1	23.6
	DECO-G	39.4	39.3	39.4	27.3	30.3	28.7	35.5	38.5	37.0	24.5	29.6	26.8
833 Qwen2.5-7B-Instruct	NL	29.7	36.2	32.6	22.6	29.4	25.5	27.5	35.7	31.2	21.0	29.0	24.4
	NL-S	28.6	39.5	33.2	20.7	31.4	24.9	26.2	39.0	31.3	19.1	31.2	23.7
	JSON	29.1	35.2	31.9	21.3	27.7	24.1	27.0	35.0	30.5	19.7	27.4	22.9
	JSON-S	32.2	36.2	34.1	23.9	28.7	26.1	29.6	36.1	32.5	21.9	28.4	24.7
	DECO-G $\gamma=2$	30.9	41.0	35.2	21.8	32.0	25.9	28.4	40.5	33.4	20.1	31.4	24.5
836 Qwen3-8B	NL	28.0	40.9	33.2	19.5	33.3	24.6	25.6	40.9	31.5	17.8	33.0	23.1
	NL-S	27.9	41.4	33.3	18.9	33.6	24.2	25.4	41.3	31.5	17.1	33.2	22.6
	JSON	27.0	38.4	31.7	18.6	31.7	23.4	24.7	38.4	30.1	16.7	31.4	21.8
	JSON-S	26.8	41.4	32.5	17.5	33.4	23.0	24.5	41.3	30.8	15.9	33.2	21.5
	DECO-G $\gamma=2$	27.9	43.6	34.0	18.9	35.1	24.6	25.5	43.4	32.1	17.3	34.8	23.1

837

G MORE EAE RESULTS

838 In Table 2, we report the *f1-scores* for each method. In Table 5, we present the full results for our
839 event argument extraction experiment.

840

H HMM DISTILLATION AND USAGE

841 For *Llama* and *Qwen*, we distill their HMMs on the LLM continuation only, since the instructions
842 from *Natural-Instructions* are human authored and should not be considered reflecting LLM distri-
843 bution. We remove the special chat tokens (e.g. `<|system|`, `<|user|`, etc.) from the responses for
844 HMM to capture the natural language distribution.

845 We tried different inputs to the HMM, including 1) regular prompt (with chat template), 2) cleaned
846 text prompt (without chat template), and 3) no prompt (empty string). In practice, their results are
847 almost identical. Nonetheless, in accordance with the distillation objective, we report scores yielded
848 from using empty input to the HMM.

849

I VISUALIZATION OF DECO-G’S STEERING PROCESS

850 Figure 3 shows an example of DECO-G steering *Llama*’s token probability to encourage the gener-
851 ation of format tokens.




Figure 3: DECO-G steers *Llama* to generate predefined template “The final answer is ...” by boosting probabilities of template tokens.




Figure 4: LLM’s token-level entropy for different models and methods. *Llama* has a more flexible token distribution as compared to *Qwen*.




Figure 5: Average retention rate (of total mass) over top-k HMM hidden states on GSM8k dataset.

Table 6: DECO-G performance on GSM8k with and without pruning.

Method	Acc (%)
<i>Llama-3.1-8B-Instruct</i>	
DECO-G w/o Pruning	85.4
DECO-G	85.2 ($\Delta=-0.2$)
<i>Qwen2.5-7B-Instruct</i>	
DECO-G w/o Pruning	89.8
DECO-G	88.6 ($\Delta=-1.2$)
<i>Qwen3-8B</i>	
DECO-G w/o Pruning	90.8
DECO-G	91.7 ($\Delta=+0.9$)

Table 7: GSM8k prompt construction and an example question.

GSM8k	
Task Instructions	<p>1. Follow the instruction to complete the task: You are a math tutor who helps students of all levels understand and solve mathematical problems. Read the last question carefully and think step by step before answering, the final answer must be only a number.</p> <p>2. Follow the instruction to complete the task: Read the last question carefully and think step by step before answering, the final answer must be only a number. You are a math tutor who helps students of all levels understand and solve mathematical problems.</p> <p>3. Follow the instruction to complete the task: Mathematical problem-solving task: - Given: A mathematical question or problem - Required: A numerical answer only - Role: You are a math tutor assisting students of all levels - Process: Think step by step to solve the problem Note: Read the question carefully before beginning your analysis.</p>
NL Format Instructions	<p>1. Provide your output in the following text format: <think step by step> The final answer is <answer></p> <p>2. Provide your output in the following text format: <Reasoning: <reasoning first>. Answer: The final answer is ...></p>
JSON Format Instructions	Provide your output in the following valid JSON format: <{"reason": "<step by step reasoning>", "answer": "<final answer>"}>
Question Example	Janet’s ducks lay 16 eggs per day. She eats three for breakfast every morning and bakes muffins for her friends every day with four. She sells the remainder at the farmers’ market daily for \$2 per fresh duck egg. How much in dollars does she make every day at the farmers’ market?

972
973
974

Table 8: SummEval prompts.

975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

SummEval	Task Instructions (Coherence)	You will be provided with a summary written for a news article. \nYour task is to rate the summary based on its coherence. \n\nPlease ensure you read and understand these instructions carefully. Keep this document open while reviewing, and refer to it as needed. \n\nEvaluation Criteria: \nCoherence (1-5): \n- 5: The summary is well-structured and organized, presenting information in a logical and seamless flow. \n- 4: The summary is mostly coherent, with minor lapses in organization or flow. \n- 3: The summary has noticeable organizational issues or lacks a smooth flow but is somewhat understandable. \n- 2: The summary is poorly structured, with significant difficulties in following its logic or flow. \n- 1: The summary is highly disjointed and lacks any meaningful structure or coherence. \nUse these criteria to assign a coherence score between 1 and 5 based on how well the summary organizes and presents information in a clear and logical manner.
	Task Instructions (Consistency)	You will be provided with a news article and a summary written for this article. \nYour task is to rate the summary based on its consistency. \n\nPlease ensure you read and understand these instructions carefully. Keep this document open while reviewing, and refer to it as needed. \n\nEvaluation Criteria: \nConsistency (1-5): \n- 5: The summary is fully factually accurate and all its statements are directly supported by the source document. \n- 4: The summary is mostly factually accurate, with only minor errors or omissions. \n- 3: The summary contains noticeable factual errors or unsupported statements but retains some alignment with the source document. \n- 2: The summary has significant factual inaccuracies or includes multiple unsupported claims. \n- 1: The summary is largely inconsistent with the source, containing numerous factual inaccuracies or fabricated details. \nUse these criteria to assign a consistency score between 1 and 5 based on how well the summary aligns factually with the source article.
	Task Instructions (Fluency)	You will be provided with a summary written for a news article. \nYour task is to rate the summary based on its fluency. \n\nPlease ensure you read and understand these instructions carefully. Keep this document open while reviewing, and refer to it as needed. \n\nEvaluation Criteria: \nFluency (1-5): \n- 5: The summary is clear and easy to read, with good grammar, spelling, and sentence structure. \n- 4: The summary is generally clear and fluent, with a few minor errors that don't interfere with understanding. \n- 3: The summary has some noticeable issues that might make it a little harder to read but still understandable overall. \n- 2: The summary has more noticeable problems that might make it challenging to follow in places. \n- 1: The summary has significant errors that make it difficult to read or understand in many parts. \nUse these criteria to assign a fluency score between 1 and 5 based on the quality of grammar, word choice, and sentence structure. \nImportant: When evaluating fluency, ignore punctuation and capitalization. Focus only on how natural and easy the language feels regardless of formatting.
	Task Instructions (Relevance)	You will be provided with a summary written for a news article. \nYour task is to rate the summary based on its relevance. \n\nPlease ensure you read and understand these instructions carefully. Keep this document open while reviewing, and refer to it as needed. \n\nEvaluation Criteria: \nRelevance (1-5): \n- 5: The summary includes all the important information from the source document with no redundancies or irrelevant details. \n- 4: The summary is mostly relevant, with only minor omissions or slight redundancies. \n- 3: The summary includes some important information but misses key points or has noticeable redundancies. \n- 2: The summary contains limited relevant information, with significant omissions or excessive irrelevant content. \n- 1: The summary is largely irrelevant, failing to capture the main points of the source document. \nUse these criteria to assign a relevance score between 1 and 5 based on how well the summary captures the important content from the source without including excess or redundant information.
	NL Format Instructions	Provide your output in the following text format: <analyze the summary>. The rating is <a number between 1 and 5>
	JSON Format Instructions	Provide your output in the following valid JSON format:\n{"analysis": "<analyze the summary>","rating": <a number between 1 and 5>}

1026
 1027
 1028
 1029
 1030
 1031
 1032
 1033
 1034
 1035
 1036
 1037
 1038
 1039
 1040
 1041
 1042
 1043

Table 9: ACE05 prompt construction and an example question.

1044
ACE05

1046	Task Instructions	You are an argument extractor designed to check for the presence of arguments regarding specific roles for an event in a sentence. \nTask Description: Identify all arguments related to the role Attacker, Target, Instrument, Place, Agent in the sentence. These arguments should have the semantic role corresponding to the given event trigger by the word span between [t] and [t].
1047		The event of interest is Conflict:Attack. The event is related to conflict and some violent physical act. Roles of interest: Attacker, Target, Instrument, Place, Agent
1048	NL Format Instructions	Provide your output in the following text format:\nThe <role_1> is: <extracted argument>\nThe <role_2> is: <extracted argument>\n...\nThe <role_n> is: <extracted argument>
1049	JSON Format Instructions	Provide your output in the following valid JSON format:\n{“<role>”: “<extracted argument>” for role in roles of interest}

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