# AUTO-DEMO PROMPTING: LEVERAGING GENERATED OUTPUTS AS DEMONSTRATIONS FOR ENHANCED BATCH PROMPTING

Anonymous authors

006

008 009 010

011 012 013

014

015

016

017

018

019

021

024

025

026

027

028

029

031 032

033 034 Paper under double-blind review

#### Abstract

Batch prompting is a common technique in large language models (LLMs) used to process multiple inputs simultaneously, aiming to improve computational efficiency. However, as batch sizes increase, performance degradation often occurs due to the model's difficulty in handling lengthy context inputs. Existing methods that attempt to mitigate these issues rely solely on batch data arrangement and majority voting rather than improving the design of the batch prompt itself. In this paper, we address these limitations by proposing "Auto-Demo Prompting," a novel approach that leverages the question-output pairs from earlier questions within a batch as demonstrations for subsequent answer inference. We provide a formal theoretical analysis of how Auto-Demo Prompting functions within the autoregressive generation process of LLMs, illustrating how it utilizes prior outputs to optimize the model's internal representations. Our method effectively bridges the gap between batch prompting and few-shot prompting, enhancing performance with only a slight compromise in token usage. Experimental results across five NLP tasks demonstrate its effectiveness in mitigating performance degradation and occasionally outperforming single prompts. Furthermore, it opens new avenues for applying few-shot learning techniques, such as demonstration selection, within batch prompting, making it a robust solution for real-world applications.

### 1 INTRODUCTION

Large language models (LLMs), such as GPT (Brown et al., 2020), and PaLM (Chowdhery et al., 2023), have demonstrated an extraordinary ability to perform in-context learning (ICL), where they utilize provided examples or contextual information to adapt and solve a wide range of downstream tasks. This capability enables LLMs to generalize from few-shot or even zero-shot examples without requiring task-specific fine-tuning, significantly enhancing their versatility across diverse applications (Song et al., 2023). The success of ICL in these models highlights their potential as powerful tools for natural language processing and as adaptable frameworks for learning in dynamic, dataconstrained environments, offering broader implications for machine learning and AI research.

Recently, the batch prompting method has attracted growing research interest (Cheng et al., 2023; 043 Lin et al., 2024; Fan et al., 2024), where models are presented with a set of homogeneous questions 044 - queries that share similar structure or content - within a single prompt. The main objective of batch prompting is to enhance interaction efficiency with LLMs by reducing computational costs, 046 especially by minimizing the number of tokens processed. By grouping similar questions, this tech-047 nique streamlines the model's handling of multiple tasks, optimizing resource usage while ensuring 048 consistent performance across repeated or related queries. As illustrated in Figure 1 (a), the model processes multiple sentences with similar structures in a single batch, identifying and correcting 050 grammatical errors for each input sentence. This demonstrates how batch prompting reduces to-051 ken usage when working with structurally similar tasks. Moreover, advancements in hardware and algorithmic design have further expanded the capacity of LLMs to retain and process longer in-052 put contexts, enabling more effective batch prompting (Munkhdalai et al., 2024; Chen et al., 2023; Ainslie et al., 2023).

077

079

054		
055	Batch Prompt	Auto-Demo Prompt
056		
057	sentence grammar check.	sentence grammar check.
058	Requirement: Generate the answer following the format of the	Requirement: Repeat the input data and generate the answer
059	examples below.	following the format of the examples below.
060	 [input]	 [input]
061	Input 1: XXXXX Input 2: XXXXX	Input 1: XXXXX Input 2: XXXXX
062	-	
063	Label1: XXXXX	[output] {Input 1: XXXXX, Label1: XXXXX}
064	Label2: XXXXX	{Input 2: XXXXX, Label2: XXXXX}
065	Please generate labels for the [Batch_Size] data given the instructions.	Please generate labels for the [Batch_Size] data given the instructions.
066	Input 1. The plan was enpressed of hy my methor	Input 4. The plan was approved of hy my methor
067	Input 2: They'll be leaving.	Input 2: They'll be leaving.
068	Input 3: Boston was flown to.	Input 3: Boston was flown to.
069	Input 5: The coat does not fit you.	Input 5: The coat does not fit you.
070	Input 6: Lee never left. Input 7: The roof is leak.	Input 6: Lee never left. Input 7: The roof is leak.
071		
072	in the United States.	in the United States.
073		
074		

Figure 1: Example: a) Batch Prompting and b) Auto-Demo Prompting

While batch prompting provides significant advantages in terms of efficiency, it also presents a fundamental challenge, as handling lengthy context inputs is particularly difficult for LLMs based on 081 transformers. Transformer models experience significant performance degradation when accessing relevant information from the middle of long contexts, leading to a decline in overall effectiveness 083 (Liu et al., 2024). This issue becomes more pronounced in tasks that require LLMs to process long 084 contextual inputs. The root of the challenge lies in the quadratic complexity of the self-attention 085 mechanism within transformers, where computational costs increase dramatically with input length. 086 As batch prompting is applied, the combined input length grows, exacerbating this issue and further degrading model performance. Therefore, developing a more effective batch prompting strategy 087 is crucial for mitigating these performance limitations and unlocking the potential of large-scale 088 applications in in-context learning, particularly for tasks involving long sequences of input data. 089

From the internal representations of LLMs, we observe that batch prompting shares similarities with few-shot prompting, a technique that has proven crucial in enhancing the performance of LLMs by providing a few example demonstrations (Brown et al., 2020). However, an intriguing contradiction emerges: while few-shot prompting typically boosts performance, batch prompting often results in performance degradation. It was found that batch prompting can sometimes outperform single prompts in tasks with smaller batch sizes (Cheng et al., 2023), yet this line of investigation remains underexplored. This raises a key question: Could we bridge the gap between batch prompting and few-shot prompting to leverage the benefits of both?

Interestingly, previous research has found that even when incorrect labels are present in few-shot demonstrations during in-context learning, the decrease in accuracy is generally minimal, typically between 0% and 5% (Min et al., 2022). This suggests that using model-generated answers as demonstrations in few-shot prompts is a viable approach, regardless of the potential for model hallucination or limitations that may lead to incorrect answers.

In this paper, we tackle the performance degradation problem from the perspective of connecting
 few-shot prompting with batch prompting. A key insight emerges: the outputs generated during
 earlier iterations of autoregressive LLMs can, with proper design, be automatically recognized
 as demonstrations for subsequent text generation, without needing to explicitly include them
 in the prompt. Building on this, we propose "Auto-Demo Prompting," a novel batch prompting
 technique that instructs the model to repeat each question before answering it. The model then au-

tomatically treats the resulting question-answer pairs as demonstrations for the following questions in the batch, eliminating the need to manually pack them into the input prompt.

As illustrated in Figure 1 (b), Auto-Demo Prompting directs the model to produce a sequence of question-answer pairs instead of just answers. This structure ensures that, during the autoregressive generation process, each subsequent question in the batch has access to the generated questionanswer pairs from the previous iterations. Consequently, all questions in the auto-demo prompt effectively receive 0 to N - 1 additional demonstrations, potentially mitigating the performance degradation associated with long context inputs.

By integrating demonstrations directly into the batch prompting process, our method not only enhances the model's capacity to manage longer inputs but also improves overall task performance, aligning batch prompting more closely with the proven success of few-shot prompting. To assess the effectiveness of Auto-Demo Prompting, we conducted extensive experiments using mainstream models, GPT-40 and GPT-40-mini, across a variety of downstream tasks. The results demonstrate that the proposed method significantly improves model performance in both an efficient and interpretable way. Our contributions can be summarized as follows:

- The proposed "Auto-Demo Prompting" represents a pioneering effort in constructing fewshot demonstrations directly within the LLM's inference. While previous research has focused on various external factors associated with batch prompts, our approach uniquely modifies the batch prompt design to enable the model to generate demonstrations as part of its output process.
  - 2. Through extensive comparisons between few-shot prompting and batch prompting, we proposed and validated the hypothesis that Auto-Demo Prompting is approximately equivalent to batch prompting supplemented with few-shot demonstrations. All experimental results showed that Auto-Demo Prompting outperformed traditional batch prompting, supporting our hypothesis and addressing the limitation of performance degradation in batch prompt.
  - 3. We discovered that batch data selection in Auto-Demo Prompting influences the selection of demonstrations (question-answer pairs) during the autoregressive generation steps of inference. By employing a common demonstration selection technique, our results showed significant improvements compared to random batch data selection.
    - 4. The experimental results show that "Auto-Demo Prompting + Batch Data Selection," with batch sizes of 16 and 32, outperforms the single prompts with a batch size of 1. This high-lights the considerable potential of integrating additional few-shot prompting techniques into Auto-Demo Prompt to serve as an efficient alternative to traditional single prompts with enhanced accuracy and efficiency.
- 142 143

129

130

131

132

133

134

135

136

137

138

139

140

141

## 2 AUTO-DEMO PROMPTING

144 145

Prompt engineering is essential for unlocking the capabilities of large language models (Marvin 146 et al., 2023). Typically, a standard prompt consists of a task description and a single data point, 147 leading to the development of various prompting techniques to enhance LLM performance on down-148 stream tasks (Sahoo et al., 2024; Besta et al., 2024; Wang et al., 2023; Wei et al., 2022). However, 149 relying on single prompts limits inference to just one data point, making this approach less efficient 150 for large-scale datasets or real-world applications. In contrast, Batch Prompting allows for the pro-151 cessing of multiple data points in a single inference. Cheng et al. (2023) suggests that while batch 152 prompting performs well with smaller batch sizes, there is a tendency for performance to decline as 153 batch sizes grow. Additionally, Lin et al. (2024) noted that varying the order of batch data can yield different results, proposing the integration of these outcomes through majority voting to enhance 154 overall performance. 155

Unlike previous works, Auto-Demo Prompting aims to maximize the potential of batch prompting and enhance its performance during single inference of large language models. This method guides LLMs to generate a question-answer pair for each input question in the batch, rather than simply providing an answer. By innovating the design of batch prompts with a new output control format, Auto-Demo Prompting ensures that LLMs maintain a consistent output structure across all questions. A key factor in this process is the autoregressive generation mechanism of the decoder in LLMs, which facilitates coherent and contextually relevant output generation.



Figure 2: Auto-Demo Prompting: Single prompts are combined into a batch prompt with a special output control for generating question-answer pairs, along with optional batch data selection. This prompt is fed into the autoregressive generation process of a decoder-only LLM, forming demonstrations for subsequent generation.



2.1 Autoregressive Generation Process in Large Language Model

Large Language Models can be classified into three categories: encoder-only, encoder-decoder, and
 decoder-only. The mainstream LLMs are primarily decoder-only and encoder-decoder models. The
 GPT series developed by OpenAI, including notable examples such as GPT-3, GPT-4, and GPT-40, consists entirely of decoder-only models and marks significant milestones in the evolution of

generative language models. The Llama series represents open-source LLMs that are also decoder-only. In contrast, examples of encoder-decoder models include T5 (Raffel et al., 2020) and BART (Lewis, 2019). This paper focuses specifically on decoder-only LLMs, where the decoder generates one token at each step of the autoregressive generation process. A single inference of an LLM typically requires multiple steps to produce a complete answer, with the token generated at the previous step added to the current input, serving as the new input for the next token. This iterative process of token generation is known as autoregressive generation.

Based on the generation mechanism of LLMs, we introduce a key insight into the autoregressive generation process: the outputs from earlier steps can serve as prompts or demonstrations for generating subsequent tokens. By employing batch prompting, the model can process multiple input data sequentially, enabling consistent and repetitive responses. This periodicity allows for the use of answers from earlier questions as demonstrations for later ones, thereby enhancing coherence and context retention throughout the generation process.

- 229
- 230 231

### 2.2 Autoregressive Generation of Demonstrations

As illustrated in Figure 2, Auto-Demo Prompting creatss a loop where demonstrations are generated in an autoregressive manner during the inference of LLMs. This technique guides LLMs to produce question-answer pairs that align with the format of demonstrations used in context learning. Due to the autoregressive nature of token generation, the previously generated output is appended to the current input for the next step, serving as a demonstration for all subsequent questions in the batch. When the batch size is N, there will be  $\{1, 2, 3, ..., N - 1\}$  demonstrations formed in the process.

Suppose the batch size of Auto-Demo Prompting is denoted by N. The batch of questions can be represented as  $Q = \{q_1, q_2, \dots, q_N\}$ , and the corresponding answers are denoted by  $A = \{a_1, a_2, \dots, a_N\}$ . Let  $\mathcal{F}$  denote the inference of the LLM. To illustrate the relationships among current batch prompting, Auto-Demo Prompting, and few-shot demonstrations in in-context learning, we compare the formulations of the LLM inference process in these prompting approaches.

#### Method 1 Few-Shot Prompting

Formulation:  $\mathcal{F}(a_n | Q_{1:n-1}, \{(q_i, a_i)\}_{i=1}^{n-1}) \quad \forall n \in \{0, 1, \dots, N\}$ Example: 1:  $\mathcal{F}(a_1 | q_1)$ 2:  $\mathcal{F}(a_2 | (q_1, a_1))$ 3: ... 4:  $\mathcal{F}(a_n | (q_1, a_1), (q_2, a_2), \dots, (q_{n-1}, a_{n-1}))$ 

254

255

256

257

258

259

264

265

266

267

268

269

243

244 245

246

247

248

249

### Method 2 Batch Prompting (BP)

Formulation:  $\mathcal{F}(a_n | \mathbf{BP} + Q_{1:n}, \{a_i\}_{i=1}^{n-1}) \quad \forall n \in \{0, 1, ..., N\}$ Example: 1:  $\mathcal{F}(a_1 | \mathbf{BP} + q_1)$ 2:  $\mathcal{F}(a_2 | \mathbf{BP} + q_1 q_2, a_1)$ 3: ... 4:  $\mathcal{F}(a_n | \mathbf{BP} + q_1 q_2 ... q_n, a_1 a_2 ... a_{n-1})$ 

 Method 3 Auto-Demo Prompting (ADP)

 Formulation:  $\mathcal{F}((q_n, a_n) | ADP + Q_{1:n-1}, \{(q_i, a_i)\}_{i=1}^{n-1}) \quad \forall n \in \{0, 1, \dots, N\}$  

 Example:

 1:  $\mathcal{F}((q_1, a_1) | ADP + q_1)$  

 2:  $\mathcal{F}((q_2, a_2) | ADP + q_1q_2, (q_1, a_1))$  

 3: ...

 4:  $\mathcal{F}((q_n, a_n) | ADP + q_1q_2 \dots q_{n-1}, (q_1, a_1), (q_2, a_2), \dots, (q_{n-1}, a_{n-1}))$ 

270 Based on the comparison, we illustrate that Auto-Demo Prompting effectively combines conven-271 tional batch prompting and few-shot prompting by modifying the output format to enhance in-272 context learning capabilities. We can express this relationship using an approximate equality: 273 274 Auto-Demo Prompting  $\approx$  Batch Prompting + Few-Shot Demonstrations (1)275 276 This approximate equality underscores the significance of Auto-Demo Prompting. As shown in Figure 1 (b), it adds just two additional lines to the Batch Prompt, representing a minor modification. 277 278 Nevertheless, this small change greatly enhances the autoregressive generation of demonstrations during LLM inference and plays a crucial role in improving the performance of batch prompting. 279 280 281 2.3 BATCH DATA SELECTION 282 In the proposed Auto-Demo Prompting framework, batch data selection - a conventional technique 283 for improving batch prompting performance - can be viewed as demonstration selection. As illus-284 trated in Method 3, the batched data (i.e., questions) is also present within the question-answer pairs 285 generated by the LLM, which serve as demonstrations for subsequent questions. Therefore, we pro-286 pose the following hypothesis: the selection of batched data in Auto-Demo Prompting achieves a 287 similar effect to the selection of demonstrations in few-shot prompting. 288 289 Batch Data Selection  $\approx$  Demonstration Selection (2)290 291 Based on this hypothesis, demonstration selection methods can be effectively adapted for batch data 292 selection in Auto-Demo Prompting. Current approaches to demonstration selection are diverse, uti-293 lizing criteria such as similarity, mutual information, perplexity, and diversity (Yang et al., 2023). 294 Notably, demonstration selection based on text similarity has proven effective for text pair classifi-295 cation and multiple-choice tasks (Peng et al., 2024; Su et al., 2024). 296 Inspired by demonstration selection method, we designed the "Batch Data Selection with Retrieval" 297 algorithm to identify similar questions within a single batch. As detailed in Algorithm 1, this algo-298 rithm features a data retrieval loop aimed at gathering similar data into one batch. For each batch 299 with size N, a target data point is randomly selected. Subsequently, the N-1 most similar data 300 points are identified using an embedding model to calculate the pairwise similarity. 301 302 Algorithm 1 Batch Data Selection with Retrieval 303 1: Input: batch size = N, dataset  $D = \{d_1, d_2, \dots, d_{|D|}\}$ , batched data  $BD = \{\}$ 304 2: Output: BD 305 3: for  $d_i \in D$  do 306  $temp\_batch = \{d_i\}$ 4: for  $d_i \in D$  s.t  $d_i \neq d_i$  do 307 5: Calculate similarity between  $d_i$  and  $d_j$  using embedding model 6: 308 if  $d_i$  is among the most similar (N-1) then 7: 309 8:  $temp\_batch.add(d_i)$ 310 9: end if 311 10: end for 312 11:  $BD.add(temp\_batch)$ 313 12: end for 314 315 316 3 EXPERIMENTS 317 318 3.1 EXPERIMENTAL SETTINGS 319

We conducted comparative experiments between Auto-Demo Prompting and conventional batch
 prompting across various NLP tasks, including question answering (BoolQ), mathematical reason ing (GSM8K, SVAMP), textual entailment (RTE), and paraphrase detection (Quora Question Pairs,
 QQP). Specifically, the batch data selection experiments were focused on the RTE and QQP datasets to assess the effectiveness of our approach.

# 324 3.1.1 DATASETS

BoolQ: The Boolean Questions (BoolQ) dataset comprises 15,942 yes/no questions derived from
 real-world queries to the Google search engine (Clark et al., 2019). Each example includes a question, passage, and answer, with an optional page title for context. The dataset is split into 9,427
 training examples and 3,270 validation examples, noted for its complexity requiring advanced reasoning and entailment inference.

RTE: The Recognizing Textual Entailment (RTE) dataset, compiled from the RTE1, RTE2, RTE3,
 and RTE5 challenges (Poliak, 2020), includes 2,490 training samples, 277 validation samples, and
 3,000 testing samples. It features premise-hypothesis pairs, each labeled to indicate their entailment
 relationship, making it a valuable resource for advancing natural language inference research.

GSM8K: The Grade School Math 8K (GSM8K) dataset consists of 8,500 diverse grade school math word problems designed for multi-step reasoning tasks (Cobbe et al., 2021). It is divided into 7,500 training problems and 1,000 test problems, with each problem requiring 2 to 8 steps to solve. Solutions are provided in a step-by-step format, and the dataset is available in both standard and "Socratic" formats, which include meta-reasoning prompts.

QQP: The Quora Question Pairs (QQP) dataset contains over 400,000 question pairs from the Quora platform, each labeled to indicate semantic equivalence. This dataset is essential for tasks like paraphrasing and duplicate question identification, enabling researchers to explore various machine learning techniques for semantic textual similarity (Wang et al., 2017).

SVAMP: The SVAMP dataset features 1,000 simple math word problems aimed at assessing NLP models, designed for up to fourth-grade level and involving single-variable equations (Patel et al., 2021). Created to challenge models with variations of existing problems, it tests their contextual understanding and problem-solving accuracy.

348 349 350

### 3.1.2 IMPLEMENTATIONS

We evaluate Auto-Demo Prompt with varying batch sizes over GPT-4o-mini and GPT-4o. The prompts used for the different datasets are provided in Appendix A.

Embedding Model for Batch Data Selection: The embedding model utilized is
"iic/nlp\_corom\_sentence-embedding\_english-base," available on the ModelScope platform.
This model is a dual-tower text representation architecture that employs the CoROM model as
its foundation for the pre-trained language model. The training data is derived from the official
open-source MS MARCO Passage Ranking dataset (Bajaj et al., 2018).

Parameters: All experiments were conducted in an in-context learning setting, without any model training or fine-tuning. The temperature for model inference was consistently set to 0 to ensure uniformity in responses. Batch sizes were set to 1/16/32 for the gpt-4o-mini, which has a maximum output token limit of 16k, and to 8 or 16 for the gpt-4o, with a maximum output token limit of 8k.

363

## 364 3.2 RESULTS AND DISCUSSIONS

366 As shown in Figure 3, the proposed method, highlighted in red, consistently outperforms Batch 367 Prompt in most experiments, allowing LLMs to generate longer outputs without compromising per-368 formance. This highlights the beneficial impact of Auto-Demo Prompting, where the generated question-answer pairs enhance the accuracy of subsequent questions within the same batch. No-369 tably, when applied with GPT-40 on the GSM8K dataset, Auto-Demo Prompt achieved an accuracy 370 of 95.7% with a batch size of 16, surpassing the 95.3% accuracy of the single prompt (batch size 371 = 1). A similar trend was observed with the SVAMP dataset, demonstrating the effectiveness of 372 Auto-Demo Prompting in improving the performance of conventional batch prompts. 373

Our findings align with previous research that highlights few-shot prompting as an effective strategy for enhancing model accuracy, particularly in datasets that necessitate sophisticated reasoning steps (Wei et al., 2022). The proposed method demonstrated significant effectiveness on the BoolQ dataset, which consists of naturally occurring yes/no questions that often require complex reasoning. In contrast, the effectiveness of Auto-Demo Prompt is less pronounced in simpler datasets like



Figure 3: Experimental results: accuracy for Auto-Demo Prompt ("w") vs. Batch Prompt ("w/o") across different models and batch sizes.

RTE and QQP, as the GPT-40 model performs relatively well on these tasks without the need for additional prompting.

3.2.1 BATCH DATA SELECTION EXPERIMENTS

405

410

411 412 413

414

To validate the hypothesis stated in Equation 2, we conducted experiments on the RTE and QQP datasets to evaluate the effectiveness of selecting similar questions for batch data selection within the Auto-Demo Prompting framework. As shown in Figure 4, the 'Batch Data Selection with Retrieval' algorithm consistently improves accuracy compared to randomly selected batched data. The results of these experiments closely align with the effectiveness of few-shot demonstration selection based on similarity, as both approaches lead to significant performance improvements.

421 Surprisingly, the results indicate that experiments utilizing larger batch sizes significantly outper-422 form those employing a single prompt (i.e., batch size 1). For the RTE dataset, the Auto-Demo 423 Prompting approach with a batch size of 48 using the gpt-40-mini model achieved an accuracy of 424 89.5%, surpassing the 88.8% accuracy obtained with a single prompt. Similarly, the gpt-40 model 425 with a batch size of 16 achieved an accuracy of 90.3%, exceeding the 88.4% accuracy of the single 426 prompt approach. In the QQP dataset, the batch size of 32 with the gpt-40 model yielded an accu-427 racy of 87.3%, which is 2.0% higher than the result from the single prompt. These findings suggest 428 that the proposed method not only addresses the limitation of performance degradation associated 429 with larger batch sizes, but further enhances overall performance by leveraging previous questions as demonstrations in the autoregressive generation process. This improvement aligns closely with the 430 effects of many-shot learning and underscores the importance of incorporating multiple examples in 431 the prompting process.



Figure 4: Experimental results for Batch Data Selection: Auto-Demo Prompting ("w") vs. Batch Prompting ("w/o"); "+bds" signifies the application of batch data selection with retrieval.

### 3.2.2 DISCUSSION

448 The results of our experiments underscore the significant potential of Auto-Demo Prompting. These 449 findings corroborate our earlier observations regarding equations 1 and 2, demonstrating that Auto-450 Demo Prompting consistently enhances Batch Prompting across most experimental scenarios. It 451 is particularly noteworthy the performance improvement observed when batch data selection is ap-452 plied. By grouping similar data within each batch, we found that Auto-Demo Prompting with larger 453 batch sizes outperformed configurations utilizing a batch size of one, which is the standard prompt 454 used in mainstream LLM applications. This advancement signifies a promising avenue for achieving 455 both enhanced performance and increased efficiency.

456 Furthermore, recent research by (Agarwal et al., 2024) demonstrates that many-shot demonstrations 457 can outperform few-shot demonstrations, a trend observed across various domains (Jiang et al., 458 2024; Moayedpour et al., 2024). As the input and output length limits of LLMs continue to expand, 459 the adoption of a more effective batch data selection method is likely to enhance these performance 460 gains, further solidifying the benefits of our approach. Given the diverse and numerous demonstra-461 tion selection methods available for in-context learning, the optimal choice can vary across different 462 datasets. This variability encourages future research to evaluate their effectiveness and to develop 463 novel approaches that are better suited for the Auto-Demo Prompting.

464 465

466 467

468

443

444

445 446 447

### 4 RELATED WORK

### 4.1 BATCH PROMPTING

469 Batch prompting has emerged as a significant area of research in LLMs to facilitate efficient batch 470 processing of many data points simultaneously. Cheng et al. (2023) first introduced the concept of batch prompting to reduce computational costs, with experiments focusing on batch sizes of 471 fewer than six demonstrating comparable performance to standard prompting. Building on this, Lin 472 et al. (2024) proposed "BatchPrompt," which highlights the variability in results caused by different 473 data orders and introduces a self-consistency method to address these discrepancies. With its cost-474 saving benefits, batch prompting has been applied in various NLP applications; for instance, Fan 475 et al. (2024) proposed a framework utilizing a covering-based demonstration selection strategy for 476 entity resolution, effectively balancing matching accuracy and computational cost, while Zhang et al. 477 (2023) explored the impact of batch size on several data preprocessing tasks.

478 479

480

### 4.2 DEMONSTRATION SELECTION IN IN-CONTEXT LEARNING

To improve the performance of in-context learning, previous studies have explored the optimization of the selection and arrangement of few-shot demonstrations (Rubin et al., 2022; Zhang et al., 2022; Wu et al., 2023; Fu et al., 2023; Zhou et al., 2023). It was found that the choice of demonstrations is both data-dependent and model-dependent, leading to the proposal of a selection method that considers both factors (Peng et al., 2024). Empirical research by (Min et al., 2022) has demonstrated that preserving the structured format of demonstrations, which consists of text-label pairs, is

essential for optimal performance. They also found that randomly altering the labels within these
demonstrations has a negligible effect on performance, providing a foundation for our research.

Our work diverges from previous research in several key aspects. While existing studies have explored batch prompting and its applications, they often overlook the relationship between few-shot prompting and batch prompting, which we believe presents an emerging opportunity for directly enhancing batch prompting. To address this gap, we introduced the Auto-Demo Prompt, which not only improves performance but also offers greater interpretability by establishing a close alignment with the principles of in-context learning and demonstration selection techniques.

494 495 496

### 5 CONCLUSION

497 In this paper, we introduce the Auto-Demo Prompt, a novel batch prompting method aimed at im-498 proving the performance of batch prompting. We provide a comprehensive overview of its oper-499 ational mechanism and clarify the key concepts that establish the relationship between few-shot 500 demonstrations in in-context learning and batch prompting. Importantly, we show that 'Batch Data 501 Selection' can be conceptualized as 'Demonstration Selection' within the framework of Auto-Demo 502 Prompting, facilitating the transfer of insights from few-shot learning research to our approach. Our extensive experiments validate the effectiveness of Auto-Demo Prompting, underscoring its signifi-504 cance as the input and output lengths of LLMs increase. As these lengths expand, the importance of 505 efficient and high-performing batch prompting methods, such as Auto-Demo Prompting combined with Batch Data Selection, will become increasingly pronounced. Further research is encouraged 506 to investigate demonstration selection in greater detail, as well as to explore the potential syner-507 gies between Auto-Demo Prompting and other emerging techniques such as prompt learning and 508 explainable AI. By delving deeper into these areas, we can uncover new insights that may enhance 509 the adaptability and effectiveness of LLMs, ultimately paving the way for advancements in their 510 applications toward achieving artificial general intelligence (AGI).

511 512

### 513 **REFERENCES**

- Rishabh Agarwal, Avi Singh, Lei M Zhang, Bernd Bohnet, Luis Rosias, Stephanie C.Y. Chan, Biao Zhang, Aleksandra Faust, and Hugo Larochelle. Many-shot in-context learning. In *ICML 2024 Workshop on In-Context Learning*, 2024. URL https://openreview.net/forum?id= goi7DFHlqS.
- Joshua Ainslie, James Lee-Thorp, Michiel de Jong, Yury Zemlyanskiy, Federico Lebron, and Sumit Sanghai. GQA: Training generalized multi-query transformer models from multi-head checkpoints. In *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023. URL https://openreview.net/forum?id=hmOwOZWZYE.
- Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, Mir Rosenberg, Xia Song, Alina Stoica, Saurabh Tiwary, and Tong Wang. Ms marco: A human generated machine reading comprehension dataset, 2018. URL https://arxiv.org/abs/1611.09268.
- Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Michal Podstawski, Lukas Gianinazzi, Joanna Gajda, Tomasz Lehmann, Hubert Niewiadomski, Piotr Nyczyk, et al. Graph of thoughts: Solving elaborate problems with large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 17682–17690, 2024.
- 531 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhari-532 wal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, 534 Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz 535 Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec 536 Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), Advances in Neural Information Processing Systems, volume 33, pp. 1877–1901. Curran Associates, Inc., 538 2020. URL https://proceedings.neurips.cc/paper\_files/paper/2020/ file/1457c0d6bfcb4967418bfb8ac142f64a-Paper.pdf.

567

568

- Shouyuan Chen, Sherman Wong, Liangjian Chen, and Yuandong Tian. Extending context window of large language models via positional interpolation, 2023. URL https://arxiv.org/abs/2306.15595.
- Zhoujun Cheng, Jungo Kasai, and Tao Yu. Batch prompting: Efficient inference with large language model APIs. In Mingxuan Wang and Imed Zitouni (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pp. 792–810, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023. emnlp-industry.74. URL https://aclanthology.org/2023.emnlp-industry.74.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm:
  Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240): 1–113, 2023.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina
   Toutanova. Boolq: Exploring the surprising difficulty of natural yes/no questions. *arXiv preprint arXiv:1905.10044*, 2019.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Meihao Fan, Xiaoyue Han, Ju Fan, Chengliang Chai, Nan Tang, Guoliang Li, and Xiaoyong Du.
   Cost-effective in-context learning for entity resolution: A design space exploration. In 2024 IEEE 40th International Conference on Data Engineering (ICDE), pp. 3696–3709. IEEE, 2024.
- Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. Complexity-based prompting
   for multi-step reasoning. In *The Eleventh International Conference on Learning Representations*,
   2023. URL https://openreview.net/forum?id=yflicZHC-19.
  - Yixing Jiang, Jeremy Andrew Irvin, Ji Hun Wang, Muhammad Ahmed Chaudhry, Jonathan H Chen, and Andrew Y. Ng. Many-shot in-context learning in multimodal foundation models. In *ICML* 2024 Workshop on In-Context Learning, 2024. URL https://openreview.net/forum? id=j2rKwWXdcz.
- M Lewis. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*, 2019.
- Jianzhe Lin, Maurice Diesendruck, Liang Du, and Robin Abraham. Batchprompt: Accomplish
   more with less. In *The Twelfth International Conference on Learning Representations*, 2024.
   URL https://openreview.net/forum?id=Agyicd577r.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173, 2024. doi: 10.1162/tacl\_a\_00638. URL https://aclanthology.org/2024.tacl-1.9.
- Ggaliwango Marvin, Nakayiza Hellen, Daudi Jjingo, and Joyce Nakatumba-Nabende. Prompt engineering in large language models. In *International conference on data intelligence and cognitive informatics*, pp. 387–402. Springer, 2023.
- Sewon Min, Xinxi Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke
   Zettlemoyer. Rethinking the role of demonstrations: What makes in-context learning work? In
   Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (eds.), Proceedings of the 2022 Conference
   *on Empirical Methods in Natural Language Processing*, pp. 11048–11064, Abu Dhabi, United
   Arab Emirates, December 2022. Association for Computational Linguistics. doi: 10.18653/v1/
   2022.emnlp-main.759. URL https://aclanthology.org/2022.emnlp-main.759.
- Saeed Moayedpour, Alejandro Corrochano-Navarro, Faryad Sahneh, Alexander Koetter, Jiří
   Vymětal, Lorenzo Kogler Anele, Pablo Mas, Yasser Jangjoo, Sizhen Li, Michael Bailey, Marc
   Bianciotto, Hans Matter, Christoph Grebner, Gerhard Hessler, Ziv Bar-Joseph, and Sven Jager.
   Many-shot in-context learning for molecular inverse design. In *ICML 2024 AI for Science Workshop*, 2024. URL https://openreview.net/forum?id=vJMp5Ms0Ba.

625

626

627

641

Tsendsuren Munkhdalai, Manaal Faruqui, and Siddharth Gopal. Leave no context behind: Efficient infinite context transformers with infini-attention. *arXiv preprint arXiv:2404.07143*, 2024.

Arkil Patel, Satwik Bhattamishra, and Navin Goyal. Are NLP models really able to solve simple math word problems? In Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tur, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou (eds.), *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 2080–2094, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.168. URL https://aclanthology.org/2021.naacl-main.168.

- Keqin Peng, Liang Ding, Yancheng Yuan, Xuebo Liu, Min Zhang, Yuanxin Ouyang, and Dacheng
  Tao. Revisiting demonstration selection strategies in in-context learning. In Lun-Wei Ku, Andre
  Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 9090–9101, Bangkok, Thailand,
  August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.492.
  URL https://aclanthology.org/2024.acl-long.492.
- Adam Poliak. A survey on recognizing textual entailment as an NLP evaluation. In Steffen
  Eger, Yang Gao, Maxime Peyrard, Wei Zhao, and Eduard Hovy (eds.), *Proceedings of the First Workshop on Evaluation and Comparison of NLP Systems*, pp. 92–109, Online, November
  2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.eval4nlp-1.10. URL
  https://aclanthology.org/2020.eval4nlp-1.10.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi
  Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text
  transformer. *Journal of machine learning research*, 21(140):1–67, 2020.
- 618
   619
   619
   619
   619
   620
   620
   621
   621
   621
   622
   622
   623
   624
   625
   626
   626
   627
   628
   629
   629
   620
   620
   620
   621
   621
   622
   622
   623
   624
   625
   626
   626
   627
   628
   629
   629
   620
   620
   620
   621
   622
   622
   623
   624
   625
   625
   626
   626
   627
   628
   628
   629
   629
   620
   620
   621
   622
   622
   623
   624
   625
   625
   626
   626
   627
   628
   629
   629
   629
   620
   621
   622
   622
   623
   624
   625
   625
   626
   627
   628
   629
   629
   629
   620
   620
   621
   622
   622
   623
   623
   624
   625
   625
   626
   626
   627
   628
   629
   629
   629
   629
   620
   620
   621
   622
   622
   623
   623
   624
   625
   625
   626
   627
   628
   629
   629
   629
   629
   620
  - Pranab Sahoo, Ayush Kumar Singh, Sriparna Saha, Vinija Jain, Samrat Mondal, and Aman Chadha. A systematic survey of prompt engineering in large language models: Techniques and applications. arXiv preprint arXiv:2402.07927, 2024.
- Yuanfeng Song, Yuanqin He, Xuefang Zhao, Hanlin Gu, Di Jiang, Haijun Yang, Lixin Fan, and
   Qiang Yang. A communication theory perspective on prompting engineering methods for large
   language models, 2023. URL https://arxiv.org/abs/2310.18358.
- Yi Su, Yunpeng Tai, Yixin Ji, Juntao Li, Yan Bowen, and Min Zhang. Demonstration augmentation for zero-shot in-context learning. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of the Association for Computational Linguistics ACL 2024*, pp. 14232–14244, Bangkok, Thailand and virtual meeting, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.846. URL https://aclanthology.org/2024. findings-acl.846.
- Kuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=1PL1NIMMrw.
- Zhiguo Wang, Wael Hamza, and Radu Florian. Bilateral multi-perspective matching for natural language sentences. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pp. 4144–4150, 2017. doi: 10.24963/ijcai.2017/579. URL https://doi.org/10.24963/ijcai.2017/579.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
   Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. Advances in neural information processing systems, 35:24824–24837, 2022.

648 649 650	Zhiyong Wu, Yaoxiang Wang, Jiacheng Ye, and Lingpeng Kong. Self-adaptive in-context learning An information compression perspective for in-context example selection and ordering. In Ann Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), <i>Proceedings of the 61st Annual Mee</i>	g: 1a 1 <i>t</i> -
651 652 653	ing of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 1423–1430 Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023 acl-long.79. URL https://aclanthology.org/2023.acl-long.79.	5, 3.
654 655 656 657 658 659 660	Zhao Yang, Yuanzhe Zhang, Dianbo Sui, Cao Liu, Jun Zhao, and Kang Liu. Representative demonstration selection for in-context learning with two-stage determinantal point process. I Houda Bouamor, Juan Pino, and Kalika Bali (eds.), <i>Proceedings of the 2023 Conference on En pirical Methods in Natural Language Processing</i> , pp. 5443–5456, Singapore, December 2022, Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.331. UR https://aclanthology.org/2023.emnlp-main.331.	/e [n 1- 3. L
661 662 663	Haochen Zhang, Yuyang Dong, Chuan Xiao, and Masafumi Oyamada. Large language models a data preprocessors, 2023. URL https://arxiv.org/abs/2308.16361.	ıs
664 665 666 667 668	Yiming Zhang, Shi Feng, and Chenhao Tan. Active example selection for in-context learning. If Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (eds.), <i>Proceedings of the 2022 Conference</i> <i>on Empirical Methods in Natural Language Processing</i> , pp. 9134–9148, Abu Dhabi, United Ara Emirates, December 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022 emnlp-main.622. URL https://aclanthology.org/2022.emnlp-main.622.	n ce 1b 2.
670 671 672 673	Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, an Jimmy Ba. Large language models are human-level prompt engineers. In <i>The Eleventh International Conference on Learning Representations</i> , 2023. URL https://openreview.netforum?id=92gvk82DE	ιd α- /
U/		
675 676 677 678	A AUTO-DEMO PROMPT DESIGN	
675 676 677 678 679	A AUTO-DEMO PROMPT DESIGN Table 1: The auto-demo prompt design for BoolQ dataset	
675 676 677 678 679 680 681 682	A AUTO-DEMO PROMPT DESIGN Table 1: The auto-demo prompt design for BoolQ dataset Auto-Demo Prompt for BoolQ Dataset	
675 676 677 678 679 680 681 682 683 683 684 685 686 686 687	<ul> <li>A AUTO-DEMO PROMPT DESIGN</li> <li>Table 1: The auto-demo prompt design for BoolQ dataset</li> <li>Auto-Demo Prompt for BoolQ Dataset</li> <li>Instruction You are a professional NLP expert at Question Answering annotation. Please generate labels given instructions. You will be given [BATCH-SIZE] passages with questions each time, as input. Each input includes a 'passage' and a 'question' about the passage. Your goal is to determine whether the answer to the question is yes or no and classify, as below:</li> </ul>	
675 676 677 678 679 680 681 682 683 684 685 684 685 686 687 688 689 690 691	A AUTO-DEMO PROMPT DESIGN Table 1: The auto-demo prompt design for BoolQ dataset Auto-Demo Prompt for BoolQ Dataset Instruction You are a professional NLP expert at Question Answering annotation. Please generate labels given instructions. You will be given [BATCH-SIZE] passages with questions each time, as input. Each input includes a 'passage' and a 'question' about the passage. Your goal is to determine whether the answer to the question is yes or no and classify, as below: Possible Answer: [class 0]: if the answer is 'No' [class 1]: if the answer is 'Yes'.	
675 676 677 678 679 680 681 682 683 684 685 684 685 686 687 688 689 690 691 692 693	A AUTO-DEMO PROMPT DESIGN Table 1: The auto-demo prompt design for BoolQ dataset Auto-Demo Prompt for BoolQ Dataset Instruction You are a professional NLP expert at Question Answering annotation. Please generate labels given instructions. You will be given [BATCH-SIZE] passages with questions each time, as input. Each input includes a 'passage' and a 'question' about the passage. Your goal is to determine whether the answer to the question is yes or no and classify, as below: Possible Answer: [class 0]: if the answer is 'No' [class 1]: if the answer is 'Yes'. You will be given [BATCH-SIZE] inputs each time.	
675 676 677 678 679 680 681 682 683 684 685 686 685 686 687 688 689 690 691 692 693 694 695	A AUTO-DEMO PROMPT DESIGN Table 1: The auto-demo prompt design for BoolQ dataset Auto-Demo Prompt for BoolQ Dataset Instruction You are a professional NLP expert at Question Answering annotation. Please generate labels given instructions. You will be given [BATCH-SIZE] passages with questions each time, as input. Each input includes a 'passage' and a 'question' about the passage. Your goal is to determine whether the answer to the question is yes or no and classify, as below: Possible Answer: [class 0]: if the answer is 'No' [class 1]: if the answer is 'Yes'. You will be given [BATCH-SIZE] inputs each time. ====================================	
675 676 677 678 679 680 681 682 683 684 685 686 685 686 687 690 691 692 693 694 695 696 697 698	A AUTO-DEMO PROMPT DESIGN         Table 1: The auto-demo prompt design for BoolQ dataset         Auto-Demo Prompt for BoolQ Dataset         Instruction You are a professional NLP expert at Question Answering annotation. Please generate labels given instructions. You will be given [BATCH-SIZE] passages with questions each time, as input. Each input includes a 'passage' and a 'question' about the passage. Your goal is to determine whether the answer to the question is yes or no and classify, as below:         Possible Answer:         [class 0]: if the answer is 'No'         [class 1]: if the answer is 'Yes'.         You will be given [BATCH-SIZE] inputs each time.         ====================================	

	Table 2: The auto-demo prompt design for GSM8K dataset
A	uto-Demo Prompt for GSM8K Dataset
_	
In	struction: You will be given [BATCH-SIZE] math problems. These problems
tw	een 2 and 8 steps to solve, and solutions primarily involve performing a sequence
m	entary calculations using basic arithmetic operations to reach the final answer.
Tl	he answer is [numeric result]
Y	ou will be given [BATCH-SIZE] inputs each time.
== R/	======================================
of	the examples below.
	1
{I	nput 1: xxxxx, Reasoning: xxxxx, Answer: The answer is [number]}
{I	nput 2: xxxxx, Reasoning: xxxxx, Answer: The answer is [number]}
== D1	======================================
ГI Pl	ease ensure the final sentence is "The answer is xxx" where the answer should be a
1 I Pl	ease make sure to generate [BATCH-SIZE] labels
	ease make sure to generate [DATEN SIZE] labers.
	Table 3: The auto-demo prompt design for OOP dataset
	Table 3: The auto-demo prompt design for QQP dataset
A	Table 3: The auto-demo prompt design for QQP dataset uto-Demo Prompt for QQP Dataset
A	Table 3: The auto-demo prompt design for QQP dataset         uto-Demo Prompt for QQP Dataset         Instruction: You are a professional NLP expert at duplicate question detection. You
Au In to	Table 3: The auto-demo prompt design for QQP dataset <b>uto-Demo Prompt for QQP Dataset</b> <b>istruction:</b> You are a professional NLP expert at duplicate question detection. Yo determine whether two questions are duplicates of each other.
A In to	Table 3: The auto-demo prompt design for QQP dataset         uto-Demo Prompt for QQP Dataset         Instruction: You are a professional NLP expert at duplicate question detection. Yo determine whether two questions are duplicates of each other.
A In to	Table 3: The auto-demo prompt design for QQP dataset         uto-Demo Prompt for QQP Dataset         estruction: You are a professional NLP expert at duplicate question detection. Yo determine whether two questions are duplicates of each other.         pssible Answer:
A In to [c	Table 3: The auto-demo prompt design for QQP dataset <b>uto-Demo Prompt for QQP Dataset</b> <b>estruction:</b> You are a professional NLP expert at duplicate question detection. Yo determine whether two questions are duplicates of each other. <b>pssible Answer:</b> lass 1]: if they have the same meaning (semantically equivalent).
A In to [c [c	Table 3: The auto-demo prompt design for QQP dataset <b>uto-Demo Prompt for QQP Dataset</b> <b>estruction:</b> You are a professional NLP expert at duplicate question detection. Yo determine whether two questions are duplicates of each other. <b>pssible Answer:</b> lass 1]: if they have the same meaning (semantically equivalent). lass 0]: if they do NOT have the same meaning.
A In to [c [c Y	Table 3: The auto-demo prompt design for QQP dataset <b>uto-Demo Prompt for QQP Dataset</b> <b>estruction:</b> You are a professional NLP expert at duplicate question detection. Yo determine whether two questions are duplicates of each other. <b>pssible Answer:</b> lass 1]: if they have the same meaning (semantically equivalent). lass 0]: if they do NOT have the same meaning. pu will be given [BATCH-SIZE] question pairs each time.
And In to [c [c Yo	Table 3: The auto-demo prompt design for QQP dataset <b>uto-Demo Prompt for QQP Dataset</b> <b>istruction:</b> You are a professional NLP expert at duplicate question detection. You determine whether two questions are duplicates of each other. <b>issible Answer:</b> lass 1]: if they have the same meaning (semantically equivalent). lass 0]: if they do NOT have the same meaning. bu will be given [BATCH-SIZE] question pairs each time.
$\mathbf{A}_{\mathbf{I}}$ $\mathbf{I}_{\mathbf{I}}$ $\mathbf{I}_{\mathbf{C}}$	Table 3: The auto-demo prompt design for QQP dataset <b>uto-Demo Prompt for QQP Dataset</b> <b>estruction:</b> You are a professional NLP expert at duplicate question detection. You determine whether two questions are duplicates of each other. <b>possible Answer:</b> lass 1]: if they have the same meaning (semantically equivalent). lass 0]: if they do NOT have the same meaning. pu will be given [BATCH-SIZE] question pairs each time. 
$- \mathbf{I} \mathbf{n}$	Table 3: The auto-demo prompt design for QQP dataset <b>uto-Demo Prompt for QQP Dataset</b> <b>estruction:</b> You are a professional NLP expert at duplicate question detection. You determine whether two questions are duplicates of each other. <b>possible Answer:</b> lass 1]: if they have the same meaning (semantically equivalent). lass 0]: if they do NOT have the same meaning. bu will be given [BATCH-SIZE] question pairs each time. <b>EXECUTE:</b> <b>equirement:</b> Repeat the input data and generate the answer following the form amples below.
$\mathbf{A}_{\mathbf{I}}$ $\mathbf{I}_{\mathbf{I}}$ $\mathbf{I}_{\mathbf{C}}$	Table 3: The auto-demo prompt design for QQP dataset         uto-Demo Prompt for QQP Dataset         istruction: You are a professional NLP expert at duplicate question detection. You determine whether two questions are duplicates of each other.         pssible Answer:         lass 1]: if they have the same meaning (semantically equivalent).         lass 0]: if they do NOT have the same meaning.         pu will be given [BATCH-SIZE] question pairs each time.         equirement: Repeat the input data and generate the answer following the form amples below.
$\mathbf{A}_{\mathbf{I}}$ $\mathbf{I}_{\mathbf{I}}$ $\mathbf{I}_{\mathbf{C}}$	Table 3: The auto-demo prompt design for QQP dataset         uto-Demo Prompt for QQP Dataset         istruction: You are a professional NLP expert at duplicate question detection. You determine whether two questions are duplicates of each other.         pssible Answer:         lass 1]: if they have the same meaning (semantically equivalent).         lass 0]: if they do NOT have the same meaning.         pu will be given [BATCH-SIZE] question pairs each time.         equirement: Repeat the input data and generate the answer following the form amples below.         Question pair 0: (Question1: xxxxx; Question2: xxxxx), Answer: [class X]}
$- \mathbf{A} \mathbf{i}$ $- \mathbf{I} \mathbf{n}$ $\mathbf{f} \mathbf{C}$ $- \mathbf{I} \mathbf{n}$	Table 3: The auto-demo prompt design for QQP dataset         uto-Demo Prompt for QQP Dataset         estruction: You are a professional NLP expert at duplicate question detection. You determine whether two questions are duplicates of each other.         pssible Answer:         lass 1]: if they have the same meaning (semantically equivalent).         lass 0]: if they do NOT have the same meaning.         pu will be given [BATCH-SIZE] question pairs each time.         equirement: Repeat the input data and generate the answer following the form tamples below.         Question pair 0: (Question1: xxxxx; Question2: xxxxx), Answer: [class X]}
$\mathbf{A}_{\mathbf{I}}$ $\mathbf{I}_{\mathbf{I}}$	Table 3: The auto-demo prompt design for QQP dataset         uto-Demo Prompt for QQP Dataset         estruction: You are a professional NLP expert at duplicate question detection. You determine whether two questions are duplicates of each other.         pssible Answer:         lass 1]: if they have the same meaning (semantically equivalent).         lass 0]: if they do NOT have the same meaning.         pu will be given [BATCH-SIZE] question pairs each time.         equirement: Repeat the input data and generate the answer following the form tamples below.         Question pair 0: (Question1: xxxx; Question2: xxxxx), Answer: [class X]}         question pair 1: (Question1: xxxx; Question2: xxxxx), Answer: [class X]}
	Table 3: The auto-demo prompt design for QQP dataset uto-Demo Prompt for QQP Dataset struction: You are a professional NLP expert at duplicate question detection. Yo determine whether two questions are duplicates of each other. possible Answer: lass 1]: if they have the same meaning (semantically equivalent). lass 0]: if they do NOT have the same meaning. pu will be given [BATCH-SIZE] question pairs each time. ====================================

A	uto-Demo Prompt for RTE Dataset
In w ea	<b>Astruction:</b> You are a professional NLP expert at sentence pair relationship annotation. You ill be given [BATCH-SIZE] sentence pairs from the Textual Entailment Recognition data tech time, as input. Each data includes a sentence pair, "Premise" and "Hypothesis". You are a single to classify the sentence pair into two classes as below:
Po [c [c	ossible Answer: lass 1]: the given Hypothesis and Premise are logical and following (entailment) to each lass 0]: the given Hypothesis and Premise are NOT following (entailment) to each other
Y	ou will be given [BATCH-SIZE] sentence pairs each time.
R ex	<b>equirement:</b> Repeat the input data and generate the answer following the format of t camples below.
{ <b>!</b> { <b>!</b>	Sentence pair 0: Premise: xxxxx, Hypothesis: xxxxx, Label: [class X]} Sentence pair 1: Premise: xxxxx, Hypothesis: xxxxx, Label: [class X]}
== P1	ease make sure each generated label is in the format of [class X].
	Table 5: The auto-demo prompt design for SVAMP dataset
A	Table 5: The auto-demo prompt design for SVAMP dataset uto-Demo Prompt for SVAMP Dataset
A In qu	Table 5: The auto-demo prompt design for SVAMP dataset uto-Demo Prompt for SVAMP Dataset astruction: You will be given [BATCH-SIZE] math problems. Each one has a body and restion, please read them and give the equation and answer.
A In qu Ya	Table 5: The auto-demo prompt design for SVAMP dataset         uto-Demo Prompt for SVAMP Dataset         istruction: You will be given [BATCH-SIZE] math problems. Each one has a body and lestion, please read them and give the equation and answer.         ou will be given [BATCH-SIZE] inputs each time.
<b>A</b> <b>In</b> qu Yo <b>R</b> of	Table 5: The auto-demo prompt design for SVAMP dataset uto-Demo Prompt for SVAMP Dataset astruction: You will be given [BATCH-SIZE] math problems. Each one has a body and testion, please read them and give the equation and answer. bu will be given [BATCH-SIZE] inputs each time. 
A' In qu Yo == R of {1	Table 5: The auto-demo prompt design for SVAMP dataset uto-Demo Prompt for SVAMP Dataset istruction: You will be given [BATCH-SIZE] math problems. Each one has a body and restion, please read them and give the equation and answer. ou will be given [BATCH-SIZE] inputs each time. 
A <sup>1</sup> In qu Yo == R of {1 {1 == Pl	Table 5: The auto-demo prompt design for SVAMP dataset uto-Demo Prompt for SVAMP Dataset istruction: You will be given [BATCH-SIZE] math problems. Each one has a body and testion, please read them and give the equation and answer. ou will be given [BATCH-SIZE] inputs each time. 

### **B** EXPERIMENTAL RESULTS

Table 6: Experimental results: "w" denotes the Auto-Batch Prompt, "w/o" denotes the conventional batch prompt, and "bs" refers to the batch size.

Dataset	Model (Method)	bs=1	bs=8	bs=16	bs=32	bs=48
	gpt-4o-mini (w)	0.926	-	0.890	0.887	-
CSM8k	gpt-4o-mini (w/o)	0.913	-	0.853	0.847	-
GSINIOK	gpt-40 (w)	0.953	0.947	-	0.957	-
	gpt-4o (w/o)	0.953	0.947	-	0.927	-
	gpt-4o-mini (w)	0.900	-	0.877	0.857	-
SVAMP	gpt-4o-mini (w/o)	0.887	-	0.817	0.803	-
SVANI	gpt-4o (w)	0.940	0.900	-	0.940	-
	gpt-4o (w/o)	0.920	0.878	-	0.900	-
	gpt-4o-mini (w)	0.870	-	0.859	0.863	0.866
	gpt-4o-mini (w/o)	0.888	-	0.845	0.845	0.837
	gpt-4o-mini (Data Selection, w)	-	-	0.881	0.870	0.895
DTE	gpt-4o-mini (Data Selection, w/o)	-	-	0.874	0.867	0.859
NIL	gpt-40 (w)	0.895	0.895	-	0.899	-
	gpt-4o (w/o)	0.884	0.899	-	0.899	-
	gpt-40 (Data Selection, w)	-	0.899	-	0.903	-
	gpt-40 (Data Selection, w/o)	-	0.899	-	0.884	-
	gpt-4o-mini (w)	0.897	-	0.893	0.897	-
BoolO	gpt-4o-mini (w/o)	0.887	-	0.880	0.870	-
DuoiQ	gpt-40 (w)	0.920	0.917	-	0.920	-
	gpt-4o (w/o)	0.917	0.907	-	0.917	-
	gpt-4o-mini (w)	0.837	-	0.803	0.823	-
	gpt-4o-mini (w/o)	0.837	-	0.807	0.823	-
	gpt-40-mini (Data Selection, w)	-	-	0.820	0.830	-
OOP	gpt-4o-mini (Data Selection, w/o)	-	-	0.817	0.820	-
l	gpt-40 (w)	0.857	0.870	-	0.847	-
	gpt-4o (w/o)	0.853	0.850	-	0.843	-
	gpt-40 (Data Selection, w)	-	-	0.877	0.873	-
	gpt-40 (Data Selection, w/o)	-	-	0.863	0.860	-