

# ParaICL: Towards Parallel In-Context Learning

Xingxuan Li<sup>1,2\*</sup> Xuan-Phi Nguyen<sup>3\*</sup> Shafiq Joty<sup>2,3</sup> Lidong Bing<sup>1\*†</sup>

<sup>1</sup>Shanda AI Research Institute <sup>2</sup>Nanyang Technological University <sup>3</sup>Salesforce Research  
{xingxuan.li, lidong.bing}@shanda.com {nguyenxu002, srjoty}@ntu.edu.sg

## Abstract

Large language models (LLMs) have become the norm in natural language processing (NLP), excelling in few-shot in-context learning (ICL) with their remarkable abilities. Nonetheless, the success of ICL largely hinges on the choice of few-shot demonstration examples, making the selection process increasingly crucial. Existing methods have delved into optimizing the quantity and semantic similarity of these examples to improve ICL performances. However, our preliminary experiments indicate that the effectiveness of ICL is limited by the length of the input context. Moreover, varying combinations of few-shot demonstration examples can significantly boost accuracy across different test samples. To address this, we propose a novel method named parallel in-context learning (ParaICL) that effectively utilizes all demonstration examples without exceeding the manageable input context length. ParaICL employs parallel batching to distribute demonstration examples into different batches according to the semantic similarities of the questions in the demonstrations to the test question. It then computes normalized batch semantic scores for each batch. A weighted average semantic objective, constrained by adaptive plausibility, is applied to select the most appropriate tokens. Through extensive experiments, we validate the effectiveness of ParaICL and conduct ablation studies to underscore its design rationale. We further demonstrate that ParaICL can seamlessly integrate with existing methods. Our code is available at <https://github.com/xingxuanli/paraicl.git>.

## 1 Introduction

In recent years, scaling up the parameters of generative language models has significantly enhanced

\*Work done while at DAMO Academy, Alibaba Group. Xingxuan Li was under the Joint Ph.D. Program between Alibaba and Nanyang Technological University.

†Corresponding author.




Figure 1: Results of Mistral-7B-Instruct-v0.2 on 100 test samples from GSM8K and WinoGrande using different numbers of few-shot demonstration examples. Increasing the number of demonstration examples does not necessarily improve the performance consistently.

their language generation capabilities (Radford et al., 2019; Brown et al., 2020; OpenAI, 2023). Large language models (LLM) have demonstrated their adeptness across a wide range of tasks via few-shot in-context learning (ICL) (Cheng et al., 2023; Zhao et al., 2023; Li et al., 2024). In few-shot ICL, models are expected to generate outputs directly based on a sequence of given examples without any parameter modifications, making it the most efficient method for adapting to new tasks.

However, the effectiveness of ICL is notably influenced by the few-shot demonstration examples used (Chen et al., 2023). Various methods have been developed to select the most effective few-shot demonstration examples for ICL. Hao et al. (2022a) demonstrated that scaling up the number of demonstration examples can improve the ICL performance. Rubin et al. (2022) and Liu et al. (2022) proposed to utilize the most relevant examples for each test sample during inference. Nevertheless, existing few-shot ICL methods mostly focus on either increasing the number of demonstration examples or selecting a few that are similar to the test samples to improve performance.

We have designed two preliminary experiments to identify key factors for effectively using demon-




Figure 2: Results of Llama-2-7B-Chat on 100 WinoGrande test samples using different combinations of 10-shot demonstration examples. Different combinations improve the model’s accuracy on various test samples.

stration examples. Firstly, we increase the number of demonstration examples. The results, as shown in Figure 1, reveal that the performance of Mistral-7B-Instruct-v0.2 on GSM8K and WinoGrande does not consistently improve with more examples. This is partly because longer input lengths, resulting from more examples, can lead to suboptimal results in LLMs (Liu et al., 2023; Li et al., 2023b). Therefore, controlling the input context length is essential in ICL, making the number of few-shot demonstration examples a critical factor.

Secondly, we focus on the selection of demonstration examples. We experiment with Llama-2-7B-Chat on 100 test samples from WinoGrande with 32 different 10-shot example combinations. As shown in Figure 2, we notice that varying combinations lead to different accuracy improvements. These combinations enabled the model to accurately answer 80% of the test samples, a significant increase compared to the 50.9% average accuracy of each individual 10-shot combination. This highlights the fact that varying demonstration examples can enhance the model’s accuracy on different test samples. Therefore, we should leverage all available demonstration examples when possible.

To combine the best of both worlds, we introduce parallel in-context learning (ParaICL). ParaICL effectively utilizes the maximum number of demonstration examples without extending the input context length, thus avoiding the potential reduction in model performance due to larger context sizes. Given a set of question–answer pairs as demonstration examples, ParaICL first assigns them into various batches based on the semantic similarity between the demonstrations’ questions and the sample test question. Consequently, each batch maintains a controlled context length while utilizing all demonstration examples. These batches are then processed by a causal model in parallel to obtain the

next token distribution for each batch. Afterward, a weighted average of these distributions is calculated, considering the semantic relation of each batch to the test question. The final step involves selecting the token with the highest weighted average probability for continued generation.

We conduct extensive experiments across a range of reasoning, natural language inference, and coding tasks to validate the effectiveness of ParaICL. We further demonstrate that ParaICL is compatible with both open- and closed-source causal language models. Our study includes abundant ablation studies and analyses to justify the design of ParaICL and demonstrate how it can be integrated with other ICL methods. In summary, our main contributions are the following: (1) We introduce ParaICL, a simple but effective method that leverages all available demonstration examples while maintaining the input context length manageable. (2) We conduct thorough experiments to prove the effectiveness of our method, along with ablation studies to justify its design. (3) We illustrate how our method can enhance and work in conjunction with other methods.

## 2 Related Work

ICL has surged as a transformative approach in the NLP domain, with its significance evident across various applications, including knowledge grounding (Zhao et al., 2023; Li et al., 2024), code generation (Li et al., 2023b), and other industrial applications (Cheng et al., 2023; Chen et al., 2024). Despite its promise, challenges such as sensitivity to the prompts and context window length limitations, highlight areas for further research and development in making ICL more robust and versatile.

As aforementioned, various studies have underscored that ICL exhibits significant sensitivity to the quality of prompts, particularly concerning the demonstration examples provided (Li and Qiu, 2023a; Gupta et al., 2023; Raventos et al., 2023). Chen et al. (2023) discovered that increasing the number of examples only leads to slight improvements. Liu et al. (2022) proposed to select examples that have the highest semantic similarity to the test question. Contrarily, Levy et al. (2023) found that leveraging a diverse set of demonstration examples could improve in-context compositional generalization. Qin et al. (2023) proposed iterative demonstration selection, which considers both the diversity and similarity dimensions of ICL demon-

stration selection for LLMs. Additionally, Chia et al. (2023) explored the potential of contrastive examples in improving the reasoning capabilities of causal language models. Nevertheless, these methods primarily focus on choosing a subset of demonstration examples from a pool of candidates based on certain perspectives. Furthermore, these methods necessitate a substantial number of candidates to efficiently select the best examples, which constrains their applicability across all situations.

Another line of research is dedicated to utilizing extended context in ICL. Ratner et al. (2023) proposed parallel context window (PCW), a method that alleviates the context window limitations for any open-source LLMs without necessitating additional training. However, PCW overlooked the semantic connections between the demonstration examples and the test samples. As evidenced by Chen et al. (2023), merely increasing the quantity of demonstration examples does not correlate with a substantial improvement in performance. In fact, this strategy might detract from overall performance due to the inclusion of incorrect or misleading content within some of the examples.

The implementation of our ParaICL method is also inspired by the ensemble of models (Ganaie et al., 2022), which use multiple slightly different models to contribute probability “votes” before producing the final answer by (weighted) averaging the votes’ probabilities. Vastly different from the ensemble of models, ParaICL uses the same language model prompted with different sets of in-context exemplars to produce varying vote probabilities before aggregating the output distributions to produce the final answer. As such, ParaICL does not require many models and is efficient to compute.

### 3 Methodology

We first formulate the problem setting for few-shot ICL in Section 3.1. Following this, in Section 3.2, we introduce ParaICL, a novel method that efficiently processes demonstration examples in batches and effectively utilizes a novel parallel semantic decoding strategy for generation. Subsequently in Section 3.3, we introduce the weighted average semantic objective alongside the adaptive plausibility constraint. Finally we define the parallel semantic decoding method, which employs the weighted average semantic objective subject to the plausibility constraint. A demonstration of ParaICL can be found in Figure 3.

#### 3.1 Few-Shot In-Context Learning

Few-shot ICL focuses on understanding and executing tasks with a set of demonstration examples (Brown et al., 2020). It leverages a number of selected examples, known as “shots”, to quickly adapt to new tasks. Specifically, within the framework of  $k$ -shot ICL, the model is provided with  $k$  demonstration examples, represented as  $\mathcal{D} = \{(x_1, y_1), \dots, (x_k, y_k)\}$ , incorporated into the input prompt for context.

#### 3.2 Parallel Batching

The preliminary experiments in Section 1 demonstrate the importance of employing varied combinations of demonstration examples without extending the length of the input context. Consequently, we proceed by organizing the demonstration examples  $\mathcal{D} = \{(x_1, y_1), \dots, (x_k, y_k)\}$  into batches. Previous studies have shown that the selection of semantically significant demonstration examples enhances ICL performances (Liu et al., 2022; Luo et al., 2023). Therefore, for each test question  $\hat{x}_i \in \hat{\mathcal{D}} = \{(\hat{x}_1, \hat{y}_1), \dots, (\hat{x}_n, \hat{y}_n)\}$ , we first sequence the demonstration examples in  $\mathcal{D}$  by their question semantic similarities to the test question  $\hat{x}_i$ . This results in an ordered sequence  $\mathcal{D}_{sorted}^i = \{(x_{(1)}^i, y_{(1)}^i), \dots, (x_{(k)}^i, y_{(k)}^i)\}$ , where the similarity function  $f_{sim}$  determines the order based on the cosine similarity of the input embeddings:

$$f_{sim}(x_{(1)}^i, \hat{x}_i) \geq \dots \geq f_{sim}(x_{(k)}^i, \hat{x}_i). \quad (1)$$

$f_{sim}$  is formulated as below,

$$f_{sim}(t_1, t_2) = \frac{f_{emb}(t_1) \cdot f_{emb}(t_2)}{\|f_{emb}(t_1)\| \|f_{emb}(t_2)\|}, \quad (2)$$

where  $t_1$  and  $t_2$  are the input texts, and  $f_{emb}(\cdot)$  is a model to compute the sentence semantic embedding of the input text.

With  $m$  defined as the divisor of  $k$  that indicates the number of examples per batch, we form  $s$  batches, where  $s = \frac{k}{m}$ . These sorted demonstration examples  $\mathcal{D}_{sorted}^i$  are then divided into  $s$  parallel batches, denoted as  $\mathcal{B}^i = \{\mathbf{b}_1^i, \dots, \mathbf{b}_z^i, \dots, \mathbf{b}_s^i\}$ , with each  $\mathbf{b}_z^i$  including demonstration examples  $\{(x_{((z-1)m+1)}^i, y_{((z-1)m+1)}^i), \dots, (x_{(zm)}^i, y_{(zm)}^i)\}$ , where  $1 \leq z \leq s$ .

For each batch, we calculate normalized batch similarity scores  $\mathcal{O}^i = \{o_1^i, \dots, o_z^i, \dots, o_s^i\}$ , where each  $o_z^i$  represents the semantic similarity of the




Figure 3: Our proposed parallel in-context learning (ParaICL) method. Colored squares with **black** borders denote demonstration samples. Squares filled in **grey** with matching borders denote test sample  $\hat{x}_i$ .

batch’s demonstration questions to the test question, determined by:

$$o_z^i = \frac{\sum_{r=1}^m f_{sim}(x_{((z-1)m+r)}^i, \hat{x}_i)}{\sum_{z=1}^s \sum_{r=1}^m f_{sim}(x_{((z-1)m+r)}^i, \hat{x}_i)}. \quad (3)$$

These scores measure the semantic similarity between the demonstration examples in each batch and the test question.

Finally, we compile the input prompts for each batch by incorporating the test question, creating  $\mathcal{U}^i = \{\mathbf{u}_1^i, \dots, \mathbf{u}_z^i, \dots, \mathbf{u}_s^i\}$ , where  $\mathbf{u}_z^i = \{\mathbf{b}_z^i, \hat{x}_i\}$ .

### 3.3 Parallel Semantic Decoding

We first define a generative language model  $f_{lm}(\cdot)$ . For given inputs  $\mathbf{u}_{prev}^i$ , the model generates a continuation  $\mathbf{u}_{cont}^i$  according to the formula:

$$f_{lm}(\mathbf{u}_{cont}^i | \mathbf{u}_{prev}^i) = \prod_{j=1}^{q^i} f_{lm}(u_j^i | u_{<j}^i, \mathbf{u}_{prom}^i), \quad (4)$$

where  $u_j^i$  is a generated token in  $\mathbf{u}_{cont}^i$ ,  $q^i$  represents the total number of generated tokens, and  $\mathbf{u}_{prom}^i$  is the provided input prompt.

**Weighted average semantic objective** As demonstration examples have varying semantic contributions to the test sample, we propose the weighted average semantic (WAS) objective. The WAS objective for test sample  $\hat{x}_i$  is defined as:

$$\mathcal{L}_{WAS}(\mathbf{u}_{cont}^i, \mathcal{O}^i, \mathcal{U}^i) = \sum_{z=1}^s o_z^i \cdot f_{lm}(\mathbf{u}_{cont}^i | \mathbf{u}_z^i). \quad (5)$$

The WAS objective rewards batches with demonstration examples that exhibit greater semantic similarity to the test question, while simultaneously

leveraging all examples to enrich the generation process comprehensively. This advantage makes ParaICL superior to PCW, which arranges examples randomly. However, certain batches may contain noise that could adversely affect the performance. To address this problem, we adopt the adaptive plausibility constraint from the contrastive decoding method (Li et al., 2023a).

**Adaptive plausibility constraint** The adaptive plausibility constraint  $\mathcal{V}_{head}$  leverages the confidence level in the foremost batch to mitigate the impact of potentially less relevant demonstration batches. Explicitly, the adaptive plausibility constraint for test sample  $\hat{x}_i$  is defined as:

$$\begin{aligned} \mathcal{V}_{head}(u_{<j}^i) &= \{u_j^i \in \mathcal{V} : f_{lm}(u_j^i | u_{<j}^i, \mathbf{u}_1^i) \\ &\geq \alpha \max_{w^i} f_{lm}(w^i | u_{<j}^i, \mathbf{u}_1^i)\}, \end{aligned} \quad (6)$$

where  $\mathcal{V}$  represents the model’s vocabulary,  $\mathbf{u}_1^i$  is the first batch (*i.e.*, the most semantically aligned with the test question  $\hat{x}_i$ ), and  $\alpha$  is a hyperparameter between  $[0, 1]$ <sup>1</sup>. A higher  $\alpha$  value signifies a preference for tokens with higher generation probabilities, whereas smaller  $\alpha$  allows tokens of lower probabilities to be generated.

**Final method** The final parallel semantic decoding method combines the WAS objective with the adaptive plausibility constraint to optimize the generation process:

$$\begin{aligned} \max_{\mathbf{u}_{cont}^i} \mathcal{L}_{WAS}(\mathbf{u}_{cont}^i, \mathcal{O}^i, \mathcal{U}^i) \\ \text{subject to } u_j^i \in \mathcal{V}_{head}(u_{<j}^i), \forall u_j^i \in \mathbf{u}_{cont}^i. \end{aligned} \quad (7)$$

<sup>1</sup>Following the contrastive decoding method (Li et al., 2023a), we set  $\alpha$  as 0.1 for all experiments.



Given the complexity at the sequence level, we simplify the optimization to the token level as follows:

$$\begin{aligned} \mathcal{L}_{\text{WAS}}(\mathbf{u}_{\text{cont}}^i, \mathcal{O}^i, \mathcal{U}^i) \\ &= \sum_{z=1}^s o_z^i \cdot f_{lm}(\mathbf{u}_{\text{cont}}^i | \mathbf{u}_z^i) \\ &= \prod_{j=1}^{q^i} \text{WAS-score}(u_j^i, u_{<j}^i, \mathcal{O}^i, \mathcal{U}^i), \end{aligned} \quad (8)$$

where  $\text{WAS-score}(u_j^i, u_{<j}^i, \mathcal{O}^i, \mathcal{U}^i)$  is the token level score formulated as:

$$\begin{aligned} \text{WAS-score}(u_j^i, u_{<j}^i, \mathcal{O}^i, \mathcal{U}^i) \\ &= \begin{cases} \sum_{z=1}^s o_z^i \cdot f_{lm}(u_j^i | u_{<j}^i, \mathbf{u}_z^i) & \text{if } u_j^i \in \mathcal{V}_{\text{head}}(u_{<j}^i), \\ -\inf & \text{otherwise.} \end{cases} \end{aligned} \quad (9)$$

We first apply plausibility constraints  $\mathcal{V}_{\text{head}}(u_{<j}^i)$  to filter tokens, discarding those that do not reach the required probability threshold within the most semantically pertinent demonstration batch. Subsequently, the surviving tokens are evaluated using the weighted average semantic scores derived from all batches. Consequently, this process allows for the selection of a token that incorporates information from every batch of examples.

## 4 Experiments

### 4.1 Datasets

ParaICL is a versatile approach applicable to a broad range of tasks that can be framed into the few-shot ICL setting. We evaluate ParaICL extensively on several task categories, including reasoning, natural language inference (NLI), and coding. For reasoning tasks, we evaluate on three datasets: (1) **GSM8K** (Cobbe et al., 2021), a mathematical reasoning dataset consisting of a collection of high-quality math word problems. (2) **WinoGrande** (Sakaguchi et al., 2019), a commonsense reasoning dataset of improved complexity beyond the Winograd Schema Challenge benchmark. (3) **ARC** (Clark et al., 2018), a knowledge reasoning dataset consisting of grad-school level science questions in a multiple-choice format. We adopt the challenge set of ARC. For the NLI tasks, we select **Hel-laSwag** (Zellers et al., 2019), a commonsense NLI task that examines a model by predicting the most logical continuation of a described event. In the

coding category, we assess using **MBPP** (Austin et al., 2021), a benchmark consisting of Python programming problems designed to be solvable by entry-level programmers.

### 4.2 Baselines

To provide a more comprehensive overview of where our framework stands, we use the following baselines: (1) **Standard few-shot (Standard)** (Brown et al., 2020): Directly generating the results based on few-shot demonstration examples. (2) **Semantically sorted few-shot (Sorted)** (Chen et al., 2023): Utilizing the same few-shot demonstration examples, this approach organizes the examples by the semantic similarity between each example’s question and the test question. **Sorted+** indicates examples sorted in ascending similarity and **Sorted-** in descending similarity. (3) **Coverage-based selection (CBS)** (Gupta et al., 2023): CBS selects examples from more salient aspects such as reasoning patterns of the test input by using BERTScore-Recall. (4) **Filter-then-search (LENS)** (Li and Qiu, 2023b): LENS utilizes InfoScore to evaluate and progressively selects examples. (5) **Parallel context window (PCW)** (Ratner et al., 2023): Carving a long context into batches, PCW restricts the attention mechanism to apply within each batch. To ensure a fair comparison, we maintain the same number of batches as used in ParaICL. (6) **Structured prompting (SP)** (Hao et al., 2022b): In SP, demonstration examples are separately encoded with well-designed position embeddings and jointly attended by the test example using a rescaled attention mechanism.

### 4.3 Experiment Setup for ParaICL

When setting up experiments for ParaICL, we maintain consistency with the demonstration examples used in baseline methods to ensure a fair comparison. To demonstrate the effectiveness of ParaICL, we conduct experiments with varying numbers of total demonstration examples, specifically 3, 9, and 15, while keeping the count of parallel batches at three<sup>2</sup>. Consequently, this results in 1, 3, and 5 demonstration examples per batch across the three settings. The impact of batch numbers will be elaborated in Section 5.1. Each experimental run is carried out with three random seeds for the selection of demonstration examples, and the results

<sup>2</sup>ParaICL focuses on optimally using any number of available examples, and as such, we do not experiment with hundreds or thousands of examples.

Method	Llama-2-7B-Chat					Mistral-7B-Instruct-v0.2				
	GSM8K	WinoGrande	ARC-C	HellaSwag	MBPP	GSM8K	WinoGrande	ARC-C	HellaSwag	MBPP
3-shot										
Standard	19.9 ± 0.9	52.7 ± 1.1	54.1 ± 1.2	32.9 ± 1.7	19.5 ± 0.6	40.7 ± 1.9	58.0 ± 2.1	70.5 ± 0.6	58.6 ± 1.2	30.2 ± 2.1
Sorted+	20.1 ± 0.8	53.1 ± 1.3	53.9 ± 0.8	33.6 ± 1.1	18.7 ± 1.9	39.2 ± 1.4	58.8 ± 0.7	72.1 ± 1.8	58.9 ± 1.1	31.5 ± 0.8
Sorted-	18.4 ± 1.6	52.3 ± 0.6	54.3 ± 0.9	33.2 ± 0.5	19.6 ± 2.4	40.9 ± 1.2	57.7 ± 1.3	71.2 ± 0.9	58.5 ± 0.5	30.6 ± 0.9
CBS	17.3 ± 1.2	51.9 ± 1.0	52.8 ± 1.1	31.5 ± 1.6	17.2 ± 0.7	38.1 ± 1.5	56.6 ± 1.1	69.9 ± 1.4	57.7 ± 1.3	30.9 ± 1.4
LENS	16.8 ± 0.7	52.5 ± 0.9	53.4 ± 1.3	32.1 ± 0.7	16.8 ± 1.8	39.5 ± 1.7	55.3 ± 1.4	69.4 ± 1.8	57.1 ± 1.6	29.3 ± 1.3
PCW	17.6 ± 1.1	<u>53.4 ± 0.7</u>	52.2 ± 0.7	33.5 ± 0.8	16.2 ± 1.3	40.5 ± 0.8	58.9 ± 1.2	70.1 ± 1.3	58.2 ± 0.8	31.9 ± 1.7
SP	19.4 ± 1.5	51.6 ± 1.5	51.6 ± 1.9	31.9 ± 1.9	19.1 ± 1.5	37.4 ± 2.5	56.1 ± 1.1	70.6 ± 1.1	57.5 ± 0.9	30.4 ± 0.8
ParaICL	22.1 ± 1.3	52.9 ± 0.5	<u>55.3 ± 0.6</u>	<u>33.7 ± 2.2</u>	<b>20.9 ± 0.6</b>	41.1 ± 1.3	59.2 ± 0.9	71.9 ± 0.6	<b>59.2 ± 1.5</b>	<u>32.6 ± 1.1</u>
9-shot										
Standard	21.9 ± 0.5	50.8 ± 0.7	57.0 ± 1.9	31.7 ± 1.5	18.5 ± 1.0	41.5 ± 1.5	59.3 ± 0.8	70.5 ± 1.5	54.5 ± 1.2	31.9 ± 0.9
Sorted+	22.3 ± 0.9	51.4 ± 1.1	55.3 ± 1.3	32.1 ± 0.8	18.8 ± 3.5	41.2 ± 0.6	58.9 ± 1.2	71.1 ± 1.7	55.6 ± 0.5	30.5 ± 0.6
Sorted-	22.1 ± 0.7	49.9 ± 1.6	55.8 ± 0.5	32.4 ± 2.1	20.1 ± 2.3	40.9 ± 0.5	60.1 ± 1.1	70.8 ± 0.6	55.2 ± 0.9	31.2 ± 1.7
CBS	18.5 ± 1.3	52.1 ± 1.8	56.2 ± 2.1	32.3 ± 0.9	18.7 ± 1.0	38.7 ± 2.5	59.9 ± 1.2	70.2 ± 2.3	51.9 ± 2.9	31.7 ± 1.6
LENS	21.5 ± 1.0	52.3 ± 1.4	57.4 ± 1.3	30.9 ± 2.1	17.4 ± 3.2	39.6 ± 1.1	59.4 ± 2.3	69.1 ± 2.5	54.7 ± 2.2	33.6 ± 1.2
PCW	23.5 ± 1.2	51.7 ± 0.8	54.9 ± 0.7	32.8 ± 1.6	19.6 ± 1.7	41.3 ± 0.9	61.5 ± 1.6	71.3 ± 1.2	52.1 ± 1.4	32.8 ± 1.1
SP	23.4 ± 1.4	53.1 ± 1.5	55.6 ± 0.9	30.8 ± 1.4	17.5 ± 2.9	40.1 ± 2.2	62.7 ± 1.9	69.9 ± 2.4	53.3 ± 3.3	32.4 ± 2.1
ParaICL	<b>25.4 ± 0.8</b>	<b>54.3 ± 0.9</b>	<b>58.1 ± 2.3</b>	<b>33.9 ± 1.7</b>	<u>20.5 ± 0.4</u>	42.8 ± 0.5	<b>63.2 ± 0.9</b>	<b>72.9 ± 1.3</b>	55.8 ± 3.8	<b>34.8 ± 1.3</b>
15-shot										
Standard	21.2 ± 1.0	50.3 ± 0.9	54.8 ± 3.6	29.0 ± 2.2	18.1 ± 1.3	40.4 ± 0.5	61.9 ± 2.7	60.6 ± 4.7	49.9 ± 2.5	32.9 ± 0.6
Sorted+	22.4 ± 0.8	49.8 ± 0.6	55.4 ± 2.1	30.1 ± 1.5	19.3 ± 0.9	41.7 ± 1.6	62.5 ± 1.3	60.3 ± 2.1	52.7 ± 0.9	33.0 ± 0.9
Sorted-	21.8 ± 1.1	50.9 ± 1.7	55.2 ± 1.8	30.6 ± 2.9	18.7 ± 0.6	39.4 ± 0.7	62.1 ± 2.1	63.5 ± 3.4	51.2 ± 1.8	33.2 ± 2.1
CBS	20.1 ± 1.0	50.2 ± 0.8	52.2 ± 3.3	29.7 ± 0.6	19.0 ± 1.9	40.3 ± 0.5	61.2 ± 2.0	61.4 ± 1.3	50.8 ± 1.3	30.4 ± 2.2
LENS	21.3 ± 2.3	48.5 ± 3.5	54.8 ± 2.2	32.1 ± 3.2	18.8 ± 1.6	42.4 ± 2.6	62.2 ± 2.4	66.0 ± 2.2	52.6 ± 1.4	31.7 ± 2.3
PCW	23.8 ± 0.9	49.7 ± 1.4	55.1 ± 0.9	<u>32.9 ± 0.7</u>	19.9 ± 1.1	41.8 ± 0.3	62.7 ± 1.5	<u>66.2 ± 1.3</u>	50.1 ± 0.7	32.8 ± 0.8
SP	22.6 ± 1.4	50.1 ± 1.1	54.9 ± 3.2	31.8 ± 2.3	18.6 ± 1.4	41.5 ± 0.9	60.9 ± 3.1	65.7 ± 2.9	53.1 ± 2.6	32.9 ± 1.4
ParaICL	<u>24.9 ± 0.6</u>	<u>51.0 ± 2.1</u>	<u>56.3 ± 1.3</u>	31.3 ± 0.9	<u>20.6 ± 0.5</u>	<b>42.9 ± 1.1</b>	<b>63.2 ± 3.2</b>	65.8 ± 0.8	<u>53.7 ± 1.0</u>	<u>33.5 ± 1.5</u>

Table 1: Experimental results of Llama-2-7B-Chat and Mistral-7B-Instruct-v0.2 on various reasoning, natural language inference, and coding benchmarks. Underlined indicates the highest scores for each shot group and **bold** indicates overall highest scores.

Method	GSM8K	HellaSwag
Retrieval (5-shot)	22.1	28.3
ParaICL (15-shot)	24.9	31.3

Table 2: ParaICL vs. retrieval-based ICL methods on GSM8K and HellaSwag.

are averaged for reporting. Our main experiments are conducted using open-source models<sup>3</sup>, including Llama-2-7B-Chat and Mistral-7B-Instruct-v0.2. We utilize supervised SimCSE (Gao et al., 2021) with BERT base to compute sentence embeddings.

#### 4.4 Experiment Results

**ParaICL consistently outperforms baseline methods** We present the experimental results of all datasets from Llama-2-7B-Chat and Mistral-7B-Instruct-v0.2 in Table 1. ParaICL consistently outperforms baseline methods on all datasets, demonstrating the effectiveness of our method. On reasoning tasks, the average improvements with Llama-2-7B-Chat are 1.2%, 2.7%, and 2.0% for 3-shot, 9-shot, and 15-shot settings, respectively. Mistral-7B-

Instruct-v0.2 shows improvements of 1.0%, 2.9%, and 3.0% for the same settings, demonstrating ParaICL’s adaptability across different large language models. On commonsense NLI and coding tasks, ParaICL presents an average performance boost of 1.8% and 2.0% for Llama-2-7B-Chat, and 1.9% and 2.0% for Mistral-7B-Instruct-v0.2, respectively. Moreover, ParaICL consistently surpasses CBS and PCW in all settings and on virtually all datasets, underscoring the advantages of utilizing semantic information in our approach. An observation is made regarding the performance similarity between the Sorted+ and Sorted- methods. Sorted+ sometimes outperforms Sorted- and vice versa. This finding is consistent with prior research by Levy et al. (2023) and Liu et al. (2022). With ParaICL, this issue is lessened due to the reduced number of demonstration examples per batch.

**Number of demonstration examples** As shown in Table 1, ParaICL consistently exhibits enhanced performance across multiple datasets under 3-shot, 9-shot, and 15-shot settings. However, the performance gain in the 3-shot setting is less marked compared to the 9-shot and 15-shot settings. This variation is attributed to the fact that the quantity of

<sup>3</sup>Due to cost concern, the validation of ParaICL’s effectiveness on closed-source models is conducted on a more compact dataset in Section 5.4.




Figure 4: Results of Mistral-7B-Instruct-v0.2 on GSM8K using different batches of five-shot demonstration examples.

demonstration examples in each batch sets a limit on ParaICL’s maximum performance. It is also observed that on certain datasets, an increase in demonstration examples can actually degrade performance. For instance, in the HellaSwag dataset, the Standard method’s performance decreases as the number of shots increases: 32.9% for 3-shot, 31.7% for 9-shot, and 29.0% for 15-shot. This decline is likely due to potential label bias within the demonstration examples, a phenomenon noted in various studies (Wang et al., 2023). Additionally, HellaSwag demands commonsense knowledge. Earlier research has indicated that merely increasing the number of demonstrations does not guarantee improved the performance (Yao et al., 2023; Li et al., 2024).

#### ParaICL vs. retrieval-based ICL methods

Retrieval-based ICL methods are designed to select the most relevant demonstration examples from a pool of candidates, in contrast, ParaICL utilizes all available demonstration examples. We compare the performance of ParaICL with retrieval-based ICL methods in a 15-shot setting on the GSM8K and HellaSwag datasets. Retrieval-based ICL methods are restricted to choosing the optimal demonstration examples based on sentence similarity (Liu et al., 2022) from a set of 15 candidates. As shown in Table 2, ParaICL surpasses the retrieval-based method on both GSM8K and HellaSwag datasets. The necessity for a large candidate pool to select from, which retrieval-based ICL methods rely on, hampers their adaptability when only a limited number of demonstration examples are available. ParaICL, however, is capable of efficiently utilizing any quantity of demonstration examples, showcasing its superior versatility.

Method	GSM8K	HellaSwag
w/o S.S. (9-shot)	21.1	32.2
w/ S.S. (9-shot)	25.4	33.9

Table 3: With vs. without semantic sorting on GSM8K and HellaSwag. S.S. stands for semantic sorting.

Method	Llama-2	Mistral
w/ $\mathcal{V}_{head}$ (9-shot)	25.4	42.8
w/o $\mathcal{V}_{head}$ (9-shot)	2.9	3.1

Table 4: Results of ParaICL on GSM8K with and without  $\mathcal{V}_{head}$ .

## 5 Ablation Studies and Analysis

### 5.1 Number of Batches

ParaICL keeps the context length for each batch manageable, allowing for the increasing of batch numbers as hardware capabilities allow. In this experiment, we set the number of demonstration examples in each batch as five, and progressively increase the number of batches. The results of Mistral-7B-Instruct-v0.2 on GSM8K is in Figure 4. With the increment in batch numbers, ParaICL’s improvements tend to converge at five batches. Further increase in batch numbers leads to instability in results. We attribute this to the higher likelihood of incorporating batches that negatively affect token selection. Consequently, we set the batch number to three in our main experiments to ensure stable performance improvements.

### 5.2 Parallel Batching with vs. without Semantic Sorting

The demonstration examples are sequenced by their question semantic similarities to the test question before parallel batching. This ensures that semantically similar examples are batched together, allowing them to share the same batch similarity score during decoding. According to results shown in Table 3, parallel batching with semantic sorting significantly outperforms without semantic sorting. In fact, Parallel batching without semantic sorting even performs worse than the standard approach on GSM8K. This is because, without semantic sorting, a single batch can contain both relevant and irrelevant demonstration examples. This mix of examples reduces the overall decoding probability from the relevant examples in the batch.

Method	Llama-2-13B-Chat	Llama-3.1-70B-Instruct	Mixtral-8x7B-Instruct-v0.1	gpt-3.5-turbo-instruct
Standard (9-shot)	23.1	86.4	58.1	62.0
ParaICL (9-shot)	25.6	87.1	59.2	66.0

Table 5: Experimental results of ParaICL on GSM8K using more open- and closed-source LLMs.

Method	Llama-2	Mistral
Standard (9-shot)	21.9	41.5
ParaICL (9-shot)	25.4	42.8
ParaICL w. C.D. (9-shot)	26.1	43.1
Standard (15-shot)	21.1	40.4
ParaICL (15-shot)	24.9	42.9
ParaICL w. C.D. (15-shot)	25.2	43.5

Table 6: Results of ParaICL with contrastive decoding on GSM8K. C.D. stands for contrastive decoding.

### 5.3 Adaptive Plausibility Constraint

The adaptive plausibility constraint plays a crucial role in our approach, akin to its importance in contrastive decoding as noted by Li et al. (2023a). To assess the impact of this constraint, we carried out an ablation study by eliminating it from our method. The outcomes clearly show a substantial decline in performance, as detailed in Table 4. This observation is consistent with findings from the contrastive decoding paper.

### 5.4 More Open- and Closed-Source LLMs

In this section, we demonstrate that ParaICL is effective in both open- and closed-source LLMs.

We first conduct experiments on GSM8K using three additional open-source models spanning a range of parameter counts, including Llama-2-13B-Chat, Llama-3.1-70B-Instruct, and Mixtral-8x7B-Instruct-v0.1. As shown in Table 5, ParaICL consistently enhances the performances, highlighting its adaptability to open-source models.

Our method necessitates the next token probabilities from the model (*i.e.*, the softmax of the raw scores generated by the final layer) for computing the semantic weighted average. However, closed-source LLMs, due to their proprietary nature, do not make these probabilities available. For instance, OpenAI’s gpt-3.5-turbo-instruct model only allows access to a maximum of five tokens that have top log probabilities, which represents the extent of information available for our use. We employ the same steps as shown in Section 3.3 to execute parallel semantic decoding using the provided log probabilities from OpenAI models. Concerns regarding API costs lead us to limit our experimentation to a randomly chosen set of 50

data points from the GSM8K datasets. Table 5 illustrates that ParaICL enhances performance beyond the standard method, further showcasing its adaptability even with limited information on token distributions. In contrast, PCW is restricted to open-source models as it requires modifications to the attention mask.

### 5.5 Integration with Other Methods

In this section, we demonstrate that ParaICL is compatible with other methods. Specifically, we explore its integration with contrastive decoding (CD) (Li et al., 2023a). Building upon the concept of contrastive objectives introduced by Li et al. (2023a), which leverages the differential signals between larger and smaller language models for decoding, we incorporate contrastive batches into ParaICL. This integration involves calculating the weighted average distributions for both positive and contrastive batches individually, then applying the contrastive objective by subtracting the logarithmic values of these two distributions. This process helps in selecting tokens generated from the positive batches that are least similar to those from the contrastive batches, thus refining the selection for more plausible outcomes. We experiment on GSM8K. We adopt the contrastive chain-of-thought as outlined by Chia et al. (2023), creating a batch consisting of five reasoning failure cases. These cases encompass invalid reasoning, incoherent objects, incoherent language, irrelevant objects, and irrelevant language. The specifics of these demonstration examples are provided in Appendix A.2. The integration of ParaICL with CD has shown to enhance performance on both the Llama-2-7B-Chat and Mistral-7B-Instruct-v0.2 models, particularly in 9- and 15-shot settings, as illustrated in Table 6. This further evidences the flexibility and improved effectiveness of ParaICL when combined with other methodologies.

### 5.6 Majority Voting vs. Standard Average vs. Weighted Average

As aforementioned in Section 3.3, demonstration examples have varying semantic contribution to the test sample. As such, we utilize weighted average



Method	GSM8K	HellaSwag
Majority Voting (9-shot)	20.3	32.7
Standard average (9-shot)	23.6	32.8
Weighted average (9-shot)	25.4	33.9

Table 7: Majority Voting vs. Weighted Average on GSM8K and HellaSwag.

semantic objective during the parallel semantic decoding. We study the effectiveness of the weighted average method by comparing it with the majority voting and standard average methods. According to the results presented in Table 7, utilizing weighted average during parallel semantic decoding substantially surpasses the performance achieved through majority voting and standard average. Notably, the application of majority voting on the GSM8K dataset results in performance that is even worse than the standard method. This can be attributed to irrelevant batches contributing votes that may dominate over the more desired tokens. In contrast, the weighted average method ensures that batches with the highest relevance have the greatest impact on the selection of subsequent tokens, leading to more accurate outcomes.

### 5.7 Performance-Efficiency Tradeoff

ParaICL introduces additional inference time due to the computation of embeddings for both the demonstrations and test examples. Nevertheless, this overhead is minimal when computations are conducted in parallel. Across various datasets, the average time to compute embeddings using SimCSE (Gao et al., 2021) is merely 5.1 seconds.

## 6 Conclusions

In this work, we introduce a novel approach known as parallel in-context learning (ParaICL), designed to enhance the effectiveness of few-shot in-context learning. ParaICL aims to fully leverage all available demonstration examples while keeping within the limits of a manageable input context size. It starts by executing parallel batching, grouping demonstration examples into various batches based on the semantic similarities between the questions in the demonstrations and the test questions. Afterward, for each batch, normalized semantic scores are calculated. The process culminates in the decoding of the final tokens, optimizing a weighted average semantic objective under an adaptive plausibility constraint. ParaICL has been proven to yield consistent performance improvements across

a diverse set of benchmarks and exhibits a high degree of compatibility for integration with other methodologies.

## Limitations

We only study the effectiveness of ParaICL in decoder-only language models. In future works, we aim to extend our evaluation of ParaICL to language models that utilize different structures, including those with encoder-only and encoder-decoder structures. In this paper, we only demonstrate that ParaICL can be easily integrated with contrastive objectives. We would like to explore the integration of ParaICL with other methods, such as retrieval-augmented generation.

## Ethical Impact

We do not foresee any potential ethical issues with our proposed method.

## Acknowledgements

This work was supported by DAMO Academy through DAMO Academy Research Intern Program.

## References

- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and Charles Sutton. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. In *Proceedings of NeurIPS*.
- Hailin Chen, Fangkai Jiao, Xingxuan Li, Chengwei Qin, Mathieu Ravaut, Ruochen Zhao, Caiming Xiong, and Shafiq Joty. 2024. Chatgpt’s one-year anniversary: Are open-source large language models catching up? *arXiv preprint arXiv:2311.16989*.
- Jiuhai Chen, Lichang Chen, Chen Zhu, and Tianyi Zhou. 2023. How many demonstrations do you need for in-context learning? In *Findings of EMNLP*.
- Liyang Cheng, Xingxuan Li, and Lidong Bing. 2023. Is gpt-4 a good data analyst? *arXiv preprint arXiv:2305.15038*.
- Yew Ken Chia, Guizhen Chen, Luu Anh Tuan, Soujanya Poria, and Lidong Bing. 2023. Contrastive chain-of-thought prompting. *arXiv preprint arXiv:2311.09277*.

- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457v1*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- M.A. Ganaie, Minghui Hu, A.K. Malik, M. Tanveer, and P.N. Suganthan. 2022. Ensemble deep learning: A review. *arXiv preprint arXiv:2022.105151*.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. Simcse: Simple contrastive learning of sentence embeddings. In *Proceedings of EMNLP*.
- Shivanshu Gupta, Matt Gardner, and Sameer Singh. 2023. Coverage-based example selection for in-context learning. In *Findings of EMNLP*.
- Yaru Hao, Yutao Sun, Li Dong, Zhixiong Han, Yuxian Gu, and Furu Wei. 2022a. Structured prompting: Scaling in-context learning to 1,000 examples. *arXiv preprint arXiv:2212.06713*.
- Yaru Hao, Yutao Sun, Li Dong, Zhixiong Han, Yuxian Gu, and Furu Wei. 2022b. Structured prompting: Scaling in-context learning to 1,000 examples. *arXiv preprint arXiv:2212.06713*.
- Itay Levy, Ben Bogin, and Jonathan Berant. 2023. Diverse demonstrations improve in-context compositional generalization. In *Proceedings of ACL*.
- Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke Zettlemoyer, and Mike Lewis. 2023a. Contrastive decoding: Open-ended text generation as optimization. *arXiv preprint arXiv:2210.15097*.
- Xiaonan Li and Xipeng Qiu. 2023a. Finding support examples for in-context learning. In *Findings of EMNLP*.
- Xiaonan Li and Xipeng Qiu. 2023b. Finding support examples for in-context learning. In *Findings of EMNLP*.
- Xingxuan Li, Liying Cheng, Qingyu Tan, Hwee Tou Ng, Shafiq Joty, and Lidong Bing. 2023b. Unlocking temporal question answering for large language models using code execution. *arXiv preprint arXiv:2305.15014*.
- Xingxuan Li, Ruochen Zhao, Yew Ken Chia, Bosheng Ding, Shafiq Joty, Soujanya Poria, and Lidong Bing. 2024. Chain-of-knowledge: Grounding large language models via dynamic knowledge adapting over heterogeneous sources. In *Proceedings of ICLR*.
- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2022. What makes good in-context examples for gpt-3? In *Proceedings of ACL*.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023. Lost in the middle: How language models use long contexts. *arXiv preprint arXiv:2307.03172*.
- Man Luo, Xin Xu, Zhuyun Dai, Panupong Pasupat, Mehran Kazemi, Chitta Baral, Vaiva Imbrasaite, and Vincent Y Zhao. 2023. Dr.icl: Demonstration-retrieved in-context learning. *arXiv preprint arXiv:2305.14128*.
- OpenAI. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Chengwei Qin, Aston Zhang, Anirudh Dagar, and Wenming Ye. 2023. In-context learning with iterative demonstration selection. *arXiv preprint arXiv:2310.09881*.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. In *Arxiv*.
- Nir Ratner, Yoav Levine, Yonatan Belinkov, Ori Ram, Inbal Magar, Omri Abend, Ehud Karpas, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. Parallel context windows for large language models. *arXiv preprint arXiv:2212.10947*.
- Allan Raventos, Mansheej Paul, Feng Chen, and Surya Ganguli. 2023. Pretraining task diversity and the emergence of non-bayesian in-context learning for regression. In *Proceedings of NeurIPS*.
- Ohad Rubin, Jonathan Herzig, and Jonathan Berant. 2022. Learning to retrieve prompts for in-context learning. In *Proceedings of NAACL-HLT*.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2019. Winogrande: An adversarial winograd schema challenge at scale. *arXiv preprint arXiv:1907.10641*.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. Self-consistency improves chain of thought reasoning in language models. In *Proceedings of ICLR*.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. React: Synergizing reasoning and acting in language models. In *Proceedings of ICLR*.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*.

Ruochen Zhao, Xingxuan Li, Shafiq Joty, Chengwei Qin, and Lidong Bing. 2023. Verify-and-edit: A knowledge-enhanced chain-of-thought framework. In *Proceedings of ACL*.

## A Appendix

### A.1 Datasets

GSM8K (Cobbe et al., 2021) and HellaSwag (Zellers et al., 2019) are under the MIT License. WinoGrande (Sakaguchi et al., 2019) and MBPP (Austin et al., 2021) are under the Apache-2.0 License. ARC (Clark et al., 2018) is under the GPL-3.0 License.

### A.2 Contrastive demonstration examples for GSM8K

**Invalid reasoning:** **Question:** 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? **Answer:** Janet sells  $16 - 3 - 4 = <<16-3-4=9>>9$  duck eggs a day. She makes  $9 + 2 = \$<<9+2=11>>11$  every day at the farmer’s market. ##### 11

**Incoherent objects:** **Question:** 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? **Answer:** Janet sells  $9 * 2 = <<9*2=18>>18$  duck eggs a day. She makes  $16 - 3 - 4 = \$<<16-3-4=9>>9$  every day at the farmer’s market.##### 18

**Incoherent language:** **Question:** 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? **Answer:** Janet sells  $16 - 3 - 4 = <<16-3-4=9>>9$  duck eggs a day. She makes  $9 - 2 = \$<<9-2=7>>7$  every day at the farmer’s market.##### 7

**Irrelevant objects:** **Question:** 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? **Answer:** Janet sells  $17 - 5 - 6 = <<17-5-6=6>>6$  duck eggs a day. She sells  $6 * 3 = \$<<6*3=18>>18$  every day at the duck eggs.##### 18

**Irrelevant language:** **Question:** 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? **Answer:** Janet sells  $16 - 3 - 4 = <<16-3-4=9>>9$  duck eggs a day. She wants her hair to be  $9 * 2 = \$<<9*2=18>>18$  inches long when she cuts it.##### 18