

# How Does the Textual Information Affect the Retrieval of Multimodal In-Context Learning?

Yang Luo, Zangwei Zheng, Zirui Zhu, Yang You  
School of Computing, National University of Singapore  
{yangluo, zangwei, zirui, youy}@comp.nus.edu.sg

## Abstract

The increase in parameter size of multimodal large language models (MLLMs) introduces significant capabilities, particularly multimodal in-context learning, where MLLMs enhance task performance without updating pre-trained parameters. However, this effectiveness hinges on the appropriate selection of in-context examples, a process currently biased towards visual data, overlooking textual information. More importantly, the area of supervised retrievers for retrieval of multimodal in-context learning, crucial for optimal in-context example selection, continues to be uninvestigated. Our study provides an in-depth evaluation of the impact of textual information on the unsupervised selection of in-context examples in multimodal contexts, uncovering a notable sensitivity of retriever performance to the employed modalities. Based on the above finding, we introduce a novel supervised MLLM prompt retriever **MSIER** that leverages a trained retriever based on MLLM’s confidence to select examples, which enhances multimodal in-context learning efficiency. This approach is validated through extensive testing across three different tasks, demonstrating the method’s effectiveness. Additionally, we investigate the influence of modalities on our supervised retrieval method’s training and explore the transferability of the supervised prompt retriever. This exploration paves the way for future advancements, highlighting the potential for refined in-context learning in MLLMs through the strategic use of multimodal data. The public code is available at <https://github.com/NUS-HPC-AI-Lab/Multimodal-ICL-Retriever>.

## 1 Introduction

The capability of large language models (LLMs) for in-context learning (ICL) has acquired significant attention and exhibited remarkable efficacy across a diverse range of downstream tasks (An et al., 2023; Min et al., 2022b; Wang et al., 2023). ICL aims to adapt the pre-trained model to achieve

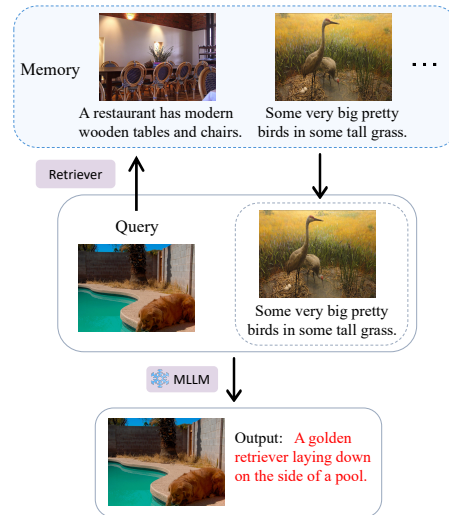


Figure 1: An overview of multimodal in-context example retrieval: This process involves receiving an image or an image-text query from the test dataset, and then using a retrieval mechanism to find similar examples in a training dataset. These examples and the original query (collectively called the prompt) are then inputted into a MLLM to generate the output.

high performance on downstream tasks without any update of parameters that require extra computing resources for the further training process. To be specific, the primary advantage of in-context learning lies in its capacity to assimilate and adapt to new information from its context. This is achieved by utilizing a series of question-and-answer pairs from the training dataset, known as in-context examples, which enable learning without necessitating updates to the model’s parameters. Significantly, recent developments in multimodal large language models (MLLMs), which are constructed based on the foundations of large language models, have similarly exhibited capabilities in in-context learning (Alayrac et al., 2022; Peng et al., 2023; Zhao et al., 2023; Yin et al., 2023), termed as multimodal in-context learning (M-ICL).

A notable advantage of M-ICL lies in its ability

to utilize a singular model across a variety of language comprehension tasks. Nevertheless, (Chen et al., 2023a; Wu et al., 2023; Ye et al., 2023) revealed that the efficacy of this approach in downstream applications can differ significantly, influenced by the selection of in-context examples. This revelation has fueled interest in the concept of in-context example retrieval, as illustrated in Figure 1. In this process, training examples (*memory*) for a given test instance (*query*) are selected for the prompt, guided by a certain similarity metric.

In contrast to Natural Language Processing (NLP) (Dong et al., 2023; Coda-Forno et al., 2023), where the primary input is text, multimodal tasks predominantly utilize a single image (image-text pair) as input and the memory for retrieval consists of multimodal examples. Current methodologies for unsupervised retrieval of in-context examples in multimodal large language models (MLLMs) predominantly focus on visual information (Alayrac et al., 2022; Yang et al., 2022; Gu et al., 2023), thereby overlooking the significance of linguistic data. This oversight prompts inquiries regarding the impact of textual content on retrieval processes and whether the inclusion of additional modalities could potentially enhance the efficacy of MLLMs’ in-context learning capabilities. Moreover, the supervised retrieval framework for M-ICL has received minimal attention so far.

In this work, our attention is centered on M-ICL, an emerging concept in the MLLMs field with limited existing research on its practical application. Specifically, we undertook a thorough investigation into the effects of textual information for retrieval of in-context examples in M-ICL. Through this research, we have developed a supervised retriever that aligns with existing knowledge from MLLMs. This framework achieves high efficiency and demonstrates excellent transferability across various datasets and sizes of MLLMs. The contributions of this paper are summarized as follows:

- We present the first thorough analysis into the role of textual modal in both unsupervised and supervised retrieval of in-context examples for M-ICL and demonstrate that the addition of textual modal plays a crucial role in improving M-ICL performance.
- By considering both visual and textual information when selecting in-context examples, we design a Multimodal Supervised In-

context Examples Retrieval (MSIER) framework with more efficient example selection performance by using a foundation MLLM scorer to evaluate the relevance and suitability of potential examples from multimodal data.

- Extensive experiments on three typical multimodal tasks demonstrate the high efficiency of our constructed supervised retrieval method that achieves the best performance. Besides, we further provide valuable insights regarding the importance and transferability of the supervised method in selecting highly relevant in-context examples for M-ICL.

## 2 Related Work

### In-context Examples Retrieval in NLP and CV

The field of natural language processing has identified that the choice of in-context examples significantly influences performance, as evidenced by (Agrawal et al., 2022) and (Min et al., 2022b). Furthermore, the construction of these in-context examples, often referred to as prompts, including aspects such as relevance and diversity of retrieved examples, has been reported to impact performance as well. This understanding has guided the community toward investigating optimal methods for selecting effective in-context examples for large language models. In their study, (Liu et al., 2021) posited that effective in-context examples should bear semantic similarity to query sentences. Based on this hypothesis, they advocated for the selection of the nearest neighbors in the training set, as determined by a sentence encoder like RoBERTa (Liu et al., 2019). (Rubin et al., 2022) initially employed an unsupervised approach to gathering potential candidates, from which the most suitable examples were selected using a supervised method.

Correspondingly, (Zhang et al., 2023) explored the application of a supervised retriever within the scope of visual in-context learning, demonstrating enhanced performance outcomes. Nevertheless, the significance of textual information is overlooked within unsupervised and supervised retrievals for M-ICL, and the impact of various modalities on the development of the supervised retriever continues to be an area requiring investigation.

### Multimodal Retrieval-augmented Generation

Retrieval augmentation for multimodal models enhances their capabilities by infusing externally retrieved information into their workflow. (Chen et al., 2022) proposed MuRAG, which accesses an

external multimodal memory to augment language generation. (Yasunaga et al., 2023) presented the first retrieval-augmented multimodal model that enables a base model to refer to relevant text and images fetched from external memory. (Ramos et al., 2023) introduced a new approach to image captioning that generates sentences given the input image and retrieved captions. The retrieval process in RAG is designed to retrieve supplementary information to enhance the context for text generation and to produce more informed and contextually rich outputs. On the other hand, the retrieval of in-context examples in M-ICL aims to find examples that demonstrate how to perform a given task, providing explicit guidance to the model on how to generate its output. However, the role of textual information in the task of in-context example retrieval has remained unexplored so far.

**In-context Learning** In-context learning was first proposed by (Brown et al., 2020) in their paper introducing GPT-3. It is a significant departure from the traditional learning method based on stochastic gradient descent and does not involve any parameter updates: the model needs to be provided with a few examples of the task as part of the context for text generation. The size of the model and training data were thought to be key to training a model with in-context learning capabilities. More recently, there has been more research on the exact causes of in-context learning. (Min et al., 2022a) has proposed MetaICL, a meta-training framework to elicit in-context learning capabilities in text-only language models. MetaICL conditions each example with related in-context examples during training. In the context of MLLMs, ICL has been extended to more modalities, leading to Multimodal ICL (Alayrac et al., 2022; Li et al., 2023; Chen et al., 2023b).

## 3 Method

### 3.1 Multimodal In-Context Learning

Within the realm of multimodal tasks such as image captioning, a pre-trained MLLM utilizes an image and its corresponding caption as an in-context example to delineate requirements of the task — essentially instructing the model on the expected form of input and output. Subsequently, when presented with a new image, the model can produce a more precise caption based on these learned patterns. Crucially, M-ICL maintains the parameters of the pre-trained model unchanged, thereby offer-

ing a more resource-efficient approach for tackling downstream tasks.

Regularly, given a memory (training dataset)  $D = [x_i, y_i]_{i=1}^N$  comprising  $N$  pairs of multimodalities information and their corresponding labels, alongside a query example  $x_q$  from the test dataset and the pre-trained model  $f$ , the process of in-context learning can be described as follows:

$$y_q = f([C_p, x_q]) \quad (1)$$

where the context prompt  $C_p$  consists of a sequence of input-output pairs from  $D$ , e.g., the input is an image and the output is its caption for the captioning task. Specifically, prompt  $C_p$  serves as a contextual guide, directing the model to generate the optimal  $y_q$  corresponding to query  $x_q$ , while avoiding any alterations to the parameters of the larger model.

### 3.2 Importance of text information for unsupervised MLLM retrieval

As for the framework of the Multimodal Unsupervised In-context Examples Retrieval (MUIER) which is a prevalent retrieval methodology, our analysis initially focuses on the contribution of various modalities. Traditionally, existing MLLMs that present outstanding M-ICL capability have primarily utilized image data to facilitate M-ICL tasks. For instance, Flamingo (Alayrac et al., 2022) employed the Retrieval-based In-Context Example Selection (RICES) strategy (Yang et al., 2022) for the identification of appropriate in-context examples, predominantly assessing the similarity between the query image and the images stored in the memory (training dataset). Nonetheless, this approach overlooks the significant role of textual information, which indeed influences the efficacy of M-ICL.

To validate the effectiveness of text-augmented MUIER, we performed a comprehensive comparison across diverse configurations of unsupervised retrievers, encompassing three main settings specifically designed for the image captioning task: (1) **Q-I-M-I** (Query-Image-Memory-Image) indicates the case where only image information of image-text pairs preserved in the memory was applied for the retrieval of context. As described in RICES, we only consider the cosine similarity between the query image and the memory image based on vision features extracted by the unsupervised retriever. (2) **Q-I-M-IT** describes the standard setting for MUIER by pairing the memory image with its

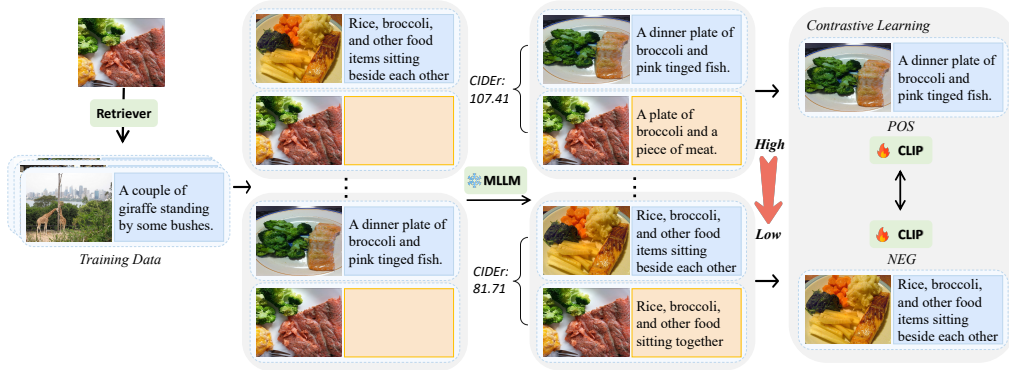


Figure 2: Overview of the MSIER Method: The fundamental principle involves assessing the in-context learning performance for each source instance, thereafter identifying and choosing those instances exhibiting the most favorable or least favorable outcomes. These selected instances are then utilized to form a dataset, categorized as either positive or negative, which is essential for the facilitation of contrastive learning. The examples with high CIDEr scores (corresponding to **low NLL loss** during the scoring process) are selected as positive samples.




Figure 3: The introduction of textual information in the unsupervised method leads to a higher M-ICL performance across all numbers of shots, demonstrating the importance of text modality.

corresponding text label. Specifically, when we use an image as a query, we will calculate its similarity between the query image  $i_q$  and the memory image  $i_m$  as well as its caption  $c_m$ , which means the final similarity  $S = \cos(i_q, i_m) + \cos(i_q, c_m)$ .

Figure 3 presents the M-ICL performance in different unsupervised retriever settings, given the same query (Q) and memory (M) configuration. Compared with the standard setting, the Q-I-M-IT setting exerts a positive impact on the M-ICL performance to varying degrees, which verifies the substantial influence of textual content on M-ICL performance.

### 3.3 Multimodal Supervised Prompt Retriever

The unsupervised methodology depicted earlier is not explicitly designed for multimodal in-context learning applications. Instead, its efficacy relies on the preliminary training phase of the feature extraction mechanism, where the objective function utilized may not align with the demands of multimodal in-context learning scenarios. In contrast, we propose a novel strategy that is based on supervised in-context examples retrieval, assuming the availability of labeled training data as shown in Figure 2. The primary aim of this approach is to refine the original retriever to ensure that the chosen in-context example(s) contribute effectively towards enhancing the log-likelihood maximization.

It is important to acknowledge that evaluating each candidate within the dataset for MLLM training entails a significant expenditure of time and computational resources. This is attributed to the extensive size of existing MLLMs and the substantial overhead associated with preprocessing image data. Consequently, in alignment with the methodologies outlined in (Rubin et al., 2022), we initially narrowed down the Top- $N$  ( $N=50$ ) candidates by assessing the cosine similarity of multimodal information through the application of the unsupervised retriever. Drawing upon the findings presented in 3.2, our approach integrates both visual and textual data for the identification of Top- $N$  examples, which are subsequently employed in the training of our proposed multimodal in-context examples

retriever.

**Scoring** In the process of scoring retrieved candidates for each training instance via an unsupervised retriever, we initially devise effective prompts as illustrated in Table 1 following (Awadalla et al., 2023). The MLLM scorer is furnished with a specific image or image-text prompt as input and proceeds to make predictions for the designated orange segment. These predictions are subsequently leveraged to compute the NLL loss, which serves as the performance metric for the given in-context instance. Consequently, the candidates are reorganized based on their scores to facilitate the subsequent selection of positive and negative samples for contrastive learning.

**Training** Adopting the contrastive learning framework as mentioned by (Rubin et al., 2022), we introduced a query encoder designed to process inputs comprising image or image-question pairs, and a context encoder tasked with handling candidate prompts represented as image-prompt pairs, for feature extraction aimed at subsequent optimization. Both encoders are initialized with vision encoders and text encoders of CLIP model (Radford et al., 2021). During each iteration, a minibatch  $B$  is selected from the training dataset to fine-tune the CLIP model. Furthermore, for every instance within the assembled memory  $B$  with  $N$  image-text pair, both a positive and a negative example are independently sampled from the Top- $K$  positive and negative candidates. This sampling strategy is implemented to develop a similarity metric. Specifically, this metric should ensure that for a given test example  $x_q$ , it exhibits similarity to training instances that facilitate the decoding of  $y_q$ . The formulation of the contrastive loss is executed as follows:

$$\mathcal{L} = -\log \frac{e^{\cos(x_q, e^+)}}{e^{\cos(x_q, e^+)} + \sum_{i=1}^{2N-1} e^{\cos(x_q, e_i^-)}} \quad (2)$$

where  $\cos(\cdot, \cdot)$  measures the cosine similarity.  $e^+$  denotes the feature representation of a positive example, and  $e^-$  denotes the feature representation of a negative example.

### 3.4 Importance of Textual Information in Supervised Retriever

In Section 3.2, the significance of various modalities is examined within the context of unsupervised




Figure 4: Impact of texts on proposed MSIER method. ‘T’ denotes the Training setting and ‘E’ denotes the Evaluation setting.

retrieval, yet the impact of textual data on MSIER has not been investigated. To thoroughly assess the contribution of text in the training of supervised retrievers, we conducted a detailed comparison of two scenarios involving fine-tuned CLIP models, distinguishing them based on their approach to incorporating textual information within the image captioning task.

We decoupled the training and evaluation settings to investigate the significance of textual data in our proposed supervised retrieval framework. Specifically, we employed distinct query-memory configurations during the retriever’s training phase and evaluation phase. As depicted in Figure 4, two training configurations – **Q-I-M-I** and **Q-I-M-IT** – were implemented, and their performance was assessed under two corresponding settings to ensure an equitable comparison. Within the context of the Q-I-M-I setting, textual data during the training phase has a negligible influence on the retriever’s performance. In contrast, incorporating textual information during the training of the supervised multimodal retriever markedly enhances its effectiveness, resulting in a significant performance improvement.

## 4 Experiment

### 4.1 Datasets

Following (Alayrac et al., 2022), we focus on three representative multimodal tasks and the details about the datasets used for these tasks as follows. MS COCO (Lin et al., 2015) for image captioning, OK-VQA (Marino et al., 2019) for Visual Question Answering, and HatefulMemes (Kiela et al., 2021) for rank classification. Accuracy on the test split

Tasks	Prompts
Image Captioning	<image> Output: [caption]
VQA	<image> Question: [question] Short answer: [answer]
Rank Classification	<image> is an image with: '[text]' written on it. Is it hateful? Answer: [answer]

Table 1: Prompts used for different tasks. Orange parts are used for calculating scores for each candidate in memory.

is measured for OK-VQA. In MSCOCO, evaluation utilizes CIDEr scores (Vedantam et al., 2015) on the Karpathy-test split. For HatefulMemes, the AUC ROC is calculated. All results are averaged with three runs. For further details of these downstream tasks and corresponding datasets, please refer to Appendix A.3 and A.4. The experiments conducted herein utilize the OpenFlamingo-3B framework (Awadalla et al., 2023). Comprehensive details regarding OpenFlamingo can be found in Appendix A.2.

## 4.2 Compared Methods

The analysis evaluates various methodologies, including the Random baseline, RICES (Q-I-M-I), Q-T-M-T textual variant, Multimodal Unsupervised In-context Examples Retrieval (MUIER) using all multimodality aspects, and the proposed approach: Multimodal Supervised In-context Examples Retrieval (MSIER) enhancing CLIP’s dual encoders. MSIER undergoes 30 epochs of training with the AdamW optimizer (Kingma and Ba, 2017), warming up to a peak learning rate  $1e-5$ , subject to reduction using the cosine annealing rule. For further details of these methods, please refer to Appendix A.1.

## 4.3 Main Results

As shown in Table 2, compared with Random and RICES, MUIER method shows further enhancements, underscoring the benefits of integrating images and text in a retrieval strategy without explicit supervision. This approach enables the model to better understand in-context examples, with the textual content offering additional direction, thus improving the retrieval of more relevant examples. MSIER method outperforms all with a significant rise in performance, evidenced by a 5.52 increase in the CIDEr score of MS COCO dataset. This suggests that a refined retrieval mechanism, informed by the foundational knowledge of multimodal large language models (MLLMs) and tailored to select correct image-text pairings, yields superior learning outcomes. Furthermore, MSIER with only 4

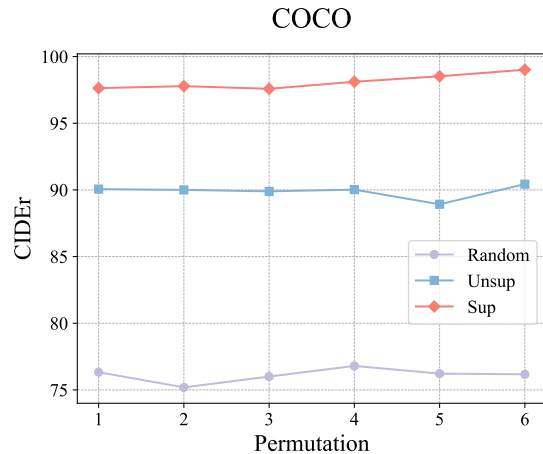


Figure 5: Impact of the order of retrieved multimodal in-context examples.

shots outperforms random selection with 32 shots. Thus, employing our method can enable MLLMs to achieve comparable M-ICL performance with limited memory usage.

## 4.4 Further Analysis

**How does the MSIER improve M-ICL specifically?** To address this query, we employed a multimodal analysis of the in-context examples identified by MUIER and MSIER as depicted in Figure 6. Our examination concentrates on image captioning by selecting one example from the MS COCO dataset where the number of shots is 4. Concretely, the columns from left to right delineate the retrieved in-context example with RICES, MUIER, and MSIER separately. The blue rows are the retrieved multimodal in-context examples, that is, a pair comprising input and output, whereas the subsequent orange rows present the query image alongside the model’s prediction. In the given figure, examples identified by MSIER exhibit a closer semantic similarity to the queries than RICES and MUIER. Consequently, the MLLM made a more accurate prediction that includes the caption "a chocolate covered peanut", which is not captured by previous predictions using RICES and MUIER methods for retrieving in-context examples.

**How does the order of in-context examples af-**

	MS COCO				OK-VQA					HatefulMemes				
	Random	RICES	MUIER	MSIER	Random	RICES	Q-T-M-T	MUIER	MSIER	Random	RICES	Q-T-M-T	MUIER	MSIER
4 shots	77.19	91.43	92.78	<b>100.58</b>	30.32	31.09	30.66	32.46	<b>33.64</b>	50.56	65.74	63.05	68.78	<b>70.77</b>
8 shots	85.97	96.91	98.93	<b>105.09</b>	31.08	32.36	31.90	34.82	<b>36.16</b>	52.67	67.28	63.64	70.13	<b>72.53</b>
16 shots	90.51	102.22	103.65	<b>108.41</b>	30.93	33.58	32.57	36.04	<b>37.12</b>	48.45	68.41	62.55	70.34	<b>72.59</b>
32 shots	92.47	106.30	107.24	<b>110.58</b>	31.02	35.41	33.32	36.79	<b>38.35</b>	50.91	68.36	63.27	71.50	<b>73.80</b>
Avg	86.54	99.22	100.65	<b>106.17</b>	30.84	33.11	32.12	35.03	<b>36.32</b>	50.65	67.45	63.13	70.19	<b>72.42</b>

Table 2: Results of M-ICL performance of random selection, RICES, MUIER, and MSIER method on MS COCO, OK-VQA, and HatefulMemes dataset.



Figure 6: Multimodal in-context examples retrieved by RICES, MUIER and MSIER. In each grid, the first four blue rows contain the prompt while the orange row contains the query image and prediction. Further outcomes are detailed in the Appendix B.

**fect M-ICL?** To understand if changing the order of multimodal in-context examples makes a difference, we fixed the number of in-context examples to 3, and evaluate all possible permutations. As shown in Figure 5, the standard deviation is generally small, so the order is not a concern as long as high-quality examples are chosen.

**How is the transferability of proposed MSIER?** The compositional characteristics of natural language and images are general, meaning the retriever may exploit similar knowledge in different tasks or scoring MLLMs. This motivates us to explore whether the proposed MSIER trained on one dataset or MLLM scorer can be directly transferred to others without further tuning. This is a practical research question as training a retriever for each dataset or MLLM scorer can be costly in real applications.

- **Datasets** To measure the transferability of supervised retrievers between datasets, we evaluated the performance of MSIER trained on

OK-VQA and HatefulMemes datasets on MS COCO dataset with different numbers of shots of retrieved in-context examples. The result in Table 3 manifests some transferability of our proposed MSIER between multiple multimodal datasets. Enhanced performance of MSIER with training on the OK-VQA dataset suggests that data form and volume significantly influence MSIER’s effectiveness.

- **MLLM Scorer** Regarding the impact of MLLM scorer on the transferability of MSIER, our methodology involves deploying a retriever trained utilizing the OpenFlamingo-3B model as the scoring model, subsequently applied to the inference processes of the OpenFlamingo-9B model. The outcomes in Table 4 demonstrate that the MSIER approach, when utilizing the 3B model as a scorer, manifests superior transferability and outperforms the MUIER-9B method. This implies that supervised retrieval methods, initially trained

on smaller-scale models, retain effectiveness upon application to larger models, obviating the necessity for separate training for larger models and thus enhancing the cost-efficiency of inference processes in larger-scale models.

Method	4 shots	8 shots	16 shots	32 shots	Avg
MUIER (Q-I-M-IT)	92.78	98.93	103.65	107.24	100.65
HMM (Q-I-M-IT)	91.92	97.87	103.82	107.30	100.23
OKVQA (Q-I-M-IT)	93.28	100.72	105.95	109.02	102.24

Table 3: M-ICL performance of MSIER trained using HMM and OK-VQA datasets on MS COCO dataset.

Method	4 shots	8 shots	16 shots	32 shots	Avg
Random	89.03	96.30	98.78	99.91	96.01
RICES (Q-I-M-I)	93.28	99.89	105.71	108.79	101.92
MUIER-9B (Q-I-M-IT)	94.53	100.48	106.11	109.95	102.77
MSIER-3B (Q-I-M-IT)	100.73	103.93	107.50	110.60	105.69

Table 4: Results of M-ICL performance of random selection, RICES, and MUIER-9B and MSIER-3B method on MS COCO dataset using OpenFlamingo-9B for inference.

#### 4.5 Ablation Study

This section presents an extensive ablation study based on the OKVQA and MS COCO datasets to verify the best setting for our proposed model’s components.

**Impact of Different Modality Encoders in CLIP** Unlike ICL retrieval in NLP and CV, M-ICL retrieval employs multiple encoders in its backbone retriever to extract features from different modalities. To investigate the roles of the text and image encoders within the base CLIP model in our proposed MSIER framework, we selectively freeze one encoder during MSIER training. Table 5 demonstrates that the update on image encoder has a more significant effect on M-ICL retrieval performance improvement because freezing the image encoder results in a notable performance drop compared to the original MSIER (1.39 ↓). In contrast, freezing only the text encoder yields similar or even improved retrieval performance, suggesting that the original trained MSIER may have overfitted to some degree on the text encoder.

**Impact of Number of Candidates** We proceed to assess the performance of MSIER across a range of candidate numbers. In MSIER, we passed the candidates to a scoring LM and label the top- $K$  and the bottom- $K$  as positive and negative examples

Method	4 shots	8 shots	16 shots	32 shots	Avg
MUIER (CLIP)	92.78	98.93	103.65	107.24	100.65
MSIER (CLIP)	100.58	105.09	108.41	110.58	106.17
MSIER (CLIP-Freeze T)	100.82	105.90	109.17	110.24	106.53
MSIER (CLIP-Freeze I)	98.13	103.37	107.59	110.01	104.78

Table 5: Comparison of M-ICL performance of MSIER method on MS COCO dataset with freezing different modality encoders of the backbone retriever.

respectively. If the number of candidates is not large enough, it will not hold sufficient information for contrastive learning; conversely, if we set the bank with a large size, it will contain a quantity of irrelevant items. The results in Table 6 present  $K=5$  as the best choice on OK-VQA dataset.

Method	4 shots	8 shots	16 shots	32 shots	Avg
$K = 1$	33.57	35.93	36.92	38.07	36.12
$K = 5$	33.64	36.16	37.12	38.35	36.32
$K = 10$	33.74	35.58	36.85	37.94	36.03

Table 6: Comparison of M-ICL performance of different settings of  $K$  for MSIER on OK-VQA dataset.

**Impact of Textual Information in Retrieved In-context Examples** The role of textual information during the evaluation of M-ICL remains unexplored for our proposed MSIER. In Table 7, we find that compared to the standard scenario, the replacement of texts in the in-context examples hugely impacts the M-ICL performance. To be specific, employing a mask on half of the captions through the random selection method results in a significant decrease in the CIDEr score. Correspondingly, the side effect of the removal of texts decreases from MUIER method to MSIER method compared with random retrieval.

Method	4 shots	8 shots	16 shots	32 shots	Avg
Random	77.19	85.97	90.51	92.47	86.54
Random w/ mask	2.27	2.28	2.40	2.79	2.44
MUIER (Q-I-M-IT)	92.78	98.93	103.65	107.24	100.65
MUIER (Q-I-M-IT) w/ mask	70.32	74.34	77.68	78.62	75.24
MSIER (Q-I-M-IT)	100.58	105.09	108.41	110.58	106.17
MSIER (Q-I-M-IT) w/mask	77.62	82.00	85.83	86.77	83.06

Table 7: Comparison of M-ICL performance of random selection, MUIER, and MSIER method with masked text in-context examples on MS COCO dataset.

**Impact of Backbone Retriever** To investigate whether using a different backbone model could improve performance, we further evaluated retrieval methods, MUIER and MSIER, on the image captioning benchmark by utilizing an alternative backbone: EVA-02-CLIP (Sun et al., 2023), a model



enhances the original CLIP model with advanced pre-training techniques. The results, reported in Table 8, show that MUIER and MSIER methods employing EVA-02-CLIP(-L/14) as the backbone retriever demonstrated better performance. This suggests that the EVA-02-CLIP model offers retrieval benefits for the M-ICL of the OpenFlamingo model, with improved feature extraction for both textual and visual information. Furthermore, the proposed MSIER method exhibited a corresponding performance improvement when using the enhanced retriever, explicitly validating MSIER’s excellent transferability and ability to achieve better results as multimodal encoders advance. Additional experiments can be found in Appendix A.5.

Method	4 shots	8 shots	16 shots	32 shots	Avg
MUIER (CLIP)	92.78	98.93	103.65	107.24	100.65
MSIER (CLIP)	100.58	105.09	108.41	110.58	106.17
MUIER (EVA-02-CLIP)	94.02	101.70	107.06	109.71	103.12
MSIER (EVA-02-CLIP)	104.20	107.44	111.61	113.07	109.08

Table 8: Comparison of M-ICL performance of MUIER and MSIER method on MS COCO dataset with EVA-CLIP2 model as the backbone retriever.

## 5 Conclusion

In this study, we conducted an extensive evaluation of textual information’s role in unsupervised and supervised in-context example retrieval for multimodal in-context learning. We introduce a novel MSIER methodology that incorporates MLLMs’ self-contained information to train a supervised prompt retriever. Our experiments across three multimodal tasks demonstrate that integrating text modality and the foundational knowledge of MLLMs significantly enhances the efficiency of example selection, yielding substantial improvements over existing approaches. Moreover, MSIER demonstrates high efficiency by outperforming benchmarks with far fewer examples, while also exhibiting strong transferability by remaining effective for larger models after training on smaller ones, thereby improving cost-efficiency for large-scale inference. Future research could delve into optimizing inter-modal interactions, leveraging our findings for more effective multimodal in-context learning retrieval strategies.

## 6 Limitations

Our research focuses exclusively on the integration of visual and textual data for multimodal in-context

learning. Given the expanding use of additional modalities (such as video and audio) in this domain, the development of a comprehensive framework capable of consolidating these varied modalities into a unified representation is increasingly imperative.

We observed that enhancements in multimodal in-context learning (M-ICL) performance on the MS COCO dataset surpass those on other datasets. This discrepancy may stem from the necessity for more meticulously constructed prompts in visual question answering (VQA) and rank classification tasks, which demand more than mere questions or captions for effective in-context example retrieval.

Moreover, both diversity and relevance play critical roles in the selection of in-context examples for M-ICL. Future investigations might, therefore, benefit from examining how the relationships between retrieved examples can inform the design of an improved supervised retriever for M-ICL.

## Acknowledgements

Yang You’s research group is being sponsored by NUS startup grant (Presidential Young Professorship), Singapore MOE Tier-1 grant, ByteDance grant, ARCTIC grant, SMI grant and Alibaba grant.

## References

- Sweta Agrawal, Chunting Zhou, Mike Lewis, Luke Zettlemoyer, and Marjan Ghazvininejad. 2022. [In-context examples selection for machine translation](#).
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian Borgeaud, Andrew Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karen Simonyan. 2022. [Flamingo: a visual language model for few-shot learning](#).
- Shengnan An, Zeqi Lin, Qiang Fu, Bei Chen, Nanning Zheng, Jian-Guang Lou, and Dongmei Zhang. 2023. [How do in-context examples affect compositional generalization?](#)
- Anas Awadalla, Irena Gao, Josh Gardner, Jack Hessel, Yusuf Hanafy, Wanrong Zhu, Kalyani Marathe, Yonatan Bitton, Samir Gadre, Shiori Sagawa, Jenia Jitsev, Simon Kornblith, Pang Wei Koh, Gabriel Ilharco, Mitchell Wortsman, and Ludwig Schmidt. 2023. [Openflamingo: An open-source framework for training large autoregressive vision-language models](#). *arXiv preprint arXiv:2308.01390*.

- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#).
- Shuo Chen, Zhen Han, Bailan He, Mark Buckley, Philip Torr, Volker Tresp, and Jindong Gu. 2023a. [Understanding and improving in-context learning on vision-language models](#).
- Wenhu Chen, Hexiang Hu, Xi Chen, Pat Verga, and William W. Cohen. 2022. [Murag: Multimodal retrieval-augmented generator for open question answering over images and text](#).
- Yi-Syuan Chen, Yun-Zhu Song, Cheng Yu Yeo, Bei Liu, Jianlong Fu, and Hong-Han Shuai. 2023b. [Sinc: Self-supervised in-context learning for vision-language tasks](#).
- Julian Coda-Forno, Marcel Binz, Zeynep Akata, Matthew Botvinick, Jane X. Wang, and Eric Schulz. 2023. [Meta-in-context learning in large language models](#).
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, Lei Li, and Zhifang Sui. 2023. [A survey on in-context learning](#).
- Jindong Gu, Ahmad Beirami, Xuezhi Wang, Alex Beutel, Philip Torr, and Yao Qin. 2023. [Towards robust prompts on vision-language models](#).
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yunhsuan Sung, Zhen Li, and Tom Duerig. 2021. [Scaling up visual and vision-language representation learning with noisy text supervision](#).
- Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ringshia, and Davide Testuggine. 2021. [The hateful memes challenge: Detecting hate speech in multimodal memes](#).
- Diederik P. Kingma and Jimmy Ba. 2017. [Adam: A method for stochastic optimization](#).
- Bo Li, Yuanhan Zhang, Liangyu Chen, Jinghao Wang, Jingkang Yang, and Ziwei Liu. 2023. [Otter: A multimodal model with in-context instruction tuning](#).
- Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. 2015. [Microsoft coco: Common objects in context](#).
- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2021. [What makes good in-context examples for gpt-3?](#)
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).
- Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. 2019. [Ok-vqa: A visual question answering benchmark requiring external knowledge](#).
- Sewon Min, Mike Lewis, Luke Zettlemoyer, and Han-naneh Hajishirzi. 2022a. [Metaicl: Learning to learn in context](#).
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022b. [Rethinking the role of demonstrations: What makes in-context learning work?](#) In *EMNLP*.
- Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. 2023. [Kosmos-2: Grounding multimodal large language models to the world](#).
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Rita Ramos, Desmond Elliott, and Bruno Martins. 2023. [Retrieval-augmented image captioning](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3666–3681, Dubrovnik, Croatia. Association for Computational Linguistics.
- Ohad Rubín, Jonathan Herzig, and Jonathan Berant. 2022. [Learning to retrieve prompts for in-context learning](#).
- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. 2022. [Laion-5b: An open large-scale dataset for training next generation image-text models](#).
- Quan Sun, Yuxin Fang, Ledell Wu, Xinlong Wang, and Yue Cao. 2023. [Eva-clip: Improved training techniques for clip at scale](#).
- Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. [Cider: Consensus-based image description evaluation](#).
- Lean Wang, Lei Li, Damai Dai, Deli Chen, Hao Zhou, Fandong Meng, Jie Zhou, and Xu Sun. 2023. [Label words are anchors: An information flow perspective for understanding in-context learning](#).

- Zhiyong Wu, Yaoxiang Wang, Jiacheng Ye, and Lingpeng Kong. 2023. [Self-adaptive in-context learning: An information compression perspective for in-context example selection and ordering.](#)
- Zhengyuan Yang, Zhe Gan, Jianfeng Wang, Xiaowei Hu, Yumao Lu, Zicheng Liu, and Lijuan Wang. 2022. [An empirical study of gpt-3 for few-shot knowledge-based vqa.](#)
- Michihiro Yasunaga, Armen Aghajanyan, Weijia Shi, Rich James, Jure Leskovec, Percy Liang, Mike Lewis, Luke Zettlemoyer, and Wen tau Yih. 2023. [Retrieval-augmented multimodal language modeling.](#)
- Jiacheng Ye, Zhiyong Wu, Jiangtao Feng, Tao Yu, and Lingpeng Kong. 2023. [Compositional exemplars for in-context learning.](#)
- Shukang Yin, Chaoyou Fu, Sirui Zhao, Ke Li, Xing Sun, Tong Xu, and Enhong Chen. 2023. [A survey on multimodal large language models.](#)
- Yuanhan Zhang, Kaiyang Zhou, and Ziwei Liu. 2023. [What makes good examples for visual in-context learning?](#)
- Haozhe Zhao, Zefan Cai, Shuzheng Si, Xiaojian Ma, Kaikai An, Liang Chen, Zixuan Liu, Sheng Wang, Wenjuan Han, and Baobao Chang. 2023. [Mmicl: Empowering vision-language model with multi-modal in-context learning.](#)
- Wanrong Zhu, Jack Hessel, Anas Awadalla, Samir Yitzhak Gadre, Jesse Dodge, Alex Fang, Youngjae Yu, Ludwig Schmidt, William Yang Wang, and Yejin Choi. 2023. [Multimodal c4: An open, billion-scale corpus of images interleaved with text.](#)

## A More Experimental Details

### A.1 Compared Methods

The comparative analysis predominantly concentrates on the effectiveness of various methodologies as follows:

- The baseline methodology, referred to as **Random**, encompasses the arbitrary selection of in-context examples from the originating training dataset.
- The **RICES** (Retrieval-based In-Context Example Selection) strategy completes unsupervised retrieval predicated on the image similarity between the query and in-context examples from memory, denominated as Q-I-M-I.
- The **Q-T-M-T** variant, an adaptation of Q-I-M-I, elects in-context examples predicated on only textual similarity.
- The standard unsupervised approach used in this research, termed Multimodal Unsupervised In-context Examples Retrieval (**MUIER**), employs readily accessible features to identify the most similar examples via search, based on all multimodality aspects.
- Multimodal Supervised In-context Examples Retrieval (**MSIER**) represents our secondary proposal, which involves the enhancement of CLIP’s dual encoders. CLIP ViT-L-14 is applied for all experiments. This enhancement process is facilitated through direct optimization to improve M-ICL performance. The supervised model undergoes training for 30 epochs utilizing the AdamW optimizer. The initial learning rate is established at 0.00001, subject to reduction according to the cosine annealing rule.

### A.2 Multimodal Large Language Model

The architecture of OpenFlamingo encompasses a fixed large language model featuring a decoder-only configuration (for instance, MPT [31]), accompanied by a static visual encoder (such as CLIP-ViT (Radford et al., 2021)), which is succeeded by a trainable perceiver resampler. To facilitate the integration of visual and linguistic data, trainable cross-attention layers are strategically interspersed among the pre-trained language model layers. OpenFlamingo underwent pre-training on the

LAION-2B (Schuhmann et al., 2022) and Multimodal C4 (Zhu et al., 2023) datasets and its performance is on par with that of Flamingo (Alayrac et al., 2022), demonstrating its competitive edge.

Some well-known MLLMs such as LLaVA and Minigt-4 do not discuss their in-context learning ability or present any experimental results on in-context learning in their papers. In contrast, the Flamingo model is one of the first and most renowned MLLMs to include experiments on in-context learning in its paper. Since the original Flamingo model is not open-sourced, we opt for the open-source version, OpenFlamingo, as the backbone. We use its in-context learning performance on three representative downstream tasks in its paper for a fair comparison.

### A.3 Downstream Tasks

- **Image Captioning** is a task where a model generates a textual description of an image. It combines elements of computer vision and natural language processing to interpret the contents of an image and articulate them in human language. The challenge is recognizing the objects within the image and describing their attributes and the relations between them in a coherent sentence structure.
- **Visual Question Answering (VQA)** In this task, a model is given an image along with a natural language question about the image, and it must provide an appropriate answer. To provide an appropriate answer, the model must understand both the visual content of the image and the semantics of the question text. It often involves aspects of object detection, attribute recognition, and language understanding to correctly answer questions such as "What color is the car?" or "How many people are in the room?"
- **Rank Classification** is a task that involves ordering or prioritizing a set of items or entities based on specified criteria, which may be derived from multiple input modalities like text and images. For example, in a multimodal setting, this could involve analyzing text and images together to rank products in an e-commerce setting based on relevance to a search query or user preference. More broadly, rank classification could involve any task where items need to be ordered or pri-

oritized, and it can be extended to scenarios like sentiment analysis, where responses are classified into ranked categories of sentiment intensity.

### A.4 Datasets

A detailed explanation of used datasets corresponding to evaluated downstream tasks is provided:

- **MSCOCO** is a large-scale dataset for object detection, segmentation, and captioning. It provides diverse images with complex scenes and multiple objects in context. Annotations include object segmentation, recognition, and image captions. The dataset contains approximately 330K images, with 1.5 million labeled instances across 80 object categories. We utilize Karparthy's split: training (83K images), validation (5K images), and test (5K images) sets.
- **OK-VQA** is a dataset designed for open-ended Visual Question Answering that requires external knowledge beyond image content. It features questions demanding multimodal knowledge, combining visual cues with general world knowledge. The dataset comprises over 14,000 images sourced from the MSCOCO dataset, with around 14,000 questions spanning 10 categories. We leverage its structured format, which includes a balanced mix of 9,009 training and 5,046 testing questions, to assess the capability of models to integrate visual understanding with external knowledge sources.
- **HatefulMemes** A dataset constructed for the detection and classification of hate speech in multimodal content. It uniquely combines textual and visual elements to challenge models in understanding complex, nuanced expressions of hate speech. The dataset consists of 10,000+ meme images, annotated for hate speech detection, with a distribution of 8,500 for training and 3,000 for testing. The Hateful Memes dataset provides a significant challenge in discerning subtle contextual cues and cultural references, requiring advanced multimodal analysis capabilities.

### A.5 Retrieval Methods

In our main experiments, MUIER and MSIER utilized all modalities to calculate similarity scores

and directly selected the most similar items as in-context examples. Recently, some newly proposed retrieval methods have emerged, such as Mixed Modality In-Context Example Selection (Chen et al., 2023a), which uses visual features to retrieve a set of top-N candidate examples and then employs textual information to retrieve the final in-context examples (with varying numbers of examples, such as 4, 8, 16, or 32). We also validated the efficiency of MSIER in this scenario, and Table 9 shows that, although we did not train a new MSIER specifically designed for the new retrieval method, it still demonstrated better performance on the VQA task. This further highlights the generality of our proposed MSIER framework.

Method	4 shots	8 shots	16 shots	32 shots	Avg
MMICES-CLIP	32.96	33.60	35.89	36.13	34.65
MMICES-MSIER	33.08	34.36	35.92	36.84	35.05

Table 9: Comparison of M-ICL performance of MMICES with CLIP as retriever, and MMICES with MSIER as retriever on MS COCO dataset. We chose the top 50 candidates for further selection of in-context examples in this experiment.

**Impact of Backbone Retriever** We further evaluated the retrieval methods MUIER and MSIER using a different backbone: ALIGN (Jia et al., 2021), a model pre-trained on noisy image-text pairs using contrastive learning. The results, reported in Table 10, show that the MUIER and MSIER methods did not demonstrate explicit advantages over random selection. This suggests that the ALIGN model does not offer retrieval benefits for the M-ICL of the OpenFlamingo model. This can be attributed to the fact that OpenFlamingo utilizes the CLIP model as its vision encoder, resulting in poor contextual information from examples primarily retrieved by the ALIGN model, leading to little enhancement in M-ICL performance.

Method	4 shots	8 shots	16 shots	32 shots	Avg
Random	77.02	84.13	88.96	92.32	85.61
MUIER (ALIGN)	78.78	84.10	89.22	92.07	86.04
MSIER (ALIGN)	78.17	83.79	89.38	94.02	86.34

Table 10: Comparison of M-ICL performance of random selection, MUIER, and MSIER method on MS COCO dataset with ALIGN model as the backbone retriever.

## B More Evaluation Examples

More evaluation examples using different retrieval methods are presented here:

<b>RICES</b>	<b>MUIER</b>	<b>MSIER</b>
A male tennis player celebrating a point.	A male tennis player celebrating a point on the tennis court.	A man holding a tennis racket with a ball in the air on the tennis court.


**Test example**



**GROUND TRUTH:** *A tennis player is standing on a tennis court with a racquet in his hand.*

<b>RICES</b>	<b>MUIER</b>	<b>MSIER</b>
A man in a blue shirt and tie.	A man wearing a blue shirt with a black tie.	A man in a blue shirt sits at a desk in front of a computer monitor.

**Test example**



**GROUND TRUTH:** *A man with long hair and a blue shirt sitting in front of a computer.*

<b>RICES</b>	<b>MUIER</b>	<b>MSIER</b>
A surfer standing on a surfboard.	A surfer in a wetsuit holding a surfboard.	A man wearing a wet suit standing on the beach holding a surfboard.

**Test example**



**GROUND TRUTH:** *A man in black wet suit holding a surfboard under his arm.*