

# Enhancing E-commerce Product Title Translation with Retrieval-Augmented Generation and Large Language Models

Bryan Zhang  
Amazon  
Pittsburgh, USA  
bryzhang@amazon.com

Taichi Nakatani  
Amazon  
Pittsburgh, USA  
taichina@amazon.com

Stephan Walter  
Amazon  
Berlin, Germany  
sstwa@amazon.de

## Abstract

E-commerce stores enable multilingual product discovery which require accurate product title translation. Multilingual large language models (LLMs) have shown promising capacity to perform machine translation tasks, and it can also enhance and translate product titles cross-lingually in one step. However, product title translation often requires more than just language conversion because titles are short, lack context, and contain specialized terminology. This study proposes a retrieval-augmented generation (RAG) approach that leverages existing bilingual product information in e-commerce by retrieving similar bilingual examples and incorporating them as few-shot prompts to enhance LLM-based product title translation. Experiment results show that our proposed RAG approach improve product title translation quality with chrF score gains of up to 15.3% for language pairs where the LLM has limited proficiency.

## CCS Concepts

• **Information systems** → *Mediators and data integration*; • **Computing methodologies** → *Machine translation*.

## Keywords

LLM, retrieval-augmented generation, e-commerce, data selection, machine translation

### ACM Reference Format:

Bryan Zhang, Taichi Nakatani, and Stephan Walter. 2024. Enhancing E-commerce Product Title Translation with Retrieval-Augmented Generation and Large Language Models. In *Proceedings of ACM CIKM Workshop on Data-Centric AI (CIKM DCAI'24)*. ACM, New York, NY, USA, 6 pages.

## 1 Introduction

As e-commerce shopping websites become localized worldwide, more customers are provided with options to browse products in their preferred language other than the primary language of the store. To accomplish this, modern e-commerce stores enable multilingual product discovery [2, 12, 15, 17, 20, 21] as well as localizing product information such as titles using machine translation (MT) systems [5, 25–27, 29]. Product titles play a crucial role in conveying key details about the products, and it is essential that this information is accurately presented in the localized language.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

CIKM-DCAI Conference'24, Oct. 2024, Boise, ID, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

E-commerce product title localization traditionally uses a fleet of bilingual neural machine translation (NMT) systems. However, recent advancements in multilingual large language models (LLMs) have demonstrated promising performance in machine translation as a single model for high resource languages [9], which makes LLMs as a viable alternative to conventional NMT for translating product titles. Furthermore, LLMs have shown capabilities to optimize the length and the enrich the content of product titles within a single language, which suggest that they may be able to tackle title enhancement and translation in a cross-lingual manner as a one-step unified process [28]. As a result, LLMs are increasingly becoming a more prominent and integral component of e-commerce stores' product title localization strategy.

However, there are challenges in using LLMs for product title localization in e-commerce: (1) product titles tend to be short, and proper title translation also often require bilingual product-specific terminology and catalog domain knowledge. For example, in the title *"Dance your cares away - greeting card"*, the core phrase *"Dance your cares away"* needs to be preserved in the translation, and the formality and style of the titles also must be maintained; (2) e-commerce is dynamic in nature, with new products rapidly emerging worldwide and requiring the model to have up-to-date product-specific knowledge; (3) a large number of language pairs need to be supported in the e-commerce domain but LLMs may not possess sufficient language proficiency for certain language pairs [4], which limits their effectiveness; (4) newer and more capable LLMs are continuously arising, prompting the need to flexibly and quickly replace the LLMs used for product title localization in order to maintain high translation quality.

To address these challenges and improve product title localization in e-commerce, this study proposes a retrieval-augmented generation (RAG) approach. Our approach utilizes the constantly-growing bilingual catalog of products as source of product domain-specific knowledge and multilingual support, which is independent and provides flexibility to work with different LLMs. This method retrieves bilingual product information (e.g. titles, bulletpoints and descriptions) that are similar to the source title and incorporates them as few-shot examples in the prompt to enhance product title translation with LLMs. Finally, the study analyzes the product title translations generated from LLMs with the RAG approach to understand its impact on translation quality.

Our experiment shows that our proposed LLM-based retrieval-augmented generation (RAG) approach for product title translation can significantly enhance the title translation quality with up to 15.3% chrF score improvement for language pairs where the LLM does not primarily support.

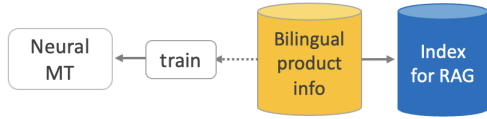


Figure 1: Index large volume of bilingual product information for RAG

## 2 RAG approach for product title translation

### 2.1 Utilizing bilingual product information for RAG

In the e-commerce industry, we can leverage the large and constantly growing volumes of bilingual product information data, including titles, bullet points, and descriptions that can be acquired. These multilingual product titles, descriptions and bulletpoints are commonly used to train smaller neural machine translation (NMT) models for product title translation tasks. Given the dynamic nature of e-commerce, where product information localization is an ongoing process, such bilingual product content continues to expand over time. The bilingual product content can be further utilized to build a search index serving as a source of product domain-specific knowledge and multilingual support for LLM RAG approaches as illustrated in Figure 1.

### 2.2 Using retrieved similar product information as few-shot examples

We propose harnessing this ever-increasing bilingual product information data to power a sustainable and scalable retrieval-augmented generation (RAG) approach to enhance product title translation. By retrieving similar bilingual product information examples from this data and using them as few-shot prompts, we can guide the LLMs to generate higher quality and more contextually appropriate translations for product titles as shown in Figure 2. This RAG-based approach allows us to continuously leverage the benefits of the existing and expanding bilingual product information to improve the performance of our cross-lingual title translation capabilities.

Our proposed RAG approach for product title translation is as follows. Given a large volume of bilingual product information data (e.g. titles),

$$D_{bil} = \{(t_{src}^0, t_{tgt}^0), (t_{src}^1, t_{tgt}^1) \dots (t_{src}^n, t_{tgt}^n)\}$$

where  $t_{src}$  and  $t_{tgt}$  are product information text (e.g. titles) in the source and target languages from a bilingual language pair  $L$ , we propose to build an index  $I_L$  on the product information (e.g. titles) in the source language and cache the corresponding target language counterparts. We recommend using text-based retrieval frameworks such as BM25, as these can better focus on textual patterns and are easy to scale and maintain, previous study has also shown that BM25 is an effective retrieval mechanism for RAG as it can achieve the high RAG performance with a number of LLMs [6]. As newer products and language pairs emerge, they can be readily incorporated into the index.

At inference time, given a product title  $t'$  in the source language for the language pair  $L$ , we retrieve from the index  $I_L$  the top- $k$




Figure 2: Proposed RAG approach for product title translation using LLM

bilingual product information (e.g. titles)

$$D_{topk} = \{(t_{src}^0, t_{tgt}^0), (t_{src}^1, t_{tgt}^1) \dots (t_{src}^k, t_{tgt}^k)\}$$

based on the similarity scoring  $S(t', t_{src})$  of the source product information (e.g. titles) where,  $D_{topk} \subset D_{bil}$ ,  $S(t', t_{src}^0) \geq S(t', t_{src}^1) \geq \dots \geq S(t', t_{src}^k) \geq \dots \geq S(t', t_{src}^n)$  and  $n = |I_L| = |D_{bil}|$ . We then integrate these retrieved top- $k$  bilingual product information (e.g. titles) into the prompts as few-shot examples. The retrieved product information (e.g. titles) share similar textual patterns, language, and semantics with the given title, allowing the LLM to leverage these references to generate a translation that is more aligned with the desired objectives.

The following is the prompt template A for title translation task without RAG, brackets (<>) are placeholders and substituted with relevant text.

```
Below is a source text, it is a title of a product in <source language e.g. English>: <title in the source language>
Your task is to translate the source text into <target language>.
Here are instructions on how to translate:
- Overall, make sure the translation is accurate, particularly pay attention to the terminology.
- Brands are generally preserved in the translation. If you are very confident there is a different brand term in <target language>, use it in the translation.

Return your response in JSON format with the 'translation' key containing your translation as {"translation": <translation>}
Please make sure the output is in JSON format, don't add other information.
```

The following is the prompt template B for title translation task with few-shot examples, we integrate the top- $k$  retrieved similar bilingual title as few-shot examples in the prompt. The brackets (<>) are placeholders and substituted with relevant text.

```
Below is a source text, it is a title of a product in <source language e.g. English>:
<title in the source language>
Your task is to translate the source text into <target language>.
Below are a few translation examples:
Example 1: source: <source title>,
translation: <title translation>
Example 2: source: <source title>,
translation: <title translation>
Example 3: source: <source title>,
translation: <title translation>
Example 4: source: <source title>,
translation: <title translation>
Example 5: source: <source title>,
translation: <title translation>

Here are instructions on how to translate:
- Overall, make sure the translation is accurate, particularly pay attention to the terminology.
- Brands are generally preserved in the translation. If you are very confident there is a different brand term in <target language>, use it in the translation.
```

```

- Refer to the provided translation examples if you could find any
  useful information. For example: any terminology
  translations, phrases translation, how the brand term is
  handled etc. only if you find it useful!

Return your response in JSON format with the 'translation' key
  containing your translation as {"translation": <translation>}
Please make sure the output is in JSON format, don't add other
  information.

```

### 3 Experiment

**Language pairs:** We experiment with 7 language pairs English-Dutch (EN-NL), English-German (EN-DE), Italian-German (IT-DE), Turkish-German (TR-DE), German-Czech (DE-CS), English-Polish (EN-PL) and English-Swedish (EN-SE).

**Test data and MT metric:** For each language pair, we sample 2000 product titles as our test set and evaluate translation quality using chrF [18].<sup>1</sup>

**RAG framework:** We use Lucene 8.11.3 and OKAPI BM25 probabilistic retrieval framework to build each index. For each bilingual product information data pair, we index the source text and cache the target text.

**RAG data:** We sample large volumes of bilingual product (titles, descriptions and bulletpoints) for each language pair to build the index for our RAG approach. For each language pair, we build four indices using sampled bilingual product information data of different domains: product titles (TTL), bullet points (BP), descriptions (PD) and a combination of aforementioned domains (T.B.D.). Language pairs EN-NL, EN-PL, EN-DE and IT-DE have over 10 million bilingual pairs in their T.B.D. index, and the rest have millions.

**Experiment configurations:** We set up the following experiment configurations using different prompts for title translation generation with LLM:

**Baseline Generation:** The prompt using the *prompt template A* presented in section 2.2, which does not incorporate retrieved or randomly sampled few-shot examples.

**RAND 1-shot and RAND 5-shot:** The prompt using the *prompt template B* presented in section 2.2, which incorporates 1 and 5 randomly-chosen examples from the respective domain (TTL, BP, PD, T.B.D.) as few-shot prompts. These serve as the lower-threshold of our RAG approach.

**RAG 1-shot and RAG 5-shot:** The prompt using the *prompt template B* in section 2.2, which incorporates 1 and 5 most similar examples retrieved from their respective domain index (TTL, BP, PD, T.B.D.) as few-shot prompts.

**LLM:** We use Mixtral-8x7B-Instruct<sup>2</sup>, a publicly available LLM, to translate source product title into the target language through

<sup>1</sup>We use chrF instead of BLEU because the average length of a title in our testsets is 11.3 raw words, making BLEU's four-gram statistics sparse and unreliable even at the corpus level [14] while the chrF operates at a character level which is more suitable for short texts such as titles.

<sup>2</sup><https://huggingface.co/mistralai/Mixtral-8x7B-Instruct-v0.1>

prompting using the *prompt template* presented in section 2.2.

## 4 Results and Analysis

	TTL	BP	PD	T.B.D.
<b>EN-PL</b>				
RAND 1-shot	-1.0%	-0.7%	+0.2%	-0.4%
RAND 5-shot	-1.7%	-0.8%	-3.0%	-2.6%
RAG 1 shot	+8.3%	+2.6%	+10.7%	+10.9%
RAG 5 shot	+11.8%	+4.6%	+14.9%	<b>+15.3%</b>
<b>EN-NL</b>				
RAND 1-shot	+1.9%	+1.1%	+3.4%	+2.9%
RAND 5-shot	+1.7%	+2.0%	+5.1%	+4.8%
RAG 1 shot	+3.0%	+2.3%	+7.2%	+7.4%
RAG 5 shot	+6.2%	+3.3%	<b>+12.0%</b>	+11.9%
<b>EN-SE</b>				
RAND 1-shot	+1.8%	+1.6%	+0.9%	+1.6%
RAND 5-shot	+2.8%	+2.4%	+1.7%	+2.4%
RAG 1 shot	+8.8%	+4.6%	+11.2%	+10.8%
RAG 5 shot	+11.5%	+6.4%	<b>+14.0%</b>	+6.4%
<b>DE-CS</b>				
RAND 1-shot	-1.0%	-0.2%	-1.7%	-2.6%
RAND 5-shot	-2.4%	+1.3%	+0.0%	-1.4%
RAG 1 shot	+3.2%	+2.0%	+1.0%	+3.2%
RAG 5 shot	+5.3%	+4.3%	+2.7%	<b>+5.9%</b>
<b>EN-DE</b>				
RAND 1-shot	-1.1%	-0.9%	-0.9%	-0.8%
RAND 5-shot	-0.6%	-1.7%	-4.0%	-0.7%
RAG 1 shot	+1.5%	+1.1%	-0.9%	+0.9%
RAG 5 shot	<b>+2.0%</b>	+1.4%	-1.6%	+1.7%
<b>IT-DE</b>				
RAND 1-shot	+1.8%	+0.0%	+1.0%	+1.0%
RAND 5-shot	+0.2%	+0.0%	+0.5%	+0.3%
RAG 1 shot	+0.1%	+0.2%	+2.6%	+2.7%
RAG 5 shot	-0.6%	+0.3%	+3.4%	<b>+3.6%</b>
<b>TR-DE</b>				
RAND 1-shot	+0.3%	+1.1%	+0.5%	+0.3%
RAND 5-shot	+0.7%	+0.2%	+0.1%	+1.2%
RAG 1 shot	+1.7%	+1.4%	+1.6%	<b>+2.3%</b>
RAG 5 shot	+1.6%	+0.9%	+0.5%	+1.4%

**Table 1:  $\Delta$  chrF% of title translation experiments using against baseline generation. "RAG {1,5} shot" indicate top-1 and 5 retrieved example(s) based on Okapi BM25 similarity score. "RAND {1,5} shot" indicates 1 and 5 randomly chosen example(s). TTL/BP/PD/T.B.D. refer to index built using bilingual titles, bulletpoints, descriptions, and all three domains respectively.**

Table 1 shows the chrF deltas relative to the *baseline generation*. We observe that RAG 1-shot and RAG 5-shot overall improve title translation quality across different data index domains (TTL, BP,

PD, and T.B.D.) used for retrieving few-shot examples. In contrast, configurations which retrieve random examples (RAND 1-shot and RAND 5-shot) generally show decreases in chrF, which indicate a negative impact on title translation quality. Where improvements are seen amongst RAND 1 and 5-shot configurations, the magnitude is much smaller than that of the RAG 1-shot and RAG 5-shot configurations.

Results show RAG 5-shot with examples from all domains (T.B.D. index) has the largest or second largest title translation quality improvement across all language pairs except EN-SE. For language pairs where the target language is not the main language of the LLM, RAG 5-shot demonstrate chrF improvements ranging from +5.9% to +15.3% across DE-CS, EN-SE, EN-NL, and EN-PL. For TR-DE, where the source language is not a main language of the LLM, RAG 1-shot show improvements by up to 2.3% chrF.

RAG 1-shot and RAG 5-shot configurations also demonstrate improvements in chrF for language pairs where both languages are the main supported languages of the LLM. Those language pairs have higher chrF for baseline generation, and also have chrF improvement up to 3.6% for IT-DE and 2.0% for EN-DE.

Notably, all language pairs show the greatest increase in chrF when examples are retrieved from the T.B.D. index which contains all three product information domains (title, bullet points, and description). This suggests that a combined index can provide more useful contextual cues to guide the LLM’s product title translation task.

**Retrieved examples similarity:** We compute the textual similarity between the test source title and examples used as few-shot examples in the prompts as shown in Table 3. Textual similarity is computed using chrF. For RAND-5 and RAG-5 configurations, we report the average of chrF scores across the 5 examples. We observe (1) RAG-1 and RAG-5 retrieved examples with higher similarity scores compared to RAND-1 and RAND-5 which use randomly chosen examples; (2) Given a language pair and one configuration, we observe generally that the RAG data domain with higher similarity score has larger translation quality improvement. However, we also observe RAG-5 with examples from domains T.B.D. do not have highest similarity scores but have the largest or second largest title translation quality improvement for all language pairs except EN-SE. This indicates that retrieved examples from more diverse product information domains can also have positive impact on the title translation generation improvement.

**Title translation improvement:** Table 2 shows examples of generated titles by the RAG LLM method. The three examples show how the approach can improve title translation quality:

**(a) Improved title translation from better understanding of products:** Many product titles for products such as signs, T-shirts, mugs have text printed on the products, and such texts need to stay intact in the translation.

Example A in Table 2 is for a sign with the message “*Posted No Trespass Hunting Fishing Trapping Under Penalty*”. In this example, the message should not be translated since the product includes the phrase verbatim. Without RAG, the LLM translates the entire title to German including the phrase. With RAG, the retrieved example

Example A (EN-DE)	
src	Posted No Trespass Hunting Fishing Trapping Under Penalty 10x7 Aluminum sign
ref	“Posted No Trespass Hunting Fishing Trapping Under Penalty ” 25,4 x 17,8 cm Aluminium-Schild
bl	Angebrachtes No Trespassing-Schild für Jagd, Fischfang und Fallenstellen unter Strafe 10x7 Aluminium
rag5	Aluminiumschild Posted No Trespass Hunting Fishing Trapping under Penalty 10x7, Metall
rag src	No Hunting Sign - Private Property Posted Hunter Keep Out Signs - Aluminum Metal
rag tgt	Schild “No Hunting”, für Privatbesitz, um Jäger abzuhalten, Aluminium, Metall
Example B (EN-DE)	
src	KYB Shock Absorber, Part Number: 343441
ref	KYB Stoßdämpfer, Teilenummer: 343441
bl	Dämpferstabilisator KYB, Teilenummer: 343441
rag5	KYB Stoßdämpfer, Teilenummer: 343441
rag src	KYB 554105, KYB Front Shock Absorber (Single Unit) 554105
rag tgt	KYB 554105, KYB-Stoßdämpfer vorne (Einzeleinheit) 554105
Example C (EN-DE)	
src	mumbi ultra slim case for Nokia 6 2018, protective case in mandala design, transparent pink (ultra slim 0.70 mm)
ref	mumbi UltraSlim Hülle für Nokia 6 2018 Schutzhülle im Mandala-Design, transparent rosa (Ultra Slim - 0,70 mm)
bl	Hülle für Nokia 6 2018 von mumbi, ultra schlanker Schutzcase mit Mandala-Muster, transparente Rosatönung (ultra dünn 0,70 mm)
rag5	mumbi Ultra Slim Hülle für Nokia 6 von 2018, Schutzhülle im Mandala Design, transparente Rosa (ultra slim 0,70 mm)
rag src	mumbi Ultra-Slim Phone Case Ultra-Slim for Huawei P9 Lite Case in Mandala Design Transparent Pink 0.55 mm.
rag tgt	mumbi UltraSlim Hülle für Huawei P9 Lite Schutzhülle im Mandala Design transparent rosa (Ultra Slim - 0.55 mm)

**Table 2: Translation examples using LLM with RAG (rag5) and baseline generation (bl) with no RAG. rag src and rag tgt refer to the retrieved top-one example used as few-shot in the prompt. src, ref refer to the title in the source language and reference title translation respectively**

is for a similar product which illustrates that the message “*No Hunting*” should be preserved in the translation, and the generated title preserves the message in the German translation.

A similar issue also occurs with named entities such as character names, brands etc., where some named entities may overlap with the common vocabulary. An example of such case is character name “*Peppa Pig*” in the product title “*Peppa Pig House Tea Playset*”. Without RAG, the LLM translates it into “*Peppa Schwein*” (German) while the character name is preserved in the translation with RAG.

**(b) Improved translation for product-specific terminology:** For example B in Table 2, the expected German translation of the term “*Shock Absorber*” is “*Stoßdämpfer*”. Without the RAG approach, the LLM instead chooses a less appropriate translation “*Dämpferstabilisator*”. However, when the top-k retrieved examples contain

		TTL	BP	PD	T.B.D
<b>EN-PL</b>					
	RAND 1-shot	14.3	11.0	<u>16.4</u>	11.0
	RAND 5-shot	13.0	11.6	<u>15.0</u>	13.2
	RAG 1-shot	46.4	35.2	<u>52.4</u>	49.1
	RAG 5-shot	40.9	32.6	<u>43.6</u>	39.0
<b>EN-NL</b>					
	RAND 1-shot	<u>16.2</u>	11.1	4.0	4.0
	RAND 5-shot	<u>15.7</u>	14.2	10.2	13.3
	RAG 1-shot	42.5	32.0	<u>54.2</u>	42.6
	RAG 5-shot	38.0	29.2	<u>45.2</u>	37.5
<b>EN-SE</b>					
	RAND 1-shot	6.1	6.2	<u>6.3</u>	6.1
	RAND 5-shot	8.4	<u>8.5</u>	8.3	8.4
	RAG 1-shot	46.4	35.6	<u>52.6</u>	47.8
	RAG 5-shot	40.9	32.9	<u>44.1</u>	39.3
<b>DE-CS</b>					
	RAND 1-shot	12.2	10.0	<u>12.9</u>	10.5
	RAND 5-shot	<u>12.5</u>	8.1	10.8	10.6
	RAG 1-shot	<u>37.9</u>	37.1	35.1	32.1
	RAG 5-shot	32.1	31.6	<u>31.6</u>	29.3
<b>EN-DE</b>					
	RAND 1-shot	<u>13.3</u>	10.2	4.4	10.2
	RAND 5-shot	<u>13.2</u>	10.4	7.9	10.5
	RAG 1-shot	<u>49.6</u>	41.2	24.6	41.1
	RAG 5-shot	<u>43.1</u>	38.1	23.4	34.9
<b>IT-DE</b>					
	RAND 1-shot	<u>16.6</u>	16.4	14.2	15.8
	RAND 5-shot	<u>14.8</u>	12.5	13.7	13.7
	RAG 1-shot	45.7	42.5	<u>49.2</u>	46.5
	RAG 5-shot	40.8	39.3	<u>41.5</u>	40.5
<b>TR-DE</b>					
	RAND 1-shot	10.6	9.4	9.0	<u>10.7</u>
	RAND 5-shot	<u>11.0</u>	8.9	7.4	9.5
	RAG 1-shot	27.5	27.9	<u>29.9</u>	25.3
	RAG 5-shot	23.9	23.8	<u>26.9</u>	23.8

**Table 3: Textual similarity (chrF) between the source titles and the source few-shot examples, TTL/BP/PD/T.B.D. refer to index built using bilingual titles, bulletpoints, descriptions, and all three domains respectively.**

the mapping of “Shock Absorber” to “Stoßdämpfer”, the LLM with RAG can then select the expected and more accurate term in the translation.

**(c) Improved formality and style of title translation:** We have also observed that the RAG approach can improve title quality by conforming to product title formality and style preferred in catalogs. In example C in Table 2, the retrieved example is for a phone case of the same brand but for a different model. The generated

title translation conforms to style conventions used in the retrieved example, such as (1) placing the brand name at the beginning of the title, (2) appending the specific phone model after the brand name, and (3) including details like color. By leveraging these stylistic cues from the retrieved bilingual title pairs, the LLM can generate translations that not only convey the accurate meaning but also adhere to the preferred formatting and conventions expected in product catalogs. This helps ensure a consistent and professional-looking title presentation across the multilingual offerings.

## 5 Related work

Retrieval Augmented Generation (RAG) has proved to be an effective prompting strategy to enhance the accuracy and reliability of large language models (LLMs) [3, 11, 24]. This is due to the inherent limitations of LLMs in accessing specialized knowledge [19]. RAG enables LLMs to combine their generative abilities with additional context and factual information retrieved from external knowledge sources, which has been explored in several prior studies [1, 8, 13, 16, 22].

Retrieval-augmented methods have been successfully applied to various tasks using LLMs with domain-specific specialized knowledge, such as question answering [6, 10, 23]. Previous work has also proposed using RAG with a “text book” of retrieved language-specific usage, syntax, and vocabulary examples to improve translation tasks with LLMs [7]. However, their study focused on lower-resource language pairs, and to the best of our knowledge, there has been no prior work on applying RAG techniques to machine translation of e-commerce product titles and languages commonly used in the e-commerce industry.

## 6 Conclusion

This study presents a retrieval-augmented generation (RAG) approach to enhance e-commerce product title translation using large language models (LLMs). By leveraging the growing bilingual catalog of product content, RAG retrieves similar examples as few-shot prompts to guide the LLM in translating product terminology more accurately and generating higher quality, contextually appropriate translations. Experiments across 7 language pairs demonstrated the RAG approach’s effectiveness, achieving chrF score improvements up to 15.3%, especially for language pairs where the LLM has limited proficiency. The RAG approach can enhance translation by more accurately translating or preserving product-specific terminology, maintaining brands, and adhering to expected formatting and deliver product tile translations essential for multilingual product discovery in e-commerce. This scalable, sustainable approach also allows rapid adaptation as e-commerce continuously introduces new products and language pairs.

## References

- [1] Sweta Agrawal, Chunting Zhou, Mike Lewis, Luke Zettlemoyer, and Marjan Ghazvininejad. 2022. In-context Examples Selection for Machine Translation. arXiv:2212.02437 [cs.CL] <https://arxiv.org/abs/2212.02437>
- [2] Tianchi Bi, Liang Yao, Baosong Yang, Haibo Zhang, Weihua Luo, and Boxing Chen. 2020. Constraint Translation Candidates: A Bridge between Neural Query Translation and Cross-lingual Information Retrieval. arXiv:2010.13658 [cs.CL]
- [3] Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng Wang, and Haofen Wang. 2024. Retrieval-Augmented

- Generation for Large Language Models: A Survey. arXiv:2312.10997 [cs.CL] <https://arxiv.org/abs/2312.10997>
- [4] Nuno M. Guerreiro, Duarte Alves, Jonas Waldendorf, Barry Haddow, Alexandra Birch, Pierre Colombo, and André F. T. Martins. 2023. Hallucinations in Large Multilingual Translation Models. arXiv:2303.16104 [cs.CL] <https://arxiv.org/abs/2303.16104>
- [5] Jyoti Guha and Carmen Heger. 2014. Machine translation for global e-commerce on eBay. In *Proceedings of the AMTA*, Vol. 2. 31–37.
- [6] Gauthier Guinet, Behrooz Omidvar-Tehrani, Anoop Deoras, and Laurent Callot. 2024. Automated Evaluation of Retrieval-Augmented Language Models with Task-Specific Exam Generation. arXiv:2405.13622 [cs.CL] <https://arxiv.org/abs/2405.13622>
- [7] Ping Guo, Yubing Ren, Yue Hu, Yunpeng Li, Jiarui Zhang, Xingsheng Zhang, and Heyan Huang. 2024. Teaching Large Language Models to Translate on Low-resource Languages with Textbook Prompting. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, Nicoletta Calzolari, Min-Yen Kan, Veronique Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue (Eds.), ELRA and ICCL, Torino, Italia, 15685–15697. <https://aclanthology.org/2024.lrec-main.1362>
- [8] Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. REALM: Retrieval-Augmented Language Model Pre-Training. arXiv:2002.08909 [cs.CL] <https://arxiv.org/abs/2002.08909>
- [9] Amr Hendy, Mohamed Abdelrehim, Amr Sharaf, Vikas Raunak, Mohamed Gabr, Hitokazu Matsushita, Young Jin Kim, Mohamed Affy, and Hany Hassan Awadalla. 2023. How Good Are GPT Models at Machine Translation? A Comprehensive Evaluation. arXiv:2302.09210 [cs.CL] <https://arxiv.org/abs/2302.09210>
- [10] Jennifer Hsia, Afreen Shaikh, Zhiruo Wang, and Graham Neubig. 2024. RAGGED: Towards Informed Design of Retrieval Augmented Generation Systems. arXiv:2403.09040 [cs.CL] <https://arxiv.org/abs/2403.09040>
- [11] Wenqi Jiang, Shuai Zhang, Boran Han, Jie Wang, Bernie Wang, and Tim Kraska. 2024. PipeRAG: Fast Retrieval-Augmented Generation via Algorithm-System Co-design. arXiv:2403.05676 [cs.CL] <https://arxiv.org/abs/2403.05676>
- [12] Zhuolin Jiang, Amro El-Jaroudi, William Hartmann, Damianos Karakos, and Lingjun Zhao. 2020. Cross-lingual Information Retrieval with BERT. In *Proceedings of the workshop on Cross-Language Search and Summarization of Text and Speech (CLSSTS2020)*. European Language Resources Association, Marseille, France, 26–31. <https://aclanthology.org/2020.clssts-1.5>
- [13] Zhengbao Jiang, Frank F. Xu, Luyu Gao, Zhiqing Sun, Qian Liu, Jane Dwivedi-Yu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. Active Retrieval Augmented Generation. arXiv:2305.06983 [cs.CL] <https://arxiv.org/abs/2305.06983>
- [14] Tom Kocmi, Christian Federmann, Roman Grundkiewicz, Marcin Junczys-Dowmunt, Hitokazu Matsushita, and Arul Menezes. 2021. To Ship or Not to Ship: An Extensive Evaluation of Automatic Metrics for Machine Translation. In *Proceedings of the Sixth Conference on Machine Translation*, Loic Barrault, Ondrej Bojar, Fethi Bougares, Rajen Chatterjee, Marta R. Costa-jussa, Christian Federmann, Mark Fishel, Alexander Fraser, Markus Freitag, Yvette Graham, Roman Grundkiewicz, Paco Guzman, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Tom Kocmi, Andre Martins, Makoto Morishita, and Christof Monz (Eds.), Association for Computational Linguistics, Online, 478–494. <https://aclanthology.org/2021.wmt-1.57>
- [15] Mike Lowndes and Aditya Vasudevan. 2021. Market Guide for Digital Commerce Search. (2021).
- [16] Hongyin Luo, Yung-Sung Chuang, Yuan Gong, Tianhua Zhang, Yoon Kim, Xixin Wu, Danny Fox, Helen Meng, and James Glass. 2023. SAIL: Search-Augmented Instruction Learning. arXiv:2305.15225 [cs.CL] <https://arxiv.org/abs/2305.15225>
- [17] Jian-Yun Nie. 2010. Cross-language information retrieval. *Synthesis Lectures on Human Language Technologies* 3, 1 (2010), 1–125.
- [18] Maja Popović. 2015. chrF: character n-gram F-score for automatic MT evaluation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*. Association for Computational Linguistics, Lisbon, Portugal, 392–395. <https://doi.org/10.18653/v1/W15-3049>
- [19] Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. In-Context Retrieval-Augmented Language Models. arXiv:2302.00083 [cs.CL] <https://arxiv.org/abs/2302.00083>
- [20] Andreas Rücklé, Krishnkant Swarnkar, and Iryna Gurevych. 2019. Improved Cross-Lingual Question Retrieval for Community Question Answering. In *The World Wide Web Conference (San Francisco, CA, USA) (WWW '19)*. Association for Computing Machinery, New York, NY, USA, 3179–3186. <https://doi.org/10.1145/3308558.3313502>
- [21] Shadi Saleh and Pavel Pecina. 2020. Document Translation vs. Query Translation for Cross-Lingual Information Retrieval in the Medical Domain. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 6849–6860. <https://doi.org/10.18653/v1/2020.acl-main.613>
- [22] Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Rich James, Mike Lewis, Luke Zettlemoyer, and Wen tau Yih. 2023. REPLUG: Retrieval-Augmented Black-Box Language Models. arXiv:2301.12652 [cs.CL] <https://arxiv.org/abs/2301.12652>
- [23] Shamane Siriwardhana, Rivindu Weerasekera, Elliott Wen, Tharindu Kaluarachchi, Rajib Rana, and Suranga Nanayakkara. 2023. Improving the Domain Adaptation of Retrieval Augmented Generation (RAG) Models for Open Domain Question Answering. *Transactions of the Association for Computational Linguistics* 11 (01 2023), 1–17. [https://doi.org/10.1162/tacl\\_a\\_00530](https://doi.org/10.1162/tacl_a_00530) arXiv:[https://direct.mit.edu/tacl/article-pdf/doi/10.1162/tacl\\_a\\_00530/2067834/tacl\\_a\\_00530.pdf](https://direct.mit.edu/tacl/article-pdf/doi/10.1162/tacl_a_00530/2067834/tacl_a_00530.pdf)
- [24] Weihang Su, Yichen Tang, Qingyao Ai, Zhijing Wu, and Yiqun Liu. 2024. DRAGIN: Dynamic Retrieval Augmented Generation based on the Information Needs of Large Language Models. arXiv:2403.10081 [cs.CL] <https://arxiv.org/abs/2403.10081>
- [25] Haifeng Wang, Hua Wu, Zhongjun He, Liang Huang, and Kenneth Ward Church. 2021. Progress in Machine Translation. *Engineering* (2021).
- [26] Andy Way. 2013. Traditional and emerging use-cases for machine translation. *Proceedings of Translating and the Computer* 35 (2013), 12.
- [27] Bryan Zhang, Salah Danial, and Stephan Walter. 2023. Enhancing Arabic Machine Translation for E-commerce Product Information: Data Quality Challenges and Innovative Selection Approaches. In *Proceedings of ArabicNLP 2023*, Hassan Sawaf, Samhaa El-Beltagy, Wajdi Zaghouni, Walid Magdy, Ahmed Abdellali, Nadi Tomeh, Ibrahim Abu Farha, Nizar Habash, Salam Khalifa, Amr Keleg, Hatem Haddad, Imed Zitouni, Khalil Mrini, and Rawan Almatham (Eds.), Association for Computational Linguistics, Singapore (Hybrid), 150–157. <https://doi.org/10.18653/v1/2023.arabicnlp-1.13>
- [28] Bryan Zhang, Taichi Nakatani, Daniel Vidal Hussey, Stephan Walter, and Liling Tan. 2024. Don't Just Translate, Summarize Too: Cross-lingual Product Title Generation in E-commerce. In *Proceedings of the Seventh Workshop on e-Commerce and NLP @ LREC-COLING 2024*, Shervin Malmasi, Besnik Fetahu, Nicola Ueffing, Oleg Rokhlenko, Eugene Agichtein, and Ido Guy (Eds.), ELRA and ICCL, Torino, Italia, 58–64. <https://aclanthology.org/2024.ecnlp-1.6>
- [29] Mingyang Zhou, Runxiang Cheng, Yong Jae Lee, and Zhou Yu. 2018. A Visual Attention Grounding Neural Model for Multimodal Machine Translation. *CoRR* abs/1808.08266 (2018). arXiv:1808.08266 <http://arxiv.org/abs/1808.08266>