

MIND: Multimodal Shopping Intention Distillation from Large Vision-language Models for E-commerce Purchase Understanding

Anonymous ACL submission

Abstract

Improving user experience and providing personalized search results in E-commerce platforms heavily rely on understanding purchase intention. However, existing methods for acquiring large-scale intentions bank on distilling large language models with human annotation for verification. Such an approach tends to generate product-centric intentions, overlook valuable visual information from product images, and incurs high costs for scalability. To address these issues, we introduce MIND, a multimodal framework that allows Large Vision-Language Models (LVLMs) to infer purchase intentions from multimodal product metadata and prioritize human-centric ones. Using Amazon Review data, we apply MIND and create a multimodal intention knowledge base, which contains 1,264,441 million intentions derived from 126,142 co-buy shopping records across 107,215 products. Extensive human evaluations demonstrate the high plausibility and typicality of our obtained intentions and validate the effectiveness of our distillation framework and filtering mechanism. Additional experiments reveal that our obtained intentions significantly enhance large language models in two intention comprehension tasks.

1 Introduction

Understanding customers' intentions behind their purchase behaviors remains crucial in E-commerce as it potentially benefits several downstream tasks, such as product recommendation (Grbovic et al., 2015; Zhao et al., 2014; Li et al., 2020) and search query answering (Zhao et al., 2019; Hirsch et al., 2020). Unlike traditional factual knowledge related to products, intentions are implicit mental states of customers, which typically require commonsense knowledge to understand and reason upon (Bratman, 1984). For example, in Figure 1, the intentions of purchasing a mouse and a keyboard can be *they are very useful to computer users*, which

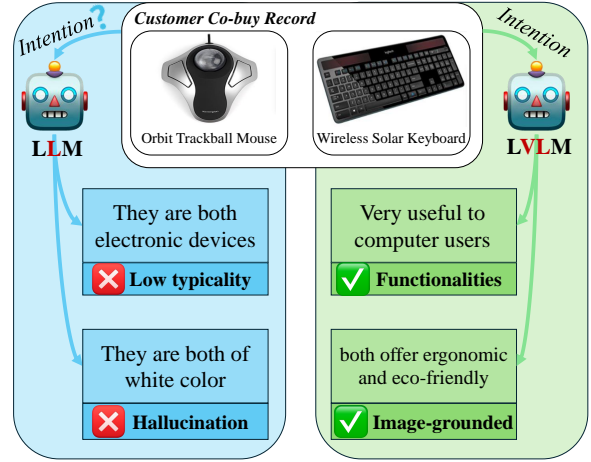


Figure 1: Examples showing the process of distilling purchase intentions from **large language models** and **large vision-language models**. Without product images, **large language models** tend to generate intentions with low typicality and hallucinated facts, while leveraging **large vision-language models** resolve such issue.

is not mentioned either in the customer's query or products' metadata. Thus, due to such implicitness, it is infeasible to perform large-scale automatic extraction from text to obtain them.

To combat this, Yu et al. (2023) proposed to distill purchase intentions from large language models, such as OPT (Zhang et al., 2022), by prompting them with real purchasing records and relevant product metadata. Human-in-the-loop annotations are also carried out to verify the plausibility and typicality of the generated intentions and train a discriminator for large-scale critic filtering. Yu et al. (2024) further entangled human annotations with instruction tuning to align the distilled intentions with a human-centric perspective. While these works provide a straightforward approach to intention acquisition, several limitations still persist.

First, previous works on E-commerce intention knowledge base construction have solely focused on the text modality, thereby sacrificing significant supervision signals from visual modalities, such as

product images. This oversight hinders the model from obtaining a more comprehensive understanding of the product, consequently compromising the quality of the generated intentions, as demonstrated in the left lower part of Figure 1. Furthermore, recent work has shown that intentions derived using current distillation methods exhibit bias towards product-centric aspects, excessively emphasizing product properties and metadata (Zhou et al., 2024). Consequently, interactions between the products and customers, including potential use cases and features of interest to customers, are absent from the derived intentions, despite being fundamental in facilitating customers’ shopping experience. Finally, human annotations are heavily deployed in current intention collection methods, which serve as a critical step in controlling the quality of the generated results. This poses a challenge towards constructing scalable yet diverse intention knowledge bases with minimum human supervision cost.

To address these issues, we propose MIND, a **Multimodal Shopping Intention Distillation** framework. MIND instructs Large Vision-Language Models (LVLMs) to generate purchase intentions in a three-step manner, based on real user co-buy records and product metadata. Specifically, we select LLaVa (Liu et al., 2023a) as a representative LVLM and incorporate both visual information from the product images and text information from the product name into the generation process. To better align the generated raw intentions with human preferences and alleviate human annotation costs for further quality control, we propose a human-centric role-aware mechanism. This mechanism first instructs LLaVa to discover similar features between the products and then imitates a customer agent to decide whether the products would be bought together under previously generated intentions.

By applying MIND to a subset of the Amazon Review Dataset (Ni et al., 2019), we construct a multimodal intention knowledge base. It features 1.26 million of intentions over 126,142 co-buy shopping records across 107,215 products. Human evaluations further confirm: (1) the exceptional quality of our generated intentions, which have higher plausibility and typicality than previous generation methods, and (2) the effectiveness of our proposed human-centric role-aware mechanism. Furthermore, we apply our generated intentions to two downstream tasks in the IntentionQA

benchmark (Ding et al., 2024), which evaluates a language model’s abilities to discriminate and utilize purchase intentions. Extensive experiments show that distilling our generated intentions into large language models’ provide substantial benefits on both tasks via fine-tuning. Further analyses reveal the strengths and mechanisms behind of our innovative filter module and confirm the superiority of MIND when compared to previous intention acquisition methods.

2 Related Works

2.1 Shopping Intention in E-commerce

Shopping intention is an implicit mental state that motivates purchase-related behaviors from the customer’s perspective (Koo and Ju, 2010). Various studies have been conducted to examine the impact of consumer shopping intentions on downstream applications (Dai et al., 2006; Zhang et al., 2016; Hao et al., 2022). Recently, Ni et al. (2019) suggested using customer reviews to investigate the underlying purchase intentions in consumer purchase behavior and created a large-scale review dataset based on Amazon. Building upon this, Yu et al. (2023) proposed FolkScope, which aims to guide LLMs in generating user co-buy intentions for different product pairs by grounding them in ConceptNet relations (Speer et al., 2017). While human evaluations confirmed its effectiveness, Zhou et al. (2024) argued that it not only remains expensive to scale up but also fails to align the resulting shopping intentions with human preferences, which encompass a wide range of factors beyond product properties and similarities. To tackle these issues, in our work, we propose MIND, a framework that undermines online co-buy intentions and aligns better with human perceptions.

2.2 Multimodal Knowledge Distillation

Since VLMs have yield significant advance recently (Liu et al., 2023d; Li et al., 2023; Zhu et al., 2023), distilling domain-specific knowledge from them has become an effective yet cost-friendly trend in multimodal studies (Liu et al., 2023c; Lu et al., 2024; Jin et al., 2021). Liu et al. (2023c) proposed a framework that applies self-distillation to stimulate the pre-train process of BERT to improve its performance in E-commerce product understanding tasks. Lu et al. (2024) similarly instructed MiniGPT4 (Zhu et al., 2023) to generate user intention form social media posts text and its

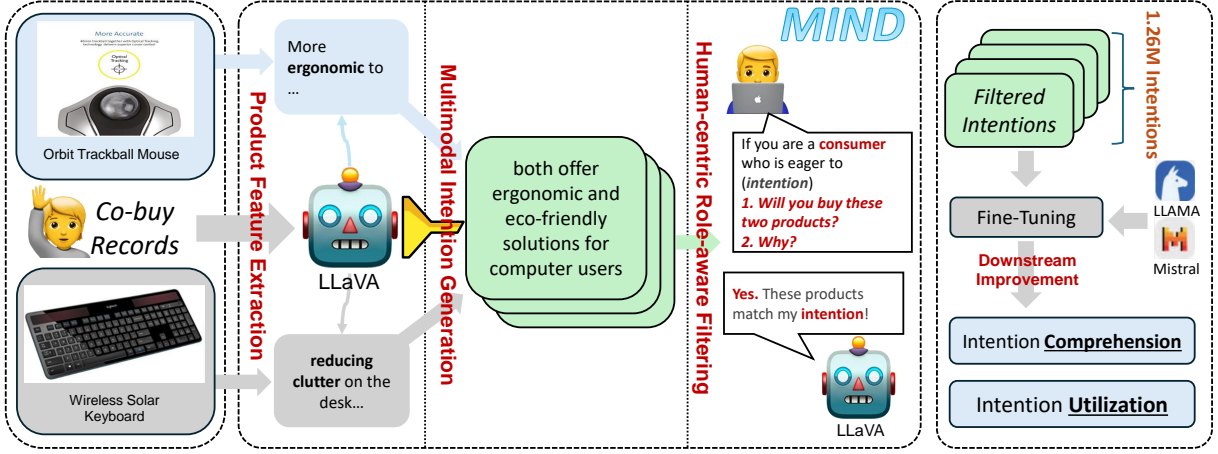


Figure 2: An overview of MIND. We first extract features from products in real-world co-buy records, generate intentions multimodally, and apply a human-centric role-aware filter for quality optimization.

associated images. Jin et al. (2021) also designed a framework to instruct the student model to imitate teacher model’s behavior, which successfully preserved the teacher model’s capabilities with fewer parameters. In our work, we share the same aspiration and leverage distillation as a tool for data collection that provides downstream benefits in the E-commerce domain. Specifically, we designed a framework to distill E-commerce intentions from LLaVa (Liu et al., 2023a) and construct a comprehensive intention knowledge base based on the resulted generations.

3 The MIND Framework

3.1 Overview of MIND

Following Yu et al. (2023), the objective of MIND is formulated as a text generation task. Given a record that shows a customer’s co-buy (purchasing together) of two products, along with the detailed metadata of both products, MIND aims to generate the intentions behind such purchase behaviors that best align with the customer’s mental state during the purchase, which includes their beliefs, desires, and intents (Georgeff et al., 1999).

Formally, for a given co-buy record, we define the two products as p^1 and p^2 , along with their associated images $\{p_i^1, p_i^2\}$, and features and attributes $\{p_f^1, p_f^2\}$. MIND aims to leverage a LVLM F to generate the intentions $I(p^1, p^2)$ of purchasing both products based on a pre-defined commonsense relation r , denoted as $I(p^1, p^2, r) = F(p_f^1, p_f^2, p_i^1, p_i^2, r)$. In this paper, we follow Yu et al. (2023, 2024) and use relations from ConceptNet (Speer et al., 2017) to model the intentions. LLaVa-1.5-13b (Liu et al., 2023a) is used as the

LVLM F .

To achieve this objective, we design three sequentially connected steps within MIND, which are shown in Figure 2. These steps are termed as: (1) product feature extraction; (2) co-buy intention generation; (3) human-centric role-aware filtering. Together, they form a collective pipeline for systematic intention acquisition without the need for human supervision and quality filtering.

3.2 Source Data Collection

We utilize the Amazon Review Data released by Ni et al. (2019), which contains millions of products from 18 domains. Each product is accompanied by detailed reviews, co-buy records, and metadata, including its product title, features, attributes, and images provided by the retailer. Following Yu et al. (2023), we select products from the *Electronics* and *Clothing, Shoes and Jewelry* domains as representative products to demonstrate the effectiveness of MIND. To fit our framework, we filter out products without accessible images that may have been removed from the Amazon platform.

3.3 Product Feature Extraction

We begin processing the collected products by first extracting key features with the aid of LVLMs. This is motivated by our observations that product descriptions and attributes, input by retailers, tend to be noisy and unorganized, probably for promotion and style organization purposes. Thus, we explicitly instructs LVLMs to augment source product metadata by extracting implicit features from each product’s image and title by leveraging a zero-shot prompt:

Prompt Template for Product Feature Extraction

Visual Input: p_i

Textual Input: <Instruction>. Given the product shown in the image: p_f , generate additional features by focusing on the product’s attribute, design, and quality.

where <Instruction> is a detailed task instruction, and p_i, p_f are the respective image and details (title, descriptions, etc.) of the product. This enables LVLM to comprehend the product from both visual and textual modalities, thereby providing us with a richer set of features that complements those provided by the retailers.

3.4 Co-buy Intention Generation

Then, for each co-buy pair of products (p^1, p^2), we provide LVLM with the acquired features together with all details of both products, and instructs it again to reason the intentions for purchasing them simultaneously. Specifically, we follow Yu et al. (2023) and leverage 20 commonsense relations from ConceptNet (Speer et al., 2017) as waymarks to lead LVLM in generating purchase intentions with controllable commonsense groundings. Similar to the previous step, a zero-shot prompt is used:

Prompt Template for Intention Generation

Visual Input: p_i^1, p_i^2

Textual Input: <Instruction>. A customer purchased a pair of products, as shown in the images. They are: p_f^1, p_f^2 . Act as the customer and infer a potential intention behind such purchase. Start the intention with <Relation>.

Where <Instruction> is a detailed task instruction and <Relation> is the corresponding text template of a commonsense relation from ConceptNet. For every relation, we generate only one intention per pair of products due to the large amount of products and co-buy records. However, this is not restricted and can easily scale up.

3.5 Human-centric Role-aware Filtering

To effectively manage a large amount of purchase intentions, quality control measures have become imperative. While previous works relied on human annotations for this purpose, recent works (Zhou et al., 2024) show that co-buy intentions generated by LLMs, despite undergoing human filtering, still fail in capturing the customers’ mental states but

rather focus on factual similarities of the products, as demonstrated in Figure 1. This phenomenon, referred to as “product-centric,” restricts the potential downstream applications of the generated intentions. To address both issues, inspired by recent works on theory-of-mind (Kosinski, 2023), we propose to incorporate a filtering module, powered by a LVLM, after the generation process. We instruct the LVLM to assume the role of an E-commerce customer and provide it with a generated intention as the objective in the customer’s mental state. Based on this intention, we present the LVLM with a pair of products and ask it to first determine whether the intention successfully motivates the purchase behavior and then generate a rationale to support its decision. This process simulates a real-world scenario where the LVLM functions as a customer, making purchase decisions. By filtering intentions that result in a positive response for purchasing, we obtain intentions that are “human-centric” in the sense that they satisfy the mental state of an agent that is aware of its role as a customer. We term this approach as *human-centric role-aware filtering*, which serves as an automatic filter to replace manual annotations. We apply this module to all the intentions we collected in previous steps and select the product-intention pairs that are accepted by the module as the final outcomes of our framework. Detailed prompts are provided in Appendix A.

4 Intrinsic Evaluations

By applying MIND to products we collected from Amazon Reviews (Ni et al., 2019), we construct a multimodal intention knowledge base, with statistics shown in Table 1. In total, 1.26 million intentions are preserved after applying our proposed filtering module, spanning across 20 relations. Therefore, in this section, we first evaluate MIND intrinsically by examining the quality of the generated intentions and the effectiveness of our proposed filter module through human annotation.

4.1 Annotation Setup

We hire human annotators from the Amazon Mechanical Turk platform to evaluate the generated intentions. For strict quality control, we only invite workers satisfying the following requirements: 1) at least 1K HITs approved, and 2) at least 95% approval rate. Then, we conduct two rounds of qualification rounds using a qualification question

set crafted by authors of this paper, which includes both straightforward and tricky questions. Over 600 workers participated and only 90 (15%) of them are deemed qualified by achieving over 87% accuracy. For a generated intention, we task each worker to evaluate four aspects:

- **Plausibility** refers to the degree to which an intention of a co-buy purchase appears correct and reasonable given both products.

- **Typicality** evaluates how well the intention reflects a specific feature that causes the user behavior, which emphasizes on *informativeness* and *causality* (Yu et al., 2023).

- **Human-centric** evaluates the extent to which the intention considers and aligns with the mental state and preferences of a human customer.

- **Filter rationale** evaluates the correctness of the reasoning or justification provided by the filtering module for accepting or rejecting a product-intention pair.

For each aspect, we ask the annotators to rate them as a binary classification task. A random sample of 5,000 generated intentions are annotated, and the final vote is determined by the majority vote from three annotators.

4.2 Annotation Results

The results of the annotations are presented in Table 1. The annotators achieved a pairwise agreement of 73.1% and a Fleiss’s κ (Fleiss, 1971) of 0.56, indicating satisfactory internal agreement. The results reveal that MIND effectively generates purchase intentions that are both highly plausible (94% on average) and typical (90% on average) across all relations. This indicates the strong product understanding and intention reasoning capabilities of MIND. Additionally, our proposed human-centric role-aware filter correctly identifies 82% of intentions on average, with 80% of them having appropriate justifications for filtering. These high percentages further validate the effectiveness of our proposed method, which serves as a cost-efficient and highly reliable quality control measure, replacing the need for human annotations.

5 Experiments and Analyses

In this section, we first study the downstream benefits brought by intentions generated by MIND. Then, we conduct in-depth analyses of the human-centric role-aware filter in MIND and compare MIND with previous intention acquisition

Relation	#Int.	Pla.	Typ.	Fil.	Rat.
Effect	97,047	0.90	0.83	0.73	0.70
MannerOf	50,563	0.93	0.89	0.83	0.82
isA	62,069	0.94	0.88	0.82	0.80
Other	545	0.94	0.90	0.79	0.75
MadeOf	40,593	0.95	0.92	0.85	0.82
SimilarTo	63,558	0.94	0.87	0.83	0.80
UsedFor	52,383	0.94	0.88	0.81	0.79
Can	90,392	0.95	0.91	0.82	0.78
CauseDesire	95,097	0.94	0.90	0.82	0.80
RelatedTo	64,152	0.93	0.89	0.81	0.79
PartOf	81,230	0.92	0.87	0.79	0.77
Open	122,296	0.93	0.89	0.83	0.82
CreatedBy	35,723	0.94	0.88	0.78	0.76
DeriveFrom	60,347	0.95	0.89	0.80	0.77
DefinedAs	51,680	0.96	0.92	0.84	0.84
PropertyOf	57,947	0.97	0.90	0.83	0.82
CapableOf	86,772	0.95	0.90	0.82	0.82
Cause	61,860	0.95	0.92	0.83	0.82
SymbolOf	64,477	0.95	0.92	0.84	0.82
DistinctFrom	27,710	0.94	0.89	0.84	0.83
Total	1,264,441	0.94	0.90	0.82	0.80

Table 1: Statistics of the intention knowledge base constructed via MIND and human annotation results.

method.

5.1 Evaluation Setup

We explore the effectiveness of MIND on the IntentionQA benchmark (Ding et al., 2024), a comprehensive multiple-choice question answering dataset comprising two challenging subtasks that require language models to comprehend and utilize intentions in E-commerce scenarios accurately. The first task assesses LLMs’ capability in accurately inferring the intention given a co-buy product pair together with 3 distractors sampled from other product pairs, while the second task evaluates LLMs’ capability in utilizing purchase intention to make reasonable product recommendation by selecting the product that best aligns with the user’s intention from four choices. While existing results show that language models struggle with both tasks, we aim to examine whether MIND can enhance LLMs’ intention understanding capabilities through fine-tuning. Specifically, from all intentions generated by MIND, we transform them into instruction-following format via natural language templates following Zhou et al. (2023). Then, we fine-tune LLAMA2-7B-chat (Jiang et al., 2023) and Mistral-7B-Instruct-v0.2 (Touvron et al., 2023) on the retrieved data as a type of knowledge injection. They are then evaluated in a zero-shot manner by being prompted to select the most plausible choice for every QA pair in IntentionQA. Accuracy is used as the evaluation metric.

Methods	Backbone	INTENTION UNDERSTANDING				INTENTION UTILIZATION			
		Easy	Medium	Hard	Avg.	Easy	Medium	Hard	Avg.
Random	-	25.00	25.00	25.00	25.00	25.00	25.00	25.00	25.00
Majority Vote	-	26.37	25.24	26.27	26.15	25.97	28.57	28.57	26.60
PTLM	RoBERTa-Large 214M	41.46	41.98	38.98	41.43	54.95	35.06	30.08	49.84
	DeBERTa-v3-Large 435M	36.40	38.72	37.62	36.90	26.52	29.35	32.33	27.39
	T5-v1.1-xxl 11B	24.84	25.47	25.42	24.99	26.71	26.23	25.56	26.55
	Flan-T5-xxl 11B	75.98	73.58	63.56	74.88	79.26	81.82	81.95	79.89
	T0-pp 11B	71.70	68.87	64.41	70.78	77.11	76.10	78.20	76.99
Commonsense	HyKAS 435M	71.81	67.17	46.69	69.61	47.02	45.97	48.12	46.90
	CAR 435M	73.69	71.46	54.38	72.20	36.18	43.12	44.36	37.94
	CANDLE 435M	74.34	70.75	52.54	72.52	35.94	43.90	43.61	37.84
	VERA 11B	69.82	70.52	61.02	69.49	59.20	58.18	64.66	59.36
	VERA-CANDLE 11B	70.59	71.33	63.41	70.02	62.18	60.13	66.13	61.81
Open LLM	LLAMA2-7B-chat	64.98	66.54	53.85	64.61	59.90	54.86	47.37	58.04
	LLAMA2-13B-chat	69.63	63.96	60.78	68.06	45.53	41.95	39.71	44.52
	Gemma-7B-instruct	65.55	64.31	52.04	64.61	33.18	36.01	41.51	34.20
	Mistral-7B-Instruct-v0.2	76.57	74.53	63.56	75.50	59.78	62.60	65.41	60.64
MIND Distilled	LLAMA2-7B-chat	65.78	64.61	55.75	66.15	59.43	57.13	60.03	59.04
	Mistral-7B-Instruct-v0.2	78.57	74.31	80.89	76.97	61.14	65.42	62.16	62.02
LLM API	ChatGPT	75.06	73.76	68.64	74.48	80.74	76.62	68.42	79.23
	ChatGPT (CoT)	76.07	74.53	63.56	75.12	78.89	75.32	78.20	78.21
	ChatGPT (CoT-SC)	76.51	73.82	63.56	75.32	85.72	77.14	82.71	83.99
	GPT 4	78.12	75.41	66.10	76.97	86.03	82.34	84.96	85.30
	GPT 4 (CoT)	78.12	75.41	66.10	76.97	86.03	82.34	84.96	85.30
	GPT 4 (CoT-SC)	78.80	72.88	65.25	76.97	84.00	80.78	84.96	83.48
Human	-	89.96	90.00	80.96	89.33	95.50	85.19	100.0	94.00

Table 2: Evaluation results (Accuracy%) of various language models on both tasks of the IntentionQA benchmark.

5.2 Baselines

For both tasks, we first incorporate random and majority voting to reflect the characteristics of the benchmark. Five Pre-Trained Language Models (PTLMs) are included: RoBERTa (Liu et al., 2019), DeBERTa-v3 (He et al., 2023), T0 (Sanh et al., 2022), T5 (Raffel et al., 2020), and Flan-T5 (Chung et al., 2022). Then, performances by five commonsense-injected PTLMs are also reported, including HyKAS (Ma et al., 2021), CAR (Wang et al., 2023a), VERA (Liu et al., 2023b), CANDLE (Wang et al., 2024), and VERA-CANDLE. We also report the performances of several LLMs, such as LLAMA2 (Touvron et al., 2023), Gemma (Mesnard et al., 2024), Mistral (Jiang et al., 2023), ChatGPT (OpenAI, 2022), and GPT-4 (OpenAI, 2023). For the latter two, we also adopt Chain-of-Thought (CoT; Wei et al., 2022) and CoT with Self-Consistency (CoT-SC; Wang et al., 2023b) prompting. Prompts are shown in Appendix A.

5.3 Results

The results are presented in Table 2, demonstrating significant improvements in both tasks when LLMs

are fine-tuned on intentions generated by MIND. For instance, LLAMA2 achieves accuracy gains of 1.54% and 1.00% for both tasks, respectively. Notably, Mistral yields a remarkable performance gain that even becomes comparable to GPT-4, despite having a significantly lower number of parameters. However, for the intention utilization task, while both fine-tuned LLMs show performance improvements, they still fall behind GPT-4. One potential reason for this gap could be the misalignment between the fine-tuning objective and the evaluated ability of the task, which involves generating intentions for a pair of products and selecting a product based on a given intention. Nevertheless, these results underscore the effectiveness and efficiency of MIND in enhancing LLMs’ capabilities in E-commerce intention comprehension and utilization.

5.4 Analyses

In this section, we first study the role of our proposed human-centric filtering module by analyzing the error types of the discarded intentions and studying its ablation. Then, we conduct empirical and case comparisons to demonstrate the effective-




Figure 3: Ablation results on IntentionQA tasks by Mistral-7B distilled on intentions generated by MIND with/without filtering.




Figure 4: Relation-wise comparison of typicality scores across all relations between MIND and FolkScope.

ness of MIND over FolkScope.

5.4.1 Impact of Human-centric Filtering

While human annotation results in Section 4.2 show that, after filtering, most of the remaining intentions are highly plausible and typical, we observe that only 46.7% generations passed our proposed filtering module as the last step of MIND. Thus, in this section, we first study the role of such human-centric filtering by looking into the causes of why the intentions get discarded, and further seek insights to resolve such a high filtering loss. To achieve this, we randomly sample 200 intentions that are abandoned by MIND during the last step and manually annotate the reasons behind based on the rationale provided by the LVLM. Three types of errors are observed and they are categorized as:

- 81.0% of the filtered intentions, while plausible, do not provide strong enough evidence to motivate a LVLM agent to execute the purchase behavior for two products. For example, the intention “*they both are related to home audio systems*” for purchasing a pair of audio adapters lacks customer interaction and solely focuses on the products themselves. A more appropriate intention, for example, “*they both are able to help in connecting audio devices,*” would retain a stronger bond between the products and customers by aligning with their functionalities.
- 13.0% of the intentions result from misjudgment by the LVLM, where the agent fails to make the correct decision despite the intention being sufficiently plausible and typical. This highlights the need for future improvements, including a more refined filter to enhance our framework.
- 6.0% of the intentions are discarded due to being implausible or containing factual errors that do not align with the products.

Overall, 87% of intentions are being properly discarded, which is considerably high for an automatic filter without human supervision.

5.4.2 Ablation Study

We then study the ablation of MIND by focusing on the role of our proposed human-centric role-aware filter mechanism in its impact toward quality control of the generated intentions. Specifically, we leverage the IntentionQA (Ding et al., 2024) as the evaluation benchmark and separately train two models on (1) intentions that are filtered by our proposed mechanism (*w. filter*) and (2) intentions without filtering (*w.o. filter*). All setups follow the same as described in Section 5.1, and we use Mistral-instruct-7B-v0.2 as the backbone and train it using a unified hyper-parameter setting. The results are plotted in Figure 3. We observe that, without filtering, performances on both tasks across all difficulty levels drop significantly, which is possibly due to the inclusion of more noisy intentions in the training data. This shows that our proposed filtering module is indeed functioning well in controlling high-quality intentions and is beneficial to downstream tasks.

5.4.3 Relation-wise Comparisons Against FolkScope

We then compare MIND against FolkScope, the previous state-of-the-art method for large-scale intention acquisition, by analyzing the typicality distribution of intentions across all relations. Specifically, we adopt the same annotation protocols designed by Yu et al. (2023) and transfer our annotation results into a four-point Likert scale (Joshi et al., 2015). Then, for each relation, we compute the average typicality scores among all intentions and plot them for comparison, as shown in Fig-

Item1	Item2	Intentions
Samsung SmartCam HD Pro	Samsung SmartThings Smart Home Hub	They are designed to work together in a smart home system They are derived from the same category .
Clarks Women’s Ankle Bootie	The Sak Kendra Hobo Shoulder Bag	The consumer is looking for a stylish and functional combination for their daily activities. They both are a manner of ‘Women’s Shoes’ and ‘Women’s Handbags’ respectively.
Western Party Mustaches	Forum Novelties Adult Cowboy Costume Vest	They are both part of a costume or a themed party . They both are a part of the ‘Adult Costume’ category .
Columbia Women’s Loveland Shorty Omni-Heat Snow Boot	Columbia Sportswear Women’s Thermarator Glove	They are designed to keep the wearer warm and comfortable during cold weather conditions They both are a part of the Columbia brand .
Banded Arc Welded Waterproof Backpack Polyester	Banded Deluxe UFS Fleece Face Mask	They are both used for outdoor activities and protection from harsh weather conditions . They are both used for outdoor activities .
Novelty Video Games Sarcastic Mens Very Funny T Shirt	Super Mario All Over Character 100 Score Print Socks	They both appeal to gamers and video game enthusiasts . They both are of high quality and are really cool .
Women’s Exaggerative Vintage Leafs Shape Necklace Sets	Plated Oil Drip Rhinestone Flower Necklace Earring Sets	They both have a vintage-inspired design and feature colorful flowers and beads They both have a property of ‘High Quality’ .
Xbox 360 4gb Kinect Bundle	Controller Charger - Xbox 360	They both cater to the needs of Xbox 360 gamers . They both are made of plastic .

Table 3: Case studies of purchase intentions generated by MIND and FolkScope. Intentions generated by MIND are highlighted in **blue** and those generated by FolkScope are marked in **green**.

ure 4. From the plot, we observe that intentions generated by MIND exhibit higher typicality scores across nearly all relations compared to those generated by FolkScope, which further demonstrates the superiority of MIND.

5.4.4 Case Studies Against FolkScope

Aside from empirical analyses above, we further show the advantages of MIND over FolkScope through additional case studies to highlight key benefits of MIND. To this end, we randomly selected 7 pairs of co-buy products and compared the intentions generated by both frameworks, as shown in Table 3. Our findings from the table indicate that MIND-generated intentions exhibit a stronger focus on the usage and functionalities that potentially fulfill customers’ needs and intentions when purchasing these products. Conversely, intentions generated by FolkScope tend to be biased towards properties and features that can be easily inferred from the product titles, which are of lesser interest to customers’ shopping intentions. Take the second row in Table 3 as an example. The intentions *both are “Women’s Shoes” and “Women’s Handbags”* generated by FolkScope merely represent an aggregation of the product categories for the two items. In contrast, MIND produces intentions

such as *looking for stylish and functional combination for daily activities*, which better captures a customer’s intention when shopping for both products. This example further reinforces our previous conclusions that MIND can generate intentions that are more human-centric and better reflect the customers’ intentions as mental states.

6 Conclusions

In this work, we present MIND, a multimodal distillation framework for enhancing E-commerce purchase understanding by automating the pipeline of intention generation and quality filtering via multiple-step instructions over LVLMS. By applying MIND to real-world E-commerce data, we construct the very first multimodal purchase intention knowledge base featuring over 1.2 million intentions. These intentions have been proven to be invaluable in distilling student models that exhibit improved performance in E-commerce intention comprehension and utilization tasks. Further analyses reveal the effectiveness of MIND by validating the proposed filtering mechanism and highlighting the strengths of MIND in comparison to FolkScope. Our work sheds light on improving large-scale E-commerce intention acquisition and application.

Limitations

First, MIND generates intention by leveraging several zero-shot prompts without additional exemplars. This decision is made as we observe that few-shot prompts may “guide” LVLM to generate intentions that tend to be similar to the provided exemplars, which harms diversity. However, it remains an open question whether more advanced prompting methods (Song et al., 2023; Parnami and Lee, 2022) would help in the generation process. It’s also worth noting that the LVLM used in our work may be outdated as new products show up on E-commerce platforms. However, switching LLaVa to more up-to-date LVLMs, preferably pre-trained on E-commerce data, can address this concern. Finally, MIND utilizes an automatically functioning filter as quality control. While we have shown its effectiveness, it remains challenging to effectively regulate the filter mechanism to be either lenient or strict. Further investigation is required to provide insights into the alignment between the values of VLMs and the real world, enhancing our understanding of them.

Ethics Statement

To avoid generating harmful intentions and toxic filter rationales in MIND, we recruit 4 expert annotators who are graduate students specializing in multimodality and natural language processing to evaluate the generated intentions and rationales. We ask all experts to go through 200 sampled data and no harmful contents are reported. The crowd-sourced annotators are paid a wage that complies with the local law. The expert annotators involved in this research are knowledgeable about the annotation protocol and the intended utilization of their annotations. They are willingly to contribute without expecting any compensation. The training and evaluation datasets utilized in this study are publicly available, anonymized, and shared under open-access licenses for research purposes, adhering to their intended usage. Thus, we believe this paper does not yield any ethical issue.

References

- Michael Bratman. 1984. Two faces of intention. *The Philosophical Review*, 93(3):375–405.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang,

Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Y. Zhao, Yanping Huang, Andrew M. Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models](#). *CoRR*, abs/2210.11416.

Honghua (Kathy) Dai, Lingzhi Zhao, Zaiqing Nie, Jirong Wen, Lee Wang, and Ying Li. 2006. [Detecting online commercial intention \(OCI\)](#). In *Proceedings of the 15th international conference on World Wide Web, WWW 2006, Edinburgh, Scotland, UK, May 23-26, 2006*, pages 829–837. ACM.

Wenxuan Ding, Weiqi Wang, Sze Heng Douglas Kwok, Minghao Liu, Tianqing Fang, Jiaxin Bai, Xin Liu, Changlong Yu, Zheng Li, Chen Luo, Bing Yin, Junxian He, and Yangqiu Song. 2024. [Intentionqa: A benchmark for evaluating purchase intention comprehension abilities of large language models in e-commerce](#).

Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.

Michael Georgeff, Barney Pell, Martha Pollack, Milind Tambe, and Michael Wooldridge. 1999. The belief-desire-intention model of agency. In *Intelligent Agents V: Agents Theories, Architectures, and Languages: 5th International Workshop, ATAL’98 Paris, France, July 4–7, 1998 Proceedings 5*, pages 1–10. Springer.

Mihajlo Grbovic, Vladan Radosavljevic, Nemanja Djuric, Narayan Bhamidipati, Jaikit Savla, Varun Bhagwan, and Doug Sharp. 2015. [E-commerce in your inbox: Product recommendations at scale](#). In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Sydney, NSW, Australia, August 10-13, 2015*, pages 1809–1818. ACM.

Zhenyun Hao, Jianing Hao, Zhaohui Peng, Senzhang Wang, Philip S. Yu, Xue Wang, and Jian Wang. 2022. [Dy-hien: Dynamic evolution based deep hierarchical intention network for membership prediction](#). In *WSDM ’22: The Fifteenth ACM International Conference on Web Search and Data Mining, Virtual Event / Tempe, AZ, USA, February 21 - 25, 2022*, pages 363–371. ACM.

Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2023. [Debertav3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.

Sharon Hirsch, Ido Guy, Alexander Nus, Arnon Dagan, and Oren Kurland. 2020. [Query reformulation in e-commerce search](#). In *Proceedings of the 43rd International ACM SIGIR conference on research and*

- Jianmo Ni, Jiacheng Li, and Julian J. McAuley. 2019. [Justifying recommendations using distantly-labeled reviews and fine-grained aspects](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 188–197. Association for Computational Linguistics.
- OpenAI. 2022. [Chatgpt: Optimizing language models for dialogue](#). OpenAI.
- OpenAI. 2023. [GPT-4 technical report](#). *CoRR*, abs/2303.08774.
- Archit Parnami and Minwoo Lee. 2022. [Learning from few examples: A summary of approaches to few-shot learning](#). *CoRR*, abs/2203.04291.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal V. Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Févry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M. Rush. 2022. [Multi-task prompted training enables zero-shot task generalization](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Yisheng Song, Ting Wang, Puyu Cai, Subrota K. Mondal, and Jyoti Prakash Sahoo. 2023. [A comprehensive survey of few-shot learning: Evolution, applications, challenges, and opportunities](#). *ACM Comput. Surv.*, 55(13s):271:1–271:40.
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. [Conceptnet 5.5: An open multilingual graph of general knowledge](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 4444–4451. AAAI Press.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *CoRR*, abs/2307.09288.
- Weiqi Wang, Tianqing Fang, Wenxuan Ding, Baixuan Xu, Xin Liu, Yangqiu Song, and Antoine Bosselut. 2023a. [CAR: conceptualization-augmented reasoner for zero-shot commonsense question answering](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 13520–13545. Association for Computational Linguistics.
- Weiqi Wang, Tianqing Fang, Chunyang Li, Haochen Shi, Wenxuan Ding, Baixuan Xu, Zhaowei Wang, Jiaxin Bai, Xin Liu, Jiayang Cheng, Chunkit Chan, and Yangqiu Song. 2024. [CANDLE: iterative conceptualization and instantiation distillation from large language models for commonsense reasoning](#). *CoRR*, abs/2401.07286.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V. Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023b. [Self-consistency improves chain of thought reasoning in language models](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Changlong Yu, Xin Liu, Jefferson Maia, Tianyu Cao, Laurence Yang Li, Yifan Gao, Yangqiu Song, Rahul Goutam, Haiyang Zhang, Bing Yin, et al. 2024. [Cosmo: A large-scale e-commerce common sense knowledge generation and serving system at amazon](#). In *Proceedings of the 2024 International Conference on Management of Data, SIGMOD 2024*.
- Changlong Yu, Weiqi Wang, Xin Liu, Jiaxin Bai, Yangqiu Song, Zheng Li, Yifan Gao, Tianyu Cao, and Bing Yin. 2023. [Folkscope: Intention knowledge graph construction for e-commerce commonsense discovery](#). In *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada*,

July 9-14, 2023, pages 1173–1191. Association for Computational Linguistics.

Chenwei Zhang, Wei Fan, Nan Du, and Philip S. Yu. 2016. Mining user intentions from medical queries: A neural network based heterogeneous jointly modeling approach. In *Proceedings of the 25th International Conference on World Wide Web, WWW 2016, Montreal, Canada, April 11 - 15, 2016*, pages 1373–1384. ACM.

Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona T. Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer. 2022. OPT: open pre-trained transformer language models. *CoRR*, abs/2205.01068.

Jiashu Zhao, Hongshen Chen, and Dawei Yin. 2019. A dynamic product-aware learning model for e-commerce query intent understanding. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM 2019, Beijing, China, November 3-7, 2019*, pages 1843–1852. ACM.

Wayne Xin Zhao, Yanwei Guo, Yulan He, Han Jiang, Yuxin Wu, and Xiaoming Li. 2014. We know what you want to buy: a demographic-based system for product recommendation on microblogs. In *The 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '14, New York, NY, USA - August 24 - 27, 2014*, pages 1935–1944. ACM.

Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023. Instruction-following evaluation for large language models. *CoRR*, abs/2311.07911.

Wendi Zhou, Tianyi Li, Pavlos Vougiouklis, Mark Steedman, and Jeff Z. Pan. 2024. A usage-centric take on intent understanding in e-commerce. *CoRR*, abs/2402.14901.

Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. Minigpt-4: Enhancing vision-language understanding with advanced large language models. *CoRR*, abs/2304.10592.




Figure 5: The rate of preserved intentions after filtering under different relations.

Appendices

A Prompts

In this section we show the instructions used in feature extraction, intention generation, and human-centric role-aware filtering stages. The prompts are shown in Table 4.

B Case Studies

In this section, we showcase various co-buy intentions for the same product pairs generated under different relations. The examples are provided in Table 3.

It is evident from the table that the intentions consistently capture the key aspect of the co-buy intention. i.e., for young kids, for costume, pirate. Though for certain relations the intention doesn’t follow the instruction strictly in terms of format, the quality of the intention remains reasonable and informative. The content of these intentions is still aligned with the intended purpose of the designed relation.

C Relation-wise Filter Case Study

In this section, we present the Relation-wise Filter Preserve Rate (RFP Rate) of MIND, which represents the proportion of intentions that passed filtering among all intentions for every relation. We show our result in Figure 5.

Our observations indicate that the open relation has the lowest RFP Rate at 0.17 yet other relations demonstrate RFP Rates ranging from 0.2 to 0.4.

Task	Prompt
Feature Extraction	The $[IMAGE_1, IMAGE_2]$ contains a product and name of it is $[PROD_NAME]$. Please analyze the product image, together with the product name, provide a detailed description focusing on the product’s features, design, and apparent quality. Highlight any unique characteristics or visible elements that distinguish this product from similar items. Additionally, speculate on the potential uses and benefits of this product for a consumer, based on its appearance or any information in the image and the name.
Intention Generation	The two $[Image_1, Image_2]$ are two different products. The product name of the upper image is $[Prod_A_Name]$. The product detail and the potential purchase intention is $Prod_A_Desc$. The product name of the lower image is $Prod_B_Name$. The product detail and the potential purchase intention is $Prod_B_Desc$. Based on information provided, together with the product images, what could be the potential intention for people buying these two products in one purchase simultaneously based on the relation of $[Relation_Prompt[Relation]]$, take the image features into consideration, limit your word count within 120 words. Start with the potential co-buy intention could be $Relation_Prompt[Relation]$
Human-centric Role-aware Filtering	The two images $[Image_1, Image_2]$ are two different products. The product name of the upper image is $[Prod_A_Name]$. The product detail and the potential purchase intention is $[Prod_A_Desc]$. The product name of the lower image is $[Prod_B_Name]$. The product detail and the potential purchase intention is $[Prod_B_Name]$. Under the relation of $[Relation_Prompt[Relation]]$, the potential co-buy intention would be $[Intention]$. If you are a consumer who are eager to buy product a or product b, would this intention encourage you to buy the two products simultaneously? be critical on your choice, output yes or no together with the reason for your answer. For example, the output should be Yes, ... or No, ...

Table 4: Prompts used for evaluating LLM baselines across various tasks in a zero-shot scenario.

We hypothesize that the under-performance of open relation generation could be attributed to its less specific instruction adopted in generation process. The lack of specific information in the instruction may hinder the LVLM’s ability to generate persuasive and informative intentions for the filter LVLM, resulting in the low preserve rate.

This finding emphasizes the importance of future intention mining research. It suggests that solely relying on the expressive power of LVLMs to undermine potential intentions is not feasible. Instead, a meticulous instruction constraint aligns with research purpose is required. Specifically, incorporating detailed relation information during intention mining is indispensable in E-commerce co-buy behavior understanding domain. This could improve the intention mining process, leading to a better construction of a credible and comprehensive intention knowledge base.





Item 1	Item 2	Relation	Intention
Girls Prewalker Toddler Cute Flower Bowtie Antiskid Shoes Sneaker 	Fisher-Price Brilliant Basics Rock-a-Stack 	symbolOf can capableOf cause isA	They both represent the early stages of a child's development. They both help to develop explore and develop children's skills. They both provide young children with a safe and engaging environment. The person wants to purchase both products as gifts for a young child. They both cater to the needs of young children.
Rubies 18th Century Colonial Man Wig Adult One Size 	Pirate Boot Toppers - Fun Costume Accessory 	usedFor symbolOf isA cause	They are both used for costume or theatrical performances. They both symbolize a pirate or colonial theme. They are both costume accessories for a pirate-themed outfit. The person wants to create a complete and authentic pirate costume.

Figure 6: MIND co-buy intentions generated under different relations.