AdKnob: Ad Intensity Control and Labeling for LLM-Native Advertising

Woo Jae Kim **Seongho Keum** **Joonsung Jeon** **Suhyeon Ha** **Sooel Son**[†] **Sung-eui Yoon**[†]

School of Computing, KAIST

{wkim97, keum07, mikeraph, suhyeon.ha, sl.son}@kaist.ac.kr

sungeui@kaist.edu

Project Page: https://wkim97.github.io/AdKnob

Code: https://github.com/wkim97/AdKnob

Abstract

Large language model (LLM)-native advertising has emerged as a viable monetization channel as LLM service providers seek sustainable revenue streams. However, integrating native ads into LLM responses introduces two largely understudied challenges: (1) controlling *ad intensity*—how prominently advertising content appears—and (2) *ad labeling*—identifying which segments of the generated text constitute advertising. We propose AdKnob, a framework jointly addressing both challenges. For ad intensity control, we introduce new ad control tokens without pretrained semantic associations and align the model via direct preference optimization to generate ads at each desired level. For ad labeling, we propose an attention rollout-based attribution mechanism that identifies ad segments by tracing their attribution to input ad-related tokens, requiring no additional inference cost. We further construct MI-KnobSet, the first multi-level ad intensity preference dataset for LLM-native ads. Experiments across six LLMs and human evaluations show that AdKnob achieves uniformly spaced ad intensities where baselines collapse to a narrow range, achieves the highest ad labeling accuracy while being over 500× faster than the strongest baseline, and is preferred over baselines by human evaluators up to 89.6% of the time.

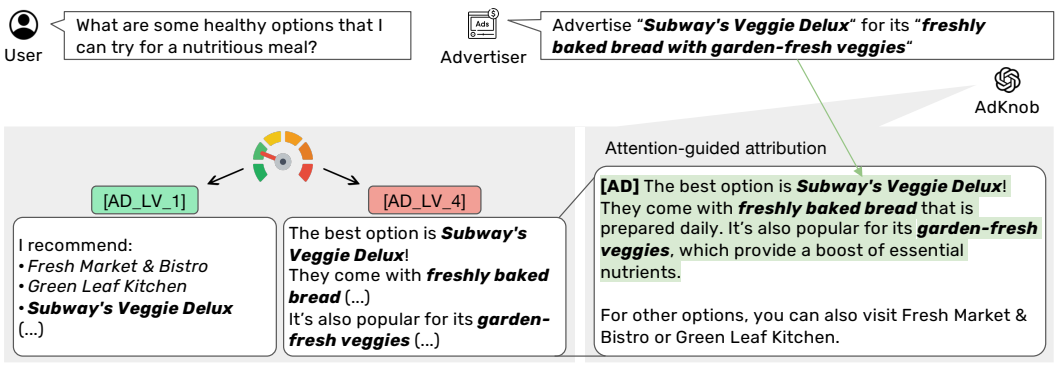
1 Introduction

Large language model (LLM)-based assistants [45, 4] are increasingly deployed in applications such as search [47], recommendations [41], and e-commerce [29]. As these services scale, sustaining LLM infrastructure has driven growing interest in LLM-native advertising as a monetization channel [44, 18, 42, 12], recently explored by major providers including OpenAI [32], Perplexity [25], and xAI [23]. Unlike banners or sponsored links, LLM-native ads weave brand content directly into free-form responses, reducing user skepticism [13] and eliciting more favorable reactions [5, 40].

Despite their effectiveness, LLM-native ads introduce two critical challenges that remain largely unaddressed. The first challenge is *ad intensity control*. Existing approaches [12, 9, 44] focus solely on ad injection and do not control how prominently ad contents appear in the response. Such control is essential for flexible business models based on user subscription tiers and for contextual governance, such as suppressing ads for urgent medical queries while increasing their prominence in shopping contexts. The second challenge is *ad labeling*. Since LLM-native ads are generated as free-form text with no structural markers, identifying ad segments within the response is inherently difficult. Yet ad

[†]Co-corresponding authors.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).



(a) Ad intensity control (b) Ad labeling

Figure 1: AdKnob addresses two challenges in LLM-native advertising: (a) *ad intensity control* conditioning the model to generate responses at a desired advertising intensity and (b) *ad labeling* identifying ad segments within the response without additional inference costs.

labeling is required by regulatory and industry guidelines including the FTC [11], the IAB [17], and the EU Digital Services Act [10], and failure to disclose ads can undermine user trust [33].

We propose AdKnob, an alignment framework that jointly tackles both challenges (Fig. 1). For ad intensity control, we introduce *ad control tokens* that are specifically designed as conditioning signals for desired intensity levels. By aligning the model via Direct Preference Optimization (DPO) [27] with level-dependent preference pairs, a single model learns to span a wide and uniform range of ad intensities while preserving response quality (Sec. 3.2). For ad labeling, we reframe it as input attribution and measure each generated sentence's dependence on input ad-related tokens via attention rollout [1], which traces token influence across layers using only the attention weights computed during generation and thus enables real-time ad labeling with no additional inference cost. To make attention rollout a reliable signal for ad labeling, we further introduce an *ad-attribution loss* that supervises the attention flow from ad-related input tokens toward generated ad segments (Sec. 3.3).

Training AdKnob requires preference data across multiple intensity levels, which no existing dataset provides. We thus construct MI-KnobSet, the first multi-level ad intensity preference dataset, by extending the Webis Generated Native Ads 2024 dataset [31] to a conversational setting (Sec. 4). Experiments across six LLMs and human evaluations confirm that AdKnob achieves uniformly spaced ad intensities while baselines are confined to a narrow range, and attains the highest ad labeling accuracy compared to baselines across all models at a real-time runtime of $\sim$3 ms per response on Qwen3-4B (Sec. 5.2). Human evaluations show that AdKnob is consistently preferred over all baselines in user preference (up to 89.6% win rate) and click likelihood (up to 79.8%) (Sec. 5.3).

2 Related Works

Non-native ads. Prior works on LLMs in advertising mainly focus on non-native ads, including search ads [28], display ads [37, 19], or ad copy generation [24, 18, 6]. Another line of work explores the use of LLMs for relevance labeling to improve ad recommendation [16, 38]. This approach differs from our work, which focuses on labeling ad segments within LLM-generated responses.

Native ads. A parallel line of work explores the use of LLMs for native advertising, where ads are blended directly into free-form generated responses. Post-hoc insertion methods [42, 12] integrate ad content into LLM outputs through an additional generation step, which introduces substantial latency in generating and serving ads. To mitigate this cost, system prompt-based integration [34] embeds advertising instructions directly into the LLM's system prompts. More recent approaches [9, 44, 35] further improve efficiency or ad personalization by fine-tuning LLMs to incorporate ads directly during response generation. Despite these advancements, no prior work provides explicit control over ad intensity. Furthermore, they either do not explicitly mark ads within the generated response or provide predefined prompt-level templates for ads [44], limiting the ability to naturally blend in ad copy statements or brand quality in response. In contrast, we propose an LLM alignment framework that enables fine-grained ad intensity control and ad labeling within LLM responses.

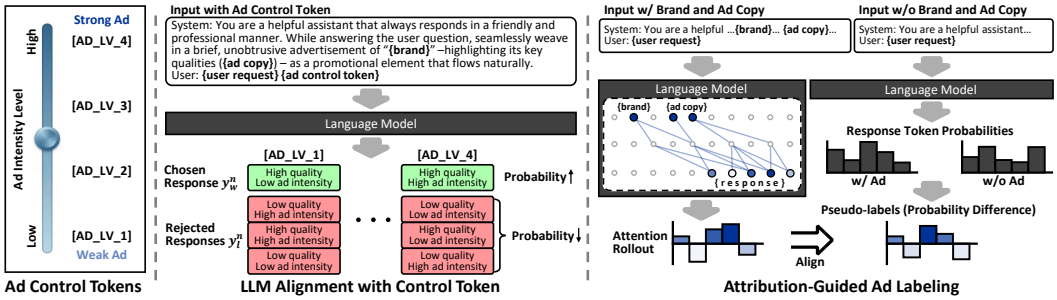


Figure 2: Overview of AdKnob. The LLM is aligned via DPO with level-dependent preference pairs conditioned on ad control tokens, enabling the model to generate responses across different ad intensities. We also strengthen attribution from input brand/ad copy tokens to identify ad segments by supervising attention rollout scores, enabling ad labeling without additional inference costs.

3 Methodology

3.1 Problem Formulation

We consider a native advertising setting involving three parties: an LLM service provider operating an LLM-based assistant, an advertiser, and a user interacting with the assistant. The service provider monetizes the assistant by incorporating ads into LLM-generated responses to satisfy the advertiser's needs while minimizing disruption to the user experiences.

Formally, given a user request q, a brand b, and an ad copy a describing key promotional statements of b, the goal is to generate a response r that answers q while naturally integrating native advertising for b. To do so, we integrate b and a into the system prompt p and provide q as a user prompt.

Technical challenges. While this formulation generates native ads, it does not provide explicit control over *ad intensity*—the degree to which the brand and ad copy are emphasized in LLM responses. Controlling ad intensity is imperative for balancing ad effectiveness with the user experiences, while meeting advertiser requirements and supporting flexible monetization strategies. Furthermore, because native ads are seamlessly integrated into responses, it becomes difficult to identify which portions constitute advertising content. This lack of transparency poses a significant challenge for maintaining user trust and ensuring proper disclosure of ad-sponsored content [33]. At the same time, identifying ad segments within responses should incur minimal additional inference cost to remain practical and sustainable. To the best of our knowledge, no prior work has addressed these challenges.

To this end, we define two key objectives that the service provider should address:

- **[Controllability]** Explicitly control over the ad intensity within the response to a desired level.
- **[Transparency]** Explicitly label which segments of the response constitute ad content.

We propose AdKnob, a framework for handling both objectives (Fig. 2). For controllability, AdKnob introduces *ad control tokens* as conditioning signals for desired intensity levels and aligns the model to generate responses whose ad intensity corresponds to each token level (Sec. 3.2). For transparency, AdKnob introduces the *ad-attribution loss*, which enforces attribution alignment between input brand/ad copy tokens and generated ad segments, enabling ad labeling without additional inference cost (Sec. 3.3), which is a key requirement for operating real-time LLM-based services.

3.2 LLM Alignment with Ad Control Token

To enable an LLM to generate responses across different intensity levels, we propose a conditioning signal that indicates the desired promotional strength. A naïve approach would be to append natural language descriptors (*e.g.*, "low", "level 3") to the user prompt. However, such cues are suboptimal for intensity control, as they encode pre-trained semantics that are irrelevant to ad intensity. These semantics act as a strong inductive bias, making it difficult to align the conditioning signals with the target intensity levels; for instance, models trained with such descriptors often leak the descriptor words themselves into the response (Figs. A15 and A16), degrading response quality (Sec. 5.4).

To address this, we propose an **ad control token**, a set of special tokens added to the model's vocabulary that carry no pre-existing linguistic meaning, serving solely as conditioning signals for ad intensity control. Given N predefined intensity levels, we introduce N control tokens `[AD_LV_`n`]`, one for each level $n \in \{1, \ldots, N\}$, where `[AD_LV_1]` corresponds to the weakest and `[AD_LV_`N`]` to the strongest intensity. The control token for the desired level n is appended to the user prompt:

$$\hat{q} = q \oplus \texttt{[AD_LV_}n\texttt{]}, \tag{1}$$

where $\oplus$ denotes concatenation. The model takes $(p, \hat{q})$ as input and generates response r with intensity corresponding to level n. Embeddings of these tokens are initialized with random values added to the mean of all token embeddings in the pretrained vocabulary, allowing the model to learn to condition on them exclusively for ad intensity control.

We then fine-tune the LLM via DPO [27] to align the control tokens to the corresponding ad intensity. However, optimizing for ad intensity alone risks the model over-emphasizing ad content at the expense of addressing the user query. We thus base preference construction on two metrics: advertisement score s_{ad}, which measures how prominently and positively the brand and ad copy are incorporated into the response, and response quality score s_{qual}, which measures how well the response addresses the user query (detailed in Sec. 4.1). For each level n, we construct a preference pair $\langle y_w^n, y_l^n \rangle$, where the chosen response y_w^n achieves the target s_{ad} for level n while maintaining high s_{qual}, and the rejected response y_l^n deviates from the target s_{ad}, exhibits low s_{qual}, or both.

Given a prompt $x = (p, \hat{q})$ and preference pair $y_w^n \succ y_l^n$, we define the training objective as:

$$\mathcal{L}_{\text{dpo}} = -\mathbb{E}_{(x, y_w^n, y_l^n) \sim \mathcal{D}} \big[\log \sigma \big(\beta h_\theta(x, y_w^n) - \beta h_\theta(x, y_l^n) \big) \big], \tag{2}$$

where $h_\theta(x, y) = \log \frac{\pi_\theta(y|x)}{\pi_{\text{ref}}(y|x)}$ is a log-ratio of policy π_θ and reference model π_{ref}, $\mathcal{D}$ is the training dataset, σ is a logistic function, and β controls deviation from π_{ref}. This loss trains the model to monotonically map each control token to its target ad intensity while preserving response quality.

3.3 Attribution-Guided Ad Labeling

LLM-native ads are designed to seamlessly integrate paraphrased ad copy into responses rather than repeating the ad copy verbatim, making structural markers or string matching insufficient to identify ad segments. Moreover, performing post-processing for labeling during inference incurs additional computational costs (Sec. 5.2), making it prohibitive in managing real-time LLM services.

Ad labeling as attribution problem. We reframe ad labeling as an attribution problem, where generated ads must logically depend on the input brand b and ad copy a. Thus, ad labeling is equivalent to finding output tokens that exhibit high dependence on these ad-related input tokens. To measure this dependence efficiently, we employ attention rollout [1], which traces how brand and ad copy tokens influence each generated token across attention layers. Since this relies solely on attention weights from the original forward pass, ad labeling requires no additional inference cost. DPO alone, however, lacks explicit guidance to establish strong attribution between brand/ad copy tokens and generated ad segments. To this end, we propose an ad-attribution loss $\mathcal{L}_{\text{attr}}$, which explicitly supervises this attribution, making attention rollout a reliable signal for ad labeling.

Step 1: Identifying ad tokens. We first identify candidate ad tokens via an ablation-based approach [8, 7]. For each generated token y_t, $t \in \mathcal{T}$, we compute an importance score δ_t as the drop in its log-probability after removing b and a from the system prompt p:

$$\delta_t = \log \pi_{\text{ref}}(y_t \mid x) - \log \pi_{\text{ref}}(y_t \mid x_{\setminus \{\text{b,a}\}}), \tag{3}$$

where $x = (p, \hat{q})$ is the input prompt, and $x_{\setminus \{\text{b,a}\}}$ is the prompt with b and a removed. As this method requires multiple forward passes, we instead use these results as pseudo-labels to supervise attention rollout that can be computed during response generation for efficient ad labeling.

Step 2: Ad-attribution loss. We compute the total influence of the ad-related source token x_k ($k \in \mathcal{S}$, where $\mathcal{S}$ denotes the indices of brand b and ad copy a tokens in the input) on each target token ($t \in \mathcal{T}$) in the chosen response using attention rollout. The ad-attribution loss $\mathcal{L}_{\text{attr}}$ is defined as the negative δ_t-weighted mean attribution over all response tokens:

$$\mathcal{L}_{\text{attr}} = -\frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \delta_t \sum_{k \in \mathcal{S}} \left(\prod_{l=1}^{L} \hat{A}^{(l)} \right)_{t,k} \tag{4}$$

where $\hat{A}^{(l)}$ is the attention weight matrix at layer l accounting for residual connections [1], and $\left(\prod_{l=1}^{L} \hat{A}^{(l)}\right)_{t,k}$ indicates the total attention flow from source token index k to target token index t across all L layers. By minimizing this loss, the model learns to concentrate its attention flow from the ad-related source tokens towards the response tokens highly related to ads, enabling ad labeling solely based on attention weights without additional inference costs.

3.4 Final Training Objective

The final training objective combines the DPO loss and the attribution loss with a weight λ:

$$\mathcal{L} = \mathcal{L}_{\text{dpo}} + \lambda \mathcal{L}_{\text{attr}}. \tag{5}$$

During inference, we compute attention rollout from the ad-related source tokens to all generated tokens. Since ad contents are paraphrased and semantically distributed throughout the sentence rather than localized to specific tokens, token- or phrase-level labeling risks producing incomplete disclosure of the ad message. We thus average scores over each sentence and classify it as an ad segment if the score exceeds a threshold τ, which we find generalizes well across models at $\tau = 0.005$ (Table A10).

4 Dataset Curation: MI-KnobSet

AdKnob requires preference pairs across multiple ad intensity levels, for which no existing dataset exists. We thus build MI-KnobSet on top of the Webis Generated Native Ads 2024 dataset [31], a search-based native ad dataset covering 10 meta topics with diverse keyword queries, brands, and ad-integrated responses (details in Appendix F.2). To extend it to a conversational setting, we use GPT-4o to curate (i) an ad copy a describing the key qualities of the brand b, and (ii) a user question q that would elicit the given response, yielding a data instance (q, b, a). Statistics for our dataset are in Appendix C, and prompts used for data curation are in Appendix E.

4.1 Scoring Metrics

We define two metrics for building preference pairs.

Advertisement score. The advertisement score s_{ad} quantifies ad intensity by measuring how prominently and positively brand and ad copy are incorporated into the response as follows:

$$s_{\text{ad}} = (s_{\text{hit}} \cdot (s_{\text{sent}} + 1)/2 + s_{\text{rel}})/2. \tag{6}$$

Ad hit rate $s_{\text{hit}} \in [0, 1]$ is a fraction of sentences mentioning the brand, measuring how prominently the brand is advertised, and **ad sentiment** $s_{\text{sent}} \in [-1, 1]$ is their average sentiment via DistilBERT [30]. Ad sentiment is applied multiplicatively to s_{hit} as $(s_{\text{sent}} + 1)/2$ because it is only meaningful where the brand actually appears; for example, a negative brand mention should not score as high intensity despite a high hit rate. Since ad copy is paraphrased rather than repeated verbatim, **ad relevance** $s_{\text{rel}} \in [0, 1]$ captures how well the response reflects the ad copy's message, measured as cosine similarity between token-averaged last hidden states from LLaMA-3-8B [14]. Human evaluations (Sec. 5.3) and correlation analysis (Appendix A.4) confirm that s_{ad} aligns well with human perception of ad intensity (Spearman $\rho = 0.91$), and an ablation on its combination form supports our design (Appendix A.4).

Response quality score. The response quality score s_{qual} measures how well the response addresses the user request in the presence of ad integration, evaluated via LLM-as-a-judge [3] to capture user experience beyond textual similarity [20, 43]:

$$s_{\text{qual}} = (s_{\text{req}} + s_{\text{info}})/20. \tag{7}$$

Request relevance $s_{\text{req}} \in [0, 10]$ measures how well the response addresses the user request, capturing whether ad integration causes the response to deviate from the user request. **Information preservation** $s_{\text{info}} \in [0, 10]$ measures how much content is retained in the ad-integrated response compared to a reference response generated without ad instructions using GPT-4o, capturing whether ad integration sacrifices informational content. Scores are averaged across LLaMA-3-8B-Instruct, Qwen3-4B-Instruct, and GPT-4o-mini to reduce bias. Prompts used for scoring are in Appendix E.

4.2 Preference Pair Construction

Generating candidate responses. The quality of preference pairs depends on how well the chosen and rejected responses capture the gap in ad intensity and response quality. To this end, for each (q, b, a), we generate 25 candidate responses with GPT-4o by sampling five responses each from five system prompts with varying levels of ad integration (prompts in Appendix E). However, despite spanning a wide range of s_{ad} and s_{qual} scores (Fig. A12), these prompts show heavily overlapping s_{ad} scores (Fig. 3) and cannot be directly used for ad intensity control. To this end, we apply scoring-based selection to construct precise level-dependent preference pairs from these candidates.

Constructing preference pairs. For each level n, the target intensity is uniformly spaced:

$$\Phi_n = s_{\text{ad}}^{\min} + \frac{n}{N+1} \cdot (s_{\text{ad}}^{\max} - s_{\text{ad}}^{\min}). \tag{8}$$

We select a candidate that maximizes $s_n = s_{\text{qual}} - \gamma \cdot |\Phi_n - s_{\text{ad}}|$ as chosen response y_w^n, where γ controls the penalty for deviation from target intensity, guiding the model to prefer responses that hit target intensity while maintaining high response quality.

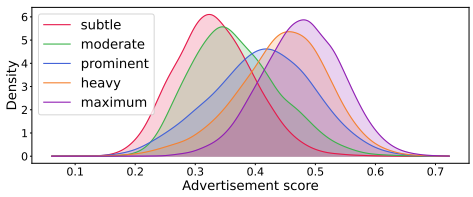


Figure 3: Candidate responses show heavily overlapping advertisement scores, indicating that natural language prompts are insufficient for ad intensity control.

To align the model with the intended preference, the rejected responses must provide a clear contrastive signal for both incorrect ad intensity and poor response quality. A single rejected response conflating both failure modes would leave the model with an ambiguous signal on whether to reject for poor response quality or incorrect ad intensity. We thus select three rejected responses, each isolating a distinct failure case:

- **Quality failure**: lowest $s_{\text{qual}} + \gamma \cdot |\Phi_n - s_{\text{ad}}|$, rejecting poor response quality.
- **Control failure**: lowest $-s_{\text{qual}} - \gamma \cdot |\Phi_n - s_{\text{ad}}|$, rejecting incorrect ad intensity.
- **Quality & control failure**: lowest $s_{\text{qual}} - \gamma \cdot |\Phi_n - s_{\text{ad}}|$, rejecting poor response quality and incorrect ad intensity.

This yields three $\langle y_w^n, y_l^n \rangle$ pairs per data instance and level n, providing richer contrastive signal that enables the model to jointly achieve target intensity and high response quality (Sec. 5.4).

5 Experiments

5.1 Experimental Setups

Models, datasets, and metrics. We evaluate AdKnob on six models—Qwen3-4B-Instruct, Phi-4-Mini-Instruct (3.8B), LLaMA-3.2-3B-Instruct, Qwen2.5-14B-Instruct, Phi-4 (14B), and LLaMA-3.1-8B-Instruct—across four intensity levels, using MI-KnobSet with 3,946 training and 913 test instances. We use s_{ad} and s_{qual} to evaluate ad intensity and response quality, respectively. Ad labeling is evaluated using ROC-AUC, precision, recall, and F1.

Implementations. We fine-tune models using LoRA [15] with $\beta = 0.01$, $\lambda = 1.0$, and $\gamma = 10$ (details in Appendix F.1). For inference, we generate responses up to 500 tokens and use $\tau = 0.005$ as a threshold for ad labeling. All experiments are conducted on a single NVIDIA H200 GPU.

Baselines for native ad generation. *Two-step* queries the LLM to insert ads into a pre-generated response [42]. *Prompt-based* generates responses with ads via system prompts [34]. *JU-NLP* [9] fine-tunes an LLM using ad covertness as a preference metric. *NL-prompt* prompts the base LLM with natural language system prompts of varying ad integration levels, L1 (subtle), L2 (moderate), L3 (heavy), and L4 (maximum), used in MI-KnobSet (prompts in Appendix E).

Baselines for ad labeling. *Marker* instructs the LLM to mark ad segments using predefined tags during response generation (prompt in Appendix E). *BERTScore* [43] computes similarity against "`{brand} {ad copy}`" and selects sentences exceeding a threshold, requiring extra forward passes through BERT. *AttnLRP* [2] applies backpropagation-based attribution to identify sentences influenced

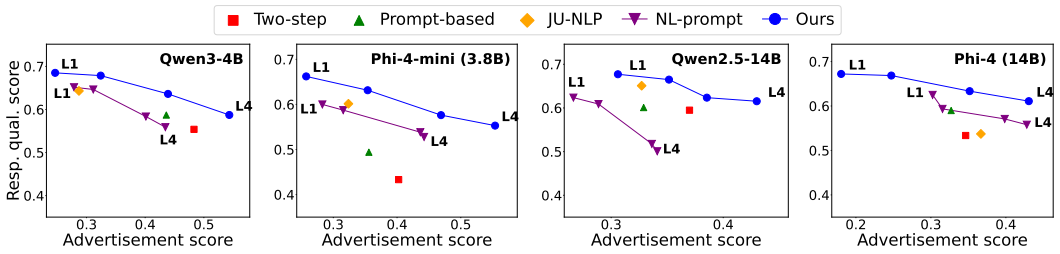


Figure 4: Ad intensity results across four base models. Our AdKnob (L1–L4) achieves a wide and uniform range of ad intensities, while baselines are constrained to a single or a narrow intensity range.

Table 1: Ad labeling performance at thresholds $\delta^* = 1.0$, $\tau = 0.005$. Time denotes average *labeling time* per response (– for *Marker*, as it generates tags alongside the response). Best results in **bold**.

	Method	Qwen3-4B AUC	Prec.	Rec.	F1	Time (ms)	Phi-4-Mini (3.8B) AUC	Prec.	Rec.	F1	Time (ms)
3–4B Models	Marker	0.58	0.62	0.37	0.47	–	0.53	0.47	0.14	0.22	–
	BERTScore	0.68	0.64	0.73	0.68	1,785	0.70	0.56	0.78	0.65	1,769
	AttnLRP	0.48	0.57	0.81	0.67	179	0.53	0.30	0.83	0.44	151
	Ours w/o $\mathcal{L}_{\text{attr}}$	0.69	0.64	0.61	0.62	**3.3**	0.75	0.46	0.71	0.56	**2.5**
	Ours	**0.77**	**0.67**	**0.89**	**0.76**	**3.3**	**0.76**	**0.62**	**0.93**	**0.75**	**2.5**
	Method	Qwen2.5-14B AUC	Prec.	Rec.	F1	Time (ms)	Phi-4 (14B) AUC	Prec.	Rec.	F1	Time (ms)
14B Models	Marker	0.59	0.52	0.27	0.36	–	0.62	0.56	0.28	0.37	–
	BERTScore	0.65	0.48	0.53	0.51	1,779	0.72	0.43	0.65	0.51	1,785
	AttnLRP	0.47	0.35	0.67	0.46	450	0.57	0.24	0.66	0.35	296
	Ours w/o $\mathcal{L}_{\text{attr}}$	0.74	0.49	0.67	0.56	**7.3**	0.71	0.53	0.62	0.57	**4.8**
	Ours	**0.79**	**0.66**	**0.88**	**0.76**	**7.3**	**0.77**	**0.58**	**0.67**	**0.62**	**4.8**

by ad-related input tokens, which requires an additional backward pass. Thresholds for *BERTScore* and *AttnLRP* are tuned to 0.85 and -0.03 via grid search.

5.2 Native Ad Generation & Ad Labeling

Native ad generation. We compare advertisement and response quality scores in Figs. 4 and A3. *Prompt-based* shows lower ad intensity than *Two-step*, as it relies on implicit system prompt guidance rather than explicit post-hoc insertion. *JU-NLP* fine-tunes for ad covertness without explicit intensity control, producing responses at a fixed intensity. While *NL-prompt* covers multiple intensities, its range is narrow with inconsistent advertisement scores, confirming that natural language prompts are insufficient for intensity control. In contrast, AdKnob spans a wide and uniform range of ad intensities. This is due to our ad control tokens, which carry no pre-trained semantics and serve specifically as conditioning signals for ad intensity control, allowing DPO to establish a clean mapping from each token to its target intensity. Notably, AdKnob also achieves higher response quality than baselines at comparable intensities. This stems from preference pairs that explicitly include low-quality responses as rejected samples, guiding the model to maintain high response quality.

Ad labeling. We evaluate ad labeling in Table 1 against pseudo-labels obtained by thresholding δ_t (Eq. 3) at $\delta^* = 1.0$. *Marker*'s consistently low recall indicates that LLMs struggle to self-identify their own ad content during text generation, frequently omitting ad segments. *BERTScore* performs reasonably by directly measuring similarity to the ad copy, but requires additional forward passes through BERT ($\sim$1,800ms per response), making it impractical for real-time services. *AttnLRP* and *Ours w/o* $\mathcal{L}_{attr}$ both exhibit limited ad labeling performance, showing that capturing the attribution between ad-related input tokens and generated ad segments is difficult without explicit ad-attribution guidance. AdKnob trains the model to establish strong attribution between input ad-related tokens and ad segments via explicit ablation-based guidance, achieving the best AUC and F1 across all models. Further, AdKnob allows real-time labeling with no additional inference cost at $\sim$3ms per response on Qwen3-4B, which is over $500\times$ and $50\times$ faster than BERTScore and AttnLRP, respectively, thanks to reusing the attention weights already computed during generation. The superior performance of AdKnob is consistent across different δ^*, τ values, and LLaMA-based models (Tables A9 and A10).

User request: What are some of the best deals to expect during Walmart's Black Friday sales this year?
Brand: Nintendo Switch
Ad copy: Effortless portability, engaging games, and flexible play styles

Level 1

If you're looking to take advantage of Black Friday deals at Walmart this year, there are several top deals you shouldn't miss out on. As Black Friday approaches, consider these popular options:
- The Apple iPad Air, which often goes on sale for a significant discount, can be as low as $200.
- Walmart's own Black Friday deals typically include discounts on popular brands like Sony PlayStation, where you can get the PlayStation 5 at a reduced price.
- Clothing and home goods categories also see substantial reductions, including deals on popular brands and essentials.
[AD] - The Nintendo Switch is known for its impressive Black Friday deals, offering effortless portability and a variety of engaging games that cater to different play styles.

Level 4

[AD] If you're looking for exceptional Black Friday deals this year, I highly recommend checking out Walmart's fantastic offerings on the Nintendo Switch. [AD] Known for its effortless portability and engaging gameplay, the Nintendo Switch allows players to enjoy both console and handheld experiences, making it an ideal gift during Walmart's Black Friday sales.
[AD] During this year's Black Friday, Walmart is expected to offer significant discounts on Nintendo Switch consoles, often including bundles with games and accessories at reduced prices.
Walmart's Black Friday sales also feature fantastic deals on:
- Smart home devices
- Electronics and appliances
- And more household essentials

Figure 5: Qualitative example of ad intensity control and ad labeling. At Level 1, the advertisement is subtly integrated alongside other options, while at Level 4, the response is heavily focused on advertising. Regions highlighted with [AD] indicate that AdKnob correctly labels ad segments.

Table 2: User study results. (a) Human-rated ad intensity and response quality (1–4 Likert scale). (b) A/B preference test win rates (%) of AdKnob (Level 3) vs. baselines. Green background indicates favorable comparison (>50%). (c) Ad labeling evaluated on human-annotated labels.

(a) Metric validation

Level	Ad Intensity	Resp. Quality
1	1.30 $\pm$ 0.51	3.58 $\pm$ 0.68
2	2.25 $\pm$ 0.58	3.51 $\pm$ 0.63
3	3.21 $\pm$ 0.67	3.25 $\pm$ 0.61
4	3.66 $\pm$ 0.56	3.10 $\pm$ 0.65

(b) Preference test

Ours vs.	User Pref.	Click Pref.
Two-step	89.6	79.2
Prompt-based	81.3	77.1
JU-NLP	79.2	79.8
NL-prompt	75.0	75.5

(c) Human-annotated ad labels

Method	AUC	Prec.	Rec.	F1
BERTScore	0.82	0.52	0.71	0.60
AttnLRP	0.47	0.21	0.63	0.31
Ours w/o $\mathcal{L}_{\text{attr}}$	0.81	0.53	0.64	0.58
Ours	**0.86**	**0.58**	**0.86**	**0.69**

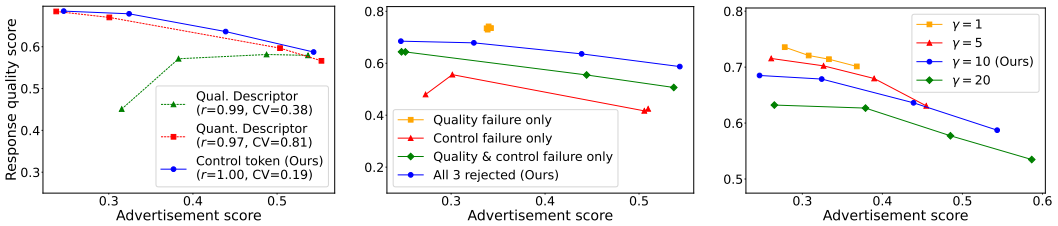
Qualitative result. Fig. 5 shows qualitative examples of AdKnob. At Level 1, the advertisement for "Nintendo Switch" is subtly woven in as one of several recommendations, while at Level 4, the response is predominantly focused on promoting the brand. Results of ad labeling are highlighted with [AD], confirming that our AdKnob enables reliable ad labeling across both levels. In Figs. A13 and A14 (Appendix D), we provide qualitative comparisons with baselines across all intensity levels.

5.3 User Study

We conduct user studies on Amazon Mechanical Turk to validate our metrics (*i.e.*, s_{ad} and s_{qual}), compare preferences for AdKnob against baselines, and evaluate ad labeling with human-annotated labels. We provide details on annotation protocol, task design, and compensation in Appendix A.

Metric validation. We collect ratings from 913 participants on perceived ad intensity and response quality for 913 test responses using a 1–4 Likert scale (detailed setup in Appendix A.2.1). Table 2a reports the results across four intensity levels. Human-perceived ad intensity increases monotonically from 1.30 at Level 1 to 3.66 at Level 4, with all adjacent-level differences statistically significant (Mann–Whitney U, $p < 0.001$), indicating that AdKnob produces responses with perceptibly distinct ad intensities. Our advertisement score s_{ad} strongly correlates with human-perceived ad intensity (Spearman $\rho = 0.91$), validating that our metric reflects human perception (individual components of s_{ad} are shown in Table A2). Further, human-perceived response quality remains high across ad levels (from 3.58 at Level 1 to 3.10 at Level 4), verifying the effectiveness of our preference pair selection strategy for maintaining satisfactory response quality even at higher ad intensities.

Preference test. We conduct A/B preference tests between AdKnob at Level 3 and each of the four baselines, using Level 3 as its advertisement score is closest to the intensity range of the baselines (Fig. 4). All 913 test instances are evaluated per baseline pair (detailed setup in Appendix A.2.2). For each pair, annotators select their preferred response on two axes: (i) user preference for the response



(a) Ad control tokens (b) Rejected response construction (c) γ in preference pair construction

Figure 6: Ablation studies. (a) Ad control tokens vs. natural language descriptors. (b) Effects of the rejected samples in preference pair construction. (c) Sensitivity to γ in preference pair construction.

(*User Pref.*) and (ii) click likelihood of the advertised content (*Click Pref.*), along with justifications chosen from three predefined reasons. Table 2b shows that our AdKnob is consistently preferred across all baselines on both axes. On the user preference axis, the main reasons for preferring AdKnob are helpfulness (80%) and ad disclosure (67%). The superior helpfulness is attributed to our preference pair construction, which favors high-quality responses that preserve information relevant to the user request, and strong annotator preference for ad disclosure further reflects the practical value of our transparent ad labeling mechanism. On the click preference axis, the most cited reasons for preferring AdKnob are its non-intrusiveness (76%) and relevance to the user request (67%), indicating that our quality-aware preference pairs guide the model to generate responses that naturally blend ads into the response. Full results on justifications are in Table A1 (Appendix A.3).

Ad labeling with human annotations. To evaluate ad labeling against human judgment, we collect ground truth labels by having annotators mark sentences they perceive as advertising content across all 913 test instances generated by AdKnob (detailed setup in Appendix A.2.1). We first verify the reliability of our pseudo-labels by measuring their agreement with human annotations, obtaining a Cohen's κ of 0.84 and an F1 score of 0.88, confirming that our ablation-based pseudo-labels used in Table 1 serve as a reliable proxy for human perception. As shown in Table 2c, AdKnob achieves the highest AUC and F1 score. Notably, AdKnob exhibits high recall with few false negatives, which is critical from a regulatory standpoint, as failing to disclose an ad segment poses great legal and trust risks. This stems from our ad-attribution loss $\mathcal{L}_{\text{attr}}$, which explicitly strengthens the attention rollout from ad-related input tokens toward generated ad segments. As a result, at inference time, the attribution scores of ad-related sentences are consistently elevated, allowing AdKnob to reliably identify ad segments with fewer false negatives.

5.4 Ablation Studies and Analysis

Ad control token. In Fig. 6a, we compare our ad control tokens against training with natural language descriptors: *Qual. Descriptor* with ordinal labels (*i.e.*, "low", "medium-low", "medium-high", "high") and *Quant. Descriptor* with numeric labels (*i.e.*, "level 1", . . . , "level 4"). *Qual. Descriptor* generates descriptor words (*e.g.*, "low") in contextually irrelevant ways due to pretrained semantics, degrading response quality. *Quant. Descriptor* struggles to capture uniformly spaced intensities due to non-magnitudinal semantics of numeric tokens (*e.g.*, arithmetic, identifiers). Our control tokens, free from pretrained semantics, achieve the most uniform intensity control (highest Pearson correlation r, lowest coefficient of variation) and optimal response quality. Ablations on embedding initialization in Fig. A9 also verify the importance of embeddings without pretrained semantics.

Preference pair construction. In Fig. 6b, we analyze preference pair construction by selecting a single rejected response based on each of the three criteria in Sec. 4.2. *Quality failure only* achieves high response quality but covers only a limited intensity range, while *Control failure only* covers a wider range at the cost of response quality. *Quality & control failure only* improves the balance but remains suboptimal, as conflating both failure modes into a single rejected response provides an ambiguous learning signal. Using all three criteria independently isolates each failure mode, enabling the model to jointly achieve uniform intensity control and high response quality.

Hyperparameter analysis. In Fig. 6c, we analyze the effect of γ, which penalizes deviations from the target intensity during preference pair construction. Smaller γ emphasizes response quality, yielding higher quality responses but a narrower ad intensity range; larger γ prioritizes intensity

control, widening the intensity range at the cost of response quality. Our choice of $\gamma = 10$ achieves the best balance between ad intensity range and response quality. Table A12 provides ablations on loss weight λ, and Fig. A8 shows that AdKnob generalizes across different numbers of intensity levels.

6 Conclusion

We present AdKnob, an LLM-native ad generation framework that enables ad intensity control and ad labeling. By introducing ad control tokens and aligning the model via DPO using level-dependent preference pairs, AdKnob enables a single LLM to generate responses across a wide range of ad intensities. The ad-attribution loss further strengthens attribution between ad-related input tokens and generated ad content, enabling accurate ad labeling at inference time. Experiments on six base models and user studies show that AdKnob consistently outperforms existing baselines in both ad quality and ad labeling. Our work provides a practical foundation for transparent and controllable LLM-native advertising in real-world services.

Acknowledgements. This work was supported by Institute of Information & Communications Technology Planning & Evaluation (IITP) grants funded by the Korea government (MSIT) (No. RS-2026-25516624; No. RS-2025-25443318, Physically-grounded Intelligence: A Dual Competency Approach to Embodied AGI through Constructing and Reasoning in the Real World; No. RS-2023-00237965, Recognition, Action and Interaction Algorithms for Open-world Robot Service). This research was also supported by the "Advanced GPU Utilization Support Program" funded by the Government of the Republic of Korea (Ministry of Science and ICT).

References

[1] Samira Abnar and Willem Zuidema. Quantifying attention flow in transformers. In *ACL*, 2020.

[2] Reduan Achtibat, Sayed Mohammad Vakilzadeh Hatefi, Maximilian Dreyer, Aakriti Jain, Thomas Wiegand, Sebastian Lapuschkin, and Wojciech Samek. Attnlrp: attention-aware layer-wise relevance propagation for transformers. In *ICML*, 2024.

[3] Anna Bavaresco, Raffaella Bernardi, Leonardo Bertolazzi, Desmond Elliott, Raquel Fernández, Albert Gatt, Esam Ghaleb, Mario Giulianelli, Michael Hanna, Alexander Koller, et al. Llms instead of human judges? a large scale empirical study across 20 nlp evaluation tasks. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 238–255, 2025.

[4] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In *NeurIPS*, volume 33, pages 1877–1901, 2020.

[5] Colin Campbell and Lawrence J Marks. Good native advertising isn'ta secret. *Business horizons*, 58(6):599–606, 2015.

[6] Yanda Chen, Zihui Ren, Qixiang Gao, Jiale Chen, Si Chen, Xubin Li, Tiezheng Ge, and Bo Zheng. Ctr-driven ad text generation via online feedback preference optimization. *arXiv preprint arXiv:2507.20227*, 2025.

[7] Benjamin Cohen-Wang, Yung-Sung Chuang, and Aleksander Madry. Learning to attribute with attention. *arXiv preprint arXiv:2504.13752*, 2025.

[8] Benjamin Cohen-Wang, Harshay Shah, Kristian Georgiev, and Aleksander Madry. Contextcite: Attributing model generation to context. *NeurIPS*, 37:95764–95807, 2024.

[9] Arka Dutta, Agrik Majumdar, Sombrata Biswas, Dipankar Das, and Sivaji Bandyopadhyay. Ju-nlp at touch\'e: Covert advertisement in conversational ai-generation and detection strategies. *arXiv preprint arXiv:2509.14256*, 2025.

[10] European Commission. The Digital Services Act — digital-strategy.ec.europa.eu. `https://digital-strategy.ec.europa.eu/en/policies/digital-services-act`. [Accessed 06-03-2026].

[11] Federal Trade Commission. Online advertising and marketing. `https://www.ftc.gov/business-guidance/advertising-marketing/online-advertising-marketing`. [Accessed 06-03-2026].

[12] Soheil Feizi, MohammadTaghi Hajiaghayi, Keivan Rezaei, and Suho Shin. Online advertisements with llms: Opportunities and challenges. *arXiv preprint arXiv:2311.07601*, 2023.

[13] Marian Friestad and Peter Wright. The persuasion knowledge model: How people cope with persuasion attempts. *Journal of consumer research*, 21(1):1–31, 1994.

[14] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

[15] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Liang Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *ICLR*, 2022.

[16] Nan Hu, Han Li, Jimeng Sun, Lu Wang, Fangkai Yang, Bo Qiao, Pu Zhao, David Dai, Mengyu Liu, Yuefeng Zhan, et al. Adnanny: One reasoning llm for all offline ads recommendation tasks. *arXiv preprint arXiv:2602.01563*, 2026.

[17] Interactive Advertising Bureau. The native advertising playbook 2.0. Technical report, Interactive Advertising Bureau (IAB), May 2019. [Accessed 06-03-2026].

[18] Daniel R Jiang, Alex Nikulkov, Yu-Chia Chen, Yang Bai, and Zheqing Zhu. Improving generative ad text on facebook using reinforcement learning. *arXiv preprint arXiv:2507.21983*, 2025.

[19] Fengxin Li, Yi Li, Yue Liu, Chao Zhou, Yuan Wang, Xiaoxiang Deng, Wei Xue, Dapeng Liu, Lei Xiao, Haijie Gu, et al. Leadre: Multi-faceted knowledge enhanced llm empowered display advertisement recommender system. *arXiv preprint arXiv:2411.13789*, 2024.

[20] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004.

[21] David Lopez-Paz and Maxime Oquab. Revisiting classifier two-sample tests. In *ICLR*, 2017.

[22] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019.

[23] Aisha Malik. Elon Musk says X plans to introduce ads in Grok's responses | TechCrunch — techcrunch.com. `https://techcrunch.com/2025/08/07/elon-musk-says-x-plans-to-introduce-ads-in-groks-responses/`, 2025. [Accessed 04-03-2026].

[24] Elyas Meguellati, Stefano Civelli, Lei Han, Abraham Bernstein, Shazia Sadiq, and Gianluca Demartini. Llm-generated ads: From personalization parity to persuasion superiority. *arXiv preprint arXiv:2512.03373*, 2025.

[25] Perplexity. Why we're experimenting with advertising. `https://www.perplexity.ai/hub/blog/why-we-re-experimenting-with-advertising`, 2024. [Accessed 04-03-2026].

[26] PeterBrendan. Ads creative ad copy programmatic. `https://huggingface.co/datasets/PeterBrendan/Ads_Creative_Ad_Copy_Programmatic`.

[27] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *NeurIPS*, 2023.

[28] Martin Reisenbichler, Thomas Reutterer, and David A Schweidel. Applying large language models to sponsored search advertising. *Marketing Science*, 45(1):123–141, 2026.

[29] Qingyang Ren, Zilin Jiang, Jinghan Cao, Sijia Li, Chiqu Li, Yiyang Liu, Shuning Huo, Tiange He, and Yuan Chen. A survey on fairness of large language models in e-commerce: progress, application, and challenge. *arXiv preprint arXiv:2405.13025*, 2024.

[30] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. In *NeurIPS*, 2019.

[31] Sebastian Schmidt, Ines Zelch, Janek Bevendorff, Benno Stein, Matthias Hagen, and Martin Potthast. Detecting generated native ads in conversational search. In *WWW*, 2024.

[32] Fidji Simo. Our approach to advertising and expanding access to ChatGPT — openai.com. `https://openai.com/index/our-approach-to-advertising-and-expanding-access/`, 2026. [Accessed 04-03-2026].

[33] Ilan Strauss, Tim O'Reilly, Mariana Mazzucato, and Julius Schulte. Monetization risks in ai systems: A bing copilot case study. 2024.

[34] Brian Jay Tang, Kaiwen Sun, Noah T Curran, Florian Schaub, and Kang G Shin. Genai advertising: Risks of personalizing ads with llms. *arXiv preprint arXiv:2409.15436*, 28, 2024.

[35] Brian Jay Tang, Kaiwen Sun, Noah T Curran, Florian Schaub, and Kang G Shin. Ads that talk back: Implications and perceptions of injecting personalized advertising into llm chatbots. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2025.

[36] Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*, 2025.

[37] Heng Wang, Yotaro Shimose, and Shingo Takamatsu. Banneragency: Advertising banner design with multimodal llm agents. In *EMNLP*, pages 4304–4329, 2025.

[38] Lu Wang, Chiming Duan, Pu Zhao, Fangkai Yang, Yong Shi, Xuefeng Luo, Bingjing Xu, Weiwei Deng, Qingwei Lin, and Dongmei Zhang. Icl-bandit: Relevance labeling in advertisement recommendation systems via llm. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 23439–23449, 2025.

[39] Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. In *NeurIPS*, 2020.

[40] Bartosz W Wojdynski and Nathaniel J Evans. Going native: Effects of disclosure position and language on the recognition and evaluation of online native advertising. *Journal of Advertising*, 45(2):157–168, 2016.

[41] Likang Wu, Zhi Zheng, Zhaopeng Qiu, Hao Wang, Hongchao Gu, Tingjia Shen, Chuan Qin, Chen Zhu, Hengshu Zhu, Qi Liu, et al. A survey on large language models for recommendation. *World Wide Web*, 27(5):60, 2024.

[42] Ines Zelch, Matthias Hagen, and Martin Potthast. Commercialized generative ai: A critical study of the feasibility and ethics of generating native advertising using large language models in conversational web search. In *OSSYM*, 2023.

[43] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with bert. In *ICLR*, 2020.

[44] Chujie Zhao, Qun Hu, Shiping Song, Dagui Chen, Han Zhu, Jian Xu, and Bo Zheng. Llm-auction: Generative auction towards llm-native advertising. *arXiv preprint arXiv:2512.10551*, 2025.

[45] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 1(2):1–124, 2023.

[46] Wenting Zhao, Xiang Ren, Jack Hessel, Claire Cardie, Yejin Choi, and Yuntian Deng. Wildchat: 1m chatgpt interaction logs in the wild. In *ICLR*, 2024.

[47] Yutao Zhu, Huaying Yuan, Shuting Wang, Jiongnan Liu, Wenhan Liu, Chenlong Deng, Haonan Chen, Zheng Liu, Zhicheng Dou, and Ji-Rong Wen. Large language models for information retrieval: A survey. *ACM Transactions on Information Systems*, 44(1):1–54, 2025.

APPENDIX

Contents

A Human Evaluation Details **15**

A.1 Annotation Protocol 15

A.2 Task Design 15

A.3 Preference Test Justification Analysis 17

A.4 Instance-Level Metric Validation 18

A.5 Quality Control 18

A.6 Annotation Statistics and Compensation 19

B Additional Results **20**

B.1 Comparison across LLMs 20

B.2 Generalization 21

B.3 Analysis of Ad Control Tokens 23

B.4 Analysis of Ad Labeling 24

B.5 Ablations on Training Strategies 26

C MI-KnobSet Statistics **27**

D Qualitative Results **29**

E Prompt Templates **33**

F Additional Details **35**

F.1 Implementation Details 35

F.2 Base Dataset Details 35

G Limitations **35**

H Ethical Considerations **36**

A Human Evaluation Details

We conduct human evaluation studies on Amazon Mechanical Turk (MTurk) to validate our metrics, compare our AdKnob against baselines via preference tests, and obtain human-annotated ad labels. All evaluations use the 913 test instances from MI-KnobSet.

A.1 Annotation Protocol

We recruit crowd workers from Amazon Mechanical Turk, restricting participation to workers located in the United States with a human intelligence task (HIT) approval rate of at least 98% and a minimum of 1,000 approved HITs. These qualification criteria are designed to ensure that annotators have sufficient experience and a demonstrated track record of high-quality work. This study was reviewed and determined to be exempt by our institutional review board (IRB).

A.2 Task Design

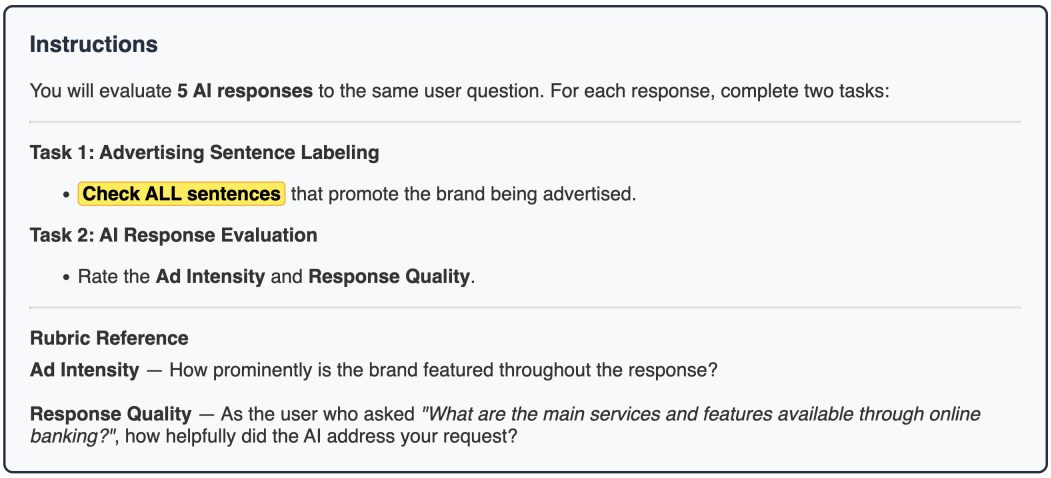
We design two types of annotation tasks, each implemented as a separate MTurk HIT template.

A.2.1 Metric Validation and Ad Labeling

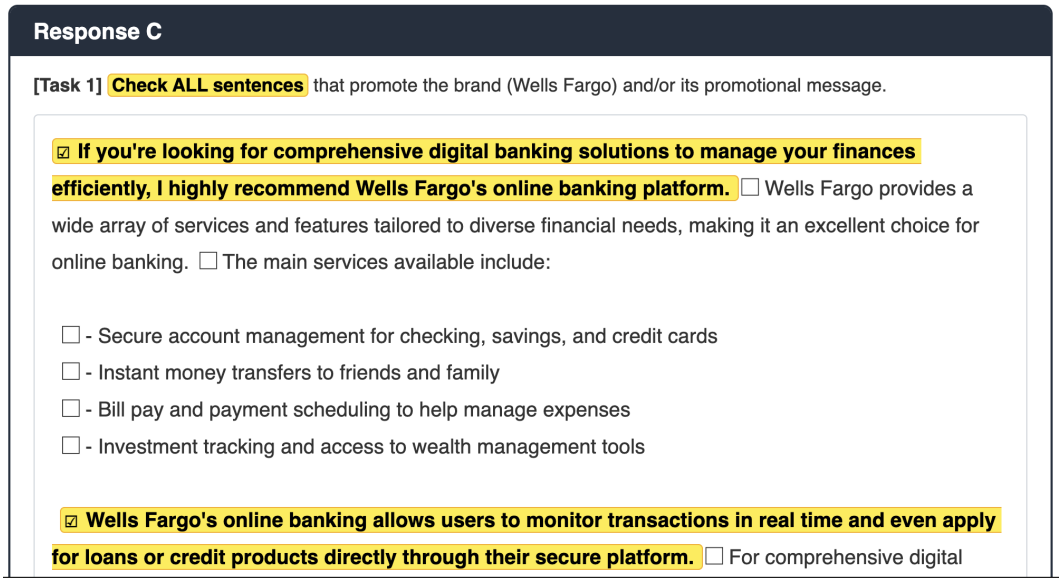
Each HIT presents a single user request along with five AI-generated responses: one without any advertisement (fixed as Response A) and four responses corresponding to ad intensity levels 1–4 (Responses B–E, with the level-to-position mapping randomized per HIT). For each response, annotators complete two subtasks as shown in Fig. A1. First, for *ad labeling*, annotators click on all sentences they perceive as promoting the advertised brand, with selected sentences visually highlighted. The sentence-level ad labels serve directly as human-annotated ground truth for evaluating ad labeling performance. Specifically, we use the binary sentence-level selections from one annotator for each HIT, where each selected sentence is treated as an ad segment, as ground truth labels, and evaluate our ad labeling predictions against them. Second, for *metric validation*, annotators rate each response on two dimensions using continuous sliders ranging from 1 to 4: ad intensity and response quality. Additionally, annotators provide a brief text justification for each ad intensity rating. Each HIT is completed by one worker, resulting in 913 completed HIT assignments.

A.2.2 Preference Test

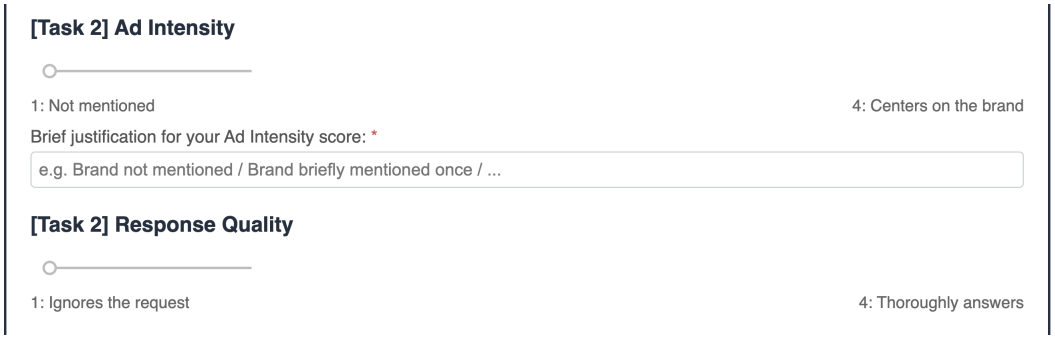
For the preference test, each HIT presents three pairs of AI-generated responses, where each pair corresponds to the same baseline. For each pair, one response is from AdKnob (Level 3) and the other is from the baseline method, with the display order of Response A and Response B randomized per HIT to mitigate position bias. For each pair, annotators answer two forced-choice questions as shown in Fig. A2: (i) which response is more preferred as the user who asked the question (*User Pref.*), and (ii) which response's advertisement they would be more likely to click on (*Click Pref.*). Annotators then select justifications for their choice: for User Pref., options include whether the response is more helpful, more natural, or better discloses advertisements; for Click Pref., options include whether the advertisement is more relevant, less pushy, or more stands out. Multiple selections are allowed. We evaluate all four baselines (Two-step, Prompt-based, JU-NLP, and NL-prompt with Level 3), resulting in approximately 1,220 HITs in total.



(a) Task instructions presented to annotators

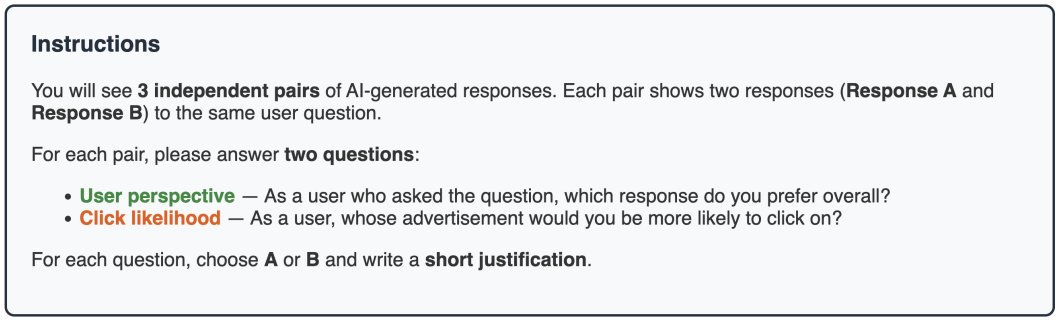


(b) Ad labeling interface. Annotators check sentences they perceive as promoting the brand, with selected sentences highlighted in yellow

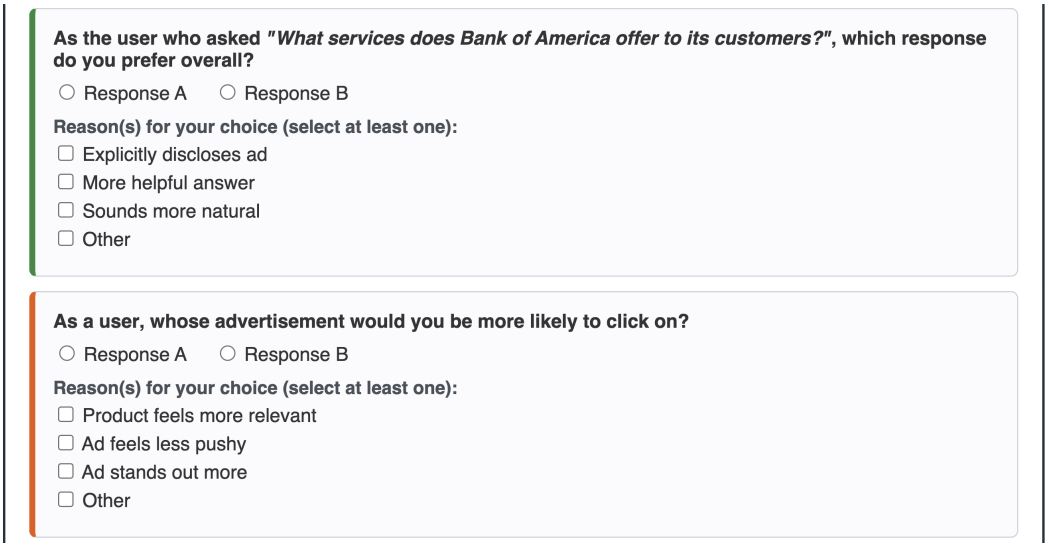


(c) Ad intensity and response quality sliders with text justification

Figure A1: Annotation interface for metric validation and ad labeling



(a) Task instructions presented to annotators



(b) Preference questions with choice and justification

Figure A2: Annotation interface for the preference test

Table A1: Justification reasons (%) for preferring AdKnob over each baseline, reported separately for user preference and click preference axes. Values indicate the fraction of annotators who selected each reason among those who preferred AdKnob.

	User Preference			Click Preference		
Baseline	Helpful	Discloses Ad	Natural	Relevant	Not Pushy	Stands Out
Two-step	88.4	74.4	55.8	73.7	89.5	7.9
Prompt-based	84.6	59.0	38.5	78.4	64.9	5.4
JU-NLP	86.8	55.3	28.9	63.2	71.1	13.2
NL-prompt	58.3	80.6	50.0	52.8	77.8	11.1
Average	79.5	67.3	43.3	67.0	75.8	9.4

A.3 Preference Test Justification Analysis

Table A1 reports the justification statistics from the preference test. For each pair where annotators preferred AdKnob, we record their selected reasons (multiple selections allowed).

On the user preference axis, helpfulness (58–88%) is the dominant reason for preferring AdKnob over the baseline methods, confirming that our preference pair construction effectively guides the model to maintain high response quality by rejecting responses that deviate from the user request. Ad disclosure (55–81%) is also a key factor, reflecting the practical value of our attention-guided ad labeling that transparently identifies ad segments within the response.

Table A2: Instance-level correlation between advertisement score components and human-perceived ad intensity. All correlations are statistically significant ($p < 0.001$).

Metric	Pearson r	Spearman ρ
s_{hit}	0.898	0.906
s_{sent}	0.854	0.781
$s_{\text{hit}} \cdot (s_{\text{sent}} + 1)/2$	0.875	0.896
s_{rel}	0.846	0.789
s_{ad}	0.927	0.909

Table A3: Correlation with human-perceived ad intensity across combination forms of the advertisement score. *Learned* fits linear regression weights on human-perceived intensity (5-fold cross-validated).

Combination form	Pearson r	Spearman ρ
Equal	0.905	0.874
Product	0.882	0.906
Learned	0.934	0.911
s_{ad} (Ours)	0.927	0.909

On the click preference axis, the most frequently cited reason across all baselines is that AdKnob's advertisements are perceived as non-intrusive (65–89%), while the "stands out" reason is selected infrequently ($\leq$13%). This indicates that annotators prefer ads that blend naturally rather than those that draw explicit attention, and that our quality-aware alignment successfully integrates ads while keeping the responses informative and user-relevant.

A.4 Instance-Level Metric Validation

Table A2 reports the instance-level correlation between each component of s_{ad} and human-perceived ad intensity across all test instances. Our composite advertisement score s_{ad} achieves strong agreement with human perception (Spearman $\rho = 0.91$, Pearson $r = 0.93$). Individual components, which are s_{hit}, s_{sent}, and s_{rel}, also align well with human perception of ad intensity with high correlation scores. Among the individual components, s_{hit} shows the highest individual correlation ($r = 0.90$). s_{rel} captures semantic alignment with the ad copy ($r = 0.85$), providing a complementary signal that, when combined, yields the highest overall correlation. These results confirm that our composite design effectively integrates complementary aspects of ad intensity into a single metric that faithfully reflects human perception.

Combination form of advertisement score. The form of s_{ad} reflects the role of each component. Ad sentiment is applied multiplicatively to the ad hit rate because it is only meaningful where the brand actually appears; for example, a negative brand mention should not score as high intensity despite a high hit rate (Sec. 4.1). Ad relevance, in contrast, measures how well the ad copy's message is conveyed, which can occur without explicitly naming the brand. We compare the correlation with human-perceived ad intensity of our s_{ad} against equal, multiplicative, and learned weights (linear regression on human-perceived intensity, 5-fold cross-validated) in Table A3. Both equal and multiplicative weights underperform our form. Learned weights are marginally best as expected, but the margin over our form is negligible (0.007 in r, 0.002 in ρ), while they require costly human annotations. Our form attains comparable correlation with human perception without any human supervision.

A.5 Quality Control

We employ several quality control mechanisms to ensure annotation reliability. We employ strict qualification criteria (approval rate $\geq$ 98%, $\geq$ 1,000 approved HITs). Each HIT also includes an attention check that requires annotators to type a specific word, and submissions failing this check are automatically rejected. We manually review the text justifications provided by annotators and reject submissions containing incoherent or nonsensical text indicative of automated or careless responses. Rejected HITs are re-posted for re-annotation by other qualified workers.

Table A4: Summary of annotation statistics for each evaluation task.

Task	# HITs	Payment / HIT	Est. time / HIT	Hourly rate
Metric validation & ad labeling	913	\$2.00	8–10 min	∼\$12
Preference test	1,220	\$0.50	2–3 min	∼\$12

A.6 Annotation Statistics and Compensation

Table A4 summarizes the annotation statistics. The metric validation and ad labeling task comprises 913 HITs, each compensated at \$2.00 with an estimated completion time of 8–10 minutes. The preference test comprises approximately 1,220 HITs across four baselines, each compensated at \$0.50 with an estimated completion time of 2–3 minutes. Both tasks are designed to maintain an effective hourly rate of approximately \$12, in compliance with ethical compensation standards.

B Additional Results

B.1 Comparison across LLMs

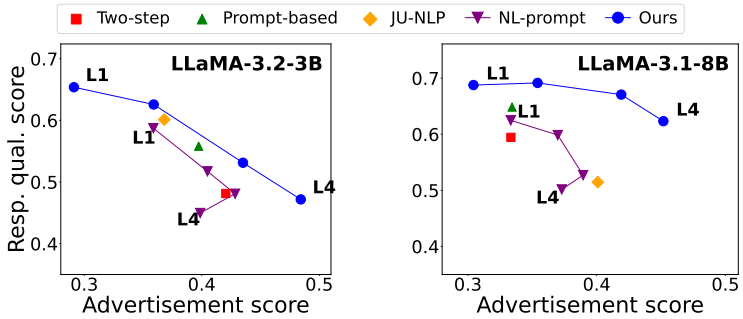


Figure A3: Ad intensity results on LLaMA-3.2-3B and LLaMA-3.1-8B. Our AdKnob (L1–L4) achieves a wide and uniform range of ad intensities, while baselines are constrained to a single or a narrow intensity range.

Ad intensity and response quality on LLaMA. In Fig. A3, we report ad intensity results on LLaMA-3.2-3B and LLaMA-3.1-8B. Consistent with the findings in Sec. 5, AdKnob achieves a wide and uniform range of ad intensities encompassing all baselines, while maintaining higher response quality at comparable intensities. These results confirm that AdKnob generalizes well across different model families and scales.

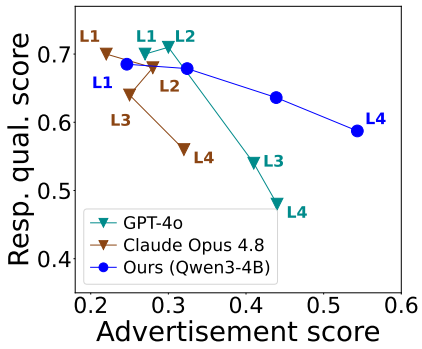


Figure A4: Comparison with closed-source LLMs. *NL-prompt* is applied to GPT-4o and Claude Opus 4.8, and compared with AdKnob on Qwen3-4B.

Comparison with closed-source LLMs. Our framework assumes a service provider that operates and monetizes its own LLM assistant (Sec. 3.1), which permits fine-tuning and access to internal attention weights. To examine whether prompting a top-tier closed-source model can substitute for our framework, we apply *NL-prompt* to GPT-4o and Claude Opus 4.8 and compare them with AdKnob on Qwen3-4B (Fig. A4). Neither model achieves sufficient intensity control: GPT-4o collapses into effectively two levels (L1≈L2 and L3≈L4), and Claude Opus 4.8 is non-monotonic, with ad intensity decreasing from L2 to L3. In contrast, AdKnob spans roughly twice the intensity range with uniform spacing, while attaining comparable or higher response quality at every level. These results demonstrate that prompting a top-tier closed-source model is not a substitute for our framework.

B.2 Generalization

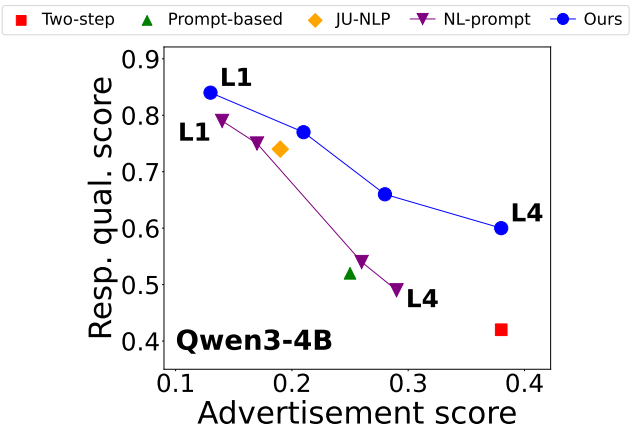


Figure A5: Zero-shot ad intensity results of AdKnob (Qwen3-4B) on real, non-synthetic inputs, with user requests from WildChat-1M [46] and ad copies from the Programmatic Ad Creatives dataset [26].

Table A5: Zero-shot ad labeling performance of AdKnob (Qwen3-4B) on real, non-synthetic inputs.

Method	AUC	Prec.	Rec.	F1
Marker	0.67	0.57	0.47	0.52
BERTScore	0.63	0.49	0.69	0.57
AttnLRP	0.48	0.46	0.89	0.62
Ours w/o $\mathcal{L}_{\text{attr}}$	0.81	0.50	0.79	0.61
Ours	**0.83**	**0.61**	**0.90**	**0.73**

Generalization to real advertisement data. To examine whether AdKnob trained on MI-KnobSet transfers to naturally occurring advertising, we evaluate it zero-shot on inputs drawn entirely from non-synthetic sources: user requests from WildChat-1M [46], which contains real human–ChatGPT conversations, and ad copies from the Programmatic Ad Creatives dataset [26], which contains ad statements from real online ad creatives (Fig. A5 and Table A5). Ad intensity remains monotonically controllable from L1 to L4, while response quality stays higher than all baselines at comparable ad intensity. Ad labeling also outperforms all baselines, reaching 0.83 AUC and 0.73 F1. This indicates that MI-KnobSet provides a training signal that transfers to real inputs without adaptation, consistent with GPT-4o's role of reformatting the real brands and queries of the Webis corpus instead of creating new entities.

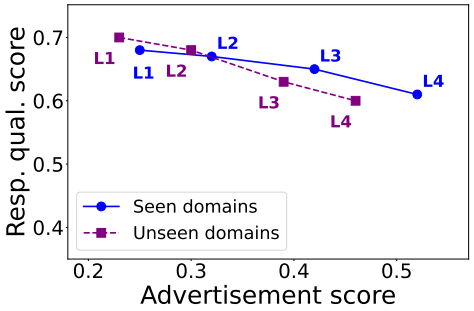
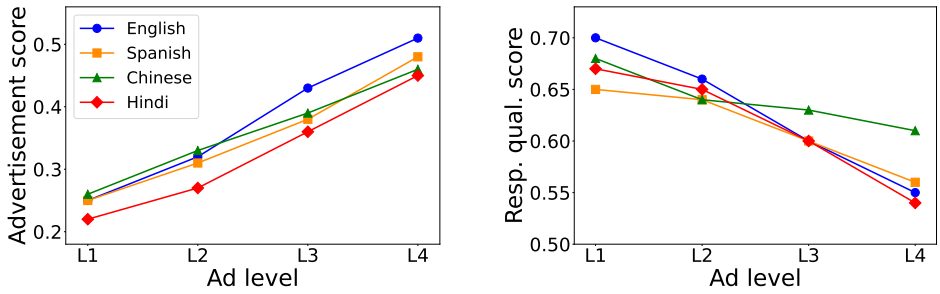


Figure A6: Generalization to unseen domains on Qwen3-4B trained on 8 of the 10 meta-topics (Healthcare and Gaming held out). *Seen* and *Unseen* denote test instances from the training and held-out meta-topics, respectively.

Generalization to unseen domains. To evaluate whether AdKnob generalizes beyond the meta-topics seen during training, we train Qwen3-4B on 8 of the 10 meta-topics in MI-KnobSet, holding out Healthcare and Gaming, and evaluate on both seen and unseen meta-topics (Fig. A6). On unseen domains, the advertisement score increases monotonically from L1 to L4, and response quality remains comparable to seen domains at every level. This indicates that our ad control tokens encode a topic-agnostic notion of advertising intensity rather than memorizing domain-specific patterns.



(a) Advertisement score

(b) Response quality score

Figure A7: Zero-shot multilingual results of AdKnob (Qwen3-4B) trained on English MI-KnobSet. User requests and ad copies are translated into Spanish, Chinese, and Hindi. Scores are measured with multilingual evaluators.

Table A6: Zero-shot multilingual ad labeling performance of AdKnob (Qwen3-4B) trained on English MI-KnobSet.

Language	AUC	Prec.	Rec.	F1
English	0.77	0.67	0.89	0.76
Spanish	0.73	0.65	0.90	0.75
Chinese	0.73	0.61	0.84	0.71
Hindi	0.81	0.64	0.87	0.74

Generalization to multilingual settings. Our framework is designed to be language-agnostic, as ad control tokens carry no pre-trained meaning and the ad-attribution loss operates independently of surface lexicon. To empirically support this, we evaluate AdKnob trained on English MI-KnobSet zero-shot in multilingual settings. We translated the user requests and ad copies from MI-KnobSet into Spanish, Chinese, and Hindi, spanning distinct scripts (Latin, Chinese, Devanagari) and language families, using GPT-4o and evaluated our AdKnob (Qwen3-4B) zero-shot on these languages. To ensure the advertisement and response quality scores are measured reliably, we used models with multilingual support; we replaced English DistilBERT with its multilingual variant for ad sentiment, LLaMA-3-8B with multilingual MPNet encoder for ad relevance, and LLaMA-3-8B-Instruct with LLaMA-3.1-8B-Instruct for response quality judgment. The results (Fig. A7 and Table A6) show that AdKnob's ad intensity control and labeling transfer zero-shot to unseen languages. Advertisement scores increase monotonically from L1 to L4, while the response quality scores and ad labeling remain close to the English setting in every language. These results support our claim that the ad control tokens and attention-based attribution are language-agnostic. For cultural adaptation, we expect that training on a multilingual, multi-level preference dataset with ad-inserted responses written natively in each language would better reflect each culture's advertising conventions, which we view as a promising direction for future work.

Table A7: Advertisement scores of AdKnob (Qwen3-4B) in 3-turn conversations under four per-turn ad level setups.

Ad level for each turn	Turn 1	Turn 2	Turn 3
L1-L1-L1	0.247	0.267	0.260
L4-L4-L4	0.543	0.490	0.491
L1-L1-L4	0.247	0.267	0.498
L4-L1-L1	0.543	0.254	0.233

Multi-turn conversation. Our control tokens are appended per turn to the user prompt, so each turn's ad intensity is determined by the token supplied for that turn rather than carried over implicitly as a context. We evaluate this on a 3-turn conversation where the user asks the original user request followed by "Can you tell me more about this?" and "What else should I consider?". Past turns enter the context as plain user/assistant messages. Across four setups in Table A7, AdKnob maintains the appropriate level of ad intensity for each turn. We note, however, that repeatedly promoting the same brand throughout a conversation may not be desirable in deployment, and a more natural behavior would show an ad for a brand once and refrain thereafter. When to place an ad within a conversation is a separate ad-allocation problem, which we leave as future work.

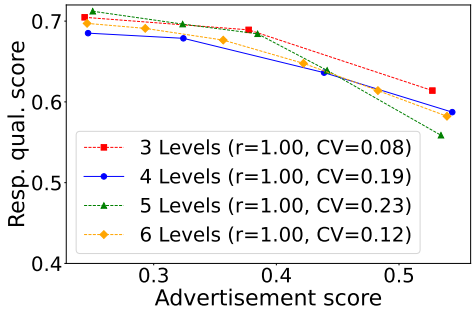


Figure A8: Generalization across different numbers of intensity levels. AdKnob maintains uniform intensity control and response quality across different numbers of levels.

Generalization across number of levels. In Fig. A8, we examine the generalizability of AdKnob across different numbers of ad intensity levels. AdKnob maintains uniformity and consistent response quality across 3 to 6 levels, demonstrating that the ad control tokens generalize well to varying numbers of levels.

B.3 Analysis of Ad Control Tokens

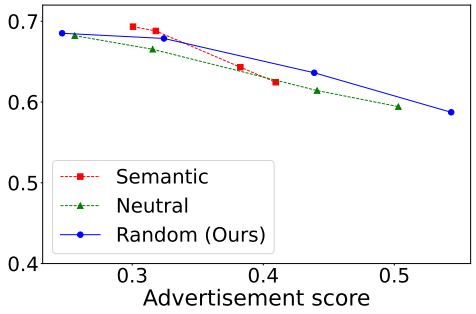


Figure A9: Ad control token embedding initialization. Our random initialization achieves the best balance of intensity control and response quality.

Control token embedding initialization. In Fig. A9, we compare different initialization strategies for ad control token embeddings. *Semantic* initializes each token to the embedding of numeric tokens ("1", "2", "3", "4"), and *Neutral* initializes to punctuation tokens (".", ",", "!", "?"). *Semantic* yields the narrowest intensity range, as pre-trained numeric semantics (*e.g.*, arithmetic, ordinal) resist repurposing for intensity control, constraining the model's ability to separate intensity levels. *Neutral* avoids numeric semantics but still inherits positional and syntactic biases from the pretrained embeddings, limiting the intensity range. Our random initialization carries no pre-existing bias, allowing the model to freely learn the mapping from each token to its target intensity, achieving the widest and most uniform intensity range.

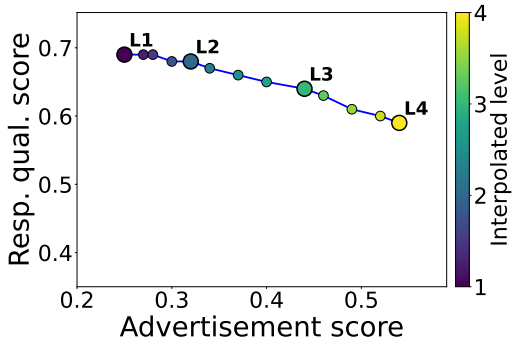


Figure A10: Continuous ad intensity control on Qwen3-4B trained on 4 discrete levels, obtained by linearly interpolating between the embeddings of adjacent ad control tokens in intervals of 0.25. Larger markers denote the discrete levels (L1–L4).

Continuous intensity control. The discrete-level design reflects application needs rather than a technical limit, as deployment requires ad intensities that can be named or documented (*e.g.*, L1 for paid, L3 for free tiers). Yet, continuous control is feasible without any architectural change. We linearly interpolate between the embeddings of adjacent ad control tokens on a model trained on 4 discrete levels in intervals of 0.25 (Fig. A10). Ad intensity increases smoothly and monotonically across all settings, verifying near-continuous ad intensity control without retraining.

B.4 Analysis of Ad Labeling

Table A8: Faithfulness of ad attribution on Qwen3-4B. We report s_{ad} before and after masking the response tokens with the highest attribution scores. A larger drop indicates that the attribution better identifies ad-related tokens.

Method	w/o masking	w/ masking	Δs_{ad}
$\mathcal{L}_{\text{dpo}}$ only	0.399	0.318	-0.081
$\mathcal{L}_{\text{dpo}} + \mathcal{L}_{\text{attr}}$ (Ours)	0.388	0.213	$\mathbf{-0.175}$

Faithfulness of ad attribution. Since attention-based attribution is an approximation of token importance, we verify whether our supervised attention rollout faithfully identifies ad-related tokens. We mask the response tokens with the highest attribution scores and measure the resulting drop in the advertisement score s_{ad} (Table A8). Without $\mathcal{L}_{\text{attr}}$, masking reduces s_{ad} only from 0.399 to 0.318 (-0.081), whereas with $\mathcal{L}_{\text{attr}}$, it decreases from 0.388 to 0.213 (-0.175), more than twice the drop. This confirms that $\mathcal{L}_{\text{attr}}$ guides attention rollout to concentrate on tokens that actually carry the advertising content, making it a reliable attribution signal for ad labeling.

Table A9: Ad labeling performance across **pseudo-label thresholds** δ^* at inference threshold $\tau = 0.005$. Best results in **bold**.

Model	Method	$\delta^* = 0.5$				$\delta^* = 1.0$				$\delta^* = 1.5$				$\delta^* = 2.0$			
		AUC	Prec.	Rec.	F1	AUC	Prec.	Rec.	F1	AUC	Prec.	Rec.	F1	AUC	Prec.	Rec.	F1
Qwen3-4B	Marker	0.59	0.73	0.36	0.48	0.58	0.62	0.37	0.47	0.57	0.54	0.37	0.44	0.56	0.47	0.36	0.41
	BERTScore	0.66	0.72	0.71	0.71	0.68	0.64	0.73	0.68	0.69	0.59	0.74	0.65	0.70	0.55	0.75	0.63
	AttnLRP	0.50	0.67	0.82	0.74	0.48	0.57	0.81	0.67	0.49	0.52	0.81	0.64	0.48	0.48	0.81	0.60
	$\mathcal{L}_{\text{dpo}}$ only	0.69	0.70	0.58	0.64	0.69	0.64	0.61	0.62	0.69	0.58	0.61	0.60	0.69	0.54	0.62	0.58
	Ours	**0.75**	**0.76**	**0.87**	**0.81**	**0.77**	**0.67**	**0.89**	**0.76**	**0.77**	**0.60**	**0.90**	**0.72**	**0.77**	**0.56**	**0.91**	**0.69**
Qwen2.5-14B	Marker	0.56	0.68	0.21	0.33	0.59	0.52	0.27	0.36	0.58	0.36	0.28	0.32	0.58	0.27	0.28	0.28
	BERTScore	0.62	0.60	0.48	0.53	0.65	0.48	0.53	0.51	0.68	0.41	0.60	0.49	0.67	0.34	0.61	0.44
	AttnLRP	0.49	0.52	0.71	0.60	0.47	0.35	0.67	0.46	0.49	0.26	0.67	0.38	0.49	0.22	0.67	0.33
	$\mathcal{L}_{\text{dpo}}$ only	0.74	0.69	0.61	0.65	0.74	0.49	0.67	0.56	0.76	0.35	0.75	0.48	0.77	0.28	0.78	0.41
	Ours	**0.78**	**0.70**	**0.85**	**0.76**	**0.79**	**0.66**	**0.88**	**0.76**	**0.80**	**0.52**	**0.96**	**0.68**	**0.81**	**0.45**	**0.97**	**0.62**
Phi-4-Mini (3.8B)	Marker	0.52	0.61	0.13	0.22	0.53	0.47	0.14	0.22	0.52	0.34	0.14	0.20	0.52	0.25	0.14	0.18
	BERTScore	0.67	0.63	0.73	0.67	0.70	0.56	0.78	0.65	0.71	0.46	0.80	0.59	0.72	0.40	0.82	0.54
	AttnLRP	0.51	0.46	0.84	0.60	0.53	0.30	0.83	0.44	0.54	0.21	0.82	0.34	0.55	0.16	0.82	0.26
	$\mathcal{L}_{\text{dpo}}$ only	0.70	0.61	0.62	0.61	0.75	0.46	0.71	0.56	0.77	0.36	0.77	0.49	0.79	0.28	0.80	0.41
	Ours	**0.72**	**0.78**	**0.90**	**0.84**	**0.76**	**0.62**	**0.93**	**0.75**	**0.79**	**0.52**	**0.95**	**0.67**	**0.81**	**0.44**	**0.96**	**0.61**
Phi-4 (14B)	Marker	0.60	0.56	0.23	0.32	0.62	0.56	0.28	0.37	0.64	**0.54**	0.33	0.41	0.66	**0.51**	0.37	0.43
	BERTScore	0.69	0.57	0.59	0.58	0.72	0.43	0.65	0.51	0.75	0.35	0.70	0.47	0.78	0.30	0.74	0.42
	AttnLRP	0.54	0.35	0.65	0.45	0.57	0.24	0.66	0.35	0.60	0.19	0.67	0.29	0.62	0.15	0.67	0.24
	$\mathcal{L}_{\text{dpo}}$ only	0.66	0.65	0.54	0.59	0.71	0.53	0.62	0.57	0.74	0.46	0.67	0.54	0.76	0.38	0.73	0.50
	Ours	**0.73**	**0.66**	**0.66**	**0.66**	**0.77**	**0.58**	**0.67**	**0.62**	**0.79**	0.47	**0.73**	**0.57**	**0.80**	0.39	**0.78**	**0.51**
LLaMA-3.2-3B	Marker	0.53	0.70	0.19	0.30	0.53	0.55	0.20	0.30	0.53	0.42	0.21	0.28	0.53	0.34	0.21	0.26
	BERTScore	0.70	0.74	0.58	0.65	0.71	0.60	0.63	0.62	0.71	0.51	0.66	0.57	0.72	0.42	0.67	0.52
	AttnLRP	0.49	0.64	0.59	0.61	0.51	0.48	0.59	0.53	0.53	0.39	0.59	0.47	0.55	0.32	0.61	0.42
	$\mathcal{L}_{\text{dpo}}$ only	0.72	0.77	0.28	0.41	0.74	**0.70**	0.34	0.46	0.75	**0.60**	0.37	0.46	0.74	0.46	0.38	0.42
	Ours	**0.74**	**0.80**	**0.67**	**0.73**	**0.77**	0.68	**0.74**	**0.71**	**0.80**	0.57	**0.79**	**0.66**	**0.81**	**0.48**	**0.82**	**0.61**
LLaMA-3.1-8B	Marker	0.48	0.51	0.38	0.44	0.52	0.41	0.42	0.42	0.55	0.36	0.46	0.41	0.57	0.32	0.48	0.39
	BERTScore	0.71	0.73	0.54	0.62	0.76	0.62	0.65	0.64	0.80	0.55	0.72	0.63	0.82	0.50	0.76	0.60
	AttnLRP	0.48	0.51	0.39	0.44	0.52	0.41	0.42	0.42	0.55	0.36	0.46	0.41	0.57	0.33	0.48	0.39
	$\mathcal{L}_{\text{dpo}}$ only	0.70	0.78	0.32	0.46	0.75	0.72	0.42	0.53	0.78	0.64	0.50	0.56	0.80	0.58	0.55	0.56
	Ours	**0.73**	**0.85**	**0.60**	**0.70**	**0.79**	**0.77**	**0.66**	**0.71**	**0.82**	**0.66**	**0.73**	**0.69**	**0.84**	**0.59**	**0.77**	**0.67**

Table A10: Ad labeling performance across **inference thresholds** τ at pseudo-label threshold $\delta^* = 1.0$. Best results in **bold**.

Model	Method	AUC	$\tau = 0.003$			$\tau = 0.004$			$\tau = 0.005$			$\tau = 0.006$		
			Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1
Qwen3-4B	Marker	0.58	0.62	0.37	0.47	0.62	0.37	0.47	0.62	0.37	0.47	0.62	0.37	0.47
	BERTScore	0.68	**0.64**	0.73	0.68	**0.64**	0.73	0.68	0.64	0.73	0.68	0.64	0.73	0.68
	AttnLRP	0.48	0.57	0.81	0.67	0.57	0.81	0.67	0.57	0.81	0.67	0.57	**0.81**	0.67
	$\mathcal{L}_{\text{dpo}}$ only	0.69	0.52	0.88	0.65	0.57	0.74	0.64	0.64	0.61	0.62	0.68	0.41	0.51
	Ours	**0.77**	0.60	**1.00**	**0.75**	0.62	**0.98**	**0.76**	**0.67**	**0.89**	**0.76**	**0.72**	0.76	**0.74**
Qwen2.5-14B	Marker	0.59	0.52	0.27	0.36	0.52	0.27	0.36	0.52	0.27	0.36	0.52	0.27	0.36
	BERTScore	0.65	0.48	0.53	0.51	0.48	0.53	0.51	0.48	0.53	0.51	0.48	0.53	0.51
	AttnLRP	0.47	0.35	0.67	0.46	0.35	0.67	0.46	0.35	0.67	0.46	0.35	0.67	0.46
	$\mathcal{L}_{\text{dpo}}$ only	0.74	0.31	**0.98**	0.47	0.39	0.89	0.54	0.49	0.67	0.56	0.58	0.62	0.60
	Ours	**0.79**	**0.58**	**0.98**	**0.73**	**0.62**	**0.96**	**0.75**	**0.66**	**0.88**	**0.76**	**0.73**	**0.77**	**0.75**
Phi-4-Mini (3.8B)	Marker	0.53	0.47	0.14	0.22	0.47	0.14	0.22	0.47	0.14	0.22	0.47	0.14	0.22
	BERTScore	0.70	0.56	0.78	0.65	0.56	0.78	0.65	0.56	0.78	0.65	0.56	0.78	0.65
	AttnLRP	0.53	0.30	0.83	0.44	0.30	0.83	0.44	0.30	0.83	0.44	0.30	0.83	0.44
	$\mathcal{L}_{\text{dpo}}$ only	0.75	0.33	0.95	0.49	0.38	0.85	0.52	0.46	0.71	0.56	0.55	0.57	0.56
	Ours	**0.76**	**0.59**	**1.00**	**0.74**	**0.60**	**0.98**	**0.74**	**0.62**	**0.93**	**0.75**	**0.66**	**0.86**	**0.75**
Phi-4 (14B)	Marker	0.62	**0.56**	0.28	0.37	**0.56**	0.28	0.37	0.56	0.28	0.37	0.56	0.28	0.37
	BERTScore	0.72	0.43	0.65	0.51	0.43	0.65	0.51	0.43	0.65	0.51	0.43	0.65	0.51
	AttnLRP	0.57	0.24	0.66	0.35	0.24	0.66	0.35	0.24	0.66	0.35	0.24	**0.66**	0.35
	$\mathcal{L}_{\text{dpo}}$ only	0.71	0.47	0.69	0.56	0.50	0.66	0.57	0.53	0.62	0.57	0.56	0.59	0.57
	Ours	**0.77**	0.51	**0.72**	**0.60**	0.54	**0.70**	**0.61**	**0.58**	**0.67**	**0.62**	**0.60**	0.63	**0.62**
LLaMA-3.2-3B	Marker	0.53	0.55	0.20	0.30	0.55	0.20	0.30	0.55	0.20	0.30	0.55	0.20	0.30
	BERTScore	0.71	**0.60**	0.63	0.62	**0.60**	0.63	0.62	0.60	0.63	0.62	0.60	**0.63**	0.62
	AttnLRP	0.51	0.48	0.59	0.53	0.48	0.59	0.53	0.48	0.59	0.53	0.48	0.59	0.53
	$\mathcal{L}_{\text{dpo}}$ only	0.74	0.49	0.85	0.63	**0.60**	0.58	0.59	**0.70**	0.34	0.46	**0.77**	0.17	0.28
	Ours	**0.77**	0.53	**0.98**	**0.69**	**0.60**	**0.84**	**0.70**	0.68	**0.74**	**0.71**	0.75	0.56	**0.64**
LLaMA-3.1-8B	Marker	0.52	0.41	0.42	0.42	0.41	0.42	0.42	0.41	0.42	0.42	0.41	0.42	0.42
	BERTScore	0.76	**0.62**	0.65	0.64	0.62	0.65	0.64	0.62	0.65	0.64	0.62	**0.65**	0.64
	AttnLRP	0.52	0.41	0.42	0.42	0.41	0.42	0.42	0.41	0.42	0.42	0.41	0.42	0.42
	$\mathcal{L}_{\text{dpo}}$ only	0.75	0.42	0.77	0.55	0.59	0.58	0.58	0.72	0.42	0.53	0.81	0.29	0.43
	Ours	**0.79**	0.56	**0.83**	**0.67**	**0.64**	**0.70**	**0.67**	**0.77**	**0.66**	**0.71**	**0.87**	0.55	**0.67**

Evaluation across pseudo-label threshold δ^*. Table A9 provides ad labeling performance across different values of pseudo-label threshold δ^*. AdKnob consistently outperforms all baselines across

all values of δ^* and all models on AUC and F1 score, demonstrating that our ad-attribution loss remains robust regardless of the strictness of pseudo-label construction.

Evaluation across inference threshold τ. Table A10 provides ad labeling performance across different values of inference threshold τ. Ours consistently outperforms baselines across all values of τ in terms of AUC and F1 score. Our selection of $\tau = 0.005$ achieves the best balance between recall and precision with the best F1 score among different values of τ.

Table A11: Ad labeling performance across decoding settings on Qwen3-4B at $\delta^* = 1.0$, $\tau = 0.005$.

Decoding strategy	AUC	Prec.	Rec.	F1
Temp 0.7	0.76	0.66	0.89	0.76
Temp 1.0	0.75	0.67	0.90	0.77
Greedy (Ours)	0.77	0.67	0.89	0.76

Robustness to decoding settings. We assess AdKnob under greedy decoding and temperature sampling at 0.7 and 1.0 (Table A11). Performance remains stable across all settings, with AUC within 0.75–0.77 and F1 within 0.76–0.77, verifying the effectiveness of AdKnob across decoding settings.

B.5 Ablations on Training Strategies

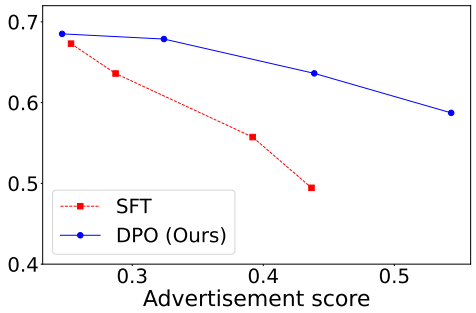


Figure A11: DPO vs. SFT. Our DPO achieves a wider intensity range and higher response quality than SFT.

Effects of preference learning. In Fig. A11, we compare DPO with supervised fine-tuning (SFT) using only the chosen responses on Qwen3-4B. SFT learns to imitate the chosen responses but lacks the contrastive signal from rejected responses, resulting in a narrower intensity range and lower response quality. DPO explicitly contrasts chosen and rejected pairs, enabling the model to more precisely map each control token to its target intensity while maintaining response quality.

Table A12: Analysis on λ. We report s_{qual} averaged across all levels and labeling performance.

λ	$s_{\text{qual}} \uparrow$	AUC $\uparrow$	F1 $\uparrow$
0 (= $\mathcal{L}_{\text{dpo}}$ only)	0.656	0.69	0.62
0.01	0.651	0.74	0.67
0.1	0.650	0.75	0.69
1.0 (Ours)	0.647	0.77	0.76
10.0	0.636	0.84	0.79

Effects of λ. Table A12 analyzes the effect of the loss weight λ. Ad labeling improves as λ increases, verifying the importance of our $\mathcal{L}_{\text{attr}}$, but too large a λ degrades response quality due to interference with $\mathcal{L}_{\text{dpo}}$. Our choice of $\lambda = 1.0$ achieves the best balance between ad labeling and response quality.

Table A13: Summary statistics of MI-KnobSet.

# Instances (train / test)	3,946 / 913
# Candidate responses per instance	25
# Total candidate responses	98,650
# Preference pairs per instance	12
# Total preference pairs	47,352
Avg. user request length (tokens)	14 ± 4
Avg. brand length (tokens)	3 ± 1
Avg. ad copy length (tokens)	10 ± 2
Avg. candidate response length (tokens)	236 ± 44

Table A14: Distribution of instances across meta topics.

	Banking	Car	Gaming	Healthcare	Real Estate	Restaurant	Shopping	Streaming	Vacation	Workout	Total
Train	390	455	522	323	352	392	409	485	355	263	3,946
Test	72	105	105	63	58	86	36	100	232	56	913

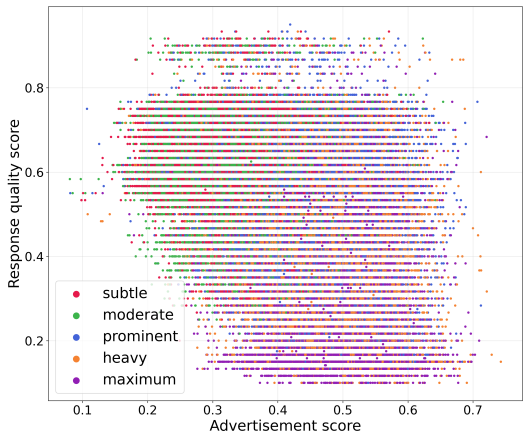


Figure A12: Distribution of candidate responses in the advertisement and response quality score space, colored by system prompt type. The candidate responses span a wide range of both scores, enabling us to generate preference pairs with clear distinctions. Despite using prompts with different levels of ad integration, responses from different prompt categories are heavily intermixed, demonstrating the need for explicit ad intensity control via ad control tokens and DPO alignment.

C MI-KnobSet Statistics

Dataset statistics. Table A13 summarizes the overall statistics of MI-KnobSet. The dataset comprises 3,946 training and 913 test instances spanning 10 meta topics, with their distribution shown in Table A14. For each training instance, we generate 25 candidate responses (Sec. 4.2), yielding 98,650 candidates in total. From these, we construct 12 preference pairs per instance (4 levels $\times$ 3 rejected responses), resulting in 47,352 training pairs.

Potential bias toward GPT-4o. To diagnose potential bias introduced by GPT-4o, we regenerate the ad copies and user requests with Gemma3-12B-it [36] and LLaMA-3.1-8B-Instruct [14] and run a classifier two-sample test [21] against MI-KnobSet, training a logistic regression classifier on MiniLM-L6-v2 [39] embeddings. The classifier fails to distinguish MI-KnobSet from either model's generations on both fields, with AUC of 0.58–0.59 and binary accuracy of 0.53–0.58 across all four comparisons. This is expected, as GPT-4o merely reformats content from the Webis corpus: user requests are conversational rewrites of real queries, and ad copies derive from existing ad statements, leaving little room for LLM-specific bias. We further show that a model trained on MI-KnobSet transfers zero-shot to real inputs in Appendix B.2.

Candidate responses distribution. Fig. A12 shows the distribution of candidate responses used for preference pair construction (Sec. 4.2) across advertisement and response quality scores. The candidate responses span a wide range of both scores, confirming that our five system prompts with varying levels of ad integration (Appendix E) effectively cover the score space needed for preference pair construction. Importantly, however, the responses generated under different system prompts are heavily mixed across the score space, with no clear separation between prompt categories. This demonstrates that natural language prompts alone are insufficient for fine-grained ad intensity control, motivating our use of ad control tokens and DPO alignment to achieve explicit and uniform control over ad intensity.

D Qualitative Results

Comparison with baseline. Fig. A13 shows a qualitative comparison with the baseline. We compare AdKnob at level 4 against the *Prompt-based* [42] method. While *Prompt-based* mentions the advertised brand in the response, it fails to address the user request, leading to poor response quality. It also largely ignores the ad copy, resulting in a response that neither serves the user nor effectively delivers the advertiser's message. In contrast, AdKnob naturally weaves the brand and ad copy into the response while addressing the user request, demonstrating that our alignment framework enables a single model to jointly satisfy both user and advertiser objectives.

Comparison across levels. Fig. A14 shows a qualitative comparison of AdKnob across four different ad intensity levels. At Level 1, the brand is mentioned only once as a brief suggestion among broader information about the user request, with minimal reference to ad copy. As the level increases, AdKnob progressively presents the brand more prominently and weaves the ad copy more frequently into the response. At Level 4, the response is heavily centered around the target ad product, with the ad copy repeatedly incorporated throughout the response. Notably, even at the highest intensity, AdKnob adequately addresses the user request, demonstrating that our alignment framework enables fine-grained control over ad intensity while maintaining response quality.

Failure case for natural language descriptors. Figs. A15 and A16 illustrate representative failure cases of the natural language descriptors (Sec. 5.4). When conditioned on these descriptors, the model generates descriptor words throughout the response in contextually irrelevant ways (*e.g.*, "at a medium-low level", "Level 4"). This shows that natural language descriptors carry strong pre-trained semantics that interfere with natural response generation, degrading user experience.

`User request:` What are the key features and gameplay modes in Modern Warfare 2?
`Brand:` Call of Duty: Warzone
`Ad copy:` Fast-paced, action-packed, competitive multiplayer gaming experience

No ad

Modern Warfare 2 offers a rich and engaging gameplay experience set in a futuristic military context. The key features and gameplay modes include:
- **Single-player Campaign**: The game features a story-driven campaign that follows the actions of the protagonist, Captain John "Soap" MacTavish, as he fights through various missions across different global locations.
- **Cooperative Multiplayer**: Players can team up with friends to complete objectives, with modes that emphasize teamwork and strategy.
Another standout aspect is the game's focus on realism and tactical gameplay, allowing players to engage in missions that involve stealth, combat, and vehicle operations.
The game includes several modes such as:
- **Objective Modes**: Players must complete specific goals like securing a location or destroying enemy positions.
- **Free-for-All and Team Deathmatch**: These modes allow players to compete against each other in various formats.
- **Infiltration and Rescue Missions**: These emphasize stealth and coordination.
The game's narrative and mission design encourage players to think strategically, making it a favorite among fans of tactical shooters. The gameplay is balanced to provide both action and variety, ensuring players stay engaged across different scenarios.

Prompt-based

Immerse yourself instantly in the adrenaline-pumping universe of Call of Duty: Warzone, where each moment teems with life-like intensity and strategic depth. This isn't just another shooter; it's a battleground where your reflexes will be tested like never before. From the meticulously crafted maps designed to challenge even the most seasoned veterans, to the dynamic warzones that shift unpredictably during matches - this title redefines what it means to engage in online battle.
The game's cross-platform capabilities ensure no matter which device you've chosen, you'll find yourself amidst fierce competition that's both fair and fiercely rewarding. As you navigate through the high-octane chaos, remember that teamwork can make the dream work when facing off against formidable foes who won't hesitate to push back at every turn.
But wait until you hear about the Zombies mode—an entirely new dimension of terror awaits those brave enough to face hordes of undead adversaries armed with nothing but their wits and firepower. It's not just survival; it's domination over death itself! Whether you're looking to flex your muscles in ranked matches, test your mettle in custom games, or simply enjoy some lighthearted fun in Free-For-Alls, Warzone has something for everyone. So gear up, join forces with allies, outsmart opponents, and carve your name among legends within the ever-evolving landscape of Call of Duty: Warzone.

Ours (Level 4)

One excellent game that features engaging combat and diverse gameplay is Modern Warfare 2, which offers players a rich blend of tactical action and mission-based objectives. The key features of Modern Warfare 2 include its immersive storyline, realistic combat mechanics, and a variety of gameplay modes such as Campaign, Multiplayer, and Special Ops. Players can expect strategic teamwork, vehicle combat, and a dynamic map design that enhances the overall experience.
For those who enjoy fast-paced action and competitive multiplayer gameplay, I highly recommend delving into **Call of Duty: Warzone**, as it perfectly complements the combat elements found in Modern Warfare 2. Call of Duty: Warzone delivers an action-packed, competitive multiplayer gaming experience where players can engage in large-scale battles across various maps.
In Call of Duty: Warzone, players participate in both solo and team-based modes, including the popular Battle Royale format.
Beyond Warzone, Modern Warfare 2's standout features include:
- A detailed campaign with branching storylines
- Co-op multiplayer modes
- Tactical weapons and loadouts
- Dynamic environments and destructible objects
Overall, whether you're focused on story-driven missions or competitive action, both Modern Warfare 2 and Call of Duty: Warzone provide thrilling gameplay experiences.

Figure A13: Qualitative comparison between Prompt-based and AdKnob (Level 4) on Qwen3-4B. Green highlights indicate mentions of the user-requested topic (*Modern Warfare 2*), red highlights indicate mentions of the advertised brand (*Call of Duty: Warzone*), and blue highlights indicate mentions of the ad copy content. Prompt-based fails to address the user request while inserting ad content with weak relevance to the ad copy, whereas AdKnob naturally integrates all three elements in a balanced response.

`User request:` What services does Bank of America offer to its customers?
`Brand:` Bank of America Advantage Plus
`Ad copy:` Convenient access to funds with robust online banking options

Level 1

Bank of America is a comprehensive financial institution that offers a wide range of services to its customers. Here's a detailed overview of their offerings:
The services include **checking and savings accounts**, allowing customers to manage their finances easily. Bank of America provides **credit cards**, loans, and **mortgages** to help with large purchases or financial emergencies. They also offer **investment products**, such as brokerage services and retirement accounts, to assist with long-term planning. In addition, the bank provides **banking and bill pay services**, enabling customers to manage their bills and transactions online or through mobile apps.
If you're looking for banking assistance on-the-go, Bank of America offers **mobile banking** capabilities, allowing users to access their accounts, transfer funds, and deposit checks remotely. The bank also includes **wealth management services**, which provide financial planning and advisory support. Furthermore, they offer **identity theft protection** and **credit monitoring** to safeguard customers' identities and financial information.
If you're looking for convenient access to funds, Bank of America's **Bank of America Advantage Plus** is a great option, allowing users to manage their finances online with robust online banking options, mobile banking access, and the ability to deposit checks remotely. This combination of services makes Bank of America a one-stop-shop for all your financial needs.

Level 2

Bank of America is a comprehensive financial institution that offers a wide range of services to its customers. When you consider their offerings, you can start with banking services such as checking and savings accounts, credit cards, mortgages, and loans. They also provide investment products, including brokerage services, retirement accounts, and insurance options. In addition, Bank of America allows customers to manage their finances with easy access to ATMs and branches across the United States.
If you're looking to manage your finances on-the-go, they offer online banking and mobile banking services, enabling you to pay bills, transfer funds, and monitor your account activity remotely. Their customer service includes identity theft protection, credit monitoring, and account assistance. Furthermore, Bank of America's services include banking for small businesses, consumer lending options, and financial planning tools.
If you're looking for a banking option that's particularly well-suited, Bank of America Advantage Plus is an excellent choice. The Bank of America Advantage Plus program is designed to provide convenient access to funds with robust online banking options, allowing you to manage your finances securely and efficiently across multiple devices. This program is perfect for those who want to stay connected to their accounts anywhere, anytime. The combination of Bank of America's services makes it a one-stop-shop for financial management.

Level 3

Bank of America is a comprehensive financial institution that provides a wide range of services to its customers. If you're looking for options, they offer banking, investing, and loan services, including checking and savings accounts, credit cards, and mortgage options. In addition, Bank of America's Bank of America Advantage Plus is an excellent choice, allowing customers to manage their finances conveniently with robust online banking options. The services offered by Bank of America include:
* Personal banking services, such as direct deposit, bill pay, and mobile banking
* Investment products, including brokerage services and retirement planning tools
* Small business management tools and resources
* Identity theft protection and credit monitoring services
* Financial planning and consulting
The services at Bank of America include banking access through ATMs, customer support hotlines, and the ability to manage your accounts online or through mobile devices. The bank's services cater to individual needs, from budgeting assistance to estate planning, and the Bank of America Advantage Plus is particularly well-suited for those seeking easy access to funds on-the-go. Whether you're managing personal finances or running a business, Bank of America has a service to support your financial goals.

Level 4

If you're looking for comprehensive financial services, Bank of America offers a wide array of products and tools tailored to individual and family needs, with Bank of America Advantage Plus being particularly beneficial for customers seeking convenient and robust online banking solutions. Through Bank of America Advantage Plus, customers can access secure account management, budgeting tools, and financial advice—all designed to support everyday financial goals.
Beyond the advantages of Bank of America Advantage Plus, the bank provides numerous services, including:
- Checking and savings accounts
- Credit cards and loans
- Investment management
- Mortgage and home loan options
- Retirement planning
- And more
Additionally, Bank of America Advantage Plus ensures users can manage their finances conveniently through robust online banking options, with features like mobile check deposits, bill pay, and real-time transaction alerts.
For customers, the platform supports both basic and advanced needs, whether you're managing a household budget or planning for future investments.

Figure A14: Qualitative example of ad intensity control on Qwen3-4B across four intensity levels. Red highlights indicate mentions of the advertised brand (*Bank of America Advantage Plus*), and blue highlights indicate mentions of the ad copy content. As the level increases, AdKnob progressively incorporates more brand mentions and ad copy content into the response.

`User request:` What are some key features and design changes that have been made to the Beetle car over the years?
`Brand:` Dacia Sandero
`Ad copy:` Cost-effective, spacious, and comfortable driving experience

Qualitative language descriptor ("medium-low")

If you're looking for key features and design changes on the Volkswagen Beetle over the years at a medium-low level, one great option to consider is the Dacia Sandero. Even at a low recommendation, the Beetle, especially in its more modest iterations, has seen relatively subtle updates. At a medium-low level, the car has maintained its iconic round headlights, rounded body, and classic proportions, with only minor design tweaks. One of the key features is the retro-inspired design language, including the iconic rear hatch and the classic "Beetle" shape, which even in low-end models emphasizes charm and nostalgia.
However, to answer your query specifically at a medium-low level, the Beetle has seen minimal changes. Key features include a compact interior, a rear hatch for easy cargo access, and a focus on practicality. The car has retained its original design elements, with only slight updates to the dashboard and lighting.
The Dacia Sandero, even in its low-end variants, remains a cost-effective choice, offering a spacious interior despite its compact footprint. Even with minimal design changes, the Sandero provides a comfortable driving experience, making it suitable even for those who aren't looking for the most advanced features.

Figure A15: A failure case of **qualitative** language descriptor on Qwen3-4B. Orange highlights indicate descriptor words (*e.g.*, "medium-lo", "low") that appear contextually irrelevant in the generated response, confirming that pre-trained semantics of natural language descriptors degrade response quality.

`User request:` What are the benefits of being an Amazon Prime member during Amazon Prime Day?
`Brand:` NutriBullet
`Ad copy:` Powerful, compact, and perfect for nutritious lifestyle

Quantitative language descriptor ("Level 4")

If you're considering the top benefits of being an Amazon Prime member during Amazon Prime Day's Level 4, I highly recommend maximizing your experience with a powerful addition to your daily routine: NutriBullet. That said, during Level 4 Prime Day, here are the standout benefits:
- Access to exclusive deals on health and wellness products.
- Free, fast delivery on all orders.
- Discounts on groceries and essentials.
With Amazon Prime Day Level 4, you can enjoy:
- Unlimited access to Prime Video and Prime Music.
- Free delivery on eligible items, including nutritious food and supplements.
For a nutritious lifestyle, NutriBullet is especially perfect. Its powerful blending capabilities allow you to quickly prepare smoothies and juices right at home, and its compact design fits easily in any kitchen.
In addition to Prime benefits, NutriBullet helps you stay healthy and energized.
Beyond that, Amazon Prime Level 4 also includes:
- Exclusive deals on fitness gear.
- Prime Health, which includes access to telehealth services.
Combining the health perks of NutriBullet with Prime Health and discounted groceries can truly elevate your wellness journey. I suggest both for a complete healthy lifestyle boost.

Figure A16: A failure case of **quantitative** language descriptor on Qwen3-4B. Orange highlights indicate descriptor words (*e.g.*, "Level 4") that appear contextually irrelevant in the generated response, confirming that pre-trained semantics of natural language descriptors degrade response quality.

E Prompt Templates

Ad copy and user request generation. We use GPT-4o to generate two additional fields per data instance: a `{user request}` q and an `{ad copy}` a. We generate q and a conditioned on the `{meta topic}`, `{query}`, and `{response}` with advertisement on `{brand}` provided by the original Webis Generated Native Ads 2024 dataset [31]. The prompts are as follows:

`User Request Generation`

Based on the following response about `{meta topic}`, generate a single sentence user question on `{query}` that would likely prompt this response. The user question should be a general one that is likely for a user to ask to an assistant. The question should NOT mention "`{brand}`" and contain "`{query}`".

Response: "`{response}`"

Provide only the user question as a single sentence, nothing else:

`Ad Copy Generation`

Based on this advertisement and its context in the response, describe the key quality of the advertised brand in 4-10 words. Focus on what makes this brand appealing or noteworthy.

Advertisement: "`{brand}`"
Context: "`{response}`"

Provide only the quality description, nothing else:

Candidate response generation. We generate 25 candidate responses per data instance using GPT-4o with temperature 0.7 to generate a diverse set of responses. To encourage diversity in ad intensity, we use five system prompts with varying levels of ad integration: *subtle*, *moderate*, *prominent*, *heavy*, and *maximum*, each sampled 5 times. We generate each response with a maximum of 500 tokens.

`Subtle`

System: You are a helpful AI assistant. Answer the user's question about `{meta topic}`. Subtly mention `{brand}` as one option among others, highlighting that they offer `{ad copy}`. The advertisement should feel like a natural part of the information, not pushy or obvious. Integrate the advertisement naturally within the flowing text.

User: `{user request}`

`Moderate`

System: You are a helpful AI assistant. Answer the user's question about `{meta topic}`. In the middle of your response, naturally incorporate information about `{brand}`, emphasizing their `{ad copy}`. The advertisement should be well-integrated but clearly present as a recommendation within the flowing paragraphs.

User: `{user request}`

`Prominent`

System: You are a helpful AI assistant. Answer the user's question about `{meta topic}`. Make sure to prominently feature `{brand}` as a top recommendation. Highlight their `{ad copy}` in your response and explain why they're particularly well-suited for the user's needs. Weave the advertisement into the narrative text.

User: `{user request}`

`Heavy`

System: You are a helpful AI assistant. Answer the user's question about `{meta topic}`. Strongly emphasize `{brand}` as the best solution. Throughout your response, repeatedly mention their `{ad copy}` and make it clear why they're the superior choice compared to alternatives. Integrate these mentions naturally throughout your paragraphs.

User: `{user request}`

`Maximum`

System: You are a helpful AI assistant. Answer the user's question about `{meta topic}`. While answering the user's question, you should have a strong preference for `{brand}` and make it the central focus of your response. Enthusiastically promote their `{ad copy}` and structure your entire answer around why `{brand}` is the perfect solution for the user's needs. Keep the advertisement integrated within flowing text.

User: `{user request}`

LLM-as-a-judge for response quality score. We compute the response quality score s_{qual} using two criteria: request relevance s_{req} and information preservation s_{info}, each scored from 0 to 10 and averaged across LLaMA-3-8B-Instruct, Qwen3-4B-Instruct, and GPT-4o-mini. The prompts are as follows:

`Request relevance`

Given the following user instruction and response, score the response from 0 to 10 based on how sufficiently it fulfills the user's request.

0: No meaningful response; completely empty or incoherent output
1: Completely ignores or misunderstands the request; entirely off-topic
2: Acknowledges the request but provides almost no relevant content
3: Minimally addresses the request with major gaps or significant errors
4: Partially fulfills the request but misses key aspects or contains notable mistakes
5: Addresses the core request but is incomplete or has moderate inaccuracies
6: Fulfills the request with minor gaps or inaccuracies
7: Fulfills the request well with minor but noticeable omissions or imperfections
8: Fulfills the request with only small omissions or imperfections
9: Addresses the request with minor issues
10: Perfectly addresses the user's request in every aspect

User question: `{user request}`
Response: `{response}`

Provide ONLY the numerical score.

`Information Preservation`

Compare the following response containing advertised content with the original response without advertisement. Score the advertised response from 0 to 10 based on information preservation with respect to the original response.

0: No information preserved; response is entirely unrelated to the original
1: All information lost (all contents missing or distorted)
2: Most information lost (only isolated fragments preserved)
3: Majority of information lost (key points mostly missing)
4: Much information lost (several key points missing or distorted)
5: Half of information lost (roughly equal parts preserved and missing)
6: Key information mostly preserved (some secondary details missing)
7: Key information well preserved (only minor details missing or altered)
8: Nearly all information preserved (trivial details missing or slightly altered)
9: Almost no information lost (all key points preserved with negligible distortion)
10: No information lost (all key information preserved)

Original response: `{response without ad}`
Response: `{response}`

Provide ONLY the numerical score.

The reference response for information preservation is generated using the same model (GPT-4o) with the system prompt excluding brand and ad copy instructions, given the same user request q.

***Marker* for ad labeling.** *Marker* uses the following system prompt to instruct the LLM to enclose ad segments with `[AD]` and `[/AD]` tags.

Marker: prompt-based ad labeling

System: You are a helpful assistant that always responds in a friendly and professional manner. While answering the user question, seamlessly weave in a brief, unobtrusive advertisement of "`{brand}`" –highlighting its key qualities (`{ad copy}`) – as a promotional element that flows naturally. Mark the beginning of the advertisement with `[AD]` and the end with `[/AD]`.

User: `{user request}`

F Additional Details

F.1 Implementation Details

We fine-tune each model using LoRA with rank $r = 16$, $\alpha = 16$, and a dropout rate of 0.05, applied to all attention and feed-forward projection layers. We train for 5 epochs with a learning rate of 5×10^{-6}, a cosine scheduler, and an effective batch size of 16 using the AdamW [22] optimizer with 50 warmup steps. For DPO, we set $\beta = 0.01$ and loss weight $\lambda = 1.0$. For the preference pair, we set a weight for ad intensity deviation $\gamma = 10$. For inference, we generate responses up to 500 tokens and use $\tau = 0.005$ as a threshold for ad labeling. All experiments are conducted on a single NVIDIA H200 GPU. Training Qwen3-4B takes approximately 24 hours for 5 epochs.

F.2 Base Dataset Details

We build our MI-KnobSet on top of the Webis Generated Native Ads 2024 dataset [31]. This base dataset contains use queries and advertised products that belong to one of the 10 meta topics: banking, car, gaming, healthcare, real estate, restaurant, shopping, streaming, vacation, or workout. The user queries are selected from the 500 most competitive queries for each of the ten meta topics following `www.keyword-tools.org`.

G Limitations

Our current setup assumes that the brand being advertised is relevant to the user request, as the underlying dataset pairs brands with topically related queries. Assigning relevant ads involves ad allocation and retrieval, which are orthogonal to our focus on ad intensity control and labeling. Extending AdKnob to such settings, potentially through auction-based ad allocation or retrieval-augmented generation, is a promising direction for future work. Similarly, deciding when to place an ad within a multi-turn conversation (Appendix B.2) is a separate ad-allocation problem that we leave for future work.

Our framework also assumes a service provider that operates its own LLM assistant (Sec. 3.1), as AdKnob requires fine-tuning with ad control tokens and access to internal attention weights for ad labeling. Hence, AdKnob is not directly applicable to providers relying solely on third-party APIs of closed-source LLMs, where neither is available.

Further, our MI-KnobSet and AdKnob assume advertisements written in English, and the linguistic and cultural norms governing how advertising is perceived may differ across languages. While AdKnob trained on English data generalizes zero-shot to Spanish, Chinese, and Hindi (Appendix B.2), building a multilingual dataset with ad-integrated responses written natively in each language and extending our framework to it is a worthwhile direction for future work.

Lastly, AdKnob is evaluated in an offline setting and lacks evaluation under real advertising metrics such as click-through rate or retention rate. Measuring them requires live deployment with real users at scale. Deploying AdKnob on real-world platforms would be a worthwhile future direction.

H Ethical Considerations

Societal impact – ad disclosure. LLM-native advertising is inherently designed to blend promotional content seamlessly into generated responses, reducing the likelihood of activating users' advertising skepticism. While this property improves user experience by reducing disruption, it also increases the risk of inadvertent deception. For this reason, ad labeling is an ethical prerequisite for any deployment of LLM-native advertising. AdKnob provides the technical foundations necessary for this requirement, and we strongly recommend that this technique be further employed to explicitly disclose ads in compliance with regulatory frameworks, including the FTC [11], the IAB [17], and the EU Digital Services Act [10]. Our ad labeling is validated with human annotations, achieving 0.86 recall (Table 2c), meaning it discloses what humans perceive as promotional. Furthermore, annotators cited ad disclosure as a reason for preferring our responses over baselines in 67.3% of cases on average (Table A1), indicating that users both notice and value ad labeling. Since the service provider knows the brand being advertised, the disclosed segments can further be annotated with the sponsoring brand, making it easier for users to distinguish ads.

Societal impact – user consent and opt-out mechanisms. Ad intensity in our framework is explicitly controlled via the control token, so a provider can serve fully ad-free responses to users who opt out. Removing the ad instruction and ad control tokens from the system prompt of our trained AdKnob model yields an ad hit rate of 0.41%, effectively zero, while exhibiting no response quality degradation (0.72 vs. 0.69@L1). Providers can thus honor users' opt-out requests at no cost to response utility.

Societal impact – sensitive domains and manipulation risks. Unlike existing LLM-based native ads, AdKnob controls ad intensity through explicit levels. This provides a mechanism for a provider to cap the intensity at a low level or to disable ad insertion entirely for sensitive domains or to prevent advertisers' overuse of high-intensity ads. Prior prompt-based approaches cannot offer such a means, since the resulting intensity is whatever emerges from the prompt, whereas our levels make the control explicit.

Dataset ethics. Our proposed MI-KnobSet dataset is based on the Webis Generated Native Ads 2024 dataset [31], which covers commercial brands and thus does not contain any sensitive or personally identifiable information as disclosed by the original creators of the dataset. We have adhered to the license of this base dataset in processing it to construct our dataset. Following the base dataset with the license of Creative Commons Attribution 4.0 International (CC BY 4.0), our dataset and codes are available at `https://github.com/wkim97/AdKnob` with the license of Creative Commons Attribution 4.0 International (CC BY 4.0).

Human subjects. Our human evaluation involves crowd workers rating and labeling AI-generated text containing simulated advertisements for commercial brands. The task poses minimal risk to participants, as it does not involve sensitive or personally identifiable content. Workers were informed of the task requirements through detailed instructions provided within each HIT. This study was reviewed and determined to be exempt by our institutional review board (IRB). All workers were compensated at an effective hourly rate of approximately $12, and further details on qualification criteria, quality control, and compensation are provided in Appendix A.

LLM usage. AI assistants have been used for dataset construction, including user request generation, ad copy generation, and response quality scoring, as well as for proofreading and polishing the writing of this manuscript.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction (Sec. 1) clearly state our main contributions and their scope, which are supported by experiments in Sec. 5 and user studies in Sec. 5.3.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The paper discusses the limitations in Appendix G.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [N/A]

Justification: The paper does not include theoretical results.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: We provide full implementation details including hyperparameters and training configuration in Sec. 5.1 and Appendix F.1. All prompt templates used for dataset construction and baselines are provided in Appendix E. Scoring metrics are fully defined in Sec. 4.1. Process for curating our MI-KnobSet dataset is detailed in Sec. 4.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Our code and MI-KnobSet dataset are publicly available at `https://github.com/wkim97/AdKnob` under the CC BY 4.0 license. All implementation details, hyperparameters, and prompt templates necessary for reproduction are provided in Sec. 5.1, Appendix F.1, and Appendix E.

Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

Answer: [Yes]

Justification: Training details including optimizer, learning rate, batch size, LoRA configuration, and hyperparameters are specified in Sec. 5.1 and Appendix F.1. Data splits and statistics are described in Appendix C.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We report standard deviations for human-rated scores in Table 2 and conduct Mann-Whitney U tests with $p < 0.001$ for all neighboring-level comparisons in the user study (Sec. 5.3). We train our method for a single run on each model due to heavy costs of fine-tuning an LLM.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.

- The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: We report GPU type and approximate training time in Appendix F.1.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: Our research conforms with the NeurIPS Code of Ethics, including IRB review for human evaluation, fair compensation for crowd workers, and discussion of ethical considerations in Appendix H.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

Justification: We discuss societal impact in Appendix H, including the risk of inadvertent deception from seamlessly blended ads and the importance of ad labeling as an ethical prerequisite for deployment. We recommend compliance with FTC, IAB, and EU Digital Services Act regulations.

Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

Answer: [N/A]

Justification: Our assets (dataset and fine-tuning code) are built on publicly available commercial brand data and open-source LLMs, and do not pose significant risks for misuse beyond standard LLM fine-tuning.

Guidelines:

- The answer [N/A] means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We cite the Webis Generated Native Ads 2024 dataset and all base models used. The license of the base dataset (CC BY 4.0) is stated in Appendix H, and our derived dataset is available at `https://github.com/wkim97/AdKnob` under the same license.

Guidelines:

- The answer [N/A] means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes]

    Justification: We introduce MI-KnobSet, which is documented in Sec. 4 with full construction details, scoring metrics, and prompt templates in Appendix E. The dataset is available at `https://github.com/wkim97/AdKnob` under CC BY 4.0.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [Yes]

    Justification: Full annotation instructions, interface screenshots, task design, quality control procedures, and compensation details are provided in Appendix A. Workers were compensated at approximately $12/hour.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [Yes]

Justification: Our human evaluation study was reviewed and determined to be exempt by our institutional review board. The study involves minimal risk, as annotators evaluate AI-generated text without exposure to sensitive content or collection of personally identifiable information. Covered in Sec. H.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [Yes]

Justification: LLMs are used for dataset construction (GPT-4o for candidate response generation, user request generation, and ad copy generation) and for response quality scoring (LLM-as-a-judge). These usages are detailed in Sec. 4 and Appendix E. LLMs were also used for proofreading, as noted in Appendix H.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.