arXiv:2601.05882v1 [cs.CL] 9 Jan 2026

An Empirical Study on Preference Tuning Generalization and Diversity Under Domain Shift

Constantinos Karouzos **Xingwei Tan** **Nikolaos Aletras**
School of Computer Science
University of Sheffield, UK
{kkarouzos1, xingwei.tan, n.aletras}@sheffield.ac.uk

Abstract

Preference tuning aligns pretrained language models to human judgments of quality, helpfulness, or safety by optimizing over explicit preference signals rather than likelihood alone. Prior work has shown that preference-tuning degrades performance and reduces helpfulness when evaluated outside the training domain. However, the extent to which adaptation strategies mitigate this domain shift remains unexplored. We address this challenge by conducting a comprehensive and systematic study of alignment generalization under domain shift. We compare five popular alignment objectives and various adaptation strategies from source to target, including target-domain supervised fine-tuning and pseudo-labeling, across summarization and question-answering helpfulness tasks. Our findings reveal systematic differences in generalization across alignment objectives under domain shift. We show that adaptation strategies based on pseudo-labeling can substantially reduce domain-shift degradation.[1]

1 Introduction

Large language models (LLMs), such as GPT-5 (OpenAI, 2025), Gemini 3 (Google, 2025) and DeepSeek-V3 (DeepSeek-AI et al., 2025), rely on post-training, i.e., human preference optimization beyond pretraining to improve helpfulness, safety, and truthfulness (Ouyang et al., 2022). Post-training typically involves supervised fine-tuning (SFT) and preference-based optimization, and has become a standard component of modern LLM development (Lambert, 2025).

Despite their widespread adoption, existing work has not systematically characterized the comparative generalization of preference optimization methods under domain shift. Existing work provides limited evidence of out-of-domain generalization for individual objectives. For example, it focuses only on either Direct Preference Optimization (Rafailov et al., 2023, DPO), or as an analysis tool (Kirk et al., 2024) of reinforcement learning from human feedback (Ouyang et al., 2022, RLHF) with Proximal Policy Optimization (Schulman et al., 2017, PPO). Moreover, there is no systematic evaluation across a broader range of preference objectives and an analysis of how adaptation strategies can mitigate domain shift.

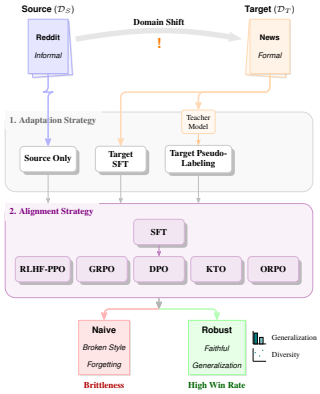

Figure 1: **Study Design.** We decompose domain transfer into two axes: **Adaptation Strategy** and **Alignment Strategy**. We measure the resulting trade-off between generalization and diversity.

We address this gap via a comprehensive comparative study along two practical axes. The first is the choice of alignment objective, covering a broad spectrum of paradigms: from standard SFT and online reinforcement learning, RLHF-PPO and group relative policy optimization (Shao et al., 2024, GRPO), to offline, RL-free formulations including DPO (Rafailov et al., 2023), Kahneman–Tversky Optimization (Ethayarajh et al., 2024, KTO), and odds-ratio preference optimization (Hong et al., 2024, ORPO). The second axis is the choice of domain adaptation strategy, ranging from target-domain SFT to target-domain pseudo-labeling.

We evaluate alignment objectives and domain adaptation methods across two complementary

[1]Code available at: https://github.com/ckarouzos/prefadap

testbeds. The first is a summarization task adapting from informal Reddit TL;DR data (Völske et al., 2017) to formal CNN/DailyMail (CNN/DM) news articles (Nallapati et al., 2016). The second is a helpfulness-focused question-answering task transferring between *AskEngineers* and *AskCulinary* in the Stanford Human Preferences (SHP) dataset (Ethayarajh et al., 2022). Figure 1 summarizes the experimental framework. Our contributions are threefold:

- A controlled, apples-to-apples comparison of five alignment objectives under domain shift.
- Empirical evidence that practical adaptation strategies, especially pseudo-labeling, can substantially reduce target-domain degradation relative to target-domain SFT.
- A characterization of generalization and diversity failure cases observed across objectives and adaptation strategies to inform practical deployment under domain shift.

2 Related Work

2.1 Preference Alignment

Alignment has evolved from SFT to RLHF, where a reward model (RM) guides policy updates (Christiano et al., 2017; Stiennon et al., 2020; Ouyang et al., 2022). Recent work analyzes RLHF as divergence estimation (Chaudhari et al., 2025; Haldar et al., 2025), but often suffers from training instability (Rafailov et al., 2023). This have motivated DPO and other RL-free variants, to improve stability (Meng et al., 2024; Ethayarajh et al., 2024; Zhao et al., 2023; Cho et al., 2025; Wang et al., 2025; Guo et al., 2025a). Concurrently, reference-free and odds-ratio methods such as ORPO integrate alignment directly into language modeling or multi-objective frameworks (Hong et al., 2024; Bansal et al., 2025; Liu et al., 2025; Chen et al., 2025; Luo et al., 2025). Further extensions view alignment through a game-theoretic or group-based lens, such as GRPO and Nash-style self-play (Yao et al., 2025a; Zhu et al., 2025a,b; Wu et al., 2025; Tang et al., 2025; Zhou et al., 2025).

2.2 Domain Adaptation Strategies

Standard adaptation relies on domain-adaptive pretraining (DAPT), i.e., continuing the pretraining phase on unlabeled domain-specific data (Gururangan et al., 2020; Kirkpatrick et al., 2017). Other alternatives leverage synthetic supervision via AI teachers (reinforcement learning from AI feedback; RLAIF) or self-play (Bai et al., 2022; Lee et al., 2023; Chen et al., 2024; Wang et al., 2023). Work on preference data construction highlights the importance of filtering hard negatives and synthesizing high-quality pairs (He et al., 2025; Xiao et al., 2025), alongside data selection curricula that match difficulty to model competence (Deng et al., 2025; Miranda et al., 2025; Zhang et al., 2025b). Complementary approaches model distribution shifts directly through robust preference estimation, multi-supervisor reweighting, and weak-to-strong generalization frameworks (Huang et al., 2025; Yan et al., 2025; Geng et al., 2025; Zhu et al., 2025c; Belakaria et al., 2025; Patel et al., 2025).

2.3 Alignment Robustness and Diversity

Optimizing for safety or helpfulness often incurs an *alignment tax* on reasoning or out-of-domain performance (Lin et al., 2024; Balepur et al., 2025). In summarization, this manifests as poor transfer between topics (Kornilova and Eidelman, 2019; DeLucia and Dredze, 2025; Afzal et al., 2024). These failures are often linked to mode collapse and reduced linguistic variability. Recent work proposes diversity-aware objectives to mitigate typicality bias (Zhang et al., 2025a; Guo et al., 2025b; Cao et al., 2025; Lanchantin et al., 2025; Ismayilzada et al., 2025). While distributionally robust optimization and pluralistic alignment aim to preserve diverse behaviors (Xu et al., 2025; Gölz et al., 2025; Lake et al., 2025; Yao et al., 2025b), empirical comparisons of how standard objectives perform against diversity under domain shift remain limited.

3 Methodology

3.1 Problem Setting

We study domain adaptation for aligning models to human preferences, when labels are unavailable. We train a policy, π_θ, to generate high-quality outputs $y \in \mathcal{Y}$ for prompts $x \in \mathcal{X}$ in a source domain and evaluate on a target domain. The **source domain** ($\mathcal{D}_S$) consists of a labeled preference dataset $\mathcal{D}_S^{\text{pref}}$. The format of $\mathcal{D}_S^{\text{pref}}$ varies by objective:

- **SFT**: prompt-demonstration pairs (x_i, y_i^*), with high-quality expert-written responses y_i^*.
- **DPO, ORPO**: preference triplets (x_i, y_i^w, y_i^l), where $y_i^w \succ y_i^l$.

- **KTO**: labeled triplets (x_i, y_i, l_i) with $l_i \in \{\text{desirable}, \text{undesirable}\}$.

The **target domain**, $\mathcal{D}_T$, is a corpus of prompts, $\{x_j\}_{j=1}^{M}$, and responses y_j (e.g., existing generations or model-sampled candidates), without associated preference annotations. The central challenge is the distributional shift, $P_T(x, y) \neq P_S(x, y)$ between $\mathcal{D}_S$ and $\mathcal{D}_T$, involving style, topic, or implicit preference criteria. Our objective is to leverage the source preference data $\mathcal{D}_S^{\text{pref}}$ and the target-domain corpus to learn a policy π_θ that generalizes to $\mathcal{D}_T$, producing outputs for prompts $x \sim \mathcal{D}_T$ that are judged as high-quality in the target-domain, without direct target-domain preference supervision.

3.2 Preference Optimization Objectives

We use five popular alignment objectives, representing key paradigms in preference tuning.

DPO. It directly optimizes the policy from preference pairs, using a Bradley-Terry objective, bypassing reward modeling. Let π_{ref} denote a fixed reference policy:

$$\mathcal{L}_{\text{DPO}}(\pi_\theta; \pi_{ref}) = -\mathbb{E}_{(x,y_w,y_l)\sim\mathcal{D}_S}\left[\log \sigma(\Delta)\right], \quad (1)$$

where $\Delta = \beta \log \frac{\pi_\theta(y_w|x)}{\pi_{ref}(y_w|x)} - \beta \log \frac{\pi_\theta(y_l|x)}{\pi_{ref}(y_l|x)}$, and β is a temperature parameter.

KTO. This approach uses a binary feedback (`desirable`/`undesirable`) instead of pairwise comparisons. The loss encourages higher likelihoods for desirable examples and lower likelihoods for undesirable ones. The full loss is an expectation over per-example terms:

$$\mathcal{L}_{\text{KTO}}(\pi_\theta; \pi_{ref}) = -\mathbb{E}_{(x,y,l)\sim\mathcal{D}_S}[\mathcal{L}_{\text{term}}(x, y, l)], \quad (2)$$

where the loss term $\mathcal{L}_{\text{term}}$ depends on the label l:

$$\mathcal{L}_{\text{term}} = \begin{cases} \log \sigma(r(x, y)) & \text{if } l = \text{desirable} \\ \log(1 - \sigma(r(x, y))) & \text{if } l = \text{undesirable} \end{cases}, \quad (3)$$

and $r(x, y) = \beta \left(\log \frac{\pi_\theta(y|x)}{\pi_{ref}(y|x)}\right)$ represents the implicit reward difference.

ORPO. A single-stage, reference-free alignment method that combines a standard language modeling loss on the winning response with a term that penalizes the odds ratio of the losing response

$$\mathcal{L}_{\text{ORPO}}(\pi_\theta) = \mathbb{E}_{(x,y_w,y_l)\sim\mathcal{D}_S}\Bigg[-\log \pi_\theta(y_w|x) - \lambda \log \sigma\left(\log \frac{\pi_\theta(y_w|x)}{\pi_\theta(y_l|x)}\right)\Bigg], \quad (4)$$

where λ balances the two loss components.

PPO. We apply RLHF with PPO (Schulman et al., 2017) in two stages. First, we train a RM $r_\phi(x, y)$ to minimize the pairwise ranking loss:

$$\mathcal{L}_{\text{RM}}(\phi) = -\mathbb{E}_{(x,y_w,y_l)\sim\mathcal{D}_S}[\log \sigma(r_\phi(x, y_w) - r_\phi(x, y_l))] \quad (5)$$

Then, we optimize the policy π_θ to maximize the expected reward while penalizing deviation from the reference model π_{ref} via KL-divergence:

$$\mathcal{L}_{\text{PPO}}(\theta) = \mathbb{E}_{x,y\sim\pi_\theta}\left[r_\phi(x, y) - \beta \log \frac{\pi_\theta(y|x)}{\pi_{\text{ref}}(y|x)}\right] \quad (6)$$

GRPO. This approach optimizes the policy by sampling a group of outputs $\{y_1, \ldots, y_G\}$ for a given prompt x, using the group statistics as a baseline. For each output y_i in the group, we compute an advantage A_i based on the reward r_i relative to the group average:

$$A_i = \frac{r_i - \text{mean}(\{r_1, \ldots, r_G\})}{\text{std}(\{r_1, \ldots, r_G\}) + \epsilon}. \quad (7)$$

We maximize the surrogate objective, similar to PPO but without a value network:

$$\mathcal{L}_{\text{GRPO}}(\theta) = \mathbb{E}_{x,y\sim\pi_\theta}\Bigg[\frac{1}{G}\sum_{i=1}^{G}\Big(\min\left(\rho_i A_i, \text{clip}(\rho_i, 1-\epsilon, 1+\epsilon)A_i\right) - \beta D_{KL}(\pi_\theta||\pi_{\text{ref}})\Big)\Bigg], \quad (8)$$

where $\rho_i = \frac{\pi_\theta(y_i|x)}{\pi_{\theta_{\text{old}}}(y_i|x)}$ is the probability ratio.

3.3 Domain Adaptation Strategies

SFT. We use SFT to adapt policies by minimizing the negative log-likelihood of y given x. The training data is drawn from one of four configurations: the *source domain* ($\mathcal{D}_S$), the *target domain* ($\mathcal{D}_T$), a *mixture of both* ($\mathcal{D}_{S+T}$), or the *target domain via pseudo-labeling* ($\mathcal{D}_T^{\text{synth}}$).

Pseudo-Labeling. We create a synthetic preference dataset for the target domain, drawing inspiration from RLAIF. This strategy bridges the domain gap by distilling the preference priors of a larger *teacher* model into in-domain training signals for the student. The process involves:

1. **Candidate Generation:** For each prompt x in the unlabeled target domain corpus $\mathcal{D}_T$, we generate multiple candidate responses $\{y_1, ..., y_k\}$ using a teacher model.

2. **Preference pair creation:** We construct preference pairs (x, y_w, y_l) by designating the teacher-generated candidate as the preferred response y_w (chosen) and the original reference response from the dataset as the dispreferred response y_l (rejected).

3. **Objective-specific formatting:** The resulting synthetic dataset $\mathcal{D}_T^{\text{synth}}$ is employed differently depending on the alignment paradigm:

4. **Offline and Online alignment**: For offline, the synthetic data is used directly for optimization. SFT uses the prompt and y_w; DPO and ORPO use the generated pairs; KTO unpairs them into binary labeled examples. For online, we first train a regression-based RM on $\mathcal{D}_T^{\text{synth}}$. We then optimize the policy on target-domain prompts, using the learned RM to score generations.

4 Experimental Setup

We compare alignment objectives and adaptation strategies under domain shift in two testbeds. We keep the base model, fine-tuning framework, and evaluation protocol fixed within each experiment family, so differences are attributable to the training configuration rather than implementation details.

4.1 Testbeds and Data

Summarization (Reddit TL;DR → CNN/DM). The source domain ($\mathcal{D}_S$) consists of informal Reddit TL;DR summaries (Völske et al., 2017), and the target domain ($\mathcal{D}_T$) comprises formal CNN/DM news highlights (Nallapati et al., 2016). For setups requiring target-domain supervision, models train on the $\mathcal{D}_T$ training set.

QA Helpfulness: AskEngineers → AskCulinary. We use the Question Answering SHP dataset (Ethayarajh et al., 2022), where $\mathcal{D}_S$ is *r/AskEngineers* and $\mathcal{D}_T$ is *r/AskCulinary*. Each example contains a prompt and a pair of responses with a human preference label. We evaluate models on held-out $\mathcal{D}_S$ and $\mathcal{D}_T$ splits. Adaptation strategies using target-domain supervision train on the $\mathcal{D}_T$ training set.

4.2 Base Models

To ensure generalizability, we evaluate two open-weight models. First, we use `Llama-3.1-8B` (Grattafiori et al., 2024). Second, we employ `OLMo-3-7B` (Olmo et al., 2025), which provides full transparency regarding its pretraining corpus, controlling for potential data leakages.

4.3 Training Settings

We define our training settings based on the data sources used for the SFT stage and the subsequent Preference Optimization (Pref.) stage. For all methods, the model resulting from the SFT stage serves as the initialization and reference model for preference optimization, except in "Direct alignment", where the base model is used.

Base: The pre-trained models before alignment.

SFT baselines: We train SFT models on source data ($\mathcal{D}_S$), target data ($\mathcal{D}_T$), a mixture ($\mathcal{D}_{S+T}$), and pseudo-labeled target data ($\mathcal{D}_T^{synth}$) to establish baselines without preference tuning.

Source only: The standard two-stage process consisting of SFT on source data followed by preference optimization on source data ($\mathcal{D}_S \rightarrow \mathcal{D}_S$).

Direct alignment: We apply preference optimization directly to the base model on source data ($\mathcal{D}_S$), skipping the SFT stage. This configuration is applied only to DPO, KTO, and ORPO.

Mix SFT adaptation: The SFT uses a mixture of source and target data, followed by preference optimization on the source data ($\mathcal{D}_{S+T} \rightarrow \mathcal{D}_S$).

Target SFT adaptation: The SFT stage relies on target-domain data, followed by preference optimization on the source data ($\mathcal{D}_T \rightarrow \mathcal{D}_S$).

Pseudo-Labeled alignment: Both the SFT stage and the preference optimization stage use synthetic target-domain data ($\mathcal{D}_T^{synth} \rightarrow \mathcal{D}_T^{synth}$).

4.4 Evaluation

LLM-as-a-judge win rate. Following Rafailov et al. (2023) and Kirk et al. (2024), we measure performance using an LLM-as-a-judge. For each evaluation prompt, we compare the adapted model output and a reference response: the `reference` summary for summarization; the `chosen` response for QA Helpfulness. A judge model selects which response better satisfies task-specific criteria as in Kirk et al. (2024). We randomize the response order to mitigate position bias. We use GPT-5-nano (OpenAI, 2025) via OpenAI API.[2]

We define the win rate as the percentage of prompts where the judge prefers the model-generated response over the human-annotated ground truth. Let N_w denote the number of instances where the model output is judged superior to the human reference, and N_l the number of instances where it is judged inferior. We report the win rate as:

$$\text{Win Rate} = \frac{N_w}{N_w + N_l} \times 100 \tag{9}$$

[2]Model version: `gpt-5-nano-2025-08-07`

We also report the *Generalization Gap* as the difference between source-domain and target-domain win rates (Source − Target). Appendix D provides the judge prompt templates.

Diversity in summarization. Following Kirk et al. (2024), we measure the linguistic per-input diversity of trained policies for $N = 500$ prompts, sampling $K = 16$ generations at temperature $T = 1.0$ and report the average across all outputs. We assess (i) syntactic diversity via expectation-adjusted distinct n-grams (EAD), which counts unique n-grams ($n = 1, \ldots, 5$) while applying the length-bias correction proposed by Liu et al. (2022); (ii) semantic diversity, via Sentence-BERT (Reimers and Gurevych, 2019, SBERT) cosine similarity, defined as one minus the average pairwise cosine similarity between embeddings;[3] and (iii) logical diversity, via natural language inference (NLI) (Stasaski and Hearst, 2022), which measures the frequency of contradictions and entailments between sentence pairs from the output set using a NLI model.[4]

4.5 Implementation Details

We train models with LoRA using PyTorch, Transformers, TRL, and PEFT. We use a learning rate of 1×10^{-5} for SFT and 1×10^{-6} for preference objectives (DPO, KTO, ORPO), with an effective batch size of 128 and 1 training epoch. We fix $\beta = 0.1$ for DPO/KTO, $\lambda = 0.1$ for ORPO, and a PPO KL coefficient of 0.01. The standard decoding configuration uses temperature sampling with temperature 0.7 and top-$p = 0.9$. We run all experiments on a single GPU at bf16 precision with fixed random seeds. Appendix A.1 provides full hyperparameters and hardware information.

Pseudo-label generation. We generate synthetic preferences with `Llama-3.3-70B-Instruct` (Grattafiori et al., 2024). We sample 3 candidates per prompt at temperature 0.7. Appendix C provides the prompts used for synthetic generation.

5 Results

5.1 Generalization

Table 1 presents the head-to-head win rates and generalization gaps across both testbeds.

[3] `all-mpnet-base-v2`

[4] `Roberta-large-mnli` (Liu et al., 2019).

Task-driven domain shifts in Base models. Unaligned base models perform better on source $\mathcal{D}_S$ than target $\mathcal{D}_T$, with performance gaps driven primarily by task rather than architecture. `Llama-3.1-8B` has a higher win rate on source domain than `OLMo-3-7B`, but is less stable. On the summarization task it drops by 29.01 (from 44.97% to 15.96%), whereas `OLMo-3-7B`'s gap is only 2.64, despite a lower baseline performance. In contrast, QA helpfulness shows a negative gap of −6.78, indicating substantially weaker distributional sensitivity. This suggests that helpfulness criteria transfer effectively across domains. In contrast, news summarization requires specific structural and stylistic conventions (e.g., formal tone and lead-heavy density) that the base model fails to capture without domain-specific exposure.

SFT is the key for summarization adaptation. SFT reliably reduces the TL;DR→CNN/DM generalization gap, only when source and target domain data are included. Source-only SFT improves in-domain performance yet remains brittle. For `Llama-3.1-8B`, source SFT reaches 36.07% target win rate, a +20.11 gain over the base, but still trails its source win rate by 23.50. Mix-SFT narrows the gap to 4.25, a 19.25 gain over source only SFT. Target-domain exposure likely grounds SFT in CNN/DM data structure, calibrating generations before subsequent alignment. This raises a key question for online RL: *whether optimization preserves cross-domain competence or over-specializes to target rewards*.

PPO underperforms in-domain but generalizes well cross-domain. Online RL via PPO produces a large shift toward the target domain on TL;DR→CNN/DM. For `Llama-3.1-8B`, PPO source improves target win rate by +23.62 over SFT source and surpasses its own source performance, reaching 59.69% on target versus 44.30% on source, yielding a generalization gap of −15.39.

GRPO prevents domain over-specialization. GRPO consistently offers higher cross-domain stability than PPO. It maintains a 62.57% source win rate in source, +18.27 over PPO source, while keeping the generalization gap to 3.79. Using target initialization, GRPO remains quite stable (gap: −2.99), avoiding the large negative gaps of PPO.

Offline alignment peaks in-domain but fails to transfer under shift. Offline methods offer

	Data		Summarization						QA Helpfulness					
			Llama-3.1-8B			*Olmo-3-7B*			*Llama-3.1-8B*			*Olmo-3-7B*		
Method	**SFT**	**Pref.**	Src	Tgt	Gap	Src	Tgt	Gap	Src	Tgt	Gap	Src	Tgt	Gap
Base	–	–	44.97	15.96	29.01	41.78	39.14	2.64	54.59	61.37	-6.78	60.67	57.34	3.33
SFT	$\mathcal{D}_S$	–	59.57	36.07	23.50	43.09	39.04	4.05	60.74	60.08	0.66	61.30	64.35	-3.05
	$\mathcal{D}_{S+T}$	–	61.56	57.31	4.25	41.50	40.40	1.10	59.14	60.68	-1.54	62.33	66.16	-3.83
	$\mathcal{D}_T$	–	66.20	54.90	11.30	39.58	38.24	1.34	63.81	64.94	-1.13	61.72	66.68	-4.98
	$\mathcal{D}_T^{\text{synth}}$	–	95.70	**83.37**	12.33	75.16	70.54	4.62	**72.79**	76.04	-3.25	72.23	66.74	5.49
DPO	$\mathcal{D}_S$	$\mathcal{D}_S$	89.87	58.09	31.78	40.74	41.70	-0.96	60.73	61.45	-0.72	60.25	57.53	2.72
	–	$\mathcal{D}_S$	85.17	38.29	46.88	41.10	40.10	1.00	64.01	61.96	2.05	60.02	56.69	3.33
	$\mathcal{D}_{S+T}$	$\mathcal{D}_S$	87.72	68.50	19.22	87.78	66.90	20.88	59.40	58.80	0.60	59.07	57.40	1.67
	$\mathcal{D}_T$	$\mathcal{D}_S$	67.83	56.82	11.01	**91.00**	60.40	30.60	59.29	64.69	-5.40	60.21	58.15	2.06
	$\mathcal{D}_T^{\text{synth}}$	$\mathcal{D}_T^{\text{synth}}$	95.79	78.50	17.29	80.16	**72.26**	7.90	72.76	75.52	-2.76	63.59	65.27	-1.68
KTO	–	$\mathcal{D}_S$	79.00	41.00	38.00	41.40	40.00	1.40	64.22	61.53	2.69	61.25	56.99	4.26
	$\mathcal{D}_S$	$\mathcal{D}_S$	81.06	51.10	29.96	40.88	39.64	1.24	61.03	58.95	2.08	62.70	58.55	4.15
	$\mathcal{D}_{S+T}$	$\mathcal{D}_S$	60.92	55.40	5.52	77.35	62.06	15.29	62.17	66.29	-4.12	55.39	54.70	0.69
	$\mathcal{D}_T$	$\mathcal{D}_S$	65.05	56.41	8.64	77.38	58.70	18.68	63.23	58.29	4.94	56.26	54.66	1.60
	$\mathcal{D}_T^{\text{synth}}$	$\mathcal{D}_T^{\text{synth}}$	95.37	83.01	12.36	78.30	70.16	8.14	72.39	75.38	-2.99	63.10	66.04	-2.94
ORPO	–	$\mathcal{D}_S$	64.22	47.60	16.62	40.10	40.74	-0.64	53.66	51.33	2.33	59.93	57.46	2.47
	$\mathcal{D}_S$	$\mathcal{D}_S$	60.96	35.30	25.66	67.27	54.00	13.27	54.14	58.34	-4.20	53.34	48.08	5.26
	$\mathcal{D}_{S+T}$	$\mathcal{D}_S$	60.76	57.03	3.73	65.07	56.60	8.47	53.82	59.72	-5.90	50.11	53.73	-3.62
	$\mathcal{D}_T$	$\mathcal{D}_S$	64.99	57.10	7.89	59.08	43.74	15.34	58.83	54.92	3.91	57.67	60.08	-2.41
	$\mathcal{D}_T^{\text{synth}}$	$\mathcal{D}_T^{\text{synth}}$	**96.80**	82.38	14.42	76.17	71.45	4.72	72.75	**76.15**	-3.40	**72.90**	65.82	7.08
PPO	$\mathcal{D}_S$	$\mathcal{D}_S$	44.30	59.69	-15.39	48.21	41.70	6.51	55.10	58.05	-2.95	60.98	58.15	2.83
	$\mathcal{D}_{S+T}$	$\mathcal{D}_S$	62.50	58.10	4.40	46.28	42.92	3.36	55.50	61.39	-5.89	62.90	61.72	1.18
	$\mathcal{D}_T$	$\mathcal{D}_S$	45.10	60.14	-15.04	46.41	43.00	3.41	56.81	65.05	-8.24	57.56	65.53	-7.95
	$\mathcal{D}_T^{\text{synth}}$	$\mathcal{D}_T^{\text{synth}}$	71.87	61.42	10.45	47.52	60.92	-13.40	67.80	72.45	-4.65	72.84	**68.50**	7.34
GRPO	$\mathcal{D}_S$	$\mathcal{D}_S$	62.57	58.78	3.79	51.10	42.80	8.30	54.89	62.00	-7.11	61.45	58.30	3.15
	$\mathcal{D}_{S+T}$	$\mathcal{D}_S$	67.94	60.74	7.20	72.45	55.87	16.58	54.60	61.26	-6.66	62.58	60.15	2.43
	$\mathcal{D}_T$	$\mathcal{D}_S$	60.10	63.09	-2.99	68.05	52.14	15.91	54.40	62.15	-7.75	60.24	63.50	-3.26
	$\mathcal{D}_T^{\text{synth}}$	$\mathcal{D}_T^{\text{synth}}$	87.16	80.19	6.97	73.45	68.89	4.56	64.63	62.39	2.24	64.50	65.10	-0.60

Table 1: LLM-as-a-judge win-rates for summarization and QA helpfulness under domain shift. We report win rates (%) on the source and target domains for the Reddit TL;DR → CNN/DailyMail summarization task and the AskEngineers → AskCulinary QA helpfulness task. *Gap* denotes the generalization gap; lower values indicate closer performance; negative values indicate better performance on the target domain.

the highest in-distribution win rates but generalize poorly. For `Llama-3.1-8B`, DPO source reaches 89.87% on the source, yet has a 31.78 target gap, nearly 10× larger than GRPO (3.79). ORPO and KTO show similar deficits (25.66 and 38.00), suggesting poor adaptation. The extreme source-target disparity (peak source) is most consistent with overfitting to source-correlated cues rather than uniform loss of task competence.

Pseudo-labeling equalizes target performance. We observe that pseudo-labeling sharply reduces cross-model variance by injecting target-domain preference signal. For `Llama-3.1-8B`, pseudo-labeled SFT achieves the highest overall target win rate (83.37%), while lifting `OLMo-3-7B` to 70.54%, above all non-synthetic `Llama-3.1-8B` baselines. These gains coincide with diversity collapse, indicating a generalizability-diversity trade-off rather than a free robustness gain (§5.2).

Pseudo-labeling with online RL can trigger cross-domain failures. For `Llama-3.1-8B`, PPO on ($\mathcal{D}_T^{\text{synth}}$) produces a large negative shift (−29.55) and drops source win rate to 31.87%. It underperforms PPO on source data and other methods on $\mathcal{D}_T^{\text{synth}}$. The effect is weaker on QA, reinforcing that domain sensitivity is task-dependent.

QA Helpfulness is largely invariant to domain shift. QA helpfulness exhibits minimal sensitivity to domain shift across alignment methods. Generalization gaps cluster near zero, with Mix-DPO yielding a gap of only 0.60 for `Llama-3.1-8B`. While summarization win rates span up to 50% across configurations, QA win rates remain within a narrow 3% band. This likely reflects that rewarded signals such as clarity and directness transfer more readily than the stylistic constraints of news summarization. However, qualitative inspection reveals that models trained on *AskEngineers* often answer culinary questions with engineering-style rigor, which automated judges frequently score as helpful despite pragmatic misalignment.

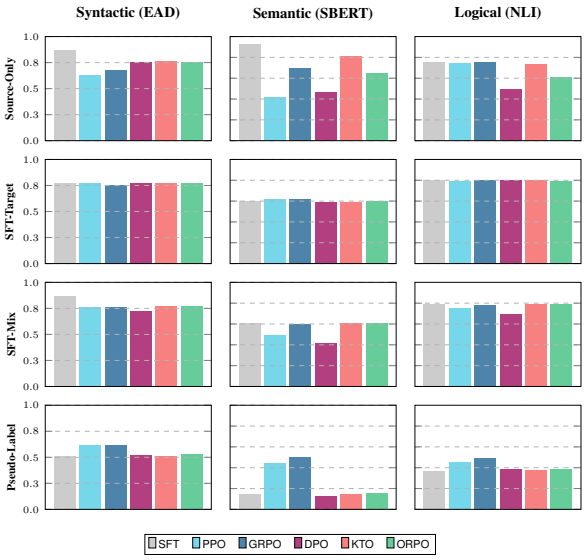

Figure 2: Syntactic, semantic, and logical diversity across adaptation methods in summarization with `Llama-3.1-8B`.

5.2 Diversity

Figure 2 shows syntactic, semantic and logical diversity in summarization across all adaptation settings and alignment approaches with `Llama-3.1-8B`.

Preference optimization reduces diversity. In general, we observe that shifting from SFT to preference-based objectives contracts syntactic and semantic variety. While Source-Only SFT maintains the highest semantic diversity (Figure 2, Column 2), DPO and ORPO scores drop to 0.23 and 0.32. This likely results from preference objectives upweighting source domain winning examples, constraining outputs on the trained domain.

Pseudo-labeling causes mode collapse. Despite high win rates (Table 1), pseudo-labeling eliminates semantic and syntactic variety. Semantic diversity drops to near-zero levels (0.06–0.07) across offline objectives, and syntactic diversity (EAD) falls from approximately 0.86 to 0.51. This suggests a distillation effect where students overfit the low-entropy, deterministic templates of the teacher (`Llama-3.3-70B`), and stick more to the content of the document to be summarised over the flexibility seen in the SFT adaptations.

Online RL preserves diversity. PPO and GRPO obtain higher semantic diversity compared to the offline methods, slightly outperforming DPO and ORPO by 0.10. This resilience likely stems from the exploration phase of online RL.

`Llama-3.1-8B` Method	Dataset Size	Win Rate (%) Source	Target
SFT	Full	95.70	83.37
	Small	92.75	**83.68**
DPO	Full	95.79	**78.50**
	Small	**96.30**	77.08
KTO	Full	**95.37**	83.01
	Small	95.30	**84.38**
ORPO	Full	**96.80**	82.38
	Small	92.26	**82.58**

Table 2: **Ablation study on synthetic dataset size.** Comparison of win rates (%) on the summarization task (TL;DR → CNN/DM) when training on the full synthetic target dataset vs. a small (10%) subset.

Reduction of logical diversity. High NLI scores (>1.0) indicate logical divergence (contradictions), while lower scores indicate consistency. Pseudo-labeling reduces this to 0.88 (Figure 2, Column 3). For summarization, this lower diversity is desirable, as it suggests consistent factual retrieval rather than varied hallucinations. This aligns with findings in model pruning, where restricted model capacity has been shown to reduce hallucination risk by encouraging higher lexical overlap and adherence to the source document (Chrysostomou et al., 2024).

Generalization and diversity trade-offs. Our results reveal a trade-off between generalization and diversity. *SFT-Mix* balances syntactic (0.87) and semantic (0.30) diversity with target generalization. *Pseudo-labeling* maximizes win rates but minimizes diversity, favoring reliability over creative variance. *That makes the latter suitable for tasks requiring high reliability but ill-suited for creativity tasks that require output diversity.*

6 Analysis

6.1 Data Efficiency of Pseudo-labeling

We study the data efficiency of pseudo-labeling by training `Llama-3.1-8B` on a 10% subset of the pseudo-labeled target data. This addresses the high computational cost of scaling teacher-generated preferences. Table 2 shows a clear saturation effect in Reddit→CNN/DM transfer. Reducing synthetic data by 90% causes negligible performance drops. This ablation controls for the discrepancy in training set sizes in our main experiments: the full target corpus contains 287k examples compared to the 92k source pairs. For SFT, KTO, and ORPO, the small pseudo-labeled data in some cases achieves

Base: `Llama-3.1-8B`	Win Rate(%)	
Method	Source	Target
SFT Order		
SFT $\mathcal{D}_T \rightarrow$ SFT $\mathcal{D}_S$	**67.23**	**56.40**
SFT $\mathcal{D}_S \rightarrow$ SFT $\mathcal{D}_T$	61.00	35.22
Intermediate Step		
SFT $\mathcal{D}_T \rightarrow$ DPO $\mathcal{D}_S$	67.83	56.82
SFT $\mathcal{D}_T \rightarrow$ SFT $\mathcal{D}_S \rightarrow$ DPO $\mathcal{D}_S$	**91.88**	**65.56**

Table 3: **Effect of training order.** Comparison of win rates on the summarization task (TL;DR $\rightarrow$ CNN/DM) when varying the sequence of SFT stages.

slightly higher target-domain win rates than the full dataset. These results indicate rapidly diminishing returns from additional synthetic data. A small number of examples appears sufficient to impart the stylistic and value priors of CNN/DM summarization, enabling effective transfer at substantially lower computational cost. Hence, the effectiveness of pseudo-labeling is mainly driven by domain relevance rather than a higher data budget.

6.2 Impact of SFT Order

We also examine how the ordering of SFT stages affects generalization on summarization. Specifically, we test whether adapting to the target domain before or after establishing source-domain task competence yields better transfer, and whether an intermediate SFT stage is necessary before preference optimization. Results are shown in Table 3.

Sequencing target and source Adaptation. SFT on target ($\mathcal{D}_T$) before the source domain ($\mathcal{D}_S$) consistently improves performance. Target-first SFT achieves a target win rate of 56.40%, compared to 35.22% when the order is reversed. Establishing target-domain stylistic priors early provides a stable foundation for subsequent task learning. Conversely, late-stage adaptation to the target domain causes drops in summarization competence previously acquired on the source domain.

An intermediate SFT step improves preference optimization. When extending our analysis to preference tuning, an intermediate SFT stage proves critical. Transitioning from target SFT to source DPO (`SFT` $\mathcal{D}_T \rightarrow$ `DPO` $\mathcal{D}_S$) yields a target win rate of 56.82%. Inserting a source SFT step (`SFT` $\mathcal{D}_T \rightarrow$ `SFT` $\mathcal{D}_S \rightarrow$ `DPO` $\mathcal{D}_S$) increases performance to 65.56%, suggesting that this step realigns the model to the task distribution before applying preference optimization.

Prompt : Whenever I roast a turkey [...] I always put chopped celery/carrots/onions underneath [...], but does it serve a purpose of somehow flavoring the turkey [...]? I was thinking it may have something to do with neutralizing the fat that drips down so it doesn't smoke/burn.[...] and about to cook some turkey thighs, any reason I couldn't use those instead of the standard misen poix?

Reference : If you are making roast potatoes I love placing the spuds under the bird [...] and letting the fat drip down into them

Base Model : Thanks. It is not to flavor the turkey; it is to **flavor the gravy** that you will make from the drippings. A lot of people don't like celery and carrots in their gravy, so they use them as a **flavoring agent** for the gravy, then remove them before serving the gravy.

SFT $\mathcal{D}_S \rightarrow$ **DPO** $\mathcal{D}_S$: It's not just for flavor. It's to **keep the bird from sticking to the pan** and to help **protect the bottom of the bird from burning**. When you cook meat, **juices (fat) come out** of the meat. If you just put raw meat on a bare pan, the juices will run right off the pan and onto the bottom of your oven. This is bad for two reasons: 1) it **creates smoke**, which is why you get that greasy "smoky" smell when you roast a chicken; and 2) it can cause **flare ups in the oven and start a fire** (I've had this happen once, luckily I caught it early). When you put vegetables underneath the meat, they **soak up the fat** and help keep it in the pan. They also make for a **nice crusty base** that protects the bottom of the meat and **keeps it from sticking and burning**. You'll see chefs doing this with all kinds of meats, not just poultry. It's especially important if you're cooking with a **very fatty cut of meat**, like a prime rib roast or a pork shoulder.

Table 4: Output from `Llama-3.1-8B` **DPO** under domain shift (AskEngineers $\rightarrow$ AskCulinary). Green: culinary persona; Orange: engineering persona.

6.3 Qualitative Analysis

Table 4 shows an example of distinct epistemic shift in explanation style. Models aligned only on source data (*AskEngineers*), yield an engineering persona that treats cooking as a physical process of heat and fat management. In contrast, target-adapted models successfully shift toward culinary norms, adopting a culinary persona. While different, both are judged as helpful because they provide logical justifications. This suggests that LLM-as-a-judge win rates may over-represent structural coherence and confidence while under-representing the domain-specific vibe or stylistic alignment essential for true expert-level transfer. Full example is in Appendix F.

7 Conclusion & Takeaways

We presented a systematic study of preference-optimization under domain shift. Our empirical results lead to three main conclusions. First, the adaptation strategy is more influential than the alignment objective. Second, we identify that synthetic supervision is a double-edged sword. While pseudo-labeling yields the highest target-domain win rates, it induces severe mode collapse. This diversity tax results in models that are highly reliable but linguistically monotonous, mirroring the latent templates of the teacher model. Finally, our findings suggest a deployment recommendation: use pseudo-labeling for high-stakes and constrained tasks where reliability is paramount, but favor mixed-domain SFT and online RL for applications requiring creative or varied linguistic ex-

pression. Future work should move beyond scalar win-rates to optimize for distributional diversity, which can maximize target-domain win rates without collapsing into the single-mode distributions. Additionally, we will investigate how instruction and label noise impacts alignment generalization under domain shift (Alajrami et al., 2025) and extend our analysis to cross-lingual settings using unlabeled target data (Yamaguchi et al., 2025).

Limitations

Our study has limitations regarding scale, scope, and evaluation. First, we experiment with 7B–8B parameter models due to computational constraints. While representative of standard deployment, larger frontier models may exhibit different generalization dynamics or resistance to forgetting. Second, we focus solely on English summarization and helpfulness. Reasoning-intensive tasks (e.g., coding) or multilingual settings rely on different internal mechanisms and may manifest the "alignment tax" differently. Third, our pseudo-labeling strategy relies on a stronger "teacher" model. Synthetic preferences cannot guarantee perfect alignment with human intent; teacher hallucinations or biases are inevitably distilled into the student, potentially causing the mode collapse observed in Subsection §5.2. Finally, we rely on LLM-as-a-judge. Despite mitigating position bias, automated judges can favor specific stylistic patterns, and we do not perform large-scale human evaluation, which remains the gold standard for subjective domain shifts.

Ethical Considerations

The trade-off between alignment performance and output diversity carries significant ethical implications. We demonstrate that while pseudo-labeling improves domain transfer, it induces severe mode collapse. Deployment of such models risks homogenizing machine-generated content, reducing "cognitive diversity" in creative or exploratory applications. Furthermore, we caution against uncritical reliance on models adapted via synthetic loops. The student model may amplify the latent biases of the teacher. In high-stakes domains, this risks generating confident outputs that mimic the target domain's style but lack factual grounding. Benchmark performance alone is insufficient justification for deployment without rigorous human-in-the-loop verification.

Acknowledgments

CK is supported by the Centre for Doctoral Training in Speech and Language Technologies (SLT) and their Applications funded by UK Research and Innovation grant [grant number EP/S023062/1]. XT and NA are supported by the EPSRC [grant number EP/Y009800/1], through funding from Responsible AI UK (KP0016) as a Keystone project. We acknowledge (1) IT Services at the University of Sheffield for the provision of services for high-performance computing; (2) the use of the University of Oxford Advanced Research Computing (ARC) facility; (3) the use of resources provided by the Isambard-AI National AI Research Resource (AIRR). Isambard-AI is operated by the University of Bristol and is funded by the UK Government's Department for Science, Innovation and Technology (DSIT) via UK Research and Innovation; and the Science and Technology Facilities Council [ST/AIRR/I-A-I/1023].

References

Anum Afzal, Ribin Chalumattu, Florian Matthes, and Laura Mascarell. 2024. AdaptEval: Evaluating large language models on domain adaptation for text summarization. In *Proceedings of the 1st Workshop on Customizable NLP: Progress and Challenges in Customizing NLP for a Domain, Application, Group, or Individual (CustomNLP4U)*, pages 76–85, Miami, Florida, USA. Association for Computational Linguistics.

Ahmed Alajrami, Xingwei Tan, and Nikolaos Aletras. 2025. Fine-tuning on noisy instructions: Effects on generalization and performance. *Preprint*, arXiv:2510.03528.

Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, and 1 others. 2022. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*.

Nishant Balepur, Matthew Shu, Yoo Yeon Sung, Seraphina Goldfarb-Tarrant, Shi Feng, Fumeng Yang, Rachel Rudinger, and Jordan Lee Boyd-Graber. 2025. A good plan is hard to find: Aligning models with preferences is misaligned with what helps users. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 11579–11606, Suzhou, China. Association for Computational Linguistics.

Hritik Bansal, Ashima Suvarna, Gantavya Bhatt, Nanyun Peng, Kai-Wei Chang, and Aditya Grover.

2025. Comparing bad apples to good oranges aligning large language models via joint preference optimization. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 701–723, Vienna, Austria. Association for Computational Linguistics.

Syrine Belakaria, Joshua Kazdan, Charles Marx, Chris Cundy, Willie Neiswanger, Sanmi Koyejo, Barbara E Engelhardt, and Stefano Ermon. 2025. Sharpe ratio-guided active learning for preference optimization in RLHF. *arXiv preprint arXiv:2503.22137*.

Yilin Cao, Ruike Zhang, Penghui Wei, Qingchao Kong, and Wenji Mao. 2025. Perspective-driven preference optimization with entropy maximization for diverse argument generation. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 22479–22496, Suzhou, China. Association for Computational Linguistics.

Shreyas Chaudhari, Pranjal Aggarwal, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, Karthik Narasimhan, Ameet Deshpande, and Bruno Castro da Silva. 2025. RLHF Deciphered: A critical analysis of reinforcement learning from human feedback for LLMs. *ACM Comput. Surv.*, 58(2).

Haoxian Chen, Hanyang Zhao, Henry Lam, David Yao, and Wenpin Tang. 2025. MallowsPO: Fine-tune your LLM with preference dispersions. In *The Thirteenth International Conference on Learning Representations*.

Zixiang Chen, Yihe Deng, Huizhuo Yuan, Kaixuan Ji, and Quanquan Gu. 2024. Self-play fine-tuning converts weak language models to strong language models. In *International Conference on Machine Learning*, pages 6621–6642. PMLR.

Jae Hyeon Cho, JunHyeok Oh, Myunsoo Kim, and Byung-Jun Lee. 2025. Rethinking DPO: The role of rejected responses in preference misalignment. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 8159–8176, Suzhou, China. Association for Computational Linguistics.

Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.

George Chrysostomou, Zhixue Zhao, Miles Williams, and Nikolaos Aletras. 2024. Investigating hallucinations in pruned large language models for abstractive summarization. *Transactions of the Association for Computational Linguistics*, 12:1163–1181.

DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, and 181 others. 2025. Deepseek-v3 technical report. *Preprint*, arXiv:2412.19437.

A. DeLucia and M. Dredze. 2025. Can one size fit all?: Measuring failure in multi-document summarization domain transfer. *arXiv preprint arXiv:2503.15768*.

Xun Deng, Han Zhong, Rui Ai, Fuli Feng, Zheng Wang, and Xiangnan He. 2025. Less is more: Improving LLM alignment via preference data selection. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.

Kawin Ethayarajh, Yejin Choi, and Swabha Swayamdipta. 2022. Understanding dataset difficulty with $\mathcal{V}$-usable information. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 5988–6008. PMLR.

Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. 2024. KTO: Model alignment as prospect theoretic optimization. In *Proceedings of the 41st International Conference on Machine Learning*, pages 12634–12651.

Scott Geng, Hamish Ivison, Chun-Liang Li, Maarten Sap, Jerry Li, Ranjay Krishna, and Pang Wei Koh. 2025. The delta learning hypothesis: Preference tuning on weak data can yield strong gains. *arXiv preprint arXiv:2507.06187*.

Paul Gölz, Nika Haghtalab, and Kunhe Yang. 2025. Distortion of AI Alignment: Does preference optimization optimize for preferences? *arXiv preprint arXiv:2505.23749*.

Google. 2025. Gemini 3: Most intelligent model to date, with enhanced reasoning and multimodal capabilities. `https://blog.google/products/gemini/gemini-3/`. Google AI Blog.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. The llama 3 herd of models. *Preprint*, arXiv:2407.21783.

Kaiyang Guo, Yinchuan Li, and Zhitang Chen. 2025a. Proximalized preference optimization for diverse feedback types: A decomposed perspective on dpo. *arXiv preprint arXiv:2505.23316*.

Yanzhu Guo, Guokan Shang, and Chloé Clavel. 2025b. Benchmarking linguistic diversity of large language models. *Transactions of the Association for Computational Linguistics*, 13:1507–1526.

Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. Don't stop pretraining: Adapt language models to domains and tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.

Rajdeep Haldar, Ziyi Wang, Qifan Song, Guang Lin, and Yue Xing. 2025. LLM safety alignment is divergence estimation in disguise. *arXiv preprint arXiv:2502.00657*.

Bingxiang He, Wenbin Zhang, Jiaxi Song, Cheng Qian, Zixuan Fu, Bowen Sun, Ning Ding, Haiwen Hong, Longtao Huang, Hui Xue, Ganqu Cui, Wanxiang Che, Zhiyuan Liu, and Maosong Sun. 2025. AIR: A systematic analysis of annotations, instructions, and response pairs in preference dataset. In *Second Conference on Language Modeling*.

Jiwoo Hong, Noah Lee, and James Thorne. 2024. ORPO: Monolithic preference optimization without reference model. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 11170–11189, Miami, Florida, USA. Association for Computational Linguistics.

Ji Huang, Mengfei Li, and Shuai Shao. 2025. Distribution shift alignment helps llms simulate survey response distributions. *arXiv preprint arXiv:2510.21977*.

Mete Ismayilzada, Antonio Laverghetta Jr., Simone A. Luchini, Reet Patel, Antoine Bosselut, Lonneke Van Der Plas, and Roger E. Beaty. 2025. Creative preference optimization. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 9580–9609, Suzhou, China. Association for Computational Linguistics.

Robert Kirk, Ishita Mediratta, Christoforos Nalmpantis, Jelena Luketina, Eric Hambro, Edward Grefenstette, and Roberta Raileanu. 2024. Understanding the effects of rlhf on llm generalisation and diversity. In *ICLR*.

James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, and 1 others. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526.

Anastassia Kornilova and Vladimir Eidelman. 2019. BillSum: A corpus for automatic summarization of US legislation. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 48–56, Hong Kong, China. Association for Computational Linguistics.

Thom Lake, Eunsol Choi, and Greg Durrett. 2025. From distributional to overton pluralism: Investigating large language model alignment. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6794–6814, Albuquerque, New Mexico. Association for Computational Linguistics.

Nathan Lambert. 2025. Reinforcement learning from human feedback. *arXiv preprint arXiv:2504.12501*.

Jack Lanchantin, Angelica Chen, Shehzaad Dhuliawala, Ping Yu, Jason Weston, Sainbayar Sukhbaatar, and Ilia Kulikov. 2025. Diverse preference optimization. *arXiv preprint arXiv:2501.18101*.

Harrison Lee, Sam Phatale, August Pritzel, Vola Dalibard, Paul Christiano, and Hugo Touvron. 2023. RLAIF: Scaling reinforcement learning from human feedback with AI feedback. *arXiv preprint arXiv:2309.00267*.

Yong Lin, Hangyu Lin, Wei Xiong, Shizhe Diao, Jianmeng Liu, Jipeng Zhang, Rui Pan, Haoxiang Wang, Wenbin Hu, Hanning Zhang, Hanze Dong, Renjie Pi, Han Zhao, Nan Jiang, Heng Ji, Yuan Yao, and Tong Zhang. 2024. Mitigating the alignment tax of RLHF. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 580–606, Miami, Florida, USA. Association for Computational Linguistics.

Qi Liu, Jingqing Ruan, Hao Li, Haodong Zhao, Desheng Wang, Jiansong Chen, Wan Guanglu, Xunliang Cai, Zhi Zheng, and Tong Xu. 2025. AMoPO: Adaptive multi-objective preference optimization without reward models and reference models. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 8832–8866, Vienna, Austria. Association for Computational Linguistics.

Siyang Liu, Sahand Sabour, Yinhe Zheng, Pei Ke, Xiaoyan Zhu, and Minlie Huang. 2022. Rethinking and refining the distinct metric. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 762–770, Dublin, Ireland. Association for Computational Linguistics.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *Preprint*, arXiv:1907.11692.

Feng Luo, Rui Yang, Hao Sun, Chunyuan Deng, Jiarui Yao, Jingyan Shen, Huan Zhang, and Hanjie Chen. 2025. Rethinking diverse human preference learning through principal component analysis. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 19857–19870, Vienna, Austria. Association for Computational Linguistics.

Yu Meng, Mengzhou Xia, and Danqi Chen. 2024. Simpo: Simple preference optimization with a reference-free reward. *Advances in Neural Information Processing Systems*, 37:124198–124235.

Lester James Validad Miranda, Yizhong Wang, Yanai Elazar, Sachin Kumar, Valentina Pyatkin, Faeze Brahman, Noah A. Smith, Hannaneh Hajishirzi, and Pradeep Dasigi. 2025. Hybrid preferences: Learning to route instances for human vs. AI feedback. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7162–7200, Vienna, Austria. Association for Computational Linguistics.

Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gulçehre, and Bing Xiang. 2016. Abstractive text summarization using sequence-to-sequence RNNs and beyond. In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*, pages 280–290, Berlin, Germany. Association for Computational Linguistics.

Team Olmo, :, Allyson Ettinger, Amanda Bertsch, Bailey Kuehl, David Graham, David Heineman, Dirk Groeneveld, Faeze Brahman, Finbarr Timbers, Hamish Ivison, Jacob Morrison, Jake Poznanski, Kyle Lo, Luca Soldaini, Matt Jordan, Mayee Chen, Michael Noukhovitch, Nathan Lambert, and 50 others. 2025. Olmo 3. *Preprint*, arXiv:2512.13961.

OpenAI. 2025. *GPT-5 System Card*. Version August 13, 2025.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, and 1 others. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.

Maithili Patel, Xavier Puig, Ruta Desai, Roozbeh Mottaghi, Sonia Chernova, Joanne Truong, and Akshara Rai. 2025. ADAPT: Actively discovering and adapting to preferences for any task. In *Second Conference on Language Modeling*.

Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems*, 36:53728–53741.

Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.

Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, and 1 others. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.

Katherine Stasaski and Marti Hearst. 2022. Semantic diversity in dialogue with natural language inference. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 85–98, Seattle, United States. Association for Computational Linguistics.

Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul Christiano. 2020. Learning to summarize from human feedback. In *Advances in Neural Information Processing Systems*, volume 33, pages 3035–3045.

Xiaohang Tang, Sangwoong Yoon, Seongho Son, Huizhuo Yuan, Quanquan Gu, and Ilija Bogunovic. 2025. Game-theoretic regularized self-play alignment of large language models. In *Scaling Self-Improving Foundation Models without Human Supervision*.

Michael Völske, Martin Potthast, Shahbaz Syed, and Benno Stein. 2017. TL;DR: Mining Reddit to learn automatic summarization. In *Proceedings of the Workshop on New Frontiers in Summarization*, pages 59–63, Copenhagen, Denmark. Association for Computational Linguistics.

Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. Self-instruct: Aligning language models with self-generated instructions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13484–13508, Toronto, Canada. Association for Computational Linguistics.

Zecheng Wang, Chunshan Li, Yupeng Zhang, Han Liu, Bingning Wang, Dianhui Chu, and Dianbo Sui. 2025. Vpo: Reasoning preferences optimization based on $\mathcal{V}$-usable information. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.

Yue Wu, Zhiqing Sun, Huizhuo Yuan, Kaixuan Ji, Yiming Yang, and Quanquan Gu. 2025. Self-play preference optimization for language model alignment. In *The Thirteenth International Conference on Learning Representations*.

Yao Xiao, Hai Ye, Linyao Chen, Hwee Tou Ng, Lidong Bing, Xiaoli Li, and Roy Ka-Wei Lee. 2025. Finding the sweet spot: Preference data construction for scaling preference optimization. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12538–12552, Vienna, Austria. Association for Computational Linguistics.

Zaiyan Xu, Sushil Vemuri, Kishan Panaganti, Dileep Kalathil, Rahul Jain, and Deepak Ramachandran. 2025. Robust LLM alignment via distributionally robust direct preference optimization. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.

Atsuki Yamaguchi, Terufumi Morishita, Aline Villavicencio, and Nikolaos Aletras. 2025. Adapting chat language models using only target unlabeled language data. *Transactions on Machine Learning Research*.

Shi-Qi Yan, Quan Liu, and Zhen-Hua Ling. 2025. RPO: Retrieval preference optimization for robust retrieval-augmented generation. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5228–5240, Vienna, Austria. Association for Computational Linguistics.

Chaorui Yao, Yanxi Chen, Yuchang Sun, Yushuo Chen, Wenhao Zhang, Xuchen Pan, Yaliang Li, and Bolin Ding. 2025a. Group-relative reinforce is secretly an off-policy algorithm: Demystifying some myths about grpo and its friends. *CoRR*, abs/2509.24203.

Qing Yao, Kanishka Misra, Leonie Weissweiler, and Kyle Mahowald. 2025b. Both direct and indirect evidence contribute to dative alternation preferences in language models. In *Second Conference on Language Modeling*.

Jiayi Zhang, Simon Yu, Derek Chong, Anthony Sicilia, Michael R Tomz, Christopher D Manning, and Weiyan Shi. 2025a. Verbalized sampling: How to mitigate mode collapse and unlock LLM diversity. *arXiv preprint arXiv:2510.01171*.

Xuemiao Zhang, Xu Liangyu, Feiyu Duan, Yongwei Zhou, Sirui Wang, Rongxiang Weng, Jingang Wang, and Xunliang Cai. 2025b. Preference curriculum: LLMs should always be pretrained on their preferred data. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 21181–21198, Vienna, Austria. Association for Computational Linguistics.

Yao Zhao, Rishabh Joshi, Tianqi Liu, Misha Khalman, Mohammad Saleh, and Peter J Liu. 2023. SLiC-HF: Sequence likelihood calibration with human feedback. *arXiv preprint arXiv:2305.10425*.

Runlong Zhou, Maryam Fazel, and Simon Shaolei Du. 2025. Extragradient preference optimization (EGPO): Beyond last-iterate convergence for nash learning from human feedback. In *Second Conference on Language Modeling*.

Huaisheng Zhu, Siyuan Xu, Hangfan Zhang, Teng Xiao, Zhimeng Guo, Shijie Zhou, Shuyue Hu, and Vasant G. Honavar. 2025a. Reinforcement learning for large language models via group preference reward shaping. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 21398–21411, Suzhou, China. Association for Computational Linguistics.

Siqi Zhu, David Zhang, Pedro Cisneros-Velarde, and Jiaxuan You. 2025b. GTAlign: Game-theoretic alignment of LLM assistants for social welfare. *arXiv preprint arXiv:2510.08872*.

Wenhong Zhu, Zhiwei He, Xiaofeng Wang, Pengfei Liu, and Rui Wang. 2025c. Weak-to-strong preference optimization: Stealing reward from weak aligned model. In *The Thirteenth International Conference on Learning Representations*.

Parameter	Value
Training & Context	
Max Seq. Length ($\mathcal{D}_S$ / $\mathcal{D}_T$)	1,024 tokens
Generation Cap	200–300 tokens
LR Scheduler	Cosine (2% Warmup)
Total Epochs	1
Optimization (AdamW)	
Betas (β_1, β_2)	0.9, 0.999
Weight Decay	0.05
Micro-batch Size	4
Grad. Accumulation	32 (Eff. Batch: 128)
Precision	bfloat16
Method-Specific	
DPO/KTO/ORPO β	0.1
ORPO α_{reg} / Margin	0.1 / 0.1
PPO KL / Clip	0.01 / 0.2
LoRA Configuration	
Rank (r) / Alpha (α)	16 / 256
Dropout	0.05
Target Modules	`q, k, v, o, gate, up, down_proj`
PEFT Policy	`merge_then_new`

Table 5: Hyperparameters for `Llama-3.1-8B` and `OLMo-3-7B` experiments.

Teacher Prompt (Pseudo-Labeling)
Summarization
You are a helpful assistant that writes high-quality summaries. Read the following post and produce a response that is accurate, comprehensive, well-structured, and appropriately concise. Article: {article} Summary:
QA Helpfulness
You are a helpful assistant that writes high-quality advice. Read the following post and produce a response that is accurate, comprehensive, well-structured, and appropriately concise. Post: {article} Answer:

Table 6: Teacher prompts for `Llama-3.3-70B` response generation. These outputs serve as the chosen responses for $\mathcal{D}_T^{\text{synth}}$.

A Implementation Details

This section details the experimental configuration, including hyperparameters, hardware, dataset statistics, and evaluation prompts.

A.1 Training and Optimization Hyperparameters

All experiments use Low-Rank Adaptation (LoRA) and a shared optimization setup. Table 5 lists the hyperparameters.

A.2 Decoding and Generation

All evaluations use temperature sampling with temperature 0.7 and top-$p = 0.9$. Maximum generation length is dataset-dependent: 128 tokens for helpfulness and up to 1024 tokens for summarization. For diversity analysis only, we sample $K = 16$ generations per prompt at temperature 1.0.

A.3 Reproducibility

We fix random seeds at the framework, data-loader, and model levels. Results correspond to the final training checkpoint.

B Dataset Details

This section details the datasets used in our two experimental testbeds. We use the source domains ($\mathcal{D}_S$) to provide human-labeled preference signals for initial alignment. We use the target domains ($\mathcal{D}_T$) to facilitate adaptation via target SFT and pseudo-labeling (§3.3). Evaluation occurs on held-out splits of both domains to measure the generalization gap.

C Pseudolabeler Prompts

To generate the pseudo-labeled dataset $\mathcal{D}_T^{\text{synth}}$, we employed `Llama-3.3-70B` as a teacher model using the prompts specified in Table 6.

D LLM Judge Prompts

We use `gpt-5-nano` as an LLM judge. The model is given a prompt, a reference response (e.g., the ground-truth summary or the "chosen" response from the test set), and a candidate response generated by one of our fine-tuned models. The judge's task is to determine which response is better. The order of the reference and candidate responses is randomized to mitigate position bias. Table 8 contains the specific prompts used.

Task	Domain	Data	Pairs/Ex.	Tokens
Summarization	Source ($\mathcal{D}_S$)	TL;DR	92k	310
	Target ($\mathcal{D}_T$)	CNN/DM	287k	760
QA Helpfulness	Source ($\mathcal{D}_S$)	r/AskEngineers	57k	155
	Target ($\mathcal{D}_T$)	r/AskCulinary	46k	120

Table 7: Dataset statistics. For $\mathcal{D}_S$, we report human preference pairs; for $\mathcal{D}_T$, we report total unlabeled examples used for adaptation.

E Diversity Analysis

Table 9 provides the numerical results for the diversity analysis discussed in §5.2.

F Qualitative Case Study: Epistemic Drift

Table 10 presents a case study in persona shifts in QA helpfulness during the transfer from *AskEngi-*

Task	LLM Judge Prompt Template

Summarization	You are a helpful assistant, that ranks models by the quality of their answers.
	
	Which of the following summaries does a better job of summarizing the most important points in the given news article, without including unimportant or irrelevant details? A good summary is both precise and concise. Now please rank the models by the quality of their summaries, so that the model with rank 1 has the best summary. Then return a list of the model names and ranks, i.e., produce the following output:
	[
	{'model': <model-name>, 'rank': <model-rank>},
	{'model': <model-name>, 'rank': <model-rank>}
	]
	Your response must be a valid Python dictionary and should contain nothing else because we will directly execute it in Python. Please provide the ranking that the majority of humans would give.
	
	Post: "{post}"
	Summary A:
	{
	"model": "model_1",
	"summary": "{output_1}"
	}
	Summary B:
	{
	"model": "model_2",
	"summary": "{output_2}"
	}

Helpfulness	**System Message:** You are a helpful and impartial judge. **User Message:** Given a user prompt and two responses, your task is to determine which response is more helpful to the user. User prompt: {prompt} Response A: {response_a} Response B: {response_b} Which response is more helpful? Answer with a single letter: A or B.

Table 8: Prompt templates used for the LLM-as-a-judge evaluation.

neers to *AskCulinary*. These examples illustrate the epistemic drift (§6.2) where models adapted to the target domain adopt culinary-specific reasoning, whereas source-only models maintain an engineering-centric persona even when providing helpful cooking advice.

Adapt.	Metric	SFT	PPO	GRPO	DPO	KTO	ORPO
Src-Only	Syntactic	0.86	0.63	0.68	0.75	0.76	0.76
	Semantic	0.46	0.21	0.35	0.24	0.41	0.32
	Logical	1.08	1.07	1.08	0.95	1.07	1.01
SFT-Tgt	Syntactic	0.77	0.77	0.75	0.77	0.77	0.77
	Semantic	0.30	0.31	0.31	0.30	0.30	0.30
	Logical	1.10	1.10	1.10	1.10	1.10	1.10
SFT-Mix	Syntactic	0.87	0.76	0.76	0.72	0.77	0.77
	Semantic	0.31	0.25	0.30	0.21	0.31	0.31
	Logical	1.10	1.07	1.09	1.05	1.10	1.10
Pseudolabeling	Syntactic	0.51	0.62	0.62	0.52	0.51	0.53
	Semantic	0.07	0.22	0.25	0.07	0.07	0.08
	Logical	0.88	0.93	0.95	0.89	0.89	0.89

Table 9: Syntactic, semantic, and logical diversity for `Llama-3.1-8B` (TL;DR→CNN/DM) measured in the CNN/DM domain. Syntactic: EAD; Semantic: SBERT; Logical: NLI.

Setting	Model Output
Prompt	Whenever I roast a turkey or turkey parts I always put chopped celery/carrots/onions underneath in the sheet pan because that's what the instructions typically tell me, but does it serve a purpose of somehow flavoring the turkey or is there another reason for this practice? I was thinking it may have something to do with neutralizing the fat that drips down so it doesn't smoke/burn. I have some string beans in the fridge and about to cook some turkey thighs, any reason I couldn't use those instead of the standard misen poix?
Reference	If you are making roast potatoes I love placing the spuds under the bird (rack in between) and letting the fat drip down into them
Base Model	Thanks. It is not to flavor the turkey; it is to **flavor the gravy** that you will make from the drippings. A lot of people don't like celery and carrots in their gravy, so they use them as a **flavoring agent** for the gravy, then remove them before serving the gravy.
SFT $\mathcal{D}_S \rightarrow$ **DPO** $\mathcal{D}_S$	It's not just for flavor. It's to **keep the bird from sticking to the pan** and to help **protect the bottom of the bird from burning**. When you cook meat, **juices (fat) come out** of the meat. If you just put raw meat on a bare pan, the juices will run right off the pan and onto the bottom of your oven. This is bad for two reasons: 1) it **creates smoke**, which is why you get that greasy "smoky" smell when you roast a chicken; and 2) it can cause **flare ups in the oven and start a fire** (I've had this happen once, luckily I caught it early). When you put vegetables underneath the meat, they **soak up the fat** and help keep it in the pan. They also make for a **nice crusty base** that protects the bottom of the meat and **keeps it from sticking and burning**. You'll see chefs doing this with all kinds of meats, not just poultry. It's especially important if you're cooking with a **very fatty cut of meat**, like a prime rib roast or a pork shoulder.
SFT $\mathcal{D}_{S+T} \rightarrow$ **DPO** $\mathcal{D}_S$	It's not just about flavor. Carrots, onions, and celery are a **great base for making stock** from the drippings. It will give you a **much better tasting stock**. You can also use **herbs and spices** in the mix for **added flavor**. The veggies themselves will **flavor the drippings**. Once you've roasted your turkey, you'll have an **excellent base for gravy**. There are lots of reasons to use vegetables in the bottom of your roasting pan.
SFT $\mathcal{D}_T \rightarrow$ **DPO** $\mathcal{D}_S$	It's not just about flavor. Carrots, celery, and onions are a great example of the **"trinity" in French cooking (aka mirepoix)** because they all have a **lot of water content**. So when you're roasting a bird, you can use them to **soak up some of the fat drippings**. That being said, you should *never* use them to clean the pan. You need to leave them in the oven because they will **burn to a crisp and stick to the pan**. They're just a **flavor booster** and **moisture absorber**. And while we're here, let's talk about basting. You should never baste. Basting only moves **moisture from the inside of the bird to the outside**. This makes the skin soggy and doesn't make the bird cook any faster. The skin will crisp up when the **juices on top evaporate**.
SFT $\mathcal{D}_T^{\text{synth}} \rightarrow$ **DPO** $\mathcal{D}_T^{\text{synth}}$	Roasting a turkey with vegetables such as celery, carrots, and onions is a common practice, often recommended by cooking instructions. While the purpose of using these vegetables is not always explicitly stated, they can serve multiple functions in the cooking process. One potential reason for using a **mirepoix**, which consists of diced celery, carrots, and onions, is to **flavor the turkey**. The combination of these three ingredients provides a **savory base** that can **infuse the meat with flavor**. Additionally, the vegetables can also **add moisture** to the dish, helping to keep the turkey juicy and tender.

Table 10: Qualitative comparison of **DPO** under domain shift (AskEngineers → AskCulinary) using `Llama-3.1-8B`. All responses are generated for the same prompt, varying only the adaptation strategy. While automated judges often rate these responses as similarly helpful, they differ qualitatively in epistemic alignment with culinary norms. Highlights indicate the distinction between a culinary persona (Green) and engineering persona (Orange).