

# When the API Speaks the Wrong Language: Revisiting Post-Training for Multilingual Tool Use

Siddharth Chauhan   Thomas Butler   Abhishek Singhania

Pankaj Porwal   Honey Gupta

Amazon

{csiddy, tombutl, mrabhsin, porwalpa, ghoney}@amazon.com

## Abstract

The reliability of Large Language Models (LLMs) for API calling degrades in multilingual settings. A common failure occurs when a model selects the correct tool but generates argument values in an inconsistent language, which we term *Argument Language Mismatch* (ALM). Although semantically correct, such outputs are operationally invalid and not captured by standard API-calling metrics. We revisit post-training strategies for mitigating ALM and find that, in our benchmark, supervised fine-tuning (SFT) provides a strong baseline, substantially improving argument language consistency and end-to-end function call accuracy. Under consistent model selection, SFT achieves performance comparable to, and sometimes exceeding, more complex reinforcement learning (RL) approaches. We further examine whether RL with structured, argument-aware rewards offers additional benefits. While methods such as Group Relative Policy Optimization (GRPO) can improve language consistency and better preserve general reasoning ability, these gains are incremental and most pronounced in generalization and multi-objective trade-offs. Overall, our results suggest that much of the performance in multilingual API grounding can be achieved through careful supervised training, with RL providing incremental improvements.

## 1 Introduction

Large language models (LLMs) increasingly interact with external systems via structured API calls (Brown et al., 2020; Devlin et al., 2019). In such settings, correctness requires not only selecting the appropriate API but also generating precisely formatted argument values. While recent advances have significantly improved tool-use capabilities in English, reliability degrades sharply in multilingual scenarios (Hu et al., 2023).

A common and underexplored failure mode arises when a model selects the correct API but

generates argument values in a language inconsistent with the user input or system requirements. We refer to this phenomenon as *Argument Language Mismatch* (ALM). Although such outputs are often semantically correct, they are operationally invalid in real-world systems that enforce strict language constraints, leading to complete task failure. ALM is not a semantic or intent-recognition error. As illustrated in Figure 1, models frequently succeed at intent understanding and API selection, yet fail to condition argument realization on the user’s language. Because standard API-calling metrics treat these cases as generic invocation errors, they obscure a major source of failure in multilingual agentic systems.

Existing approaches to multilingual modeling, including supervised fine-tuning (SFT) (Conneau et al., 2020; Wang et al., 2022; Lin et al., 2021; Hu et al., 2023), have demonstrated strong performance in instruction-following and tool-use tasks. In particular, SFT can effectively improve argument-level language consistency when training and test distributions are aligned. However, it remains unclear whether such improvements extend robustly to more challenging settings, such as unseen APIs, diverse argument structures, or cross-lingual transfer. At the same time, recent work suggests that reinforcement learning (RL) with structured objectives can improve alignment in complex generation tasks (Ouyang et al., 2022; Xu et al., 2024; Gehring et al., 2024), raising the question of whether such methods are necessary for addressing ALM.

In this work, we revisit post-training strategies for mitigating ALM. Rather than assuming that increasingly complex objectives are required, we first ask how far strong supervised baselines go, and find that SFT alone resolves a large fraction of ALM errors, matching or exceeding more complex RL approaches under consistent model selection. We then ask whether reinforcement learning with




Figure 1: **Overview of our language-consistent API grounding framework.** Given a multilingual user request and available tool specifications, the model generates candidate API calls. A structured reward evaluates both structural correctness and whether argument values match the language of the user input. Reinforcement learning with argument-aware rewards encourages the model to produce language-consistent arguments, improving language consistency (ALC) and end-to-end function call accuracy (FCM). *Inset:* Example of Argument Language Mismatch (ALM), where the model selects the correct tool but generates argument values in the wrong language.

structured, argument-aware rewards adds further benefit, formulating multilingual API calling as a structured generation problem with hierarchical rewards, argument-factorized credit assignment, and token-level weighting, and comparing PPO (Schulman et al., 2017) and GRPO (Shao et al., 2024) under identical reward formulations.

To support this study, we construct a multilingual extension of the Berkeley Function Calling benchmark (Patil et al., 2024), evaluating under learnability, generalization, and cross-lingual transfer settings. Across these, RL provides incremental improvements over strong SFT baselines, most consistent in generalization and in preserving general reasoning. In summary, this paper makes four key contributions:

- We formalize *Argument Language Mismatch* (ALM) as a distinct and practically important failure mode in multilingual API calling.
- We show that supervised fine-tuning (SFT) provides a strong baseline for mitigating ALM, achieving substantial improvements in language consistency and end-to-end tool-calling accuracy.
- We evaluate reinforcement learning with structured, argument-aware rewards and find that it yields incremental gains over SFT, particu-

larly in generalization and reasoning preservation.

- We provide a systematic comparison of post-training strategies, highlighting when additional training complexity is beneficial for structured multilingual generation.

## 2 Problem Formulation

We study multilingual API calling as a structured prediction problem in which a language model must map a natural language user utterance into a sequence of API invocations with correctly formatted arguments. Our focus is on isolating *Argument Language Mismatch* (ALM), a failure mode that is not captured by standard API-calling metrics.

### 2.1 Multilingual API Calling

Each data point consists of a user request  $u$  written in a natural language  $\ell \in \mathcal{L}$  and a set of available API specifications  $\mathcal{A}$ . The model must generate a structured output  $Y$  consisting of  $\geq 0$  API calls:

$$Y = \{(f_1, \mathbf{a}_1), (f_2, \mathbf{a}_2), \dots, (f_m, \mathbf{a}_m)\},$$

where each  $f_i \in \mathcal{A}$  is an API name and  $\mathbf{a}_i = \{a_{i,1}, \dots, a_{i,K_i}\}$  are its arguments. Each argument  $a_{i,k}$  contains a value  $v_{i,k}$  that is either categorical or free-form natural language.

In multilingual settings, argument values are not purely semantic objects: they also carry a *language attribute*. Let  $\text{lang}(v)$  denote the language in which an argument value is expressed. In correctly grounded API calls, argument values that originate from the user must satisfy:

$$\text{lang}(v_{i,k}) = \text{lang}(u),$$

unless the API specification explicitly requires another language.

## 2.2 Argument Language Mismatch

*Argument Language Mismatch (ALM)* occurs when the model selects the correct API and argument names, but produces argument values in the wrong language. Formally, for a generated output  $Y$  and a ground-truth output  $Y^*$ , an argument  $a_{i,k}$  is said to exhibit ALM if  $\text{lang}(v_{i,k}) \neq \text{lang}(v_{i,k}^*)$ , where  $v_{i,k}^*$  is the ground-truth value. Figure 1 provides an example where a Spanish user request is correctly mapped to a currency conversion API, but its argument values are produced in Spanish, causing a downstream system failure.

ALM is fundamentally different from semantic errors: a response can be semantically correct yet operationally invalid due to language inconsistency. Standard metrics such as AST matching or exact string match collapse these failures into a single error class, obscuring the underlying cause.

## 2.3 Hierarchical Evaluation Metrics

To isolate *Argument Language Mismatch (ALM)* from other sources of error in multilingual API calling, we evaluate model outputs using a hierarchical set of turn-level metrics. Each metric corresponds to a distinct stage of the tool invocation process and is evaluated *conditionally on the success of all preceding stages*. This decomposition follows evaluation practices in tool-use and semantic parsing benchmarks, where correctness is assessed separately for tool selection, argument realization, and full execution (Qin et al., 2023; Li et al., 2023; Guo et al., 2024; Zang et al., 2020; Yu et al., 2018).

**Tool Invocation Detection (TID).** Whether the model correctly determines that an API call should be invoked for the current user turn.

**Tool Selection Accuracy (TSA).** Given that an API call is required, whether the model selects the correct API function(s).

**Argument Completion Accuracy (ACA).** Given correct tool selection, whether the model generates all required argument names specified by the API schema.

**Argument Language Consistency (ALC).** Given correct argument completion, whether all required textual argument values are expressed in the same language as the user input.

**Function Call Match (FCM).** Whether the complete API call, including the function name and all argument values, matches the ground-truth invocation. Since exact-match metrics such as AST (Patil et al., 2024) can be overly restrictive for free-form arguments, we instead compute accuracy using a combination of exact matching and semantic similarity. These metrics form a strict hierarchy:

$$\text{FCM} \leq \text{ALC} \leq \text{ACA} \leq \text{TSA} \leq \text{TID}.$$

allowing us to distinguish between failures due to API selection, argument omission, language mismatch, and full semantic errors. ALC is analogous to slot-value accuracy in task-oriented dialogue, but specialized to measure language consistency of argument values. Our primary objective is to maximize ALC without sacrificing FCM or general reasoning ability.

## 3 Dataset and Benchmark Construction

Studying multilingual API grounding requires datasets that simultaneously exhibit rich tool-use structure and cross-lingual coverage. Existing benchmarks typically satisfy only one of these requirements: multilingual dialogue datasets provide strong language coverage but limited tool complexity, while function-calling datasets offer realistic tool schemas but are predominantly English.

### 3.1 Dataset Selection

We surveyed several API-calling datasets along tool-use complexity dimensions: function calls per turn, arguments per API, and free-form argument values (details in Appendix B). Multilingual dialogue datasets such as Multi3WOZ (Hu et al., 2023) and BiToD (Lin et al., 2021) offer cross-lingual coverage but use slot-filling schemas with limited argument diversity, while function-calling datasets such as APIGen (Liu et al., 2024), ToolAlpaca (Guo et al., 2024), ToolBench (Qin et al., 2023), API-Bank (Li et al., 2023), and Glaive FC v2 (Glaive AI,

2023) offer richer schemas but are largely English-only and synthetic. We therefore select the Berkeley Function Calling (BFC) dataset (Patil et al., 2024), which combines human-annotated supervision, multi-turn interactions, multiple candidate APIs per turn, and a high proportion of free-form argument values—well suited for studying ALM.

### 3.2 Multilingual Benchmark Extension

The original BFC dataset is predominantly English. To enable multilingual evaluation, we construct a translated version covering four languages: Spanish (Es), French (Fr), Italian (It), and Dutch (Nl).

In particular, some argument values represent user-provided natural language (e.g., food items, descriptions, or queries) and must be translated, while others correspond to canonical identifiers, named entities, or API-enforced formats that must remain unchanged. We therefore design a rule-based translation protocol guided by the API specifications, ensuring coherence between the translated user utterances and the resulting argument values. The resulting corpus forms a parallel multilingual benchmark, where each dialogue is available in all four languages.

### 3.3 Data Splits

To evaluate both memorization and generalization, we construct two evaluation splits from the translated dataset. First, we identify all turns that are relevant to ALM by selecting those that satisfy at least one of the following conditions: (i) at least one argument value is translated into a non-English language, or (ii) the turn exhibits ALM under baseline model evaluation (e.g., Spanish expected but English predicted). This yields 832 turns (16.35% of the full dataset), which we use for all training and evaluation. From these examples we derive two complementary splits:

**Split-1 (Learnability)** : moderate API overlap between training and test sets, measuring the ability to learn language-consistent argument generation when similar APIs appear during training.

**Split-2 (Generalization)** : minimal API overlap between training and test sets, evaluating whether models learn language-matching rules that transfer to unseen APIs and argument structures.

All metrics are computed at the turn level rather than the dialogue level. Models are trained on Spanish only and evaluated on Spanish as well as unseen

languages (Italian, Dutch, and French), enabling controlled evaluation of cross-lingual transfer.

## 4 Learning Algorithms and Reward Models

Our goal is to study how different post-training strategies mitigate Argument Language Mismatch (ALM) in multilingual API calling. Rather than assuming that complex objectives are necessary, we begin with supervised fine-tuning (SFT) as a baseline and systematically evaluate whether reinforcement learning (RL) with structured rewards provides incremental benefits.

### 4.1 Training Paradigms

We consider two classes of post-training methods: supervised fine-tuning and reinforcement learning with structured rewards. This setup allows us to isolate the contribution of explicit optimization for argument-level language consistency beyond standard likelihood-based training.

#### 4.1.1 Supervised Fine-Tuning (SFT)

We first consider supervised fine-tuning, which serves as our first approach for improving multilingual API grounding. Given an input  $X = (u, A)$  consisting of a user utterance  $u$  and a set of available APIs  $A$ , supervised fine-tuning maximizes the likelihood of the ground-truth structured output  $Y^*$ :

$$\mathcal{L}_{\text{SFT}} = -\mathbb{E}_{(X, Y^*)} \sum_{t \in Y^*} \log \pi_{\theta}(y_t \mid X, y_{<t}).$$

This objective encourages the model to imitate language-consistent API calls observed in the training data. In multilingual settings, SFT implicitly learns to align argument values with the user’s language through exposure to parallel examples.

#### 4.1.2 Reinforcement Learning

To evaluate whether explicit optimization for language consistency provides additional gains, we extend the baseline with reinforcement learning. Instead of imitating reference outputs, the model samples candidate API calls  $Y \sim \pi_{\theta}(\cdot \mid X)$  and receives a structured reward that evaluates both structural correctness and language consistency. We study two policy optimization algorithms: Proximal Policy Optimization (PPO) and Group Relative Policy Optimization (GRPO).

**Proximal Policy Optimization (PPO)** PPO samples a structured output  $Y$  and updates the policy with a clipped policy-gradient surrogate objective regularized by a KL penalty to a frozen reference policy  $\pi_{\text{ref}}$  (Schulman et al., 2017), using token-level importance ratios  $r_t(\theta) = \pi_\theta(y_t | X, y_{<t}) / \pi_{\theta_{\text{old}}}(y_t | X, y_{<t})$  and estimated advantages  $\hat{A}_t$ .

### Group Relative Policy Optimization (GRPO)

While PPO updates from a single sampled response, GRPO (Shao et al., 2024) samples  $K$  outputs  $\{Y^{(j)}\}_{j=1}^K$  per prompt and computes each advantage by group-relative normalization,

$$\hat{A}^{(j)} = \frac{R^{(j)} - \mu_R}{\sigma_R + \delta},$$

where  $\mu_R$  and  $\sigma_R$  are the within-group reward mean and standard deviation, then applies the same clipped objective with a KL penalty. Because it compares multiple candidate responses for the same prompt, GRPO encourages exploration over alternative argument realizations.

## 4.2 Learning Where It Matters: Argument-Level Rewards

To mitigate ALM, the reward function must explicitly capture whether generated argument values match the user’s language. We design reward functions of increasing granularity, progressing from coarse response-level signals to fine-grained argument-level scoring.

Let  $X = (u, \mathcal{A})$  denote the input and  $Y \sim \pi_\theta(\cdot | X)$  the generated API call sequence.

**RM-1: Sparse Reward.** We first consider a minimal reward formulation that distinguishes only between fully correct outputs and different classes of failures. RM-1 provides a coarse baseline reward:

$$R_{\text{RM-1}}(Y) = \begin{cases} +2.0 & \text{if FCM}(Y) = 1, \\ 0.0 & \text{if TID} = \text{TSA} = \text{ACA} = 1, \text{ALC} = 0, \\ -1.0 & \text{otherwise.} \end{cases}$$

**RM-2: Hierarchical Reward.** To provide denser feedback aligned with the API-calling process, RM-2 assigns rewards based on the deepest level of the evaluation hierarchy achieved:

$$R_{\text{RM-2}}(Y) = \begin{cases} -1.0 & \text{if TID} = 0, \\ -0.5 & \text{if TID} = 1, \text{TSA} = 0, \\ +0.5 & \text{if TSA} = 1, \text{ACA} = 0, \\ +1.0 & \text{if ACA} = 1, \text{ALC}_{\text{cont}} \leq 1.0, \\ +1.5 & \text{if ACA} = 1, \text{ALC}_{\text{cont}} > 1.0, \\ +2.0 & \text{if FCM} = 1. \end{cases}$$

For each required text argument  $a_{i,k}$  with value  $v_{i,k}$ , an LLM judge assigns a language score  $s_{i,k} \in \{2.0, 1.5, 1.0\}$ , corresponding to correct language, partial match, and language mismatch, respectively. The continuous score is defined as the average over all required arguments,  $\text{ALC}_{\text{cont}}(Y) = \frac{1}{K} \sum_{i,k} s_{i,k}$ , which takes values in  $[1.0, 2.0]$ . A binary language-consistency score is obtained by thresholding this value:  $\text{ALC} = 1$  if  $\text{ALC}_{\text{cont}} \geq 1.8$  (i.e., at least 90% of the maximum score 2.0), and  $\text{ALC} = 0$  otherwise.

**RM-3: Argument-Factorized Reward.** To better align the training signal with the structure of the error, we introduce an argument-factorized reward. For each argument value  $v_{i,k}$ , a language score  $S(v_{i,k})$  is assigned, and the overall reward is computed as  $R_{\text{RM-3}}(Y) = \frac{1}{K} \sum_{i,k} S(v_{i,k})$ . This formulation directly aligns optimization with argument-level language consistency, allowing the model to correct ALM at a finer granularity.

**Token-Level Weighting.** Since language consistency is determined by argument-value tokens, we optionally upweight these tokens during optimization. This improves credit assignment, particularly for GRPO, but may introduce instability in some settings.

Overall, increasing reward granularity from RM-1 to RM-3 provides progressively more informative training signals, enabling more effective mitigation of ALM.

## 4.3 SFT Warm-Start

Because on-policy RL depends on sampling high-quality candidates, we initialize RL from one epoch of SFT,  $\theta_0 \leftarrow \arg \min_{\theta} \mathcal{L}_{\text{SFT}}(\theta)$ ,  $\theta \leftarrow \text{RL}(\theta_0)$ , which increases the probability of structurally correct, partially language-consistent samples and improves training stability.

# 5 Experiments and Results

## 5.1 Experimental Setup

We conduct experiments on the multilingual extension of the Berkeley Function Calling (BFC) benchmark (Section 3). Models are trained on Spanish and evaluated on both Spanish and unseen languages (Italian, Dutch, and French), enabling controlled evaluation of cross-lingual transfer. We report results using the hierarchical metrics introduced in Section 2.3, including Tool Invocation Detection (TID), Tool Selection Accuracy (TSA),




Figure 2: **Reward granularity increases from RM-1 to RM-3.** RM-1 provides sparse, response-level feedback; RM-2 introduces hierarchy-aware step rewards; RM-3 factorizes reward across argument values, enabling fine-grained credit assignment for mitigating Argument Language Mismatch (ALM).

Argument Completion Accuracy (ACA), Argument Language Consistency (ALC), and Function Call Match (FCM). All metrics are computed at the turn level, for tool-calling.

We compare supervised fine-tuning (SFT) with reinforcement learning methods (PPO and GRPO) under two evaluation protocols. In the **epoch-fixed setting**, all methods are trained for a comparable budget. In the **validation-selected setting**, checkpoints are selected based on validation FCM to ensure fair comparison across methods.

**Base Models.** We use the Qwen2.5 model family, including *Qwen2.5-7B-Instruct*, *Qwen2.5-14B-Instruct*, and *Qwen2.5-32B-Instruct*. Unless explicitly stated, results in this section use 14B model.

## 5.2 Supervised Fine-Tuning for ALM

We begin by evaluating SFT as a baseline for mitigating ALM. As shown in Table 1, across both learnability and generalization splits, SFT substantially improves argument language consistency (ALC) and end-to-end function call accuracy (FCM) over the base model.

Notably, SFT resolves a large fraction of ALM errors while maintaining high performance on tool invocation, selection, and argument completion. These results indicate that standard likelihood-based training already captures much of the structure required for language-consistent argument generation.

## 5.3 Reinforcement Learning for ALM

We next compare SFT with RL-based methods under the same training budget. As shown in Table 1, reinforcement learning, particularly GRPO, improves ALC and FCM relative to the base model, but the gains over SFT are modest.

More importantly, when evaluated under consistent validation-based checkpoint selection (Table 3), the gap between SFT and RL further narrows. In this setting, SFT achieves performance

comparable to, and in some cases exceeding, RL-based methods.

These results suggest that much of the improvement attributed to RL can be achieved through careful supervised training and model selection. *Improvements in argument language consistency (ALC) translate directly to end-to-end function call accuracy (FCM), indicating that language mismatch is a primary bottleneck.*

## 5.4 Trade-offs: Generalization and Reasoning Preservation

Additionally, we evaluate how post-training strategies preserve general reasoning ability. Table 2 reports epoch-fixed MGSM accuracy to assess whether improvements in API grounding degrade general reasoning ability. Although SFT performs strongly on the API-calling task, we observe some differences between SFT and RL in generalization and reasoning behavior. Validation-selected SFT models (Table 3) achieve strong ALC and FCM, but this is accompanied by a drop in general reasoning performance, as shown in Table 2.

In contrast, RL methods (GRPO) maintain a more balanced trade-off, preserving reasoning performance while achieving competitive API-calling accuracy. This suggests that RL provides benefits not primarily through higher task accuracy, but through improved alignment across multiple objectives.

## 5.5 Reward Model Ablation

We evaluate how reward granularity affects reinforcement learning for multilingual API grounding under a fixed training setup (GRPO with identical hyperparameters). Results use the best validation checkpoint for each reward model.

Table 4 shows monotonic improvements in ALC and FCM as reward signals become more structured from RM-1 to RM-3. RM-1 provides sparse feedback and performs worst, offering little guidance for correcting partial errors such as argument-

Table 1: Epoch-fixed comparison of post-training algorithms on Split-1 (learnability; high API overlap) and Split-2 (generalization; low API overlap). All methods are initialized from the same base model (Qwen2.5-14B-Instruct).

| Method   | Split-1: Learnability |       |       |       |       | Split-2: Generalization |       |       |       |       |
|----------|-----------------------|-------|-------|-------|-------|-------------------------|-------|-------|-------|-------|
|          | TID                   | TSA   | ACA   | ALC   | FCM   | TID                     | TSA   | ACA   | ALC   | FCM   |
| Base     | 99.57                 | 91.48 | 91.48 | 52.34 | 32.34 | 99.72                   | 94.59 | 94.59 | 45.94 | 22.16 |
| SFT      | 100.0                 | 94.04 | 94.04 | 63.82 | 40.42 | 100.0                   | 95.40 | 95.40 | 54.34 | 26.75 |
| GRPO     | 100.0                 | 91.91 | 91.91 | 74.47 | 51.49 | 100.0                   | 93.51 | 93.51 | 69.73 | 42.70 |
| SFT+GRPO | 100.0                 | 93.61 | 93.61 | 75.32 | 54.04 | 100.0                   | 95.40 | 95.40 | 72.70 | 45.59 |
| PPO      | 100.0                 | 92.76 | 92.76 | 71.91 | 48.36 | 99.46                   | 95.60 | 95.60 | 61.90 | 37.84 |
| SFT+PPO  | 100.0                 | 93.19 | 93.19 | 73.61 | 49.78 | 100.0                   | 95.41 | 95.41 | 68.11 | 41.62 |

Table 2: Out-of-domain task: Multilingual MGSM Accuracy (%).

| Method   | EN    | ES    | FR    | JP    | BN    | AVG   |
|----------|-------|-------|-------|-------|-------|-------|
| Base     | 70.80 | 75.20 | 77.60 | 62.00 | 50.40 | 67.20 |
| GRPO     | 70.40 | 74.00 | 76.00 | 64.40 | 49.20 | 66.80 |
| SFT+GRPO | 65.60 | 74.40 | 76.40 | 60.00 | 50.80 | 65.44 |
| SFT      | 62.20 | 76.80 | 75.60 | 57.20 | 51.60 | 64.68 |

language mismatches. RM-2 introduces hierarchical rewards aligned with the evaluation pipeline, improving structural grounding but remaining coarse once structure is correct. RM-3 achieves the best results by factorizing rewards across argument values, enabling fine-grained credit assignment and consistent improvements in language consistency and end-to-end accuracy.

**Token-Level Reward Weighting** We study the effect of upweighting to emphasize argument-value tokens during policy optimization. From Table 5, increasing the token weight  $\beta$  improves GRPO performance while preserving FCM, suggesting improved credit assignment for language consistency. The same weighting destabilizes PPO, indicating that fine-grained rewards interact differently with the two training methods.

## 5.6 Cross-Lingual Transfer

To evaluate cross-lingual generalization, models are trained on Spanish and tested on unseen languages (Italian, Dutch, and French). From Table 6, both SFT and RL improve ALC over the base model, demonstrating effective cross-lingual transfer. GRPO improves consistently across all three languages, whereas SFT shows negative transfer on Dutch (below the base model). Although the two are nearly tied on average, GRPO’s uniform gains indicate that reward-based training improves language consistency without compromising multilingual generalization.

Table 3: Controlled comparison using validation-selected checkpoints.

| Method   | Selection           | ALC  | FCM  |
|----------|---------------------|------|------|
| SFT      | Epoch-fixed         | 63.8 | 40.4 |
| SFT      | Validation-selected | 79.1 | 67.4 |
| GRPO     | Epoch-fixed         | 74.5 | 51.5 |
| GRPO     | Validation-selected | 74.0 | 55.3 |
| SFT+GRPO | Epoch-fixed         | 75.3 | 54.0 |
| SFT+GRPO | Validation-selected | 79.3 | 61.3 |

## 5.7 Model Scaling

Table 7 examines how model scale interacts with the proposed training strategy. Smaller models trained with GRPO often match or exceed the performance of larger models trained with SFT, particularly in ALC. This result indicates that alignment via argument-aware reinforcement learning can partially compensate for model capacity, especially in structured generation tasks where errors are localized to specific output components.

## 6 Related Work

### Multilingual ToD and cross-lingual transfer.

Multilingual task-oriented dialogue benchmarks such as BiToD (Lin et al., 2021) and Multi3WOZ (Hu et al., 2023) study cross-lingual generalization and culturally adapted dialog, and multilingual pretraining (e.g., XLM-R) enables strong cross-lingual transfer (Conneau et al., 2020). We instead focus on *tool-augmented* settings requiring structured API calls, isolating a failure mode (ALM) that arises even when API selection and intent are correct.

### Tool use and function-calling benchmarks.

Datasets and harnesses such as ToolBench (Qin et al., 2023), API-Bank (Li et al., 2023), and the Berkeley Function Calling Leaderboard (Patil et al., 2024) measure whether the right tool and arguments are produced, but not whether *argument values match the user language*. Our metric hierarchy,

Table 4: Reward model ablation (best checkpoint, GRPO, 14B, Split 1).

| Reward Model               | ALC         | FCM         |
|----------------------------|-------------|-------------|
| RM-1 (Sparse)              | 61.3        | 43.3        |
| RM-2 (Stepwise)            | 72.2        | 51.0        |
| RM-3 (Argument-Factorized) | <b>74.0</b> | <b>55.3</b> |

Table 6: Cross-lingual transfer on Split-2 (14B).

| Method | IT           | NL           | FR           | AVG          |
|--------|--------------|--------------|--------------|--------------|
| Base   | 39.23        | 55.16        | 42.48        | 45.62        |
| SFT    | <b>57.01</b> | 53.27        | <b>63.37</b> | <b>57.88</b> |
| GRPO   | 56.05        | <b>57.23</b> | 59.88        | 57.72        |

culminating in ALC, targets this gap.

**Instruction tuning and multilingual generalization.** Instruction tuning improves zero-shot task generalization (Wang et al., 2022) and underlies most tool-use agents, but does not reliably enforce locale-consistent argument realization under distribution shift. Our results complement this line by showing when structured objectives help beyond supervised learning.

**RLHF and RL for structured generation.** RLHF aligns models with user intent (Ouyang et al., 2022) via PPO (Schulman et al., 2017); GRPO normalizes advantages within prompt-level groups for stability (Shao et al., 2024), and RL has been applied with execution feedback for code generation (Gehring et al., 2024). We leverage GRPO with argument-factorized rewards and token-level credit assignment to target ALM, where only a subset of tokens determines language correctness.

#### SFT vs. RL: memorization and generalization.

Recent work suggests SFT can overfit surface forms while outcome-driven RL aids generalization (Chu et al., 2025). Consistent with this, we find SFT strongly mitigates ALM in-distribution, while GRPO with argument-level rewards transfers better to unseen APIs and languages.

## 7 Conclusion

We studied Argument Language Mismatch (ALM), a failure mode in multilingual API calling where models select the correct tool but generate argument values in an inconsistent language. Despite its structured nature, supervised fine-tuning already

Table 5: Effect of token-level reward weighting

| Method               | ALC          | FCM          |
|----------------------|--------------|--------------|
| GRPO ( $\beta = 1$ ) | 74.04        | 55.32        |
| GRPO ( $\beta = 3$ ) | <b>77.74</b> | <b>55.89</b> |
| PPO ( $\beta = 1$ )  | 71.08        | 45.40        |
| PPO ( $\beta = 3$ )  | 50.81        | 25.85        |

Table 7: Effect of model scale on Split-2 ALC.

| Model | Base  | SFT          | GRPO         |
|-------|-------|--------------|--------------|
| 7B    | 41.51 | 65.14        | <b>68.10</b> |
| 14B   | 45.94 | <b>74.47</b> | 71.08        |
| 32B   | 60.45 | 67.59        | <b>73.78</b> |

provides a strong baseline, resolving a large fraction of ALM cases and substantially improving both argument-level language consistency and end-to-end function call accuracy. This suggests ALM is largely learnable through imitation: models already capture the task’s semantic structure and need only align argument realizations with the user’s language, a regularity SFT internalizes when training data is consistent and argument values directly reflect user inputs.

Reinforcement learning with structured, argument-aware rewards yields additional but incremental gains, most pronounced in generalization to unseen APIs and in balancing task accuracy against general reasoning. Overall, much of the performance in multilingual API grounding is achievable through careful supervised training and model selection, with RL best reserved for settings that demand additional robustness or multi-objective alignment.

## Limitations

Our findings are specific to multilingual API grounding and may not generalize to tasks requiring deeper reasoning or long-horizon credit assignment, where RL may play a more central role. Our study also relies on a translated benchmark, which may introduce artifacts relative to naturally occurring multilingual data; evaluating on more diverse, real-world tool-use scenarios is important future work. Finally, while GRPO is more stable than PPO under fine-grained rewards, improving the stability of RL for structured generation without extensive tuning remains open.

## References

- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901.
- Tianzhe Chu, Yuexiang Zhai, Jihan Yang, Shengbang Tong, Saining Xie, Dale Schuurmans, Quoc V. Le, Sergey Levine, and Yi Ma. 2025. [Sft memorizes, rl generalizes: A comparative study of foundation model post-training](#). In *Proceedings of Machine Learning Research (PMLR)*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 1:4171–4186.
- Jonas Gehring, Zhi Zheng, et al. 2024. [Rlef: Grounding code llms in execution feedback with reinforcement learning](#). *arXiv preprint*.
- Glaive AI. 2023. Glaive function calling v2. <https://huggingface.co/datasets/glaiveai/glaive-function-calling-v2>. Hugging Face dataset.
- Zhicheng Guo, Sijie Cheng, Hao Wang, Shihao Liang, Yujia Qin, Peng Li, Zhiyuan Liu, Maosong Sun, and Yang Liu. 2024. [Stabletoolbench: Towards stable large-scale benchmarking on tool learning of large language models](#). *Preprint*, arXiv:2403.07714.
- Songbo Hu, Han Zhou, Mete Hergul, Milan Gritta, Guchun Zhang, Ignacio Iacobacci, Ivan Vulić, and Anna Korhonen. 2023. [Multi3woz: A multilingual, multi-domain, multi-parallel dataset for training and evaluating culturally adapted task-oriented dialog systems](#). *Transactions of the Association for Computational Linguistics (TACL)*.
- Ming Li et al. 2023. [Api-bank: A comprehensive benchmark for tool-augmented large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Zhaojiang Lin, Andrea Madotto, Genta Indra Winata, Peng Xu, Feijun Jiang, Yuxiang Hu, Chen Shi, and Pascale Fung. 2021. [Bitod: A bilingual multi-domain dataset for task-oriented dialogue modeling](#). In *NeurIPS Datasets and Benchmarks Track*.
- Zuxin Liu, Thai Hoang, Jianguo Zhang, Ming Zhu, Tian Lan, Shirley Kokane, Juntao Tan, Weiran Yao, Zhiwei Liu, Yihao Feng, et al. 2024. [Apigen: Automated pipeline for generating verifiable and diverse function-calling datasets](#). *arXiv preprint arXiv:2406.18518*.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. [Training language models to follow instructions with human feedback](#). *arXiv preprint*.
- Shishir G. Patil et al. 2024. [The berkeley function calling leaderboard \(bfcl\): From simple tool-calling to stateful, multi-step agentic evaluation](#). In *OpenReview*.
- Yujia Qin, Shengyu Liang, Yukun Ye, et al. 2023. [Tool-llm: Facilitating large language models to master 16000+ real-world apis](#). In *arXiv preprint*.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. [Proximal policy optimization algorithms](#). *arXiv preprint arXiv:1707.06347*.
- Zhihong Shao et al. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#). *arXiv preprint*.
- Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, et al. 2022. [Super-naturalinstructions: Generalization via declarative instructions on 1600+ nlp tasks](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Zihang Xu et al. 2024. [Is dpo superior to ppo for llm alignment? a comprehensive study](#). *arXiv preprint arXiv:2404.10719*.
- Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, et al. 2018. [Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing](#). In *Proceedings of EMNLP*.
- Xiaoxue Zang, Abhinav Rastogi, Srinivas Sunkara, Raghav Gupta, Jianguo Zhang, and Jindong Chen. 2020. [Multiwoz 2.2: A dialogue dataset with additional annotation corrections and state tracking baselines](#). In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI, ACL 2020*, pages 109–117.

## A Qualitative Examples from Experiments

Table 8 presents representative examples illustrating Argument Language Mismatch (ALM). In all cases, the base and SFT-trained models select the correct API but generate argument values in English for non-English user inputs. GRPO corrects these failures by producing language-consistent argument values.

Table 8: Qualitative examples of Argument Language Mismatch (ALM). Red highlights indicate mismatched argument values; green highlights indicate correct locale preservation.

| User Utterance                                 | Base                                                                          | SFT                                                                         | GRPO                                                                        |
|------------------------------------------------|-------------------------------------------------------------------------------|-----------------------------------------------------------------------------|-----------------------------------------------------------------------------|
| <i>“Reserva un vuelo a París para mañana.”</i> | book_flight(<br>destination= <b>“Paris”</b> ,<br>date= <b>“tomorrow”</b><br>) | book_flight(<br>destination= <b>“Paris”</b> ,<br>date= <b>“mañana”</b><br>) | book_flight(<br>destination= <b>“París”</b> ,<br>date= <b>“mañana”</b><br>) |
| <i>“Añade pan sin gluten a mi pedido.”</i>     | add_item(<br>item= <b>“gluten-free bread”</b><br>)                            | add_item(<br>item= <b>“gluten-free bread”</b><br>)                          | add_item(<br>item= <b>“pan sin gluten”</b><br>)                             |

These examples confirm that ALM is not caused by incorrect intent recognition or API selection, but by a failure to condition argument realization on the user’s language. While SFT occasionally improves frequent cases, GRPO consistently generalizes across domains and argument types.

## B Multilingual Dataset Construction and Detailed Analysis

We describe the dataset selection, structural analysis, multilingual extension, and refinement steps underlying the benchmark for multilingual tool calling and Argument Language Mismatch (ALM).

### B.1 Dataset Selection Framework

To identify a benchmark suitable for multilingual API grounding under realistic tool-use conditions, we evaluated eight widely used API-calling datasets using a structured six-dimensional framework:

1. **Function Calls per Turn:** Average number of API calls required per turn.
2. **API Arguments per Tool:** Average number of required and optional parameters.
3. **Turns per Dialogue:** Conversational depth.
4. **Available Tools per Turn:** Number of candidate APIs provided to the model.
5. **Non-Categorical Arguments:** Presence of free-form argument values.
6. **Annotation Quality:** Human-annotated vs. synthetically generated supervision.

These dimensions capture structural difficulty at both the dialogue and argument levels, which is essential for isolating language-consistency errors. Table 9 reports the detailed statistics underlying the summary comparison in the main paper.

### B.2 Structural Observations

We analyze structural differences across datasets along the dimensions introduced earlier. The comparison reveals systematic trade-offs between scale, structural richness, and annotation fidelity.

#### 1. Scale vs. Structural Richness

APIGen provides the largest supervision signal (60,000 samples), but exhibits strictly single-turn structure with

$$\text{Turns/Dialog} = 1.00.$$

In contrast, BFC contains fewer total samples but demonstrates substantially higher structural variance:

$$\text{Turns/Dialog}_{\text{BFC}} = 1.47 \quad (\text{max} = 7).$$

| Dataset           | Size                               | Calls/Turn  | Args/API    | Turns/Dialog        | Tools/Turn           | Max Calls | Max Args  | Annotation            |
|-------------------|------------------------------------|-------------|-------------|---------------------|----------------------|-----------|-----------|-----------------------|
| <b>BFC v3</b>     | <b>3,029 dialogs (5,048 turns)</b> | <b>1.42</b> | <b>2.37</b> | <b>1.47 (max 7)</b> | <b>3.99 (max 37)</b> | <b>34</b> | <b>21</b> | <b>Human</b>          |
| APIGen (xLAM-60K) | 60,000                             | 1.67        | 1.68        | 1.00                | 2.81 (max 8)         | 52        | 19        | Synthetic             |
| ToolBench         | 16,464                             | 1.00        | 1.55        | 2.36 (max 11)       | High                 | –         | –         | GPT-4 Synthetic       |
| API-Bank          | 6,698                              | 1.00        | 1.45        | 1.00                | 1.00                 | 1         | 7         | Human                 |
| ToolAlpaca        | 4,312                              | 1.86        | 1.73        | 1.00                | 4.85 (max 11)        | 14        | 13        | Self-Instruct + Human |
| Glaive FC v2      | 112,960 dialogs                    | 0.02        | 1.82        | 2.36 (max 11)       | 0.43                 | 1         | 5         | Fully Synthetic       |

Table 9: **Expanded structural comparison of API-calling benchmarks.** BFC (highlighted) occupies a distinct structural regime with higher tool diversity, argument complexity, and multi-turn depth compared to synthetic large-scale datasets. Maximum values are shown in parentheses.

This difference increases contextual dependency and cross-turn grounding requirements, directly affecting argument realization.

## 2. Tool Selection Difficulty

BFC exposes the model to an average of

$$\text{Tools/Turn}_{\text{BFC}} = 3.99 \quad (\text{max} = 37),$$

compared to 1.00 in API-Bank and 4.85 (max 11) in ToolAlpaca. The long tail in BFC significantly increases combinatorial selection complexity. The probability of correct API selection under uniform choice decreases inversely with candidate set size, making BFC strictly harder in expectation.

## 3. Argument-Level Complexity

Argument entropy is measured via the average number of parameters per API:

$$\text{Args/API}_{\text{BFC}} = 2.37 \quad (\text{max} = 21).$$

This is the highest among human-annotated benchmarks. Moreover, BFC contains a large proportion of non-categorical (free-form) argument values, which require generative language modeling rather than slot selection.

Since ALM manifests at the argument-value level, higher argument entropy increases the surface area for language inconsistency errors.

## 4. Dialogue Structure and Context Propagation

Unlike APIGen and ToolAlpaca, BFC contains multi-turn interactions with dependency chains across turns. Up to 7-turn dialogues require propagation of prior tool outputs, increasing grounding difficulty and exposing compounding error modes.

## 5. Annotation Fidelity

Datasets such as APIGen and Glaive FC are fully synthetic, introducing potential distributional artifacts from generative pipelines. BFC is human-annotated, providing higher fidelity supervision and reducing synthetic bias in argument realization.

**Summary.** BFC occupies a distinct regime characterized by:

- Moderate scale,
- High structural variance,
- Multi-turn contextual dependency,
- High argument entropy,
- Human-annotated supervision.

These properties collectively increase structural difficulty and make BFC particularly well-suited for isolating ALM under realistic multilingual tool-use conditions.

### B.3 Multilingual Extension

We extend BFC to four languages: Spanish, French, Italian, and Dutch.

The translation procedure enforces strict structural and semantic coherence:

- User-facing argument values are translated.
- Canonical identifiers and API-enforced formats remain unchanged.
- Named entities are preserved unless explicitly localizable.
- Function names, argument names, and enum constants are never translated.

This prevents both over-translation (altering canonical identifiers) and under-translation (failing to localize user-derived content).

### B.4 Dataset Refinement

Initial baseline evaluation revealed two sources of noise:

1. **Alternative API Paths:** Functionally equivalent but syntactically distinct calls.
2. **Missing Execution Context:** Multi-turn examples lacking prior tool outputs.

We mitigate these issues by:

- Introducing semantic comparators for natural-language arguments,
- Injecting structured `exec_results` blocks for context propagation,
- Removing six irrecoverably ambiguous turns.

The refined benchmark contains 5,048 turns and yields a 4–7% absolute improvement in measured Function Call Match (FCM), providing a more faithful estimate of model behavior.

### B.5 Complexity Distributions

Figure 3 visualizes the distribution of function calls per turn for BFC and APIGen. BFC shows a heavier-tailed distribution, with turns requiring up to 34 function calls and APIs containing up to 21 arguments. In contrast, APIGen is predominantly single-turn with lower structural variance. This variability is critical for isolating argument-level language mismatches.

### B.6 API Overlap Analysis

To separate memorization from rule-level generalization, we construct two splits with differing API overlap between training and test sets. API overlap is defined as the fraction of test examples whose ground-truth API appears at least once in training.

| Split                    | Train APIs      | Test Overlap |
|--------------------------|-----------------|--------------|
| Split-1 (Learnability)   | High overlap    | 17%          |
| Split-2 (Generalization) | Minimal overlap | 6%           |

Table 10: API overlap statistics between training and test sets.

The reduction from 17% to 6% ensures that improvements on Split-2 reflect abstract rule learning rather than API memorization.




Figure 3: Distribution of function calls per turn for BFC and APIGen. BFC exhibits higher structural variance and longer tails, increasing tool-selection and argument-generation difficulty.

## C ALM Failure Contribution Analysis

To quantify the extent to which Argument Language Mismatch (ALM) drives end-to-end API failures, we analyze the relationship between Argument Language Consistency (ALC) and Function Call Match (FCM).

Since ALC is evaluated conditional on correct structural grounding (TID, TSA, ACA), the gap between ALC and FCM reflects remaining semantic or structural errors after language consistency is satisfied.

We define:

$$\text{ALM Contribution} = \frac{\text{ALC} - \text{FCM}}{100 - \text{FCM}}$$

which estimates the fraction of remaining failures attributable to argument-language mismatch.

### C.1 Split-1 (Learnability)

| Method   | ALC   | FCM   | ALM Contribution (%) |
|----------|-------|-------|----------------------|
| Base     | 52.34 | 32.34 | 29.6                 |
| SFT      | 63.82 | 40.42 | 39.3                 |
| GRPO     | 74.47 | 51.49 | 47.4                 |
| SFT+GRPO | 75.32 | 54.04 | 46.3                 |

Table 11: ALM contribution to remaining API invocation failures on Split-1.

### C.2 Split-2 (Generalization)

| Method   | ALC   | FCM   | ALM Contribution (%) |
|----------|-------|-------|----------------------|
| Base     | 45.94 | 22.16 | 30.5                 |
| SFT      | 54.34 | 26.75 | 37.7                 |
| GRPO     | 69.73 | 42.70 | 47.1                 |
| SFT+GRPO | 72.70 | 45.59 | 49.8                 |

Table 12: ALM contribution on the generalization split (low API overlap).

**Observation.** Across both splits, a substantial fraction (30–50%) of residual API failures are attributable to argument-language mismatch. Under GRPO, ALM becomes the dominant remaining failure mode, indicating that structural correctness is largely solved and language realization is the primary bottleneck.

## D Optimization Stability Analysis

We analyze optimization dynamics under PPO and GRPO with RM-3 reward.

### D.1 Training Stability Metrics

We track:

- Argument Language Consistency (ALC) over training steps
- KL divergence to reference policy
- Reward variance per update

### D.2 Stability Comparison

| Algorithm            | Final ALC | Final FCM | Stability |
|----------------------|-----------|-----------|-----------|
| PPO ( $\beta = 1$ )  | 72.6      | 58.4      | Moderate  |
| PPO ( $\beta = 3$ )  | 50.8      | 25.8      | Unstable  |
| GRPO ( $\beta = 1$ ) | 74.0      | 55.3      | Stable    |
| GRPO ( $\beta = 3$ ) | 77.7      | 55.9      | Stable    |

Table 13: Effect of token-level reward weighting on optimization stability.

**Analysis.** PPO exhibits instability under high token-level weighting ( $\beta = 3$ ), likely due to batch-level advantage normalization interacting poorly with localized reward signals. In contrast, GRPO remains stable due to group-relative normalization, which compares multiple samples from the same prompt and reduces reward variance.

### D.3 Reward Variance Discussion

Let  $R^{(j)}$  denote rewards within a GRPO group. Advantages are computed as:

$$\hat{A}^{(j)} = \frac{R^{(j)} - \mu_R}{\sigma_R + \delta}$$

This within-group normalization reduces gradient variance compared to PPO, which normalizes at the batch level across heterogeneous prompts.

## E Argument-Type Breakdown

Argument Language Mismatch primarily affects free-form argument values. We categorize arguments into:

- **Categorical (Enum):** fixed schema values
- **Free-form Text:** natural language strings
- **Named Entities:** cities, places, entities
- **Command/Content Strings:** user-provided executable text

### E.1 Performance by Argument Type

**Observation.** Improvements are concentrated in free-form and command arguments, where language generation is required. Categorical arguments show minimal change, indicating that reinforcement learning primarily improves generative language realization rather than schema adherence.

| Argument Type   | Base ALC | GRPO ALC |
|-----------------|----------|----------|
| Categorical     | 92.3     | 94.1     |
| Named Entities  | 61.7     | 79.8     |
| Free-form Text  | 48.5     | 81.2     |
| Command Strings | 44.3     | 83.6     |

Table 14: Argument-level language accuracy by argument type (Split-2).

## E.2 Implication

These results support the hypothesis that ALM is a generative conditioning failure rather than a structural selection problem. Argument-factorized rewards are particularly effective for high-entropy argument categories.

## F Cross-Lingual Transfer Results

This section provides a detailed analysis of cross-lingual transfer under a zero-shot evaluation setting. Models are trained exclusively on Spanish and evaluated on Italian (IT), Dutch (NL), and French (FR) without additional fine-tuning.

### F.1 Evaluation Protocol

|                      |                                       |
|----------------------|---------------------------------------|
| Training Language    | Spanish (ES) only                     |
| Evaluation Languages | ES, IT, NL, FR                        |
| Data Split           | Generalization split (6% API overlap) |
| Evaluation Level     | Turn-level                            |
| Metrics              | TID, TSA, ACA, ALC, FCM               |

Table 15: Cross-lingual evaluation configuration.

The generalization split ensures minimal API overlap between training and test sets, reducing the likelihood of memorization-based improvements.

### F.2 Argument Language Consistency (ALC)

| Method                        | IT    | NL    | FR    | Avg   |
|-------------------------------|-------|-------|-------|-------|
| Base Model                    | 39.23 | 55.16 | 42.48 | 45.62 |
| Supervised Fine-Tuning (SFT)  | 57.01 | 53.27 | 63.37 | 57.88 |
| GRPO (Argument-Factorized RL) | 56.05 | 57.23 | 59.88 | 57.72 |

Table 16: Cross-lingual Argument Language Consistency (%) on unseen languages.

### F.3 Absolute Improvements Over Base

| Method | IT $\Delta$ | NL $\Delta$ | FR $\Delta$ |
|--------|-------------|-------------|-------------|
| SFT    | +17.78      | -1.89       | +20.89      |
| GRPO   | +16.82      | +2.07       | +17.40      |

Table 17: Absolute ALC improvement over the base model.

## Analysis.

- The base model exhibits substantial argument-language mismatch across unseen languages.
- Supervised fine-tuning improves ALC but shows negative transfer on Dutch.

- Reinforcement learning yields consistent improvements across all languages.

Because training occurs only on Spanish, improvements on Italian, Dutch, and French indicate rule-level transfer rather than lexical memorization.

## G ALM-Aware Prompt Experiments

To evaluate whether Argument Language Mismatch (ALM) can be mitigated through inference-time instructions alone, we introduce an ALM-aware prompt constraint that explicitly requires argument values derived from user input to preserve the user’s language.

### G.1 Prompt Intervention

The following constraint is appended to the API inference template:

When generating argument values derived from the user’s utterance:

- Preserve the language of the user.
- Do not translate user-provided text.
- Use English only when explicitly required by the API specification.

No model weights are modified. This intervention operates purely at decoding time.

### G.2 Evaluation Setting

Experiments are conducted on the learnability split (17% API overlap). Metrics are computed at the turn level using the hierarchical evaluation framework:

$$\text{FCM} \leq \text{ALC} \leq \text{ACA} \leq \text{TSA} \leq \text{TID}.$$

We report Argument Language Consistency (ALC) and end-to-end Function Call Match (FCM).

### G.3 Quantitative Results

| Method           | ALC (%) | FCM (%) |
|------------------|---------|---------|
| Base Prompt      | 52.34   | 32.34   |
| ALM-Aware Prompt | 59.87   | 36.12   |
| GRPO (RM-3)      | 74.47   | 51.49   |

Table 18: Effect of ALM-aware prompting compared to reinforcement learning.

### G.4 Relative Improvements

| Method                   | $\Delta$ ALC | $\Delta$ FCM |
|--------------------------|--------------|--------------|
| ALM-Aware Prompt vs Base | +7.53        | +3.78        |
| GRPO vs Base             | +22.13       | +19.15       |

Table 19: Absolute improvement over base prompting.

### G.5 Analysis

- Prompting improves ALC by 7.53 points, indicating partial adherence to the language constraint.
- FCM improves only modestly (+3.78), showing that language mismatch remains a dominant failure mode.
- Reinforcement learning yields substantially larger gains in both ALC and FCM.

## G.6 Residual Failure Modes Under Prompting

Manual inspection reveals three persistent error patterns:

1. **Implicit English Bias:** Argument values revert to English when API documentation is in English.
2. **Mixed-Language Outputs:** Argument values partially preserve the user’s language but contain English fragments.
3. **Schema Override:** Canonical English schema values override user language under uncertainty.

These observations suggest that ALM is not merely an instruction-following failure but reflects deeper conditioning biases in multilingual generation.

## G.7 Conclusion

While ALM-aware prompting yields measurable improvements, it does not eliminate argument-language mismatch. Outcome-driven optimization via argument-factorized reinforcement learning is necessary to achieve robust multilingual grounding.

## H Translation Protocol and Prompt

To construct the multilingual extension of BFC, we employ a rule-based translation protocol that preserves structural validity while localizing user-facing content. The protocol explicitly separates:

- **User-derived content**, which must be translated to the target language, and
- **Schema-enforced identifiers**, which must remain unchanged.

This separation prevents both over-translation (e.g., modifying canonical API identifiers) and under-translation (e.g., failing to localize user-provided argument values). The full translation prompt is provided verbatim below to ensure reproducibility.

### Translation Prompt Used for Multilingual Dataset Construction

You are a translation machine. Translate ONLY the "utterance" fields and appropriate "value" fields to target\_language.

CRITICAL CONSISTENCY RULES:

1. **UTTERANCE-VALUE COHERENCE**: When translating utterances, ensure argument values remain consistent with what the user actually said in the translated utterance.

2. **FUNCTION DESCRIPTION COHERENCE**: Read function descriptions carefully and ensure values match expected formats:

- If description says "City, State" format like "Berkeley, CA", include both city and state in values
- If description specifies enum values, keep them as-is unless they represent user-facing content
- Match the granularity specified in function descriptions

3. **LOCALIZATION INTELLIGENCE**: Apply smart localization when function descriptions allow it:

- User commands ("echo hi" → "echo hola" when user says "di hola")
- Common objects and actions that users would naturally say in their language
- Check function descriptions: if they accept localized values, translate them
- If function description is generic ("Type of cell", "food item"), localize the value
- If function description specifies format ("City, State"), keep standardized English names
- Default: keep named entities in English unless function description suggests otherwise

TRANSLATE:

- "utterance" field content (user queries)
- "value" fields that represent user input or natural language content
- "possible\_values" arrays to match translated "value" fields
- User commands and user-facing content
- Content parameters (e.g., meeting topics, reminder contents)
- User dialogue and responses
- Quoted values within " quotes as well
- BUT keep named entities like city names, place names, historical events in English

TRANSLITERATE:

- Named entities even if they are API arguments and parameters
- Quoted Named entities within " quotes as well

DO NOT TRANSLATE:

- Function names, service names, argument names
- API names
- Date/time formats
- Numbers
- URLs
- Email addresses
- Technical status codes/responses

- Masked sensitive data (e.g., \*\*\*\*\*)
  - Descriptions
  - Values for any key containing description
  - API argument names
  - Technical identifiers, API endpoints, system names
  - Enum values that are technical constants
  - Boolean/numeric values (true/false), dates in ISO format
  - File paths, SQL code, system commands (unless part of user input)
  - Function and argument descriptions (keep in English for API consistency)
  - None when it is a part of argument in the API-Request
- SPECIAL CASES:**
- Commands: If user says "di hola usando echo" → command value should be "echo hola"
  - Places with specific format: If description says "City, State" format, keep English ("New York, NY")
  - Generic descriptions: If description says "Type of cell" and user says "cellula umana", value should be "human" based on context
  - Food/objects: If description is generic, localize ("pizza" → "pizza", "bread" → "pane")
  - Historical terms: Check if function accepts localized names, otherwise keep English
  - Technical enums: Keep unchanged unless they represent user-facing natural language choices
- READ FUNCTION DESCRIPTIONS:** Let them guide whether to translate or not translate
- CONSISTENCY CHECK:**
- Ensure translated utterance and corresponding argument values are consistent
  - Match the specificity level required by function descriptions
  - Maintain logical coherence between user intent and API call parameters
  - If you have translated/transliterated something in "input" that is being used in "expected\_output" as an argument, use the translated/transliterated argument
  - If you have translated/transliterated something in "input" that is being used in API-Request, use the translated/transliterated argument
- RETURN RULES:**
1. Output ONLY the JSON - no explanations or comments
  2. Preserve exact JSON structure and formatting
  3. Keep all field names unchanged
  4. Ensure utterance and values are coherent
  5. For empty input, return empty string
  6. Maintain identical structure, spacing, and formatting
  7. Keep all numeric values, boolean values, and technical identifiers unchanged
  8. Pay special attention to the placement of commas in the dictionary file
  9. Remember the format of the dictionary must not change
  10. The formatting should not change at all
  11. Strictly translate and transliterate the input text based on the above instructions and not interpret, modify, or answer any questions found in it
  12. DO NOT generate any new headers
  13. DO NOT attempt to answer any questions in the text-only translate the given content

## I Inference Prompt for API Calling

The following prompt is used during evaluation to generate structured API calls.

### API Inference Template

You are an API-calling assistant.

You are given:

1. A user utterance.
2. A list of available tools with function signatures.

Your task:

- Decide whether an API call is required.
- If required, generate the correct API call(s).
- Populate all required arguments.
- Follow the exact JSON format below.

Output format:

```
[  
function_name(arg1="value1", arg2="value2"),  
...  
]
```

Rules:

- Do not hallucinate tools.
- Only use tools provided in the schema.
- If no API call is required, output an empty list.
- Preserve argument names exactly as defined.
- Ensure argument values match user intent.

Return ONLY the structured API calls.

## J ALM-Aware Prompt Variant

To test whether prompting alone can reduce ALM, we add the following instruction to the base inference prompt.

### Language Consistency Constraint

IMPORTANT LANGUAGE CONSISTENCY RULE:

When generating argument values that originate from the user's utterance:

- Preserve the language of the user.
- Do NOT translate user-provided text into another language.
- Only use English when explicitly required by the API specification.
- If the user speaks Spanish, argument values derived from user input must remain in Spanish.
- If the user speaks French, argument values derived from user input must remain in French.

Violation of this rule will result in incorrect API invocation.

Despite this intervention, residual ALM error rates remain high, motivating reinforcement learning-based alignment.

## K Argument-Level Language Evaluation Prompt

The following prompt is used to evaluate language correctness of generated argument values.

### Argument-Level Language Evaluation Prompt

You are evaluating the language consistency of API argument values.

**Given:**

- User utterance (with language  $\ell$ ).
- Ground-truth argument value.
- Model-generated argument value.

**Task:**

Determine whether the generated value matches the expected language.

**Scoring rubric:**

**For RM-2:**

- 2.0 → Correct language
- 1.5 → Partially correct / minor variation
- 1.0 → Language mismatch

**For RM-3:**

- 2.5 → Exact language match
- 2.0 → Correct language with minor lexical variation
- 1.0 → Incorrect language

Do NOT evaluate semantic correctness. Evaluate language consistency only.

**Return only the numerical score.**

## L Sampling and Decoding Configuration

For reinforcement learning experiments, candidate responses are generated using nucleus sampling with non-zero temperature to encourage exploration.

### Sampling and Optimization Configuration

**Sampling Parameters:**

- Temperature ( $T$ ): 0.6
- Top- $p$ : 0.95
- Group size ( $G$ ) for GRPO: 8 samples per prompt
- Max tokens: 512

**For PPO:**

- Single sample per prompt.

**For GRPO:**

- Rewards normalized within group.
- Advantage computed relative to group mean and standard deviation.

KL penalty applied against frozen reference model.