

LLM-Supported Natural Language to Bash Translation

Finnian Westenfelder^{1,2}, Erik Hemberg¹, Miguel Tulla¹,
Stephen Moskal¹, Una-May O’Reilly¹, Silviu Chiricescu³

¹ALFA Group MIT-CSAIL, ²Draper Scholar, ³Charles Stark Draper Laboratory
{finnw,ehemberg,mtulla,smoskal,unamay}@mit.edu silviu@draper.com

Abstract

The Bourne-Again Shell (Bash) command-line interface for Linux systems has complex syntax and requires extensive specialized knowledge. Using the natural language to Bash command (NL2SH) translation capabilities of large language models (LLMs) for command composition circumvents these issues. However, the NL2SH performance of LLMs is difficult to assess due to inaccurate test data and unreliable heuristics for determining the functional equivalence of Bash commands. We present a manually verified test dataset of 600 instruction-command pairs and a training dataset of 40,939 pairs, increasing the size of previous datasets by 441% and 135%, respectively. Further, we present a novel functional equivalence heuristic that combines command execution with LLM evaluation of command outputs. Our heuristic can determine the functional equivalence of two Bash commands with 95% confidence, a 16% increase over previous heuristics. Evaluation of popular LLMs using our test dataset and heuristic demonstrates that parsing, in-context learning, in-weight learning and constrained decoding can improve NL2SH accuracy by up to 32%. Our findings emphasize the importance of dataset quality, execution-based evaluation and translation method for advancing NL2SH translation. Our code is available at <https://github.com/westenfelder/NL2SH>.

1 Introduction

The default command-line interface (CLI) for interacting with Linux systems is the Bourne-Again Shell (Bash) (Shotts, 2019). Bash commands allow computer users to control processes, interact with the file system and manage the network. However, using Bash requires knowledge of numerous utilities, each with unique parameters and complex syntax (Ramey and Fox, 2024). Moreover, the reference documentation for these utilities, called manual pages, can be cumbersome and confusing

(Kerrisk, 2024). This makes the CLI a barrier for inexperienced users and increases the chance of errors for experienced users (Agarwal et al., 2021).

Language models that convert natural language to command-line instructions, a task referred to as *NL2SH*, *NL2CMD* or *NL2Bash Translation*, offer a promising solution to this problem (Agarwal et al., 2021). We use the term *NL2SH model* to refer to models trained specifically for the task of NL2SH, as well as the NL2SH capabilities of general-purpose large language models (LLMs). Figure 1 shows an example of natural language to Bash command translation. NL2SH models are well suited for CLIs because they are designed for text-based interactions. NL2SH models can simplify human-computer interactions by allowing users to interact with Linux systems through natural language on the command line. This advancement enhances usability by reducing the need for syntax memorization (Sammet, 1966).

Input: Natural Language
List files in the /workspace directory that were accessed over an hour ago.

Output: Bash Command
`find /workspace -type f -amin +60`

Figure 1: Natural language to Bash command translation example from our NL2SH-ALFA dataset.

The use of NL2SH models necessitates benchmarks to measure task accuracy (Lloyd et al., 2024; Sadykov, 2024; Rosenkilde et al., 2024; Microsoft, 2024; Services, 2024). A NL2SH benchmark requires test data consisting of natural language prompts and ground truth commands (referred to as instruction-command pairs). Given a natural language prompt, a NL2SH model generates a Bash command. A benchmark must then use a heuristic to determine if the model command is functionally equivalent to the ground truth command. Determining functional equivalence of commands is difficult because there are multiple possible correct com-

mands for a given task, due to a wide range of interchangeable utilities. Further, command execution may not result in identical outputs, neutralizing evaluation with string comparison. Current benchmarks do not accurately measure NL2SH model performance due to errors in assessment data and inaccurate heuristics for determining the functional equivalence of commands (Yang et al., 2023; Song et al., 2024; Aggarwal et al., 2024). This makes it difficult to assess model capabilities and measure methods for improving model performance.

To address these challenges, we investigate the following research questions: (1) How can we validate NL2SH datasets to ensure models are evaluated using accurate assessments? (2) How can we design a functional equivalence heuristic that accurately measures the quality of model translations? (3) How can we improve the accuracy of NL2SH models as measured by a reliable benchmark?

Our contributions are summarized as follows: (1) We create a manually verified test dataset of 600 instruction-command pairs and a training dataset of 40,939 pairs, increasing the size of previous test and training datasets by 441% and 135%, respectively. (2) We present a novel functional equivalence heuristic that combines command execution with LLM evaluation of command outputs, capable of determining the functional equivalence of two Bash commands with 95% confidence, a 16% increase over previous heuristics. (3) We evaluate popular LLMs using our test data and heuristic and demonstrate that parsing, in-context learning, in-weight learning and constrained decoding can improve NL2SH accuracy by up to 32%.

2 Background

NL2SH translation falls under the broader domain of machine translation, where models automatically translate text or speech from one language to another. LLMs are well suited for this task, enabling translation that was impossible with previous methods (Zhu et al., 2024b). Evaluating NL2SH models requires determining the functional correctness of generated commands. Functional correctness is defined as whether the code produces the correct output for each input, as specified, or as compared to ground truth (Chen et al., 2024). Functional correctness does not consider the diversity of code generated, or other factors such as run time and memory consumption (Chon et al., 2024).

Ensuring the functional correctness of code is

difficult because validation methods are error-prone and take an impractical amount of time for large volumes of code (Chen et al., 2024). There are two main validation techniques: static and dynamic analysis. Static analysis checks code without execution, using parsers, lexical analysis or control flow checking (Shaikhelislamov et al., 2024). Dynamic analysis evaluates code outputs and runtime behavior using an execution environment (Yang et al., 2023). Some frameworks combine static and dynamic analysis (Aggarwal et al., 2024).

Determining the functional correctness of Bash commands translated from natural language is a sub-problem of validating code correctness. We assess the functional correctness of Bash commands by comparing the generated (model) command with a ground truth Bash command. We define the term "*functional equivalence heuristic*" (FEH) to describe a heuristic that performs this comparison and determines the functional correctness of a Bash command. Due to varying definitions in this field, we provide notation in Table 1 defining the terms used in this paper.

Table 1: Definition of terms.

Term	Definition
Natural Language Task	$t \in T$ English
Ground Truth Command	$b \in B$ Bash-5.2
Model Command	$b' \in B$
Functionally Equivalent Command	$\hat{b} \in B$
Docker Environment	$e \in E$ Linux
Command Output and Side Effects	$o \in O$ stdout and system state
Model Weights	$\theta \in \mathbb{R}$
Translation	$f : T \times \mathbb{R} \rightarrow B \quad f(t, \theta) = b'$
Execution	$g : B \times E \rightarrow O$ $g(b, e) = o \quad g(b', e) = o'$
Ideal FEH	$m : T \times O \times O \rightarrow \{0, 1\}$ $m(t, o, o') = \begin{cases} 1 : o \approx o' \\ 0 : o \neq o' \end{cases}$
Training Dataset	$D_T : \{(t_i, b_i) \mid i = 1, 2, \dots, x\}$
Test Dataset	$D_H : \{(t_i, b_i, \hat{b}_i) \mid i = 1, 2, \dots, y\}$ $m(t, g(b, e), g(\hat{b}, e)) = 1 \quad \text{ideal } m$ $D_T \cap D_H = \emptyset$
Benchmark	(D_H, m)

3 Related Work

NL2SH translation is a well-studied natural language processing (NLP) task. Table 2 summarizes the contributions of previous work by listing the names of datasets, functional equivalence heuristics (FEHs), and models used in this field. The table is sparsely populated because the majority of previous work focused on improving a NL2SH dataset, FEH, or model in isolation.

Table 2: Summary of NL2SH datasets, FEHs, and models created in previous work.

Citation	Datasets	FEHs	Models
Lin et al. (2018)	NL2Bash	-	-
Gros (2019)	-	-	ALinux
Agarwal et al. (2021)	-	NL2CMD	Tellina
Fu et al. (2021)	-	-	Magnum
Ramesh (2022)	NL2CMD	-	-
Bharadwaj et al. (2022)	-	-	AST
Jenson and Liu (2022)	-	-	T5, GPT2
Shi et al. (2023)	-	-	ShellGPT
Yang et al. (2023)	InterCode-Bash	InterCode	-
Mali (2023)	text_to_bash	-	-
Cassano et al. (2023)	MultiPL-E	Unit Tests	-
Song et al. (2024)	-	TSED	-
Aggarwal et al. (2024)	CodeSift	CodeSift	-
Vo et al. (2024)	IBM_Instana	Podman	-
Romit (2024)	LinuxCmds	-	-
Lloyd et al. (2024)	-	-	Warp AI
Sadykov (2024)	-	-	shell-gpt
Rosenkilde et al. (2024)	-	-	Copilot CLI
Services (2024)	-	-	CodeWhisperer
Microsoft (2024)	-	-	AI Shell
Chatterjee et al. (2024)	-	-	ScriptSmith
Joshi (2024)	-	-	CodeLlama2
Ours (2025)	NL2SH-ALFA	IC-ALFA	Llama, Qwen, GPT

The 2020 NeurIPS NLC2CMD Competition formalized the task of NL2SH translation by providing a human-curated NL2Bash dataset of 9,305 instruction-command pairs and the NL2CMD benchmark for evaluating submitted models (Lin et al., 2018; Agarwal et al., 2021). The competition resulted in numerous NL2SH models and showed that fine-tuning a pre-trained foundation model could outperform dedicated transformer (Fu et al., 2021), recurrent neural network (Lin, 2017), abstract syntax tree (Bharadwaj et al., 2022), and sequence to sequence (Gros, 2019) based models for the task of NL2SH translation (Shi et al., 2023).

The NL2CMD benchmarks’s FEH parses commands and assigns a similarity score based on the utilities used, order of utilities, and number of utility flags. This heuristic outperforms conventional string comparison techniques, such as edit distance, for measuring the functional equivalence of commands. However, Agarwal et al. (2021) state that the NL2CMD FEH could be improved by executing commands and measuring the similarity of the outputs and side effects. Verification by execution is preferable because Bash is a Turing-complete language, so verifying the equivalence of two commands before execution is undecidable due to side effects (Churchill et al., 2019). Despite this known shortcoming, the NL2CMD benchmark is widely used for model evaluations (Fu et al., 2023).

Yang et al. (2023) address this shortcoming with the InterCode-Bash benchmark. The benchmark uses a subset of 224 instruction-command pairs from the NL2Bash dataset for testing. In-

terCode’s FEH executes the model command and ground truth command in identical Docker containers (Merkel, 2014). The results of command execution are then compared using three checks. First, the pre and post-execution states of each container are compared using git-diff. Second, the file contents of each container are compared using MD5 hashes. Third, the standard output of both commands are vectorized and compared using the term frequency, inverse document frequency (TFIDF) method (Sparck Jones, 1988). If every check finds the execution results identical, the model and ground truth command are considered functionally equivalent (Yang et al., 2023). Although this method is more accurate than previous heuristics, it will fail to identify a valid model command that has syntactically different output from the ground truth command.

Huang et al. (2022) and Vijayaraghavan et al. (2024) present similar execution-based frameworks for Jupyter Notebooks and VHDL code, respectively. Vo et al. (2024) describe an execution-based framework similar to the InterCode-Bash benchmark using Podman containers. Unfortunately, the code for their FEH and the 50 instruction-command pairs in their test dataset are not public.

Focusing on the FEH, Song et al. (2024) present a novel benchmark that uses OpenAI’s GPT-4 model to determine functional equivalence. Their FEH passes the model command and ground truth command to GPT-4 with the prompt, "Given 2 Bash commands, please generate a similarity score from 0 to 1." Song et al. (2024) find this method fails to determine functional equivalence because current LLMs are unable to emulate command execution. Maveli et al. (2024) and Naik (2024) confirm this finding with broader evaluations of LLM’s ability to determine semantically equivalent or different pairs of programs. They find LLMs show a lack of depth in understanding code semantics.

Aggarwal et al. (2024) attempt to advance the work presented by Song et al. (2024), by addressing the scalability constraints of execution-based frameworks. Their FEH, CodeSift, uses an LLM to convert the model command to a natural language description. Then they compare this natural language description with the original natural language task using an LLM. While they find CodeSift to be more effective than conventional FEHs, their work lacks comparison with execution based heuristics. Further, their FEH introduces uncertainty by requiring accurate Bash to natural

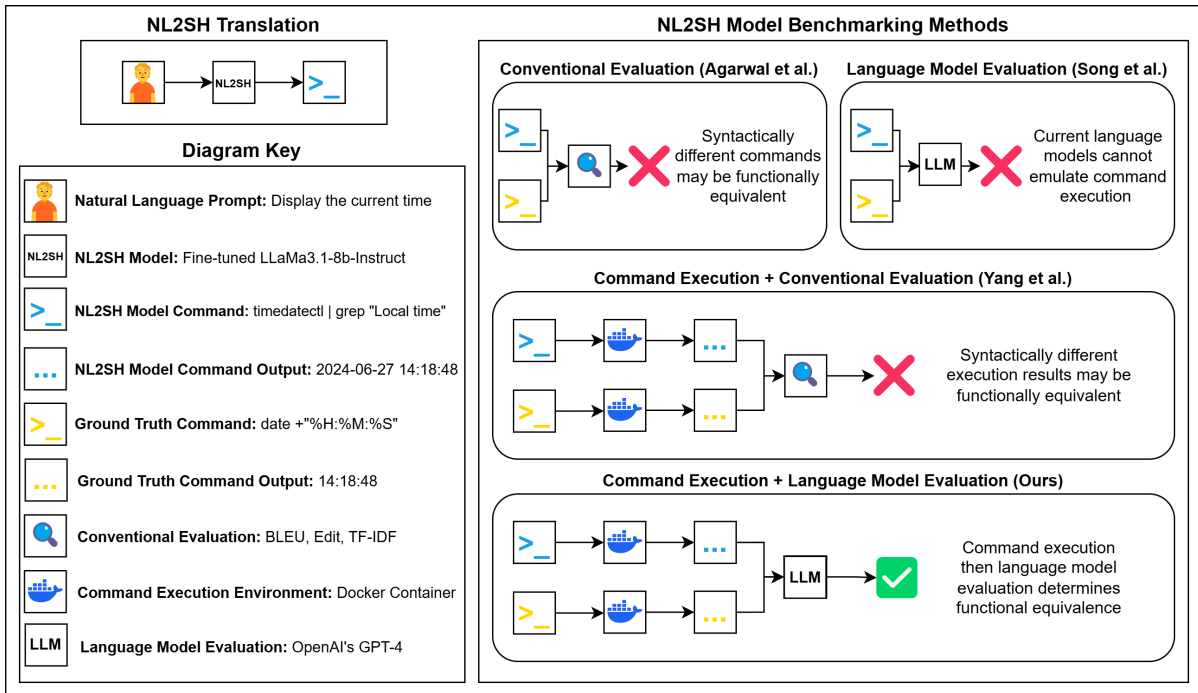


Figure 2: A diagram of NL2SH translation with a comparison of functional equivalence heuristics.

language translation, which is equally as challenging as the task they aim to measure, natural language to Bash translation. Their follow on work concludes that *"executing [Bash] scripts within a controlled environment would offer more reliable assessments"* (Chatterjee et al., 2024).

The MultiPL-E benchmark is widely used for evaluating code generation models, containing 540 Bash scripting tasks (Cassano et al., 2023). This benchmark uses unit tests to determine if a generated script produces the expected output for a given input. The use of unit tests is sufficient for this benchmark because its tasks, such as string manipulation and math calculations, are deterministic and result in simple outputs. This method fails to assess file manipulation, system administration and network management tasks because they produce more complex outputs than what can reasonably be assessed with unit tests.

Current SOTA NL2SH models use general purpose LLMs for translation. In practice, users can either accept, reject, or edit the model translation (Lloyd et al., 2024; Rosenkilde et al., 2024). A human-in-the-loop approach is necessary because the models may produce incorrect translations (Hui et al., 2024). Efforts to improve model performance include fine-tuning and prompt engineering. Jen-son and Liu (2022) fine-tune the BART, T5, and GPT-2 models on the NL2Bash dataset and find

model performance improves as measured by the NL2CMD benchmark. Joshi (2024) conducts a similar study, fine-tuning the CodeLlama2 model on the NL2Bash dataset. Unfortunately, neither of these studies evaluate their models using a reliable execution-based benchmark.

We begin by creating verified and expanded NL2SH datasets starting from multiple datasets presented in previous work. Next, we combine the InterCode execution FEH presented by Yang et al. (2023) with the language model evaluation presented by Song et al. (2024). Using our new datasets and FEH, we evaluate methods for improving model performance. Figure 2 summarizes the shortcomings of FEHs in previous work.

4 Methodology

4.1 Dataset Creation

Bash is considered a low-resource programming language due to the limited availability of NL2SH data (Joel et al., 2024). We aim to augment NL2SH datasets and begin with an evaluation of InterCode. All 224 commands in the InterCode dataset were manually curated from the NL2Bash dataset presented by Lin et al. (2018), containing 9,305 instruction-command pairs. The InterCode dataset is significantly smaller than the NL2Bash dataset because a Docker environment is configured for each command, enabling execution. We manually

verify all 224 instruction-command pairs and find that over half of the InterCode dataset is erroneous.

Table 3 shows the number of errors organized by type. We define three types of errors: invalid prompt, invalid command, and invalid environment. An invalid prompt error refers to a natural language instruction that describes an impossible task or does not give enough information to accomplish the task. An invalid command error refers to a Bash command that does not accomplish the task described in the prompt or does not execute. An invalid environment error refers to an incorrect Docker configuration such as a missing file, environment variable, or utility that prevents a valid command from accomplishing the task. Our manual verification reveals 102 instruction-command pairs with one or more errors and 11 duplicate pairs.

Table 3: InterCode dataset errors.

Error Type	Count	Percentage
Duplicate	11	4.9%
Invalid Prompt	17	7.6%
Invalid Cmd	24	10.7%
Invalid Env	18	8.0%
Invalid Prompt and Cmd	29	12.9%
Invalid Prompt and Env	0	0.0%
Invalid Cmd and Env	3	1.3%
Invalid Prompt, Cmd and Env	11	4.9%
Invalid Total	113	50.4%
Valid Total	111	49.6%

We fix 82 of these errors by correcting natural language prompts, Bash commands, and Docker configuration files. We remove 11 duplicate and 20 irreparable pairs from the dataset, resulting in 193 verified pairs. We create 117 more verified pairs by referencing Bash tutorials and books, such as The Linux Command Line by Shotts (2019) and the Linux Command Line and Shell Scripting Bible by Blum and Bresnahan (2021).

Additionally, for our 300 verified pairs, we create a second Bash command that accomplishes the task described in the prompt. Our final test dataset contains two functionally equivalent, ground truth Bash commands for each natural language instruction, for a total of 600 instruction-command pairs. This is an increase of 441% over the 111 valid commands in the InterCode dataset. Our annotated corrections for the InterCode dataset errors can be found on HuggingFace¹.

We collect training data by combining the NL2Bash dataset with three publicly available NL2SH datasets (Ramesh, 2022; Romit, 2024; Mali, 2023). Further, we scrape the tldr-pages,

a collection of example Bash commands, as a new data source (Krishna et al., 2024). We combine these data sources and deduplicate with exact matching. Then, we use the bashlex parser to remove unparseable commands (Kamara, 2016).

We de-conflict our training and test dataset using exact matching and semantic similarity, removing 917 rows from the training data. First, we remove rows from the training data that exactly match instructions or commands in the test data. Next, we remove pairs from the training data with a natural language prompt that is syntactically similar to a prompt in the test data using the mxbai-embed-large-v1 embedding model and a cosine similarity threshold of 0.9 (Lee et al., 2024). Our final training dataset contains 40,939 instruction-command pairs, an increase of 135% over the previous largest dataset. Figure 3 shows the relationships between data sources used to create our datasets. Our final datasets can be found on HuggingFace².




Figure 3: Relationships between NL2SH datasets.

4.2 Functional Equivalence Heuristic (FEH)

Our evaluation of related work in Section 3 reveals the InterCode benchmark is more accurate than previous NL2SH benchmarks because its FEH uses execution-based evaluation. However, its TFIDF method for comparing command outputs may fail to determine functional equivalence because syntactically different outputs may convey the same information to the end user.

For example, consider the prompt "Print the disk usage of the current directory", a ground truth command of "du -s ." and a model command of "du -d 0 -h". The first command outputs the number of bytes and the second command outputs the number of bytes in human-readable format. The two commands are functionally equivalent, conditioned on the prompt. However, their outputs contain different characters, resulting in a low similarity score using TfIdfVectorizer. Similar issues arise when comparing commands that print hardware information or system time in different formats, display

¹InterCode-Corrections HuggingFace

²NL2SH-ALFA HuggingFace

text with line numbers or other delimiters, and use non-deterministic utilities, such as those that interact with the network. The difficulty of determining functional equivalence is exacerbated by ambiguity in natural language prompts, which is an inherent problem with human inputs.

To address this problem, we replace the TfidfVectorizer method for comparing command outputs with an LLM. Our intuition is that an LLM can determine more complex cases of functional equivalence by evaluating the semantics of command outputs with relation to the prompt. Replacing TfidfVectorizer with an LLM increases the computational cost of the FEH. Additionally, since LLMs are stochastic, our FEH has inherent variability. We compare our FEH with previous heuristics in Section 5.1, finding it achieves superior performance.

We present our FEH and test dataset as a new version of the InterCode benchmark, InterCode-ALFA. Our benchmark and datasets are released under MIT licenses. In addition to the dataset and FEH modifications, we add error handling and update the Docker configuration files to use stable Linux releases. We also identify and fix an error in the benchmark’s Docker reset script that causes the filesystem structure of the two execution environments to diverge. We publish the benchmark source code on GitHub³ and provide a Python package on PyPI⁴ for ease of use. Our benchmark and dataset can be configured with 10 lines of code, simplifying the process for evaluating new models.

4.3 Translation Methods

Using our benchmark, we evaluate the NL2SH performance of the Llama, Qwen and GPT model families (Llama, 2024; Hui et al., 2024; OpenAI, 2024). We find the models have poor baseline performance and identify three translation failure modes: incorrect output format, incorrect utility and syntactically incorrect Bash command.

Incorrect output format refers to a translation with extraneous information, such as an explanation of the translation, or additional text formatting, such as markdown code blocks. Incorrect utility refers to a translation with a utility that cannot accomplish the task described in the prompt. Syntactically incorrect Bash command refers to a translation that is not valid Bash syntax.

To address these failure modes, we evaluate four methods for improving model performance:

³InterCode-ALFA GitHub

⁴InterCode-ALFA PyPI

markdown parsing 4.3.1, constrained decoding (CD) 4.3.2, in-context learning (ICL) 4.3.3, and in-weight learning (IWL) 4.3.4. Our results are listed in Section 5.2.

4.3.1 Markdown Parser

Despite prompting models with *"You will not output markdown or other formatting"*, translations often include markdown formatting, likely due to instruct tuning. We implement a markdown parser to extract the Bash command from the first code block in model outputs, discarding additional text.

4.3.2 Constrained Decoding

We inspect the token probabilities for each ground truth Bash command in our test dataset using the Llama3.1-8b-Instruct model. We find the average relative probability of the first token is four orders of magnitude smaller than the following tokens. In our case, the first token of each command is a Bash utility. This indicates the model is unlikely to select the correct utility as the first token. However, if it does select the correct utility, the following flags and arguments are correct with high probability. We address this by constraining the first tokens of the model output to a list of Bash utilities using grammar-constrained decoding Geng et al. (2023).

4.3.3 In-Context Learning

In-context learning can improve model performance for a variety of tasks (Brown et al., 2020; Alves et al., 2023). We select 50 representative instruction-command pairs from our training dataset as ICL examples. We create embeddings for the commands using the mx-bai-embed-large-v1 model and cluster the embeddings using k-means clustering. The closest instruction-command pair to each centroid is selected as an ICL example. We append these pairs to our translation prompt as show in Figure 9. We evaluate the performance of Llama3.1-8b-Instruct with the number of appended pairs ranging from 1-50 and find the optimal number to be 25, with performance saturating as more pairs are added. We use 25 example instruction-command pairs for all ICL evaluations.

4.3.4 In-Weight Learning

We use our training dataset to perform a LoRA fine-tune of the Llama and Qwen models (Hu et al., 2021). We experiment with common hyperparameters within our hardware constraint of a single Nvidia RTX A6000. We find that training each model for 10 epochs with an adapter rank of 64,

adapter alpha of 32, adapter dropout of 0.1, batch size of 32 and learning rate of $1e - 5$ results in the best performance. We do not fine-tune the GPT models due to financial constraints and the inability to control training hyper-parameters.

5 Experiments

5.1 Evaluation of FEHs

We compare our FEH with the heuristics presented by Sparck Jones (1988), Papineni et al. (2002), Agarwal et al. (2021), Yang et al. (2023) and Song et al. (2024) in previous work using our test dataset. A FEH should return true given two functionally equivalent Bash commands, and false given two non-equivalent Bash commands. We record the precision, recall, F1 score and accuracy of each FEH and report our results in Table 4.

Our test dataset, described in Section 4.1, is structured {nl, bash, bash2}, providing 300 pairs of functionally equivalent commands. To create a set of non-equivalent commands, we arbitrarily rotate the third column of the dataset by ten positions. The result is 600 pairs of Bash commands. Explicitly, each FEH is tested using 300 functionally equivalent pairs $m(t, g(b, e), g(\hat{b}, e)) = 1$ where $g(b, e) \approx g(\hat{b}, e)$ and 300 non-equivalent pairs $m(t, g(b, e), g(\hat{b}, e)) = 0$ where $g(b, e) \neq g(\hat{b}, e)$.

For the bleu and nl2cmd FEHs we use a threshold of 0.75 for functional equivalence. For the tfidf and mxbai-embed FEHs we calculate the cosine similarity of the resulting embeddings and use a threshold of 0.75 for functional equivalence. For the llama-3.1-8b-inst, gpt-3.5-t-0125 and gpt-4-0613 FEHs we pass the tasks and commands to each model using the prompt in Figure 5. For the exec + tfidf and exec + mxbai-embed FEHs, we pass the stdout of command execution to the models, calculate the cosine similarity of the resulting embeddings and use a threshold of 0.75 for functional equivalence. Finally, for the exec + llama-3.1-8b-inst, exec + gpt-3.5-t-0125 and exec + gpt-4-0613 FEHs, we pass the tasks, commands and stdouts of command execution to each model using the prompt in Figure 6. We use a temperature of zero and a static seed value of 123 for all LLMs.

5.2 Evaluation of Translation Methods

We evaluate the impact of parsing, constrained decoding, in-context learning and in-weight learning on the NL2SH performance of the Llama, Qwen and GPT model families. All models are evaluated

Table 4: Evaluation of Bash functional equivalence heuristics. Heuristics were tested on a dataset comprising 300 pairs of equivalent commands and 300 pairs of non-equivalent commands. We find execution paired with LLM evaluation significantly increases recall. Bold indicates the highest F1 score and accuracy.

Heuristic	Prec.	Rec.	F1	Acc.
bleu (Papineni et al., 2002)	0.99	0.39	0.56	0.69
nl2cmd (Agarwal et al., 2021)	0.98	0.20	0.33	0.60
tfidf (Sparck Jones, 1988)	0.99	0.46	0.63	0.73
exec + tfidf (Yang et al., 2023)	0.99	0.65	0.79	0.82
mxbai-embed (Lee et al., 2024)	0.84	0.88	0.86	0.85
exec + mxbai-embed (Ours)	0.97	0.83	0.90	0.90
llama-3.1-8b-inst (Llama, 2024)	1.00	0.05	0.10	0.53
exec + llama-3.1-8b-inst (Ours)	0.88	0.74	0.80	0.82
gpt-3.5-t-0125 (Brown et al., 2020)	1.00	0.37	0.54	0.69
exec + gpt-3.5-t-0125 (Ours)	0.98	0.60	0.75	0.80
gpt-4-0613 (Song et al., 2024)	1.00	0.51	0.68	0.76
exec + gpt-4-0613 (Ours)	0.99	0.91	0.95	0.95

using version 0.3.6 of the InterCode-ALFA benchmark with the execution + mxbai-embed FEH. Our results are summarized in Table 5.

We use the prompt in Figure 7 for the baseline evaluation. For the constrained decoding evaluation, we use the prompt in Figure 8 and constrain the first tokens of the model output to a list of Bash utilities, as described in Section 4.3.2. For the parser evaluation, we use the prompt in Figure 8 and pass model outputs to a markdown parser, as described in Section 4.3.1. For the in-context learning evaluation, we use the prompt in Figure 9. Finally, for the in-weight learning evaluation, we use the prompt in Figure 8 and the fine-tuned models described in Section 4.3.4. We use a temperature of zero and a static seed value of 123 for all model evaluations.

Table 5: Impact of constrained decoding, parsing, in-context learning and in-weight learning on the NL2SH performance of Llama, Qwen and GPT models. Accuracy is measured using the exec + mxbai-embed FEH. Translation method can improve performance up to 32% over the baseline, but model size remains the dominant factor. The highest accuracy in each row and column is indicated with bold and an underline, respectively.

Model	Base	CD	Parse	ICL	IWL
llama-3.2-1b-instruct	0.12	0.19	0.32	0.34	0.37
llama-3.2-3b-instruct	0.17	0.39	0.49	0.47	0.47
llama-3.1-8b-instruct	0.46	<u>0.51</u>	0.53	0.56	0.40
qwen2.5-coder-0.5b-instruct	0.10	0.05	0.35	0.36	0.27
qwen2.5-coder-1.5b-instruct	0.21	0.06	0.50	0.44	0.19
qwen2.5-coder-3b-instruct	0.26	0.06	0.58	0.50	0.51
qwen2.5-coder-7b-instruct	0.61	0.08	0.62	0.62	<u>0.51</u>
gpt-3.5-turbo-0125	0.58	-	0.67	0.69	-
gpt-4o-mini-2024-07-18	0.71	-	0.72	0.71	-
gpt-4o-2024-08-06	0.74	-	<u>0.73</u>	<u>0.73</u>	-
gpt-4-0613	0.68	-	0.68	<u>0.73</u>	-

6 Discussion

6.1 Dataset

Our first research question aims to validate NL2SH datasets to ensure models are evaluated using valid translations. Manual verification of the InterCode dataset identified over half of the instruction-command pairs as erroneous. We find human verification of data is important for reliable evaluations.

Manual creation and verification of our test dataset took over 100 hours, highlighting the need for more efficient means to verify larger datasets. Further, since the InterCode dataset is sampled from the NL2Bash dataset, there is a risk the NL2Bash dataset contains a significant percentage of erroneous data. This is concerning because the NL2Bash dataset is commonly used to train NL2SH models (Fu et al., 2021; Lin, 2017; Shi et al., 2023; Bharadwaj et al., 2022; Joshi, 2024).

We are confident our training dataset contains valid data due to our filtering process to remove invalid Bash commands. Moreover, fine-tuning the Llama and Qwen models using our dataset results in an average performance increase of 11%.

6.2 Functional Equivalence Heuristic

Our second research question aims to design an FEH that accurately measures the quality of model translations. We find that command execution paired with LLM evaluation of command outputs can determine the functional equivalence of Bash commands with 95% accuracy. The ability of LLMs to condition command outputs on natural language inputs is an advancement that was not possible with previous heuristics. Further, command execution improves performance across methods and the use of an LLM significantly increases recall. Broadly, LLM evaluation of execution outputs is a promising advancement for measuring the functional correctness of generated code and more investigation is warranted.

In accordance with Maveli et al. (2024), we find that without execution, current LLMs are poor arbiters of functional equivalence, achieving similar performance when compared to conventional evaluation methods. Non-execution methods likely fail because two commands can be syntactically similar and yield different results when executed. For example, changing a single flag can result in vastly different command outputs. Further, two commands with no syntactic similarity can yield identical results when executed. For example, the

awk and sed utilities can accomplish identical text processing tasks but use different domain-specific languages, requiring different syntax. Notably, the low recall of bleu and nl2cmd FEHs indicates these methods cannot identify cases where syntactically different commands are functionally equivalent.

6.3 Translation Methods

Our third research question aims to improve the accuracy of NL2SH models as measured by a reliable benchmark. We find that constrained decoding, parsing, in-context learning and in-weight learning can improve model performance by up to 32%. Our baseline evaluation shows that model performance is correlated with number of parameters.

We find that constrained decoding is model dependent, with performance increases for Llama models and significant performance decreases for Qwen models. Parsing and ICL provide performance increases across Llama and Qwen models, with average increases of 21% and 19%, respectively. However, these methods have a decreasing impact as model size increases. This is evidence that incorrect output format is the dominant failure mode for small (less than 7b parameter) models.

With IWL, llama-3.2-3b-instruct and qwen2.5-coder-0.5b-instruct achieve the baseline performance of llama-3.1-8b-instruct and qwen2.5-coder-3b-instruct, respectively. Despite performance gains for small models, fine-tuning decreases the performance of llama-3.1-8b-instruct and qwen2.5-coder-7b-instruct. This is likely due to computational constraints on the size of our LoRA adapters, which we are unable to scale with model size.

We find that gpt-4o-2024-08-06 achieves SOTA performance on our benchmark, correctly translating 74% of test cases. From the total of our experiences, we find that NL2SH translation is a difficult task for current models, necessitating improvements before models can be used in practice.

7 Conclusion

In this paper, we explore applications for LLMs in NL2SH translation and benchmarking. We identify issues with current benchmarks, including inaccurate datasets and unreliable functional equivalence heuristics. To address these problems, we correct and expand NL2SH datasets and create a new heuristic to determine the functional equivalence of Bash commands. We assess our heuristic and find that Bash command execution paired with

language model evaluation of command outputs can determine the functional equivalence of commands more accurately than previous heuristics. Using our dataset and heuristic, we evaluate how constrained decoding, parsing, in-context learning and in-weight learning impact the performance of Llama, Qwen and GPT models. We find that parsing and in-context learning reliably improve the performance of open and closed-source LLMs for the task of NL2SH translation. Ultimately, we find that NL2SH translation is a difficult task for LLMs, necessitating further research in this field. In future work, we plan to investigate efficient means to verify our training dataset and conduct more fine-tuning experiments.

8 Limitations

This work presents a verified and expanded NL2SH test dataset. However, due to the time and effort required to manually configure an execution environment for each command, the dataset remains small, with only 600 test cases. In contrast, our training dataset is too large for manual verification, and we are unable to guarantee its correctness. Our datasets are limited to English prompts and one-line Bash commands. We do not consider other natural languages or scripting languages.

Although improved over previous methods, our functional equivalence heuristic has inherent variability due to the use of an LLM, requiring multiple runs to assess model performance. The use of an LLM also increases the computational cost of running our heuristic compared to conventional methods. Finally, despite improved model performance with constrained decoding, parsing, ICL and IWL, the accuracy of SOTA LLMs for NL2SH translation remains low, motivating further research.

9 Ethical Considerations

Due to the low performance of current NL2SH models, using these models in practice could result in invalid commands that have unintended effects on a system. We recommend that model-generated commands are never used without human verification. Further, we recommend that users test commands using a sand-boxed environment, such as try (Zhu et al., 2024a), before running them on personal systems. Figure 4 shows an example of a dangerous command observed during model testing. The hallucinated `rm -f /dev/null` command corrupted our benchmark’s Docker container. While

our benchmark automatically creates a new Docker container to handle this type of error, the command could have corrupted a user’s system.

Natural Language Prompt:
Recursively remove all empty folders from the /system/temp folder.

Ground Truth Command:
`find /system/temp -type d -empty -delete`

Llama3.1-8b-Instruct Command:
`find /system/temp -type d -empty -delete -print; rm -f /dev/null 2>&1`

Figure 4: Dangerous translation observed in testing.

Natural language to Bash translation aims to increase computer accessibility by simplifying interactions with the command-line interface. Unfortunately, good and bad actors can benefit from increased accessibility. Models could be used to generate malicious Bash commands. This risk is difficult to mitigate because malicious use depends on the intent of the user. For example, a command to delete files could be used for a legitimate or harmful purpose. We do not believe current NL2SH models pose any risks beyond those of other readily available Bash resources.

References

- Mayank Agarwal, Tathagata Chakraborti, Quchen Fu, David Gros, Xi Victoria Lin, Jaron Maene, Kartik Talamadupula, Zhongwei Teng, and Jules White. 2021. [Neurips 2020 nlc2cmd competition: Translating natural language to bash commands](#). In *Proceedings of the NeurIPS 2020 Competition and Demonstration Track*, volume 133 of *Proceedings of Machine Learning Research*, pages 302–324. PMLR.
- Pooja Aggarwal, Oishik Chatterjee, Ting Dai, Prateeti Mohapatra, Brent Paulovicks, Brad Blancett, and Arthur De Magalhaes. 2024. [CodeSift: An LLM-Based Reference-Less Framework for Automatic Code Validation](#). In *2024 IEEE 17th International Conference on Cloud Computing (CLOUD)*, pages 404–410, Los Alamitos, CA, USA. IEEE Computer Society.
- Duarte Alves, Nuno Guerreiro, João Alves, José Pomal, Ricardo Rei, José de Souza, Pierre Colombo, and Andre Martins. 2023. [Steering large language models for machine translation with finetuning and in-context learning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 11127–11148, Singapore. Association for Computational Linguistics.
- Shikhar Bharadwaj, Shirish Shevade, and Marine Carpuat. 2022. [Efficient constituency tree based encoding for natural language to bash translation](#). In *Proceedings of the 2022 Conference of the North*

- American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3159–3168, Seattle, United States. Association for Computational Linguistics.
- Richard Blum and Christine Bresnahan. 2021. *Linux Command Line and Shell Scripting Bible*. Wiley.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Federico Cassano, John Gouwar, Daniel Nguyen, Sydney Nguyen, Luna Phipps-Costin, Donald Pinckney, Ming-Ho Yee, Yangtian Zi, Carolyn Jane Anderson, Molly Q Feldman, Arjun Guha, Michael Greenberg, and Abhinav Jangda. 2023. [MultiPL-E: A Scalable and Polyglot Approach to Benchmarking Neural Code Generation](#). *IEEE Transactions on Software Engineering*, 49(07):3675–3691.
- Oishik Chatterjee, Pooja Aggarwal, Suranjana Samanta, Ting Dai, Prateeti Mohapatra, Debanjana Kar, Ruchi Mahindru, Steve Barbieri, Eugen Postea, Brad Blancett, and Arthur De Magalhaes. 2024. [Scriptsmith: A unified llm framework for enhancing it operations via automated bash script generation, assessment, and refinement](#). *Preprint*, arXiv:2409.17166.
- Liguo Chen, Qi Guo, Hongrui Jia, Zhengran Zeng, Xin Wang, Yijiang Xu, Jian Wu, Yidong Wang, Qing Gao, Jindong Wang, Wei Ye, and Shikun Zhang. 2024. [A survey on evaluating large language models in code generation tasks](#). *Preprint*, arXiv:2408.16498.
- Heejae Chon, Seonghyeon Lee, Jinyoung Yeo, and Dongha Lee. 2024. [Is functional correctness enough to evaluate code language models? exploring diversity of generated codes](#). *Preprint*, arXiv:2408.14504.
- Berkeley Churchill, Oded Padon, Rahul Sharma, and Alex Aiken. 2019. [Semantic program alignment for equivalence checking](#). In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2019*, pages 1027–1040, New York, NY, USA. Association for Computing Machinery.
- Quchen Fu, Zhongwei Teng, Marco Georgaklis, Jules White, and Douglas Schmidt. 2023. [Nl2cmd: An updated workflow for natural language to bash commands translation](#). *Journal of Machine Learning Theory, Applications and Practice*, 1.
- Quchen Fu, Zhongwei Teng, Jules White, and Douglas C. Schmidt. 2021. [A transformer-based approach for translating natural language to bash commands](#). In *2021 20th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 1245–1248.
- Saibo Geng, Martin Josifoski, Maxime Peyrard, and Robert West. 2023. [Grammar-constrained decoding for structured NLP tasks without finetuning](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10932–10952, Singapore. Association for Computational Linguistics.
- David Gros. 2019. [Ainix: An open platform for natural language interfaces to shell commands](#). Undergraduate Honors Thesis, Computer Science Department, University of Texas at Austin.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#). *Preprint*, arXiv:2106.09685.
- Junjie Huang, Chenglong Wang, Jipeng Zhang, Cong Yan, Haotian Cui, Jeevana Priya Inala, Colin Clement, and Nan Duan. 2022. [Execution-based evaluation for data science code generation models](#). In *Proceedings of the Fourth Workshop on Data Science with Human-in-the-Loop (Language Advances)*, pages 28–36, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Binyuan Hui, Jian Yang, Zeyu Cui, Jiayi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajuan Zhang, Bowen Yu, Kai Dang, et al. 2024. [Qwen2. 5-coder technical report](#). *arXiv preprint arXiv:2409.12186*.
- Daniel Jenson and Yingxiao Liu. 2022. [Translating natural language to bash commands using deep neural networks](#). https://web.stanford.edu/class/archive/cs/cs224n/cs224n.1224/reports/custom_116997097.pdf. Accessed: 2024-10-09.
- Sathvik Joel, Jie JW Wu, and Fatemeh H. Fard. 2024. [A survey on llm-based code generation for low-resource and domain-specific programming languages](#). *Preprint*, arXiv:2410.03981.
- Anish Joshi. 2024. [Bash scripting assistant](#). <https://github.com/AnishJoshi13/Bash-Scripting-Assistant>. Accessed: 2024-10-09.
- Idan Kamara. 2016. [bashlex: Python parser for bash](#). <https://github.com/idank/bashlex>. Accessed: 2024-09-09.
- Michael Kerrisk. 2024. [Linux manual pages](#). <https://man7.org/linux/man-pages/index.html>. Accessed: 2024-09-09.
- K.B.Dharun Krishna, Sebastiaan Speck, Owen Voke, and Darío Herenu. 2024. [Tldr: Collaborative cheat-sheets for console commands](#). <https://github.com/tldr-pages/tldr>. Accessed: 2024-10-09.

- Sean Lee, Aamir Shakir, Darius Koenig, and Julius Lipp. 2024. [Open source strikes bread - new fluffy embeddings model](#).
- Xi Victoria Lin. 2017. [Program synthesis from natural language using recurrent neural networks](#).
- Xi Victoria Lin, Chenglong Wang, Luke Zettlemoyer, and Michael D. Ernst. 2018. [NL2Bash: A corpus and semantic parser for natural language interface to the linux operating system](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Meta Llama. 2024. [The llama 3 herd of models](#). Preprint, arXiv:2407.21783.
- Zach Lloyd, Michelle Lim, and Alope Desai. 2024. [Warp: Your terminal, reimagined](#). <https://www.warp.dev/>. Accessed: 2024-09-09.
- Yogesh Mali. 2023. [text_to_bash](#).
- Nickil Maveli, Antonio Vergari, and Shay B. Cohen. 2024. [What can large language models capture about code functional equivalence?](#) Preprint, arXiv:2408.11081.
- Dirk Merkel. 2014. Docker: lightweight linux containers for consistent development and deployment. *Linux journal*, 2014(239):2.
- Microsoft. 2024. [Microsoft ai shell](#). <https://learn.microsoft.com/en-us/powershell/utility-modules/aishell/overview?view=ps-modules>.
- Atharva Naik. 2024. [On the limitations of embedding based methods for measuring functional correctness for code generation](#). Preprint, arXiv:2405.01580.
- OpenAI. 2024. [Gpt-4 technical report](#). Preprint, arXiv:2303.08774.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics, ACL '02*, pages 311–318, USA. Association for Computational Linguistics.
- T. Ramesh. 2022. [NI2cmd](#). Accessed: 2024-08-20.
- Chet Ramey and Brian Fox. 2024. [Bash: GNU Project's Shell](#). Accessed: 2024-09-09.
- Romit. 2024. [Linuxcommands](#). Accessed: 2024-08-20.
- Johan Rosenkilde, Matt Rothenberg, and Andy Feller. 2024. [Github copilot for cli](#). <https://githubnext.com/projects/copilot-cli/>.
- Farkhod Sadykov. 2024. [Shellgpt: A command-line productivity tool powered by ai large language models like gpt-4](#). https://github.com/TheR1D/shell_gpt. Accessed: 2024-09-09.
- Jean E. Sammet. 1966. [The use of english as a programming language](#). *Commun. ACM*, 9(3):228–230.
- Amazon Web Services. 2024. [Amazon codewhisperer natural language to bash translation](#). <https://docs.aws.amazon.com/codewhisperer/latest/userguide/command-line-conversation.html>.
- Danil Shaikhelislamov, Mikhail Drobyshevskiy, and Andrey Belevantsev. 2024. [Codepatchllm: Configuring code generation using a static analyzer](#).
- Jie Shi, Sihang Jiang, Bo Xu, Jiaqing Liang, Yanghua Xiao, and Wei Wang. 2023. [Shellgpt: Generative pre-trained transformer model for shell language understanding](#). In *2023 IEEE 34th International Symposium on Software Reliability Engineering (ISSRE)*, pages 671–682.
- William E Shotts. 2019. *The Linux Command Line: A Complete Introduction*. No Starch Press.
- Yewei Song, Cedric Lothritz, Xunzhu Tang, Tegawendé Bissyandé, and Jacques Klein. 2024. [Revisiting code similarity evaluation with abstract syntax tree edit distance](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 38–46, Bangkok, Thailand. Association for Computational Linguistics.
- Karen Sparck Jones. 1988. *A statistical interpretation of term specificity and its application in retrieval*, pages 132–142. Taylor Graham Publishing, GBR.
- Prashanth Vijayaraghavan, Luyao Shi, Stefano Ambrogio, Charles Mackin, Apoorva Nitsure, David Beymer, and Ehsan Degan. 2024. [Vhdl-eval: A framework for evaluating large language models in vhdl code generation](#). Preprint, arXiv:2406.04379.
- Ngoc Phuoc An Vo, Brent Paulovicks, and Vadim Sheinin. 2024. [Execution-based evaluation of natural language to bash and powershell for incident remediation](#). Preprint, arXiv:2405.06807.
- John Yang, Akshara Prabhakar, Karthik Narasimhan, and Shunyu Yao. 2023. [Intercode: standardizing and benchmarking interactive coding with execution feedback](#). In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23*, Red Hook, NY, USA. Curran Associates Inc.
- Ezri Zhu, Georgios Liargkovas, Michael Greenberg, and Konstantinos Kallas. 2024a. [try utility](#). <https://github.com/binpash/try>. Accessed: 2024-09-25.
- Wenhao Zhu, Hongyi Liu, Qingxiu Dong, Jingjing Xu, Shujian Huang, Lingpeng Kong, Jiajun Chen, and Lei Li. 2024b. [Multilingual machine translation with large language models: Empirical results and analysis](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2765–2781, Mexico City, Mexico. Association for Computational Linguistics.

Appendix A. Prompts

Functional Equivalence Heuristic Prompt: LLM

You will be given a task and two Bash commands. The first command is the ground truth. If the second command accomplishes the task, return true. Otherwise, return false. Only output 'true' or 'false'. Task: natural_language_prompt, Ground Truth Command: ground_truth_command, Model Command: model_command.

Figure 5: Prompt for evaluating the functional equivalence of Bash commands.

Functional Equivalence Heuristic Prompt: Execution + LLM

You will be given a task, two Bash commands, and the output of the two Bash commands. The first command is the ground truth. If the second command accomplishes the task, return true. Otherwise, return false. Only output 'true' or 'false'. Task: natural_language_prompt, Ground Truth Command: ground_truth_command, Model Command: model_command, Ground Truth Command Output: ground_truth_command_output, Model Command Output: model_command_output.

Figure 6: Prompt for evaluating the functional equivalence of Bash commands after execution. Note the addition of command outputs compared to the prompt in Figure 5.

Translation Prompt: Baseline

Your task is to translate a natural language instruction to a Bash command. You will receive an instruction in English and output a Bash command that can be run in a Linux terminal. You will not output markdown or other formatting. You will not include additional information. natural_language_prompt

Figure 7: NL2SH translation prompt used in the baseline evaluation.

Translation Prompt: Parser, Constrained Decoding and In-Weight Learning

Your task is to translate a natural language instruction to a Bash command. You will receive an instruction in English and output a Bash command that can be run in a Linux terminal. natural_language_prompt

Figure 8: NL2SH translation prompt used in the parsing, constrained decoding and in-weight learning evaluations.

Translation Prompt: In-Context Learning

Your task is to translate a natural language instruction to a Bash command. You will receive an instruction in English and output a Bash command that can be run in a Linux terminal.

Show logged-in users info
w

Print the contents of "xx.sh"
cat xx.sh

Change owner to "root" and group to "www-data" of "/foobar/test_file"
chown root:www-data /foobar/test_file

delete all the text files in the current folder
find . -type f -name "*.txt" -delete

find all the files in the /path folder and delete them
find /path -type f -delete

Print the exit status of the last executed command
echo \$?

Display a tree of processes
pstree

Display information about all CPUs
lscpu

Make an HTTPS GET request to example.com and dump the contents in 'stdout'
curl https://example.com

Display system memory
free

List all files, including hidden files
ls -a

Print a sequence from 1 to 10
seq 10

Get the properties of all the user limits
ulimit -a

List the name and status of all services
service --status-all

Display a calendar for the current month
cal

Show the environment
env

create directory TestProject
mkdir TestProject

Query the default name server for the IP address of example.com
nslookup example.com

Print Hello World
echo "Hello World"

List all bound commands and their hotkeys
bind -p

Display the openssl version
openssl version

Print current time, uptime, number of logged-in users
uptime

Print file system disk space usage
df

List all configuration values available
getconf -a

Delete empty folder 'nonsense_dir'.
rmdir nonsense_dir

natural_language_prompt

Figure 9: NL2SH translation prompt used in the in-context learning evaluation.