

# ConCodeEval: Evaluating Large Language Models for Code Constraints in Domain-Specific Languages

Mehant Kammakomati<sup>\*1</sup>, Sameer Pimparkhede<sup>\*2</sup>, Srikanth G. Tamilselvam<sup>1</sup>, Prince Kumar<sup>1</sup>, and Pushpak Bhattacharyya<sup>2</sup>

<sup>1</sup>IBM Research

<sup>2</sup>IIT Bombay

<sup>1</sup>{mehant.kammakomati2,prince.kumar12}@ibm.com

<sup>1</sup>{srikanth.tamilselvam}@in.ibm.com

<sup>2</sup>{sameerp,pb}@cse.iitb.ac.in

## Abstract

System-level programming is essential for modern enterprise infrastructure, enabling the automation and management of complex systems through declarative code. Developers write this code based on schemas, which themselves are a form of code that defines constraints like data types and required fields. These schemas help ensure operational correctness and smooth integration across systems. However, as enterprise schemas become complex, manually writing code adhering to these constraints becomes challenging for developers. Large Language Models (LLMs) have demonstrated potential in code generation and natural language understanding, particularly in zero-shot and few-shot settings. However, applying LLMs to handle constraints represented in code, essential for system-level programming rather than natural language, has not been explored. Hence, we introduce ConCodeEval, a study across two key dimensions: format and constraint efficacy, with a first-of-its-kind benchmark involving two novel experiments for code constraints across five representations (JSON, YAML, XML, Python, and natural language). Our findings suggest that conscious choice of representations can lead to optimal use of LLMs in enterprise use cases involving constraints. Nonetheless, LLMs continue to struggle significantly with code constraints, motivating the need for innovation in this direction.

## 1 Introduction

System-level programming is the backbone of modern enterprise infrastructure, enabling developers to define, manage, and automate complex systems seamlessly. Numerous enterprises use concepts like Infrastructure as Code<sup>1</sup> (IaC) to let developers write declarative code. Such code must adhere to constraints called schemas, which define

<sup>\*</sup>The first two authors contribute equally.

<sup>1</sup>[https://en.wikipedia.org/wiki/Infrastructure\\_as\\_code](https://en.wikipedia.org/wiki/Infrastructure_as_code)

rules, including data types, required fields, and valid value ranges, ensuring operational correctness and smooth integration. For instance, the schema in Listing 1 mandates an array of even numbers within specific bounds, containing 1 to 7 elements.

Listing 1: The JSON sample generated (highlighted in yellow) by the Granite 20B model does not adhere to the *minContains* and subsequent numerical constraints specified in the schema.

```
Write a JSON sample with field values as per the
JSON format schema given below.
```

```
{
  "type": "array",
  "contains": {
    "type": "number",
    "multipleOf": 2,
    "exclusiveMinimum": 0,
    "exclusiveMaximum": 65535
  },
  "minContains": 1,
  "maxContains": 7
}
```

```
JSON sample:
```

```
...
[2, 3, 4, 6, 8, 10, 12, 14]
..
```

Schemas are crucial in real-world enterprise settings. For instance, deploying a database service in an OpenShift cluster involves writing compliant code with the correct attributes, such as the number of instances, port number to expose, compute to allocate, etc. Developers write system-level code in structured Domain Specific Languages (DSLs) such as JSON, YAML, XML, or Python, adhering to strict schema constraints. However, enterprise schemas are often complex and difficult to learn, slowing development and increasing errors. As a result, the need for automated and accurate systems for system-level programming is increasing leading to products such as Ansible Lightspeed (Lig).

LLMs have shown great promise in generating coherent text and code in zero-shot and few-shot settings, making them highly appealing for system-level coding (Brown et al., 2020; Roziere et al.,

2023; Mishra et al., 2024). Using LLMs to handle constraints represented in natural language (NL) has been extensively explored for tasks like poem generation and summarization (Sun et al., 2023). However, Unlike these natural language tasks, constraints are often represented as code for system-level programming; hence, evaluating LLMs requires a different approach. In addition to assessing how well models adhere to constraints expressed in natural language, we must examine their ability to process, interpret, and generate structured formats while ensuring schema compliance. To ensure this, we evaluate LLMs under two key dimensions: **Format Efficacy** and **Constraint Efficacy**.

Format efficacy involves studying the performance of LLMs on varying constraint representations that form the input and output representations downstream enterprise use cases can consume. Specifically, we aim to answer the following research questions (RQ) for format efficacy: 1) Which format is optimally suited for constraint and output representation? 2) What is the trade-off between performance and context length cost?. While constraint efficacy involves studying LLMs’ performance on various schema constraints within a format. Precisely, we aim to answer the following research questions related to constraint efficacy: 1) How does performance vary across different types of constraints? 2) What are the ideal positions for constraints in the schema for better adherence?

We prepare first-of-its-kind benchmark test set<sup>2</sup> and conduct two experiments involving 5 schema formats (JSON, YAML, XML, Python, and NL) and 3 output formats (JSON, YAML, and XML) resulting 15 combinations of use cases to investigate the aforementioned research questions. 1) Data as Code Generation (Section 2.1. 2) Data Validation (Section 2.2). This study provides insights into leveraging LLMs effectively for system-level programming tasks involving code constraints in enterprises.

**Our contributions are:**

1. First-of-its-kind study of language models for crucial industry use case of system-level programming involving code format constraints across four key dimensions: Format and Constraint efficacy.
2. A benchmark test set consisting of 602 schema

<sup>2</sup>Dataset is made available at <https://hf.co/datasets/kmehant/concodeeval>

samples, each containing multiple instructions. Each schema sample in our test set is represented in 5 different language formats (JSON, YAML, XML, Python, and NL).

3. Comparative and qualitative analysis of state-of-the-art language models involving code generation from fine-grained schema instructions and code validation against schemas. To the best of our knowledge, we are the first to evaluate LLMs code constraint competency.

## 2 Experiments

### 2.1 Data as Code Generation in DSL




Figure 1: Uniform trend of steep decline in performance across models for constraints positioned in the middle and beginning of the JSON schema context and output for data as code generation experiment. We divide the schema into 3 portions, Begin, Middle, and End, and put the violated constraints based on their locality into either of these three buckets.

**Description.** Given the schema, the experiment (see Listing 1) aims to produce a compliant data sample in DSL code format. We draw inspiration from several use cases (see Appendix A.3), including synthesizing schema-compliant data from LLMs’ parametric memory to train and evaluate smaller-sized models (Song et al., 2020) and generating diverse sets of samples to be used in product test pipelines. For reliable DSL code generation, LLMs need to be schema-aware.

**Dataset.** We synthetically prepare 602 schemas for each of the 5 representations having combinations of various constraints (Appendix A.4). First, we prepare JSON schemas using our combinatorial tool to generate a good mix of constraints. A combinatorial data generation tool factors in constraints of interest, constraint-specific information, and combinatorial preferences to generate the schemas. We then convert each JSON schema

Model	Schema	Output Representation					
		JSON		YAML		XML	
		Gen Acc	Val Acc	Gen Acc	Val Acc	Gen Acc	Val Acc
Llama3 8B	JSON	28.2	56.0	29.2	45.0	7.9	47.0
Granite 8B		47.5	56.0	24.7	55.0	5.1	45.0
Granite 20B		50.4	52.0	37.7	44.0	10.1	53.0
Granite 34B		53.3	64.0	32.2	57.0	11.2	<b>65.0</b>
Codellama 34B		58.4	64.0	23.0	54.0	9.4	53.0
🏆 Llama3 70B		<b>62.8</b>	<b>67.0</b>	<b>40.1</b>	<b>58.4</b>	<b>18.9</b>	55.7
Llama3 8B	XML	10.2	37.0	22.5	42.0	10.2	46.0
Granite 8B		18.9	47.0	12.1	44.0	8.4	52.0
Granite 20B		24.0	37.0	12.4	47.0	8.6	57.0
Granite 34B		18.7	68.0	18.1	58.0	8.6	<b>58.0</b>
Codellama 34B		8.8	46.0	14.2	46.0	8.6	50.0
🏆 Llama3 70B		<b>28.4</b>	<b>70.3</b>	<b>24.8</b>	<b>60.1</b>	<b>16.6</b>	54.2
Llama3 8B	YAML	25.9	46.0	8.1	44.0	6.4	45.0
Granite 8B		47.0	47.0	15.7	50.0	8.6	44.0
Granite 20B		34.7	31.0	25.9	38.0	8.4	47.0
Granite 34B		52.1	68.0	26.4	61.0	8.6	<b>58.0</b>
Codellama 34B		48.0	59.0	27.9	53.0	9.1	<b>58.0</b>
🏆 Llama3 70B		<b>56.0</b>	<b>71.0</b>	<b>32.4</b>	<b>63.2</b>	<b>14.6</b>	56.9
Llama3 8B	Python	13.7	43.0	10.2	42.0	<b>11.6</b>	43.0
Granite 8B		10.2	54.0	11.9	58.0	11.1	<b>55.0</b>
Granite 20B		14.6	45.0	11.7	67.0	7.3	44.0
Granite 34B		17.7	54.0	13.9	67.0	10.6	46.0
Codellama 34B		13.7	49.0	11.6	53.0	8.4	44.0
🏆 Llama3 70B		<b>24.7</b>	<b>57.2</b>	<b>18.9</b>	<b>70.4</b>	<b>14.9</b>	52.1
Llama3 8B	NL	30.2	63.0	24.5	56.0	9.6	57.0
Granite 8B		52.3	59.0	42.1	61.0	11.1	58.0
Granite 20B		65.4	54.0	46.0	48.0	10.9	<b>60.0</b>
Granite 34B		69.7	55.0	55.1	46.0	10.9	56.0
Codellama 34B		60.4	57.0	40.6	57.0	8.69	50.0
🏆 Llama3 70B		<b>75.2</b>	<b>67.7</b>	<b>57.2</b>	<b>64.2</b>	<b>13.4</b>	58.1

Table 1: Zero shot results for both the experiments. Models scoring the highest accuracy the majority of times across all output representations for a particular schema are labeled with 🏆. Gen Acc represents the accuracy of valid samples for DSL generation experiment. Val Acc represents the accuracy of the binary classification validation experiment.

to XML and YAML schemas using openly available automatic lossless language-to-language translation tools. Further, we include resource-rich general-purpose language - Python using the Pydantic library generated using the Gemini-1.0-pro (Team et al., 2023) model as a code translation task. We extend our evaluation to NL representation generated using rule-based templates. We<sup>3</sup> ensure equivalence of the generated schemas across languages. We plan to open-source all the scripts used for data preparation. Table 3 gives details regarding schema token length.

**Evaluation metric.** Each schema-compliant code output LLM generates is awarded one point where schema compliance is checked using a schema validator tool. We then utilize the accuracy metric (Gen Acc) over all samples to benchmark performance across the models. Additionally, we also report the percentage of samples generated with the invalid root data type (RTV%) and invalid

<sup>3</sup>The schemas are manually validated by the paper’s authors.

samples (IS%) in Table 5. The root data type is the data type of the whole DSL sample. For example, the root data type of sample represented in Listing 1 is *array*. For IS and RTV metrics, the lesser the number, the better the performance.

**Experimental setup.** We report greedy decoding results since it performed slightly better than beam search with a beam width of 3. We perform inference for all the models in *bfloat16* precision and a max new token limit of 1024 tokens.

**Prompts.** We experiment with zero- and 3-shot prompting for each model. For 3-shot prompting, we identify errors from the zero-shot setting, then select shots similar to the most frequent errors. We observe that most errors made by all the models are regarding short schema and the schema having root type of array as shown in sample 1. An example of a 3-shot prompt for a DSL generation experiment is shown below. Examples of prompts are in Appendix 1.

		Output Representation					
		JSON		YAML		XML	
Model	Schema	Gen Acc	Val Acc	Gen Acc	Val Acc	Gen Acc	Val Acc
Llama3 8B	JSON	48.3	71.2	46.6	68.1	39.2	64.1
Granite 8B		51.2	69.2	52.3	66.1	47.8	65.8
Granite 20B		58.3	73.5	56.4	72.3	50.2	68.2
Granite 34B		66.3	76.2	64.5	75.4	51.3	73.2
Codellama 34B		65.1	75.1	63.4	73.2	50.6	71.2
🏆 Llama3 70B		<b>70.1</b>	<b>79.3</b>	<b>69.4</b>	<b>77.9</b>	<b>58.6</b>	<b>74.2</b>
Llama3 8B	XML	46.6	65.8	42.3	63.4	36.6	60.1
Granite 8B		46.2	64.8	44.5	63.2	34.5	57.3
Granite 20B		50.4	66.7	48.2	64.1	36.4	56.1
Granite 34B		52.3	68.5	51.1	63.4	39.2	53.2
Codellama 34B		49.2	66.2	49.2	63.2	35.1	52.1
🏆 Llama3 70B		<b>56.4</b>	<b>70.3</b>	<b>55.6</b>	<b>68.2</b>	<b>43.6</b>	<b>66.3</b>
Llama3 8B	YAML	46.7	67.2	45.3	64.2	43.5	63.2
Granite 8B		48.1	65.2	46.2	61.2	44.2	61.2
Granite 20B		52.3	68.9	49.7	66.7	47.8	65.1
Granite 34B		54.2	67.7	51.3	65.3	45.3	56.4
Codellama 34B		56.8	66.4	50.2	64.3	47.8	56.2
🏆 Llama3 70B		<b>60.4</b>	<b>76.3</b>	<b>57.3</b>	<b>69.1</b>	<b>49.6</b>	<b>68.3</b>
Llama3 8B	Python	43.2	60.1	41.1	58.9	39.2	57.6
Granite 8B		45.1	60.5	46.7	59.4	37.4	56.0
Granite 20B		48.2	57.2	45.9	57.8	38.4	58.2
Granite 34B		50.6	59.2	47.1	55.6	41.3	57.3
Codellama 34B		47.2	56.4	45.3	57.2	39.2	55.1
🏆 Llama3 70B		<b>56.2</b>	<b>65.1</b>	<b>50.7</b>	<b>64.2</b>	<b>43.4</b>	<b>60.6</b>

Table 2: Few shot results for generation (3 shots) and validation (2 shots) experiments. Models scoring the highest accuracy the majority number of times across all output representations for a particular schema are labeled with 🏆. Gen Acc represents the accuracy of valid samples for DSL generation experiment. Val Acc represents the accuracy of the binary classification validation experiment.

## 2.2 DSL Validation

**Description.** There is a growing body of work (Hada et al., 2024) on showing promising usage of LLMs as evaluators in many tasks. On similar lines, given the DSL sample and schema to validate, this experiment (see Listing 2) aims to determine the validity of the provided sample against the constraints through boolean question answering (QA). Also, the experiment is highly motivated from various use cases (see Appendix A.3) and throws light on LM’s understanding of the relation between requirements and output in various representations.

**Dataset.** We synthetically prepare 602 schemas across 5 representations having combinations of hard and soft constraints. First, we prepare JSON schemas using our combinatorial tool to generate a good mix of constraints. We then convert each JSON schema to XML and YAML schemas using automated tools to ensure equivalence across representations. Further, we include Python representation using the Pydantic library as a resource-rich general-purpose language in our evaluation generated using the Gemini-1.0-pro (Team et al., 2023) model as a code translation task. We extend our evaluation to natural language representation generated using rule-based templates over the JSON

schema. We<sup>4</sup> ensure equivalence of the generated schemas across languages by manually eyeballing the samples.

Listing 2: In the JSON sample, values for fields *stingo* and *anistic* do not adhere to schema constraints. But the Granite 34B model gives the incorrect answer (highlighted in yellow) as *yes*.

```

Question:
Does the JSON sample { "tamil": false, "baser": null
, "anistic": 1906.34, "stingo": "officiis tellus
. illum modi odit quas mattis nunc", "
pigheadedness": 52.0 } adhere to all the
constraints defined in JSON format schema
{
  "type": "object",
  "properties": {
    "tamil": { "type": "boolean" },
    "baser": { "type": "null" },
    "anistic": { "type": "number", "multipleOf": 17.0
2 },
    "stingo": { "type": "string", "maxLength": 20 },
    "pigheadedness": { "type": "number", "
exclusiveMinimum": 27.65410407394338, "
maximum": 93.85523810367313 } },
  "additionalProperties": false
}
Respond to yes or no.
Answer:
...
yes
"
```

**Evaluation metric.** Since it is a boolean QA experiment, we use Macro average F1 (see Table 6)

<sup>4</sup>The generated Python samples are manually validated by the paper’s authors.

and Accuracy (Val Acc) as evaluation metrics (see Table 1).

**Experimental setup.** The decoding strategy used here is similar to the data generation experiment as mentioned in Section 2.1. We perform inference in *bfloat16* precision and a max new token limit of 1024 tokens. For beam search decoding, we use the beam width of 3.

**Prompts.** The goal of this experiment is to answer *yes* or *no*. We experiment with zero- and few-shot prompting. With few shot prompting, we provide one example each of *yes* and *no* answers. Results for few-shot prompting and examples of prompts are given in Appendix (Table 2).

Language	Max schema tokens	Avg schema tokens
XML	3316	364.82
JSON	1954	208.23
YAML	1295	135.09

Table 3: Schema length comparison using Llama3 tokenizer

### 3 Format Efficacy

#### 3.1 Objective

To identify the most effective schema representation and output format for system-level programming while employing language models. Since schemas can be represented in various structured formats, including JSON, YAML, XML, Python, and even NL, determining which format best enables constraint adherence for language models while balancing context-length costs is critical.

#### 3.2 RQ1: Which format is optimally suited for constraint and output representation?

**Finding 1.** In the data as code generation experiment (section 2.1), models best understand (Table 1) NL across all outputs. At the same time, JSON and YAML schemas perform well (Table 2) for constraints in code despite their limited presence in pre-training data. Surprisingly, models struggle with constraints in Python, likely due to a bias toward generating general-purpose Python code rather than schema-specific patterns. In contrast, JSON and YAML schemas benefit from their rigid structures and alignment with schema-centric applications, making them easier for models to interpret.

**Finding 2.** Using the same schema and output representation does not always enhance performance. For instance, in Table 2, YAML as schema and JSON as output representation performed better than YAML for both representations.

**Finding 3.** Although NL representation excels in generation experiments, it degrades the validation performance of larger models like 70B. Like generation experiment, models perform sub-optimally when schema and output representations are the same. In line with the first experiment, XML stands as a challenging language for models. The Llama3 70B model performs best in validation as in the first experiment, with other models hovering around 50% Val Acc, likely reflecting the random choice given the binary nature of the experiment. Smaller models, particularly the Llama3-8B with natural language representation, show notable improvement, as its pre-training combines NL and code.

**Key takeaway.** NL is a favorable language for schema representation, however, since its possible that enterprises lean more toward structured languages for better interoperability in which case JSON and YAML are ideal candidates for schema representation with JSON being favourable candidate for output representation. Nonetheless, the inconsistency in performance across experiments and model sizes underscores need for better schema comprehension and improved training strategies for NLP tasks involving validation.

#### 3.3 RQ2: What is the trade-off between performance and context length cost?

**Findings.** From section 3.2 key takeaway, JSON and YAML are ideal candidates for schema representation which form the context to the LLM. From Table 3, representing schema in YAML on an average takes  $\sim 35\%$  less tokens than JSON. However, while choosing YAML would mean taking a drop of  $\sim 14\%$  in Gen Acc and  $\sim 4\%$  in Val Acc performance compared to JSON.

**Key takeaway.** Enterprises should be cognizant of such tradeoff and choose ideal representation that fits their use case. Further, better tokenizer training techniques might lead to lower token expenditure for the desired representation.

Constraint	Llama3 8B	Llama3 70B
type	302	49
exclusiveMinimum	18	44
multipleOf	170	42
minLength	47	21
contains	22	12
exclusiveMaximum	22	12
maximum	11	2
maxLength	7	19
additionalProperties	4	0
minimum	4	15

Table 4: Both the models, least and best performing, irrespective of their performance, show a similar distribution of mistakes for each constraint.

## 4 Constraint Efficacy

### 4.1 Objective

To examine how language models handle various types of constraints embedded within schemas. Enterprise schemas enforce structural (e.g., required fields, data types) and semantic (e.g., dependencies, value constraints) rules.

### 4.2 RQ1: How does performance vary across different types of constraints?

**Findings.** The analysis of the results shows that LLaMA3 8B and 70B exhibit similar patterns of missing constraints when generating JSON samples from a given schema (Table 4). In particular, constraints such as *type*, *multipleOf*, and *exclusiveMinimum* are often missing, while constraints such as *maximum*, *additionalProperties*, and *minimum* are more frequently followed. The high error rate in fundamental constraints *type* can be because training data contains many JSON-like samples where *type* is implicit rather than explicitly stated. The reason behind missing constraints like *exclusiveMinimum* and *multipleOf* may be because they involve high numerical precision. LLMs treat numbers as tokens, leading to potential rounding errors or incorrect enforcement.

**Key takeaway.** LLMs struggle with numerical constraints underscoring need for better techniques throughout the stack from tokenizer to model training. For enterprises, a rudimentary solution is to integrate constrained decoding or use post-processing validation to correct missing constraints after generation.

### 4.3 RQ2: What are the ideal positions for constraints in the schema for better adherence?

**Findings.** We categorize the constraints of the schema into three sections based on tokens: beginning (first 30%), middle (next 40%), and end (last 30%). Later, we perform a needle-in-the-haystack experiment for the data-as-code generation. The heatmap in Figure 1 shows the statistics of constraints missed at every position for JSON to JSON generation. It reveals a consistent trend where models struggle the most with constraints positioned at the beginning of the schema, followed by the middle. In contrast, constraints at the end are least frequently missed. This suggests that models may prioritize constraints appearing later in the schema, likely due to the left-to-right decoding nature of autoregressive models, causing early constraints to be overwritten or ignored. We also observe that constraints in the middle position of the schema are frequently missed. This aligns with previous findings that the middle part of the long context is often missed (Liu et al., 2024). For the data validation task, we analyze attention maps, which reveal a similar trend where the model pays less attention to the middle part of the schema (Figure 2).

**Key takeaway.** This suggests that important constraints should be placed at the end of the schema or the beginning for longer schemas, depending on the use case.

## 5 Related Work

**Generation:** There is extensive work (Muenighoff et al., 2024; Cassano et al., 2023) on

evaluating capabilities of LLMs for various code tasks such as code completion, translation, etc, for resource-rich languages like Python. Despite there being work (Cassano et al., 2023) on multi-lingual code, there is scant attention to low-resource languages such as DSLs, though having crucial importance. One notable work (He et al., 2024), studies the bearing of prompt format in DSLs with LLM performance, however, does not include impact of output formats and controllability aspect in terms of code constraints crucial for enterprises. Further, using LLMs as evaluators for low-resource languages is gaining interest, however limited, mainly focusing on languages like XML and INI (Lian et al., 2023).

**Controllability of LLMs:** While LLMs can handle coarse-grained constraints like sentiment, they struggle with fine-grained constraints, such as ending a text with a specific word (Sun et al., 2023). Code schemas often require such fine-grained control, and to our knowledge, we are the first to explore LLM controllability for constraints in code.

## 6 Conclusion

We evaluate LLMs for system-level programming across two key dimensions: Format Efficacy and Constraint Efficacy. Format efficacy examines how LLMs handle different constraint formats, while constraint efficacy assesses their performance on various schema constraints within a format. We conduct two novel experiments to study these aspects: Data as Code generation and DSL validation. We evaluate LLMs across 5 schema (YAML, JSON, Python, XML, NL) and 3 output formats (YAML, JSON, XML). Our findings reveal that model performance does not directly correlate with a language's presence in pre-training data. JSON and YAML are best suited for system-level programming, and enterprises should convert Python and XML formats to one of these for better LLM performance. We also observe that schema constraint locality affects performance, with constraints in the start and middle being most frequently violated. Placing critical constraints at the end improves reliability. We hope our work drives innovation in improving LLM capabilities for crucial industry use case of system-level programming involving code constraints.

## 7 Limitations

While we explore the DSL validation task by generating *yes* or *no*, exploring the model's reasoning can give a more comprehensive analysis of LLM's understanding. Further, one can include more complex constraints in the future for general-purpose programming languages, like coding style constraints to write code along with natural language prompts and schema.

## Ethics Statement

Custom-created datasets have been created synthetically using open-source tools. The language models, tools, and frameworks used for evaluation are open source and can be used without copyright issues.

## References

- Ansible Lightspeed with IBM watsonx Code Assistant | Red Hat Developer — developers.redhat.com. <https://developers.redhat.com/products/ansible/lightspeed>. [Accessed 21-03-2025].
- Lukas Berglund, Meg Tong, Maximilian Kaufmann, Mikita Balesni, Asa Cooper Stickland, Tomasz Korbak, and Owain Evans. 2024. *The reversal curse: Llms trained on "a is b" fail to learn "b is a"*. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- 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 of Software Engineering (TSE)*.
- Xinyun Chen, Ryan A. Chi, Xuezhi Wang, and Denny Zhou. 2024. *Premise order matters in reasoning*

- with large language models. In *Proceedings of the 41st International Conference on Machine Learning, ICML'24*. JMLR.org.
- 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.
- Rishav Hada, Varun Gumma, Adrian Wynter, Harshita Diddee, Mohamed Ahmed, Monojit Choudhury, Kalika Bali, and Sunayana Sitaram. 2024. [Are large language model-based evaluators the solution to scaling up multilingual evaluation?](#) In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 1051–1070, St. Julian's, Malta. Association for Computational Linguistics.
- Jia He, Mukund Rungta, David Koleczek, Arshdeep Sekhon, Franklin X Wang, and Sadid Hasan. 2024. [Does prompt formatting have any impact on llm performance?](#) *Preprint*, arXiv:2411.10541.
- Xinyu Lian, Yinfang Chen, Runxiang Cheng, Jie Huang, Parth Thakkar, and Tianyin Xu. 2023. [Configuration validation with large language models](#). *CoRR*, abs/2310.09690.
- Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Mayank Mishra, Matt Stallone, Gaoyuan Zhang, Yikang Shen, Aditya Prasad, Adriana Meza Soria, Michele Merler, Parameswaran Selvam, Saptha Surendran, Shivdeep Singh, et al. 2024. Granite code models: A family of open foundation models for code intelligence. *arXiv preprint arXiv:2405.04324*.
- Niklas Muennighoff, Qian Liu, Armel Randy Zebaze, Qinkai Zheng, Binyuan Hui, Terry Yue Zhuo, Swayam Singh, Xiangru Tang, Leandro von Werra, and Shayne Longpre. 2024. [Octopack: Instruction tuning code large language models](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Sameer Pimparkhede, Mehant Kammakomati, Srikanth G. Tamilselvam, Prince Kumar, Ashok Pon Kumar, and Pushpak Bhattacharyya. 2024. [DocGen: Document-based controlled code generation](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 18681–18697, Miami, Florida, USA. Association for Computational Linguistics.
- Saurabh Pujar, Luca Buratti, Xiaojie Guo, Nicolas Dupuis, Burn Lewis, Sahil Suneja, Atin Sood, Ganesh Nalawade, Matt Jones, Alessandro Morari, and Ruchir Puri. 2023. [Invited: Automated code generation for information technology tasks in yaml through large language models](#). In *2023 60th ACM/IEEE Design Automation Conference (DAC)*, pages 1–4.
- Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, et al. 2023. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950*.
- Kaitao Song, Hao Sun, Xu Tan, Tao Qin, Jianfeng Lu, Hongzhi Liu, and Tie-Yan Liu. 2020. [Lightpaff: A two-stage distillation framework for pre-training and fine-tuning](#). *Preprint*, arXiv:2004.12817.
- Jiao Sun, Yufei Tian, Wangchunshu Zhou, Nan Xu, Qian Hu, Rahul Gupta, John Wieting, Nanyun Peng, and Xuezhe Ma. 2023. [Evaluating large language models on controlled generation tasks](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3155–3168, Singapore. Association for Computational Linguistics.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Bailin Wang, Zi Wang, Xuezhi Wang, Yuan Cao, Rif A. Saurous, and Yoon Kim. 2024. Grammar prompting for domain-specific language generation with large language models. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23, Red Hook, NY, USA*. Curran Associates Inc.

## A Appendix

### A.1 Prompts

This section defines the prompts which are used for models. We report different prompts for every model tried here and report the best-performing prompt results. Generally, the model consists of a System Prompt followed by a prompt template specific to the model.

#### A.1.1 Common prompt

For zero shot inference, we use a common prompt as it is for all the models irrespective of the model's prompt format and we observe best results for Task-1 with this prompt. The prompt is as follows.

Listing 3: common prompt

```
Write an {input_representation} sample with field
values as per the {output_representation}
format schema given below.

{schema}
```



```
{output_representation} sample:  
....
```

### A.1.2 Granite model family

The granite model generally follows the question-answering format. Task-1 prompts for granite family models are as follows.

#### System prompt:

System:

You are an intelligent AI programming assistant, utilizing a Granite code language model developed by IBM. Your primary function is to assist users in code explanation, code generation and other software engineering tasks. You MUST follow these guidelines: - Your responses must be factual. Do not assume the answer is *yes* when you do not know, and **DO NOT SHARE FALSE INFORMATION**. - You should give concise answers. You should follow the instruction and provide the answer in the specified format and **DO NOT SHARE FALSE INFORMATION**.

#### Prompt 2:

Listing 4: QA-prompt-1

```
{System prompt}  
  
Question:  
Write an {input_representation} sample with field  
values as per the {input_representation} format  
schema given below.  
  
{schema}  
  
Answer:  
....
```

#### Prompt 3:

Listing 5: QA-prompt-2

```
{System prompt}  
  
Question:  
Write an {input_representation} sample with field  
values as per the {output_representation}  
format schema given below. Please wrap your  
code  
answer using ` ``` `.  
  
{schema}  
  
Answer:  
....
```

{output\_representation} and {input\_representation} are the variables where {input\_representation} take the values JSON, YAML, XML, Python, and natural language. {output\_representation} takes the values JSON, YAML, and XML.

### A.1.3 Llama family

For codellama 34B model we wrap the common prompt in [INST] and [/INST] tags. For the llama3-8B model, we use the System prompt along with user tags <sup>5</sup>.

**System prompt:** You are a helpful, respectful, and honest assistant. Always answer as helpfully as possible, while being safe. Your answers should not include any harmful, unethical, racist, sexist, toxic, dangerous, or illegal content. Please ensure that your responses are socially unbiased and positive. If a question does not make any sense or is not factually coherent, explain why instead of answering something not correct. If you don't know the answer to a question, please don't share false information.

Other than this, similar to the granite family, we try Question answering format and instruction to wrap the output in quotes (“”).

#### Few shot prompt

Listing 6: Few shot prompt

```
{System prompt}  
  
Your task is to write a JSON sample with field  
values as per JSON format schema.  
You are given a few examples demonstrating the same.  
  
JSON format schema:  
{  
  "type": "array",  
  "contains": {  
    "type": "boolean"  
  },  
  "minContains": 0  
}  
JSON sample:  
....  
[true, true, false]  
....  
JSON format schema:  
{  
  "type": "string",  
  "format": "idn-email"  
}  
JSON sample:  
....  
"hchavezexample.org""JSON format schema:"type":  
"array", "items": "type": "number" "multipleOf":  
5.82, "exclusiveMinimum": 3.069158195370172JSON` sample:"`
```

### A.2 Limitations of Constrained Decoding

This section outlines some common problems with constrained decoding and emphasizes why it cannot be a complete and viable solution for factoring in schemas to generate compliant text using language models.

<sup>5</sup><https://huggingface.co/meta-llama/Meta-Llama-3-8B-Instruct>

Example 1 (Llama3 8B):

```
{
  "type": "array",
  "contains": {
    "type": "object",
    "properties": {
      "cockboats": {
        "type": "number",
        "minimum": 6.838020040763691,
        "exclusiveMaximum":
9.08848414476941,
        "multipleOf": 8.93
      },
      "remeditate": {},
      "endocrinology": {
        "type": "null"
      },
      "subvillain": {
        "type": "boolean"
      },
      "tubularian": {
        "type": "null"
      },
      "jingoed": {
        "type": "null"
      },
      "neoplasties": {
        "type": "null"
      }
    }
  },
  "additionalProperties": true
},
  "minContains": 0
}
```

Example 1 (Llama3 70B):

```
{
  "type": "array",
  "contains": {
    "type": "object",
    "properties": {
      "cockboats": {
        "type": "number",
        "minimum": 6.838020040763691,
        "exclusiveMaximum":
9.08848414476941,
        "multipleOf": 8.93
      },
      "remeditate": {},
      "endocrinology": {
        "type": "null"
      },
      "subvillain": {
        "type": "boolean"
      },
      "tubularian": {
        "type": "null"
      },
      "jingoed": {
        "type": "null"
      },
      "neoplasties": {
        "type": "null"
      }
    }
  },
  "additionalProperties": true
},
  "minContains": 0
}
```

Figure 2: Attention maps for Llama3 8B and 70B model for Data validation experiment. The more the intensity of color, the more attention is given to that part of input by the model.

### A.2.1 Inference Performance Bottleneck

Constrained decoding often negatively affects inference throughput, widely mentioned as one of the major drawbacks in many works (Wang et al., 2024; Pimparkhede et al., 2024; Geng et al., 2023) due to involvement of token-level operations keeping track of the schema constraints and tokens generated so far. This latency can be a factor of the complexity of the schema, tokens generated so far, and the nature of the constrained decoding implementation. Further, advances such as batched inference<sup>6</sup> are not yet there for constrained decoding limiting their scalability and practical use.

### A.2.2 Complex Engineering Effort

Implementing a constrained decoding system can involve instrumenting at the decoding phase of the language model while keeping track of the tokens generated so far and structured schema adherence which can involve implementation specific to a schema representation and may not be possible to generalize to any schema representation. For instance, most of the openly available constrained decoding systems<sup>7</sup> have limited support and not generalized to various schemas such as XML and output formats such as YAML and others. It is worthwhile to note that some approaches tend to convert schemas to context free grammars, however, this approach is possible with common schema representations such as Python pydantic. Additionally, implementing such a system requires deep domain expertise.

### A.2.3 Model Performance Bottleneck

LLMs have multiple failure modes that can likely be triggered through constrained decoding. Many works show that LLMs are sensitive to the text being fed into them and often deteriorate the model's performance. Some examples being the reverse curse from (Berglund et al., 2024), where LLM understanding "A is B" may not guarantee to learn "B is A". Another work (Chen et al., 2024) shows that the order of the premises can have a substantial impact on the performance often affecting negatively. Such failures can be triggered when the natural flow of text generation is interrupted through constrained decoding over autoregressive generation. The problem can worsen when it involves mixed generation of structured output and unstructured NL text.

<sup>6</sup><https://github.com/microsoft/batch-inference>

<sup>7</sup><https://github.com/outlines-dev/outlines>

### A.2.4 Limited Scope

Since constrained decoding needs access to the decoding phase of the language model, its often not possible to apply such decoding to hosted or gated LLM deployments.

Applying constrained decoding to some common use cases is not obvious. Given  $n$  structured schemas from  $s_1$  to  $s_n$ , unstructured NL text output as  $k$  and structured output as  $u$ . Common use cases in natural language processing (NLP) such as summarization involve the following input-output relationship. For some arbitrary schema  $i$ ,  $s_i \rightarrow u$ . Further typical use cases involve factoring in  $n$  multiple schemas and generate  $m$  multiple structured outputs  $(s_1 \dots s_n) \rightarrow (k_1 \dots k_m)$ .

Employing constrained decoding in such use cases is not viable since in the first use case, tasks that output  $u$  cannot leverage constrained decoding and schema has to go into LLMs as input. When multiple schemas and structured outputs are involved, its not obvious to choose the right schema for decoding a particular structured output. Such common use cases substantially limit the scope of using constrained decoding.

## A.3 Task Motivation

### A.3.1 Data as Code Generation Task

This section describes use cases from enterprise and research points of view motivating data as code generation seed tasks in our study.

**Enterprise Use Cases:** (i) Test case structured data generation to test application interfaces such as REST API endpoints. Often, enterprises have a large number of services exposing API endpoints that have to be tested, and LLMs can be a drop-in solution to generate test case data at scale. (ii) Structured configuration data generation for a particular use case and domain. Enterprise applications such as Kubernetes use DSLs for configuration and usage, preparing them require deep domain expertise and there is increasing motivation (Pujar et al., 2023) to employ LLMs in enterprises to generate DSL code. (iii) Some more downstream tasks involving structured data, such as forms and tables often represented in a programmable format such as JSON, can leverage LLMs to generate structured data to fill forms or tables leveraging the schema.

**Research Use Cases:** (i) Since DSLs are typically low resource languages, LLMs are often em-

		Output Representation					
		JSON		YAML		XML	
Model	Schema	IS (%)	RTV (%)	IS (%)	RTV (%)	IS (%)	RTV (%)
Llama3 8B	JSON	<b>1.9</b>	50.1	1.8	49.8	<b>1.6</b>	73.9
Granite 8B		2.9	31.0	2.8	57.3	17.1	<b>70.26</b>
Granite 20B		13.9	<b>15.6</b>	2.3	<b>38.0</b>	7.9	71.92
Granite 34B		2.6	23.5	2.6	48.6	4.1	73.08
Codellama 34B		3.6	17.9	<b>1.8</b>	51.4	3.7	71.12
Llama3 8B	XML	12.9	64.1	6.1	52.8	<b>4.8</b>	73.5
Granite 8B		3.6	60.7	2.8	70.9	10.7	72.0
Granite 20B		2.1	<b>53.3</b>	1.9	73.9	12.2	<b>70.5</b>
Granite 34B		<b>1.9</b>	56.9	1.6	63.1	10.6	71.9
Codellama 34B		2.3	71.2	1.6	56.9	10.2	71.7
Llama3 8B	YAML	1.3	53.3	3.1	62.4	0.4	74.5
Granite 8B		11.2	<b>13.7</b>	1.8	63.9	12.2	70.5
Granite 20B		<b>1.6</b>	39.8	1.4	56.6	10.7	72.0
Granite 34B		3.1	14.9	<b>1.1</b>	<b>40.6</b>	10.6	71.9
Codellama 34B		7.1	24.9	1.4	50.3	12.6	71.0
Llama3 8B	Python	5.4	64.9	3.1	72.9	<b>3.1</b>	72.9
Granite 8B		2.4	73.0	2.3	70.9	10.7	72.71
Granite 20B		<b>1.6</b>	64.7	2.4	68.7	16.6	71.42
Granite 34B		2.6	<b>61.2</b>	<b>2.4</b>	<b>66.9</b>	8.9	69.35
Codellama 34B		5.6	65.1	2.9	64.1	14.1	<b>69.1</b>
Llama3 8B	NL	5.8	50.4	3.4	54.1	5.6	73.9
Granite 8B		<b>2.1</b>	28.9	2.6	29.2	8.3	69.24
Granite 20B		2.9	0.6	2.8	30.2	7.97	69.24
Granite 34B		2.3	<b>1.9</b>	<b>2.4</b>	<b>8.9</b>	9.86	<b>63.42</b>
Codellama 34B		2.8	60.4	2.9	34.5	<b>7.88</b>	65.51

Table 5: Task 1 zero shot results having IS and RTV metric values. IS denotes the percentage of invalid samples and RTV denotes the percentage of sample root data type errors. For IS and RTV, the lesser the value better the performance.

ployed (Song et al., 2020) to synthesize data from LLMs to train and evaluate smaller-sized models. (ii) This task acts as a seed for similar NLP use cases such as code translation.

### A.3.2 DSL Validation Task

This section describes use cases from an enterprise and research perspective that motivate our study’s DSL validation seed task.

**Enterprise Use Cases:** (i) Given the schema, employing LLMs to generate domain-aware suggestions over the provided structured data is not viable with traditional schema validators, which only pinpoint syntactic errors and cannot provide semantic suggestions. Such as providing optimizations over the existing resource YAML in Kubernetes while complying with resource schema. (ii) In an assistive chat system, the constraints are often in NL representation from the user, which is not machine-readable, and LLMs should be able to understand such constraints. (iii) Quick interoperability across different schema and data representation versions. Often in enterprises, schemas can be in a particular version that is incompatible with the structured data version. For instance, the schema could be in an older JSON schema version such as Draft 0 and data in Draft 7, in such cases LLMs can come

handy to perform validation at scale.

**Research Use Case:** Understanding LLMs’ capability in validating the given structured data against the schema across representations can provide seed evidence for more complex tasks such as automatically fixing data in compliance with the given schema.

### A.4 Schema Examples

This section provides schemas across 5 representations from Listings 7 to 11. All the schemas are equivalent in terms of constraints.

Listing 7: Sample schema using JSON Schema

```
{
  "type": "object",
  "properties": {
    "footbaths": {
      "type": "boolean"
    },
    "deluded": {
      "type": "null"
    },
    "bravadoing": {
      "type": "number",
      "exclusiveMaximum": 5.131849487240756
    },
    "queintise": {},
    "manucodia": {
      "type": "number"
    },
    "antagonized": {},
    "outbacker": {
      "type": "number"
    }
  }
}
```

		Output Representation		
		JSON	YAML	XML
Model	Schema	Macro-F1	Macro-F1	Macro-F1
Llama3 8B	JSON	0.55	0.37	0.40
Granite 8B		0.55	0.55	0.42
Granite 20B		0.48	0.37	0.47
Granite 34B		0.60	<b>0.56</b>	<b>0.63</b>
Codellama 34B		<b>0.64</b>	0.53	0.50
Llama3 8B	XML	0.44	0.35	0.41
Granite 8B		0.45	0.44	0.50
Granite 20B		0.24	0.45	<b>0.56</b>
Granite 34B		<b>0.52</b>	<b>0.47</b>	0.39
Codellama 34B		0.41	0.41	0.48
Llama3 8B	YAML	0.38	0.40	0.40
Granite 8B		0.45	0.50	0.44
Granite 20B		0.24	0.31	0.45
Granite 34B		0.52	<b>0.55</b>	0.47
Codellama 34B		<b>0.59</b>	0.52	<b>0.58</b>
Llama3 8B	Python	0.37	0.36	0.38
Granite 8B		0.54	0.44	<b>0.54</b>
Granite 20B		0.34	0.45	0.36
Granite 34B		<b>0.53</b>	<b>0.47</b>	0.40
Codellama 34B		0.48	0.45	0.46
Llama3 8B	NL	<b>0.63</b>	<b>0.55</b>	0.57
Granite 8B		0.45	0.51	0.39
Granite 20B		0.53	0.45	<b>0.57</b>
Granite 34B		0.45	0.46	0.38
Codellama 34B		0.52	0.54	0.42

Table 6: Task 2 zero shot Macro-F1 scores. Task 2 is a binary classification task.

```

"sphenotripsy": {
  "type": "boolean"
},
"hw": {
  "type": "null"
}
},
"additionalProperties": true,
"required": []
}

```

Listing 8: Sample schema using YAML

```

additionalProperties: true
properties:
  antagonized: {}
  bravadoing:
    exclusiveMaximum: 5.131849487240756
    type: number
  deluded:
    type: 'null'
  footbaths:
    type: boolean
  hw:
    type: 'null'
  manucodia:
    type: number
  outbacker:
    type: number
  quintise: {}
  sphenotripsy:
    type: boolean
required: []
type: object

```

Listing 9: Sample schema using Python

```

from pydantic import BaseModel, Field

class Schema(BaseModel):
    footbaths: bool
    deluded: None = Field(None, alias="null")
    bravadoing: float = Field(..., exclusive_maximum=5.131849487240756)
    quintise: None = {}

```

```

manucodia: float
antagonized: None = {}
outbacker: float
sphenotripsy: bool
hw: None = Field(None, alias="null")

```

Listing 10: Sample schema using XML

```

<?xml version="1.0" ?>
<all>
  <type type="str">object</type>
  <properties type="dict">
    <footbaths type="dict">
      <type type="str">boolean</type>
    </footbaths>
    <deluded type="dict">
      <type type="str">>null</type>
    </deluded>
    <bravadoing type="dict">
      <type type="str">number</type>
      <exclusiveMaximum type="float">5.131849487240756</exclusiveMaximum>
    </bravadoing>
    <quintise type="dict"/>
    <manucodia type="dict">
      <type type="str">number</type>
    </manucodia>
    <antagonized type="dict"/>
    <outbacker type="dict">
      <type type="str">number</type>
    </outbacker>
    <sphenotripsy type="dict">
      <type type="str">boolean</type>
    </sphenotripsy>
    <hw type="dict">
      <type type="str">>null</type>
    </hw>
  </properties>
  <additionalProperties type="bool">>true</additionalProperties>
  <required type="list"/>

```

</all>

### Listing 11: Sample schema in NL

This is a JSON schema that defines the structure of an object. Here's a breakdown of the schema:

```
# **Top-level properties**  
  
# * `type`: The type of the JSON data, which is an  
  object (`"object"`).  
# * `properties`: An object that defines the  
  properties of the object.  
# * `additionalProperties`: A boolean value that  
  indicates whether additional properties not  
  specified in the schema are allowed. In this  
  case, it is set to True  
* required: An empty array that specifies no  
  properties are required in the object.
```

**\*\*Properties object\*\***

The `properties` object defines the structure of each property in the object. Here's a brief description of each property:

```
footbaths: A boolean  
deluded: A null  
bravadoing: A number that must be strictly lesser  
  than 5.131849487240756,  
queintise: An object with no specific type or  
  constraints.  
manucodia: A number  
antagonized: An object with no specific type or  
  constraints.  
outbacker: A number  
sphenotripsy: A boolean  
hw: A null
```