

# Controlling Equational Reasoning in Large Language Models with Prompt Interventions

Jordan Meadows<sup>1</sup>, Marco Valentino<sup>2</sup>, André Freitas<sup>1, 2, 3</sup>

<sup>1</sup>Department of Computer Science, University of Manchester

<sup>2</sup>Idiap Research Institute

<sup>3</sup>National Biomarker Centre, CRUK-MI

jordan.meadows@postgrad.manchester.ac.uk

{marco.valentino, andre.freitas}@idiap.ch

## Abstract

This paper investigates how hallucination rates in Large Language Models (LLMs) may be controlled via a symbolic data generation framework, exploring a fundamental relationship between the rate of certain mathematical errors and types of input intervention. Specifically, we systematically generate data for a derivation generation task using a symbolic engine, applying targeted interventions to prompts to perturb features of mathematical derivations such as the surface forms of symbols, equational tree structures, and mathematical context. We then evaluate the effect of prompt interventions across a range of LLMs including fine-tuned T5 models, GPT, and LLaMa-based models. Our experiments suggest that T5-Large can outperform the few-shot performance of GPT-4 on various evaluation sets generated via the framework. However, an extensive evaluation based on human analysis, template-based error detection, and text generation metrics reveals model weaknesses beyond what the reference-based metrics singularly describe. We use these results to tie characteristic distributional footprints of interventions to the human evaluation of LLM derivation quality, potentially leading to significant control over fine-grained mathematical capabilities of language models with respect to specific types of errors.

## Introduction

Large Language Models (LLMs) possess the potential to accelerate mathematical discovery (Trinh et al. 2024). Yet, without incorporating symbolic approaches, their ability to derive correct mathematical results is significantly impeded by their sensitivity to input perturbations and tendency to hallucinate (Mirzadeh et al. 2024; Frieder et al. 2023). In order to understand and mitigate persisting challenges in mathematical reasoning, recent work has focused on the ability of LLMs to learn and sequentially apply symbolic operations to premise equations (Chen et al. 2024; Valentino et al. 2024), investigating whether state-of-the-art models can derive goal equations defined within prompts (Meadows et al. 2024). Such *derivation-style* equational reasoning is at the core of many applied mathematical fields, such as theoretical physics, engineering, and quantitative finance (Plaisted 1993; Premtoon, Koppel, and Solar-Lezama 2020).

In this paper, we aim to establish a deeper connection between input perturbations applied to prompts and certain

mathematical hallucinations, attempting to expose a fundamental relationship between training data, prompt intervention type, and derivational error distributions in LLM output. In particular, we are interested in whether such a relationship can be identified and potentially leveraged to control the rate of certain types of mathematical errors.

To this end, we develop a methodology supported by a symbolic data generation framework that has been effectively applied in related contexts (Valentino et al. 2024; Meadows et al. 2024), which we adapt and improve ( $\approx 15x$  faster) to construct and augment various *fine-grained* datasets for equational reasoning tasks. This framework allows us to operate at a mathematical granularity within derivations that is far greater than what is surfaced on published derivations (Mann et al. 2018; Akrobotu et al. 2022), where many steps are typically omitted or summarised. This ultimately addresses a fundamental incompleteness problem for reasoning data available for training and evaluating LLMs (Villalobos et al. 2022). Given that fine-grained workings contribute to much of the theoretical research distilled in papers, that generative models frequently hallucinate when solving domain-specific reasoning problems (Shuster et al. 2021; Taylor et al. 2022; Frieder et al. 2023; Wysocka et al. 2023; Meadows, James, and Freitas 2024), and that granular reasoning lends itself better to explainability and inference control (Hebenstreit et al. 2023; Yao et al. 2023; Yuan et al. 2023; Valentino and Freitas 2024), it is clear that methods of control over the rate of mathematical reasoning errors made by LLMs, without incorporating external solvers (Toshniwal et al. 2024; Liu and Yao 2024; Trinh et al. 2024; Jiang et al. 2022; Quan et al. 2024), are highly valuable.

By leveraging the granularity of the symbolic data generation framework, we investigate control mechanisms that involve fine-tuning LLMs and applying targeted prompt interventions that systematically manipulate inputs to the models. These interventions — here, alterations in symbolic representation, equational structure, and contextual elements — serve as levers to induce and regulate specific error types in model outputs, moving beyond methodologies which are agnostic to certain classes of mathematical hallucination and/or deal with less complex equation manipulation (Stolfo et al. 2023; Meadows et al. 2024). Crucially, our experiments suggest that certain interventions distinctly correspond to distributional footprints in the error space, opening up the possibil-

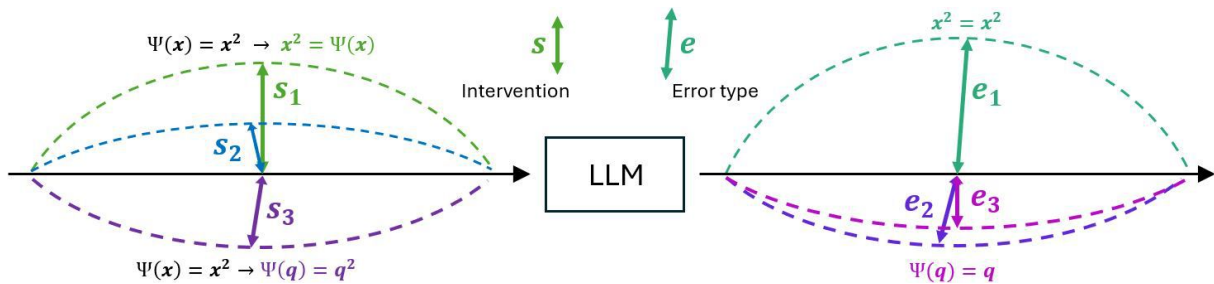


Figure 1: Mechanism relating the type and strength of a set of interventions to their corresponding mathematical errors. The black line indicates the original prompt/output, while the arrows and dashed lines indicate the extent of the input perturbation and respective output delta.

ity to better investigate the underlying nature of mathematical hallucinations in language models.

Overall, our contribution can be summarised as follows:<sup>1</sup>

1. We construct and release a dataset of 30k mathematically fine-grained prompt-derivation pairs spanning 18 operators, 155 wildcard (LaTeX) symbols, 4 targeted distribution shifts, up to 10 equations per derivation, and 160k steps — all developed using a symbolic data generation framework. We also improve the speed of this framework by  $\approx 15x$  and fix limitations involving irrelevant steps.
2. We demonstrate that T5 models fine-tuned on fine-grained mathematical derivations generated using our approach can match or surpass the in-distribution few-shot performance of GPT-4 according to all evaluation methods. In addition, we evaluate a range of open source models in a few-shot setting including LLaMa-2-7B and Llemma-7B.
3. We investigate 3 *complementary evaluation methods* to determine the mathematical proficiency of LLMs: **(1.)** reference-based text generation metrics (including 4 metrics), **(2.)** template-based detection of mathematical errors, and **(3.)** human evaluation of 750 derivations. Each method is applied to both in-distribution and out-of-distribution data augmented by interventions. While we find that template-based and human evaluations are correlated, they both *strongly disagree* with reference-based metrics, proving that the latter deliver misleading model performance rankings and inappropriate representations of the relative effect of interventions.
4. We demonstrate a fundamental underlying mechanism where the rates of certain errors are controlled by systematically varying both the strength and type of interventions on the prompt, visualised in Fig. 1. For instance, the rate of “redundant” equations (*e.g.*,  $x = x$ ) increases by up to 2000% in fine-tuned models, based on an intervention which perturbs equation symmetry. An intervention that removes integration/differentiation results from the prompt leads to a relative increase in step skipping by up to 300%, and leads GPT-4 to make 1000% more reasoning errors, according to human evaluation. Although each intervention naturally affects multiple error categories, their

<sup>1</sup>Code, data, hyperparameters, and further experimental details available at: <https://github.com/jmeadows17/deriving-equations-with-LLMs>

distributional footprint can be uniquely identified, and controlled via the magnitude of each intervention type.

### Derivation Generation with LLMs

Given a goal equation  $G$  and premises  $\mathcal{P}$ , that are arranged within some prompt template  $t(\mathcal{P}, G)$ , we aim to assess the ability of an LLM to systematically apply a set of symbolic operations to premises to generate a sequence of equations  $\hat{\mathcal{D}}$ , which represents a reasonable derivation of  $G$ . Given an LLM  $\mathcal{M}$ , a derivation is generated through  $\mathcal{M} : t(\mathcal{P}, G) \mapsto \hat{\mathcal{D}}$ . An idealised metric  $M^*(\mathcal{D}^*, \hat{\mathcal{D}})$  scores a derivation, where  $\mathcal{D}^*$  is ideally a valid human written derivation corresponding to input prompt  $t(\mathcal{P}, G)$ . We generally aim to optimise

$$M^* = \operatorname{argmax}_{\mathcal{M}}; M^*(\mathcal{D}^*, \mathcal{M}(t(\mathcal{P}, G))).$$

However, we do not have access to ideal derivations  $\mathcal{D}^*$  corresponding to templates  $t(\mathcal{P}, G)$ , nor ideal metric  $M^*$  suitable for scoring  $\hat{\mathcal{D}}$ . Instead, we employ a symbolic engine to *approximate* ground truth derivations to obtain  $\tilde{\mathcal{D}}^*$  (Alg. 1). Moreover, we evaluate over a sample of derivations. This means that, in practice, we are instead finding  $\mathcal{M}^*$  such that

$$\mathcal{M}^* = \operatorname{argmax}_{\mathcal{M}}; \frac{1}{N} \sum_{i=1}^N M(\tilde{\mathcal{D}}^*_i, \mathcal{M}(t(\mathcal{P}_i, G_i))),$$

where  $N$  is the sample size. In this work, we consider  $M$  as a reference-based generation metric to automatically evaluate derivations, but we contrast this with a human evaluation based on equation consistency and coherent operator usage, and a template-based error detection method.

### Dataset Construction

The symbolic data generation process used in this work is described by Alg. 1. Given a vocabulary of symbols  $\mathcal{V}$  and a set of computer algebra operations  $\mathcal{R}$ , our goal is to generate a mathematical derivation  $\mathcal{D}$  represented by an ordered list of steps  $s_i \in \mathcal{D}$ . In order to construct  $\mathcal{D}$ , an initial reasoning step  $s_1 = (\text{premise}, \text{annotation})$  is generated such that  $\mathcal{D} = [s_1]$ . Subsequently, an operation  $r \in \mathcal{R}$  is sampled, which in its most general form accepts two operands (arity 2). The first operand is an equation  $s_{j,1}$  from tuple  $s_j \in \mathcal{D}$ . A suitable secondary variable ( $\in \mathcal{V}$ ), expression, or equation operand  $m$  is extracted from  $\mathcal{D}$ , and the next equation is generated

---

**Algorithm 1: Derivation Generation**

---

**Input:** Vocabulary of symbols  $\mathcal{V}$ , Set of operations  $\mathcal{R}$ **Output:** Ordered list of derivation steps  $\mathcal{D}$ 

- 1: Initialize derivation  $\mathcal{D}$  with a premise step  $s_1 =$  (premise equation, annotation)
  - 2: Set  $i = 1$
  - 3: **while** desired length of  $\mathcal{D}$  not reached **do**
  - 4:   Sample an operation  $r \in \mathcal{R}$
  - 5:   Select an equation  $s_{j,1}$  from tuple  $s_j \in \mathcal{D}$
  - 6:   Extract a suitable operand  $m$  from  $\mathcal{V}$  or  $\mathcal{D}$  that matches the requirements of  $r$
  - 7:   Generate the next equation  $s_{i+1,1} = r(s_{j,1}, m)$
  - 8:   Create an annotation  $s_{i+1,2}$  representing the operation and operands:  $s_{i+1,2} = [r, j, m']$  where  $m'$  is an index or variable/expression string corresponding to  $m$
  - 9:   Append the new step to the derivation:  $\mathcal{D}.append((s_{i+1,1}, s_{i+1,2}))$
  - 10:   Increment  $i$
  - 11: **end while**
  - 12: **return**  $\mathcal{D}$
- 

Dataset	Size (k)
Training	15.3
Static Test Set (In-distribution)	3.1
Variable Renaming (VR)	2.9
Expression Exchange (EE)	3.1
Alternative Goal (AG)	3.1
Step Removal (SR)	1.0


Table 1: Sizes for the various Derivation Generation datasets.

by applying operation  $r$  through  $s_{i+1,1} = r(s_{j,1}, m)$ . The annotation  $s_{i+1,2}$  is a list containing the name of the operation, the equation index, and the secondary operand, such that  $s_{i+1,2} = [r, j, m']$  (where  $m'$  is a variable/expression string or equation index representing operand  $m$ ). Therefore, step  $s_{i+1} = (r(s_{j,1}, m), [r, j, m'])$ . If  $\mathcal{D} = [s_1]$ , then  $i = j = 1$ , and the derivation updates such that  $\mathcal{D} = [s_1, s_2]$ . This process repeats until the derivation reaches a target length.

A derivation generated from Alg. 1 is then *perturbed* according to a set of interventions to form out-of-distribution examples. Specifically, given the task of Derivation Generation instantiated via the prompt template  $t = t(\mathcal{P}, G)$  and a ground truth derivation  $\mathcal{D}$ , a static dataset  $X$  consisting of  $(t, \mathcal{D})$  is constructed using the process described above. Subsequently, a perturbed dataset  $X_n$  is formed by applying a perturbation function  $P_n$  to all  $(t, \mathcal{D}) \in X$  to form  $(t', \mathcal{D}') \in X_n$ , such that  $P_n : X \rightarrow X_n$ , and  $n$  denotes the number of perturbations considered.

Tab. 1 describes the dataset sizes generated by the symbolic framework, Fig. 2 displays the distribution of equation counts in ground truth references (*i.e.*, derivation length), and Tab. 2 shows that the operation chains responsible for forming the underlying derivation reasoning do not frequently repeat.

**On the omission of natural language.** Although the framework (through Alg. 1) outputs *both* equations and step an-

Figure 2: Length distribution  $P(L)$  of derivations.

Length (L)	Permutations	Chain	P(Chain)
4	842	$\partial \rightarrow \partial_E \rightarrow S_L$	0.0369
		$\int \rightarrow \int_E \rightarrow S_L$	0.0186
5	2850	$+ \rightarrow \partial \rightarrow \partial_E \rightarrow S_L$	0.0053
		$- \rightarrow \partial \rightarrow \partial_E \rightarrow S_L$	0.0048

Table 2: For a given derivation length  $L$ , **Permutations** describes the number of unique operation sequences present in the training data. **Chain** describes the two most frequent operation sequences based on symbols. **P(Chain)** is the probability of the chain.

notations by default, we purposefully remove annotations from the output in the specific Derivation Generation task considered in this work. Firstly, annotations give additional information on the dependency structure between equations, but they are certainly not necessary for the purpose of generating valid dependency graphs between equations with generative models. Ground truth derivations from the dataset can be clearly followed without natural language. Second, without annotations, the coherence of the derivation depends on the equations and their dependencies alone. This is more targeted than the alternative of additionally determining whether annotations match up with equations.

**Improvements to symbolic data generation framework.**

We rely on the symbolic framework proposed in Meadows et al. (2024) to support the experimental pipeline following related work in non-generative settings (Valentino et al. 2024). We improve the data generation approach in the following ways:

1. **Support for complex LaTeX symbols** (*e.g.*,  $\Psi_{nl}$ ) instead of more basic symbols (*e.g.*,  $x$ ).
2. **Removed irrelevant and disconnected equations** from derivations by including additional dependency checks between derivation steps. This improvement was crucial for eliciting the desired derivational behaviour in models via fine-tuning and in-context learning.
3. **Improved runtime efficiency by a factor of 15** by allowing derivations to both equal or exceed the target length, including timeout decorators on certain operations, and using more efficient iteration limiting. The approximate difference is **< 0.05 min/derivation** compared to **0.7**

min/derivation tested over 100 samples.

We adopt the same set of hyperparameters described in Meadows et al. (2024) using the following values:  $p\_history=10$ ,  $p\_arity\_0=5$ ,  $p\_renaming=1$ ,  $p\_arity\_1=50$ ,  $p\_evaluate=50$ ,  $p\_arity\_2=100$ ,  $p\_int\_or\_diff=1$ ,  $p\_subs=5$ .

## Prompt Interventions and Perturbations

A perturbation or intervention is a transformation applied to the input text and/or ground truth that ideally changes a single target textual aspect. We apply 4 interventions to the static test set to generate corresponding perturbed sets.

**Variable Renaming (VR).** In the training set, derivations rely on a vocabulary of 155 symbols (e.g.,  $\Psi_{nl}$ ,  $E_n$ ,  $\mathbf{J}_P$ ,  $\eta$ ,  $g_\varepsilon$ ). For each example in the static set, we uniquely map each symbol to an out-of-distribution symbol sampled from 11 Greek letters (e.g.,  $E_n = n + x$  becomes  $\alpha = \beta + \gamma$ ).

**Expression Exchange (EE).** In the training set and applied mathematics in general, there is an asymmetry with respect to premises being defined with functions on the LHS and expressions on the RHS (e.g.,  $E_n(n, x) = n + x$ ). However, operations are frequently used that can substitute LHS for RHS (and vice versa) in many cases, and both functions and operations may appear on either side of equations. Simply, we swap expressions either side of the equality for all equations in the static test set (e.g.,  $E_n = n + x$  becomes  $n + x = E_n$ ).

**Alternative Goal (AG).** For each example in the static set, we derive an alternative goal equation from the penultimate equation, by random selection of operators and operands, which equates to the synthetic data algorithm skipping its first choice goal equation for that derivation. This perturbation should not result in significant differences in model outputs as it simply applies alternative in-distribution operations that occur frequently during training or within few-shot prompts.

**Step Removal (SR).** In the training set, equations that occur as a result of evaluating differentials and integrals are included in the prompt as intermediate steps. These are used to guide model outputs. This perturbation removes such “*then derive*” equations from the prompt, which forces models to either circumvent such steps or derive them during inference.

## Empirical Evaluation

**Empirical setup.** We evaluate a range of LLMs on the Derivation Generation task including T5 (Raffel et al. 2020), FLAN-T5 (Chung et al. 2024), GPT-3.5 and 4 (Brown et al. 2020; Achiam et al. 2023), LLaMa-2-7B (Touvron et al. 2023), and Llemma-7B (Azerbaiyev et al. 2023). We fine-tune base (220M) and large (770M) variants of T5 and FLAN-T5, while using the remaining models in a few-shot setting. The evaluation occurs across 3 complementary methods: **(1.)** use of reference-based text generation metrics (Tab. 3), **(2.)** error count as determined by searching model output for surface-level mathematical errors (Tab. 4), and **(3.)** a *manual analysis* of models’ reasoning accuracy across 750 total derivations

(Tab. 5). Additional details on models, training, and metrics can be found online.<sup>2</sup>

**Prompt design.** We fine-tune and zero-shot prompt the T5 models following the template below, which corresponds to a ground truth sequence of equations:

Given  $q(a) = e^a$   
and  $G(a) = -e^a + \frac{d}{da}q(a)$ ,  
then derive  $-e^a + \frac{d}{da}q(a) = 0$ ,  
then obtain  $e^{G(a)} = 1$

To few-shot prompt the decoder-only models (i.e., GPT, LLaMa and Llemma) we use the following design, where  $n = 5$  is the number of in-context examples:

---

The following examples consist of a prompt (denoted by Prompt:) and a mathematical derivation (denoted by Derivation:). Each derivation contains LaTeX equations separated by "and".

Prompt: [Prompt 1]  
Derivation: [Derivation 1]

⋮

Prompt: [Prompt  $n$ ]  
Derivation: [Derivation  $n$ ]

Now given the following prompt, generate the derivation. Ensure equations are split by the word "and".

Prompt: [Evaluation Prompt]

---

This approach was chosen to minimise natural language in the generated output, and to force derivations into the desired format (LaTeX equations split by “and”). Notably, only the Evaluation Prompt is perturbed, ensuring that the bulk differences in scores are not caused by changes to in-context examples, and evaluation is pair-wise consistent.

## Results with Text Generation Metrics

Here we highlight the main results obtained using common text generation metrics including ROUGE (Lin 2004), BLEU (Papineni et al. 2002), BLEURT (Sellam, Das, and Parikh 2020), and GLEU (Mutton et al. 2007).

**Small fine-tuned LMs outperform few-shot GPT-4 across all generation metrics.** On 2K examples (denoted by ( $f$ )) from the static set, FLAN-T5-large outperforms all models in all metrics. This minor advantage over T5-large may stem from further instruction fine-tuning (our prompt is an instruction). However this advantage over T5 does not extend to

---

<sup>2</sup>Further experimental details: <https://github.com/jmeadows17/deriving-equations-with-LLMs>

	ROUGE					BLEU					BLEURT					GLEU				
	S	VR	EE	AG	SR	S	VR	EE	AG	SR	S	VR	EE	AG	SR	S	VR	EE	AG	SR
T5-base ( <i>f</i> )	88.6	80.2	86.2	88.3	77.0	81.3	74.8	78.2	80.9	64.2	70.5	67.7	67.3	67.4	51.5	83.4	76.1	80.4	83.1	69.2
FLAN-T5-base ( <i>f</i> )	87.3	24.4	84.3	86.7	77.7	79.4	41.1	76.0	78.8	66.6	68.9	18.7	67.0	67.9	<b>56.8</b>	81.7	44.2	78.5	81.3	71.0
T5-large ( <i>f</i> )	89.4	<b>85.0</b>	86.8	89.2	77.7	82.8	<b>79.3</b>	79.5	82.5	66.4	72.1	<b>70.8</b>	68.3	69.6	54.1	<b>80.8</b>	<b>81.5</b>	84.4	70.6	
FLAN-T5-large ( <i>f</i> )	<b>90.2</b>	83.0	<b>87.1</b>	<b>89.5</b>	<b>78.6</b>	<b>84.6</b>	<b>78.5</b>	<b>80.4</b>	<b>83.5</b>	<b>68.9</b>	<b>73.2</b>	69.0	<b>68.7</b>	<b>70.3</b>	56.1	<b>86.1</b>	79.6	<b>82.1</b>	<b>85.1</b>	<b>72.4</b>
T5-base	89.5	82.2	87.3	89.9	79.9	82.8	77.2	81.6	83.7	68.8	70.5	71.1	69.6	70.1	56.5	84.4	78.0	82.6	85.3	72.5
FLAN-T5-base	87.0	25.7	86.7	87.8	78.5	80.3	40.4	81.1	81.1	68.5	67.2	14.6	69.0	66.4	56.7	81.9	42.9	82.2	82.9	71.8
T5-large	91.0	<b>86.2</b>	87.7	<b>90.5</b>	80.6	85.1	<b>80.7</b>	82.4	84.7	71.0	72.5	<b>71.9</b>	<b>70.7</b>	<b>71.8</b>	59.6	86.4	<b>81.7</b>	<b>83.3</b>	86.1	74.1
FLAN-T5-large	<b>91.2</b>	85.1	<b>87.9</b>	90.4	<b>80.7</b>	<b>86.1</b>	79.8	<b>83.1</b>	<b>84.8</b>	<b>72.3</b>	<b>72.9</b>	71.2	70.5	71.4	<b>61.0</b>	<b>87.2</b>	80.6	<b>83.8</b>	<b>86.2</b>	<b>75.0</b>
GPT-3.5	80.3	78.8	78.8	80.6	73.3	70.8	70.2	70.7	71.4	64.2	63.1	63.9	62.1	61.7	50.9	73.5	72.7	72.9	74.3	67.7
GPT-4	82.8	81.6	80.9	82.1	75.6	72.2	71.1	68.3	70.4	61.7	62.9	64.2	61.3	61.8	50.4	75.6	74.4	72.3	74.4	67.2
LLaMa-2-7B	34.3	29.6	37.6	36.5	39.2	28.6	24.6	31.1	30.3	29.1	-18.2	-25.8	-14.6	-15.0	-13.3	30.8	27.1	34.1	32.9	35.1
Llemma-7B	75.7	73.9	73.0	74.9	62.6	63.6	63.8	61.9	63.4	52.6	59.7	63.4	59.3	58.9	49.9	67.3	66.7	65.6	67.2	56.2

Table 3: Evaluation results with both in-distribution static scores (S) and those from the interventions (VR, EE, AG, SR).

	Syntax Errors					Equality Errors					Repeating Errors					Redundant Errors					Skipped steps					Verbose				
	S	VR	EE	AG	SR	S	VR	EE	AG	SR	S	VR	EE	AG	SR	S	VR	EE	AG	SR	S	VR	EE	AG	SR	S	VR	EE	AG	SR
T5-base	20	32	7	19	25	0	3	2	2	0	7	36	35	8	2	1	15	18	1	8	69	46	44	65	155	7	73	28	4	0
FLAN-T5-base	2	43	4	3	1	1	23	0	0	0	7	33	24	5	3	2	13	25	0	10	95	196	64	92	168	5	33	22	4	0
T5-Large	11	11	14	9	15	0	3	8	0	0	5	28	33	8	3	1	21	2	4	53	41	41	53	140	9	40	33	13	2	
FLAN-T5-Large	<b>1</b>	20	7	2	3	<b>0</b>	4	0	1	0	7	21	30	7	3	2	8	26	4	<b>48</b>	36	30	50	133	11	50	38	11	2	
GPT-3.5	4	0	1	1	3	2	2	2	3	2	3	4	3	3	1	0	2	1	0	1	96	91	81	98	128	51	37	41	26	28
GPT-4	<b>1</b>	1	0	0	0	<b>0</b>	0	0	0	0	<b>0</b>	1	0	0	0	<b>0</b>	1	1	0	5	112	105	134	118	154	7	20	13	8	11
LLaMa-2-7B	6	3	5	6	9	4	8	7	6	17	14	11	28	9	35	6	2	3	1	8	99	91	115	107	219	21	17	33	15	27
Llemma-7B	13	11	13	11	35	17	12	13	9	26	94	139	153	108	243	1	11	3	1	16	114	110	103	134	110	92	125	120	92	181

Table 4: Error counts for specific equation-level and derivation-level categories. **Syntax** refers to the number of equations with unbalanced brackets. **Equality** counts the number of equations without equality (or inequality) symbols. **Repeating** is the total number of repeated equations. **Redundant** is the number of equations where the LHS exactly matches the RHS. **Skipped steps** and **Verbose** respectively count the excess or reduced number of equations in the output compared to the reference derivation.

FLAN-T5-base, which scores lower than T5-base in all metrics. This may be due to fine-tuning instability observed in T5 (Asai et al. 2022). We note that despite the success of the fine-tuned models, according to the metrics, we are not suggesting they are more suitable for equational reasoning than GPT, as the other sections of the evaluation reveal.

The scores reported in Tab. 3 (without (*f*)) are evaluated on 100 examples from the static set *explicitly containing integration and differentiation results in the prompt*, in order to fairly examine the effect of the Step Removal intervention which perturbs input by removing these results. The fine-tuned LMs score within 3 units of their previous scores, model rankings are preserved across all metrics, and we assume the GPT/LLaMa scores would report similarly for larger samples. With that said, according to all metrics, all fine-tuned models generally outperform all (5-shot) decoder models. Notably, vanilla LLaMa-2 scores are less than half of those obtained by Llemma, indicating the benefits of Llemma’s fine-tuning on mathematical corpora (Azerbayev et al. 2023). This difference is mirrored in the out-of-distribution scores.

## Results with Template-based Error Detection

In this section, we count the rate of different categories of errors by extracting equations from a model’s derivation into an ordered list and adopting a set of metrics using both comparative and reference-free methods.

We consider 6 error categories in total (Tab. 4) and sum the per-derivation error counts over all examples. For example, determining whether a model has either *Skipped steps*

or is too *Verbose* (not necessarily “errors”) occurs by comparing the length of the generated derivation with that of the ground truth reference. The number of *Repeating equations* is determined by counting the number of uniquely generated equations in the output derivation (reference-free). The number of *Equality* errors per derivation is calculated considering the number of equations that do not contain the “=” token.

**Fine-tuned FLAN-T5-Large and few-shot GPT-4 obtain the lowest rate of in-distribution errors.** Across the 6 error categories, both FLAN-T5-Large and GPT-4 possess the least *Syntax* and *Equality* errors. Results then diverge as FLAN-T5-Large repeats several equations (*Repeating*) and outputs a couple of equations where the LHS is an exact string match with the RHS (*Redundant*), whereas GPT-4 *skips around twice as many steps* as the largest fine-tuned models. This supports the claim that we are generally training models at a level of mathematical granularity *below* that of the data used to train GPT-4.

**Model rankings and perturbation difficulty based on template-based error count disagree with those determined by generation metrics.** At the top of the model rankings, on both the in-distribution and perturbed test sets, Tab. 4 shows that GPT-4 and GPT-3.5 generally leapfrog the fine-tuned models in terms of low total error count, which disagrees with the scores obtained by the generation metrics. The LLaMa-based models remain at the bottom.

In terms of perturbation difficulty, SR (Step Removal) is no longer the most challenging perturbation for fine-tuned models by error count, yet remains difficult for all other models. From this, *FLAN-T5-Large does not outperform GPT-4*

	S	VR	EE	AG	SR
GPT-4	98	96	92	100	80
FLAN-T5-Large	98	68	62	92	76
Llemma-7B	70	70	64	76	36

Table 5: Equational reasoning accuracy % from human evaluation.

or GPT-3.5 on out-of-distribution examples, but it generally does outperform LLaMa-2 and Llemma. This suggests that for the fine-tuned models, certain interventions (such as SR) correspond to a clearly identifiable distribution of mathematical error types, such that the intervention may be inferred from the results. We discuss the implications of this after the manual evaluation.

### Human Evaluation of LLM Derivation Quality

To recap, the generation metrics return scores and performance rankings via n-gram-based similarity measures between model derivations and ground truth references. In parallel, the error detection returns scores via a combination of reference-based and reference-free surface-level checks, that pick up on basic mathematical errors such as imbalanced brackets, missing equality signs, and repeating equations. What is lacking is an assessment of models’ underlying operational reasoning that is largely independent of surface-level checks, and should be as close to a reference-free evaluation as possible within the scope of the task. We aim to provide such analysis in this section. Our approach to determining whether a model’s derivation is coherent is as follows. A derivation:

- (1.) must not include any equations that are inconsistent (and can not be made consistent by substituting a number for a variable, e.g.,  $\alpha = 1$ ).
- (2.) must not include malformed equations with the exclusion of minor typos.
- (3.) must include the exact goal equation, as this is always given in the prompt.
- (4.) must apply operators correctly and in the correct order where operators are non-commutative.
- (5.) may skip numerous steps, even premises, so long as a path may be reasonably derived between consistent equations.
- (6.) may include irrelevant (but consistent) equations that do not contribute to the core path linking premises to the goal equation.
- (7.) may repeat some equations.

The above marking scheme provides a very lenient framework which *prioritises the consistency and operation-wise correctness of equations*.

**Sampling derivations.** The results in Tab. 5 were determined from a sample of 50 static derivations per model, with 4 corresponding perturbed derivations (VR, EE, AG, SR), totalling 250 per model. The 50 static derivations are sampled by ensuring the average ROUGE, BLEU, BLEURT, and GLEU scores over the sample aligns with that model’s static score (S) in Tab. 3 to within 0.1 units. For instance,

GPT-4’s sample averages scores of 82.8, 72.2, 62.9, and 75.6, and the perturbed derivations are based on the static examples from this selection.

We select derivations which are output by the best fine-tuned, GPT, and open-source models, totalling 750 examples. We convert equations into a easily readable format rendered in LaTeX, and either manually delete equations or fix any minor typos causing compilation errors. This document is available here.<sup>3</sup>

**Fine-tuned FLAN-T5-Large and GPT-4 are matched in-distribution in terms of equation consistency and coherent use of operators.** Tab. 5 illustrates how the fine-tuned and GPT models scored 98% accuracy according to the lenient marking scheme. GPT’s single incoherent derivation involved adding a variable to both sides of an equation, then later integrating and forcing that variable to be the constant of integration (breaking rule (2.)). FLAN-T5’s incoherent derivation involved a sequence of equations which implied that  $\frac{d\theta}{dq} = \left(\frac{d\theta}{dq}\right)^q$ , which is true only if  $q = 1$  (breaking rule (2.)).

**Accounting for numerous surface-level error checks approximates human evaluation.** Despite the fact that one evaluation scheme focuses on equation consistency, while the other compares surface-level errors, the respective manual and template-based error results agree that, in-distribution, GPT-4 and FLAN-T5-Large are tied (with Llemma significantly underperforming). Regarding out-of-distribution evaluation, both manual and error-based scores agree that the *fine-tuned models are less capable than all generation metrics suggest*. In terms of intervention difficulty, both agree that SR (*Step Removal*) is particularly challenging while AG (*Alternative Goal*) is the least, and both schemes agree that the fine-tuned models are less affected by SR. This is not reflected by any of the generation metrics.

Given this alignment between the manual evaluation (Tab. 5) and the template-based error detection (Tab. 4) spanning only 6 error types, and that more mathematically capable language models are less likely to hallucinate syntax errors and related trivially detectable artefacts, together this suggests that by accounting for a large ( $\gg 6$ ) number of surface-level errors, we can approximate human evaluation of LLM equational reasoning, *at least* more faithfully than many canonical generation metrics. This can be achieved via the *weighted average of counts over all surface-level categories*, where the category weights are empirically determined through comparison with rankings based on human evaluation.

### Controlling Equational Reasoning in LLMs

Supported by the manual and error-based evaluations, interventions are intrinsically linked to certain LLM hallucinations with varying degrees of association.

We start defining such relationships by considering the rate of hallucination types  $R(i, m, t)$  corresponding to an

<sup>3</sup>[https://github.com/jmeadows17/deriving-equations-with-LLMs/blob/main/Derivation\\_Analysis.pdf](https://github.com/jmeadows17/deriving-equations-with-LLMs/blob/main/Derivation_Analysis.pdf)




Figure 3: The average distributional footprint left by certain interventions.

intervention  $i$ , error type  $t$ , model  $m$ , and the static rate  $S(m, t)$  from Tab. 4. For instance, Fig. 3 is characterised by:

$$\delta(i, t) = \mathcal{N} \sum_m (R(i, m, t) - S(m, t)) \quad (1)$$

where  $\mathcal{N}$  is a normalisation factor (reciprocal number of models considered). Hence  $\delta(i, t)$  is the average error count for intervention  $i \in \{\mathbf{VR}, \mathbf{EE}, \mathbf{AG}, \mathbf{SR}\}$  of type  $t \in \{\text{Syntax, Equality, ...}\}$ . If  $\delta < 0$ , then the intervention (on average) *reduced* the rate of hallucinations of that type across the models, and vice versa.

Furthermore, certain classes of interventions (*e.g.*, injection of random noise, random token deletion) may depend on a practically continuous variable,  $s$ , that predictably varies the rate of certain errors. Hence, with some loss of information, intervention  $i$  may be represented as a vector  $\mathbf{x}_i(s) = (\mathbb{E}_t[\delta(i, t, s)], \sigma_t(\delta(i, t, s)), \dots)$ , where  $\mathbb{E}_t[\delta(i, t, s)]$  and  $\sigma_t(\delta(i, t, s))$  (etc.) are the expectation value and standard deviation of  $\delta$  over the error types.

To find  $s$  such that intervention  $i$  likely improves the output quality over most hallucination types at that strength, we can write

$$\mathbb{E}_t[\delta(i, t, s)] + \varepsilon \sigma_t(\delta(i, t, s)) < 0 \quad (2)$$

where large  $\varepsilon$  ensures that  $\delta(i, t, s) < 0$  across a greater number of hallucination types  $t$ .

The clear error distributions associated with each intervention in Fig. 3 (characterised by Eq. 1) are averaged over all evaluated models, but most closely align with the fine-tuned T5 models. For this class of approaches, the interventions have a distinct effect on specific surface-level errors such as the rate of repeating equations or incorrect syntax. If these distributions  $\mathbf{x}_i(s)$  may be further controlled by some variable  $s$ , we can define conditions for reducing surface-level error rates (*e.g.*, Eq. 2), which correlates with improved derivation quality according to human evaluation.

## Related Work

Our focus is evaluating and controlling the LLM-based (Brown et al. 2020; Ahmed and Devanbu 2022; Song et al. 2022; Ge et al. 2023; Hu et al. 2023; Yang et al. 2023; Dubey et al. 2024; Meta 2024) generation of informal mathematical reasoning that resembles step-wise detailed equation derivations. While we focus on *equation generation*, mathematical generation exists in various forms, and can be

clustered into two main categories: approaches that consider formal languages, and those that consider informal mathematical natural language (Meadows and Freitas 2023; Lu et al. 2022; Zhong, Yang, and Lin 2022). In the formal case, GPT-j (Polu and Sutskever 2020; Polu et al. 2022), LISA (Jiang et al. 2021), and Baldur (First et al. 2023) focus on Metamath and Isabelle/HOL proofs. For generation involving informal reasoning, an approach based on OpenAI’s Codex (Chen et al. 2021; Drori et al. 2022) translates university-level problems into executable code, and generates solution explanations. Minerva (Lewkowycz et al. 2022) is a PaLM (Chowdhery et al. 2022) model trained on a large corpus of mathematical text, and solves university-level problems in applied math, outputting solutions in the form of mathematical natural language. NaturalProver (Welleck et al. 2022) generates similar solutions to proofs from a curated dataset (Welleck et al. 2021), and is most similar to our present work. However, our approach differs in a number of ways. Firstly, we focus exclusively on the generation of equational chains (in contrast to the inclusion of natural language statements). Our prompts and derivations are procedurally generated valid derivations in LaTeX, and many examples are guaranteed to include reasoning which is out-of-distribution with respect to other datasets, while containing up to 10 equations with wildcard symbols (Zanibbi et al. 2016). Lastly, our use of symbolic interventions follows from a previous approach (Meadows et al. 2024), and we describe specific improvements in a later section.

## Conclusion

To control and assess the fine-grained equation derivation capabilities of LLMs via prompt interventions, we first construct a dataset comprising 30k mathematically fine-grained prompt-derivation pairs using a symbolic data generation framework. We find fine-tuned models match or surpass the *in-distribution* performance of few-shot GPT-4, despite a difference in parameter count of up to 3 orders of magnitude.

However, while all generation metrics suggest the fine-tuned models also outperform few-shot GPT-4 on perturbed data, the manual and template-based error detection methods both *strongly disagree* with reference-based metrics, suggesting the latter lead to inappropriate representations of the relative effect of interventions and model capabilities. This strong alignment between human and error-based analysis suggests that extensive human evaluation can be approximated by accounting for numerous categories of surface-level errors.

We use these results to describe how a fundamental underlying mechanism relating distribution shifts to certain surface-level errors may be leveraged by varying the type and strength of prompt interventions to mitigate hallucination rates, which in turn facilitates control over the quality of LLM-based equational reasoning.

Given a sufficiently large number of detectable mathematical hallucinations, an intervention with a variable strength that predictably controls the rate of certain errors, and an appropriate statistical condition, we can improve the quality of LLM reasoning post-training by experimentally determining an appropriate intervention strength.

## Acknowledgements

This work was partially funded by the Swiss National Science Foundation (SNSF) project NeuMath: Neural Discourse Inference over Mathematical Texts (200021\_204617).

## References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Ahmed, T.; and Devanbu, P. 2022. Few-shot training LLMs for project-specific code-summarization. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*, 1–5.
- Akrobotu, P. D.; James, T. E.; Negre, C. F.; and Mniszewski, S. M. 2022. A QUBO formulation for top- $\tau$  eigencentality nodes. *Plos one*, 17(7): e0271292.
- Asai, A.; Salehi, M.; Peters, M.; and Hajishirzi, H. 2022. ATTEMPT: Parameter-Efficient Multi-task Tuning via Attentional Mixtures of Soft Prompts. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 6655–6672. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics.
- Azerbaiyev, Z.; Schoelkopf, H.; Paster, K.; Santos, M. D.; McAleer, S.; Jiang, A. Q.; Deng, J.; Biderman, S.; and Welleck, S. 2023. Llemma: An open language model for mathematics. *arXiv preprint arXiv:2310.10631*.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Chen, M.; Tworek, J.; Jun, H.; Yuan, Q.; Pinto, H. P. d. O.; Kaplan, J.; Edwards, H.; Burda, Y.; Joseph, N.; Brockman, G.; et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.
- Chen, X.; Chi, R. A.; Wang, X.; and Zhou, D. 2024. Premise Order Matters in Reasoning with Large Language Models. *arXiv preprint arXiv:2402.08939*.
- Chowdhery, A.; Narang, S.; Devlin, J.; Bosma, M.; Mishra, G.; Roberts, A.; Barham, P.; Chung, H. W.; Sutton, C.; Gehrmann, S.; et al. 2022. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*.
- Chung, H. W.; Hou, L.; Longpre, S.; Zoph, B.; Tay, Y.; Fedus, W.; Li, Y.; Wang, X.; Dehghani, M.; Brahma, S.; et al. 2024. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70): 1–53.
- Drori, I.; Zhang, S.; Shuttleworth, R.; Tang, L.; Lu, A.; Ke, E.; Liu, K.; Chen, L.; Tran, S.; Cheng, N.; et al. 2022. A neural network solves, explains, and generates university math problems by program synthesis and few-shot learning at human level. *Proceedings of the National Academy of Sciences*, 119(32): e2123433119.
- Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- First, E.; Rabe, M. N.; Ringer, T.; and Brun, Y. 2023. Baldur: Whole-Proof Generation and Repair with Large Language Models. *arXiv preprint arXiv:2303.04910*.
- Frieder, S.; Pinchetti, L.; Griffiths, R.-R.; Salvatori, T.; Lukasiewicz, T.; Petersen, P. C.; Chevalier, A.; and Berner, J. 2023. Mathematical capabilities of chatgpt. *arXiv preprint arXiv:2301.13867*.
- Ge, Y.; Hua, W.; Ji, J.; Tan, J.; Xu, S.; and Zhang, Y. 2023. Openagi: When llm meets domain experts. *arXiv preprint arXiv:2304.04370*.
- Hebenstreit, K.; Praas, R.; Kiesewetter, L. P.; and Samwald, M. 2023. An automatically discovered chain-of-thought prompt generalizes to novel models and datasets. *arXiv preprint arXiv:2305.02897*.
- Hu, Z.; Lan, Y.; Wang, L.; Xu, W.; Lim, E.-P.; Lee, R. K.-W.; Bing, L.; and Poria, S. 2023. LLM-Adapters: An Adapter Family for Parameter-Efficient Fine-Tuning of Large Language Models. *arXiv preprint arXiv:2304.01933*.
- Jiang, A. Q.; Li, W.; Han, J. M.; and Lisa, Y. W. 2021. Language models of isabelle proofs. In *6th Conference on Artificial Intelligence and Theorem Proving*.
- Jiang, A. Q.; Welleck, S.; Zhou, J. P.; Li, W.; Liu, J.; Jamnik, M.; Lacroix, T.; Wu, Y.; and Lample, G. 2022. Draft, sketch, and prove: Guiding formal theorem provers with informal proofs. *arXiv preprint arXiv:2210.12283*.
- Lewkowycz, A.; Andreassen, A.; Dohan, D.; Dyer, E.; Michalewski, H.; Ramasesh, V.; Slone, A.; Anil, C.; Schlag, I.; Gutman-Solo, T.; et al. 2022. Solving quantitative reasoning problems with language models. *arXiv preprint arXiv:2206.14858*.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.
- Liu, H.; and Yao, A. C.-C. 2024. Augmenting math word problems via iterative question composing. *arXiv preprint arXiv:2401.09003*.
- Lu, P.; Qiu, L.; Yu, W.; Welleck, S.; and Chang, K.-W. 2022. A survey of deep learning for mathematical reasoning. *arXiv preprint arXiv:2212.10535*.
- Mann, C.-R.; Sturges, T. J.; Weick, G.; Barnes, W. L.; and Mariani, E. 2018. Manipulating type-I and type-II Dirac polaritons in cavity-embedded honeycomb metasurfaces. *Nature communications*, 9(1): 1–11.
- Meadows, J.; and Freitas, A. 2023. Introduction to mathematical language processing: Informal proofs, word problems, and supporting tasks. *Transactions of the Association for Computational Linguistics*, 11: 1162–1184.
- Meadows, J.; James, T.; and Freitas, A. 2024. Exploring the Limits of Fine-grained LLM-based Physics Inference via Premise Removal Interventions. *arXiv:2404.18384*.
- Meadows, J.; Valentino, M.; Teney, D.; and Freitas, A. 2024. A Symbolic Framework for Evaluating Mathematical Reasoning and Generalisation with Transformers. *arXiv:2305.12563*.
- Meta, A. 2024. Introducing meta llama 3: The most capable openly available llm to date. *Meta AI*.

- Mirzadeh, I.; Alizadeh, K.; Shahrokhi, H.; Tuzel, O.; Bengio, S.; and Farajtabar, M. 2024. Gsm-symbolic: Understanding the limitations of mathematical reasoning in large language models. *arXiv preprint arXiv:2410.05229*.
- Mutton, A.; Dras, M.; Wan, S.; and Dale, R. 2007. GLEU: Automatic evaluation of sentence-level fluency. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, 344–351.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 311–318.
- Plaisted, D. A. 1993. Equational reasoning and term rewriting systems. *Handbook of logic in artificial intelligence and logic programming*, 1: 273–364.
- Polu, S.; Han, J. M.; Zheng, K.; Baksys, M.; Babuschkin, I.; and Sutskever, I. 2022. Formal mathematics statement curriculum learning. *arXiv preprint arXiv:2202.01344*.
- Polu, S.; and Sutskever, I. 2020. Generative language modeling for automated theorem proving. *arXiv preprint arXiv:2009.03393*.
- Premtoon, V.; Koppel, J.; and Solar-Lezama, A. 2020. Semantic code search via equational reasoning. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*, 1066–1082.
- Quan, X.; Valentino, M.; Dennis, L. A.; and Freitas, A. 2024. Verification and Refinement of Natural Language Explanations through LLM-Symbolic Theorem Proving. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y.-N., eds., *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 2933–2958. Miami, Florida, USA: Association for Computational Linguistics.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *arXiv:1910.10683*.
- Sellam, T.; Das, D.; and Parikh, A. P. 2020. BLEURT: Learning robust metrics for text generation. *arXiv preprint arXiv:2004.04696*.
- Shuster, K.; Poff, S.; Chen, M.; Kiela, D.; and Weston, J. 2021. Retrieval augmentation reduces hallucination in conversation. *arXiv preprint arXiv:2104.07567*.
- Song, C. H.; Wu, J.; Washington, C.; Sadler, B. M.; Chao, W.-L.; and Su, Y. 2022. Llm-planner: Few-shot grounded planning for embodied agents with large language models. *arXiv preprint arXiv:2212.04088*.
- Stolfo, A.; Jin, Z.; Shridhar, K.; Schölkopf, B.; and Sachan, M. 2023. A Causal Framework to Quantify the Robustness of Mathematical Reasoning with Language Models. *arXiv:2210.12023*.
- Taylor, R.; Kardas, M.; Cucurull, G.; Scialom, T.; Hartshorn, A.; Saravia, E.; Poulton, A.; Kerkez, V.; and Stojnic, R. 2022. Galactica: A large language model for science. *arXiv preprint arXiv:2211.09085*.
- Toshniwal, S.; Moshkov, I.; Narenthiran, S.; Gitman, D.; Jia, F.; and Gitman, I. 2024. OpenMathInstruct-1: A 1.8 Million Math Instruction Tuning Dataset. *arXiv preprint arXiv:2402.10176*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Trinh, T. H.; Wu, Y.; Le, Q. V.; He, H.; and Luong, T. 2024. Solving olympiad geometry without human demonstrations. *Nature*, 625(7995): 476–482.
- Valentino, M.; and Freitas, A. 2024. On the Nature of Explanation: An Epistemological-Linguistic Perspective for Explanation-Based Natural Language Inference. *Philosophy & Technology*, 37(3): 88.
- Valentino, M.; Meadows, J.; Zhang, L.; and Freitas, A. 2024. Multi-Operational Mathematical Derivations in Latent Space. *arXiv:2311.01230*.
- Villalobos, P.; Sevilla, J.; Heim, L.; Besiroglu, T.; Hobbahn, M.; and Ho, A. 2022. Will we run out of data? An analysis of the limits of scaling datasets in Machine Learning. *arXiv:2211.04325*.
- Welleck, S.; Liu, J.; Bras, R. L.; Hajishirzi, H.; Choi, Y.; and Cho, K. 2021. Naturalproofs: Mathematical theorem proving in natural language. *arXiv preprint arXiv:2104.01112*.
- Welleck, S.; Liu, J.; Lu, X.; Hajishirzi, H.; and Choi, Y. 2022. Naturalprover: Grounded mathematical proof generation with language models. *arXiv preprint arXiv:2205.12910*.
- Wysocka, M.; Wysocki, O.; Delmas, M.; Mutel, V.; and Freitas, A. 2023. Large Language Models, scientific knowledge and factuality: A systematic analysis in antibiotic discovery. *arXiv preprint arXiv:2305.17819*.
- Yang, J.; Jin, H.; Tang, R.; Han, X.; Feng, Q.; Jiang, H.; Yin, B.; and Hu, X. 2023. Harnessing the power of llms in practice: A survey on chatgpt and beyond. *arXiv preprint arXiv:2304.13712*.
- Yao, S.; Yu, D.; Zhao, J.; Shafran, I.; Griffiths, T. L.; Cao, Y.; and Narasimhan, K. 2023. Tree of thoughts: Deliberate problem solving with large language models. *arXiv preprint arXiv:2305.10601*.
- Yuan, Z.; Yuan, H.; Tan, C.; Wang, W.; and Huang, S. 2023. How well do Large Language Models perform in Arithmetic tasks? *arXiv preprint arXiv:2304.02015*.
- Zanibbi, R.; Aizawa, A.; Kohlhase, M.; Ounis, I.; Topic, G.; and Davila, K. 2016. NTCIR-12 MathIR Task Overview. In *NTCIR*.
- Zhong, W.; Yang, J.-H.; and Lin, J. 2022. Evaluating Token-Level and Passage-Level Dense Retrieval Models for Math Information Retrieval. *arXiv preprint arXiv:2203.11163*.