

SylloBio-NLI: Evaluating Large Language Models on Biomedical Syllogistic Reasoning

Magdalena Wysocka^{1,2}, Danilo Carvalho^{1,2}, Oskar Wysocki³,
Marco Valentino³, André Freitas^{1,2,3}

¹Department of Computer Science, University of Manchester, UK
²National Biomarker Centre (NBC), CRUK Manchester Institute, UK
³Idiap Research Institute, Martigny, Switzerland

Abstract

Syllogistic reasoning is crucial for Natural Language Inference (NLI). This capability is particularly significant in specialized domains such as biomedicine, where it can support automatic evidence interpretation and scientific discovery. This paper presents SylloBio-NLI¹, a novel framework that leverages external ontologies to systematically instantiate diverse syllogistic arguments for biomedical NLI. We employ SylloBio-NLI to evaluate Large Language Models (LLMs) on identifying valid conclusions and extracting supporting evidence across 28 syllogistic schemes instantiated with human genome pathways. Extensive experiments reveal that biomedical syllogistic reasoning is particularly challenging for zero-shot LLMs, which achieve an average accuracy between 70% on generalized modus ponens and 23% on disjunctive syllogism. At the same time, we found that few-shot prompting can boost the performance of different LLMs, including Gemma (+14%) and LLama-3 (+43%). However, a deeper analysis shows that both techniques exhibit high sensitivity to superficial lexical variations, highlighting a dependency between reliability, models' architecture, and pre-training regime. Overall, our results indicate that, while in-context examples have the potential to elicit syllogistic reasoning in LLMs, existing models are still far from achieving the robustness and consistency required for safe biomedical NLI applications.

1 Introduction

Syllogistic reasoning – i.e., the process of deriving valid conclusions from premises through the systematic application of abstract reasoning schemes – is a fundamental type of inference for developing and evaluating Natural Language Inference (NLI) models that can reason over textual evidence

¹code and dataset available at: <https://github.com/neuro-symbolic-ai/SylloBio-NLI>

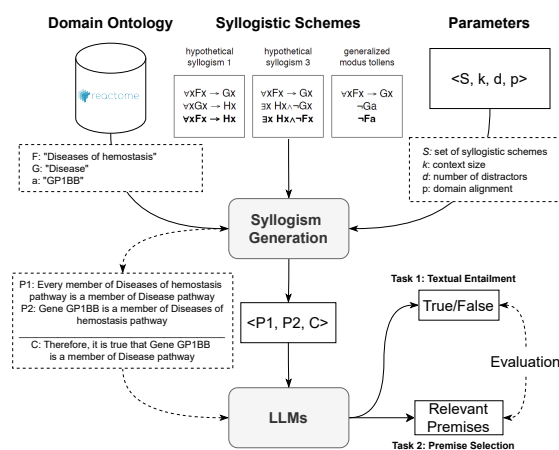


Figure 1: End-to-end diagram of the proposed methodological framework, illustrating the generation of syllogistic arguments from domain-specific ontologies, parameterized input to LLMs, and evaluation tasks including textual inference and premise selection.

at scale (MacCartney and Manning, 2009; Wu et al., 2023). within specialised domains such as biomedicine, the ability to reason over natural language can have significant practical implications for supporting complex discourse interpretation, scientific discovery, and the development of downstream biomedical and clinical applications (Jullien et al., 2023a,b), allowing for a set of formally defined patterns for controlled inference.

Moreover, determining the reliability and robustness of NLI models with regard to syllogistic reasoning is a crucial type of assessment, particularly in critical domains requiring safety guarantees (Eisape et al., 2024; Jullien et al., 2024). This type of evaluation is further motivated by two key factors. First, syllogistic reasoning is prevalent in discourse at large and, therefore, it is expected that approaches based on Large Language Models (LLMs) are exposed to a variety of reasoning schemes during pre-training (Wu et al., 2023). Second, a model that learns to perform syllogistic

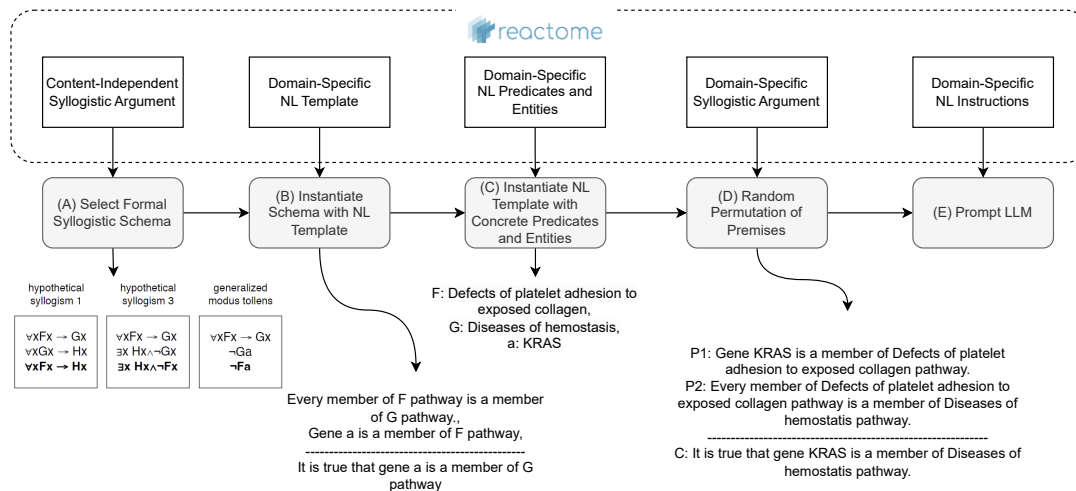


Figure 2: Domain-specific pipeline for creating natural language instances of argument schemes with multiple templating. Steps include selecting a syllogistic schema (A), applying a domain-specific template (B), instantiating with predicates and entities (C), permuting premises (D), and prompting LLMs for evaluation (E).

reasoning should intrinsically possess the ability to generalize to different domains, regardless of the specific world knowledge acquired during pre-training (Kim et al., 2024). This is because the ability to perform syllogistic reasoning should be content-independent, that is, the ability to derive logically valid conclusions is only a function of the formal logical schemes and should be independent of its concrete instantiation. However, despite the importance of syllogistic reasoning for NLI and the abundance of benchmarks involving commonsense knowledge (Yu et al., 2023), resources for assessing how systematic reasoning capabilities transfer to specialised domains are still scarce and require substantial expert-level annotation effort to guarantee quality and correctness (Zhao et al., 2024; Eisape et al., 2024; Porada et al., 2022; Wysocka et al., 2024).

This paper focuses on advancing the availability of resources for biomedical syllogistic reasoning along with our understanding of the capabilities of state-of-the-art NLI models. In particular, we propose SyllBio-NLI, a novel framework for automatically generating NLI resources for evaluating syllogistic reasoning within biomedical domains. Specifically, we demonstrate how external domain-specific resources such as ontologies and thesauri can be leveraged to instantiate a wide range of syllogistic schemes for biomedical NLI tasks, including textual entailment and evidence extraction.

The methodological framework behind SyllBio-NLI is designed to address the annota-

tion scarcity problem for the granular evaluation of complex reasoning in specialized domains. The framework minimises the human annotation effort while guaranteeing the correctness of the generated data by leveraging explicit domain knowledge in the ontologies and the systematicity of known syllogistic schemes (Fig. 1).

By instantiating SyllBio-NLI on human biological pathways using Reactome (Croft et al., 2014), we evaluate the domain-specific reasoning capabilities of 8 open-source LLMs on 28 syllogistic schemes, comparing the performance in a zero-shot (ZS) and few-shot (FS) settings. An extensive empirical evaluation led to the following findings and conclusions:

1. We determine that LLMs exhibit surprisingly low ZS performance on biomedical syllogistic arguments with an average accuracy between 70% on *generalised modus ponens* and 23% on *disjunctive syllogism* (where random performance is equal to 50%). These low performances are generally shared across models’ family (i.e., Llama-3, Mistral, Gemma, BioMistral) and pre-trained regimes (language modelling and instruction-tuning).
2. We found that the FS setting can improve the performance of different LLMs. In particular, we observe a significant boost for both Gemma and Llama-3, with an average increase in F1-score of 14% and 43% respectively. At the same time, the experiments reveal that such improvement is inconsistent across models’ families and pre-training regimes.

3. We perform a robustness analysis adopting logically equivalent variations of the same syllogistic schemes by rephrasing the arguments via negations, complex predicates, and De Morgan’s laws. Such analysis reveals that both ZS and FS techniques are highly sensitive to surface-form and lexical variations, demonstrating a shared inability to systematically abstract the underlying reasoning rules required to derive valid conclusions. These results indicate that, while FS has the potential to elicit syllogistic reasoning in LLMs, existing models are still far from achieving the robustness and consistency required for safe biomedical NLI.

Overall, the above findings suggest that, upon granular inference scrutiny, the reasoning mechanisms induced in LLMs still confound formal and material inference patterns. Moreover, while there are FS intervention mechanisms which can improve models’ performance, delivering controlled specialised syllogistic reasoning remains a challenge for LLMs at large.

To the best of our knowledge, this is the first work focusing on designing a methodology for evaluating syllogistic reasoning within specialised domains, thoroughly assessing the performance of LLMs, and releasing a domain-specific resource for supporting future work in the field. The code for the dataset generation and the evaluation pipeline is fully available online¹.

2 SyllBio-NLI

We introduce a general framework for developing resources to systematically evaluate the syllogistic reasoning capabilities of NLI models within biomedical domains. In particular, SyllBio-NLI leverages the systematicity of known syllogistic schemes (Betz et al., 2021) along with external thesauri/ontologies to scale up the generation of deductively valid arguments and to characterise inference performance across a granular set of syllogistic schemes.

To this end, we focus on seven general syllogistic schemes, including *generalized modus ponens*, *generalized contraposition*, *hypothetical syllogism 1*, *hypothetical syllogism 3*, *generalized modus tollens*, *disjunctive syllogism* and *generalized dilemma*. These base schemes are then adopted to generate 28 different variations of syllogistic argument templates applying negation, complex predicates, and De Morgan’s laws, which can then be instantiated with concrete domain knowl-

edge (see Fig. 6 in Appendix B). The overall methodology is outline by the stages in Fig. 2.

First, for each syllogistic scheme, a corresponding formal argument scheme (consisting of abstract premises and conclusion expressed in first-order logic) is created (Fig. 2A). Next, each symbolic formula in the formal scheme is individually replaced by a natural language domain-specific sentence schema (Fig. 2B). For example, the formulae for the base schema of generalized modus ponens (i.e., Premise 1: $\forall xFx \Rightarrow Gx$, Premise 2: Fa , Conclusion: Ga) can be translated into the natural language template: “*Premise 1: Every member of F is a member of G*”. Premise 2: “*a is a member of F*”. Conclusion: “*a is a member of G*”.

Subsequently, the entity and property placeholders in the natural language template are replaced argument-wise with domain-specific entities and predicates extracted from an external ontology (Fig. 2C). In this process, the syllogistic schemes provide the logical validity component of the formal inference while the ontology subdomain and its associated mapping to the schemes provide the soundness (content-based) for the premises and conclusions. Hence, we obtain instances of syllogistic arguments in natural language that are stored in a knowledge base of premises and conclusions.

Finally, the natural language syllogistic arguments can be structured to create prompts for evaluating LLMs. This involves introducing the syllogistic argument, clearly framing the premises, and specifying the instructions for the inference, as illustrated in Appendix F.

2.1 Biomedical NLI Tasks

SyllBio-NLI takes advantage of the corpus’s design, where every premise and conclusion in each argument is explicitly stated and fully transparent (Fig. 2). Once the full set of natural language syllogistic arguments is generated, the premises and conclusions in each instance can be adopted to define two different NLI tasks:

Task 1. A *textual entailment* task where the model determines whether the conclusion logically follows from the premises, with the output being ‘True’ or ‘False’. The task is designed using natural language syllogisms generated from formal schemes, with valid and invalid argument instances. The model evaluates each argument’s logical validity, requiring it to reason deductively based on the given premises. The focus is on testing the ability

to discern valid reasoning structures in a biomedical context, independent of factual correctness.

Formally, given a set of premises $\mathcal{P} = \{P_1, P_2, \dots, P_n\}$ and a conclusion C , the task is to determine whether the logical entailment holds:

$$\mathcal{P} \models C$$

such that (i) The premises and conclusion are instantiated from formal syllogistic schemes and (ii) arguments may be *valid* or *invalid*. The assessment is purely on logical validity, independent of the factual correctness of premises and conclusion.

Premises: A set $\mathcal{P} = \{P_1, P_2, \dots, P_n\}$ of premises, where each P_i is a natural language sentence derived from syllogistic schemes.

Conclusion: A natural language sentence C , representing the conclusion.

Logical Validity Indicator: A binary output $\text{Output} \in \{\text{True}, \text{False}\}$ such that:

$$\text{Output} = \begin{cases} \text{True} & \text{if } \mathcal{P} \models C, \\ \text{False} & \text{if } \mathcal{P} \not\models C. \end{cases}$$

Here, $\mathcal{P} \models C$ denotes that C is a logical consequence of \mathcal{P} .

Task 2. A *premise selection* task where the model has to identify which premises are necessary and sufficient to justify the conclusion from Task 1, noting that some premises may be irrelevant. The task involves presenting the model with a set of premises, including both relevant and distractor premises. The model must correctly select the subset of premises that logically support the conclusion, excluding those that do not contribute to the entailment. This tests the ability to filter essential information and assess the understanding of the logical relationships between premises and conclusions. Task 2 can be formalized as follows:

Objective: Identify the minimal subset of premises $\mathcal{P}' \subseteq \mathcal{P}$ that are necessary and sufficient to logically derive the conclusion C .

• **Given:**

- A set of premises $\mathcal{P} = \{P_1, P_2, \dots, P_n\}$.
- A conclusion C .

- **Task Definition:** Find the minimal subset $\mathcal{P}' \subseteq \mathcal{P}$ satisfying:

1. **Logical Sufficiency:**

$$\mathcal{P}' \models C.$$

2. **Minimality:**

$$\nexists \mathcal{P}'' \subsetneq \mathcal{P}', \text{ such that } \mathcal{P}'' \models C.$$

3. **Necessity:** Every premise in \mathcal{P}' is necessary for the entailment.

Premises: A set $\mathcal{P} = \{P_1, P_2, \dots, P_n\}$ containing both relevant and irrelevant (distractor) premises.

Conclusion: A natural language sentence C .

- **Selected Premises:** A subset $\mathcal{P}' \subseteq \mathcal{P}$ such that:

$$\mathcal{P}' \models C,$$

and \mathcal{P}' is minimal and necessary.

2.2 Human Genome Pathways for Syllogistic Reasoning Evaluation

To instantiate the SyllBio-NLI methodology, we developed a specialized dataset using Reactome², a comprehensive knowledge base containing detailed information on human biological pathways, including gene functions and interactions.

Reactome’s hierarchical structure, with its well-defined hierarchical gene-pathway membership relations, enables a systematic instantiation of the syllogistic arguments. This allows for the efficient generation of domain-specific NLI tasks that assess LLMs’ ability to reason about biological pathway membership, controlling for different levels within the hierarchy. This use case is widely relevant in the context of interpreting pathway-level interactions, disease and treatment response mechanisms at a genomics level (Fang et al., 2019).

By focusing on the Disease super-pathway, the largest and most intricate group within Reactome, we generated a diverse set of syllogistic schemes (see Fig. 2) that mirror the complex molecular interactions and regulatory mechanisms underlying disease processes. This approach ensures our evaluation reflects the depth and specificity needed to rigorously test the syllogistic reasoning properties of NLI models in this highly specialised and clinically significant domain. A comprehensive

²<https://reactome.org>

Level of pathway	Nr of unique pathway names	Nr of unique genes names
top-level	1	
sub-pathway	13	2131
sub-sub-pathway	59	2131
sub-sub-sub-pathway	345	2114
sub-sub-sub-sub-pathway	1110	1775
sub-sub-sub-sub-sub-pathway	994	1002

Table 1: Summary of a dictionary with the true taxonomic relationships between pathways and genes.

and systematic assessment of these properties is essential if LLMs are to be applied to this domain.

The resulting dataset includes 12,098 entity-gene names, which are used as substitutes for *entity* placeholders. In contrast, pathway names, derived from different levels of biological hierarchies, are treated as *predicate* placeholders because they describe relationships or actions involving the entities, rather than being entities themselves. The corpus includes 3767 complex predicates in total. The premises of the natural language argument are randomly re-ordered to mitigate potential biases during the evaluation (Fig. 2D).

2.3 Dictionary of Gene and Pathway Membership

Reactome³ (version 88—March 2024) has entries for 11 226 protein-coding genes involved in 15 212 human reactions annotated from 38 549 literature references. These reactions are grouped into 2 698 pathways collected under 29 superpathways (e.g. Immune System) that describe normal cellular functions. Each superpathway is represented as a roughly circular ‘burst,’ with the central node corresponding to the top-level of the Reactome event hierarchy and concentric rings representing increasingly specific levels of the event hierarchy (sub-pathways) (e.g. Disease → Diseases of signal transduction by growth factor receptors and second messengers → Signalling by EGFR in Cancer → Signalling by Ligand-Responsive EGFR Variants in Cancer → Constitutive Signalling by Ligand-Responsive EGFR Cancer Variants). The relationships between these pathways are captured through parent-child arcs, reflecting the ontological “is-a” relationships.

The 29 Reactome superpathways group are each organized as a roughly circular ‘burst’. However, we built the corpus based on one, largest group of pathways called Disease. The central node of the Disease burst corresponds to the uppermost

level of the Reactome event hierarchy (Table 1). Concentric rings of nodes around the central node represent successive more specific levels of the event hierarchy (e.g. Disease → Diseases of signal transduction by growth factor receptors and second messengers → Signalling by EGFR in Cancer → Signalling by Ligand-Responsive EGFR Variants in Cancer → Constitutive Signalling by Ligand-Responsive EGFR Cancer Variants). The arcs connecting nodes between successive rings within a burst represent parent–child (is-a) relationships in the event hierarchy. When a specific pathway is shared by more than one burst, arcs connect its nodes between bursts. A node’s size is proportional to the number of physical entities (proteins, complexes, chemicals) it contains.

2.4 Reasoning Challenges

Domain-specific Reasoning. Biomedical datasets differ from general datasets in that they involve complex semantic structures, such as gene-pathway relationships and hierarchies, where sentences encode intricate biological interactions and dependencies. Reasoning about statements such as “*COL1A1 is involved in the TGF-beta signalling pathway*” requires not only an understanding of specific gene functions but also the ability to navigate over hierarchical biological pathways and their associated relations. This demands a higher level of domain-specific reasoning and the capability to controllably interpret relationships within a highly specialised context.

Material Inference Component. Each sentence in the formal scheme (premises and conclusion) is linked to the domain knowledge, i.e. the membership of genes to a given level of pathway, which refers to a multi-level, hierarchy. For example, if gene X is a member of pathway Y, and that pathway is a child of a top-level pathway Z, the model must be able to infer that the gene also belongs to a top-level pathway.

Formal Inference Component. Syllogisms are examples of content-independent formal inference in which, provided the truth value of the premises, the conclusion can be derived through the application of specific logical inference rules. This form of reasoning can pose challenges to models that are incapable of abstracting inference patterns from text and that are affected by content biases in their representation (Prange et al., 2023; Kim et al., 2024). Moreover, the syllogistic schemes in this paper

³<https://reactome.org>

Model	ZS									FS						
	Non-empty output	Irrelevant text	Following instruction	Acc.	Recall	Precision	F1	Faith.	Non-empty output	Irrelevant text	Following instruction	Acc.	Recall	Precision	F1	Faith.
BioMistral-7B	0.00	-	-	-	-	-	-	-	0.00	-	-	-	-	-	-	-
Meta-Llama-3-8B	0.03	0.03	0.01	0.00	0.01	0.01	0.01	0.00	0.69	0.01	0.68	0.35	0.51	0.39	0.44	0.19
Meta-Llama-3-8B-Instruct	1.00	0.86	0.14	0.11	0.18	0.16	0.17	0.04	1.00	0.99	0.01	0.00	0.00	0.00	0.00	0.00
Mistral-7B-Instruct-v0.2	1.00	0.1	0.9	0.47	0.63	0.48	0.54	0.48	1.00	0.24	0.76	0.45	0.44	0.45	0.44	0.46
Mistral-7B-v0.1	0.88	0.88	0.00	0.00	0.00	0.00	0.00	0.00	0.69	0.69	0.00	0.00	0.00	0.00	0.00	0.00
Mixtral-8x7B-Instruct-v0.1	1.00	0.37	0.63	0.52	0.45	0.53	0.49	0.35	1.00	0.09	0.91	0.61	0.44	0.67	0.53	0.46
Gemma-7b	0.99	0.2	0.79	0.42	0.76	0.45	0.57	0.09	1.00	0.00	1.00	0.60	1.00	0.55	0.71	0.19
Gemma-7b-it	1.00	0.00	1.00	0.64	0.71	0.62	0.66	0.66	1.00	0.03	0.97	0.63	0.46	0.69	0.55	0.42

Model	ZS					FS				
	Non-empty output	Irrelevant text	Following instruction	RA	Faith.	Non-empty output	Irrelevant text	Following instruction	RA	Faith.
BioMistral-7B	0.00	0.00	-	-	-	0.00	-	-	-	-
Meta-Llama-3-8B	0.02	0.018	0.002	0.00	0.00	0.87	0.01	0.87	0.35	0.31
Meta-Llama-3-8B-Instruct	1.00	0.36	0.64	0.30	0.33	1.00	0.99	0.01	0.01	0.00
Mistral-7B-Instruct-v0.2	1.00	0.02	0.98	0.26	0.64	1.00	0.05	0.95	0.34	0.65
Mistral-7B-v0.1	1.00	1.00	0.00	0.00	0.00	0.84	0.84	0.00	0.00	0.00
Mixtral-8x7B-Instruct-v0.1	1.00	0.05	0.95	0.34	0.63	1.00	0.05	0.95	0.24	0.46
Gemma-7b	0.69	0.31	0.38	0.23	0.14	1.00	0.00	1.00	0.58	0.69
Gemma-7b-it	1.00	0.03	0.97	0.55	0.65	1.00	0.00	1.00	0.55	0.31

Table 2: Main results for Task 1 (i.e., textual entailment, top) and Task 2 (i.e., premise selection, bottom).

are designed to assess the interpretation of fine-grained logical operators including quantifiers, implications and negation, which have been proved to be particularly challenging for NLI models (Pitler et al., 2023).

3 Empirical Evaluation

3.1 Model Architectures

We used the proposed methodology and resources to assess the syllogistic NLI inference properties of eight open-source LLMs.

Specifically, We test a range of architectures, including mistralai/Mistral-7B-v0.1, mistralai/Mistral-7B-Instruct-v0.2 (Jiang et al., 2023), mistralai/Mixtral-8x7B-Instruct-v0.1 (Mistral, 2023), google/gemma-7b, google/gemma-7b-it (Gemma and Google, 2024), meta-llama/Meta-Llama-3-8B, meta-llama/Meta-Llama-3-8B-Instruct (AI@Meta, 2024), BioMistral/BioMistral-7B (Labrak et al., 2024). Details of the models, access, parameters, and the prompts used are available in the Appendix C, D, F, G.

3.2 Evaluation Metrics

Accuracy, F1-score, Recall, and Precision are used to evaluate performance in Task 1, a binary classification task where the conclusion C is labelled as either True or False. These metrics compare the predicted labels \hat{y}_i against the annotated gold labels y_i (True/False).

We used Reasoning Accuracy RA to evaluate the models’ performance in Task 2 to correctly predict the entailment label \hat{y}_i and select the appropriate subset of premises \hat{P}_i that justify this

prediction. The metric is calculated as the percentage of responses where both the predicted label \hat{y}_i matches the ground truth y_i and the selected premises \hat{P}_i match the ground truth premises P_i . In addition, we use a measure of faithfulness introduced by Jullien et al. (2024), which assesses the ability of a model to correctly change its prediction when altering the truth value of the conclusion (see Appendix E for additional details).

3.3 Main Results and Discussion

The main results on Task 1 and 2 are reported in Table 2. In particular, it is possible to derive the following main conclusions:

ZS LLMs struggle with controlled, domain-specific syllogistic inference. Overall, the results demonstrate that LLMs exhibit significant challenges when performing biomedical syllogistic reasoning in a ZS setting (Table 2). We found, in fact, that the majority of the models struggle to achieve performances that are above the 50% random accuracy on Task 1, with Gemma-7B-it being the only exception, achieving an accuracy of 64%. Similarly, in Task 2, Gemma-7B-it was again the best-performing model, achieving *reasoning accuracy* of 55%. In general, when comparing the performance on the two tasks, we observed that the models achieved lower performance in the premise selection task compared to textual entailment, indicating a shared inability of ZS models to identify the sufficient and necessary set of premises required for the entailment to hold. In addition to low performance, we found a discrepancy between faithfulness and accuracy on both tasks, which

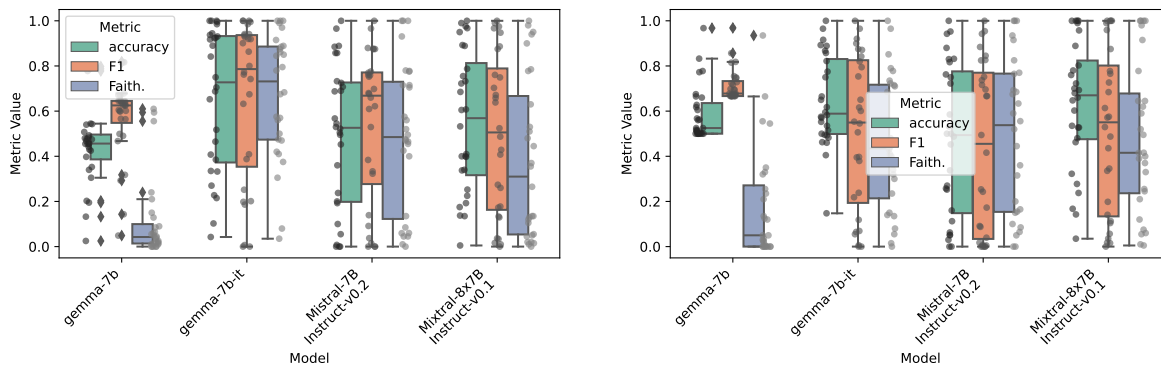


Figure 3: Comparative Analysis of accuracy, F1, and Faithfulness across two prompt types: ZS (left) and FS (right) for Task 1 for the four best models.

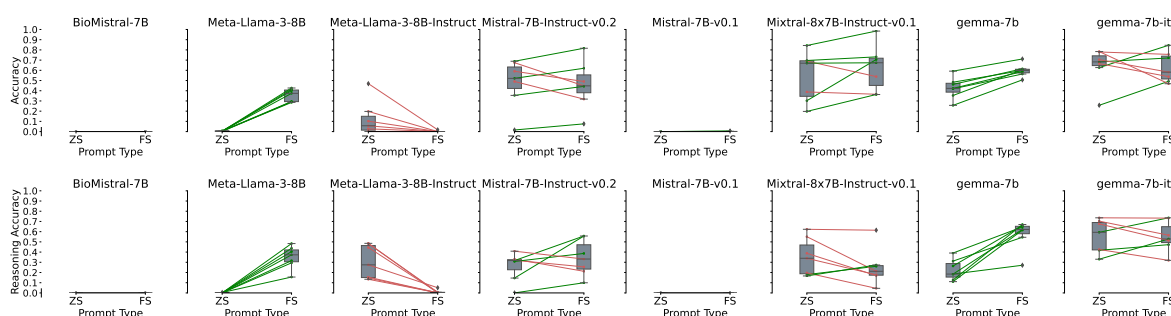


Figure 4: Accuracy across two prompt types: ZS and FS, for Task 1 (top) and Task 2 (bottom). The lines connect the average accuracy for each of the seven syllogistic argument schemes, with green lines indicating an increase and red lines indicating a decrease. Gray boxplots display the median, Q1, Q3, and minimum and maximum values.

highlights the inability to capture the underlying reasoning phenomena along with the presence of biases in the inference process. For example, we observed that for Gemma-7B, the ZS accuracy and F1-score are ≈ 0.5 - 0.6 , while faithfulness plummets below 0.1 (Fig. 3). This is due to the tendency of the model to frequently output "True" (see Figs. 8-11 in Appendix K and Appendix H) regardless of the truth value of the conclusion. In contrast, we observed that such a phenomenon is less prominent in Gemma-7B-it where faithfulness is comparable to accuracy on both Task 1 (i.e., 0.66) and Task 2 (0.65). Overall, the results reveal that instruction-tuned models can obtain higher performances in the ZS setting, while models that are pre-trained on language modelling struggle to follow the provided natural language instructions. Moreover, we observed that this trend is independent of the specific knowledge the LLM is exposed to during pre-training. For instance, BioMistral-7B, the only biomedical domain-specific model, entirely failed to follow the instructions and generate outputs that are relevant to the target tasks.

FS can improve performance, but its impact is inconsistent across models' families and pre-training regimes. We found that the FS setting can improve the performance of different LLMs. In particular, we observe a significant boost for both Gemma-7b and Meta-Llama-3-8b, with an average increase in F1-score of 14% and 43% respectively. However, we observe such improvement to be inconsistent across models' family and pre-training regime, indicating variability in how different models handle prompt types in syllogistic reasoning tasks within the biomedical domain (Fig. 4). In particular, the results show that FS improves for all schemes only for Gemma-7b and Llama-3. For Mistral-8x7B Instruct, Mistral-7B-Instruct, and Gemma-7b-it the overall *accuracy* increases only for a subset of selected schemes, while for others it drops significantly. This inconsistency is demonstrated by the fact that Mistral-8x7B Instruct achieves worse accuracy in FS compared to ZS on both tasks. Overall, the results on the FS setting highlight distinct strengths and weaknesses that are highly dependent on the prompt

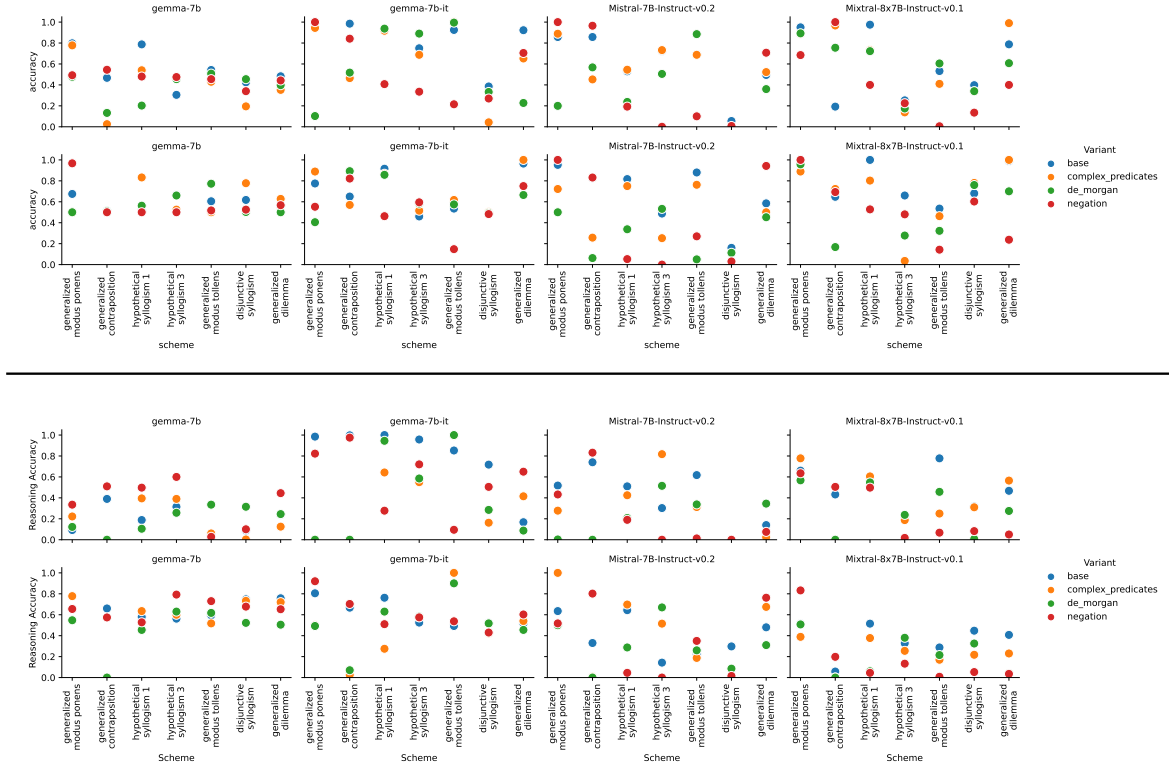


Figure 5: Robustness to lexical variations. The graphs show the accuracy by model, syllogistic schema and lexical variations for ZS and FS respectively for Task 1 (top) and Task 2 (bottom).

strategy, the pre-training setup, and the model’s family. At the same time, improvements observed in FS for some models might indicate a potential for enhanced performance when contextual information is available, although this gain is not universal across the models.

Significant variability in model accuracy is observed across different syllogistic schemes. In general, we found that LLMs exhibit variable performances depending on the specific syllogistic schema used for evaluation. The results indicate that LLMs perform better on *generalized modus ponens* (i.e., three out of four models), with *accuracy* levels ranging from 0.71 to 0.98 in the FS setting in Task 1 (Fig. 5, Table 5 in Appendix L) and with *reasoning accuracy* scores ranging from 0.56 to 0.74 in Task 2 (Table 7 in Appendix L). On the other side, In the ZS setting, the Gemma-7B-it model achieved higher performance on *generalized modus tollens* (accuracy = 0.78 in Task 1 and RA = 0.74 in Task 2), while in the FS setting, the highest performances were registered on *generalized dilemma* in Task 1 (accuracy = 0.85) and *generalized modus ponens* in Task 2 (RA = 0.74).

3.4 Robustness to Lexical Variations

We performed a robustness analysis adopting surface-form variations applied to the syllogistic schemes. This is done by rephrasing the natural language syllogistic arguments via negation, complex predicates, and De Morgan’s laws, keeping the underlying logical relation between premises and conclusion unaltered (see Fig. 6 in Appendix B). The results of the analysis are reported in Figure 5.

This intervention reveals that LLMs are highly sensitive to surface forms and lexical variations, showing a significant variability across different schemes (Fig. 5). At the same time, we found that in the FS setting, model responses demonstrated overall greater consistency, with lower fluctuations across variants.

In Task 1, *disjunctive syllogism* was generally identified as the most challenging scheme in the ZS setting for all four models, with none exceeding an accuracy of 0.35, significantly worse than random guessing. While FS slightly improved accuracy, disjunctive syllogism remained the most challenging scheme. Overall, only Gemma-7B in the FS setting effectively handled all syllogistic schemes and their variants in Task 1, demonstrating a higher level of consistency. The other models exhibited

significant variability depending on the specific surface form, which highlights a shared inability to systematically abstract the underlying reasoning rules required to derive valid conclusions.

3.5 Impact of Distractors and Factuality

Finally, we perform an analysis introducing an increasing number of distractors (i.e., irrelevant premises) in the syllogistic arguments and replacing existing genes in Reactome with synthetic gene names to generate arguments that are independent of medical content and assess the impact of factuality on reasoning (see Appendix I, J).

We found that LLMs show varied sensitivity to distractors, with models like Gemma-7b exhibiting a significant decline in reasoning accuracy as the number of distractors increases, while Mistral-7B Instruct shows improvements in some cases. The impact is scheme-dependent and reveals that more complex syllogisms are particularly affected by increasing distracting information (Figs. 12-15 in Appendix K).

Furthermore, contrary to previous work showing that LLMs exhibit higher performance on syllogisms that are in line with commonsense knowledge (Eisape et al., 2024; Kim et al., 2024), we found stable performances when intervening on the factual correctness of the biomedical arguments. These results indicate that (1) logical structure and contextual information have a greater impact on output generation than biomedical knowledge in both ZS and FS settings and (2) existing models might have limited exposition to human genome pathways information during pre-training, being unaffected by the substitution of real gene names with synthetic ones (Fig. 7 in Appendix J).

4 Related Work

Syllogistic reasoning, which involves deriving valid conclusions from given premises based on formal logical structures, has long been a central focus in the study of Natural Language Inference (NLI) (Bertolazzi et al., 2024; Kim et al., 2024).

Early transformer models like BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019) were trained on general NLI datasets, such as SNLI (Bowman et al., 2015) and MNLI (Williams et al., 2018), to address reasoning tasks. However, these models were limited in their ability to generalize abstract reasoning patterns, particularly in more specialized domains. Recent work has explored

whether larger, decoder-based language models are capable of capturing syllogistic reasoning without task-specific fine-tuning, as reasoning itself is content-independent (Bertolazzi et al., 2024).

Recent research, such as Liu et al. (Liu et al., 2023), has introduced benchmarks like GLoRE to test logical reasoning in LLMs, utilizing techniques such as Chain-of-Thought (CoT) prompting (Wei et al., 2023) and in-context learning (ICL) (Huang and Chang, 2023). These methods have been shown to improve reasoning performance, but issues remain, as models often rely on superficial patterns instead of deep logical comprehension. Eisape et al. (Eisape et al., 2024) further demonstrated that even advanced models exhibit biases when handling syllogisms, particularly when faced with invalid syllogisms, often failing to generate "nothing follows" conclusions.

While the general-domain performance of LLMs in syllogistic reasoning has been explored (Eisape et al., 2024; Dasgupta et al., 2024), research on syllogistic reasoning in the biomedical domain remains scarce. This paper addresses this gap by focusing specifically on biomedical applications, where logical reasoning over specialized knowledge is critical. We introduce a large-scale biomedical syllogism dataset, SyllBio-NLI, and provide a comprehensive evaluation of the syllogistic reasoning capabilities of state-of-the-art LLMs within this domain.

5 Conclusions

In this work, we proposed a novel methodological framework, SyllBio-NLI, designed to evaluate the syllogistic reasoning capabilities of state-of-the-art LLMs within the biomedical domain. Through comprehensive analysis across 28 syllogistic schemes, we assessed the performance of eight different models under varying conditions, including zero-shot and few-shot settings.

Our results show that both techniques exhibit high sensitivity to superficial lexical variations, highlighting a dependency between reliability, models' architecture, and pre-training regime. Overall, our evaluation indicates that, while few-shot strategies have the potential to elicit syllogistic reasoning in LLMs, existing models are still far from achieving the robustness and consistency required for safe biomedical NLI applications in specialised domains.

6 Limitations

A key challenge for scaling our approach to different domains is its dependency on high-quality external ontologies and knowledge bases. This factor limits the scope of our analyses across biomedical domains. More efficient methods for populating the natural language syllogistic arguments could be investigated in future work, involving automated NLP methods, such as those used in RepoDB (Brown and Patel, 2017), MSI, Hetionet (Himmelstein et al., 2017), DrugMechDB (Gonzalez-Cavazos et al., 2023), and INDRA (Gyori et al., 2017; Bachman et al., 2023), or synthetic data generation methods coupled with efficient quality checks. However, these approaches still face challenges in balancing precision and generalization, particularly for complex reasoning tasks in biomedicine. Further improvements are necessary to develop scalable resources and more adaptable NLP techniques for real-world applications.

Acknowledgments

This work was funded by the European Union’s Horizon 2020 research and innovation programme under grant agreement No 965397 through the cancer core Europe DART project, by the Swiss National Science Foundation (SNSF) project “NeuMath: Neural Discourse Inference over Mathematical Texts” (200021_204617), by the CRUK National Biomarker Centre, and supported by the Manchester Experimental Cancer Medicine Centre and the NIHR Manchester Biomedical Research Centre.

References

- AI@Meta. 2024. [Llama 3 model card](#).
- John A Bachman, Benjamin M Gyori, and Peter K Sorger. 2023. Automated assembly of molecular mechanisms at scale from text mining and curated databases. *Molecular Systems Biology*, 19(5):e11325.
- Leonardo Bertolazzi, Albert Gatt, and Raffaella Bernardi. 2024. [A systematic analysis of large language models as soft reasoners: The case of syllogistic inferences](#). *Preprint*, arXiv:2406.11341.
- Gregor Betz, Christian Voigt, and Kyle Richardson. 2021. [Critical thinking for language models](#). In *Proceedings of the 14th International Conference on Computational Semantics (IWCS)*, pages 63–75. Association for Computational Linguistics.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Adam S Brown and Chirag J Patel. 2017. A standard database for drug repositioning. *Scientific data*, 4(1):1–7.
- David Croft, Antonio Fabregat, Kristina S Keays, Mark Rutherford, Robin HW Van Iersel, Esteban D’Eustachio, Henning Hermjakob, Lincoln Stein, and Peter D’Eustachio. 2014. Reactome: a knowledgebase of biological pathways. *Nucleic acids research*, 42(D1):D472–D477.
- Ishita Dasgupta, Andrew K. Lampinen, Stephanie C. Y. Chan, Hannah R. Sheahan, Antonia Creswell, Dharshan Kumaran, James L. McClelland, and Felix Hill. 2024. [Language models show human-like content effects on reasoning tasks](#). *Preprint*, arXiv:2207.07051.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Tiwalayo Eisape, Michael Tessler, Ishita Dasgupta, Fei Sha, Sjoerd Steenkiste, and Tal Linzen. 2024. [A systematic comparison of syllogistic reasoning in humans and language models](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 8425–8444, Mexico City, Mexico. Association for Computational Linguistics.
- Gang Fang, Wen Wang, Vanja Paunic, Hamed Heydari, Michael Costanzo, Xiaoye Liu, Xiaotong Liu, Benjamin VanderSluis, Benjamin Oatley, Michael Steinbach, et al. 2019. Discovering genetic interactions bridging pathways in genome-wide association studies. *Nature communications*, 10(1):4274.
- Team Gemma and Deepmind Google. 2024. [Gemma: Open models based on gemini research and technology](#).
- Adriana Carolina Gonzalez-Cavazos, Anna Tanska, Michael Mayers, Denise Carvalho-Silva, Brindha Sridharan, Patrick A Rewers, Umasri Sankarlal, Lakshmanan Jagannathan, and Andrew I Su. 2023. Drugmechdb: A curated database of drug mechanisms. *Scientific Data*, 10(1):632.

- Benjamin M Gyori, John A Bachman, Kartik Subramanian, Jeremy L Muhlich, Lucian Galescu, and Peter K Sorger. 2017. From word models to executable models of signaling networks using automated assembly. *Molecular systems biology*, 13(11):954.
- Daniel Scott Himmelstein, Antoine Lizee, Christine Hessler, Leo Brueggeman, Sabrina L Chen, Dexter Hadley, Ari Green, Pouya Khankhanian, and Sergio E Baranzini. 2017. Systematic integration of biomedical knowledge prioritizes drugs for repurposing. *Elife*, 6:e26726.
- Jie Huang and Kevin Chen-Chuan Chang. 2023. [Towards reasoning in large language models: A survey](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1049–1065, Toronto, Canada. Association for Computational Linguistics.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Mael Jullien, Marco Valentino, and André Freitas. 2024. [SemEval-2024 task 2: Safe biomedical natural language inference for clinical trials](#). In *Proceedings of the 18th International Workshop on Semantic Evaluation (SemEval-2024)*, pages 1947–1962, Mexico City, Mexico. Association for Computational Linguistics.
- Mael Jullien, Marco Valentino, Hannah Frost, Paul O’Regan, Dónal Landers, and Andre Freitas. 2023a. [NLI4CT: Multi-evidence natural language inference for clinical trial reports](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 16745–16764, Singapore. Association for Computational Linguistics.
- Maël Jullien, Marco Valentino, Hannah Frost, Paul O’regan, Donal Landers, and André Freitas. 2023b. [SemEval-2023 task 7: Multi-evidence natural language inference for clinical trial data](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 2216–2226, Toronto, Canada. Association for Computational Linguistics.
- Geonhee Kim, Marco Valentino, and André Freitas. 2024. A mechanistic interpretation of syllogistic reasoning in auto-regressive language models. *arXiv preprint arXiv:2408.08590*.
- Yanis Labrak, Adrien Bazoge, Emmanuel Morin, Pierre-Antoine Gourraud, Mickael Rouvier, and Richard Dufour. 2024. [Biomistral: A collection of open-source pretrained large language models for medical domains](#). *Preprint*, arXiv:2402.10373.
- Hanmeng liu, Zhiyang Teng, Ruoxi Ning, Jian Liu, Qiji Zhou, and Yue Zhang. 2023. [Glore: Evaluating logical reasoning of large language models](#). *Preprint*, arXiv:2310.09107.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *ArXiv*, abs/1907.11692.
- Bill MacCartney and Christopher D Manning. 2009. An extended model of natural logic. In *Proceedings of the eighth international conference on computational semantics*, pages 140–156. Association for Computational Linguistics.
- AI Team Mistral. 2023. [Mixtral of experts - a high quality sparse mixture-of-experts](#).
- Emily Pitler, Chris Callison-Burch, and Ellie Pavlick. 2023. The logic of learning: Syllogistic reasoning in modern ai systems. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Ian Porada, Alessandro Sordoni, and Jackie Cheung. 2022. [Does pre-training induce systematic inference? how masked language models acquire commonsense knowledge](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4550–4557, Seattle, United States. Association for Computational Linguistics.
- Jakob Prange, Khalil Mrini, and Noah A. Smith. 2023. Challenges and opportunities in nlp for systematic generalization: Beyond compositionality. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. [Chain-of-thought prompting elicits reasoning in large language models](#). *Preprint*, arXiv:2201.11903.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Yongkang Wu, Meng Han, Yutao Zhu, Lei Li, Xinyu Zhang, Ruofei Lai, Xiaoguang Li, Yuanhang Ren, Zhicheng Dou, and Zhao Cao. 2023. [Hence, socrates is mortal: A benchmark for natural language syllogistic reasoning](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2347–2367, Toronto, Canada. Association for Computational Linguistics.

Magdalena Wysocka, Oskar Wysocki, Maxime Delmas, Vincent Mutel, and André Freitas. 2024. Large language models, scientific knowledge and factuality: A framework to streamline human expert evaluation. *Journal of Biomedical Informatics*, page 104724.

Fei Yu, Hongbo Zhang, Prayag Tiwari, and Benyou Wang. 2023. [Natural language reasoning, a survey](#). Preprint, arXiv:2303.14725.

Zehui Zhao, Laith Alzubaidi, Jinglan Zhang, Ye Duan, and Yuantong Gu. 2024. [A comparison review of transfer learning and self-supervised learning: Definitions, applications, advantages and limitations](#). *Expert Systems with Applications*, 242:122807.

A Formalization of the SyllBio-NLI Resource Generation Process

This appendix formalises the generation process of the syllogistic inference patterns.

We start by defining the main constructs (formal and linguistic artefacts and functions) of the underlying framework:

1. **Syllogistic Scheme (S):** A logical inference pattern consisting of premises and a conclusion, $S = \{P_1, P_2, \dots, P_n, C\}$, where P_i is premise i and C is the conclusion.
2. **Formal Argument Scheme (σ):** Representation of a syllogistic scheme in first-order logic (FOL), $\sigma(S) = \{\phi_1, \phi_2, \dots, \phi_n, \psi\}$, where ϕ_i corresponds to P_i and ψ corresponds to C .
3. **Natural Language Template (τ):** A natural language schema mapping each formula in $\sigma(S)$ to a sentence template, $\tau(\sigma(S)) = \{\tau_1, \tau_2, \dots, \tau_n, \sigma\}$, where τ_i is the sentence template for ϕ_i and σ is the sentence template for ψ .
4. **Ontology (O):** A domain-specific knowledge base containing entities E and predicates Π , $O = \{E, \Pi\}$, where $E = \{e_1, e_2, \dots, e_k\}$ and $\Pi = \{\pi_1, \pi_2, \dots, \pi_l\}$.
5. **Instantiation Function (I):** A function that replaces placeholders in τ with entities and predicates from O , $I : \tau(\sigma(S)) \times O \rightarrow \text{NL}$, where NL is the set of natural language sentences.
6. **Expert Mapping Function (μ_{Expert}):** A function provided by a domain expert to map placeholders to appropriate ontology terms, $\mu_{\text{Expert}} : \text{Placeholders} \rightarrow E \cup \Pi$.

7. **Knowledge Base (KB):** A collection of instantiated syllogistic arguments, $\text{KB} = \{A_1, A_2, \dots, A_m\}$, where $A_i = \{P'_1, P'_2, \dots, P'_n, C'\}$ and P'_i, C' are instantiated natural language sentences.

A.1 Process Formalisation

The process formalisation defines a systematic process for generating domain-specific syllogistic arguments by:

1. **Defining** formal representations of syllogistic schemes in first-order logic.
2. **Generating** natural language templates from these formal representations.
3. **Mapping** placeholders to domain-specific entities and predicates using an ontology and expert knowledge.
4. **Instantiating** the templates to produce logically valid and semantically sound arguments.
5. **Constructing** a knowledge base for evaluating NLI models.

This ensures that the generated arguments are both logically valid and contextually relevant to the biomedical domain.

Input: A set of syllogistic schemes: $\mathcal{S} = \{S_1, S_2, \dots, S_m\}$, an ontology: $O = \{E, \Pi\}$, an expert mapping function: μ_{Expert} .

Output: A knowledge base of instantiated arguments: KB.

Step 1: Formal Argument Scheme Selection:

For each syllogistic scheme $S_i \in \mathcal{S}$, define its formal argument scheme in first-order logic:

$$\forall S_i \in \mathcal{S}, \quad \sigma(S_i) = \{\phi_1^i, \phi_2^i, \dots, \phi_n^i, \psi^i\}.$$

Step 2: Natural Language Template Generation:

Transform each formula in $\sigma(S_i)$ into a natural language template:

$$\forall \phi_j^i \in \sigma(S_i), \quad \tau_j^i = \tau(\phi_j^i),$$

$$\sigma^i = \tau(\psi^i).$$

Step 3: Ontology Mapping and Instantiation:

Apply the expert mapping function to select appropriate entities and predicates from the ontology:

$\forall \text{Placeholder } p \in \{\mathbf{F}, \mathbf{G}, \mathbf{a}\}, \quad \mu_{\text{Expert}}(p) \rightarrow E \cup \Pi.$

Instantiate the templates:

$$P'_j = I(\tau_j^i, \mu_{\text{Expert}}),$$

$$C' = I(\sigma^i, \mu_{\text{Expert}}).$$

under the following constraints:

- **Logical Validity:** The instantiated arguments must preserve the logical structure of $\sigma(S_i)$.
- **Domain Soundness:** The selected entities and predicates must be semantically coherent within the targeted subdomain.

These constraints can be further formalised as:

Logical Validity Constraint: The instantiated argument A_i must be logically valid:

$$\{\phi'_1, \phi'_2, \dots, \phi'_n\} \models \psi',$$

where ϕ'_j corresponds to the logical form of P'_j .

Domain Soundness Constraint: The entities and predicates used must be semantically valid within the domain:

$$\forall e \in E', \pi \in \Pi', \quad \text{DomainValid}(e, \pi) = \text{True},$$

where $E' \subseteq E$ and $\Pi' \subseteq \Pi$ are entities and predicates used in A_i .

Verification of Logical Validity: Ensure that the instantiated premises logically entail the conclusion:

$$\{\phi'_1, \phi'_2, \dots, \phi'_n\} \models \psi',$$

using logical inference rules.

Verification of Domain Soundness: Confirm that:

- All entities and predicates are correctly used.
- There are no semantic contradictions.

Step 4: Knowledge Base Construction: Aggregate all instantiated arguments into the knowledge base:

$$\text{KB} = \bigcup_{i=1}^m \{A_i\},$$

where:

$$A_i = \{P'_1, P'_2, \dots, P'_n, C'\}.$$

This is summarised with the following algorithmic outline:

1. Initialize $\text{KB} \leftarrow \emptyset$.
2. For each $S_i \in \mathcal{S}$:
 - (a) Create $\sigma(S_i) = \{\phi_1^i, \phi_2^i, \dots, \phi_n^i, \psi^i\}$.
 - (b) Generate $\tau(S_i) = \{\tau_1^i, \tau_2^i, \dots, \tau_n^i, \sigma^i\}$, where $\tau_j^i = \tau(\phi_j^i)$.
 - (c) Obtain mappings μ_{Expert} for placeholders in τ_j^i and σ^i .
 - (d) Instantiate:

$$P'_j = I(\tau_j^i, \mu_{\text{Expert}}),$$

$$C' = I(\sigma^i, \mu_{\text{Expert}}).$$

- (e) Form argument $A_i = \{P'_1, P'_2, \dots, P'_n, C'\}$.

- (f) Add A_i to KB:

$$\text{KB} \leftarrow \text{KB} \cup \{A_i\}.$$

3. Return KB.

B Domain-specific pipeline for creating NL instances

Figure 6: Syllogistic argument schemes used to create a biologically factual argument corpus with domain-specific examples for *generalized modus ponens base scheme* and *disjunctive syllogism complex predicates scheme*. For each syllogistic scheme, a formal argument scheme (consisted of premises and conclusion (bold)) was provided.

C Accessing LLMs

To access these LLMs, we use the Mistral AI (mistral), and the open-source weights of the remaining models, available at the *HuggingFace Hub*⁴ repositories:

⁴<https://huggingface.co>

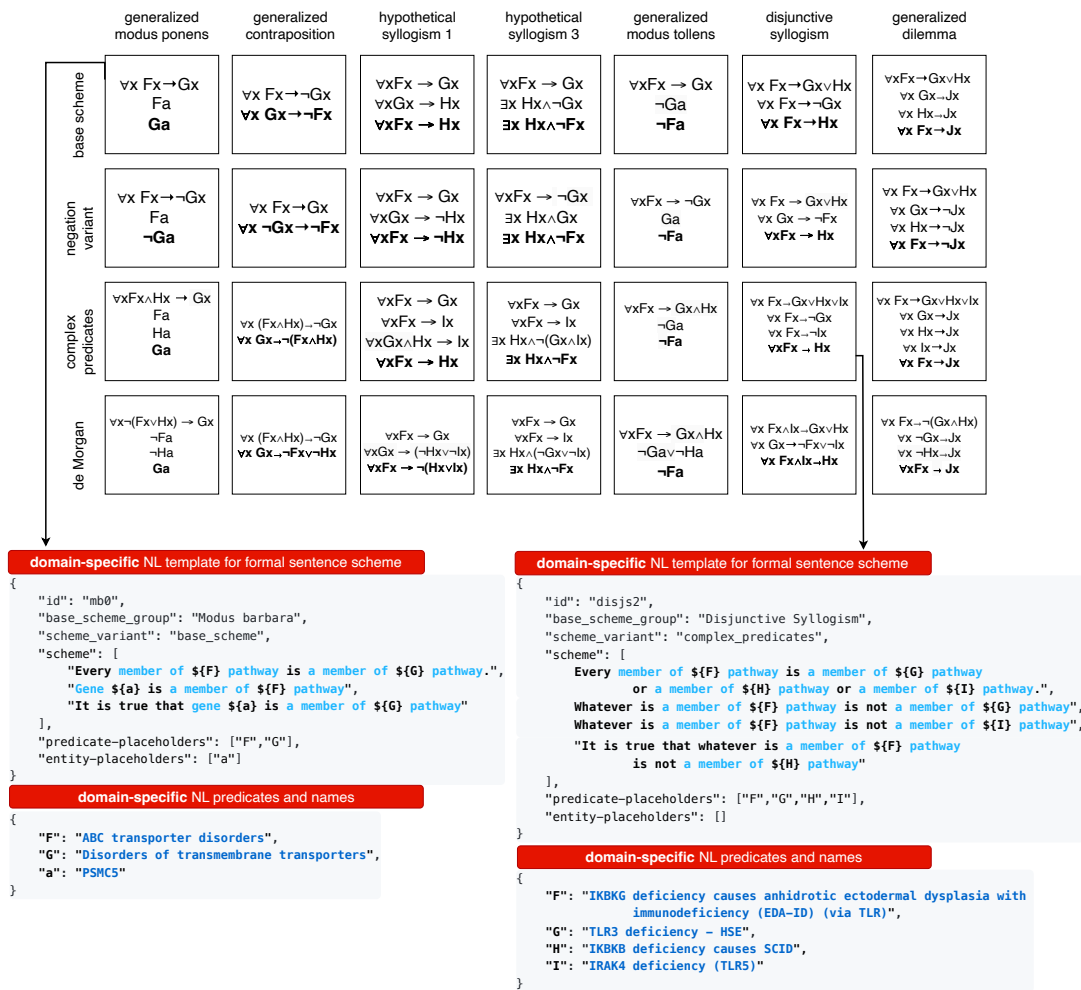


Figure 6: Syllogistic argument schemes used to create a biologically factual argument corpus with domain-specific examples for *generalized modus ponens base scheme* and *disjunctive syllogism complex predicates scheme*. For each syllogistic scheme, a formal argument scheme (consisted of premises and conclusion (bold)) was provided.

- mistralai/Mistral-7B-v0.1
- mistralai/Mistral-7B-Instruct-v0.2
- mistralai/Mixtral-8x7B-Instruct-v0.1
- google/gemma-7b
- google/gemma-7b-it
- meta-llama/Meta-Llama-3-8B
- meta-llama/Meta-Llama-3-8B-Instruct
- BioMistral/BioMistral-7B

The pretrained LLM weights are used through the *transformers*⁵ python library. All models were loaded with standard configurations and their respective default tokenizers, using the *AutoModelForCausalLM* and *AutoTokenizer* classes. Ad-

⁵<https://huggingface.co/docs/transformers/>

ditionally, the models were loaded with the options *device_map*="auto", *torch_dtype*="auto" *attn_implementation*="flash_attention_2" and *offload_buffers*=True to make the best use of GPU resources available.

For the instruction models, the inputs were passed through each model's chat template (with *tokenizer.apply_chat_template*), so that they would follow the appropriate prompt format. Responses were cleaned of special symbols for evaluation. Table 3 contains the relevant characteristics of all analyzed models.

D Experimental Details

The entire experimental setup was implemented as a *python* code package, consisting of 5 modules:

- *pathways*: defines the ontological relations and operations for biological pathways, as

Model	Type	Parameters	Training Data	Architecture Enhancements	Open-Source	Performance in Benchmarks	Efficiency	Use Cases	Year
Mistral-7B-v0.1	Base LLM	7B	Diverse (internet, academic, etc.)	Grouped-query attention, sliding-window	Yes	Strong in reasoning, real-life scenarios	Similar, optimizations possible	Math, code generation	2023
Mistral-7B-Instruct-v0.2	Instruct LLM	7B	Instruction-tuned data	Instruction-tuned architecture	Yes	High on MT-Bench	Efficient for instruction-based tasks	Customer service chatbots and virtual assistants	2023
Mixtral-8x7B-Instruct-v0.1	Instruct LLM	8x7B	Diverse data with expert focus	Mixture of Experts architecture	Yes	Exceptional performance in Arena Elo and MMLU, excellent MT-Bench score	Great at simulated dialogues and general language understanding	Customer service chatbots and interactive storytelling	2023
Gemma-7B	Base LLM	7B	6 trillion tokens (web, math, code)	Multi-query attention, RoPE, GeGLU	Yes (with terms of use)	Strong in code generation, math	Lightweight, runs on various platforms	Text generation, translation	2023
Gemma-7B-it	Instruct LLM	7B	6 trillion tokens (web, math, code)	Instruction-tuned architecture	Yes	Strong in reasoning, in MMLU and HellaSwag	Efficient for instruction-based tasks	Text comprehension and generation inference	2023
Meta-Llama-3-8B	Base LLM	8B	Mix of publicly available online data	Standard transformer	Yes	Moderate	Efficient for general tasks	General NLP tasks	2023
Meta-Llama-3-8B-Instruct	Instruct LLM	8B	Instruction-tuned data	Instruction-tuned architecture	Yes	High for Arena Elo, impressive MT Bench scores for translation, exceptional MMLU score, indicating strong reasoning and knowledge	Efficient for instruction-based tasks	High-volume applications, great for real-time interactions	2023
BioMistral-7B	Domain-Specific LLM	7B	Collection of medical LLMs resulting from further pre-training Mistral-7B Instruct on PubMed Central resources	Enhanced for domain-specific tasks	Yes	State-of-the-art performance on the multilingual medical evaluation benchmark compared to other open-source 7B models	Efficient for domain-specific tasks	Biomedical research	2024

Table 3: Detailed Characteristics of Selected LLMs.

well as the logic to retrieve and transform data from the domain ontology.

- *logic2nl*: translates logic formulas into NL statements (premises, conclusions), using parameterised scheme templates.
- *llm*: provides access to LLMs, and facilitates logging.
- *experiments*: defines all test logic, parameterisation and metrics for the experiments.
- *main*: orchestrates all the tests and aggregates results.

The LLMs were loaded and run using the HuggingFace transformers library (v4.43.3). The prompts were processed directly by the models with *generate*, without sampling.

Parameters were set as follows:

- *Number of premises* = 2: the number of valid premises per instance (factual argumentative text).
- *Max distractors* = 5: maximum number of distractors to be added per instance.
- *Subset size* = 200: maximum number of positive and negative instances for each scheme.
- *Batch size* = 20: number of instances evaluated simultaneously.
- *ICL*: Whether the in-context learning prompt would be used or not.
- *Model*: the LLM to be evaluated.

For both Task 1 and Task 2, each model was analyzed across all 28 syllogistic schemes, with responses evaluated from a total of 11,200 instances (28 schemes \times (200 positive + 200 negative)). The

exceptions were: *Generalized Modus Ponens* - complex predicates (18 instances), *Hypothetical Syllogism 1* - base (202 instances) and *Generalized Contraposition* - negation (202 instances). This is due to the number of factually possible different instances being limited by the ground-truth data in such cases. For both tasks, this amounted to a total of 211,840 instances, with performance comparisons made between ZS and ICL settings, leading to a total of 423,680 prompts.

When examining the impact of distractors, responses were analyzed across five variants of each argument scheme, reflecting different numbers of irrelevant premises (from 1 to 5 distractors). For each scheme and model, this resulted in 1,000 responses (200 prompts per scheme \times 5 variants), with the exception of the aforementioned cases. Across all schemes, this totaled 26,480 responses for one model, and 211,840 responses were analyzed for all models for each number of distractors, totaling 1,059,200 responses. Again, comparisons were made between ZS and ICL settings, doubling the previous number.

For both Task 1 and Task 2, performance was compared using actual gene entity names versus synthetic names across two selected schemes. Each scheme was tested with 400 instances (200 positive, 200 negative), resulting in a comprehensive analysis of 12,800 instances per task (200 instances \times 2 schemes \times 2 entity types \times 8 models), providing a thorough comparison of model performance with both factual and synthetic gene names.

The experiments were run on a computer with an AMD EPYC 7413 24-Core CPU, 128GB of available RAM and 2 \times NVIDIA A100-SXM4-80GB GPUs.

E Evaluation Metrics

Accuracy, F1-score, Recall, and Precision are used to evaluate performance in Task 1, a binary

classification task where the conclusion C is labeled as either True or False. These metrics compare the predicted labels \hat{y}_i against the annotated gold labels y_i (True/False).

Reasoning Accuracy RA evaluates a model's ability in Task 2 to correctly predict the entailment label \hat{y}_i and select the appropriate subset of premises \hat{P}_i that justify this prediction. The metric is calculated as the percentage of responses where both the predicted label \hat{y}_i matches the ground truth y_i and the selected premises \hat{P}_i match the ground truth premises P_i .

$$RA = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(\hat{y}_i = y_i \text{ and } \hat{P}_i = P_i)$$

In addition we used two supplementary metrics: **Non-empty output** quantifies the frequency with which the model generates any textual response. This metric is crucial, as an empty output, while not providing an incorrect answer, indicates a failure to engage with the prompt. **Irrelevant text** assesses the extent to which the model's output deviates from the expected format, such as failing to adhere to the directive to respond with a straightforward "True" or "False", or a structured statement like "The correct answer is True." This metric identifies instances where the model's response is extraneous or non-compliant with the specified guidelines.

Faithfulness F evaluates a model's ability to consistently adjust its predictions in response to meaningful changes in input. In this analysis, faithfulness is measured by examining N pairs of prompts that differ only in the 'C' conclusion: one concludes with "it is True that..." (y_i) and the other with "it is False that..." (x_i). A pair is deemed faithful if the model changes its prediction ($f(x_i), f(y_i)$) accordingly—shifting from True to False or from False to True between the two prompts. If the model's response to either prompt in a pair is empty or contains irrelevant text that does not follow the instructions, the entire pair is considered unfaithful ($g = 0$). This criterion ensures that only responses that are meaningful and adhere to the given instructions are evaluated.

$$F = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(f(x_i) \neq f(y_i)) \cdot g(f(x_i), f(y_i))$$

F Prompting LLMs - Zero-shot prompts

F.1 TASK 1

prompt = *Suppose you are a specialist with existing knowledge about a signaling and metabolic molecules and their relations organized into biological pathways and processes. Given premises marked with the letter P and the following number and the conclusion marked with the letter C, determine whether the conclusion logically follows from these premises. If the conclusion logically follows from the premises, you need to return 'True'. If the conclusion does not follow logically from the premises, you need to return 'False'. The output should be a single word <True> or <False>.*

"P1: " + Premise 1

"P2: " + Premise 2

"C:" + Conclusion

F.2 TASK 2

prompt = *Suppose you are a specialist with existing knowledge about a signaling and metabolic molecules and their relations organized into biological pathways and processes. Given premises marked with the letter P and the following number and the conclusion marked with the letter C, determine whether the conclusion logically follows from these premises. If the conclusion logically follows from the premises, you need to return 'True'. If the conclusion does not follow logically from the premises, you need to return 'False'. Specify the premises you used to determine whether the conclusion logically follows from the premises, and only these premises. The output should be a single word <True> or <False> and the numbers of the selected premises after the decimal point, like <True, P1, P2>.*

"P1: " + Premise 1

"P2: " + Premise 2

"C:" + Conclusion

G Prompting LLMs - Few-shot prompts

G.1 TASK 1

Context: *Suppose you are a specialist with existing knowledge about a signaling and metabolic molecules and their relations organized into biological pathways and processes.*

Instructions: *Given premises marked with the letter P and the following number and the conclusion marked with the letter C, determine whether the conclusion logically follows from these premises.*

Relevance: *If the conclusion logically follows from the premises, you need to return 'True'. If the conclusion does not follow logically from the premises, you need to return 'False'.*

Constraint: *The output should be a single word <True> or <False>.*

Demonstration:

"P1: " "Every member of Diseases of hemostasis pathway is a member of Disease pathway"

"P2: " "Gene GP1BB is a member of Diseases of hemostasis pathway"

"C:" "Gene GP1BB is a member of Disease pathway"

The correct answer is: True

"P1: " "Every member of Infectious disease pathway is a member of Disease pathway"

"P2: " "Gene PKQQ is a member of Infectious disease pathway"

"C:" "Gene PKQQ is a member of Infectious disease pathway"

The correct answer is: True

"P1: " "Every member of SLC transporter disorders pathway is a member of Disorders of transmembrane transporters pathway"

"P2: " "Gene AXZY is a member of SLC transporter disorders pathway"

"C:" "Gene AXZY is not a member of Disorders of transmembrane transporters pathway"

The correct answer is: False

"P1: " "Every member of HIV Life Cycle pathway is a member of HIV Infection pathway"

"P2: " "Gene MLLX is a member of HIV Life Cycle pathway"

"C:" "Gene MLLW is a member of HIV Infection pathway"

The correct answer is: False

"P1: " "Every member of ABC transporter disorders pathway is a member of Disorders of transmembrane transporters pathway"

"P2: " "Gene PSMC5 is a member of ABC transporter disorders pathway"

"C:" "It is true that Gene PSMC5 is a member

of Disorders of transmembrane transporters pathway"

G.2 TASK 2

Context: *Suppose you are a specialist with existing knowledge about a signaling and metabolic molecules and their relations organized into biological pathways and processes.*

Instructions: *Given premises marked with the letter P and the following number and the conclusion marked with the letter C, determine whether the conclusion logically follows from these premises.*

Relevance: *If the conclusion logically follows from the premises, you need to return 'True'. If the conclusion does not follow logically from the premises, you need to return 'False'. Specify the premises you used to determine whether the conclusion logically follows from the premises, and only these premises.*

Constraint: *The output should be a single word <True> or <False> and the numbers of the selected premises after the decimal point, like <True, P1, P2>.*

Demonstration:

"P1: " "Every member of Diseases of hemostasis pathway is a member of Disease pathway"

"P2: " "Every member of NS1 Mediated Effects on Host Pathways pathway is a member of Influenza Infection pathway"

"P3: " "Gene AABC is a member of Diseases of hemostasis pathway"

"C:" "Gene AABC is a member of Disease pathway"

The correct answer is: True, P1, P3

"P1: " "Every member of SARS-CoV Infections pathway is a member of Viral Infection Pathways pathway"

"P2: " "Every member of Infectious disease pathway is a member of Disease pathway"

"P3: " "Gene PKQQ is a member of Infectious disease pathway"

"C:" "Gene PKQQ is a member of Infectious disease pathway"

The correct answer is: True, P2, P3

"P1: " "Every member of SLC transporter disorders pathway is a member of Disorders of transmembrane transporters pathway"

"P2: " "Gene AXZY is a member of SLC transporter disorders pathway"

"C:" "Gene AXZY is not a member of Disorders of transmembrane transporters pathway"

The correct answer is: False

"P1: " "Every member of HIV Life Cycle pathway is a member of HIV Infection pathway"

"P2: " "Gene MLLX is a member of HIV Life Cycle pathway"

"C:" "Gene MLLW is a member of HIV Infection pathway"

The correct answer is: False

"P1: " "Every member of ABC transporter disorders pathway is a member of Disorders of transmembrane transporters pathway"

"P2: " "Gene PSMC5 is a member of ABC transporter disorders pathway"

"C:" "It is true that Gene PSMC5 is a member of Disorders of transmembrane transporters pathway"

H Results: Misaligned Instruction-Response

We observed four types of text outputs: those aligned with the instruction (regardless of correctness), empty outputs where no text was generated, incorrect text outputs such as repeated prompts or random content, and outputs resembling Chain-of-Thought (CoT) reasoning that, while potentially containing correct reasoning, did not align with the given instructions (Figs. 8-11). We noticed that BioMistral-7B generated empty outputs in 100% of the cases regardless of the specific settings, while Meta-Llama-3-8B exhibits this behaviour for ZS settings in both tasks. We attribute this observation to safety mechanisms applied during pre-training (Labrak et al., 2024), suggesting that domain-specific knowledge about human genome pathways is absent in both models. Similarly, Mistral-7B-v0.1 responses simply repeat the prompt text in 88% of the cases in the ZS settings, and 69% of the cases in FS (Table 2). Moreover, CoT outputs including phrases like e.g. "A nice logical puzzle! Let's break it down step by step..." were particularly common for Meta-Llama-3-8B Instruct, which often ignored the specific instructions to address the task. This behaviour highlights potential biases introduced during instruction-tuning which make the models unable to generalise to domains that are out-of-distribution of the training set.

I Results: Ambiguous Impact of Distractors on Reasoning

LLMs show a slight sensitivity to increasing number of distractors n_D in the prompt, with overall accuracy remaining stable (Figs. 12- 15, Table 4).

While some models struggle with increasing n_D , others can leverage few-shot learning to mitigate their impact, though the effect is scheme-dependent. Considering Task 1, in the ZS setting (Fig. 12, Table 6), Gemma-7b shows a significant decline in accuracy as n_D increases, particularly in the *generalized dilemma* ($r = -0.643, p = 0.001$), *generalized modus ponens* ($r = -0.592, p = 0.002$), and *generalized modus tollens* ($r = -0.571, p = 0.004$) schemes, indicating a moderate negative correlation. In contrast, in the ZS setting, Mistral-7B Instruct-v0.2 exhibits a moderate improvement in accuracy with higher n_D , in the *generalized modus tollens* ($r = 0.540, p = 0.006$) scheme, reflecting a weak positive correlation overall ($r = 0.333, p < 0.001$). Considering *reasoning accuracy*, in the ZS setting (Fig. 14, Table 8), the Gemma-7b model exhibited a substantial drop as the n_D increased ($r = -0.951, p < 0.001$), with an initial low accuracy of 0.3 even with $n_D = 0$. The steepest declines were observed in *hypothetical syllogism 1* ($r = -1.0, p < 0.000$) and *generalized dilemma* ($r = -1.0, p < 0.000$) schemes. For the Gemma-7b-it, the strongest negative correlation between the model and n_D was for the *hypothetical syllogism 1*, *generalized modus ponens* and *generalized contraposition* schemes ($r = -1.0, p = 0.000$). In the FS setting (Fig. 15), the Gemma-7b-it model consistently exhibited significant decreases in *reasoning accuracy* across all schemes, with the most pronounced effect in *hypothetical syllogism 3* and *generalized modus tollens* ($r = -1.0, p < 0.000$). Interestingly, the Mistral-7B Instruct model in both settings, depending on the scheme, showed a positive or negative significant correlation.

The findings underscore the substantial impact of distractors on reasoning accuracy, particularly in complex syllogistic reasoning tasks, revealing that current LLMs are highly susceptible to performance degradation as distractor complexity increases.

J Results: Models Prioritize Contextual Knowledge Over Background Knowledge

The lack of statistically significant differences (Fig. 7) in accuracy between biologically factual and artificial datasets across *generalized modus ponens* and *generalized modus tollens* schemes suggests that the models' reasoning capabilities rely more on stated contextual knowledge and logical structure than on pre-existing background knowledge. This holds true for both *accuracy* and *reasoning accuracy*, as well as in both ZS and FS settings: models that perform well on a given scheme maintain their performance even when factual gene names are replaced by synthetic names, and the same consistency is observed for models with weaker performance. This ability to maintain accuracy with synthetic gene names in the artificial set demonstrates that models can abstract and apply logical reasoning independently of their internal domain-specific knowledge.

K Supplementary Figures

L Supplementary Tables

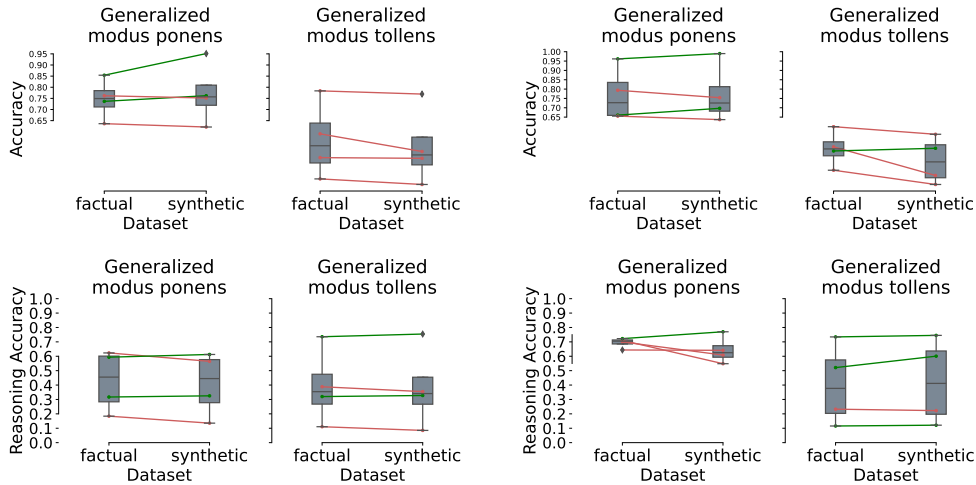


Figure 7: Accuracy comparison between two datasets: Biologically Factual vs. Artificial with Synthetic Names for Task 1 (top) and Task 2 (bottom); ZS (left) and FS (right). Lines connect the accuracy for each model, with green indicating an increase and red indicating a decrease. Gray boxplots display the median, Q1, Q3, and the range (minimum to maximum) of the data.

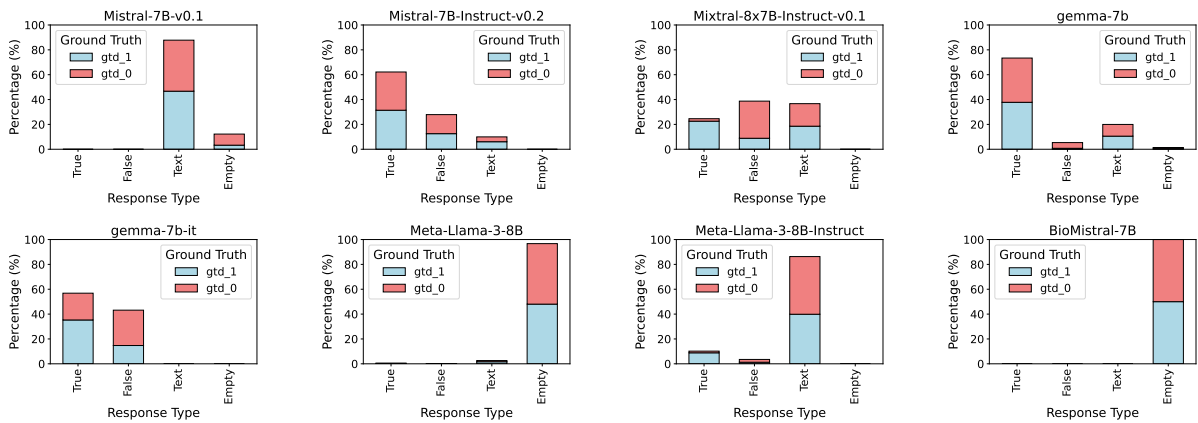


Figure 8: Percentage distribution of model response types under zero-shot settings for prompts with no distractors for the set of biologically factual argumentative texts.

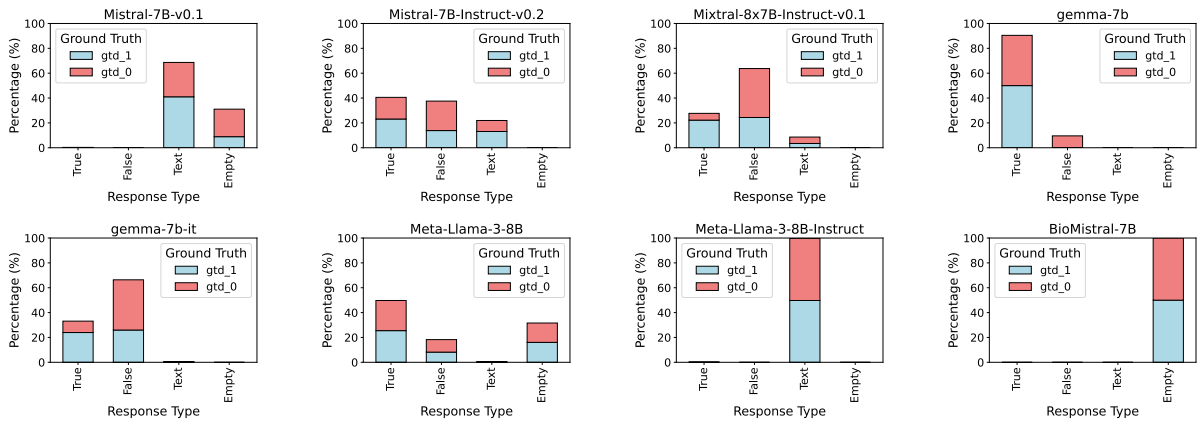


Figure 9: Percentage distribution of model response types under few-shot settings for prompts with no distractors for the set of biologically factual argumentative texts.

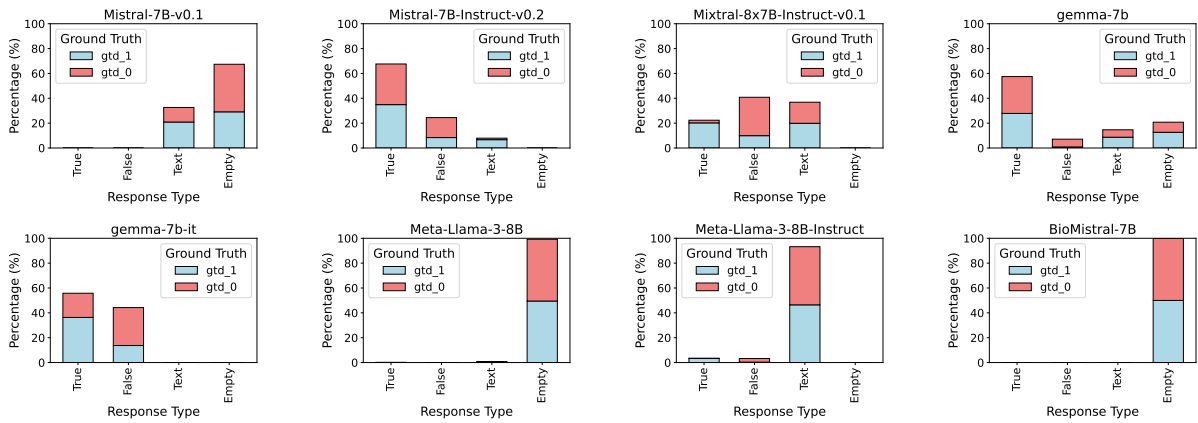


Figure 10: Percentage distribution of model response types under zero-shot settings for prompts with all distractors for the set of biologically factual argumentative texts.

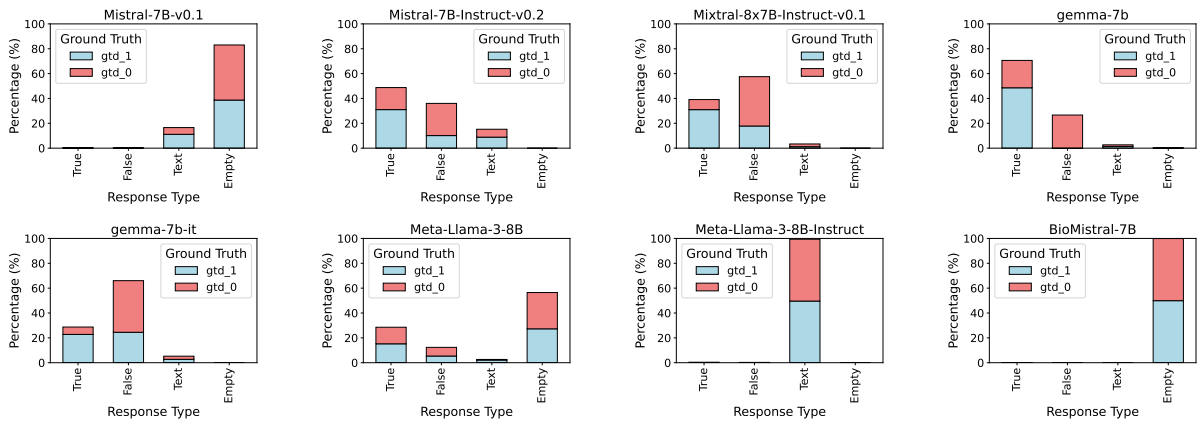


Figure 11: Percentage distribution of model response types under few-shot settings for prompts with all distractors for the set of biologically factual argumentative texts.

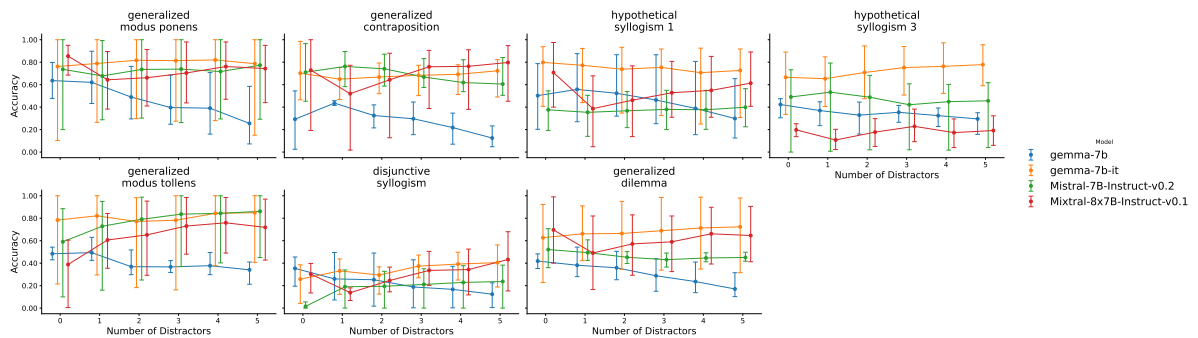


Figure 12: Task 1: Accuracy versus number of distractors and scheme in the zero-shot setting. Lines connect the average values for each model, with error bars representing the range (min-max).

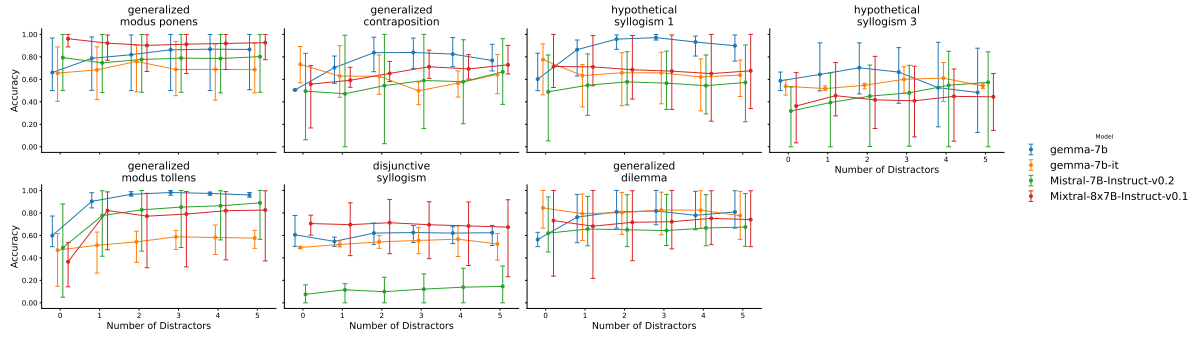


Figure 13: Task 1: Accuracy versus number of distractors and scheme in the few-shot setting. Lines connect the average values for each model, with error bars representing the range (min-max).

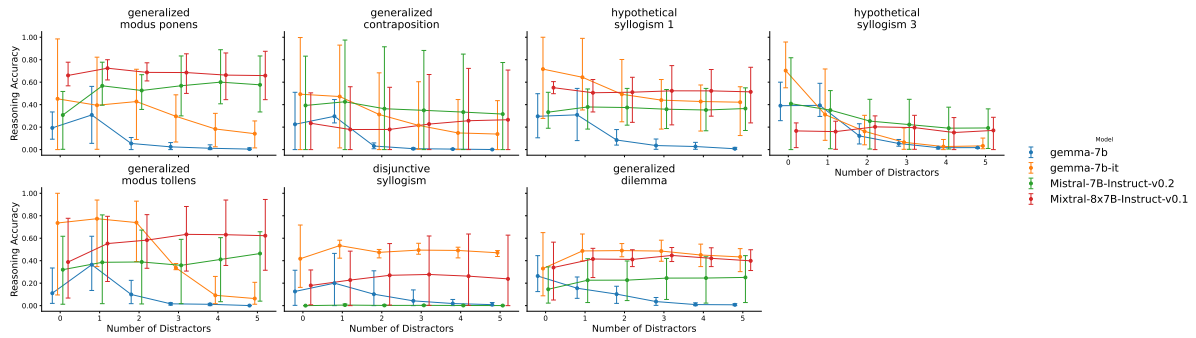


Figure 14: Task 2: Reasoning accuracy versus number of distractors and scheme in the zero-shot setting. Lines connect the average values for each model, with error bars representing the range (min-max).

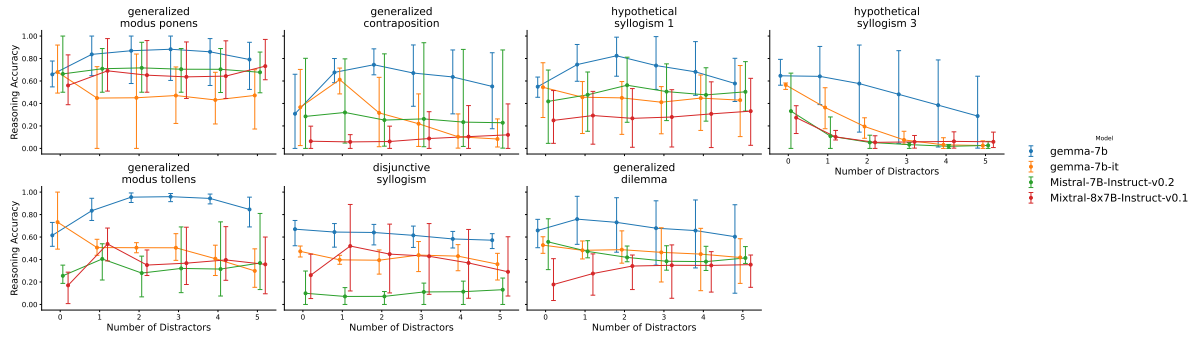


Figure 15: Task 2: Reasoning accuracy versus number of distractors and scheme in the few-shot setting. Lines connect the average values for each model, with error bars representing the range (min-max).

Model				ZS					FS							
	Non-empty output	Irrelevant text	Following instruction	Acc.	Recall	Precision	F1	Faith.	Non-empty output	Irrelevant text	Following instruction	Acc.	Recall	Precision	F1	Faith.
BioMistral-7B	0.00	-	-	-	-	-	-	-	0.00	-	-	-	-	-	-	-
Meta-Llama-3-8B	0.01	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.44	0.06	0.38	0.19	0.25	0.23	0.24	0.08
Meta-Llama-3-8B-Instruct	1.00	0.93	0.07	0.06	0.07	0.06	0.06	0.06	1.00	0.99	0.01	0.00	0.00	0.00	0.00	0.00
Mistral-7B-Instruct-v0.2	1.00	0.08	0.92	0.51	0.70	0.51	0.59	0.43	1.00	0.18	0.82	0.54	0.60	0.53	0.57	0.53
Mistral-7B-v0.1	0.33	0.33	0.00	0.00	0.00	0.00	0.00	0.00	0.17	0.16	0.01	0.00	0.00	0.00	0.00	0.00
Mixtral-8x7B-Instruct-v0.1	1.00	0.37	0.63	0.51	0.40	0.51	0.45	0.33	1.00	0.06	0.94	0.68	0.62	0.70	0.66	0.56
Gemma-7b	0.79	0.15	0.64	0.34	0.56	0.39	0.46	0.09	0.99	0.02	0.97	0.75	0.97	0.68	0.80	0.53
Gemma-7b-it	1.00	0.00	1.00	0.67	0.73	0.65	0.69	0.63	1.00	0.08	0.92	0.62	0.43	0.70	0.53	0.39

Table 4: Task 1: Distribution of model response types and performance outcomes across two experimental settings—ZS and FS, considering all distractor conditions (n distractors from 0 to 5), for all syllogistic schemes within biologically factual argumentative texts. Response types include: non-empty outputs, irrelevant text generation, and outputs adhering to the given instructions.

Scheme	Model	Recall		Precision		F1-score		Accuracy	
		ZS	FS	ZS	FS	ZS	FS	ZS	FS
disjunctive syllogism	All models	0.11	0.24	0.11	0.25	0.11	0.25	0.12	0.27
	BioMistral-7B	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Meta-Llama-3-8B	0.01	0.43	0.01	0.33	0.01	0.38	0.01	0.29
	Meta-Llama-3-8B-Instruct	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Mistral-7B-Instruct-v0.2	0.03	0.03	0.03	0.03	0.03	0.03	0.02	0.08
	Mistral-7B-v0.1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Mixtral-8x7B-Instruct-v0.1	0.11	0.42	0.18	0.98	0.13	0.59	0.30	0.71
	Gemma-7b	0.65	1.00	0.41	0.56	0.50	0.72	0.35	0.61
Gemma-7b-it	0.08	0.02	0.13	0.34	0.10	0.03	0.26	0.49	
generalized contraposition	All models	0.39	0.38	0.33	0.34	0.35	0.36	0.30	0.31
	BioMistral-7B	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Meta-Llama-3-8B	0.01	0.46	0.01	0.35	0.01	0.39	0.01	0.29
	Meta-Llama-3-8B-Instruct	0.10	0.00	0.09	0.00	0.10	0.00	0.06	0.00
	Mistral-7B-Instruct-v0.2	0.97	0.39	0.61	0.44	0.75	0.41	0.67	0.45
	Mistral-7B-v0.1	0.00	0.1	0.00	0.01	0.00	0.01	0.00	0.00
	Mixtral-8x7B-Instruct-v0.1	0.61	0.45	0.72	0.55	0.67	0.49	0.69	0.54
	Gemma-7b	0.42	1.00	0.32	0.50	0.36	0.67	0.26	0.51
Gemma-7b-it	0.98	0.74	0.61	0.71	0.75	0.73	0.68	0.72	
generalized dilemma	All models	0.37	0.48	0.32	0.41	0.34	0.45	0.28	0.40
	BioMistral-7B	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Meta-Llama-3-8B	0.01	0.61	0.01	0.45	0.01	0.52	0.00	0.43
	Meta-Llama-3-8B-Instruct	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Mistral-7B-Instruct-v0.2	0.91	0.86	0.51	0.58	0.65	0.69	0.52	0.62
	Mistral-7B-v0.1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Mixtral-8x7B-Instruct-v0.1	0.48	0.63	0.85	0.79	0.61	0.70	0.70	0.73
	Gemma-7b	0.79	1.00	0.45	0.53	0.58	0.70	0.42	0.56
Gemma-7b-it	0.82	0.79	0.59	0.89	0.69	0.84	0.63	0.85	
generalized modus ponens	All models	0.51	0.50	0.43	0.45	0.47	0.47	0.41	0.44
	BioMistral-7B	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Meta-Llama-3-8B	0.00	0.59	0.00	0.44	0.00	0.50	0.00	0.41
	Meta-Llama-3-8B-Instruct	0.75	0.03	0.48	0.03	0.59	0.03	0.47	0.02
	Mistral-7B-Instruct-v0.2	0.80	0.99	0.65	0.73	0.72	0.84	0.69	0.82
	Mistral-7B-v0.1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Mixtral-8x7B-Instruct-v0.1	0.91	1.00	0.80	0.97	0.85	0.98	0.84	0.98
	Gemma-7b	0.91	1.00	0.56	0.63	0.69	0.78	0.59	0.71
Gemma-7b-it	0.74	0.37	0.69	0.64	0.71	0.47	0.70	0.58	
generalized modus tollens	All models	0.40	0.28	0.34	0.28	0.36	0.28	0.31	0.29
	BioMistral-7B	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Meta-Llama-3-8B	0.01	0.48	0.01	0.40	0.01	0.43	0.01	0.37
	Meta-Llama-3-8B-Instruct	0.32	0.00	0.26	0.00	0.28	0.00	0.20	0.00
	Mistral-7B-Instruct-v0.2	0.70	0.42	0.57	0.49	0.63	0.45	0.59	0.49
	Mistral-7B-v0.1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Mixtral-8x7B-Instruct-v0.1	0.50	0.14	0.41	0.26	0.45	0.18	0.39	0.37
	Gemma-7b	0.92	1.00	0.49	0.55	0.64	0.71	0.48	0.60
Gemma-7b-it	0.71	0.17	0.83	0.42	0.77	0.24	0.78	0.47	
hypothetical syllogism 1	All models	0.36	0.37	0.32	0.36	0.34	0.37	0.30	0.36
	BioMistral-7B	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Meta-Llama-3-8B	0.01	0.54	0.01	0.42	0.01	0.47	0.01	0.40
	Meta-Llama-3-8B-Instruct	0.16	0.00	0.14	0.00	0.15	0.00	0.10	0.00
	Mistral-7B-Instruct-v0.2	0.59	0.42	0.40	0.44	0.48	0.43	0.35	0.44
	Mistral-7B-v0.1	0.00	0.01	0.00	0.01	0.00	0.01	0.00	0.01
	Mixtral-8x7B-Instruct-v0.1	0.55	0.36	0.72	0.97	0.63	0.53	0.67	0.67
	Gemma-7b	0.74	1.00	0.48	0.56	0.58	0.72	0.46	0.61
Gemma-7b-it	0.80	0.62	0.77	0.85	0.78	0.72	0.78	0.76	
hypothetical syllogism 3	All models	0.30	0.29	0.26	0.28	0.28	0.28	0.23	0.27
	BioMistral-7B	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Meta-Llama-3-8B	0.00	0.45	0.00	0.34	0.00	0.39	0.00	0.29
	Meta-Llama-3-8B-Instruct	0.04	0.00	0.03	0.00	0.03	0.00	0.03	0.00
	Mistral-7B-Instruct-v0.2	0.48	0.13	0.49	0.21	0.49	0.16	0.49	0.32
	Mistral-7B-v0.1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Mixtral-8x7B-Instruct-v0.1	0.15	0.23	0.16	0.31	0.15	0.27	0.20	0.36
	Gemma-7b	0.84	1.00	0.46	0.55	0.59	0.71	0.42	0.59
Gemma-7b-it	0.87	0.54	0.62	0.54	0.72	0.54	0.67	0.54	

Table 5: Results from Task 1 baseline models on the biologically factual argumentative texts set (without synthetic data and without distractors, bold - the best and worst accuracy values for each model).

Scheme	Gemma-7b		Gemma-7b-it		Mistral-7B Instruct-v0.2		Mixtral-8x7B Instruct-v0.1	
	r	p-val	r	p-val	r	p-val	r	p-val
ZS								
disjunctive syllogism	-0.497	0.013	0.432	0.035	0.405	0.050	0.418	0.042
generalized contraposition	-0.495	0.014	0.155	0.469	-0.397	0.055	0.229	0.282
dilemma	-0.643	0.001	0.203	0.342	-0.282	0.182	0.049	0.819
generalized modus ponens	-0.592	0.002	0.104	0.630	0.069	0.748	-0.026	0.902
generalized modus tollens	-0.571	0.004	0.198	0.353	0.540	0.006	0.432	0.035
hypothetical syllogism 1	-0.375	0.071	-0.268	0.205	0.115	0.594	0.012	0.954
hypothetical syllogism 3	-0.465	0.022	0.211	0.321	-0.125	0.560	0.100	0.640
FS								
disjunctive syllogism	0.279	0.187	0.370	0.075	0.197	0.356	0.116	0.588
generalized contraposition	0.511	0.011	-0.201	0.347	0.189	0.377	0.349	0.095
generalized dilemma	0.462	0.023	-0.097	0.652	0.183	0.391	-0.046	0.831
generalized modus ponens	0.455	0.025	0.051	0.812	0.011	0.961	0.023	0.914
generalized modus tollens	0.568	0.004	0.224	0.292	0.419	0.041	0.391	0.059
hypothetical syllogism 1	0.420	0.041	-0.314	0.135	0.056	0.793	-0.026	0.902
hypothetical syllogism 3	-0.122	0.571	0.199	0.351	0.422	0.040	0.116	0.588

Table 6: Values of the Spearman’s Ranked Correlation Coefficient (r) for Accuracy by Distractors and Syllogistic Scheme for evaluated Models in Task 1: Spearman’s correlation coefficients (r) and p-values for the accuracy metric are shown across various levels of distractors and syllogistic schemes for each model. Negative r values reflect a decrease in accuracy with increasing distractor complexity. The highest correlation values for each scheme are highlighted in bold, indicating the models most affected by distractors.

Scheme	BioMistral-7B		Meta-Llama-3-8B		Meta-Llama-3-8B Instruct		Mistral-7B Instruct-v0.2		Mistral-7B-v0.1		Mistral-8x7B Instruct-v0.1		Gemma-7b		Gemma-7b-it		All models	
	ZS	FS	ZS	FS	ZS	FS	ZS	FS	ZS	FS	ZS	FS	ZS	FS	ZS	FS	ZS	FS
disjunctive syllogism	0.00	0.00	0.00	0.32	0.15	0.00	0.00	0.10	0.00	0.00	0.18	0.26	0.13	0.67	0.42	0.47	0.11	0.23
generalized contraposition	0.00	0.00	0.00	0.16	0.15	0.00	0.33	0.21	0.00	0.00	0.20	0.04	0.18	0.27	0.43	0.32	0.16	0.13
generalized dilemma	0.00	0.00	0.00	0.37	0.49	0.00	0.15	0.56	0.00	0.00	0.34	0.18	0.26	0.66	0.33	0.53	0.20	0.29
generalized modus ponens	0.00	0.00	0.00	0.48	0.28	0.05	0.32	0.56	0.00	0.00	0.62	0.61	0.18	0.62	0.59	0.74	0.25	0.38
generalized modus tollens	0.00	0.00	0.00	0.44	0.48	0.01	0.32	0.26	0.00	0.00	0.39	0.17	0.11	0.62	0.74	0.73	0.25	0.28
hypothetical syllogism 1	0.00	0.00	0.00	0.40	0.45	0.00	0.31	0.39	0.00	0.00	0.55	0.21	0.31	0.54	0.68	0.51	0.29	0.26
hypothetical syllogism 3	0.00	0.00	0.00	0.29	0.13	0.00	0.41	0.33	0.00	0.00	0.17	0.27	0.39	0.65	0.70	0.57	0.23	0.26

Table 7: Reasoning accuracy results from Task 2 baseline models on the set of biologically factual argumentative texts (without synthetic data and without distractors).

Scheme	Gemma-7b		Gemma-7b-it		Mistral-7B Instruct-v0.2		Mixtral-8x7B Instruct-v0.1	
	r	p-val	r	p-val	r	p-val	r	p-val
ZS								
disjunctive syllogism	-0.943	0.005	0.029	0.957	-0.086	0.872	0.486	0.329
generalized contraposition	-0.943	0.005	-1.000	0.000	-0.829	0.042	0.429	0.397
generalized dilemma	-1.000	0.000	-0.086	0.872	1.000	0.000	0.314	0.544
generalized modus ponens	-0.943	0.005	-1.000	0.000	0.543	0.266	0.771	0.072
generalized modus tollens	-0.943	0.005	-0.829	0.042	0.829	0.042	0.771	0.072
hypothetical syllogism 1	-1.000	0.000	-1.000	0.000	0.029	0.957	-0.486	0.329
hypothetical syllogism 3	-0.886	0.019	-0.943	0.005	-0.943	0.005	0.029	0.957
FS								
disjunctive syllogism	-1.000	0.000	-0.543	0.266	0.829	0.042	-0.143	0.787
generalized contraposition	-0.086	0.872	-0.943	0.005	-0.943	0.005	0.829	0.042
generalized dilemma	-0.657	0.156	-0.943	0.005	-0.829	0.042	0.943	0.005
generalized modus ponens	0.371	0.468	-0.886	0.019	0.600	0.208	0.829	0.042
generalized modus tollens	0.486	0.329	-1.000	0.000	0.371	0.468	0.257	0.623
hypothetical syllogism 1	-0.086	0.872	-0.943	0.005	0.314	0.544	0.829	0.042
hypothetical syllogism 3	-1.000	0.000	-1.000	0.000	-0.943	0.005	-0.486	0.329

Table 8: Values of the Spearman’s Ranked Correlation Coefficient (r) for Accuracy by Distractors and Syllogistic Scheme for evaluated Models in Task 2: Spearman’s correlation coefficients (r) and p-values for the reasoning accuracy metric are shown across various levels of distractors and syllogistic schemes for each model. Negative r values reflect a decrease in accuracy with increasing distractor complexity. The highest correlation values for each scheme are highlighted in bold, indicating the models most affected by distractors.