

Evaluating Large Language Models with NeuBAROCO: Syllogistic Reasoning Ability and Human-like Biases

Risako Ando
risakochoan@keio.jp

Takanobu Morishita
morishita@keio.jp

Hirohiko Abe
hirohiko-abe@keio.jp

Koji Mineshima
minesima@abelard.flet.keio.ac.jp

Mitsuhiro Okada
mitsu@abelard.flet.keio.ac.jp

Keio University

Abstract

This paper investigates whether current large language models exhibit biases in logical reasoning, similar to humans. Specifically, we focus on syllogistic reasoning, a well-studied form of inference in the cognitive science of human deduction. To facilitate our analysis, we introduce a dataset called NeuBAROCO, originally designed for psychological experiments that assess human logical abilities in syllogistic reasoning. The dataset consists of syllogistic inferences in both English and Japanese. We examine three types of biases observed in human syllogistic reasoning: belief biases, conversion errors, and atmosphere effects. Our findings demonstrate that current large language models struggle more with problems involving these three types of biases.

1 Introduction

Syllogistic inferences and their various variants have been extensively studied since *Prior Analytics* by Aristotle in the 4th century BC. While the Aristotelean syllogism is a small part of the predicate logic and a limited inference system when compared to, for example, the formal system of logical inference rules, there has been recently a revival movement of Aristotelean syllogism and its variants, including natural logic (van Benthem, 1986; Sánchez Valencia, 1991; Moss, 2015). This renewed attention arises from the perspective of viewing syllogistic inferences as a “natural” inference rule applicable to our everyday reasoning in ordinary language.

Not only is there a re-evaluation of the significance of syllogistic inferences and their variants in relation to their usefulness in ordinary language, but they are also considered as a benchmark for various inference studies in different disciplines. For example, cognitive psychological studies of logical inferences (Stenning and Van Lambalgen, 2012), diagrammatic logical inference studies (Sato and

Mineshima, 2015), neuroscientific studies of logical inferences (Goel et al., 2000), all draw upon syllogistic reasoning as a point of reference. On the other hand, the recent developments of deep-learning-based AI-tools of natural languages, in particular, the state-of-the-art Large Language Models (LLMs), including BERT (Devlin et al., 2019) and GPT (Brown et al., 2020), are remarkable. These tools hold the potential to be useful for logical inferences. However, there is still a need for further accumulation of studies on the use of AI models for logical inferences in natural language.

In this paper, we explore the potential of LLMs for performing logical inferences, with a specific focus on using syllogistic inferences as a benchmark. We present the NeuBAROCO dataset, a new dataset consisting of syllogistic inferences in both English and Japanese. The dataset is derived from the question collection BAROCO (Shikishima et al., 2009), which has been used in various studies in Japan to evaluate human syllogistic reasoning abilities.

The field of cognitive science of human reasoning has not yet been fully integrated with the recent advancements in AI, despite its potential to provide valuable insights for AI inference research. Considering the significant attention given to biases in studies of logical inference within cognitive psychology and cognitive science (Evans, 1989; Evans et al., 1993), our primary focus lies in assessing whether LLMs exhibit the biases observed in human logical inferences. We focus on three types of human biases that have been studied in cognitive science, namely, belief biases, conversion errors, and atmosphere effects (see Section 4). We explore the extent to which currently available LLMs for natural language inference can effectively address syllogistic inferences that are susceptible to errors resulting from these biases.

The contributions of the paper are summarized as follows.

1. We present the NeuBAROCO dataset, specif-

ically designed for syllogistic inferences, which serves as a valuable resource for examining human-biases in language models.

2. Using this dataset, we evaluate the logical reasoning ability of several recent LLMs both for English and Japanese.
3. Our evaluation reveals that the current LLMs exhibit significant shortcomings when faced with problems that are prone to errors resulting from the three biases of interest.

The dataset will be made publicly available for research on human and machine understanding of logical inferences.

2 Background and Related Work

2.1 Syllogism

A syllogism consists of the following four types of categorical sentences described as **A**, **E**, **I**, and **O**.

Type	Form	Description
A	All S are P	Universal affirmative
E	No S are P	Universal negative
I	Some S are P	Particular affirmative
O	Some S are not P	Particular negative

Table 1: Four types of categorical sentences

Each syllogism is composed of two premises (**P1**, **P2**) and one conclusion (**C**). Thus, each type of syllogisms is identified by the types of three sentences. The following is an example of EIO-type syllogism (called *Ferio* in the traditional mnemonic), which is a valid pattern of syllogism.

P1: No B are C. (**E**-type)
P2: Some A are B. (**I**-type)

C: Some A are not C. (**O**-type)

As outlined in Section 1, syllogisms have been extensively studied in the fields of logic and psychology. To assess the logical reasoning capabilities of current language models in natural language inferences, we use a dataset that encompasses various types of syllogistic inferences.

2.2 Machine learning and logical inference

In recent years, syllogistic inferences and their variants have been used to examine the reasoning ability of machine learning based models for natural language inferences (NLI).

Richardson et al. (2020) examined the capacities of NLI models to perform various types of logical

inferences that involve boolean operators, quantifiers, comparatives, conditionals, negation, and counting expressions. The study utilized synthetically generated data. Their findings indicate that models fine-tuned with NLI datasets perform well, suggesting that NLI models enhance their accuracy when provided with additional datasets as input. However, due to the artificial nature of the data, the majority of the inferences in each type of inferences share a similar structure, making it relatively easy for machine learning algorithms to solve such similar problems. In contrast, our study utilizes human-generated data, with a specific emphasis on inferences that elicit human-like biases. Furthermore, our study focuses on the capacity of LLMs to solve syllogistic inferences without requiring fine-tuning with NLI training examples.

Yanaka et al. (2019) examined monotonicity inferences, those logical inferences that are licensed by substituting general terms in quantified sentences. A testset for monotonicity inference was created semi-automatically and then tested on various language models. One interesting finding is that all the models performed poorly on a class of inferences involving negative contexts (the so-called downward monotone inferences). Monotonicity inferences are structurally simpler than syllogistic inferences in that they consist of single-premises; by contrast, syllogisms are composed of multiple premises. Combination of quantifiers in syllogisms such as *no* and *some* can be more challenging than monotonicity inferences.

Schlegel et al. (2022) conducted an empirical study to investigate the detection of formally valid inference within controlled fragments of natural language. These fragments were specifically designed to increase the complexity of the satisfiability problem. In their study, each fragment consisted of artificially generated sets of English sentences that incorporated determiners such as *every*, *some*, *no*, negation *not*, and relative clauses, within the context of a vocabulary comprising count nouns and transitive verbs. The findings of Schlegel et al. (2022) indicate that transformer-based language models fine-tuned with training data tend to exhibit overfitting to superficial patterns present in the data, rather than acquiring the logical principles that govern reasoning within these fragments. In other words, the models seem to focus on surface-level features rather than grasping the underlying logical principles. Furthermore, according to Schlegel et al.

(2022), the ability of neural networks to learn and solve the various satisfiability problems does not appear to align with the complexity classes associated with the elicited fragments. Our study focuses on a small yet manually controlled dataset, as opposed to a large corpus of artificially generated data. Specifically, we manually annotate problems that are susceptible to human-like biases with corresponding labels. This approach facilitates meaningful comparisons between the syllogistic inference capabilities of humans and models.

Closest to our work is Dasgupta et al. (2022), which reveals that large language models show content effects (i.e., what we called *belief biases*) in syllogism reasoning as well as humans. They introduced new datasets of abstract logical inferences including syllogisms. Each syllogism is annotated with the information about whether or not a proposition is consistent with human beliefs and knowledge. They found that when the conclusion of an inference contradicts reality, the language model exhibits a strong bias towards classifying the argument as invalid, regardless of its logical validity. Our experimental results provide further support of these findings and demonstrate that similar effects are observed not only in English but also in Japanese, a typologically different language from English. Furthermore, we expand our focus beyond belief biases to include various types of biases such as conversion errors and atmosphere effects. We systematically examine the impact of these biases on LLMs with the ultimate objective of comparing the performance of LLMs in logical reasoning to that of humans.

3 The NeuBAROCO Dataset

3.1 Background: the BAROCO dataset

BAROCO is the collection of logical inference questions to examine subjects' ability of logical inference. The questions of BAROCO are mainly composed of syllogistic inferences and their variants. BAROCO has been used in various studies on human logical inference abilities. BAROCO was first used for behavioral genetic studies with the twin method in Shikishima et al. (2006, 2009), where the genetic factor and the environmental factor of the logical inference ability were measured. 500 twin pairs (1,000 participants) were asked to answer 100 questions about a version of BAROCO. The results were then compared with the subjects' scores on a standard IQ test that typically did not in-

clude logical inference abilities. Additionally, correlations were explored between logical inference abilities and decision-making skills in the fields of behavioral economics and cognitive sociology. For instance, Shikishima et al. (2015) investigated the relationship between logical inference abilities and Allais's decision-making task, along with other related studies.

3.2 Data construction

The full version of the original BAROCO dataset comprises a collection of 209 logical inferences divided into seven sections, each containing different types of questions. The version of BAROCO called "BAROCO-ALL" encompasses a total of 200 questions, which includes the following three sections. Examples of each section will be presented in Table 2.

- (1) Abstract syllogism inferences: This section consists of inferences where the terms used in the sentences are represented by capital letters of the alphabet.
- (2) Contentual (belief-consistent) syllogistic inferences: In this section, the inferences are constructed using concrete nouns commonly used in ordinary language.
- (3) Belief-inconsistent syllogistic inferences: This section introduces inferences where belief-inconsistent sentences may appear within the inference itself.

Most questions in BAROCO dataset consist of two premises and three options for the correct answer. In the original setup, the participants were asked to choose one logically valid conclusion from the given options. We transformed each question into a format commonly used for evaluating NLI models, where inferences are categorized as *entailment*, *contradiction*, or *neither* (which we call *neutral*). In accordance with the format of syllogisms, each inference consists of two premises and one conclusion. We manually assigned each inference with the appropriate label of *entailment*, *contradiction*, and *neutral*.

The resulting dataset obtained from this process is referred to as NeuBAROCO. In total, there are 375 inference problems in the NeuBAROCO dataset, with 122 instances labelled as *entailment*, 71 instances labelled as *contradiction*, and 182 instances labelled as *neutral*.

Type	Language	Example
Symbol	English	P1: All A are B. P2: All B are C. C: All A are C.
	Japanese	P1: すべての A は B である P2: すべての B は C である C: すべての A は C である
Consistent	English	P1: One friend of Taro is a friend of Paul. P2: All of Paul’s friends are German. C: One of Taro’s friends is German.
	Japanese	P1: 太郎のある友人はポールの友人である。 P2: ポールのすべての友人はドイツ人である。 C: 太郎のある友人はドイツ人である。
Inconsistent	English	P1: Some animals are human beings. P2: All animals are tomatoes. C: Some humans are tomatoes.
	Japanese	P1: ある動物は人間である。 P2: すべての動物はトマトである。 C: ある人間はトマトである。

Table 2: Examples of symbolic, consistent, and inconsistent syllogism in the NeuBAROCO dataset. The English sentences (**P1**, **P2**, **C**) in each example correspond to the respective Japanese sentences. The correct label for all examples is *entailment*.

The BAROCO dataset was written in Japanese. We translated each problem into English using the DeepL translation tool (<https://www.deepl.com/translator>). We manually checked and adjusted the wording of each sentence, ensuring that they conform to the patterns of categorical sentences. We normalized the quantifiers in the English sentences. We used *all* or *every* for universal quantification in **A**-type sentences, and *some*, *a certain*, or *one of* for existential quantification in **I**-type and **O**-type sentences, and *no* for universal negative in **E**-type sentences. To prevent the sentences from being interpreted as generic statements, we refrained from using the indefinite article *a* (or *an*) for existential quantification. The presence of the indefinite article can lead to a generic interpretation, such as in the sentence *A cat is an animal*. This ensures consistency and clarity in the translation of the original Japanese sentences into English.

3.3 Annotation

We annotated each inference problem in the NeuBAROCO dataset as to what type of inference it is and whether the sentences appearing in it are consistent with beliefs.

3.3.1 Types of logical inferences

There are two types of inferences in the dataset: basic syllogisms and extended syllogisms.

Basic syllogisms As explained in Section 2, basic syllogisms consist of two premises (**P1**, **P2**) and one conclusion (**C**). We annotate each basic syllogism with the types of premises and conclusion. The following is an example of **IAI**-type syllogism:

$$\begin{array}{l} \mathbf{P1:} \text{ Some A are B. (I-type)} \\ \mathbf{P2:} \text{ All B are C. (A-type)} \\ \hline \mathbf{C:} \text{ Some A are C. (I-type)} \end{array}$$

The first premise is **I**-type. The second premise is **A**-type. The conclusion is **I**-type. Therefore, the inference is labeled as **IAI**.

Extended syllogisms Extended syllogisms can be classified into two types. One is a boolean inference where conjunction *and* and *or* appear between terms. The following is an example:

$$\begin{array}{l} \mathbf{P1:} \text{ All A or B are C.} \\ \mathbf{P2:} \text{ No C are D.} \\ \hline \mathbf{C:} \text{ No B are D.} \end{array}$$

The other is a hypothetical syllogism, one of whose premises is a conditional sentence of the form *If P then Q*. Here, *P* or *Q* can be a negated sentence. The following is an example:

P1: If Hanako has blood type O,
then Hanako’s daughter has blood type B.
P2: Hanako’s daughter does not have blood type B.
C: Hanako does not have blood type O.

In the dataset, there are 318 basic syllogisms and 57 extended syllogisms.

3.3.2 Belief consistency

We also classify the inferences into three distinct types based on the types of sentences they contain: *symbolic*, *consistent*, and *inconsistent*. Table 2 shows examples of each type.

Symbolic A symbolic inference is composed of sentences where all the terms are abstract symbols (alphabets). For humans, they can be considered to be neutral with respect to beliefs.

Consistent An inference is labelled as *consistent* if all of the premises and conclusion are consistent with common-sense beliefs. In the case of the example in Table 2, all the sentences, i.e., *One friend of Taro is a friend of Paul*, *All of Paul’s friends are German*, and *One of Taro’s friends is German*, can be interpreted consistent with belief.

Inconsistent An inference is labelled as *inconsistent* if at least one of the premises and conclusion is inconsistent with common-sense beliefs, that is, it goes against what is commonly believed or accepted. In the case of the example in Table 2, the contents of two sentences, *Some humans are not living things* and *None of the animals are human* are contrary to common sense.

There are 95 instances of *symbolic*, 167 instances of *consistent*, and 102 instances of *inconsistent*. For cases where the judgment of belief consistency is unclear, we classify them as *others*. We encountered 11 instances that fell into this category.

4 Human-like Biases

Based on the above classification of the types of syllogistic inferences and sentences, we examine three types of human-like biases that can cause reasoning errors: belief biases, conversion errors, and atmosphere effects. We annotated information to each inference in the dataset to make explicit which inferences are misjudged by these biases.

4.1 Belief biases

Belief bias is one of the most well-known biases causing inference errors and has been applied to

various types of logical inferences including syllogisms and Wason’s selection task (Evans, 1989; Newstead et al., 1992; Evans et al., 1993). It is widely recognized that people tend to have trouble in determining whether an inference is valid when it includes a sentence contrary to common sense. For example, the inference that is labelled as *inconsistent* in Table 2, repeated here, has inconsistent sentences **P2** and **C**:

P1: Some animals are human beings.
P2: All animals are tomatoes.
C: Some humans are tomatoes.

Although the correct label for this problem is *entailment*, the fact that the conclusion **C** is contrary to beliefs may lead some to judge it as *contradiction* instead of *entailment*, regardless of its logical validity. Similarly, in the following example, while **P2** is contrary to our beliefs, the conclusion **C** remains consistent. Hence, one might judge the inference as *entailment* rather than *neutral* due to the belief-consistency of the conclusion.

P1: All canines are animals.
P2: All animals are robots.
C: No canine is a robot.

As mentioned in Section 3.3.2, when either one of the premises or the conclusion is inconsistent with our beliefs, we assigned the *inconsistent* label to the inference. We investigate whether or not NLI models are influenced by this type of belief biases.

Conversion errors are errors in syllogisms caused by the incorrect interpretation of terms that appear in premises. There are at least two types of errors, called *illicit conversion* (Wilkins, 1928; Newstead, 1989; Geurts, 2003):

1. The tendency to interpret *All A are B* as equivalent to *All B are A* (**A**-type)
2. The tendency to interpret *Some A are not B* as equivalent to *Some B are not A* (**O**-type).

Note that *All A are B* and *Some A are not B* mean $A \subseteq B$ and $A \cap \bar{B} \neq \emptyset$, respectively, in the standard predicate logic¹, hence terms *A* and *B* are not convertible. Table 3 shows some examples of syllogistic inference problems where conversion errors cause wrong answer.

¹In traditional syllogisms, **A**-type sentence *All A are B* implies that *A* is not empty. The BAROCO dataset follows this traditional interpretation (the *existential import* of universal expressions) when annotating the gold labels.

Type of syllogisms	Language	Example
AAA	English	P1: All B are A. P2: All B are C. C: All A are C.
	Japanese	P1: すべての B は A である P2: すべての B は C である C: すべての A は C である
AOO	English	P1: All chimpanzees are animals. P2: Some animals are not primates. C: Some primates are not chimpanzees.
	Japanese	P1: すべてのチンパンジーは動物である。 P2: ある動物は霊長類でない。 C: ある霊長類はチンパンジーでない。
OAO	English	P1: Some ghosts are not students. P2: All students are humans. C: Some humans are not ghosts.
	Japanese	P1: ある幽霊は生徒でない。 P2: すべての生徒は人間である。 C: ある人間は幽霊でない。
EAO	English	P1: No robot is human. P2: Every human being is a living organism. C: A certain living organism is not a robot.
	Japanese	P1: どのロボットも人間でない。 P2: すべての人間は生物である。 C: 生物のあるものはロボットでない。
AII	English	P1: All humans are animals. P2: Some robots are animals. C: Some humans are robots.
	Japanese	P1: すべての人間は動物である。 P2: あるロボットは動物である。 C: ある人間はロボットである。

Table 3: Examples of syllogistic inference problems where conversion errors give wrong answer. The correct label for all the examples is *neutral*. The English sentences (**P1**, **P2**, **C**) in each example correspond to the respective Japanese sentences.

We identified the syllogisms in which the correct answer is *neutral* and whose premises contain a sentence of the form *All A are B* or *Some A are not B*, and whose correct answer changes from *neutral* to *entailment* by applying conversion to either one or both premises. We annotate the *conversion* label to this type of problems. In the original dataset, there are only 10 such problems. Thus we expanded the dataset by adding more *conversion* problems that have the types not included in the original dataset. We fixed a set of schematic types to be added and obtained instances of these types by substituting abstract terms with concrete nouns in the dataset. In total, there are 70 problems to which the *conversion* label is assigned in the NeuBAROCO dataset.

4.2 Atmosphere effects

Atmosphere effects are one of the inferential biases that can be traced back to studies in the 1930s (Woodworth and Sells, 1935; Khemlani and Johnson-Laird, 2012). It can be interpreted as two principles (Chater and Oaksford, 1999):

1. *The principle of quality*: if one or both premises are negative (**E**-type or **O**-type), the conclusion should be negative; otherwise, it is positive (**A**-type or **I**-type).
2. *The principle of quantity*: if one or both premises are particular (**I**-type or **O**-type), then the conclusion will be particular; oth-

Type of syllogisms	Language	Example
AOE	English	P1: All animals are living things. P2: Some humans are not living things. C: None of the animals are human.
	Japanese	P1: すべての動物は生物である。 P2: ある人間は生物でない。 C: どの動物も人間でない。
OAI	English	P1: A certain police officer is not a public servant. P2: All human beings are public servants. C: Some police officer is a human being.
	Japanese	P1: ある警察官は公務員でない。 P2: すべての人間は公務員である。 C: ある警察官は人間である。

Table 4: Examples of syllogistic inference problems where atmosphere effects can give wrong answer. The correct label for all the examples is *neutral*. The English sentences (**P1**, **P2**, **C**) in each example correspond to the respective Japanese sentences.

erwise, it is universal (A-type or E-type).

Previous psychological experiments based on the original BAROCO data (for Japanese) have shown that O-type inference is particularly difficult for logically untrained human participants (Shikishima et al., 2009). Thus among various patterns, we focus on the cases where at least one premise is an O-type or I-type sentence. We assign the *atmosphere* label to an inference if its correct answer is *neutral* and (1) at least one of the premises is O-type and the conclusion is either E-type, I-type, or O-type or (2) at least one of the premises is I-type and the conclusion is either I-type or O-type.

Table 4 shows some representative examples of syllogisms satisfying these conditions. In total, there are 104 problems labelled as *atmosphere* in the dataset.

5 Experiments

5.1 Experimental settings

We evaluate syllogistic reasoning ability of deep neural networks, and in particular state-of-the-art large language models using our NeuBAROCO dataset. We evaluate transformer-based pre-trained language models, RoBERTa (Liu et al., 2019) and BART (Lewis et al., 2020), both being fine-tuned with the MultiNLI dataset (Williams et al., 2018). We use the models available in the transformers library, roberta-large-mnli and facebook/bart-large-mnli.² We also evaluate OpenAI’s GPT-3.5 model, an improved version of

²<https://github.com/huggingface/transformers>

GPT-3 (Brown et al., 2020). We use OpenAI’s ChatGPT API with the GPT-3.5-turbo model.³ Table 5 shows some prompt examples in English and Japanese. We chose the best one among several prompts we tried. When we use a phrase like *logical inference* or *syllogism* instead of *inference*, the performance became worse.

To test whether the models show belief biases, conversion errors, and atmosphere effects, we tested how well the models can answer correctly to the problems labelled as *inconsistent*, *conversion*, and *atmosphere*, and compared the accuracies with the total average on the NeuBAROCO dataset.

5.2 Results and discussion

5.2.1 Overall results

Table 6 shows the overall accuracy of each model and the accuracy on each correct label. Table 7 shows the accuracy of the models on basic and extended syllogisms.

RoBERTa The overall accuracy was low (34.67%). The accuracy for the *contradiction* problems was very high (74.65%), although the accuracy of the *neutral* problems was very low (0.55%). The accuracy of the *entailment* problems was between the two (62.3%). The accuracy on extended syllogism was higher than basic syllogism (54.39% and 31.13%).

BART The results on BART shows the same tendency. The overall accuracy was low (35.2%).

³<https://platform.openai.com/docs/model-index-for-researchers>

Carefully evaluate the following inference, and determine whether the premises entail or contradict the conclusion. Answer with entailment, contradiction, or neither.

Premise 1: Some A are B.

Premise 2: All B are C.

Conclusion: All A are C.

次の推論を注意深く読み、前提が結論を含意するか、矛盾するかを判定しなさい。「含意」「矛盾」「どちらでもない」のいずれかで答えなさい。

前提1：あるAはBである。

前提2：すべてのBはCである。

結論：すべてのAはCである。

Table 5: Prompt examples in English and Japanese

Language	Models	All	Entailment	Contradiction	Neutral
English	RoBERTa	34.67	62.30	74.65	0.55
	BART	35.20	55.74	83.10	2.75
	GPT-3.5	51.73	79.51	38.03	38.46
Japanese	GPT-3.5	48.27	80.33	54.93	24.18

Table 6: Accuracy (%) of the models on all the inference problems and each correct inference label.

Language	Models	Basic	Extended
English	RoBERTa	31.13	54.39
	BART	31.13	57.89
	GPT-3.5	51.57	52.63
Japanese	GPT-3.5	46.86	56.14

Table 7: Accuracy (%) of the models on basic and extended syllogisms.

Language	Models	All	Symbol	Consistent	Inconsistent	Conversion	Atmosphere
English	RoBERTa	34.67	24.21	46.11	22.55	0.0	0.0
	BART	35.20	34.74	45.51	15.69	1.43	0.96
	GPT-3.5	51.73	61.05	56.89	31.37	25.71	39.42
Japanese	GPT-3.5	48.27	55.79	56.29	25.49	21.43	22.12

Table 8: Accuracy (%) of the models on each type of syllogistic inferences and biases.

BART answered correctly in most *contradiction* cases, but less in *entailment* and still less in *neutral* (83.1%, 55.74% and 2.75%). The performance on extended syllogism was better than that on basic syllogism (57.89% and 31.13%).

GPT-3.5 The overall accuracy was 51.73%. The accuracy on the *entailment* problems was very high (79.51%), while that on the *contradiction* and *neutral* problems were low (38.03% and 38.46%). We found little difference between basic and extended syllogism (51.57% and 52.63%).

5.2.2 Results on problems concerning biases

Table 8 shows the accuracy of the models on each type of syllogistic inferences and biases.

Belief biases Among the three types of inferences, *symbol*, *consistent*, and *inconsistent*, the performance on the *inconsistent* cases was the lowest in every model. We found a significant difference among the models on which of the *symbol* or *consistent* cases they answered most accurately. GPT-3.5 performed better in *symbol* and *consistent* than *inconsistent* (61.05%, 56.89% and 31.37%). For RoBERTa, the percentage of correct answers to the

consistent cases was higher than in *inconsistent* and *symbol* cases (46.11%, 22.55% and 24.21%). The percentages of correct responses of BART were, in order of highest to lowest, *consistent*, *symbol*, and *inconsistent* (45.51%, 34.74% and 15.69%).

Overall, the results show a tendency that GPT-3.5 outperforms both RoBERTa and BART models in symbolic reasoning. Also, GPT-3.5 performed equally well on *symbol* and *consistent* problems. This suggests that *symbol* and *consistent* are relatively easy to handle in that both types of problems are not contrary to beliefs. These results contrast with the results on RoBERTa, which performed almost equally on *symbol* and *inconsistent* problems.

Conversion error Regarding the *conversion* problems, all models exhibit low performance. A striking difference exists between RoBERTa and BART, on one hand, and GPT-3.5, on the other. While RoBERTa and BART hardly answered correctly (0% and 1.43%), GPT-3.5 performed better, answering correctly almost a quarter of the cases, which is about half of the overall average (25.71%). The results show that all the models performed poorly on the problems where incorrect responses are made by conversion errors.

Atmosphere effects We found that the performances of the models were notably lower in *atmosphere* cases. While RoBERTa and BART hardly provide correct answer to *atmosphere* cases (0% and 0.96%), GPT-3.5 performed better (39.42%).

5.2.3 Results on the Japanese GPT model

The results on the Japanese GPT model shows the strikingly same tendency as the English GPT models. One notable exception is the performance on the *contradiction* problems (see Table 6), where the Japanese GPT-3.5 model performed better than the English GPT-3.5 models. By contrast, the performance on *atmosphere* problems was worse than that of the English model.

6 Conclusion

In this paper, we investigated syllogistic reasoning ability of current large language models in focusing on human-like biases that have been studied in the context of cognitive science of human reasoning. The experiments indicated that the state-of-the-art models fail for problems where errors are caused by various human-like biases, and that there is large room for improvement in deductive reasoning capabilities of large language model.

Among other things, our results on conversion errors suggest the importance of distinguishing the problems of interpreting sentences (in particular, interpreting quantifiers and negation) from the problem of performing logical inferences. For conversion cases, there is abundant room for discussion on at which level the models mistook. Further inquiry into this issue could provide insight into a better understanding of the behavior of neural models.

There remain many issues to be addressed. First, although we tested the models with no solved examples, experiments on few-shot learning will be insightful. Dasgupta et al. (2022) reported that the models performed syllogistic reasoning better with few-shot learning, whether or not the content of the inferences were consistent with common-sense beliefs. In addition, we suppose that the performances can be different depending on the wording in instructions. Further research will contribute to improvement of instructions.

Second, we showed that the models performed worse in the cases in which humans have troubles because of some well-known human biases. It is left for future work to make more detailed comparisons between humans and neural models, which is a promising research direction since the original BAROCO data on human reasoning ability is available.

Finally, it is interesting to consider extended forms of natural logic inferences other than basic syllogisms, including relational syllogism with transitive verbs and comparatives and those inferences with generalized quantifiers such as *most*. It is left for future work to examine whether the models show similar biases for such extended syllogism inferences.

Author Contributions

The three authors, Risako Ando, Takanobu Morishita, and Hirohiko Abe, made equal contributions as core authors.

Acknowledgements

We express our gratitude to the anonymous reviewers for their helpful feedback, which has contributed to the improvement of this work. We would like to express our thanks to Prof. Chizuru Shikishima for the discussion on earlier versions of the BAROCO inference tasks collection. This work is partially supported by JST, CREST grant number JPMJCR2114, MEXT-JSPS Kakenhi MKK279H

and MKK477J.

References

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Nick Chater and Mike Oaksford. 1999. The probability heuristics model of syllogistic reasoning. *Cognitive Psychology*, 38(2):191–258.
- Ishita Dasgupta, Andrew K Lampinen, Stephanie CY Chan, Antonia Creswell, Dharshan Kumaran, James L McClelland, and Felix Hill. 2022. Language models show human-like content effects on reasoning. *arXiv 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.
- Jonathan St.B. T. Evans. 1989. *Bias in Human Reasoning: Causes and Consequences*. Lawrence Erlbaum Associates, Inc.
- Jonathan St.B. T. Evans, Stephen E. Newstead, and Ruth M. J. Byrne. 1993. *Human Reasoning: The Psychology of Deduction*. Psychology Press.
- Bart Geurts. 2003. [Reasoning with quantifiers](#). *Cognition*, 86(3):223–251.
- Vinod Goel, Christian Buchel, Chris Frith, and Raymond J Dolan. 2000. Dissociation of mechanisms underlying syllogistic reasoning. *Neuroimage*, 12(5):504–514.
- Sangeet Khemlani and Philip N Johnson-Laird. 2012. Theories of the syllogism: A meta-analysis. *Psychological Bulletin*, 138(3):427.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- 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 preprint arXiv:1907.11692*.
- Lawrence S. Moss. 2015. Natural logic. In Shalom Lappin and Chris Fox, editors, *The Handbook of Contemporary Semantic Theory*, 2 edition, pages 559–592. Wiley.
- Stephen E. Newstead. 1989. Interpretational errors in syllogistic reasoning. *Journal of Memory and Language*, 28(1):78–91.
- Stephen E Newstead, Paul Pollard, Jonathan St BT Evans, and Julie L Allen. 1992. The source of belief bias effects in syllogistic reasoning. *Cognition*, 45(3):257–284.
- Kyle Richardson, Hai Hu, Lawrence Moss, and Ashish Sabharwal. 2020. Probing natural language inference models through semantic fragments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8713–8721.
- Víctor Sánchez Valencia. 1991. *Studies on Natural Logic and Categorical Grammar*. Ph.D. thesis, University of Amsterdam.
- Yuri Sato and Koji Mineshima. 2015. How diagrams can support syllogistic reasoning: an experimental study. *Journal of Logic, Language and Information*, 24:409–455.
- Viktor Schlegel, Kamen Pavlov, and Ian Pratt-Hartmann. 2022. [Can transformers reason in fragments of natural language?](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11184–11199, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Chizuru Shikishima, Juko Ando, Pierre Grialou, Ryo Takemura, and Mitsuhiro Okada. 2006. A behavioral genetic study of syllogism solving using linguistic and graphical representations: a preliminary report. In *Images and Reasoning*, pages 69–85. Keio University Press.
- Chizuru Shikishima, Kai Hiraishi, Shinji Yamagata, Juko Ando, and Mitsuhiro Okada. 2015. Genetic factors of individual differences in decision making in economic behavior: A Japanese twin study using the Allais problem. *Frontiers in Psychology*, 6:1712.
- Chizuru Shikishima, Kai Hiraishi, Shinji Yamagata, Yutaro Sugimoto, Ryo Takemura, Koken Ozaki, Mitsuhiro Okada, Tatsushi Toda, and Juko Ando. 2009. [Is g an entity? a Japanese twin study using syllogisms and intelligence tests](#). *Intelligence*, 37(3):256–267.
- Keith Stenning and Michiel Van Lambalgen. 2012. *Human Reasoning and Cognitive Science*. MIT Press.
- Johan van Benthem. 1986. *Essays in Logical Semantics*. Reidel, Dordrecht.
- Minna Cheves Wilkins. 1928. The effect of changed material on ability to do formal syllogistic reasoning. *Archives of Psychology*, 16:1–83.

- 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.
- Robert S Woodworth and Saul B. Sells. 1935. An atmosphere effect in formal syllogistic reasoning. *Journal of Experimental Psychology*, 18(4):451–460.
- Hitomi Yanaka, Koji Mineshima, Daisuke Bekki, Kentaro Inui, Satoshi Sekine, Lasha Abzianidze, and Johan Bos. 2019. Can neural networks understand monotonicity reasoning? In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 31–40.