Limits for learning with Language Models

Nicholas Asher
CNRS, IRIT
asher@irit.fr

Swarnadeep Bhar
Université de Toulouse 3
IRIT
swarnadeep.bhar@irit.fr

Akshay Chaturvedi
Université de Toulouse 3
IRIT
akshay91.isi@gmail.com

Julie Hunter
LINAGORA
jhunter@linagora.com

Soumya Paul
TELINDUS
soumya.paul@gmail.com

arXiv:2306.12213v1 [cs.CL] 21 Jun 2023

Abstract

With the advent of large language models (LLMs), the trend in NLP has been to train LLMs on vast amounts of data to solve diverse language understanding and generation tasks. The list of LLM successes is long and varied. Nevertheless, several recent papers provide empirical evidence that LLMs fail to capture important aspects of linguistic meaning. Focusing on universal quantification, we provide a theoretical foundation for these empirical findings by proving that LLMs cannot learn certain fundamental semantic properties including semantic entailment and consistency as they are defined in formal semantics. More generally, we show that LLMs are unable to learn concepts beyond the first level of the Borel Hierarchy, which imposes severe limits on the ability of LMs, both large and small, to capture many aspects of linguistic meaning. This means that LLMs will continue to operate without formal guarantees on tasks that require entailments and deep linguistic understanding.

1 Introduction

The success of large language models (LLMs) has led researchers in NLP to harness LLMs trained on vast amounts of data to solve a variety of language understanding and generation tasks, and some have claimed that LLMs can solve any task that can be specified via prompting (Brown et al., 2020). While the list of LLM successes is long, there have been several recent papers that provide empirical evidence that LLMs at least sometimes fail to capture important aspects of linguistic meaning (Kuhnle and Copestake, 2019; Sinha et al., 2020; Yuksekgonul et al., 2022; Chaturvedi et al., 2022; Kalouli et al., 2022). Those who have dabbled in "BERTology" with respect to linguistic meaning often have the feeling that fixing one LLM deficiency just leads to the discovery of new ones.

This paper provides a theoretical explanation of certain of these observed failings of LLMs. In particular, we prove that LLMs cannot learn the notions of semantic entailment or consistency as defined in formal semantics (Dowty et al., 1981) because they are incapable of mastering universal quantification. Our work builds on Siegelmann and Sontag (1992); Siegelmann (2012); Weiss et al. (2018), concerning the expressive power of neural networks, but we focus on the learnability of semantic concepts and use novel tools.

Our argument has widespread implications: not only does a general capacity to recognize semantic entailment and consistency underlie everyday conversational interactions, but the meanings of a great many common linguistic expressions depend on universal quantification. This set includes—but is certainly not limited to—a long list of quantifiers (*every, some, many, most,... every other*, ...), temporal adverbs (*always, never, eventually*) that are essential to planning (Lamport, 1980), modal operators (*possibly*, *necessarily*,...), and certain discourse connectives and adverbs (*therefore*, *if / then*, *except*, *because*, ...).

We begin in Section 2 by contextualizing our claims in terms of expectations about the linguistic capacities and applications of LLMs. In Section 3, we introduce the framework of continuation semantics, which will allow us to adapt certain notions central to truth-conditional semantics to the case of LLMs. Section 4 lays out the core of our theoretical argument, focusing first on what is needed to learn universal quantification and then generalizing our argument to a wide range of linguistic expressions. Our theoretical argument suggests that we should expect certain empirical failures from LLMs, and in Section 5, we provide evidence that our predictions are borne out. Section 6 concludes.

2 Context

Our results are particularly relevant to downstream tasks that require an agent to not only create fluent, creative and contextually relevant speech but also to act precisely based on the meaning of linguistic expressions and reliably recognize semantic inconsistency. For a robot that has been instructed (via conversation) to tighten *every* screw of a door, to *never* walk on an airplane wing, or to *stop* drilling *immediately if* certain conditions hold, acting appropriately requires being able to infer what do to based on the *linguistic meaning* of the words *every*, *never*, *stop*, *immediately* and *if*—and in these cases, getting things *mostly* right won't do, especially if lives or substantial economic loss are at risk.

An important corollary of our argument is that while it might be tempting to separate reasoning and linguistic competence (Mahowald et al., 2023), the former is in fact inextricably tied to our ability to draw inferences based on *linguistic* content—not just on, say, mathematical or real-world facts. This in turn suggests that approaches which attempt to patch up knowledge deficiencies for LLMs by giving them access to external models (Mialon et al., 2023) will fall short in developing reliable models of linguistic understanding because LLMs fail to grasp the notions that underlie the very way that sentences (and actions) are woven together in conversation.

Empirical studies like Chaturvedi et al. (2022) show that LLM failures to respect semantic entailment in question answering tasks follow from fundamental features of LLM training; thus while extensive training and large data sets may improve LLM results, performance will inevitably remain unstable and we should continue to expect hallucinations and reasoning errors in NLP tasks like question-answering and natural language inference.

3 Language models and formal semantics with continuations

3.1 LLMs and strings

We consider LLMs trained on transformer architectures over very large corpora using classic language modeling tasks, namely masked language modeling or next sentence prediction. The former involves masking certain words in a given corpus and training the model to guess the missing words, while in the latter, a context (a sentence typically) is provided to the model, which is trained to predict the sentence that follows. This unsupervised training allows language models to build rich internal representations that have been shown through probing to contain at least implicitly a large amount of linguistic information (Devlin et al., 2019; Liu et al., 2019; Tenney et al., 2018).

Formally, LLMs learn a function $f : C \times X \rightarrow [0, 1]$ that assigns a probability to a word (or string or discourse move) $x \in X$ given a context (or finite string) C. More abstractly, let V be a countable set called the vocabulary. For $i > 0$, let V^i denote the set of all length i strings in the vocabulary V and $V^{\leq i}$ denote the set of all strings V whose length is at most i. V^* denotes the set of all finite strings and V^ω the set of countably infinite strings in V. We can then rewrite f as $f\colon V^{\leq n} \rightarrow \mu$, where μ is a probability measure (which is often called its *prediction*) over V^{n+m} for $m \geq 1$. Typically, the prediction function is used on strings of length m where m is smaller than n.

By exploiting f, an LLM can extend μ to a distribution on the set of strings V^*. The most straightforward way is to follow autoregressive models that calculate the probability of strings via conditionalization. For a new sentence $s' = (w_1, w_2, ..., w_{m+1})$, and an input string s of length n provided as context, we have:

$$\mu^{n+m+1}(s'|s) = \mu^{n+1}(w_1|s) \times \mu^{n+2}(w_2\,|s, w_1) \times \ldots \times \mu^{n+m}(w_n\,|s, w_{m-1}, ..., w_1) \quad (1)$$

For any $s' \in V^*$, $\mu(s')$ represents the confidence with which an LLM predicts s', after training on strings in $V^{\leq n}$.

3.2 Linguistic meaning

In what follows, we are in interested strings that have a well formed meaning and are evaluable as true or false. Linguists use truth conditional semantics to define the meanings of strings or well formed sentences in terms of the conditions under which they are true. Thanks to the work of Tarski (1944, 1956), we can formalize the notion of truth conditions using the set-theoretic notion of a model that defines denotations or truth conditions for sentences recursively from denotations for sentential constituents (Dowty et al., 1981).

The notion of a model not only serves to define truth conditions; it also captures entailments. We

define the notion of *semantic consequence* using the notion of a model or structure $\mathfrak{A}$ as follows (Chang and Keisler, 1973):

Definition 1. ϕ is a semantic consequence of Γ (in symbols, $\Gamma \models \phi$) if and only if in every structure $\mathfrak{A}$ in which Γ is satisfied ($\mathfrak{A} \models \Gamma$), $\mathfrak{A}$ also makes true or satisfies ϕ ($\mathfrak{A} \models \phi$). That is: $\forall \mathfrak{A}$, $\mathfrak{A} \models \Gamma \Rightarrow \mathfrak{A} \models \phi$

The notion of semantic consequence integrates entailment with truth conditional meaning; two strings have exactly the same entailments just in case they are true in the same models. Accordingly we can capture the truth conditional meaning of a string in terms of the strings that it entails. *Socrates is a man*, for example, entails *Socrates is human*, *Socrates is mortal*, *Socrates is an adult* but also that *someone* is a man, human, mortal and so on. What it means for Socrates to be a man (and, indirectly, the meaning of *man*) can be captured by the full set of these entailments.

Our idea is to apply truth conditional semantics to LLMs by representing models themselves as strings. Semanticists have used strings and *continuation semantics* (Reynolds, 1974) —in which the meaning of a string s is defined in terms of its possible *continuations*, the set of longer strings S that contain s—to investigate the meaning and strategic consequences of conversational moves (Asher et al., 2017), temporal expressions (Fernando, 2004), generalized quantifiers (Graf, 2019), and the "dynamic" formal semantics of (Kamp and Reyle, 1993; Asher, 1993)(De Groote, 2006; Asher and Pogodalla, 2011). In our case, we will use strings to define models $\mathfrak{A}_s$. We will use this trick to reformulate semantic consequence: where $\|\phi\|$ is the set of strings describing models that satisfy a truth evaluable string ϕ:$\Gamma \models \phi$ iff $\|\Gamma\| \subseteq \|\phi\|$.

LLMs naturally find their place in such a framework (Fernando, 2022): given their training regime, the meaning of any natural language expression for an LLM is a function from input contexts to sets of larger strings or continuations. LLMs provide a probability distribution over possible continuations and can predict possible continuations of a given text or discourse.

4 Learning limits for semantic concepts

Semantic consequence defines linguistic entailments and importantly provides the fundamental connection between meaning and inference that ensures linguistic understanding (Montague, 1974). Crucial to $\models$ is the use of universal quantification over all possible structures—an infinite space of possible circumstances of evaluation or set of possibilities. A true grasp of semantic consequence thus requires an understanding of universal quantification at least over countably infinite domains. In Section 4.1, we show that an LLM's training regime makes it fundamentally unable to learn the concept of universal quantification. In Section 4.2, we generalize our argument to show that LLMs are incapable of learning a wide variety of everyday semantic concepts.

4.1 Learning the full meaning of *every*

To see if the set of strings that define the concept *every* is learnable for an LLM, consider (1).

(1) Every object is blue.

We will use strings of atomic formulas and their negations to define models (or more precisely their atomic diagrams) that we will use to test whether an LLM M can learn the concept of universal quantification through inductive reasoning from a series of individual trials over finite subsequences of strings representing countably infinite domains. In particular, we will ask whether it is possible to train an LLM M to judge, for a string s of arbitrary length n, whether s is consistent with (1), or equivalently, given that s defines a model $\mathfrak{A}_s$, whether given $\mathfrak{A}_s$, (1) is true. If M can reliably judge in which models $\mathfrak{A}_s$ (1) is true, we can conclude it has learned the meaning of *every*.

To this end, consider a language $\mathcal{L}$ containing negation, the predicate *is blue* and a countably infinite number of constants a_i enumerating objects of a countably infinite domain. $\mathcal{L}$ formulas are of the form a_i *is blue* and a_i *is not blue*. We use the formulas of $\mathcal{L}$ as "words" to define the set of finite strings, $V_{\mathcal{L}}^*$ and the set of countably infinite strings $V_{\mathcal{L}}^\omega$. Each such string corresponds to a finite or countably infinite model in which (1) is true or not.

In the course of training, M will be presented with finite sequences that define structures of increasing size. For each n and set of models of size n, M will form a set of hypotheses $\mathcal{H}^n$, where for $h \in \mathcal{H}^n, h : V_{\mathcal{L}}^n \rightarrow \{0, 1\}$. $\mathcal{H}^n$ corresponds to the hypothesis space of the problem; each h^n says whether a presented sequence of length n is consistent with (1). As each $h^n \in \mathcal{H}^n$ is a characteristic function of a subset of $V_{\mathcal{L}}^n$, we can identify

hypotheses with sets of strings. So for instance, $h^n_\forall$ is the set of strings in $V^n_\mathcal{L}$ that are consistent with (1) and that define models in which (1) is true. We will additionally assume that $h^n_\forall$ picks out a suitable set for each V^n, and similarly for each h^n_k.

However simply learning $h^n_\forall$ for some n will not be sufficient for M to learn the meaning of *every*. Universal quantification is a concept that applies to arbitrarily large domains. So the question, *Can M inductively learn the meaning of* every*?* becomes *Can M inductively learn hypothesis $h^\omega_\forall \in \mathcal{H}^\omega$?*

To answer this question, we first have to specify what we mean by inductive learning. Recall that an LLM M has learned from unsupervised training a function $f\colon V^{\leq n} \rightarrow \mu(V^{\leq n+m})$ with $\mu(V^{\leq n+m})$ a probability distribution over completions of length m of contexts of length n. An LLM can use this distribution to compute probability values for arbitrarily long strings or continuations using Equation 1.

In the case at hand, M needs to use this distribution over $\mathcal{L}$ strings to compute the probability that a string s is in $h_\forall$ or the probability of s given $h_\forall$. To learn inductively M must use its training data $D^{\leq n}$ to update its prior for the distribution μ^n using a rationally justifiable form of *inductive inference*; e.g., for $h \in \mathcal{H}^n$, $\mu^n(s|h) = \frac{\mu^n(h|s)\times\mu^n(s)}{\mu^n(h)}$.

Additionally, we consider two constraints on distributions to define learning in terms of an inductively inferred change in the distribution from the priors. The first constraint, *Max Ent*, says that the distribution μ prior to training should assign all hypotheses a weight based on maximum entropy or a least informative distribution. This is usual with auto-regressive models and a common assumption in other models.

The second constraint is that distributions for inductive learning should be *non-degenerate*. We have assumed that our LLM M has been trained over sequences of length n. Through Equation 1, M can extend the distribution it has learned for V^n to one over V^{n+m} for any string of finite length $n+m$. Recall that we are looking at strings of $\mathcal{L}$ that define structures; the structures defined by strings of length $n+m$ are independent of those defined in V^n and none is intuitively more likely than another. So the prior distribution over V^{n+m} should consider as equally likely all continuations $s.a \in V^{n+m}$, where $s \in V^n$, $a \in V^m$ and . is concatenation. There are also correspondingly more hypotheses in $\mathcal{H}^{n+m}$ than in $\mathcal{H}^n$, since there are $V^{|m|}$ more strings in V^{n+m} than in V^n. Thus $\mu^{n+m}(s.a|h_k) < \mu^n(s.a|h_k)$ for $s.a \in V^{n+m}, s \in V^n$ for each h_k. Non-degenerate distributions will reflect this and should make the model converge to the least general hypothesis supported by the evidence (Muggleton et al., 1992; Plotkin, 1972).

Definition 2. *M's distributions over sets of hypotheses $\mathcal{H}^n$, $\mu^n(\mathcal{H}^n)$, after training over V^n are non-degenerate* if $\forall h\ \forall\delta\ (0 < \delta \leq 1), \exists m > 0$ such that $\forall a \in V^m\ \forall s \in V^n : \mu^n(s.a|h) = \max\{0, \mu^n(s|h) - \delta\}$, where $s.a \in V^{n+m}$.

Proposition 1. Models that calculate distributions over strings using Equation 1 have non-degenerate distributions.

As continuations get longer the probability of the continuation will decrease monotonically.□

Because quantifiers like *every* and *some* are eliminable in terms of Boolean functions when we consider finite structures definable with strings in V^*, we must consider strings in V^ω to define countably infinite models that capture the full truth conditions of *every*. To extend a distribution over V^n for finite n to a distribution over V^ω, we lift the probability of a string to the set of its continuations. In V^ω, the set of strings A characterizes the set $A.V^\omega$, where $A.V^\omega$ is the set of all strings formed by concatenating a string from A with a string from V^ω. Using this correspondence, the probabilities of sets of finite strings in V^n can lifted to probabilities of sets of the form $V^n.V^\omega$. The laws of probability extend the distribution to complements, intersections and unions of such sets.

We now propose a simple but general notion of inductive learning.

Definition 3. Suppose μ_0 is M's Max Ent prior distribution and let $h \in \mathcal{H}^\beta$ for some countable β. *M effectively learns h* iff after some finite amount of training using inductive inference, there is an α, such that: for any $s \in V^\beta$, $\mu^\beta(s|h) > \alpha > \mu_0(s|h)$ iff $s \in h$.

Proposition 2. If M can effectively learn $h^n_\forall$ from sequences of $V^n_\mathcal{L}$ for arbitrarily large $n \in \omega$, then M can effectively learn $h^\omega_\forall$

Assume that M cannot effectively learn $h^\omega_\forall$ but it can effectively learn $h^n_\forall$ for arbitrarily large $n \in \omega$. Then it must admit some string $s \in V^\omega_\mathcal{L}$, such that $s \notin h^\omega_\forall$. But then at some finite stage i, s_i must

have $\neg blue(a_i)$. By hypothesis M has learned $h^i_\forall$. So it has ruled out s_i and *a fortiori* s. □

We now negatively answer our question, *Can M inductively learn hypothesis $h^\omega_\forall$?*, under either of two conditions: (i) M has non-degenerate distributions; (ii) M obeys Max Ent and inductive inference.

Proposition 3. Suppose M's distributions are non-degenerate. Then $h^\omega_\forall$ is not effectively learnable by M over $\mathcal{H}^\omega$.

Suppose M trained on $V^{\leq n}$ has effectively learned $h^n_\forall$. So $\forall s \in h^n_\forall$, $\mu^n(s|h_\forall) > \alpha$ where α is as in Definition 3. Since M's distributions are non-degenerate, $\exists m$, such that for all $s \in V^m$, $\exists \delta$: $0 < \delta \leq 1$ where $\mu^m(s|h_\forall) - \delta < \alpha$ and a continuation of s, $s.a$, such that $s.a \in h^{m+n}_\forall$ but $\mu^{m+n}(s.a|h_\forall) = \mu^m(s|h_\forall) - \delta < \alpha$. □

By Propositions 1 and 3, a basic auto-regressive model cannot learn $h_\forall$. We can generalize Proposition 3 to other :

Proposition 4. Suppose M's priors only obey Max Ent and M uses inductive inference. Then $h^\omega_\forall$ is not effectively learnable by M over $\mathcal{H}^\omega$.

Suppose M's training data $D^{\leq n} \subseteq V^{\leq n}$ and M has learned $h^n_\forall$. To learn $h^{n+m}_\forall$, M must project its distribution of $\mathcal{H}^n$ onto $\mathcal{H}^{n+m}$. But the distributions in $\mathcal{H}^n$ and $\mathcal{H}^{n+m}$ are *independent*; for one thing the cardinality of $\mathcal{H}^n$, $|\mathcal{H}^n|$, is such that $|\mathcal{H}^n| < |\mathcal{H}^{n+m}| = |\mathcal{H}^n| \times 2^m$. Our assumptions about inductive inference on $D^{\leq n}$ make it no more likely that *every* will be associated with $h_\forall$ than it is with any of the 2^m $h \in \mathcal{H}^{n+m}$, where strings in h contain the same n prefix as an $s \in h^n_\forall$ but $h \cap h^{n+m}_\forall = \varnothing$. In $\mathcal{H}^{n+m}$ these hypotheses h can be distinguished from $h_\forall$. Max Ent priors over $\mathcal{H}^{n+m}$ imply that for any $s \in V^{n+m}$, $\mu^{n+m}(s|h_\forall) = \frac{1}{2^m}\mu^n(s|h_\forall)$. □

Corollary 1. M cannot effectively learn $h^n_\forall$ from sequences in $V^n_\mathcal{L}$ for arbitrarily large $n \in \mathbb{N}$. There is some n such that $h^n_\forall$ is not effectively learnable.

While LLMs can represent any Borel function to an arbitrary degree of precision (Hornik et al., 1989), Propositions 3 and 4 shows they cannot always learn such functions, given either the constraints of inductive epistemology or the way LLMs generate probabilities for strings. In particular, given our assumptions, no LLM can effectively learn $h^\omega_\forall$. In addition, each LLM is bounded by some number n in the size of sequences for which it can learn $h^n_\forall$. LLMs do not have the capacity to learn the meaning of 'every' even over all finite domains.[1]

Even supposing that an LLM can effectively learn $h^n_\forall$ for some n, this does not amount to understanding *every*. $h^n_\forall$ can be effectively represented with quantifier free conjunctions of formulae, and these do *not* correctly approximate reasoning with a sentence like (2) that applies to arbitrarily large domains. Identifying $\forall$ with a finite conjunction of length n will make $\forall x Fx$ consistent with $\neg\forall x Fx$ in larger structures. In ω structures, for example, $\neg\forall x Fx$ is consistent with every finite subset of the Π^0_1 string $blue(0)$, $blue(1)$, $blue(2)$, ..., in $h^\omega_\forall$, making it inevitable that LLMs will reason incorrectly with *every* in large enough structures.

The situation worsens with sampling: suppose that when we present our model M a long string, M only samples some of the elements in the string; the threat of inconsistency in such a situation can become high and we have no guarantees that such inconsistencies will not arise.[2] But this reasoning is not *independent* of the meaning of *every*; as the semantics and rules of first order logic show, this reasoning is *an integral part* of the meaning of *every*. As a result, LLMs unable to grasp semantic consequence defined in terms of universal quantification; and we thus cannot provide them guarantees that they follow semantic entailments when asked to do semantic tasks. This predicts phenomena like LLM hallucinations and observed elementary reasoning errors.

4.2 Generalizing our answer to Q2

Using tools from statistical learning and the Borel Hierarchy, we now generalize Propositions 3 and 4 to other concepts beyond *every*.

Statistical learning examines the application of a learned function over a test domain and the expected loss over novel applications. The ability to bring the error over test to that over the training set is typically taken to indicate an ability to generalize (Neyshabur et al., 2017). Villa et al. (2013) define learnability in statistical learning theory via the notion of *uniform consistency*. Let μ be a dis-

[1] Unlike Hume's problem of induction (Popper, 1963) and (Wolpert et al., 1995), we exploit particularities of LLMs and the structure of a classification problem. The finite bound on learning of hypotheses goes beyond standard Humean conclusions

[2] Approximation and approximation error can also affect learnability of mathematical functions (Colbrook et al., 2022).

tribution over $\mathcal{H}$ and μ_n the update of μ after n training samples $z_i = (x_i, y_i)$. Let A_{z_n} be an algorithm for picking out a hypothesis from $\mathcal{H}$ based on n training samples. $inf_{\mathcal{H}}$ is the hypothesis in $\mathcal{H}$ with the lowest possible error (Shalev-Shwartz et al., 2010; Kawaguchi et al., 2017).

Definition 4. An algorithm A on a hypothesis space $\mathcal{H}$ is uniformly consistent if and only if
$\forall \epsilon > 0 \, lim_{n\to\infty} sup_{\mu}$
$\mu_n(\{z_n : \mathbb{E}_{\mu}(\{A_{z_n} - inf_{\mathcal{H}} \mathbb{E}_{\mu} > \epsilon\}) = 0$

In our case, the best hypothesis, $inf_{\mathcal{H}}$, for instance $h_{\forall}$, will yield 0 error. Our question is whether there is an algorithm that converges to that hypothesis given a certain $\mathcal{H}$ and certain assumptions.

Definition 5. A class of hypotheses $\mathcal{H}$ is *uniformally learnable* just in case there exists a uniformly consistent algorithm for $\mathcal{H}$.

This enables us to link learnability with a number of other features:

Theorem 1. (Anthony et al., 1999) Let $Y = \{0, 1\}$. Then the following conditions are equivalent: (i) $\mathcal{H}$ is uniformly learnable; (ii) Empirical risk minimization on $\mathcal{H}$ is uniformly consistent; (iii) $\mathcal{H}$ is a uGC-class; (iv) the VC-dimension of $\mathcal{H}$ is finite.

The Borel Hierarchy We now turn to generalize the hypotheses we are investigating. V^{ω} has a natural topology, the Cantor topology, which allows us to characterize linguistic concepts precisely. To define the topology, we first define the basic open sets to be sets of the form $A.V^{\omega}$, denoted as $\mathcal{O}(A)$, where $A \subseteq V^*$ is a set of finite strings. Importantly, $\mathcal{O}(A)$ sets are both open and closed or *clopen*, because if $A \subset V^*$ is a countable set, then the complement of $A.V^{\omega}$, $(V^* \setminus A).V^{\omega}$, is also open. And thus, $A.V^{\omega}$ is also closed. The Δ^0_1 class is at the intersection of the Σ^0_1 and Π^0_1 classes and consists of the clopen sets. Σ^0_1 sets include countable unions of Δ^0_1 sets, while Π^0_1 are complements of Σ^0_1 sets and so include countable intersections of Δ^0_1 sets.

These sets form the basis of the *Borel hierarchy* of sets that includes the Δ^0_1, Σ^0_1, and Π^0_1 sets, and more generally includes $\Sigma^0_{\alpha+1}$ as the countable union of all Π^0_{α} and Δ^0_{α} sets, and $\Pi^0_{\alpha+1}$ as the complement of $\Sigma^0_{\alpha+1}$ sets, with $\Delta^0_{\alpha} = \Sigma^0_{\alpha} \cap \Pi^0_{\alpha}$. The hierarchy is strict and does not collapse (Kechris, 1995). We will use this hierarchy to characterize linguistic concepts. Below is a picture of some simple Borel sets and their $\subseteq$ relations.

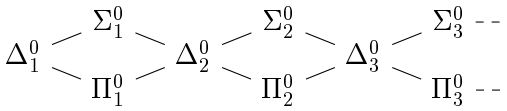

As an example, $h^{\omega}_{\forall} \subseteq V^{\omega}_{\mathcal{L}}$ of the previous section is a Π^0_1 Borel set; i.e., $h^{\omega}_{\forall} = \bigcap_{i\in\omega} B_i$ where the B_i are Δ^0_1.

We are interested in the learnability of Borel sets B with respect to a hypothesis space. The hypothesis space $\mathcal{H}^n$ for $V^{\leq n}$ and algorithem $\mathcal{A}^n$ that an LLM can consider is typically fixed by the maximal length strings it has been trained on. But we will be looking at how an LLM extends its training generalizing to longer and longer strings. More generally, we consider a countable collection of hypotheses—in the case of *every* and $V_{\mathcal{L}}$, the set consists of $h_{\forall}, h_{the first\ 2^n}$ etc. We will assume a countable hypothesis space $\mathcal{H}^{\omega}$ for the Borel sets in V^{ω}, with $|V| > 2$ we want to learn in what follows.

Definition 6. An LLM M can effectively learn a Borel set $S \subset V^{\omega}$ out of a countable set of hypotheses $\mathcal{H}$ iff M has a uniformly consistent algorithm such that $h_S = inf\mathcal{H}$, as in Definition 4, and where $h_S : V^{\omega} \to \{0, 1\}$ defines S.

Clearly if h is $inf\mathcal{H}$, and $\mathcal{A}$ is uniformly consistent, then Definition 3 is satisfied; i.e., there is some $\alpha > \epsilon$ such that $\mu(s|h) > \alpha$ iff $s \in h$.

Theorem 2. An LLM with either (a) non degenerate distribution or (b) Max Ent priors and trained on $V^{<n}$ for some finite n to learn $h \subset V^{\omega}$ via inductive inference (i) can effectively learn a Δ^0_1 set $\mathcal{O}(S) \subset V^{\omega}_{\mathcal{L}}$, where S is a finite subset of $V^{\leq n}$, given $\mathcal{H}_{\mathcal{O}(V^{<n})}$, a hypothesis space restricted to Δ^0_1 sets; but (ii) it cannot effectively learn any Π^0_1 Borel set $B \subset V^{\omega}_{\mathcal{L}}$.

We first show (i). Let $\mathcal{H} = \{\mathcal{O}(A) : A \subseteq V^{<n}\}$. Any $h \in \mathcal{H}$ is determined by a finite set of prefixes P in $V^{<n}$. There are only finitely many such sets in $V^{\leq n}$, and so M has an algorithm A that eliminates at each finite stage of training some Δ^0_1 $\mathcal{O}(P)$ sets. This enables it to converge uniformly toward 0 expected error for the set of finite prefixes that determines $\mathcal{O}(S)$ and so eventually M will have effectively learned $\mathcal{O}(S)$.

Now for (ii). We first consider the case (ii.a) where our learned model has non-degenerate distributions. Consider an arbitrary Π^0_1 complete set

B. So $B = \bigcap_{n\in\omega} \mathcal{O}(B_n)$, with $\mathcal{O}(B_{n+1}) \subset \mathcal{O}(B_n)$, where the $B_i \subset V^*$. To compute B, M needs a uniformly consistent algorithm $\mathcal{A}$ over our countable hypothesis space $\mathcal{H}$ that converges on h_B, the hypothesis defining B. Now suppose M has been trained on strings in $V^{<n}$; its algorithms $\mathcal{A}$ are thus restricted to $\mathcal{H}^{<n}$.

Suppose M trained on $V^{\leq n}$ has effectively learned $h^n_{B_n}$. Let $s \in h^n_{B_n}$. Since M's distributions are non-degenerate, $\forall\alpha \geq 0, \exists m, \delta : 0 < \delta \leq 1$ where $\mu^m(s|h_B) - \delta < \alpha$ and a continuation of s, $s.a$, such that $s.a \in h^{m+n}_{B_n}$ but $\mu^m_M(s.a|h_B) = \mu^m_M(s|h_B) - \delta < \alpha$. So there is no convergence at any finite stage n of $\mathcal{A}$ towards h_B. Non uniform learnability of $\mathcal{H}$ then follows.

(ii.b) Let's now assume that M only has Max Ent priors and learns by inductive inference. Uniform convergence of any algorithm obeying these conditions is not guaranteed as a similar argument as in Proposition 3 applies. □

Corollary 2. The hypothesis space $\mathcal{H}_B$ is not uniformly learnable. Hence the the VC-dimensions of $\mathcal{H}_B$ are not finite, and empirical risk minimization on $\mathcal{H}_B$ are not uniformly consistent.

Corollary 3. M cannot effectively learn Σ^0_1 complete Borel sets.

Assume M can effectively learn a Σ^0_1 complete set. Then it can effectively learn a Π^0_1 set that is its complement, which is impossible by Theorem 2.□

Proposition 5. An LLM M cannot effectively learn Borel sets B of higher complexity than Δ^0_1.

Proposition 2 and Corollary 2 show that M cannot effectively learn Π^0_1 or Σ^0_1 sets. But any Π^0_n or Σ^0_n complete Borel set B for n > 1 is at least a countable intersection or countable union of such sets. So B is not effectively learnable.□

Asher et al. (2017); Asher and Paul (2018) examine concepts of discourse consistency and textual and conversational coherence, which true, human-like conversational capacity requires. Using continuations in a game-theoretic setting, they show those concepts determine more complex Π^0_2 sets in the Borel Hierarchy; and intuitive measures of conversational success—like the fact that one player has more successful unrefuted attacks on an opponent's position than vice versa—determine Π^0_3 sets. Given Proposition 5, LLMs cannot learn these concepts, which are needed for full conversational mastery.

Proposition 6. For any LLM M, there is a maximally large and fixed number n such that $\mathcal{H}^n$ is uniformly learnable for M but $\mathcal{H}^{n+k}$ is not uniformly learnable, for $k > 0$.

Suppose that for M $\mathcal{H}^n$ is uniformly learnable for all n. Then, M can compute the countable intersection of sets defined by the best hypotheses in $\mathcal{H}^n$ for each n. So M can effectively learn a Π^0_1 set, which contradicts Theorem 2.□

Corollary 4. M cannot effectively learn Δ^0_1 sets of the form $\mathcal{O}(A)$ if the length of A is longer than the maximal number n such that $\mathcal{H}^n$ is uniformly learnable for M.

4.3 The importance of order

Order is important for the most elementary reasoning about linguistic content in finite domains. Let us add another predicate A to $\mathcal{L}$ to form the language $\mathcal{L}^+$. Now consider the strings in $V^\omega_{\mathcal{L}^+}$. Strings consistent with (2) may include formulae like $A(a_i)$ or $\neg A(a_i)$, paired with a choice of $blue(a_i)$ or $\neg blue(a_i)$. Even to find effectively initial segments of strings in $h^\omega_{\forall_{\mathcal{L}^+}}$, M must learn some sentence structure or word order. The negation sign has to be paired with the predicate *blue*; if it's appended to A (e.g., *large*, or some other independent term), this should count as a string in $h^n_\forall$. If s is a finite string, M does not effectively capture word order if it does not distinguish between s and permutations of elements in s.

Proposition 7. If M does not effectively capture word order, it cannot effectively learn basic sets of the form $\mathcal{O}(A)$ for $A \subset V^*$.

Let $s \in A$ be a string containing $A(a_i) \wedge \neg blue(a_i)$ but A has no string containing $\neg A(a_i) \wedge blue(a_i)$. If M does not capture word order, M cannot distinguish between s and s's permutation containing $\neg A(a_i) \wedge blue(a_i)$. □

Corollary 5. If M does not effectively capture word order, it will not reason soundly in propositional logic.

The example in Proposition 7 shows that M will not be able to reason about logical structure if it does not effectively capture word order. □

Yuksekgonul et al. (2022); Sinha et al. (2020) provide evidence that small to moderate sized LLMs do not reliably capture word order. Our empirical examples show even GPT3.5 and ChatGPT have difficulties with sentential word order, and,

worryingly, with the order of arguments in a logical operator; the example in Appendix B suggests that even ChatGPT can't be trusted to always do elementary inferences involving conditionals correctly. Thus, LLMs with their initial training do not necessarily find basic Δ_1^0 sets of the form $a.V^\omega$ where a is a single string but only sets $A.V^\omega$ where A is a set of prefixes that are permutations on a. This is surprising and poses extreme difficulties for valid reasoning with operators that have order dependent arguments.

5 Empirical investigations of LLMs with *every*

While the theoretical argument laid out in Section 4 does not hinge on empirical statistics of LLM failures, it certainly suggests that we should expect such failures. In this section, we describe some of the tests we have performed using continuations to query LLMs directly about their mastery of universal quantification.

Let us return to our simple example from above, repeated here as (2):

(2) Everything is blue.

We used finite sequences of formulas as a context, e.g., a_1 *is blue*, a_2 *is red*, a_3 *is red*,..., a_i *is blue* to determine a model. We then asked an LLM M whether (2) in this model, allowing us to gauge its behavior with respect to finite domains.

BERT-large and RoBERTa-large already failed to reliably distinguish very small models (containing 2 and 5 objects respectively) in which (2) is true from those in which it is not. To test these models, we fine-tuned BERT-large and RoBERTa-large on the CoQA dataset (Reddy et al., 2019). For finetuning, the model had 4 output heads for *yes*, *no*, *unknown*, and *span* type questions. Since the CoQA dataset provides a rationale for each question, the models were jointly trained on question answering and rationale tagging tasks to enhance their performance. We report scores on the finetuned models on CoQA for 1 epoch as we did not observe significant improvement with an increased number of epochs.

For BERT-large, we provided strings like (3) and then asked *Is everything blue?*

(3) My car is blue. My house is blue

We generated a total of 5 examples in which (2) was true and 5 examples in which (2) is false. All the examples had only 2 objects. The inconsistent examples were constructed by varying the position of the object which was inconsistent with the asked question and by trying out different combination of colours and objects.

Object Count	Pass Fraction
2	1/1
3	2/3
4	1/4
5	0/5
6	0/6
7	0/7
8	0/8
9	0/9
10	0/10

Table 1: Pass fraction on inconsistent examples for RoBERTa-large

The consistent examples were of the form:
1. The car is blue. The house is blue.
2. The car is purple. The house is purple.
3. The car is yellow. The house is yellow.
3. The shirt is violet. The table is violet.
4. The cup is black. The plate is black.

Inconsistent examples were of the form:
1. The car is blue. The house is red.
2. The car is green. The house is purple.
3. The car is yellow. The house is brown.
2. The shirt is violet. The table is brown.
3. The cup is black. The plate is white.

BERT-large was able to correctly identify the consistent examples but failed for all the inconsistent examples. As the model failed for all the inconsistent examples with 2 objects, we did not experiment with models containing more than 2 objects.

For RoBERTa-large, we generated a total of 9 consistent examples and 53 inconsistent examples. We constructed sequences ranging from 2 to 10 objects. For each number, the inconsistent examples were constructed by varying the position of the object in the string (context) which is responsible for the inconsistency. The model was able to correctly identify all the consistent examples. For models of a given size (i.e., number of objects), we defined the *pass fraction* as the ratio of the examples in which the model was able to report models inconsistent with (2) correctly to the total number of inconsistent examples. Table 1 reports the pass fraction on inconsistent examples.

While BERT's and RoBERTa's behavior was stable on the strings tested, GPT3.5 davinci and ChatGPT, while more robust, are unstable from one day to the next, even when temperature is set to 0 (on GPT3.5). This made it difficult to pin down the models' abilities, though some generalizations emerged. Typically (though not always), these models can recognize which objects in a string have a certain property, but they cannot necessarily exploit this information to answer questions about the string as a whole (see the "hats" example in Appendix A). In addition both GPT3.5 and ChatGPT will sometimes (frequently in our most recent tests) over-generalize and say that all items in a list are, say, blue if it is specified for all items but one that they are blue and it is not specified one way or the other for the remaining item (see the *fifteen hearts* example from ChatGPT in Appendix A). Thus, even these sophisticated models still fail on more complicated questions and longer strings.

Our empirical observations on LLMs like BERT and RoBERTa and probing of ChatGPT strongly support our argument that LLMs are unable to master quantification, complementing observed LLM difficulties with negation (Naik et al., 2018; Kassner and Schütze, 2019; Hossain et al., 2020; Hosseini et al., 2021) and to some extent quantifiers (Kalouli et al., 2022).

6 Conclusions

We have shown that LLMs' demonstrably inadequate grasp of the meanings of words like *every* and other linguistic constructions has a theoretical foundation and explanation: for certain expressions S, S's content should be defined via consistent sets of strings in V^ω, and LLMs cannot effectively learn certain sets in V^ω. More generally, LLMs cannot effectively learn full meanings of first order quantifiers or any Borel sets beyond the basic open sets, which means that they fail to grasp the meaning of a long list of mundane, frequently used expressions.

Many of these expressions are syncategorematic terms and express what we might call *precise concepts*. Such concepts are needed for understanding ordinary entailment across all expressions; in addition, correctly reasoning with these concepts and grasping their entailments is essential to understanding them. Reasoning and entailment are intimately tied with meanings. For us and most formal semanticists (Montague, 1974), grasping meaning and correctly reasoning with linguistically expressed concepts go hand in hand; if you cannot exploit the meanings of words in correct reasoning, you do not really know what they mean. The incorrect reasoning of LLMs exemplifies their failure to grasp semantic entailments and meaning.

Our arguments go beyond those of Bender and Koller (2020), who argue that stochastic models cannot capture linguistic meaning because they consider only form, not denotation. While we agree that denotation plays a very important role in meaning for many expressions, the meaning of most expressions, and especially that of syncategoregmatic ones, requires us to capture their semantic entailments. We have shown that we can capture these entailments within the semantic framework of LLMs using continuation semantics. But we have also shown that LLMs nevertheless fail in this task.

LLMs *can* learn certain types of Δ^0_1 sets and finite intersections and unions of learnable Δ^0_1 sets. For many open class words—including many nouns, adjectives and verbs—whose characteristic denotations can be determined given a finite sample, this probably suffices to capture their meaning or at least a very good approximation of it. In addition, many NLP tasks may not involve logical inference but an independent form of string optimization; in text summarization or translation, where given a context s, M tries to find an optimal continuation s'. If the length of $s.s'$ falls within the constraints of Corollary 4, then we can expect an LLM to succeed at such a task.

Proposition 6 and Corollary 4 generalize Corollary 1 and they all point to a general limit on learnability for LLMs. They establish that language models have strict bounds even on the Δ^0_1 sets they can effectively learn. So we cannot count on LLMs having full linguistic competence even on finite domains. Different models may have different limits; smaller models generally with lower limits. This motivates a comparative study of the limits of learnability for different LLMs, complementing Colbrook et al. (2022).

Because we do not make assumptions about memory but only about inductive processes and learning, our results hold for arbitrarily large LLMs and for any task that relies on an LLM's capacity of string prediction, even if strings are not

directly predicted.

Our research implies that full language mastery needs a different approach from one in which one seeks to build ever larger LLMs with language masking or autoregressive training. Following Raissi et al. (2017), we believe we need to inject knowledge about linguistic structure and content into our models to further constrain learning and in particular hypothesis sets and learning architectures as suggested in Steinert-Threlkeld and Szymanik (2019).

Limitations

Our paper offers an in-depth look at the limits of statistical learnability for LLMs in terms of Borel sets. Our results however rely on certain, rather mild hypotheses about learning. We do not know how our claims fare when different assumptions about learning are made.

Another limitation is that we have not looked at potential ways of restricting learning hypotheses about the meanings of linguistic expressions. Directly relevant are the constraints on generalized quantifiers of Barwise and Cooper (1981). Nor have we said much about compositionality and learning, although they are well known to interact. Ideally, we should assign a meaning to *every* which consists in a reduced hypothesis space, but that has to combine with the meanings of, say, the other words in (2).

A third limitation concerns the extent to which an LLM can leverage what it has been trained to do in learning semantic concepts on other tasks. In many experiments we see that an LLM can appear to learn X but then cannot exploit X in task Y. We leave this issue for future research, but we believe it may be key to understanding debates about LLM performance and the role of forgetting in LLM learning (Kirkpatrick et al., 2017).

Ethics Statement

We show that LLM reliability in gauging the accuracy of universal quantifiers is limited. Given how quantification is linked to concepts like semantic consequence, it means that we cannot assume LLMs to have anything like full semantic competence. This suggests that we cannot fully trust them with tasks where accuracy on reasoning is required. Thus, our work has certain ethical implications and provides a cautionary tale for those working on LLMs as general AI models.

Acknowledgments

For financial support, we thank the National Interdisciplinary Artificial Intelligence Institute ANITI (Artificial and Natural Intelligence Toulouse Institute), funded by the French 'Investing for the Future– PIA3' program under the Grant agreement ANR-19-PI3A-000, and we also thank the projects COCOBOTS (ANR-21-FAI2-0005) and DISCUTER (ANR-21-ASIA-0005). We also thank the COCOPIL "Graine" project of the Région Occitanie of France.

References

Anthony, M., Bartlett, P. L., Bartlett, P. L., et al. (1999). *Neural network learning: Theoretical foundations*, volume 9. cambridge university press Cambridge.

Asher, N. (1993). *Reference to Abstract Objects in Discourse*. Kluwer Academic Publishers.

Asher, N. and Paul, S. (2018). Strategic conversation under imperfect information: epistemic Message Exchange games. *Logic, Language and Information*, 27.4:343–385.

Asher, N., Paul, S., and Venant, A. (2017). Message exchange games in strategic conversations. *Journal of Philosophical Logic*, 46.4:355–404.

Asher, N. and Pogodalla, S. (2011). Sdrt and continuation semantics. In Onada, T., Bekki, D., and McCready, E., editors, *New Frontiers in Artificial Intelligence: JSAI-isAI 2010 Workshops, LENLS, JURISIN, AMBN, ISS, Tokyo, Japan, November 18-19, 2010, Revised Selected Papers*, pages 3–15. Springer Berlin Heidelberg, Berlin, Heidelberg.

Barwise, J. and Cooper, R. (1981). Generalized quantifiers in natural language. *Linguistics and Philosophy*, 4(1):159–219.

Bender, E. M. and Koller, A. (2020). Climbing towards NLU: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online. Association for Computational Linguistics.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Chang, C. C. and Keisler, H. J. (1973). *Model theory*. north Holland.

Chaturvedi, A., Bhar, S., Saha, S., Garain, U., and Asher, N. (2022). Analyzing semantic faithfulness of language models via input intervention on conversational question answering. *Computing Research Repository*, arXiv:2212.10696.

Colbrook, M. J., Antun, V., and Hansen, A. C. (2022). The difficulty of computing stable and accurate neural networks: On the barriers of deep learning and smale's 18th problem. *Proceedings of the National Academy of Sciences*, 119(12):e2107151119.

De Groote, P. (2006). Towards a montagovian account of dynamics. In *Semantics and linguistic theory*, volume 16, pages 1–16.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). Bert: Pre-training of deep bidirectional transformers for language understanding.

Dowty, D. R., Wall, R., and Peters, S. (1981). *Introduction to Montague semantics*. Dordrecht. Synthese Library vol. 11.

Fernando, T. (2004). A finite-state approach to events in natural language semantics. *Journal of Logic and Computation*, 14(1):79–92.

Fernando, T. (2022). Strings from neurons to language. In *Proceedings of the 3rd Natural Logic Meets Machine Learning Workshop (NALOMA III)*, pages 1–10.

Graf, T. (2019). A subregular bound on the complexity of lexical quantifiers. In *Proceedings of the 22nd Amsterdam colloquium*.

Hornik, K., Stinchcombe, M., and White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366.

Hossain, M. M., Kovatchev, V., Dutta, P., Kao, T., Wei, E., and Blanco, E. (2020). An analysis of natural language inference benchmarks through the lens of negation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9106–9118.

Hosseini, A., Reddy, S., Bahdanau, D., Hjelm, R. D., Sordoni, A., and Courville, A. (2021). Understanding by understanding not: Modeling negation in language models. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1301–1312.

Kalouli, A.-L., Sevastjanova, R., Beck, C., and Romero, M. (2022). Negation, coordination, and quantifiers in contextualized language models. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3074–3085.

Kamp, H. and Reyle, U. (1993). *From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Kluwer Academic Publishers.

Kassner, N. and Schütze, H. (2019). Negated and misprimed probes for pretrained language models: Birds can talk, but cannot fly. *arXiv preprint arXiv:1911.03343*.

Kawaguchi, K., Kaelbling, L. P., and Bengio, Y. (2017). Generalization in deep learning. *arXiv preprint arXiv:1710.05468*.

Kechris, A. (1995). *Classical descriptive set theory*. Springer-Verlag, New York.

Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526.

Kuhnle, A. and Copestake, A. (2019). The meaning of "most" for visual question answering models. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 46–55.

Lamport, L. (1980). Sometime is sometimes not never: On the temporal logic of programs. In *Proceedings of the 7th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 174–185. ACM.

Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. (2019). Roberta: A robustly optimized bert pretraining approach.

Mahowald, K., Ivanova, A. A., Blank, I. A., Kanwisher, N., Tenenbaum, J. B., and Fedorenko, E. (2023). Dissociating language and thought in large language models: a cognitive perspective. arXiv preprint arXiv:2301.06627.

Mialon, G., Dessì, R., Lomeli, M., Nalmpantis, C., Pasunuru, R., Raileanu, R., Rozière, B., Schick, T., Dwivedi-Yu, J., Celikyilmaz, A., et al. (2023). Augmented language models: a survey. *arXiv preprint arXiv:2302.07842*.

Montague, R. (1974). *Formal Philosophy*. Yale University Press, New Haven.

Muggleton, S., Feng, C., et al. (1992). Efficient induction of logic programs. *Inductive logic programming*, 38:281–298.

Naik, A., Ravichander, A., Sadeh, N., Rose, C., and Neubig, G. (2018). Stress test evaluation for natural language inference. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2340–2353.

Neyshabur, B., Bhojanapalli, S., McAllester, D., and Srebro, N. (2017). Exploring generalization in deep learning. *Advances in neural information processing systems*, 30.

Plotkin, G. (1972). *Automatic methods of inductive inference*. PhD thesis, The University of Edinburgh.

Popper, K. (1963). *Conjectures and refutations: The growth of scientific knowledge*. routledge.

Raissi, M., Perdikaris, P., and Karniadakis, G. E. (2017). Physics informed deep learning (part i): Data-driven solutions of nonlinear partial differential equations. *arXiv preprint arXiv:1711.10561*.

Reddy, S., Chen, D., and Manning, C. D. (2019). CoQA: A Conversational Question Answering Challenge. *Transactions of the Association for Computational Linguistics*, 7:249–266.

Reynolds, J. C. (1974). On the relation between direct and continuation semantics. In *International Colloquium on Automata, Languages and Programming*.

Shalev-Shwartz, S., Shamir, O., Srebro, N., and Sridharan, K. (2010). Learnability, stability and uniform convergence. *The Journal of Machine Learning Research*, 11:2635–2670.

Siegelmann, H. T. (2012). *Neural networks and analog computation: beyond the Turing limit*. Springer Science & Business Media.

Siegelmann, H. T. and Sontag, E. D. (1992). On the computational power of neural nets. In *Proceedings of the fifth annual workshop on Computational learning theory*, pages 440–449.

Sinha, K., Parthasarathi, P., Pineau, J., and Williams, A. (2020). Unnatural language inference. *arXiv preprint arXiv:2101.00010*.

Steinert-Threlkeld, S. and Szymanik, J. (2019). Learnability and semantic universals. *Semantics and Pragmatics*, 12(4).

Tarski, A. (1944). The semantic conception of truth: and the foundations of semantics. *Philosophy and phenomenological research*, 4(3):341–376.

Tarski, A. (1956). The concept of truth in formalized languages. In translated by J.H. Woodger, editor, *Logic, Semantics and Metamathematics*, pages 152–278. Oxford University Press, New York.

Tenney, I., Xia, P., Chen, B., Wang, A., Poliak, A., McCoy, R. T., Kim, N., Van Durme, B., Bowman, S. R., Das, D., et al. (2018). What do you learn from context? probing for sentence structure in contextualized word representations. In *International Conference on Learning Representations*.

Villa, S., Rosasco, L., and Poggio, T. (2013). On learnability, complexity and stability. In *Empirical Inference*, pages 59–69. Springer.

Weiss, G., Goldberg, Y., and Yahav, E. (2018). On the practical computational power of finite precision rnns for language recognition. *arXiv preprint arXiv:1805.04908*.

Wolpert, D. H., Macready, W. G., et al. (1995). No free lunch theorems for search. Technical report, Technical Report SFI-TR-95-02-010, Santa Fe Institute.

Yuksekgonul, M., Bianchi, F., Kalluri, P., Jurafsky, D., and Zou, J. (2022). When and why vision-language models behave like bag-of-words models, and what to do about it? *arXiv preprint arXiv:2210.01936*.

A GPT 3.5 and ChatGPT on *every*

Because our access to GPT 3.5 and ChatGPT is limited and the responses that we receive seem to vary depending on the day that we query the models, we qualify the results in these appendices as anecdotal. Nevertheless, we emphasize that for

each type of example that we provide, we made numerous trials in order to make sure that the results were repeated and so seemed to follow from features of the model, not just random quirks.

In the queries below, GPT3.5 on 18.01.2023, with temperature set to 0, correctly classifies both blue hats as being blue when queried directly about the color of the beret, but still fails to correctly answer *Were all of the hats blue?* (Results sometimes varied with the form of the prompt; e.g., using commas rather than *and* can impact GPT's answers).

Input Text: One guy was wearing a blue beret and another guy was wearing a pink bonnet and a third guy was wearing a blue baseball cap.

Q: Were all of the hats blue?
GPT: No, only one of the hats was blue.
Q: Which hat was blue? GPT: The blue baseball cap.
Q: Is a beret a hat? GPT: Yes, a beret is a type of hat.
Q: What color is the beret? GPT: The beret is blue.

Here is an input and query to ChatGPT on 20.01.2023, showing that it ignores underspecification

Input text: I have a shirt with a picture of ten blue bears, five blue snails, three blue guitars, fifteen hearts, 15 blue fluffy rabbits and 25 blue stars. Is everything on my shirt blue?

ChatGPT: Yes, everything on your shirt is blue.

B ChatGPT and GPT3.5 on word order and argument order

The examples below illustrate how both GPT3.5 and ChatGPT fail to recognize instructions about word order. When asked to write sentences that end with a particular word, for example, the models manage to produce sentences in which the target word appears, but in most cases, it is not in the correct position.

Instruction: Write five sentences that end with the word "house".

ChatGPT:
1. I built a house with my own two hands.
2. The storm destroyed the old house.
3. We found a beautiful house to rent for the summer.
4. The house was haunted, so we left in the middle of the night.
5. I always dreamed of living in a big white house.

GPT 3.5:
1. I can't wait to get home and relax in my own house.
2. The old house had been abandoned for years.
3. The house was filled with the smell of freshly baked cookies.
4. The house was surrounded by a tall white picket fence.
5. The house was the perfect place to spend a lazy summer day.

For GPT3.5 we set temperature to 0. These observations suggest both GPT 3.5 and ChatGPT fail to respect the *last word* constraint.

Finally, we offer this example that shows that ChatGPT can confuse the order of arguments in a logical operator, thus leading to incorrect reasoning.

Prompt: If p then q. q, therefore p.

ChatGPT:
This is a valid logical argument. It is known as modus ponens and it is a valid rule of inference in propositional logic. The statement "If p then q" is known as the conditional statement and "q, therefore p" is the conclusion that can be drawn from the conditional statement.