

Entailment Semantics Can Be Extracted from an Ideal Language Model

William Merrill

New York University

{willm, linzen}@nyu.edu

Alex Warstadt

ETH Zürich

alexanderscott.warstadt@inf.ethz.ch

Tal Linzen

New York University

linzen@nyu.edu

Abstract

Language models are often trained on text alone, without additional grounding. There is debate as to how much of natural language semantics can be inferred from such a procedure. We prove that entailment judgments between sentences can be extracted from an ideal language model that has perfectly learned its target distribution, assuming the training sentences are generated by Gricean agents, i.e., agents who follow fundamental principles of communication from the linguistic theory of pragmatics. We also show entailment judgments can be decoded from the predictions of a language model trained on such Gricean data. Our results reveal a pathway for understanding the semantic information encoded in unlabeled linguistic data and a potential framework for extracting semantics from language models.

1 Introduction

Recent advances in building computational models of language have been powered by *distributional semantics*: the idea that a text span’s surrounding context encodes its meaning (Firth, 1957). In particular, large pretrained language models (LMs; Peters et al., 2018; Devlin et al., 2019; Brown et al., 2020) have become an integral part of NLP systems: the representations that emerge from training to predict missing words in a text are empirically useful for natural language understanding tasks.

Despite this empirical progress, Bender and Koller (2020) argue LMs cannot learn to understand the semantics of sentences. This is because of a mismatch between the LM training objective—predicting missing words in text (“form”)—and Bender and Koller’s conception of meaning as the relation of a sentence to the external world. Thus Bender and Koller claim “that the language modeling task, because it only uses form as training data, cannot in principle lead to learning of meaning.”

In this paper, we argue meaning *can* be learned from form because the communicative goals of hu-

man authors encode semantic information in unlabeled text. We show how this semantic information can be extracted to resolve semantic relations between sentences (e.g., whether one sentence entails another): in this inferentialist sense, ideal LMs encode the meaning of sentences. This argument has been raised speculatively by others (Michael, 2020; Potts, 2020; Bommasani et al., 2021), but we will rigorously justify it here with formal results.

To give the simplest (and least general) illustration of our argument, we first assume training data is generated by overly idealized *uniformly truthful* speakers: agents who decide what to say by picking sentences they consider true uniformly at random.¹ This very coarsely captures human authors’ goal of being informative (rather than misleading) to their listeners (Grice, 1975). In Theorem 1, we prove a sentence x entails sentence y if and only if, after uttering x , a uniformly truthful speaker is just as likely to say y as to repeat x . Thus, entailment semantics can be extracted from probabilistic languages generated by uniformly truthful speakers.

Uniformly truthful speakers are not a realistic model of humans: while humans favor true sentences to false ones (Grice, 1975), not all true sentences are equally likely to be produced. It is a common principle in linguistic theories of pragmatics that human speakers choose their utterances in order to balance two competing objectives: (a) conveying information to their listener and (b) brevity (Levinson et al., 1983; Grice, 1975). We define a class of *Gricean* speakers who optimize for these objectives, and prove in Theorem 2 that x entails y if and only if a simple equation holds in terms of text probabilities produced by such speakers. Thus, entailment semantics can be decoded from probabilistic languages generated by Gricean speakers.

¹ Studying the ability of LMs to understand programming language semantics, Merrill et al. (2021) make a similar assumption that programmers are more likely to write true assertion statements than false ones.

The previous results assume access to a language’s ideal likelihood function, but, in practice, one only ever receives a corpus *sampled* from the language. Moving to the corpus setting, we analyze how much data allows approximately computing our derived entailment test using probabilities estimated from sentence frequencies in a corpus. We find that the corpus size needed to guarantee the entailment test holds approximately is inversely related to the likelihood of the sentences. We estimate that approximating the entailment test between 4-word sentences using corpus frequencies is possible with $\sim 10^{10}$ sentences, about the size of the GPT-3 training data (Brown et al., 2020). On the other hand, approximating the entailment test for 10-word sentences should be possible with $\sim 10^{17}$ sentences, or $\sim 10^7$ GPT-3 corpora. Thus, extracting entailment judgments using corpus frequencies requires an infeasible amount of data—even by modern NLP standards.

To overcome this limitation, one might hope to use probabilities estimated by LMs to extract entailment judgments between longer sentences that are rare even in a large corpus. With synthetic data generated by Gricean speakers, we find that entailment can be decoded from n-gram LM predictions to some extent. However, we speculate that current neural LMs may not score the probability of rare text well enough to enable decoding entailment judgments between natural language sentences.

In summary, our main contribution is to show a correspondence between the semantics of text and its likelihood, assuming the likelihood function matches models of human text production from linguistic theory. Determining whether a sentence in a probabilistic language entails another sentence can be reduced to modeling the probabilities of strings in the language. In practice, entailment judgments between very short sentences can be extracted from corpus frequencies, but this becomes infeasible for slightly longer sentences. LMs can in principle be used to extrapolate the likelihood of longer strings, but we hypothesize current LMs are not well-suited for doing so well enough to enable extracting entailment from natural language. Our theory demonstrates a formal sense in which unlabeled text data encodes linguistic meaning and makes quantitative predictions for (a) how to extract semantics from text corpora and (b) how much data this requires.

2 Definitions

2.1 Sentences and Worlds

Let \mathcal{X} be a finite set of sentences, and \mathcal{W} a countable² set of possible world states. A sentence x is a string whose denotation $\llbracket x \rrbracket$ is a *proposition*, i.e., a set of world states ($\subseteq \mathcal{W}$) where x is true. Following standard conventions in formal semantics (cf. Heim and Kratzer, 1998), the set $\llbracket x \rrbracket$ can be equivalently viewed as a function mapping a world state w to $\{0, 1\}$ that indicates whether x is true in w , which we will write as $\llbracket x \rrbracket(w)$. We imagine w to encode a partial description of the world, much like the concept of a *situation* in formal semantics (Kratzer, 2021). For simplicity, we assume an individual’s subjective belief state can be modeled as the unique, maximal w that fully describes the facts which they believe to be true.

Example $x = \text{John has at least two cats.}$ Let $\mathcal{W} = \{w_0, w_1, w_2, w_3\}$ be the set of possible worlds, where w_n denotes the state in which John has n cats. Then $\llbracket x \rrbracket = \{w_2, w_3\}$, because John has at least two cats in these worlds. Furthermore, it holds that $\llbracket x \rrbracket(w_2) = 1$, but $\llbracket x \rrbracket(w_1) = 0$.

2.2 Speakers and Texts

We refer to a sequence of sentences $z \in \mathcal{X}^*$ as a *text*.³ The meaning of a text is the set of worlds consistent with all its sentences, i.e.,

$$\llbracket z \rrbracket = \bigcap_{t=1}^{|z|} \llbracket z_t \rrbracket.$$

We will imagine that a text $z \in \mathcal{X}^*$ is produced by iteratively sampling $z_t \in \mathcal{X} \cup \{\$\}$ from a *speaker model* $p(z_t \mid z_{<t}, w)$. $p(z_t \mid z_{<t}, w)$ represents the probability of saying sentence z_t with belief state w after having said $z_1 \cdots z_{t-1}$. Let $\$ \notin \mathcal{X}$ be a special *end of sequence* token satisfying $\llbracket \$ \rrbracket = \mathcal{W}$. We refer to any text ending with $\$$ as *complete*. Given a world w , an incomplete text $z \in \mathcal{X}^*$ or complete text $z \in \mathcal{X}^*\$$ has *conditional probability*

$$p(z \mid w) = \prod_{t=1}^{|z|} p(z_t \mid z_{<t}, w).$$

The conditional probability of an incomplete text represents the probability of observing z as the

²Our results extend to uncountable sets of world states if entailment is relaxed to hold almost surely (cf. §B). Alternatively, our results apply as-is if we assume a countable set of equivalence classes over uncountably many worlds.

³Where \mathcal{X}^* denotes the Kleene star closure of \mathcal{X} .

prefix of a text written by a human with beliefs w . In contrast, the probability of a complete text represents the probability that a speaker produces z and no further text. The conditional distribution $p(z \mid w)$ cannot be observed directly by a LM, since w is a latent variable missing from the training data. Rather, a LM has access to texts that have been generated by speakers across many possible belief states. Mathematically, this can be expressed by saying a LM’s target distribution is a *marginal* distribution over $z \in \mathcal{X}^* \cup \mathcal{X}^*\$$ according to some *prior* distribution over worlds $p(w)$:

$$\begin{aligned} p(z) &= \mathbb{E}_{w \sim p(w)} [p(z \mid w)] \\ &= \mathbb{E}_{w \sim p(w)} \left[\prod_{t=1}^{\infty} p(z_t \mid z_{<t}, w) \right]. \end{aligned}$$

The prior $p(w)$ represents the probability that a speaker contributing to the corpus will have belief state w —we make no assumptions about its form besides that $p(w) > 0$ for all $w \in \mathcal{W}$, and, for every sentence, there is some world state that makes that sentence true. In contrast to $p(z)$, which corresponds to the expected corpus frequency of z , we denote by $p(\llbracket z \rrbracket)$ the probability that z is true.⁴

Example Let z be the 2-sentence text:⁵

$z_1 =$ We swung our swords.
 $z_2 =$ That was ever so long ago.

Let p be the distribution of all possible English web texts. The marginal probability $p(z)$ can be decomposed across many possible worlds. One such world w_1 might be the world where the speaker is the semi-legendary Viking hero Ragnar Loðbrók (in modern English translation); another world w_2 might be the perspective of a Reddit user reviewing a coffee maker. Each of these worlds corresponds to one term in a sum over all worlds. We expect $p(z \mid w_1)$ to be higher than $p(z \mid w_2)$ since it is more likely for a medieval literary character to utter z than a modern product reviewer. Finally, $p(z \mid w_1)$ can be factored as

$$p(z_1 \mid w_1)p(z_2 \mid z_1, w_1).$$

In contrast to $p(z)$, which counts all contexts where z is the beginning of a longer text, $p(z\$)$ measures the frequency of z_1z_2 followed by nothing else.

⁴The notation explicitly represents the probability mass assigned to the set of worlds where z is true.

⁵Text taken from the Wikipedia page for the skaldic poem *Krákumál*, written in Ragnar’s voice.

2.3 Distributional and Semantic Relations

Distributional Relations A *distributional relation* d is a relation over sentences x and y defined in terms of likelihood of different texts under some distribution p . Let $d_p(x, y)$ be the value of the distributional relation d between sentences x, y according to distribution p . If we train an LM on texts sampled from a target distribution p , the LM estimates a predictive distribution \hat{p} . Thus, any LM parameterizes $d_{\hat{p}}$: an instantiation of the distributional relation d with respect to the probabilities learned by the LM. If the LM perfectly approximates $p(x)$ for all x , then $d_{\hat{p}} = d_p$ by construction.

Example Define the distributional relation d (with respect to some distribution p) such that $d_p^>(x, y) \iff p(x) > p(y)$. $d_p^>(x, y)$ says x is more likely than y according to p . If \hat{p} represents LM predictions trained on the target distribution p , then $d_{\hat{p}}^>(x, y)$ says whether the LM predicts a sentence x is more likely than another sentence y .

Semantic Relations In contrast, a semantic relation between x and y is a relation defined in terms of their denotations $\llbracket x \rrbracket$ and $\llbracket y \rrbracket$. We will focus on the key semantic relation of entailment:

Definition 1 For two sentences $x, y \in \mathcal{X}$, x entails y if and only if $\llbracket x \rrbracket \subseteq \llbracket y \rrbracket$.

It is not clear *prima facie* if LMs can represent entailment relations. However, it could be that a semantic relation s can somehow equivalently be written as a distributional relation d_p . If so, a LM that perfectly approximates p could be understood to encode s , since s can be extracted from \hat{p} via $d_{\hat{p}}$.

Formally, we can ask if a semantic relation can be alternatively expressed as a distributional relation by analyzing if there exists an isomorphism between a semantic relation $s(\llbracket x \rrbracket, \llbracket y \rrbracket)$ and some distributional relation $d_p(x, y)$:

Definition 2 (Isomorphism) A semantic relation s is isomorphic to a distributional relation d under speaker p if and only if, for all $x, y \in \mathcal{X}$,

$$s(\llbracket x \rrbracket, \llbracket y \rrbracket) \iff d_p(x, y).$$

If Definition 2 holds under a speaker model p , then predicting whether s holds between two sentences is reducible to perfectly modeling the probabilities of texts generated by p . Our goal going forward will be to derive distributional relations isomorphic to entailment assuming p models the goals of humans when they produce text.

3 Uniformly Truthful Speakers

We start by illustrating our research question and technical approach assuming an overly simple model of humans as *uniformly truthful* speakers. A uniformly truthful speaker chooses a sentence to produce by selecting one of the true sentences that holds in their belief state uniformly at random. This very coarsely captures the property of natural language pragmatics that subjectively true sentences tend to be more likely than false ones, although it does not account for many other factors that influence human speech patterns in complex ways (Grice, 1975).⁶ Let $n(w)$ be the number of *sentences* true in world w . We can formally define a uniformly truthful speaker as follows:

Definition 3 A speaker p is *uniformly truthful* if, for all sentences $x \in \mathcal{X} \cup \{\$\}$,

$$p(x | w) = \frac{\llbracket x \rrbracket(w)}{\sum_{x'} \llbracket x' \rrbracket(w)} = \frac{\llbracket x \rrbracket(w)}{n(w)}.$$

In other words, p uniformly spreads probability mass across all sentences that are true in world w . We will show that, if the corpus consists of text written by uniformly truthful speakers, entailment can be decided by a distributional relation. The following lemma will be a core technical tool in our analysis. Informally, it is useful because it establishes a correspondence between relations over sets of worlds and probabilities.

Lemma 1 Let $\mathbb{1}_{\mathcal{S}}$ be the indicator function for set \mathcal{S} . For sets \mathcal{A}, \mathcal{B} such that $\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{W}$, and $c : \mathcal{W} \rightarrow \mathbb{R}_+$, $\mathcal{A} = \mathcal{B}$ if and only if

$$\sum_{w \in \mathcal{W}} \mathbb{1}_{\mathcal{A}}(w)c(w) = \sum_{w \in \mathcal{W}} \mathbb{1}_{\mathcal{B}}(w)c(w).$$

Proof. We will prove that $\mathcal{B} \subseteq \mathcal{A}$ by contradiction. Assume there exists $w \in \mathcal{B}$ such that $w \notin \mathcal{A}$. Then the right sum contains the positive term $c(w)$, while the left sum does not. Because all terms in the right sum are positive, the left sum must contain at least one term $c(w')$ that the right sum does not. Thus, $w' \in \mathcal{A}$ but $w' \notin \mathcal{B}$. But this has violated our assumption that $\mathcal{A} \subseteq \mathcal{B}$. \square

We now use Lemma 1 to derive a simple distributional relation that is isomorphic to entailment.

⁶LMs sometimes generate objectively false statements (Lin et al., 2022), presumably due to the occurrence of such facts in their training data. This is actually consistent with a uniform truthfulness assumption, which only requires that speakers only produce sentences they believe are true, not sentences that are actually true in some objective sense.

Theorem 1 If p is a uniformly truthful speaker, then entailment is isomorphic to a distributional relation. Specifically, for all sentences $x, y \in \mathcal{X}$,

$$\llbracket x \rrbracket \subseteq \llbracket y \rrbracket \iff p(xy) = p(xx).$$

Proof. $d_p(x, y)$ holds if and only if

$$\begin{aligned} p(xy) &= p(xx) \\ \mathbb{E}_w \left[\frac{\llbracket x \rrbracket(w) \llbracket y \rrbracket(w)}{n(w)^2} \right] &= \mathbb{E}_w \left[\frac{\llbracket x \rrbracket(w) \llbracket x \rrbracket(w)}{n(w)^2} \right] \\ \mathbb{E}_w \left[\frac{\llbracket x \rrbracket(w) \llbracket y \rrbracket(w)}{n(w)^2} \right] &= \mathbb{E}_w \left[\frac{\llbracket x \rrbracket(w)}{n(w)^2} \right]. \end{aligned}$$

An expectation in a countable space is a sum weighted by probability masses. So, by Lemma 1, this holds iff $\llbracket x \rrbracket = \llbracket xy \rrbracket = \llbracket x \rrbracket \cap \llbracket y \rrbracket$. We conclude $p(xy) = p(xx)$ if and only if $\llbracket x \rrbracket \subseteq \llbracket y \rrbracket$. \square

A similar proof suffices to show that the following isomorphism also holds:

Corollary 1.1 If p is a uniformly truthful speaker, the following isomorphism holds for all $x, y \in \mathcal{X}$:

$$\llbracket x \rrbracket \subseteq \llbracket y \rrbracket \iff p(xy) = p(x\$).$$

3.1 Discussion

Uniformly truthful speakers resemble humans in that they mimic the *tendency* of humans to tell the truth about what they believe. However, they are clearly too simple to account for human speech patterns. Most crucially, humans generally aim to produce informative speech, rather than sampling true sentences at random. More fundamentally, natural language has a countably infinite number of possible sentences, so a uniform distribution over all true sentences is not even mathematically well-defined. These limitations motivate our more involved analysis of Gricean speakers, which will adapt the technical tools used in this section.

4 Gricean Speakers

In this section, we will define a new class of speakers who pick sentences in order to be informative to their listener, while also trying to be concise. To do this, we will draw on information theory to formalize what it means for a speaker to be informative. We will then derive a distributional relation that is isomorphic to entailment for Gricean speakers, which is a generalization of the relation for uniformly truthful speakers from §3.

4.1 Definition

Information The first step towards formalizing Gricean speakers is to define a notion of the semantic information contained in a sentence. We formalize a listener $\ell(w | z)$ as the inverse of a speaker: Given a text $z \in \mathcal{X}^*$, a listener produces a distribution over possible world states. Then, in a given world w we can define the information that a text conveys to the listener as the reduction in the number of bits needed to transmit w to ℓ after they have read z compared to before they have read z .

Definition 4 The information content of a text $z \in \mathcal{X}^* \cup \mathcal{X}^*\$$ to a listener $\ell(w | z)$ is⁷

$$I_\ell(z; w) = \log \ell(w | z) - \log \ell(w).$$

In other words, the information content of a text is the reduction in ℓ 's code length for the world after having read the text compared to beforehand. We can naturally extend Definition 4 to measure the *conditional information* conveyed by sentence y given that x has already been produced:

Definition 5 The information content of $y \in \mathcal{X}^* \cup \mathcal{X}^*\$$ given $x \in \mathcal{X}^*$ to a listener $\ell(w | z)$ is

$$\begin{aligned} I_\ell(y | x; w) &= I_\ell(xy; w) - I_\ell(x; w) \\ &= \log \ell(w | xy) - \log \ell(w | x). \end{aligned}$$

Informative Speaker We now define a Gricean speaker in terms of I_ℓ . Our definition generalizes the rational speech acts model (Goodman and Frank, 2016), but makes weaker assumptions about the listener and allows a dynamic semantics where later sentences can condition on previous ones (Lewis, 1979; Kamp, 1981; Heim, 1982). We define an utterance's utility as a convex combination of its information content and its cost to produce, operationalizing the Gricean idea that speakers pick utterances by weighing their informativeness against their cost. The cost function $c : \mathcal{X}^* \cup \mathcal{X}^*\$ \rightarrow \mathbb{R}$ can be any measure of sentence complexity (e.g., length) satisfying $c(xy) = c(x) + c(y)$ for $x, y \in \mathcal{X}^* \cup \mathcal{X}^*\$$.⁸

Definition 6 A speaker p is *Gricean* if there exists a listener $\ell(w | z)$, some $\alpha > 0$, and a cost function c such that, for all $z \in \mathcal{X}^* \cup \mathcal{X}^*\$$:

$$p(z | w) \propto \exp(\alpha I_\ell(z; w) - c(z)).$$

⁷For convenience, we let $\log 0 = -\infty$ and $\infty - \infty = 0$.

⁸This is satisfied when $c(x)$ is the length of x , but also for other options like the corpus frequency of x (Goodman and Frank, 2016) or the depth of the syntactic tree of x .

Further, ℓ must satisfy the following for all $x \in \mathcal{X}^*$, $y \in \mathcal{X} \cup \{\$\}$, and $w \in \mathcal{W}$,

$$I_\ell(y | x; w) = 0 \iff \llbracket x \rrbracket(w) \rightarrow \llbracket y \rrbracket(w).$$

In other words, the speaker must be trying to convey information about the state of the world to some listener who fully absorbs the semantic information in all sentences they have already heard: clarifying already established information will not benefit the listener. We can formalize this by deriving $p(y | x, w)$ for $x \in \mathcal{X}^*$ and $y \in \mathcal{X} \cup \{\$\}$:

$$\begin{aligned} p(y | x, w) &= \frac{p(xy | w)}{p(x | w)} \\ &\propto \exp(\alpha I_\ell(y | x; w) - c(y)). \end{aligned}$$

Notably, the probability of y given x depends on the *conditional information* of y given x , which means only information conveyed by y that is nonredundant with x will make y more likely.⁹

4.2 Results

Proofs are in §C. Under a Gricean speaker, the cost of an utterance can be expressed:

Lemma 2 For any Gricean speaker p and $x \in \mathcal{X}$,

$$\frac{p(x\$)}{p(xx)} = \frac{\exp(c(x))}{\exp(c(\$))}.$$

Corollary 2.1 Under a Gricean speaker, for all $x \in \mathcal{X}$, $c(x) = \log p(x\$) - \log p(xx) + c(\$)$.

Corollary 2.1 says that a sentence is costly to the extent that it is unlikely to be repeated twice, giving an intuitive characterization of this quantity in terms of text probabilities. Now, we will use this characterization of cost to derive a distributional relation that is isomorphic to entailment.

Theorem 2 Under any Gricean speaker p , entailment is isomorphic to a distributional relation. Specifically, for all sentences $x, y \in \mathcal{X}$,

$$\llbracket x \rrbracket \subseteq \llbracket y \rrbracket \iff \frac{p(xy)}{p(x\$)} = \frac{p(yy)}{p(y\$)}.$$

If we allow our decision rule to depend on the cost function c in addition to probabilities, we can simplify Theorem 2 as follows:

⁹From a technical perspective, the \exp in Definition 6 is justified by the fact that probabilities decompose multiplicatively, i.e., $p(xy | w) = p(x | w)p(y | x, w)$, but the information content and cost of text should decompose additively across different sentences. Applying basic exponent rules shows that Definition 6 satisfies this desideratum.

Corollary 2.1 *Under any Gricean speaker p , for all sentences $x, y \in \mathcal{X}$, $\llbracket x \rrbracket \subseteq \llbracket y \rrbracket$ if and only if*

$$\log p(x\$) - \log p(xy) = c(y) - c(\$).$$

If we imagine $c(y) - c(\$) = 0$ for a uniformly truthful speaker, we see the equation in Theorem 2 is a generalization of the equation in Theorem 1.

4.3 Discussion

Gricean speakers are a general enough model of humans speakers to capture the basic pragmatic principles influencing speech production. Thus, it is notable that Theorem 2 establishes a closed-form distributional relation isomorphic to entailment.

One conceptual limitation of Gricean speakers is that their simulated listener must fully consume information, such that redundantly conveying the same information twice will not lead to any information gain the second time. This contrasts with real speech, where potential interpretation errors by the listener incentivize the speaker to be somewhat redundant (Degen et al., 2019). Mathematically, this would violate the axiom of Definition 6 that

$$I_\ell(y \mid x; w) = 0 \iff \llbracket x \rrbracket(w) \rightarrow \llbracket y \rrbracket(w).$$

Extending Theorem 2 to speakers who use redundancy to account for noise and interpretation errors is an interesting direction for future work.

Another interesting extension would be formalizing speakers who aim to be informative regarding some question under discussion, rather than being generally informative about w (cf. Goodman and Lassiter, 2015). This could encompass both “what” questions that aim to clarify some aspect of the world, and “why” questions that aim to convey explanations for established facts.

5 Decoding Entailment from Empirical Text Frequencies

We have so far shown that entailment judgments can be extracted from the sentence probabilities in the ideal distribution $p(z)$. What happens if, more practically, we estimate the probability of a sentence by its frequency in a large corpus sampled from $p(z)$? We prove this method enables feasible extraction of entailment judgments between very short sentences, but the corpus size may become intractably large for longer sentences.

Imagine we have a finite corpus of *iid* sentences $\{Z_i\}_{i=1}^n$, each sampled from $p(z)$. Let $\hat{p}(z)$ be the

empirical frequency of a text z in the corpus, i.e., if $\pi(z, z')$ returns whether text z is a prefix of text z' ,

$$\hat{p}(z) = \frac{1}{n} \sum_{i=1}^n \pi(z, Z_i).$$

Since $p(z)$ encodes entailment via our extraction rules, $\hat{p}(z)$ will encode entailment between sentences if $\hat{p}(z)$ is close to $p(z)$. A naive notion of closeness is to guarantee, for all ϵ , there exists some number of texts n such that, with high probability, $|p(z) - \hat{p}(z)| < \epsilon$. But this notion is not strict enough: if $p(z)$ is small, this difference will also be small, even if $\hat{p}(z)$ is not a good approximation of $p(z)$ on a relative scale. Instead, we want to guarantee that $\hat{p}(z)/p(z)$ converges to 1, or, equivalently, that their difference as log probabilities converges to 0. This ensures that convergence will still be meaningful for low-probability sentences, which most sentences are in natural language.

Under this standard, rarer sentences take more samples to approximate. Define the sentence complexity $\mathfrak{K}_p(z) = \frac{1}{p(z)}$. We bound the approximation error in terms of $\mathfrak{K}_p(z)$.¹⁰

Lemma 3 *For $z \in \mathcal{X}^* \cup \mathcal{X}^*\$$ and $\delta > 0$, it holds with probability at least $1 - \delta - (1 - p(z))^n$ that*

$$|\log p(z) - \log \hat{p}(z)| \leq \sqrt{\frac{\mathfrak{K}_p(z)}{\delta n}}.$$

To make this bound non-vacuous, n must be large enough to counteract $\mathfrak{K}_p(z)$ and bring $(1 - p(z))^n$ close to 0. Thus, good approximation requires fewer samples for more common sentences. To get a more concrete view of the number of samples required to extract entailment judgments from an LM, we analyze $\mathfrak{K}_p(z)$ for Gricean speakers.¹¹

Recall that we write $c(z)$ for the cost that a Gricean speaker assigns to producing a sentence z . For Gricean speakers, $\mathfrak{K}_p(z)$ is related to $c(z)$ as well as the probability z is true.

Theorem 3 *Assume that $p(z \mid w)$ is a Gricean speaker with respect to listener ℓ and $\llbracket z \rrbracket(w) = 1 \iff I_\ell(z; w) \geq 0$. Let $g_p(x, y) = \log \frac{p(xy)}{p(x\$)} - \log \frac{p(yy)}{p(y\$)}$. Let $q = 1 - \min\{p(xy), p(yy)\}$. Then, for all $x, y \in \mathcal{X}$ such that $\llbracket xy \rrbracket(p) > 0$, for all $\delta > 0$, it holds with probability at least $1 - \delta - 4q^n$ that $|g_p(x, y) - g_{\hat{p}}(x, y)|$ is at most*

$$8 \sqrt{\frac{\exp(\max\{c(xy), c(yy)\})}{p(\llbracket xy \rrbracket)} \cdot \frac{1}{\delta n}}.$$

¹⁰Omitted proofs from §5 are in §D.

¹¹§D also analyzes uniformly truthful speakers.




Figure 1: Estimated number of training sentences for guaranteeing $g_{\hat{p}}$ closely approximates g_p , where \hat{p} is estimated using empirical text frequencies.

Theorem 3 says we can use text frequencies to decode entailment between sentences x, y from a Gricean corpus, but the number of training sentences to guarantee this grows exponentially with the cost of x and y . Thus, we probably cannot expect to extract entailment judgments from text frequencies except between *very short* sentences.

We make this more quantitative in Figure 1, where we estimate the number of training sentences needed to ensure g_p and $g_{\hat{p}}$ are close on sentences of length $\leq k$ as a function of k . The main assumption behind this calculation is that a sentence’s probability vanishes exponentially in its length, where the exponential base is the perplexity of the language. §E documents the underlying assumptions in more detail. Figure 1 predicts g_p and $g_{\hat{p}}$ can be made close for length-4 sentences using $\sim 10^{10}$ training sentences: about as much data as GPT-3 was trained on. In contrast, handling (still short) sentences of length 10 can be done with $\sim 10^{17}$ training sentences, or $\sim 10^7$ GPT-3 corpora. Thus, relying solely on corpus frequencies is likely not a feasible way to extract entailment relations from text generated by Gricean speakers.

6 Decoding Entailment from LMs

We have just analyzed how many samples are necessary to decode entailment relations from the text frequencies in a finite corpus. As shown by Theorem 3, this approach will require intractably many samples for sentences of nontrivial length because longer strings will appear infrequently (if at all) in the corpus. In order to estimate the probability of rare, longer, strings what if we use an LM to estimate $\hat{p}(z)$ instead of text frequencies? Perhaps a smoothed LM should allow us to extrapolate $\hat{p}(z)$ well enough for long sentences to extract entail-

ment judgments between them. In this section, we briefly discuss some limitations of this approach.

It is tempting to take low LM perplexity as evidence that an LM estimates sentence probabilities well enough to approximately satisfy the isomorphism in Theorem 2. After all, low test perplexity implies that $\hat{p}(z)$ is, on average, a good approximation of $p(z)$: if the perplexity is bounded below ϵ , then the KL divergence $\text{KL}(p, \hat{p})$ is bounded below $\log \epsilon$. ϵ decreases with the amount of training data n at a rate between $\Omega(1/\sqrt{n})$ and $\Omega(1/n)$ (Wang et al., 2013; Li and Liu, 2021). Thus, with enough data, $\hat{p}(z)$ will closely approximate $p(z)$ for an average sentence z in the training distribution.

But low error on an average z does not establish entailment can be decoded from \hat{p} because $d_{\hat{p}}$, as derived in Theorem 2, depends on the text $z = yy$, which is very unlikely in natural language.¹² Poorly estimating $p(yy)$ has little impact on $\text{KL}(p, \hat{p})$, so LMs trained to minimize $\text{KL}(p, \hat{p})$ have no reason to estimate $p(yy)$ well unless they are imbued with strong inductive biases. Thus, we expect that LMs trained with a standard cross-entropy loss may not produce reliable entailment judgments because they poorly estimate the probability of key valid (but unlikely) texts.¹³ However, we find in the next section that they do succeed in the easier setting of small artificial languages and fully Gricean speakers.

7 Experiments: Extracting Semantics from Simulated Gricean Corpora

We test empirically whether we can extract entailment judgments from LMs trained on unlabelled text.¹⁴ Natural language corpora are unlikely to adhere exactly to our idealized assumptions about the speakers generating texts, so we generate the training corpora from a simulated Gricean speaker (see §4). To make learning semantics more tractable with limited computation, we set $|\mathcal{W}| = 3$ and restrict the vocabulary \mathcal{X} to 7 utterances, each denoting one of the 7 non-empty subsets of \mathcal{W} . Each sentence in the training corpus is generated by sampling utterances from a Gricean speaker, conditioned on a uniformly sampled world state and the

¹² yy is unlikely to be produced by a Gricean speaker because the second y conveys no information.

¹³Future work should more carefully analyze how much data is required to extract complex entailment relations from LM predictions (rather than corpus frequencies). This is beyond the scope of the current project.

¹⁴<https://github.com/viking-sudo-rm/formal-language-understanding>

previously generated utterance, until the tautological utterance is generated. The semantic value of a sentence is taken to be the conjunction over all of its utterances. We set the rationality parameter α and the cost function heuristically (details in §G).

We generate training sets varying in size from 2 texts to 10M texts, and train two types of models on each: a simple empirical text frequency as described in Section 5, and a trigram model implemented using NLTK (Bird, 2006). Then for all sentence pairs (x, y) , where x and y have 6 utterances or fewer and each denotes a non-empty proposition, we compute $g_{\hat{p}}(x, y)$ from §5. Theorem 2 shows that, under the true distribution p , $g_p(x, y) = 0$ if and only if x entails y .

The results are plotted in Figure 2. We arrive at the following conclusions:

Entailment relations can be extracted with greater-than-chance performance from LM predictions. The value of $g_{\hat{p}}(x, y)$ is much closer to 0 on average for entailed pairs than for non-entailed pairs. This is predicted by Theorem 2.

The size of the corpus needed to extract entailment grows predictably with sentence length. For entailed pairs, the average value of $g_{\hat{p}}(x, y)$ for shorter sentences approaches 0 more quickly with a large training corpus. This is in line with the predictions of Theorem 4.

Model inductive bias impacts the ease of extracting entailment. Entailed and non-entailed pairs are better distinguished by the trigram model than the text frequency model. Specifically, $g_{\hat{p}}(x, y)$ is closer to 0 for the trigram model for a given amount of data, and the trigram model’s predictions are less sensitive to sentence length.

8 Generality of Extracting Semantics

Our main result that entailment judgments can be extracted from an ideal LM assumes the corpus was produced by Gricean speakers. While pragmatic theory supports this assumption, real human speakers are undoubtedly more complex. What if we relax the assumption that speakers are Gricean? In Theorem 6 in §F, we show that any semantic relation is isomorphic to some distributional relation as long as, for any pair of possible semantics, there is some text whose probability distinguishes between the two candidate semantics.

We take it to be uncontroversial that semantics influences speech production, so we interpret Theo-

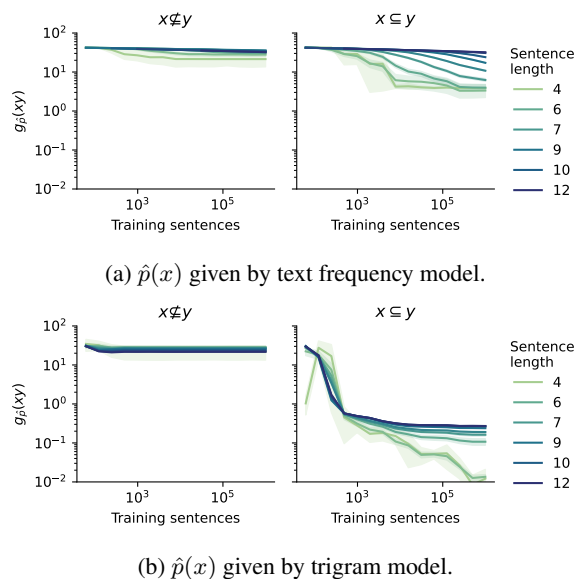


Figure 2: Plot of $g_{\hat{p}}(x, y) = \log \frac{\hat{p}(xy)}{\hat{p}(x\&)} - \log \frac{\hat{p}(yy)}{\hat{p}(y\&)}$ as a function of the number of sentences in the training corpus and the length $|xy|$. Given the true distribution p , $g_p(x, y) = 0$ iff x entails y . We exclude pairs x, y where both xy and yy are absent from the training data.

rem 6 to say all semantic relations are fully encoded in ideal LMs. In contrast to Theorem 2, however, this result is nonconstructive, so we do not know which algorithm to use to decide entailment between two sentences, even though one exists. Further, without further assumptions about the speaker, we cannot guarantee the extraction relation is efficiently computable or even computable at all.

9 Conclusion

Given a general, linguistically motivated model of human text production, we proved that entailment judgments can be decoded from the likelihood function for texts because of semantic artifacts created by human authors. We also showed empirically that entailment could be extracted n-gram LMs trained on simple formal languages. Thus, we have given one explanation for why distributional information encodes semantic information (Firth, 1957) and how semantic relations are, in principle, extractable from LMs. It is an open question whether entailment judgments might be extractable from current large LMs, but we hypothesize that the complexity of natural language makes this substantially more challenging than with our synthetic experiments, and that the loss function and inductive biases of current neural LMs are not well suited for doing so without an infeasible amount of data.

A natural next step for future work is to test this hypothesis empirically by measuring whether entailment judgments can be extracted from large LMs using our theory. Similarly, it would be interesting to think about how LMs could be modified so that they can better pick up on the semantic information encoded in their training distribution.

10 Acknowledgments

This project benefited from informal discussions with Chris Barker, Sam Bowman, Lucius Bynum, Kyunghun Cho, Tiwa Eisape, Yanai Elazar, Najaoung Kim, Alexander Koller, Vishakh Padmakumar, Chris Potts, Naomi Saphra, Sebastian Schuster, and Noah A. Smith. We also thank the members of the CAP Lab and Semantics Group at NYU for their feedback. This work was supported in part through the NYU IT High Performance Computing resources, services, and staff expertise. It was funded by NSF award 1922658, and WM was supported by an NSF graduate research fellowship.

References

- Emily M. Bender and Alexander Koller. 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.
- Steven Bird. 2006. NLTK: The Natural Language Toolkit. In *Proceedings of the COLING/ACL 2006 Interactive Presentation Sessions*, pages 69–72, Sydney, Australia. Association for Computational Linguistics.
- Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri Chatterji, Annie Chen, Kathleen Creel, Jared Quincy Davis, Dora Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khattab, Pang Wei Koh, Mark Krass, Ranjay Krishna, Rohith Kuditipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvir Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avanika Narayan, Deepak Narayanan, Ben Newman, Allen Nie, Juan Carlos Niebles, Hamed Nilforoshan, Julian Nyarko, Giray Ogut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Rob Reich, Hongyu Ren, Frieda Rong, Yusuf Roohani, Camilo Ruiz, Jack Ryan, Christopher Ré, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishnan Srinivasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. 2021. On the opportunities and risks of foundation models.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Judith Degen, Robert XD Hawkins, Caroline Graf, Elisa Kreiss, and Noah D Goodman. 2019. When redundancy is rational: A Bayesian approach to “over-informative” referring expressions. *arXiv preprint arXiv:1903.08237*.
- 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.
- John R Firth. 1957. A synopsis of linguistic theory, 1930-1955. *Studies in linguistic analysis*.
- Noah D Goodman and Michael C Frank. 2016. Pragmatic language interpretation as probabilistic inference. *Trends in cognitive sciences*, 20(11):818–829.
- Noah D Goodman and Daniel Lassiter. 2015. Probabilistic semantics and pragmatics: Uncertainty in language and thought. *The handbook of contemporary semantic theory, 2nd edition*. Wiley-Blackwell.
- Herbert P Grice. 1975. Logic and conversation. In *Speech acts*, pages 41–58. Brill.
- Irene Heim and Angelika Kratzer. 1998. *Semantics in Generative Grammar*. Blackwell.

- Irene Roswitha Heim. 1982. *The semantics of definite and indefinite noun phrases*. University of Massachusetts Amherst.
- Hans A Kamp. 1981. A theory of truth and semantic representation formal methods in the study of language, part 1, ed. by jeroen groenendijk, theo janssen and martin and stokhof. *Amsterdam: Mathematisch Centrum*.
- Angelika Kratzer. 2021. Situations in Natural Language Semantics. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*, Winter 2021 edition. Metaphysics Research Lab, Stanford University.
- Stephen C Levinson, Stephen C Levinson, and S Levinson. 1983. *Pragmatics*. Cambridge university press.
- David Lewis. 1979. Scorekeeping in a language game. In *Semantics from different points of view*, pages 172–187. Springer.
- Shaojie Li and Yong Liu. 2021. Towards sharper generalization bounds for structured prediction. *Advances in Neural Information Processing Systems*, 34.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. TruthfulQA: Measuring how models mimic human falsehoods. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252, Dublin, Ireland. Association for Computational Linguistics.
- William Merrill, Yoav Goldberg, Roy Schwartz, and Noah A. Smith. 2021. Provable Limitations of Acquiring Meaning from Ungrounded Form: What Will Future Language Models Understand? *Transactions of the Association for Computational Linguistics*, 9:1047–1060.
- Julian Michael. 2020. To dissect an octopus: Making sense of the form/meaning debate.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. 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 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Christopher Potts. 2020. Is it possible for language models to achieve understanding?
- Shaojun Wang, Russell Greiner, and Shaomin Wang. 2013. Consistency and generalization bounds for maximum entropy density estimation. *Entropy*, 15(12):5439–5463.

A Limitations

We derived a recipe for computing entailment in terms of text probabilities, hinting that entailment judgments may be decodable from LM predictions. Yet two key concerns qualify this conclusion.

Learnability We reduce entailment classification to computing probabilities in the *target distribution* of an LM, not probabilities predicted by an LM. In §6, we argue that the loss function of current LMs is not well suited to producing models from which entailment can be extracted.

Speaker Assumptions Gricean speakers capture important factors influencing speech production in pragmatic theory, but human speakers are undoubtedly more complex. Based on §8, we expect a similar isomorphism to hold under any reasonable speaker model, but the mathematical form may change and it may become harder to compute.

B Uncountable World Spaces

In this section, we assume \mathcal{W} is an uncountably infinite set with a probability density function $p(w)$. We then define “almost sure” entailment as follows:

Definition 7 For $x, y \in \mathcal{X}$, we say x almost surely entails y (i.e., $\llbracket x \rrbracket \sqsubseteq \llbracket y \rrbracket$) if and only if

$$p(\llbracket x \rrbracket \setminus \llbracket y \rrbracket) = 0.$$

Note that if \mathcal{W} is countable, then $\mathcal{A} \sqsubseteq \mathcal{B}$ reduces to $\mathcal{A} \subseteq \mathcal{B}$. We can generalize Lemma 1 as follows, which shows that all our results go through for almost sure entailment when \mathcal{W} is uncountable.

Lemma 4 Let $\mathbb{1}_{\mathcal{S}}$ be the indicator function for set \mathcal{S} . Let $f : \mathcal{W} \rightarrow \mathbb{R}$ be some function such that $\inf_{w \in \mathcal{W}} f(w) > 0$. For any sets \mathcal{A}, \mathcal{B} such that $\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{W}$, then $p(\mathcal{B} \setminus \mathcal{A}) = 0$ if and only if

$$\mathbb{E}_{w \sim p(w)} [\mathbb{1}_{\mathcal{A}}(w)f(w)] = \mathbb{E}_{w \sim p(w)} [\mathbb{1}_{\mathcal{B}}(w)f(w)].$$

Proof. If $p(\mathcal{B} \setminus \mathcal{A}) = 0$, then the condition follows by construction. We thus only need to show that the condition follows from $p(\mathcal{B} \setminus \mathcal{A}) = 0$. Let $q = p(\mathcal{B} \setminus \mathcal{A})$. By linearity of expectation, we rewrite the premise condition as

$$\begin{aligned} 0 &= \mathbb{E}_{w \sim p(w)} [(\mathbb{1}_{\mathcal{A}}(w) - \mathbb{1}_{\mathcal{B}}(w)) f(w)] \\ &= \mathbb{E}_{w \sim p(w)} [(\mathbb{1}_{\mathcal{A}}(w) - \mathbb{1}_{\mathcal{B}}(w)) f(w) \mid w \in \mathcal{B} \setminus \mathcal{A}] q \\ &\quad + \mathbb{E}_{w \sim p(w)} [(\mathbb{1}_{\mathcal{A}}(w) - \mathbb{1}_{\mathcal{B}}(w)) f(w) \mid w \notin \mathcal{B} \setminus \mathcal{A}] (1 - q) \\ &\geq \mathbb{E}_{w \sim p(w)} [f(w) \mid w \in \mathcal{B} \setminus \mathcal{A}] q. \end{aligned}$$

Letting $f^* = \inf_{w \in \mathcal{W}} f(w) > 0$, we get $0 \geq f^* q$. Since $f^* > 0$ and $q \geq 0$, $q = p(\mathcal{B} \setminus \mathcal{A}) = 0$. \square

C Gricean Speaker Proofs

Lemma 2 For any Gricean speaker p and $x \in \mathcal{X}$,

$$\frac{p(x\$)}{p(xx)} = \frac{\exp(c(x))}{\exp(c(\$))}.$$

Proof. Starting with the definition of an Gricean speaker, for any $x \in \mathcal{X}^*$ and $y \in \mathcal{X} \cup \{\$\}$,

$$p(xy) = \mathbb{E}_w [p(x \mid w)p(y \mid x, w)].$$

Now, letting $g(x, w) \triangleq p(x \mid w) / \left(\sum_{y'} \exp(\alpha I_\ell(y' \mid x; w) - c(y')) \right)$,

$$p(xy) = \exp(-c(y)) \mathbb{E}_w [\exp(\alpha I_\ell(y \mid x; w))g(x, w)].$$

We apply this identity to both sides of the fraction in the lemma statement:

$$\begin{aligned}\frac{p(x\$)}{p(xx)} &= \frac{\exp(-c(\$)) \mathbb{E}_w [\exp(\alpha I_\ell(\$ | x; w))g(x, w)]}{\exp(-c(x)) \mathbb{E}_w [\exp(\alpha I_\ell(x | x; w))g(x, w)]} \\ &= \frac{\exp(c(x))}{\exp(c(\$))} \cdot \frac{\mathbb{E}_w [\exp(\alpha I_\ell(\$ | x; w))g(x, w)]}{\mathbb{E}_w [\exp(\alpha I_\ell(x | x; w))g(x, w)]}.\end{aligned}$$

Since $\llbracket x \rrbracket \subseteq \llbracket \$ \rrbracket$ and $\llbracket x \rrbracket \subseteq \llbracket x \rrbracket$, we know that the conditional information of both $\$$ and x given x is 0, and, thus,

$$\frac{p(x\$)}{p(xx)} = \frac{\exp(c(x))}{\exp(c(\$))} \cdot \frac{\mathbb{E}_w [\exp(0)g(x, w)]}{\mathbb{E}_w [\exp(0)g(x, w)]} = \frac{\exp(c(x))}{\exp(c(\$))}.$$

□

Theorem 2 *Under any Gricean speaker p , entailment is isomorphic to a distributional relation. Specifically, for all sentences $x, y \in \mathcal{X}$,*

$$\llbracket x \rrbracket \subseteq \llbracket y \rrbracket \iff \frac{p(xy)}{p(x\$)} = \frac{p(yy)}{p(y\$)}.$$

Proof. Recall from the proof of Lemma 2 that there exists a function $g(x, w)$ such that, for all $x \in \mathcal{X}^*$ and $y \in \mathcal{X} \cup \{\$\}$,

$$p(xy) \propto \exp(-c(y)) \mathbb{E}_w [\exp(\alpha I_\ell(y | x; w))g(x, w)].$$

Thus, by Lemma 2, the proposed distributional relation can be expanded as

$$\begin{aligned}d_p(x, y) &\iff \frac{p(xy)}{p(x\$)} = \frac{p(yy)}{p(y\$)} \\ &\iff p(xy) \cdot \frac{p(y\$)}{p(yy)} = p(x\$) \cdot \frac{p(xx)}{p(xx)} \\ &\iff p(xy) \frac{\exp(c(y))}{\exp(c(\$))} = p(x\$) \frac{\exp(c(x))}{\exp(c(\$))} \\ &\iff p(xy) \exp(c(y)) = p(x\$) \exp(c(x)) \\ &\iff \mathbb{E}_w [\exp(\alpha I_\ell(y | x; w))g(x, w)] = \mathbb{E}_w [\exp(\alpha I_\ell(x | x; w))g(x, w)].\end{aligned}$$

By Lemma 1, this holds if and only if, for all w ,

$$\begin{aligned}\exp(\alpha I_\ell(y | x; w)) &= \exp(\alpha I_\ell(x | x; w)) \\ I_\ell(y | x; w) &= I_\ell(x | x; w) \\ I_\ell(y | x; w) &= 0 \\ \llbracket x \rrbracket(w) &\rightarrow \llbracket y \rrbracket(w) = 1.\end{aligned}$$

We conclude the distributional relation holds if and only if $\llbracket x \rrbracket \subseteq \llbracket y \rrbracket$. □

D Proofs for Learning Bounds

Lemma 3 *For $z \in \mathcal{X}^* \cup \mathcal{X}^*\$$ and $\delta > 0$, it holds with probability at least $1 - \delta - (1 - p(z))^n$ that*

$$|\log p(z) - \log \hat{p}(z)| \leq \sqrt{\frac{\mathfrak{K}_p(z)}{\delta n}}.$$

Proof. Without loss of generality, assume $p(z) > 0$. With probability $1 - (1 - p(z))^n$ over the draw of our sample, the random variable $\log \hat{p}(z)$ has finite variance defined by

$$\text{Var} [\log \hat{p}] = \frac{1}{n} \cdot \frac{1 - p(z)}{p(z)} \leq \frac{\mathfrak{K}_p(z)}{n}.$$

With finite variance, we can apply Chebyshev's inequality to conclude that

$$\Pr [|\log p(z) - \log \hat{p}(z)| \geq \epsilon] \leq \frac{\text{Var} [\log \hat{p}]}{\epsilon^2} \leq \frac{\mathfrak{K}_p(z)}{n\epsilon^2}.$$

Solving for $\delta \leq \Pr [|\log p(z) - \log \hat{p}(z)|]$, we get

$$\begin{aligned} \delta &\leq \frac{\mathfrak{K}_p(z)}{n\epsilon^2} \\ \therefore \epsilon &\leq \sqrt{\frac{\mathfrak{K}_p(z)}{\delta n}}. \end{aligned}$$

We conclude that that with probability $1 - \delta - (1 - p(z))^n$,

$$|\log p(z) - \log \hat{p}(z)| \leq \sqrt{\frac{\mathfrak{K}_p(z)}{\delta n}}.$$

□

We now characterize the complexity factor $\mathfrak{K}_p(z)$ for uniformly truthful speakers.

Lemma 5 For all $z \in \mathcal{X}^* \cup \mathcal{X}^*\$$ such that $\llbracket z \rrbracket(p) > 0$, it holds that

$$\mathfrak{K}_p(z) \leq \frac{|\mathcal{X}|}{p(\llbracket z \rrbracket)}.$$

Proof. We start by deriving a lower bound on $p(z)$.

$$\begin{aligned} p(z) &= \sum_w \frac{\llbracket z \rrbracket(w)}{\sum_{z'} \llbracket z' \rrbracket(w)} p(w) \\ &\geq \sum_w \frac{\llbracket z \rrbracket(w)}{|\mathcal{X}|} p(w) \\ &= \frac{\llbracket z \rrbracket(p)}{|\mathcal{X}|}. \end{aligned}$$

Applying this inequality to the definition of $\mathfrak{K}_p(z)$, we conclude that

$$\mathfrak{K}_p(z) \leq \frac{|\mathcal{X}|}{\llbracket z \rrbracket(p)}.$$

□

Lemma 5 lets us to derive the following guarantee for estimating entailment scores using a corpus produced by uniformly truthful speakers:

Theorem 4 For a uniformly truthful speaker p , let $u_p(x, y) = \log p(x\$) - \log p(xy)$. For $x, y \in \mathcal{X}$ such that $\llbracket xy \rrbracket(p) > 0$ and $\delta > 0$, it holds with probability at least $1 - \delta - 2(1 - p(xy))^n$ that

$$|u_p(x, y) - u_{\hat{p}}(x, y)| \leq 2\sqrt{\frac{|\mathcal{X}|}{p(\llbracket xy \rrbracket)}} \cdot \frac{2}{\delta n}.$$

Proof. We expand the difference in scores as follows:

$$|u_p(x, y) - u_{\hat{p}}(x, y)| \leq |\log p(x) - \log \hat{p}(x\$)| + |\log p(xy) - \log \hat{p}(xy)|.$$

We then apply Lemma 3 with $\frac{\delta}{2}$. Since $p(x\$) \geq p(xy)$, this implies that with probability $1 - \delta - 2(1 - p(xy))^n$,

$$\begin{aligned} |u_p(x, y) - u_{\hat{p}}(x, y)| &\leq \sqrt{\frac{2\mathfrak{K}_p(x\$)}{\delta n}} + \sqrt{\frac{2\mathfrak{K}_p(xy)}{\delta n}} \\ &\leq 2\sqrt{\frac{2 \max\{\mathfrak{K}_p(x\$), \mathfrak{K}_p(xy)\}}{\delta n}}. \end{aligned}$$

Finally, we apply Lemma 5 to conclude that

$$\begin{aligned} |u_p(x, y) - u_{\hat{p}}(x, y)| &\leq 2\sqrt{\frac{|\mathcal{X}|}{\min\{\llbracket x\$ \rrbracket(p), \llbracket xy \rrbracket(p)\}} \cdot \frac{2}{\delta n}} \\ &= 2\sqrt{\frac{|\mathcal{X}|}{\llbracket xy \rrbracket(p)} \cdot \frac{2}{\delta n}}. \end{aligned}$$

□

We now characterize the complexity factor for Gricean speakers.

Lemma 6 *Assume that $p(z | w)$ is a Gricean speaker with respect to listener ℓ and $\llbracket z \rrbracket(w) = 1 \iff I_\ell(z; w) \geq 0$. Then, for all $z \in \mathcal{X}^* \cup \mathcal{X}^*\$,$*

$$\mathfrak{K}_p(z) \leq \frac{\exp(c(z))}{p(\llbracket z \rrbracket)}.$$

Proof. We start by writing out the form of $p(z)$:

$$p(z) = \frac{\sum_w \exp(\alpha I_\ell(z; w)) p(w)}{\exp(c(z))}.$$

Because $z \in \mathcal{X}^* \cup \mathcal{X}^*\$,$ all terms where $\llbracket z \rrbracket(w) = 1$ contribute at least 0 information; other terms contribute negative information. Thus, we bound the information content of the “true” terms above 0, and ignore the other terms to get the lower bound

$$\begin{aligned} p(z) &\geq \frac{\sum_w \llbracket z \rrbracket(w) \exp(0) p(w)}{\exp(c(z))} \\ &= \frac{\sum_w \llbracket z \rrbracket(w) p(w)}{\exp(c(z))} \\ &= \frac{\llbracket z \rrbracket(p)}{\exp(c(z))}. \end{aligned}$$

Plugging this into $\mathfrak{K}_p(z)$, we conclude that

$$\mathfrak{K}_p(z) \leq \frac{\exp(c(z))}{\llbracket z \rrbracket(p)}.$$

□

Theorem 3 *Assume that $p(z | w)$ is a Gricean speaker with respect to listener ℓ and $\llbracket z \rrbracket(w) = 1 \iff I_\ell(z; w) \geq 0$. Let $g_p(x, y) = \log \frac{p(xy)}{p(x\$)} - \log \frac{p(yy)}{p(y\$)}$. Let $q = 1 - \min\{p(xy), p(yy)\}$. Then, for all $x, y \in \mathcal{X}$ such that $\llbracket xy \rrbracket(p) > 0$, for all $\delta > 0$, it holds with probability at least $1 - \delta - 4q^n$ that $|g_p(x, y) - g_{\hat{p}}(x, y)|$ is at most*

$$8\sqrt{\frac{\exp(\max\{c(xy), c(yy)\})}{p(\llbracket xy \rrbracket)} \cdot \frac{1}{\delta n}}.$$

Proof. We start by expanding $g_p(x, y)$:

$$\begin{aligned} g_p(x, y) &= \log \frac{p(xy)}{p(x\$)} - \log \frac{p(yy)}{p(y\$)} \\ &= \log p(xy) - \log p(x\$) - \log p(yy) + \log p(y\$). \end{aligned}$$

Thus, following Theorem 4, we can bound

$$\begin{aligned} |g_p(x, y) - g_{\hat{p}}(x, y)| &\leq |\log p(xy) - \log \hat{p}(xy)| + |\log p(x\$) - \log \hat{p}(x\$)| \\ &\quad + |\log p(yy) - \log \hat{p}(yy)| + |\log p(y\$) - \log \hat{p}(y\$)|. \end{aligned}$$

We apply Lemma 3 to each term with $\frac{\delta}{4}$. Since $p(yy) \leq p(y\$)$ and $p(xy) \leq p(x\$)$, we get that with probability at least $1 - \delta - 4q^n$,

$$\begin{aligned} |g_p(x, y) - g_{\hat{p}}(x, y)| &\leq 4\sqrt{\frac{4 \max\{\mathfrak{K}_p(xy), \mathfrak{K}_p(x\$), \mathfrak{K}_p(yy), \mathfrak{K}_p(y\$)\}}{\delta n}} \\ &= 8\sqrt{\frac{\max\{\mathfrak{K}_p(xy), \mathfrak{K}_p(x\$), \mathfrak{K}_p(yy), \mathfrak{K}_p(y\$)\}}{\delta n}}. \end{aligned}$$

Finally, we apply Lemma 6 to conclude that, with probability at least $1 - \delta - 4q^n$,

$$\begin{aligned} |g_p(x, y) - g_{\hat{p}}(x, y)| &\leq 8\sqrt{\max\left\{\frac{\exp(c(xy))}{\llbracket xy \rrbracket(p)}, \frac{\exp(c(x\$))}{\llbracket x\$ \rrbracket(p)}, \frac{\exp(c(yy))}{\llbracket yy \rrbracket(p)}, \frac{\exp(c(y\$))}{\llbracket y\$ \rrbracket(p)}\right\}} \cdot \frac{1}{\delta n} \\ &\leq 8\sqrt{\frac{\exp(\max\{c(xy), c(yy)\})}{\llbracket xy \rrbracket(p)}} \cdot \frac{1}{\delta n}. \end{aligned}$$

□

We can use Corollary 2.1 to derive a tighter version of Theorem 3 by removing the dependence on the uncommon string yy :

Theorem 5 *Let $s_p(x, y) = \log \frac{p(x\$)}{p(xy)} - c(y) + c(\$)$. Then, for all $x, y \in \mathcal{X}$ such that $\llbracket xy \rrbracket(p) > 0$, for all $\delta > 0$, the following holds with probability $1 - \delta - 2(1 - p(xy))^n$,*

$$|s_p(x, y) - s_{\hat{p}}(x, y)| \leq 2\sqrt{\frac{\exp(c(xy))}{p(\llbracket xy \rrbracket)}} \cdot \frac{2}{\delta n}.$$

The proof follows analogously to Theorem 3. The main improvement of Theorem 5 compared to Theorem 3 is that the probability the bound holds no longer depends on the unlikely probability $p(yy)$. We also get the benefit that the cost complexity factor has been reduced to only depend on $c(xy)$ and obtain better constants ($2\sqrt{2}$ instead of 8), although these changes are likely less important than removing the dependence on $p(yy)$. Of course, the drawback is that we are assuming access to the cost function $c(y)$. If we have such access, though, the improvements in the bound suggest we may be able to extract entailment from a finite corpus of Gricean text with better sample complexity than if we did not.

E Sample Complexity Estimation Details

Assuming the approximation error in Theorem 3 is $\leq \epsilon$, we aim to solve the following inequality for n :

$$\epsilon \leq 8\sqrt{\frac{\exp(\max\{c(xy), c(yy)\})}{p(\llbracket xy \rrbracket)}} \cdot \frac{1}{\delta n}.$$

Sentence Length We make the simplifying assumption that $\max\{c(xy), c(yy)\} = 2w(\ell + 1)$, where ℓ is a variable representing sentence length.¹⁵ Let Σ be the word-level vocabulary of English. We estimate the value w by assuming $q(z) = \exp(-w(|z| + 1))$ is a valid prior over Σ^* and solving for the unique value of w to satisfy this condition:

$$\begin{aligned} \sum_{z \in \Sigma^*} \exp(-w(|z| + 1)) &= 1 \\ \sum_{\ell=0}^{\infty} \frac{|\Sigma|^\ell}{\exp(w(\ell + 1))} &= 1 \\ \exp(-w) \sum_{\ell=0}^{\infty} \left(\frac{|\Sigma|}{\exp(w)}\right)^\ell &= 1 \\ \therefore w &= \log(|\Sigma| + 1). \end{aligned}$$

This reveals that w should be set ≥ 1 , but the question remains how to set $|\Sigma|$. In practice, we assume the speaker prior is defined over the support of all syntactically valid or likely strings in English, not over all possible strings as derived above. Letting S be the word-level perplexity of English, we set w according to

$$w \approx \log(S + 1).$$

We set S to the value estimated by GPT-3: ~ 20 nats/word (Brown et al., 2020). Simplifying the numerator in the bound yields

$$\exp(\log(21)(\ell + 1)) = 21^{\ell+1}.$$

Making the prior less strong, i.e., increasing $|\Sigma|$ to be greater than this perplexity estimate, would only increase the number of samples needed to extract entailment judgments.

Truth Probability We conservatively assume $p(\llbracket xy \rrbracket) = \frac{1}{2}$, although in practice it may be smaller for more informative sentences. Reducing it would lead to higher sample complexity estimates.

Final Form Putting together our estimates for sentence length and truth probability yields

$$\begin{aligned} \epsilon &\leq 8 \sqrt{\frac{2 \cdot 21^{\ell+1}}{\delta n}} \\ \therefore n &\leq \frac{128 \cdot 21^{\ell+1}}{\delta \epsilon^2}. \end{aligned}$$

The final form captures the intuition that the likelihood of a string vanishes exponentially with its length, and that the base of this decay is roughly inversely proportional to the perplexity of the language. In practice, we set $\delta = 0.1$ and $\epsilon = 1.0$. Changing the value of ϵ (the desired approximation accuracy) would shift the curve.

F General Relations and Speakers

So far, we have characterized concrete distributional relations that are isomorphic to entailment for different classes of speaker models. In this section, we analyze the conditions under which a distribution relation isomorphic to a semantic relation exists, given no assumptions about the speaker. Informally, we prove in Theorem 6 that a distributional isomorphism exists if and only if the speaker model depends on semantics “at all”. This is a very weak condition, and should be satisfied by any reasonable model of natural speakers. Thus, we take this as evidence that any speaker model—not just the ones we have considered, admits a distributional relation isomorphic to entailment.

¹⁵We write $\ell + 1$ instead of ℓ here for technical reasons: we want to guarantee that $q(z)$ can be a valid probability distribution.

We now turn to the formal presentation of this result. Let M be the function that takes a set of worlds \mathcal{W} and returns all semantic evaluation functions $\mu : \mathcal{X} \mapsto 2^{\mathcal{W}}$ over \mathcal{W} . For a semantic evaluation function $\mu = \lambda x. \llbracket x \rrbracket$, let p_μ be a speaker model parameterized by semantics μ .

Say two semantic evaluation functions μ, μ' are isomorphic with respect to s if and only if, for all x, y ,

$$S(\mu(x), \mu(y)) \iff S(\mu'(x), \mu'(y)).$$

Theorem 6 *The following are equivalent for any speaker p and semantic relation s :*

1. *There exists a distribution relation d such that, for all \mathcal{W} , for all $\mu \in M(\mathcal{W})$, s is isomorphic to d_{p_μ} .*
2. *For all $\mathcal{W}, \mathcal{W}'$, for all $\mu \in M(\mathcal{W})$ and $\mu' \in M(\mathcal{W}')$ such that μ, μ' are not isomorphic w.r.t. s , there exists $z \in \mathcal{X}^*$ such that $p_\mu(z) \neq p_{\mu'}(z)$.*

Proof. We will show that equivalence holds in both directions.

Forward Direction: We assume the second statement does not hold by way of modus tollens. Thus, there exists $\mathcal{W}, \mathcal{W}'$ with $\mu \in M(\mathcal{W})$ and $\mu' \in M(\mathcal{W}')$ with μ, μ' not isomorphic such that, for all $z \in \mathcal{X}^*$, $p_\mu(z) = p_{\mu'}(z)$. Thus, for all d and sentences x, y ,

$$d_{p_\mu}(x, y) \iff d_{p_{\mu'}}(x, y).$$

But μ and μ' are not isomorphic, so there exist x, y such that $S(\mu(x), \mu(y)) \not\iff S(\mu'(x), \mu'(y))$. Thus, we can conclude that one of the following must hold:

$$\begin{aligned} d_{p_\mu}(x, y) &\not\iff S(\mu(x), \mu(y)) \\ d_{p_{\mu'}}(x, y) &\not\iff S(\mu'(x), \mu'(y)). \end{aligned}$$

We conclude by modus tollens that the first statement implies the second.

Backward Direction: Assume the second statement holds. The function $f(\mu) = p_\mu$ is invertible up to isomorphism to s . In other words, there exists $g(p_\mu) = \mu^*$ such that, for all x, y ,

$$S(\mu^*(x), \mu^*(y)) \iff S(\mu(x), \mu(y)).$$

Then we define d according to

$$\begin{aligned} d_{p_\mu}(x, y) &\iff S(g(p_\mu)(x), g(p_\mu)(y)) \\ &\iff S(\mu^*(x), \mu^*(y)) \\ &\iff S(\mu(x), \mu(y)). \end{aligned}$$

Thus, the second statement implies the first. □

G Experimental Details

G.1 Language Description

We set the vocabulary $\mathcal{X} = \{100, 010, 001, 110, 011, 111\}$ and define $\mathcal{W} = \{1, 2, 3\}$. We refer to each three-digit binary string as an *utterance*, and define the evaluation function for an utterance x as $\llbracket x \rrbracket(w) = 1 \iff x_w = 1$. Thus, 100 is true only in world 1, while 111 is true in all worlds (i.e., is tautological). We identify 111 with the end of sequence \$.

In line with our formal definitions, we define a *text* z as a concatenation of utterances $z_1 \cdots z_n$ ending with \$. Recall that we define the evaluation function over a text as the intersection of the evaluation functions of the utterances it contains. For our language, this reduces to $\llbracket z \rrbracket(w) = 1 \iff \forall i (z_i = 1)$. Thus, 011 101 111 is true only in w_3 , and 011 101 110 111 is true in no worlds (i.e., contradictory).




Figure 3: Properties of the data generated by the speaker in our experiments, with $\alpha = 5$ and $c(x) = 0.1 \cdot |x|$.

G.2 Speaker Model Parameters

We model the listener of the informative speaker as a literal listener (Goodman and Frank, 2016), which means our informative speaker is a rational speaker of depth 1 in the language of rational speech acts.

We set $c(x) = 0.1 \cdot |x|$, where $|x|$ is the length of the string x . We set the rationality parameter $\alpha = 5$. These choices were made heuristically, by inspecting the the properties of the speaker’s output, as summarized in Figure 3. These parameters led to a relatively uniform distribution over utterances (except for the stop token 111 which is present in all texts), and a variety of text lengths without excessive redundancy. We found that larger values of α or of the coefficient for the cost function produced short texts, biasing maximally informative utterances (i.e., 100, 010, or 001); while smaller values produced long, repetitive utterances or sometimes empty utterances.

G.3 Training and Evaluation

We sample a dataset from a speaker by independently sampling n texts from the speaker model. We generate datasets of varying size from each speaker, with the number of texts n decreasing by factors of 2 from 10^7 texts down to just 2 texts.

We train models of two kinds: a text frequency model, and a trigram model. The text frequency model simply assigns a probability to a text proportional to its frequency in the training data, assigning a small $\epsilon = 10^{-20}$ probability to an unknown sequence. The trigram model is trained using NLTK’s (Bird, 2006) MLE implementation, i.e., the probabilities are unsmoothed. We do not need to use smoothing due to the small number of possible trigrams in the language.

For evaluation data, we generate pairs of texts labeled for entailments. We include all pairs where each text is 6 utterances or shorter, except for utterances that are contradictory or consist only of the end of sequence token. The total number of test pairs is about 1.1M.