

# TaxiNLI: Taking a Ride up the NLU Hill

Pratik Joshi<sup>2†\*</sup>, Somak Aditya<sup>1†</sup>, Aalok Sathe<sup>3†\*</sup>, and Monojit Choudhury<sup>1</sup>

<sup>1</sup>Microsoft Research India, 9 Lavelle Road, Vigyan, Bengaluru, India

<sup>2</sup>Google Research, Carina East Tower Bagmane Constellation Business Park, Bengaluru, India

<sup>3</sup>University of Richmond, 410 Westhampton Way, Richmond, VA, USA

pratikmjoshi123@gmail.com, aalok.sathe@richmond.edu, {t-soadit, monojitc}@microsoft.com

## Abstract

Pre-trained Transformer-based neural architectures have consistently achieved state-of-the-art performance in the Natural Language Inference (NLI) task. Since NLI examples encompass a variety of linguistic, logical, and reasoning phenomena, it remains unclear as to which specific concepts are learnt by the trained systems and where they can achieve strong generalization. To investigate this question, we propose a taxonomic hierarchy of categories that are relevant for the NLI task. We introduce TAXINLI, a new dataset, that has 10k examples from the MNLI dataset (Williams et al., 2018) with these taxonomic labels. Through various experiments on TAXINLI, we observe that whereas for certain taxonomic categories SOTA neural models have achieved near perfect accuracies—a large jump over the previous models—some categories still remain difficult. Our work adds to the growing body of literature that shows the gaps in the current NLI systems and datasets through a systematic presentation and analysis of reasoning categories.

## 1 Introduction

The Natural Language Inference (NLI) task tests whether a hypothesis (H) in text contradicts with, is entailed by, or is neutral with respect to a given premise (P) text. This 3-way classification task, popularized by Bowman et al. (2015), which was in turn inspired by Dagan et al. (2005), now serves as a benchmark for *evaluation* of natural language *understanding* (NLU) capability of models; for example, NLI datasets (Bowman et al., 2015; Williams et al., 2018) are included in all NLU benchmarks such as GLUE and SuperGLUE (Wang et al., 2018). These corpora, in turn, have been successfully used to train models such as BERT (Devlin et al.,

2019) to achieve state-of-the-art (SOTA) performance in these tasks. Despite the wide adoption of NLI datasets, a growing concern in the community has been the lack of clarity as to *which linguistic or reasoning concepts these trained NLI systems are truly able to learn and generalize* (see, for example (Linzen, 2020) and (Bender and Koller, 2020), for a discussion). Over the years, as models have shown steady performance increases in NLI tasks, many authors (Nie et al., 2019; Kaushik et al., 2019) demonstrate steep drops in performance when these models are tested against adversarially (or counterfactually) created examples by non-experts. Richardson et al. (2019) use templated examples to show trained NLI systems fail to capture essential logical (negation, boolean, quantifier) and semantic (monotonicity) phenomena.

Herein lie the central questions of our work: 1) what is the distribution of various categories of *reasoning tasks* in the NLI datasets? 2) which categories of tasks are rarely captured by current NLI datasets (owing to the nature of the task and the non-expert annotators)? 3) which categories are well-understood by the SOTA models? and 4) are there categories where Transformer-based architectures are consistently deficient?

In order to answer these questions, we first discuss why performance-specific error analysis categories (Wang et al., 2018; Nie et al., 2019), and stress testing categories (Naik et al., 2018) are inadequate. We then propose a taxonomy of the various reasoning tasks that are commonly covered by the current NLI datasets (Sec 2). Next, we annotate 10,071 P-H pairs from the MNLI dataset (Williams et al., 2018) with the lowest level taxonomic categories, 18 in total (Sec 3). Then we conduct various experiments and careful error analysis of the SOTA models—BERT and RoBERTa, as well as other baselines, such as Bag-of-words Naïve Bayes and ESIM, on their performance across these categories

† denotes equal contribution. \*Work was done while Authors were at Microsoft Research India.

(Sec 4). Our analyses indicate that while these models perform well on some categories such as linguistic reasoning, the performance on many other categories, such as those that require world knowledge or temporal reasoning, are quite poor. We also look into the embeddings of the P-H pairs to understand which of these categorical distinctions are captured well in the learnt representations, and which get conflated (Sec 5). In line with our previous finding, we observe strong correlation between the level of clustering within the representation of the examples from a category, and the performance of the models for that particular category.

## 2 A New Taxonomy for NLI

### 2.1 Necessity for a New Taxonomy

According to Wittgenstein (1922), “Language disguises the thought”, and human beings try to gauge such thought from colloquial language using “complex silent adjustments”. The journey from lexicon and syntax of “language” to the aspects of semantics and pragmatics can be thought of as a journey that portrays important milestones that an ideal NLU system should achieve. Irrespective of the order of such milestones<sup>1</sup>, we believe that NLU (and NLI) systems should be tested and analyzed with respect to fundamental linguistic and logical phenomena. Recently, different types of phenomena have been tested through 1) creating new datasets, 2) probing tasks, and 3) error-analysis categorizations. Researchers have created new datasets by recasting various NLU tasks to a large NLI dataset (Poliak et al., 2018), eliciting counter-factual examples from non-experts by considering different lexical and reasoning factors (Kaushik et al., 2019), and adversarial example (Nie et al., 2019) elicitation by letting non-experts come up with examples through interacting with SOTA systems. However, these datasets do not expose the linguistic aspects where the current systems have difficulty. Using the probing task methodology, researchers (Jawahar et al., 2019; Goldberg, 2019) observed that BERT captures syntactic structure, along with some semantics such as NER, and semantic role labels (Tenney et al., 2019b). However, BERT’s ability to reason is questioned by the observed performance degradation in MNLI (McCoy et al., 2019). Linzen (2020) also called for a pretraining-agnostic evaluation setup, where the setup is not limited to pre-trained

<sup>1</sup>“For an infant, a foreigner, or an instant-message addict, context is more important than syntax” (Sowa, 2010).

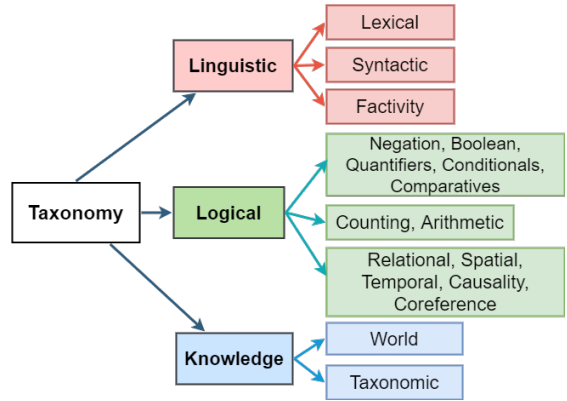


Figure 1: Taxonomic Categorization of the NLI task.

language models. Our taxonomic categorization is meant to serve as a set of necessary inferencing capabilities that one would expect a competing NLI system to possess; thereby promoting more probing tasks along unexplored categories.

Existing categorization efforts have centred around informing feature creation in the pre-Transformer era, and model-specific error analysis in more recent times. Previously, (LoBue and Yates, 2011) enumerated the type of commonsense knowledge required for NLI. Among recent error analysis efforts, the GLUE diagnostic dataset (Wang et al., 2018), inference types for Adversarial NLI (Nie et al., 2019), the new CheckList (Ribeiro et al., 2020) system and the Stress Tests (Naik et al., 2018) are mentionworthy. As we attempted to group the categorizations in Nie et al. (2019) and Wang et al. (2018) into four high-level categories (lexical, syntactic, semantic, and pragmatic)<sup>2</sup>, we observe that there is a lack of consensus, non-uniformity and repetitiveness of these categories. For example, the *Tricky* label in Nie et al. (2019) groups examples that involve “wordplay, linguistic strategies such as syntactic transformations, or inferring writer intentions from contexts”; thereby spanning aspects of syntax and pragmatics. Similarly, *Reference and Names* requires both reasoning and knowledge. The GLUE diagnostic categories (Wang et al., 2018) does not include interesting reasoning categories such as *temporal*, and *spatial*. The stress types proposed by Naik et al. (2018) are specific to mostly lexical and some semantic corner cases. This is expected, as these categorizations are analysis-oriented and often dependent on the performance of a set of models in question. Here, we propose a taxonomic categorization that delineates

<sup>2</sup>Table provided in Appendix

a set of necessary uniform inferencing capabilities for the NLI task.

## 2.2 Taxonomic Categories: Definitions and Examples

In Figure 1, we present our taxonomic categorization. Our categorization is based on the following principles. First, we take a model-agnostic approach, where we work from the first principles to arrive at a set of basic inferencing processes that are required in NLI task. These include an unrestricted variety of linguistic and logical phenomena, and may require knowledge beyond text, thus providing us with the higher-level categories: *linguistic*, *logical* and *knowledge-based*. Second, we retain categories that are non-overlapping and sufficiently represented in NLI datasets. For example, for sub-categories under *linguistic*, we prune *semantics* because necessary aspects are covered by *logical* and *knowledge-based* categories. We omit specific aspects of *pragmatics* such as implicatures and pre-suppositions, as they are rarely observed in NLI datasets (Jeretic et al., 2020). Thirdly, we aim to list a set of necessary sub-categories. For example, for logical deduction sub-categories, we take inspiration from Davis and Marcus (2015), who list the commonsense reasoning categories where systems have seen success. Lastly, since we aim to employ non-experts for collecting annotations, we decide to restrict further sub-division wherever the definitions get complicated, or pre-suppose certain expertise; for example the *lexical* category is not sub-divided further (as followed in Wang et al. (2018)). Thus, we take a pragmatic approach that is theory neutral and does not warrant coverage of all reasoning tasks, though we do believe that the taxonomy is sufficiently deep and generic that allows systematic and meaningful analysis of NLI models with respect to their reasoning capabilities.

Next we define the categories. For a full set of examples, please see Table 1.

High-Level Categories: The **Linguistic** category represents NLI examples where the inference process to determine the entailment are internal to the provided text. We classify examples as **Logical** when the inference process may involve processes external to text, such as mapping words to percepts and reason with them (Sowa, 2010). **Knowledge-based** category represents examples where some form of external, domain or commonly assumed knowledge is required for inferencing.

**Linguistic** category is further sub-divided into *lexical*, *syntactic*, and *factivity*.

1. **Lexical**: This category captures P-H pairs where the text is almost the same apart from removal, addition or substitution of some lexical items. **Example**: P: Anakin was kind. H: Anakin was cruel.
2. **Syntactic**: *Syntactic* deals with examples where syntactic variations or paraphrases are essential to detecting entailment. **Example**: P: Anakin was an excellent pilot. H: The piloting skills of Anakin were excellent.
3. **Factivity**: Here the hypothesis contains an assumed fact from the premise, mostly an assumption about the existence of an entity or the occurrence of an action (inspired from Wang et al. (2018)). **Example**: P: Padme recognized that Anakin was intelligent. H: Anakin was intelligent.

Based on commonalities, **Logical** categories are grouped under “Connectives”. “Mathematical” and “Deduction”.

1. **Connectives (Negation, Boolean, Quantifiers, Conditionals, Comparatives)**: We group the logical categories *negation*, *boolean*, *quantifier*, *conditional* and *comparative* (Salvatore et al., 2019) under the “Connectives” label. *Negation* applies when P negates one (or more) of the facts in H. We apply the category *boolean* when P is a set of statements connected by *or*, *and* and H talks about one of the statements. *Quantifier* is applied when P or H requires understanding of words denoting existential or universal quantifiers. Similarly, *conditional* is applicable where P or H has conditional statements. If P (or H) compares entities via comparative phrases, then we label it as *comparative*. **Examples**: (*boolean* and *negation*) P: Jar Jar, R2D2 and Padme only visited Anakin’s house. H: Jar Jar Binks didn’t visit Anakin’s shop.
2. **Mathematical (Counting, Arithmetic)**: This group of categories is concerned with examples that require mathematical reasoning. For brevity, we concentrate on examples that require counting and simple arithmetic operations. However, we observed exceedingly low number of examples in this category group from our pilot study on SNLI and MNL1, and hence we remove these from our final annotations.
3. **Deductions (Relational, Spatial, Temporal, Causal, Coreference)**: Motivated by predicate logic, success of qualitative representation and

Taxonomic Category	MNLI Examples	Taxonomic Category	MNLI Examples
Lexical	P: so it's stayed cold for the <u>entire</u> week H: It has been cold for the <u>whole</u> week.	Relational	P: Actually, <u>my</u> sister wrote a story on it. H: <u>My sibling</u> created a story about it.
Syntactic	P: Those in Egypt, Libya, Iraq, and Yemen were eventually overthrown by secular nationalist revolutionaries. H: Secular nationalist revolutionaries eventually overthrew them in Egypt and Libya.	Spatial	P: At the eastern end of Back Lane and turning right, Nicholas Street becomes Patrick Street, and in St. Patrick's Close is St. Patrick's Cathedral. H: Nicholas Street becomes Patrick Street after turning left at the eastern end of Back Lane.
Factivity	P: The best place to view the spring azaleas is at the Azalea Festival in the last week of April at Tokyo's Nezu shrine. H: There is an Azalea Festival at the Nezu Shrine.	Temporal	P: See you Aug. 12, or soon thereafter, we hope. H: The person told not to come <u>until December</u> .
Negation	P: They <u>post</u> loads of newspaper articles–Yahoo! H: Yahoo <u>does not post</u> any articles from newspapers.	Causal	P: Acroseon the mountainside is another terrace on which imperial courtiers and dignitaries would sit while enjoying dance performances and music recitals on the <u>hondo's broad terrace</u> . H: There is a <u>terrace</u> where musicians play.
Boolean	P: According to contemporaneous notes, at 9:55 the Vice President was still on the phone with the President advising that three planes were missing <u>and</u> one had hit the Pentagon. H: The President called the Vice President to tell him the plane hit the Pentagon.	Coreference	P: A dozen minor wounds crossed his forearms and body. H: The grenade explosion left him with a lot of wounds.
Quantifier	P: <u>Some</u> travelers add Molokai and Lanai to their itineraries. H: <u>No one</u> decides to go to Molokai and Lanai.	World	P: In this respect, bringing Steve Jobs back to save Apple is like bringing Gen. H: Steve Jobs unretired in 2002.
Conditional	P: If the revenue is transferred to the General Fund, it is recognized as nonexchange revenue in the Government-wide consolidated financial statements. H: Revenue from the General Fund is not considered in financial statements	Taxonomic	P: Benson's action picture in Lucia in London (Chapter 8)- Georgie stepped on a beautiful <u> pansy</u> . H: Georgie crushed a beautiful <u> flower</u> in Chapter 8 of Lucia in London.
Comparative	P: Load time is divided into elemental and coverage related load time. H: The coverage related load time <u>is longer than</u> elemental.		

Table 1: For each category, we provide an example from the MNLI dataset. For a full set of synthetic examples and definitions, please look at appendix.

reasoning in dealing with temporal and spatial reasoning (Gabelaia et al., 2005), and causality (Pearl, 2009), we list `relational`, `temporal`, `spatial` and `causal` under “Deductions”. The `relational` reasoning stands for the requirement to perform deductive reasoning using relations present in text. `Spatial` (and `temporal`) denotes reasoning using spatial (and temporal) properties of objects represented in text. We also consider language-inspired reasoning categories such as co-reference resolution, which is known to often require event-understanding (Ng, 2017) beyond superficial cues. **Example:** (`relational`) P: The lamp was working properly. H: The lightbulb from the lamp was not functioning.

Lastly, we define two sub-categories under **Knowledge**, namely `world` and `taxonomic`.

1. **World:** Examples that require knowledge about named entities, knowledge about historical, current events; and domain-specific knowledge. **Example:** (`world`) P: Michelle Obama stayed in the White House during 2009-17. H: Michelle was living in the White House legally during 2009.

2. **Taxonomic:** Examples that require taxonomies and hierarchies. For example, `IsA`, `hasA`, `hasProperty` relations. **Example:** (`taxonomic`) P: Norman hated all musical instruments. H: Norman loves the piano.

Note that presence of a certain lexical trigger for a category (such as negation) does not warrant the labeling with the category, unless understanding of that concept is invoked in the deduction process.

### 3 TaxiNLI: Dataset Details

We present TaxiNLI, a dataset collected based on the principles and categorizations of the aforementioned taxonomy. We curate a subset of examples from MultiNLI (Williams et al., 2018) by sampling uniformly based on the entailment label and the domain. We then annotate this dataset with fine-grained category labels.

#### 3.1 Annotation Process

**Task Design** For large-scale data collection, our aim was to propose an annotation methodology that is relatively flexible in terms of annotator qualifications, and yet results in high quality annotations. To employ non-expert annotators, we designed a simplified guideline (questionnaire/interface) for the task, that does not pre-suppose expertise in language or logic. As an overhead, the guideline requires a few rounds of one-on-one training of the annotators. Because it is expensive to perform such rounds of training in most crowdsourcing platforms, we hire and individually train a few chosen annotators. Upon conducting the previously-discussed pilot study and using the given feedback,

we created a hierarchical questionnaire which first asked the annotator to do the NLI inference on the P-H pair, and then asked targeted questions to get the desired category annotations for the datapoints. The questionnaire is shared in the Appendix.

For the MNLI datapoints with ‘neutral’ gold labels, we realized, through observation and annotator feedback, that annotating the categories were difficult, as sometimes the hypotheses could not be connected well back to their premise. Hence, we created 2 questionnaires, one for the ‘entailment/contradiction’ examples, and one for ‘neutral’ examples. For the entailment/contradiction examples, We collected binary annotations for each of the 15 categories in our NLI taxonomy, for datapoints in MNLI which had ‘entailment’ or ‘contradiction’ as gold labels. To resolve this, for the ‘neutral’ examples we specifically asked them whether the premise and hypothesis were discussing 1) the same general topic (politics, geology, etc.), and if so, 2) had the same subject and/or object of discussion (Obama, Taj Mahal, etc.). If the response to 2) was ‘yes’, then they were asked to provide the category annotations as previously defined.

**Annotator Training/Testing** We first tested our two annotators by asking them to do inference on a set of randomly selected premise-hypothesis pairs from MultiNLI. This was to familiarize them with the inference task. After giving the category annotation task, we also continuously tested and trained the two annotators. After a set of datapoints were annotated, we reviewed and went through clarification and feedback sessions with the annotators to ensure they understood the task, and the categories, and what improvements were required. More details are provided in the Appendix.

### 3.1.1 Annotation Metrics

Here, we assess the individual annotator performance and inter-annotator agreement. Since, automated metrics for individual complex category annotations are hard to define, we use an indicative metric that matches the annotated inference label with the gold label, i.e., their **inference accuracy**. We also calculated inter-annotator agreement between the two annotators for an overlapping subset of 600 examples. For agreement, we use the Fleiss’ Kappa ( $\kappa$ ) (Fleiss, 1971). We also compute another simple statistic, namely the ‘IOU’ (Intersection-Over-Union) of categories per datapoint, defined as:  $IOU = \frac{1}{N} \sum_{i=1}^N \frac{|C_{A1}^i \cap C_{A2}^i|}{|C_{A1}^i \cup C_{A2}^i|}$ , where  $C_{A_i}^j$  are the




Figure 2: Inter-annotator Agreement (IAA) values between the two annotators, plotted category-wise.

category annotations for Annotator  $i$  for datapoint  $j$ , averaged over total datapoints  $N$ . Looking at the category-wise Fleiss’  $\kappa$  values in Fig. 2, we observe that there are promising levels of agreement in most of the categories except `syntactic`, `relational`, and `world`. We observe the average inference accuracy (86.7%) is high despite known issues in MNLI example ambiguity. Similarly, both the average Fleiss’  $\kappa$  (**0.226**) and the IOU metric (**0.241**) suggest an overall reasonable inter-annotator agreement.

### 3.2 Dataset Statistics

Each datapoint in TaxiNLI<sup>3</sup> consists of a premise-hypothesis pair, the entailment label, and binary annotations for 18 features. 15 features correspond to the 15 categories discussed in the taxonomy, and 3 additional features for the ‘neutral’ gold label datapoints based on same general topic, same subject, and same object. The statistics are listed in Tab. 2.

Total datapoints	10,071
Datapoints overlapping with MNLI	2343 (train) 7728 (dev)
Avg. datapoints per domain	1007.1
Datapoints per NLI label	3375 (C), 3201 (N), 3495 (E)
Avg. categories per datapoint	1.6
Neutral example characteristics	3087 (Same general topic), 2843 (Same object), 877 (Same subject)

Table 2: TaxiNLI Statistics

**Categorical Observations** From our annotations, we observe that inferring each MNLI ex-

<sup>3</sup>The dataset is available for download from <https://github.com/microsoft/TaxiNLI>.




Figure 3: The number of datapoints annotated with each category, split by the gold label of the datapoints. For taxonomic, the gold label split is 13(E)/13(C)/8(N).

ample requires about 2 categories. Fig. 3 shows the distribution of categories in the TaxiNLI dataset. We see that a large number of P-H pairs in MNLi require lexical and syntactic knowledge to make an inference; whereas the challenges of relational, spatial, and taxonomic for inference are not adequately represented. There is a large proportion of examples in the syntactic category which have the ‘entailment’ label, and a large proportion of negation examples have the ‘contradiction’ label. Additionally, many ‘neutral’ examples were classified as requiring causal knowledge. The feedback session with annotators revealed that there were many ‘neutral’ examples where the hypothesis was essentially an unverifiable intent or detail of a certain action mentioned in the premise. An example is “**P:** Another influence was that patrician politician Franklin Roosevelt, who was, like John D. Rockefeller, the focus of Nelson’s relentless sycophancy and black-belt bureaucratic infighting. **H:** Nelson targeted Roosevelt in order to gain political favor.”.

**Categorical Correlations** Fig. 4 shows correlations among categories in our dataset. We observe that most categories show weak correlation in the MNLi dataset, hinting at a possible independence of categories with respect to each other. Relatively stronger positive correlations are seen between boolean-quantifier, and boolean-comparative categories. We specifically looked at the genre-wise split of datapoints containing boolean-quantifier and saw that

nearly 25% of them came from the ‘telephone’ genre of MNLi. An example is “**P:** have that well and it doesn’t seem like very many people uh are really i mean there’s a lot of people that are on death row but there’s not very many people that actually um do get killed **H:** Most people on death row end up living out their lives awaiting execution.”. Factivity, on the other hand, is negatively correlated with almost all the other categories, except world, which means P-H pairs labeled with factivity typically have no other categories marked.

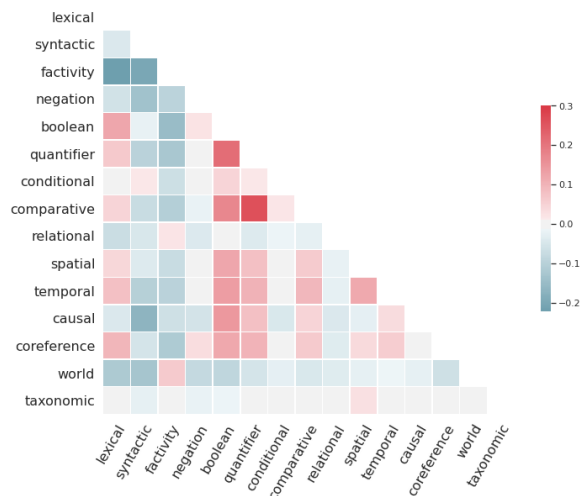


Figure 4: The correlation matrix between the taxonomic categories.

## 4 (Re)Evaluation of SOTA Models

We re-evaluate two Transformer-based and two standard baseline machine learning models on TAXINLI, under the lens of the taxonomic categories. As baselines, we choose BERT-base (Devlin et al., 2019), and RoBERTa-large (Liu et al., 2019b) as two state-of-the-art NLI systems. For our experiments, we use the pre-trained BERT-base and RoBERTa models from HuggingFace’s Transformers implementation (Wolf et al., 2019). As pre-Transformer baselines, we use the bidirectional LSTM-based Enhanced Sequential Inference model (ESIM) (Chen et al., 2017). We also train a Naive Bayes (NB) model using bag-of-words features for the P-H pairs after removing stop words<sup>4</sup>.

### 4.1 TAXINLI Error Analysis

We report the NLI task accuracy of the baseline systems on the MNLi validations sets in Table 3.

<sup>4</sup>Using NLTK’s RTEFeatureExtractor

The systems are fine-tuned on the MNLI training set using the procedures followed in Devlin et al. (2019); Liu et al. (2019b); Chen et al. (2017).

MNLI-dev	NB	ESIM	BERT <sub>BASE</sub>	RoBERTa <sub>LARGE</sub>
Matched	51.46	72.3	84.7	92.3
Mismatched	52.31	72.1	84.8	90.0

Table 3: MNLI-validation set accuracy.

We evaluate the systems on a total of 7.7k examples, which are in the intersection of TAXINLI and the validation sets of MNLI.

Figure 5 shows for each category  $c_i$ , the normalized frequency for a model predicting an NLI example of that category accurately, i.e.,  $\frac{\#c_i=1, \text{correct}=1}{\#c_i=1}$ . We observe that compared to NB, the improvements in BERT have been higher in *lexical*, *syntactic* categories compared to others. Improvements in ESIM compared to NB show a very similar trend, and show for knowledge categories the improvements are negligible. ESIM shows largest improvement on *negation*.




Figure 5: The normalized frequency of the systems predicting an example correctly, provided a category.

## 4.2 Factor Analysis

In order to quantify the precise influence of the category labels on the prediction of the NLI models, we probe into indicators and confounding factors using two methods: linear discriminant analysis (LDA) and logistic regression (LR). We use indicators for each category (0 or 1) and for two potential confounding variables (lengths of P,H), to model the correctness of prediction of the NLI system. The coefficients of these analyses on BERT are shown in Fig 6. The values for RoBERTa follow a similar trend, and are presented in the appendix. We see that presence of certain taxonomic categories strongly influence the correctness




Figure 6: Coefficients obtained through Linear Discriminant Analysis (LDA) and Logistic Regression (LR) to model the correctness of NLI prediction by BERT, given taxonomy categories and possible confound variables. Significant LR coefficients: *syntactic\*\**, *negation\*\*\**, *boolean\**, *causal\*\*\**, *world\*\*\**, *Length2\*\**; where  $p$  value is smaller than: 0.001\*\*\*, 0.01\*\*, 0.05\*.

of prediction. As we found in the analysis presented in Sec. 4, we observe that *syntactic*, *negation*, and *spatial* categories are strong indicators of correctness of prediction. On the other hand, *conditional*, *relational*, *causal*, *coreference* are harder to predict accurately. Sentence length does not play a significant role.

We also make an observation for categories such as *lexical*, *syntactic*, where the proportion of a single NLI label is high, also correlated with a high prediction accuracy (Fig. 6).

## 5 Discussion

**Visual Analysis** Section 4 paints a thorough picture by analysing the fine-grained capabilities of SOTA NLI systems at a *behavioral*<sup>5</sup> level. Whereas we can say the systems are lacking in certain aspects despite their high overall performance, it naturally also raises questions at the *understanding* level: 1) Is there any implicit knowledge acquired by the NLI-finetuned systems about the *kinds* of reasoning required in the inference task? 2) If not, do the systems simply lack the understanding of what kind of reasoning is required per example, or despite understanding that, are unable to do the

<sup>5</sup>Similar to social sciences, as a black-box system

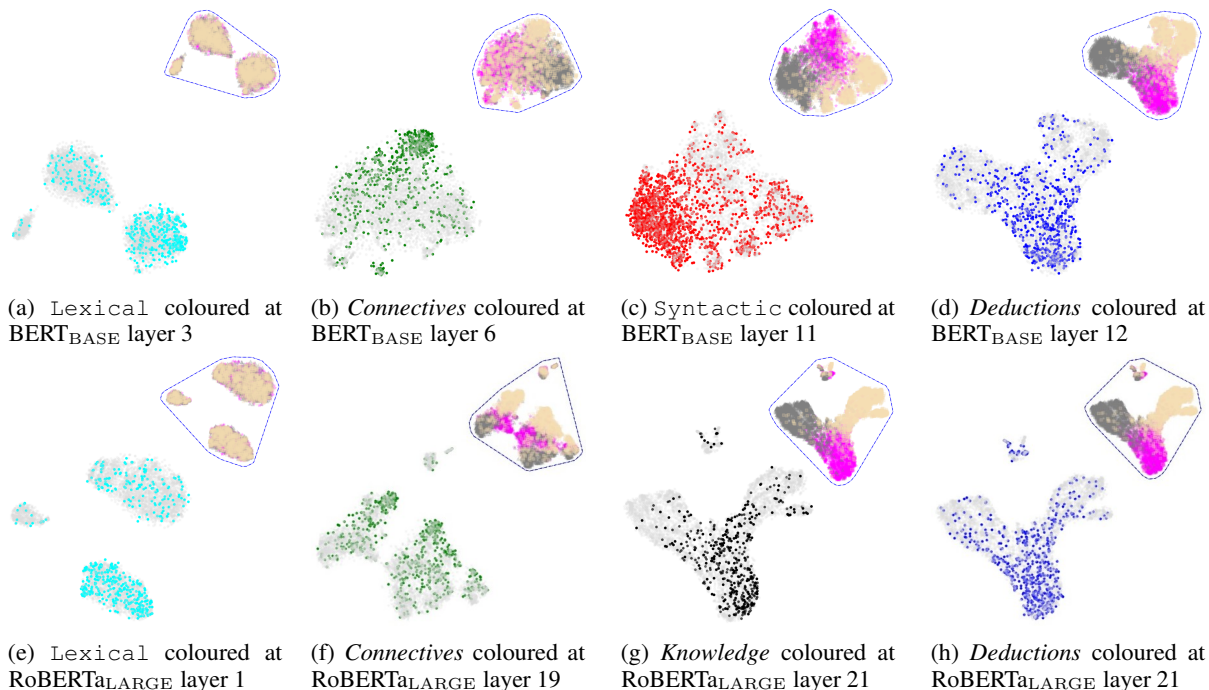


Figure 7: Layer-wise 2D t-SNE plots of pooled contextualized embeddings of the TaxiNLI examples extracted from BERT and RoBERTa finetuned on MNLI with no retraining on the taxonomic labels. Color codes represent taxonomic categories (inset shows NLI categories: **entailment** (■), **neutral** (■), **contradiction** (■)) and their combinations. Only combinations of up to two categories are included for brevity.

reasoning? 3) Can we make an argument for future work and model architecture that can more consciously use this information?

In light of recent probing task literature (Tenney et al., 2019a; Jawahar et al., 2019; Liu et al., 2019a), we specifically investigate whether representations of examples cluster meaningfully into taxonomic categories relevant to the reasoning required for NLI. We use the t-SNE (Maaten and Hinton, 2008) algorithm to visualize pooled contextualized representations of NLI examples, under the lens of our taxonomy. For an NLI example, we construct the embeddings at a Transformer layer by max-pooling hidden states over all input token positions concatenated with the [CLS] token (typically used for classification tasks) representation. The resulting visualizations (Fig. 7) reveal definitive patterns of clustering by taxonomic categories. The earliest separation is observed for the *lexical* category, at layer 3 in BERT (and layer 1 in RoBERTa), much before any other categories are realized. At layers 6 in BERT, and 19 in RoBERTa, about the same time as clustering by NLI label is seen, the *connectives* cluster is revealed. The *deductions* (see Sec. 2), and *syntactic* categories are revealed in later layers (layer 11 and 12 in BERT and layer 21 in RoBERTa). The *knowledge* cate-

gory is revealed more prominently in RoBERTa at layer 21, while BERT does not seem to show such a cluster. By the last few layers, separation into most categories becomes apparent. This means, along various layers of a NLI finetuned language model, taxonomic information is implicitly captured. Despite this, as discussed in the previous sections, SOTA models seem to be deficient in some of the categories—certain categories remain harder to perform inference on. In the latter layers, the separation along taxonomic categories also corresponds strongly with separation along NLI labels. For instance, in Fig. 7 (c), the examples categorized as *syntactic* almost entirely lie in the *entailment* cloud, which matches our intuition based on the statistics in Fig. 3.

The layer-wise separation of examples by taxonomy raises an interesting possibility to motivate model architectures that may attempt to use its discriminative power to identify such taxonomic categories, for specialized treatment to examples requiring certain reasoning capabilities.

**Recasting:** The under-representation of certain categories in the MNLI dataset raises a need for more balanced data collection. A possible alternative is to build recast diagnostic datasets for each category, and create probing tasks. Some datasets



(Zhang et al., 2019; Richardson et al., 2020) can be recast to the `syntactic` and **Logical** categories respectively, as their data creation aligns with our category definitions. However, most categories lack such aligned synthetic data, and crowd-sourced data would require manual annotation as above. This poses an avenue for future work.

## 6 Conclusion

To bridge the gap between accuracy-led performance measurement and linguistic analysis of state-of-the-art NLI systems, we propose a taxonomic categorization of necessary inferencing capabilities for the NLI task, and a re-evaluation framework of systems on a re-annotated NLI dataset using this categorization; which underscores the reasoning categories that current systems struggle with.

We would like to emphasize that unlike the case with challenge and adversarial datasets, TAXINLI re-annotates samples from existing NLI datasets which the SOTA models have been exposed to. Therefore, a lower accuracy in certain taxonomic categories in this case cannot be simply explained away by the “lack of data” and “unnatural distribution” arguments.

## Acknowledgements

We gratefully acknowledge Sandipan Dandapat and Rohit Nargunde for their help regarding annotations. We would like to thank the anonymous reviewers for their insightful comments.

## 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.
- Samuel Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642.
- Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, Hui Jiang, and Diana Inkpen. 2017. [Enhanced LSTM for natural language inference](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1657–1668, Vancouver, Canada. Association for Computational Linguistics.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2005. The pascal recognising textual entailment challenge. In *Machine Learning Challenges Workshop*, pages 177–190. Springer.
- Ernest Davis and Gary Marcus. 2015. Commonsense reasoning and commonsense knowledge in artificial intelligence. *Communications of the ACM*, 58(9):92–103.
- 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.
- JL Fleiss. 1971. [Measuring nominal scale agreement among many raters](#). *Psychological bulletin*, 76(5):378–382.
- David Gabelaia, Roman Kontchakov, Agi Kurucz, Frank Wolter, and Michael Zakharyashev. 2005. Combining spatial and temporal logics: expressiveness vs. complexity. *Journal of Artificial Intelligence Research*, 23:167–243.
- Yoav Goldberg. 2019. Assessing bert’s syntactic abilities. *arXiv preprint arXiv:1901.05287*.
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. [What does BERT learn about the structure of language?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3651–3657, Florence, Italy. Association for Computational Linguistics.
- Paloma Jeretic, Alex Warstadt, Suvrat Bhooshan, and Adina Williams. 2020. [Are natural language inference models IMPPRESsive? Learning IMPLIcature and PRESupposition](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8690–8705, Online. Association for Computational Linguistics.
- Divyansh Kaushik, Eduard Hovy, and Zachary C. Lipton. 2019. [Learning the difference that makes a difference with counterfactually-augmented data](#).
- Tal Linzen. 2020. [How can we accelerate progress towards human-like linguistic generalization?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5210–5217, Online. Association for Computational Linguistics.
- Nelson F Liu, Matt Gardner, Yonatan Belinkov, Matthew E Peters, and Noah A Smith. 2019a. Linguistic knowledge and transferability of contextual representations. *arXiv preprint arXiv:1903.08855*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b.

- Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Peter LoBue and Alexander Yates. 2011. Types of common-sense knowledge needed for recognizing textual entailment. In *Proceedings of the 49th annual meeting of the association for computational linguistics: human language technologies*, pages 329–334.
- Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Aakanksha Naik, Abhilasha Ravichander, Norman Sadeh, Carolyn Rose, and Graham Neubig. 2018. [Stress test evaluation for natural language inference](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2340–2353, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Vincent Ng. 2017. Machine learning for entity coreference resolution: A retrospective look at two decades of research. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2019. Adversarial nli: A new benchmark for natural language understanding. *arXiv preprint arXiv:1910.14599*.
- Judea Pearl. 2009. *Causality*. Cambridge university press.
- Adam Poliak, Aparajita Haldar, Rachel Rudinger, J. Edward Hu, Ellie Pavlick, Aaron Steven White, and Benjamin Van Durme. 2018. [Collecting diverse natural language inference problems for sentence representation evaluation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 67–81, Brussels, Belgium. Association for Computational Linguistics.
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. [Beyond accuracy: Behavioral testing of NLP models with CheckList](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4902–4912, Online. Association for Computational Linguistics.
- Kyle Richardson, Hai Hu, Lawrence Moss, and Ashish Sabharwal. 2020. [Probing natural language inference models through semantic fragments](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34:8713–8721.
- Kyle Richardson, Hai Na Hu, Lawrence S. Moss, and Ashish Sabharwal. 2019. Probing natural language inference models through semantic fragments. *ArXiv*, abs/1909.07521.
- Felipe Salvatore, Marcelo Finger, and Roberto Hirata Jr. 2019. A logical-based corpus for cross-lingual evaluation. In *Proceedings of the 2nd Workshop on Deep Learning Approaches for Low-Resource NLP (DeepLo 2019)*, pages 22–30.
- John F Sowa. 2010. The role of logic and ontology in language and reasoning. In *Theory and applications of ontology: philosophical perspectives*, pages 231–263. Springer.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019a. Bert rediscovers the classical nlp pipeline. *arXiv preprint arXiv:1905.05950*.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel Bowman, Dipanjan Das, et al. 2019b. What do you learn from context? probing for sentence structure in contextualized word representations. In *7th International Conference on Learning Representations, ICLR 2019*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *EMNLP 2018*, page 353.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics.
- L. Wittgenstein. 1922. *Tractatus logico-philosophicus*. London: Routledge, 1981.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.
- Yuan Zhang, Jason Baldridge, and Luheng He. 2019. [PAWS: Paraphrase adversaries from word scrambling](#). 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 1298–1308, Minneapolis, Minnesota. Association for Computational Linguistics.

	ANLI Inference Types	GLUE Diagnostic
Lexical	Standard Inference Lexical Inference	Lexical Entailment Morphological Negation Factivity, Redundancy
Syntactic	Tricky	Syntactic Ambiguity, Prepositional Phrase Alternations: Active/Passive, Genitives/Partitives, Nominalization, Datives
Semantic	Standard Inference	Propositional Structure, Intersectivity Quantifiers, Restrictivity, Quantification
	Numerical, Quantitative	Counting
	Reference and Names	Symmetry/Collectivity, Coreference, Richer Logical Structures
	Reference and Names Reasoning about Facts	Named Entities Knowledge and Commonsense
Pragmatics	Tricky	Ellipsis/Implicits

Table 4: We show existing NLI error-analysis categorizations proposed by the recent papers, and group them in higher-level categories.

## A Other Categorizations

## B Bayesian Estimate of Correctness Correlation with Categories

Since, examples are annotated with multiple categories, we capture the dependencies by defining a Bayesian Network (BN) where, each category (a boolean random variable) has a directed edge to the *correct* node (representing correctness of prediction)<sup>6</sup>. We learn the parameters by fitting this BN to the observed data. In Figure 9, we see from the Bayesian estimate again that, the improvements by BERT in categories such as *relational* reasoning has been low. It also shows, that there is a sharp decrease in accuracy for examples requiring the use of *taxonomic* knowledge. However, RoBERTa improves over NB and ESIM by large margins, albeit non-uniformly.

## C Factor Analysis of Correctness of Prediction by RoBERTa

In Fig. 9, we show the results of Linear Discriminant Analysis (LDA) and Linear Regression (LR) results for RoBERTa predictions. A very similar trend as BERT can be seen here as well.

<sup>6</sup>Additionally, we attempted to learn a Bayes Net from the data using `bnlearn` package. But the limited number of observations yield non-intuitive results.




Figure 8: We show a Bayesian Estimate of  $P(\text{correct} = 1 | \text{category} = 1)$  for different systems.

## D Annotation Questionnaire

Our annotation process went through several steps of refinement and improvement. We started with the most basic annotation flow, which was to have a manual which defines each taxonomic category in detail, and then have the annotator mark for each category. For the pilot study, we took roughly 300 examples from MNLI and asked an initial annotator to annotate. The feedback was the following:

- The manual describing each taxonomic category had a lot of information and took time to understand and digest.
- It was difficult to keep referring to the guide, although after sufficient examples, it became easier.
- There was confusion and ambiguity about the definitions, and the annotator interpreted the definitions differently than what we intended.
- Figuring out the categorical annotations for neutral examples was a challenge, as sometimes the topic or subject of what the hypothesis was discussing was separate from what the premise was discussing.

Through the analysis of these annotations, we also observed that some of the initial categories we had were either exceedingly underrepresented in the MNLI dataset, or were consistently confused with others. Thus, we revised the set of categories, setting more distinct boundaries, and ensuring independence of categories. We revised the questionnaire into a hierarchical 'if-else' multi-choice design. The questionnaire is structured as follows:

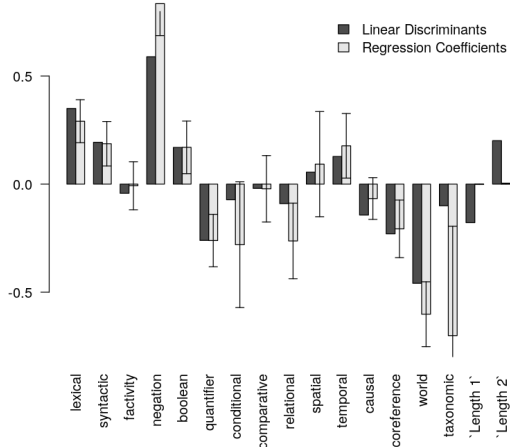


Figure 9: Coefficients obtained through Linear Discriminant Analysis (LDA) and Logistic Regression (LR) to model the correctness of NLI prediction by RoBERTa, given taxonomy categories and possible confound variables. Significant LR coefficients: `lexical**`, `negation***`, `world***`, `quantifier*`, `boolean*`, `Length1*`; where  $p$  value is smaller than: 0.001\*\*\*, 0.01\*\*, 0.05\*.

### D.0.1 Questionnaire 1

“We present you with a set of statements. Statement 1 (S1) is the truth and context. Statement 2 (S2) is a claim/hypothesis. The task is to evaluate statement 2 as true, false, can’t say.

S1: ... S2: ...

1. Can you evaluate S2 by just using the information/context given in S1? Or do you require knowledge from external documents, say history books, news articles, science books, etc.?
  - (a) Need more information
  - (b) Do not need more information
2. If yes, what kind of information did you require? (More than one answer can be ticked)
  - (a) Knowledge about certain facts from say history books, news articles, tech magazines, etc.? This is also knowledge about named entities (e.g. Obama, Taj Mahal, New York etc.). E.g:
    - S1: Barack Obama lived in the White House during 2009-17.
    - S2: Barack Obama was the President in 2009.

This is TRUE and requires external knowledge that US presidents live in the White House.

- (b) Knowledge about taxonomies and hierarchies. A few examples are animal groups (snakes are reptiles), currencies (dollar is a currency), types of activities (football is a sport, sport is an activity). Basically, how a common noun (snake) belongs to a class (say reptiles), which can belong to yet another class (animals). Do not select this if the name of one object belongs (or is a substring) of the other class of objects (e.g. green snake is a snake isn’t part of this category), or if the names of the objects are pronouns (e.g. Barack Obama - president and related examples are part of category 2a, not this one). E.g:
  - S1: Norman hated all musical instruments.
  - S2: Norman loves the piano.

This is FALSE and requires external knowledge that a piano belongs to the class of instruments, hence Norman cannot love the piano.

- (c) No extra knowledge required
3. Using just the information from S1, and given that you have the required knowledge from the above question, did you have to use some reasoning to figure out the answer, or did you just need the knowledge of words and paraphrasing, or both? (More than one can be ticked)
    - (a) Some reasoning was required, which wasn’t explicitly written down in S1, but was implicitly understood.
    - (b) Knowledge of words (e.g. synonyms, antonyms), and recognizing paraphrases. The information I needed was explicitly written down in S1.
  4. What kind of reasoning was required (if applicable) (More than one can be ticked)?
    - (a) You needed reasoning about relations in S1. You observed that there were objects/entities in S1 and there were explicit mentions of how they were related (e.g. Jack and his son went to the circus), and you used your reasoning about the nature of those relations to arrive to the

answer (e.g. Jack and a stranger went to the circus is FALSE) E.g:

- S1: Jack and his son went to the circus.
- S2: Jack and a stranger went to the circus.

This is FALSE, but you need to reason that S1 contains the relation “father of” between Jack and some person X (who is his son). S2 is false because if X is a stranger, it cannot be Jack’s son.

- (b) You needed reasoning about spatial setup. S1 contained information about relative locations of objects/entities (e.g. Jack was on the right of Jim, and Jim was on the right of John) and you needed to reason about how objects/entities were located, when it wasn’t explicitly stated (e.g. John was on the left of Jack is TRUE).
- (c) You needed reasoning about time intervals, duration, or temporal reasoning. S1 contained information about events and time information (e.g. Jack went to the shop from 8:00am to 10:00am), and you needed to reason about the timing to arrive at the answer (e.g. Jack was at the shop at 9:30am is TRUE).
- (d) You needed reasoning about cause, effect, and intent behind it. S1 contained information about an event or an action (e.g. Jack was hurt). S2 contains either a cause or an intent, and you needed to reason whether S1 and S2 were possible cause/effect or intent/effect pairs (e.g. Jack was hit by a car is CAN’T SAY)
- S1: X shot Y
  - S2: Y is hurt

This is TRUE, but you needed to reason that upon getting shot, Y should get hurt.

- (e) You needed to be able to reason about who is being referred to in the text. Better demonstrated via example:
- S1: Jane didn’t visit Janette because she didn’t want to speak with her.
  - S2: Jane didn’t want to speak with Janette.

This is TRUE. ‘She’ refers to Jane, and you needed to reason about that to get the answer.

5. Did you need logical reasoning? This applies if S1 and/or S2 consist of statements connected by logical connective words (and, or, not, every, some, only, either, neither, etc. or any synonyms of these words). These connectives were important for arriving at the answer. If so, which connectives? (Can tick multiple choices)

(a) Negation (not, no, incapable etc.), where S2 negates one of the facts in S1. E.g:

- S1: Laurie has visited Nephi, Marion has only visited Calistoga.
- S2: Laurie didn’t visit Nephi.

This is FALSE, and S2 is a negation of the first statement in S1.

(b) Boolean (or, and), where S1 is a set of statements connected by Or and AND, and S2 talks about one or more of these statements. E.g:

- S1: Jar Jar Binks, R2D2 and Padme only visited Anakin’s house.
- S2: Jar Jar Binks didn’t visit Anakin’s shop.

This is TRUE, and S1 is connected by ‘and’ statements for three entities who visited Anakin’s house. S2 talks about the sub-statement “Jar Jar Binks only visited Anakin’s house” in the ‘and’ connective, and it is true because Jar Jar Binks didn’t visit anywhere else.

(c) Quantifier (every, some, at least, at most, etc.), where S1 and S2 contain the use of these terms.

- S1: Everyone visited Anakin’s home.
- S2: Padme didn’t visit Anakin’s home.

This is FALSE, with S1 containing the quantifier ‘everyone’, and S2 stating that someone, Padme, didn’t visit.

(d) Conditionals (if-else, if-then, etc.), where S1 has if-then, if-else statements or similar.

- S1: Francisco has visited Potsdam and if Francisco has visited Potsdam then Tyrone has visited Pampa.
- S2: Tyrone has visited Pampa

This is TRUE, since there is a if-then condition in S1, and it is satisfied to make S2 true.

(e) Comparatives (e.g. as tall as, taller than, faster than, etc.) where S1 compares entities via these comparative phrases, and S2 needs knowledge about the comparisons.

- S1: John is taller than Gordon and Erik, and Mitchell is as tall as John
- S2: Gordon is taller than Mitchell.

This is FALSE. This are comparative statements “Is taller than” in S1, and S2 needs logical reasoning on who is taller than whom to get S2. In addition to this, this also needs knowledge of Boolean due to the presence of the ‘and’ in S1.

6. Finally, can you describe some word/phrase (explicitly written) properties of S1 and S2 which helped to arrive to the answer? (more than one can be ticked)

(a) S1 and S2 were almost the same, apart from the removal, addition, or substitution of a few words. If substituted, the words were synonyms or antonyms. E.g:

- S1: Anakin Skywalker was compassionate.
- S2: Anakin Skywalker was cruel.

This is FALSE, with S1 and S2 being very similarly framed statements, with the substitution of a word for it’s antonym. Thus it belongs to this category.

(b) S1 and S2 were paraphrases of each other. S2 is a paraphrase of S1 or a certain part of information mentioned in S1

- S1: Anakin was an excellent pilot.
- S2: The piloting skills of Anakin were excellent.

This is TRUE, and S1 and S2 being paraphrases of one another. Also, to note, if ‘excellent’ in S2 were replaced by ‘terrible’, it would still fit this category, but would also fit category 6a, since it would be a paraphrase with a swapped word.

(c) S2 contains an assumed fact from S1, mostly an assumption about the existence or the occurrence of an action.

- S1: Anakin found the Death Star.
- S2: The Death Star exists.

This is TRUE. The Death Star exists if Anakin has found it, thus S1 makes the assumption that it exists.

- S1: James was happy that his plane could fly.
- S2: His plane couldn’t fly.

This is FALSE. Since James was happy that the plane flew (S1 makes the assumption that it happened), it is FALSE that his plane couldn’t fly since it happened.”

The above questionnaire was given along with premise-hypothesis pairs having the gold label of ‘entailment’ or ‘contradiction’ . However, to prevent biasing the annotator, we allowed them to choose ‘neutral’ (CAN’T SAY) as well.

### D.0.2 Questionnaire 2

This questionnaire was given to the annotators after they had done a sufficient number of ‘entailment/contradiction’ samples using Questionnaire 1. For Questionnaire 2, annotators were told that the datapoints were ‘neutral’, and asked them to first answer these 3 questions:

“Given S1, there isn’t enough information to decide whether S2 is TRUE or FALSE. Please answer the following questions for each datapoint which has been annotated as CAN’T SAY.

1. Are S1 and S2 talking about the same general topic (e.g. sports, politics, religion)? [Yes/No]
2. If yes, are S1 and S2 talking about the same subject? (e.g. S1: Obama was the president of USA, S2: Obama was a nice guy, subject of sentence is Obama) [Yes/No]
3. If yes, are S1 and S2 talking about the same objects of discussion? (e.g. S1: Obama lived in the White House often, S2: Obama said the White House was huge. Here both subject (Obama) and object (White House) of discussion are the same)
4. If yes, what kinds of information are required which you used, and what kinds of information are missing? If no, what kinds of information are required which you used, and what kinds of information are missing? ”

Upon answering the above, if the answer to the second question was yes, then they proceeded with the category annotation, else they moved on to the next question. This helped eliminate the random hypotheses.

### D.0.3 Annotator Feedback

We received a lot of important feedback from our annotators during the clarification and training sessions. They are listed below:

- Many premise sentences seems out of place, and the context is still insufficient many times. As a result, the hypothesis also introduces ambiguity, making the process a bit tricky.
- There were cases where a certain name of an entity in the premise is switched for something else in the hypothesis. This created some confusion because it fell somewhere between lexical and coreference (according to the annotator).
- Another confusion arose from the quantifier category, where initially the name of the category led the annotators to believe that it referred to not just what we described (e.g. some, all), but quantities (say 5000 in the premise was swapped with 2000 in the hypothesis). This was again a middle ground between lexical and quantifier.
- Many of the premises contained incoherent, difficult to understand sentences. A lot of premises (which we later found to be from the telephone category), contained many filler words (uh, uhm, etc.) which made comprehension difficult.
- Another issue lies with an implicit rigidity of the annotation process using just the questionnaire. The targeted questions were written so as to allow annotators to generalize and apply intuitive principles along those thought lines that we try to demarcate via the questions. We wanted to prevent them completely. However, as annotators have not been exposed to the exact intentions behind the annotation (so as to prevent bias), they followed the questionnaire strictly, and did not always generalize until subsequent training/clarification sessions where we encouraged them to generalize. However, the implicit rigidity still impacts the annotations to some extent, although mitigated to a large level by the training. This remains a challenge due to the tradeoff between open interpretation of the task, as well as a rigidity of desired annotations which stem from an analysis perspective (from our side).

- Idiomatic references, metaphors, and common phrases were also a source of confusion, and although to some extent were marked as world knowledge, did leave annotators unsure about where to place them.

The above feedback only strengthened our belief in an iterative training system for a complicated task such as this. It also sheds light on how difficult a task like this is to crowdsource.