

LawngNLI: A Long-Premise Benchmark for In-Domain Generalization from Short to Long Contexts and for Implication-Based Retrieval

William Bruno and Dan Roth

University of Pennsylvania

{wwbruno, danroth}@seas.upenn.edu

Abstract

Natural language inference has trended toward studying contexts beyond the sentence level. An important application area is law: past cases often do not foretell how they apply to new situations and implications must be inferred. This paper introduces LawngNLI, constructed from U.S. legal opinions with automatic labels with high human-validated accuracy. Premises are long and multigranular. Experiments show two use cases. First, LawngNLI can benchmark for in-domain generalization from short to long contexts. It has remained unclear if large-scale long-premise NLI datasets actually need to be constructed: near-top performance on long premises could be achievable by fine-tuning using short premises. Without multigranularity, benchmarks cannot distinguish lack of fine-tuning on long premises versus domain shift between short and long datasets. In contrast, our long and short premises share the same examples and domain. Models fine-tuned using several past NLI datasets and/or our short premises fall short of top performance on our long premises. So for at least certain domains (such as ours), large-scale long-premise datasets are needed. Second, LawngNLI can benchmark for implication-based retrieval. Queries are entailed or contradicted by target documents, allowing users to move between arguments and evidence. Leading retrieval models perform reasonably zero shot on a LawngNLI-derived retrieval task. We compare different systems for re-ranking, including lexical overlap and cross-encoders fine-tuned using a modified LawngNLI or past NLI datasets. LawngNLI can train and test systems for implication-based case retrieval and argumentation.

1 Introduction

This work proposes a new natural language inference (NLI) benchmark LawngNLI constructed from U.S. legal opinions via the Caselaw Access

Project ([The President and Fellows of Harvard University, 2018](#)) that have been largely cleaned of inline citations in order to read more naturally.¹ It follows the general three-label NLI formulation: given a pair of texts (premise and hypothesis), the goal is to predict the label of whether the premise entails, is neutral toward, or contradicts the hypothesis. LawngNLI’s premises are especially long and are multigranular. Its automatic labels derive from the dataset construction using (negation-based) contradiction and (similarity-based) neutralization algorithms. These labels exhibit an accuracy of 88.8% (94.7% for high-confidence human labels) on a subset with human-validated gold labels. Examples are derived from actual inferences from past cases that judges wrote to apply usefully to new situations, relying on their expertise while avoiding crowdsourced labeling.

We conduct two sets of experiments on LawngNLI. First, regarding in-domain generalization to long contexts, we compare top long-sequence and short-sequence NLI models on LawngNLI. Models are transferred with and without intermediate fine-tuning on existing NLI benchmarks. We find that absent fine-tuning directly, models fall substantially short of top performance on our long premises. This continues to hold true when we leverage LawngNLI’s multigranularity to control for domain shift, by further fine-tuning models on our own short premises. Top performance on our long premises is by short-sequence models prepended with a standard retrieval method

¹Code for obtaining LawngNLI and unfiltered-LawngNLI2 to be released at http://cogcomp.org/page/publication_view/990. LawngNLI contains about 140 thousand twinned examples. Our script and coded ID files should reconstruct LawngNLI using 1 CPU without GPUs in under 1 day (Intel Xeon Gold 6130 CPU @ 2.10GHz). A dataset from an earlier stage in construction process, unfiltered-LawngNLI2 (not studied in the paper), contains about 10.3 million untwinned candidate examples. Another script should reconstruct unfiltered-LawngNLI2 without GPUs by multiprocessing (about 3 days or fewer with 35 CPUs).

(BM25 (Robertson and Zaragoza, 2009)) filtering across each premise. Absent their fine-tuning that uses long premises as inputs, however, these models would underperform on long premises at inference time. We provide evidence that this gap is likely robust to any deviations between the automatic and gold labels. Thus a large-scale long-premise dataset like LawngNLI is needed.

Second, regarding legal retrieval systems, we conduct a comparison of leading models by Recall@k for selecting target cases via implication/NLI-based retrieval (when a user provides arguments entailed or contradicted by target cases as queries), in our domain. Implication-based retrieval is an underexplored subtask (see, e.g., Schuster et al., 2022). Our zero-shot panel comprises Sentence-Transformers (Reimers and Gurevych, 2021) lightweight bi-encoders pre-trained using short-sequence NLI, user web queries, and general semantic relatedness. While baseline models transfer reasonably to our domain zero shot, we further compare re-ranking with models fine-tuned using several previous NLI datasets along with an adjusted retrieval version of LawngNLI. Future improvements can put more evidence within the range of human users’ cognitive reach beyond the top result, including for law which could help make legal work more affordable and equitably accessible (see Section 6).

Overall, our main contributions are:

- (1) A new NLI benchmark with multigranular premises much longer than in most existing NLI benchmarks across percentiles (see Table 2).
- (2) A benchmarking of models’ ability to generalize from short context to long context on the same domain and examples. It shows that (in our domain at minimum) short-premise NLI models fall substantially short of top performance on long-premise NLI at inference time, unless a large-scale long-premise dataset is created to fine-tune on.
- (3) A benchmarking of leading retrieval models on case retrieval with entailed or contradicted arguments as queries, comparing lexical overlap and fine-tuning using a modified LawngNLI or previous NLI datasets. LawngNLI provides a benchmark for future implication-based case retrieval systems.

2 Significance for NLI and for Law

Regarding in-domain generalization to long contexts, our work stands within a fast-growing research area on how models can learn to reason

over long text. Benchmarks for NLI, or Recognizing Textual Entailment (RTE), stretch back to Dagan et al. (2005). Recently, various “efficient” Transformer architectures have been proposed to address the obstacle of quadratic self-attention complexity in scaling to long sequences (Tay et al., 2020). LawngNLI’s long premises frequently exceed the limits of such long-sequence models with efficient attention mechanisms, although on our long premises not exceeding those limits, the long-sequence models included here are outperformed by short-sequence models both with and without filtering of premise paragraphs by relevance to the hypotheses. Most existing NLI benchmarks, meanwhile, contain largely short premises. The recent ContractNLI (Koreeda and Manning, 2021) is an exception, containing premises with lengths similar to ours though with a much smaller number of examples in a different domain (607 contracts as premises, paired with each in a common set of 17 shared hypotheses).²

Regarding legal retrieval systems, law is an important area where humans perform long-context NLI in practice. The core of legal advocacy is articulating what existing applicable law (cases, legislation, regulations, etc.) implies for a new situation. Practitioners must move between case text and the entailed and contradicted arguments that they aim to support or counter. Automatic systems are developed to help find relevant precedents, which requires filtering the millions of U.S. cases.

3 LawngNLI Dataset

We construct LawngNLI beginning with all citations with parentheticals in U.S. state and federal case opinions, via the Caselaw Access Project (The President and Fellows of Harvard University, 2018). Its hypotheses derive from the actual inferences from past cases applied by judges to decide between competing parties’ positions across a large cross-section of U.S. opinions. Thus we hope its examples capture the variety of reasoning underlying the useful arguments that legal practitioners aim to support or undermine during their daily work.

²Two other outliers are two-label DocNLI (Yin et al., 2021) and three-label ConTRoL (Liu et al., 2021). However, while their premises often exceed the usual 512 maximum sequence length, unlike LawngNLI they still largely are not near the typical maximum sequence lengths of key current long-sequence pretrained models (see Table 2).

Sample twin Entail/Contradict examples with same premise from LawngNLI	
Twin hypotheses with same premise, from "analysis" subset	<ul style="list-style-type: none"> • <i>Contradict</i>: city acted affirmatively to create or increase risk of harm on city street by ignoring residents' requests to reduce speed limit or by taking down residents' signs indicating drivers should adhere to a lower speed limit • <i>Entail</i>: city did not act affirmatively to create or increase risk of harm on city street by ignoring residents' requests to reduce speed limit or by taking down residents' signs indicating drivers should adhere to a lower speed limit
Additional hypotheses with same premise	<ul style="list-style-type: none"> • <i>Entail</i>: failing to enforce or lower the speed limit on a residential street "did not create a 'special danger' to a discrete class of individuals..[ed.: excerpted]..as opposed to a general traffic risk to pedestrians and other automobiles" • <i>Contradict</i>: traffic laws and enforcement practices did not pose "a general traffic risk to pedestrians and other automobiles"
Relevant excerpts of shared premise	<ul style="list-style-type: none"> • [ed.: Plaintiffs] ...submit that the City of Fort Thomas..violated their son's substantive due process rights by failing to act upon their request (and the requests of others) to lower the speed limit on the street..The police also removed signs posted by residents indicating that drivers should adhere to a 15 mile-per-hour speed limit.. • [ed.: Plaintiffs] ...alleged that the City's failure to maintain safe conditions on Garrison Avenue violated their son's substantive due process rights..established a "state-created danger" under DeShaney.. • ...DeShaney's holding..precludes [ed.: Plaintiffs'] argument that the Due Process Clause constitutionalizes a locality's choices about what speed limit to adopt for a given street or how to enforce that speed limit.. • There are two exceptions to the DeShaney rule..Under the second exception..a plaintiff may bring a substantive due process claim by establishing (1) an affirmative act by the State that either created or increased the risk that the plaintiff would be exposed to private acts of violence.. • [ed.: Plaintiffs] fail to satisfy any of the three requirements for establishing our circuit's "state-created danger" exception to DeShaney. First, the creation of a street and the management of traffic conditions on that street are too attenuated and indirect to count as an "affirmative act"..
Distractor excerpts of same premise	<ul style="list-style-type: none"> • ...After all, the City was told about the risks of not lowering the speed limit to 15 miles per hour (more accidents); it intentionally chose not to heed this warning (taking on the risk of more accidents); and the alleged risk came to pass when..was killed (an accident).. • ...For in one sense, it could be said that all governing bodies act with deliberate indifference when they consider and reject a traffic-safety proposal of this sort that comes with known risks..

Table 1: Sample twin Entail/Contradict examples from LawngNLI. See Appendix Tables 9 and 10 for detailed view.

3.1 Construction Steps

Key steps are outlined below. They are presented in Figure 3 and further detailed in Appendix Section A.1.

When judges cite another case in an opinion, they may highlight content or takeaways from that case in a parenthetical which we use as an initial hypothesis.³

1. We begin with Entail examples: long premises are the majority opinion cited alongside the parenthetical, and short premises are the cited opinion pages (extracted using Eyecite (Cushman et al., 2021)). Filters screen for hypotheses that are not flagged as conflicting with premises and are valid inputs for our later contradiction algorithm.
2. Stratified by the Cartesian product of geography and pivotal negation versus not, we apply a contradiction algorithm and a neutralization algorithm to convert a random 1/3 each of the remaining Entail examples into Contradict and Neutral examples, respectively. The contradiction algorithm removes or (if absent) adds pivotal negation in the hypotheses assigned Contradict, by modifying

³These explanatory parentheticals are used by, for example, the legal research platform Casetext (Arredondo, 2017).

a method from Bilu et al. (2015)). All examples are filtered for NLI difficulty. Then the neutralization algorithm re-pairs the hypotheses assigned Neutral with alternative similar premises not adjacent in the citation network.

3. The dataset is rebalanced on labels and pivotal negation, citation spans are removed from the premises so that they read more naturally, and long premises are prepended with copied paragraphs from the end (the minimum number including at least 512 tokens) to limit models from relying on cues near the start of the underlying opinion.

Table 1 shows sample examples from our dataset. Descriptive statistics (Table 2) show that its long premises skew much longer than premises in other key existing NLI datasets.

3.2 Automatic Labels and Human Assessment

LawngNLI includes only automatic NLI labels. The Entail labels were effectively "annotated" by each judge authoring the (hypothesis) parenthetical citing another case's pages, but our algorithms constructing Neutral and Contradict examples could import some error rate. So with a balanced sample of long-premise and short-premise examples without prepending, these labels were assessed for

LawngNLI	Long premises		Short premises	
	“Analysis” subset	Full	“Analysis” subset	Full
Premise length	[970, 1527, 2339, 3154, 3693]	[1285, 2179, 3692, 6044, 9238]	[301, 462, 711, 925, 1397]	[331, 498, 746, 966, 1581]
Hypothesis length	22	21	22	21
Hypothesis negation	[0.579, 0.583, 0.586]	[0.574, 0.578, 0.583]	[0.579, 0.583, 0.586]	[0.574, 0.578, 0.583]
Unique training premises	27246	47803	29004	52643
Unique training hypotheses	69545	124486	69545	124486
Training examples	71442	128520	71442	128520
Text sources	Legal case opinions			
Existing datasets	ContractNLI	anli	DocNLI	ConTRoL-dataset
Premise length	[871, 1281, 2019, 2741, 4134]	[14, 28, 63, 80, 95]	[57, 73, 115, 557, 1050]	[55.6, 138, 333, 996, 1147]
Hypothesis length	19	14	57	16
Hypothesis negation	[0.284, 0.391, 0.012]	[0.074, 0.069, 0.197]	[0.187, 0.202]	[0.094, 0.078, 0.107]
Unique training premises	423	377440	488673	1530
Unique training hypotheses	17	1168146	572368	6238
Training examples	7191	3233665	942314	6719
Text sources	Contracts	Wikipedia, news, etc, SNLI, MNLI, NLI version of FEVER	ANLI, SQuAD, CNN/DailyMail, DUC (2001), Curation	Various civil service exams

Table 2: Descriptive statistics of NLI datasets. Negation words [‘no’, ‘not’, ‘never’, ‘none’, ‘nobody’, ‘nothing’, ‘neither’, ‘nor’, ‘cannot’] or contains “n’t”. Proportions are by label: Entail/Neutral/Contradict or Entail/Not entail. *About 50% of LawngNLI’s hypotheses contain pivotal negation, even though over 50% contain negation under the keyword definition (used here for comparability across datasets). See Appendix Section A.1 on dataset construction (the twinning step leads to the larger number of unique hypotheses than unique premises). Token lengths are [10, 25, 50, 75, 90] percentiles or an average via a RoBERTa (Liu et al., 2019) tokenizer.*

accuracy.

Screened Amazon Mechanical Turk workers provided 300 gold labels for LawngNLI examples. Because our long premises are lengthy, there is a particular risk of partly random guessing by workers. Even small frequencies can erroneously multiply the estimated error rates of our labels. As such, for each example, two workers independently chose a label and confidence level (“probably” or “definitely”). Gold labels were adopted on examples where both workers chose the same label (unanimity). Where the workers chose different labels, no label carries unanimous confidence and the example is not included. This process continued until 300 gold labels were obtained. Detailed steps are outlined in Appendix Section A.3.

We find a 88.8% human-validated accuracy (94.7% for high-confidence labels, where both workers chose “definitely”). Table 4 shows human-assessed accuracies of our automatic labels.

To better understand where dataset errors arise, in Appendix Table 11 we present several examples where automatic and gold labels differ. In the first, the citation parser linked the wrong citation to the parenthetical underlying the hypothesis. So the

example was incorrectly automatically labeled as Entail, while workers correctly labeled the example as Neutral. In the second example, the hypothesis is the conjunction of two possible conditions for a non-medical source to be given the weight of an acceptable medical source under the applicable regulations. Yet along with the first condition, the premise describes the opposite of the first condition. Thus the automatic label Contradict is arguably correct, while the worker label Entail may have arisen if both workers saw that the second condition is correct while misassessing the first condition. And in the third example, the premise and hypothesis were re-paired by our neutralization algorithm and so automatically labeled Neutral. However, despite safeguards such as excluding majority opinions from adjacent cases in the citation network as candidate premises, the workers recognized that the premise arguably contradicts the hypothesis despite deriving from a case not cited by the hypothesis’s underlying parenthetical. Thus the automatic label is incorrect and the worker label is correct.

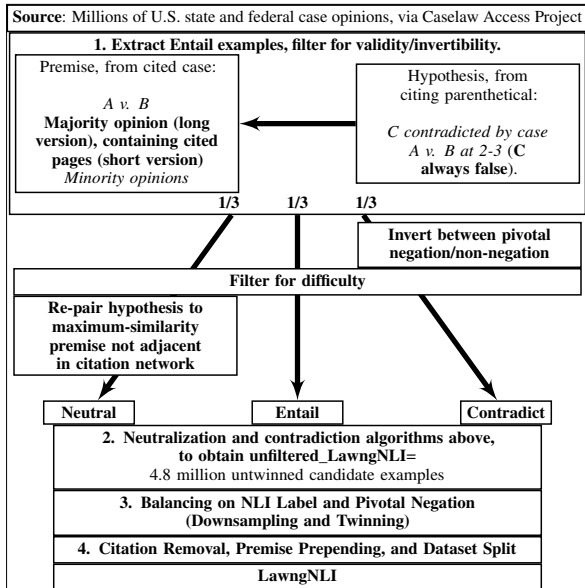


Table 3: Major steps in LawngNLI dataset construction process. Steps are detailed in Appendix Section A.1.

4 Experimental Evaluation

Our experiments demonstrate two applications of LawngNLI.⁴ See implementation details in Appendix Section A.2.

1. Generalization First, we evaluate whether models perform competitively on LawngNLI’s long premises, before and after fine-tuning on existing NLI benchmarks and/or its short premises (with LawngNLI’s multigranularity, we can evaluate on long premises within the same domain).

RQ1: Can top NLI models approach top performance on LawngNLI with long premises, absent fine-tuning directly with our long premises?

RQ1A: Is our answer to RQ1 robust to any error rate in LawngNLI’s automatic labels?

2. Retrieval Second, we evaluate retrieval and NLI models on a retrieve-and-re-rank approach to NLI-based (entailed or contradicted arguments as queries) case retrieval.

RQ2: How do models compare on implication-based case retrieval?

4.1 In-Domain Generalization to Long Contexts

To compare generalization from short to long contexts, we build a panel of 28 NLI models as follows:

⁴For all experiments, we use only LawngNLI’s “analysis” subset: with long premises at most 4096 tokens, via a RoBERTa (Liu et al., 2019) tokenizer.

LawngNLI automatic labels	Short premise		Long premise	
	All	Negation	All	Negation
Agreed-upon (gold) labels				
Accuracy	0.92	0.901	0.888	0.87
N	160	76	140	66
High-confidence agreed-upon (gold) labels				
Accuracy	0.972	0.976	0.947	0.905
N	81	39	68	31
Full assessment set				
Worker agreement	0.758	0.71	0.761	0.75
High confidence, if agreement	0.506	0.513	0.486	0.47
N	211	107	184	88

Table 4: Human assessment of a stratified random sample of LawngNLI’s “analysis” subset (sequence length of long premise at most 4096). The split refers to *pivotal* negation. Provided accuracies are balanced (macro-averaged recall). High-confidence labels are when the two workers who labeled a given example both clicked “definitely” (versus “probably”) the label.

1. We begin with 7 pretrained models that are top performing on key existing NLI benchmarks. HuggingFace (Wolf et al., 2020) model names with leaderboard positions on existing NLI benchmarks are listed in Appendix Section A.4.
2. To improve ability on general NLI, we add model versions with intermediate fine-tuning on other NLI benchmarks: three-label ANLI (Nie et al., 2020) which contains MNLI (Williams et al., 2018), three-label ConTRoL (Liu et al., 2021),⁵ and two-label (Entail, Not Entail) DocNLI (Yin et al., 2021).⁶
3. We then fine-tune on LawngNLI by adapting the code from Xiong et al. (2021).

Past NLI benchmarks have included artifacts spuriously correlated with their labels. To check for LawngNLI, we fine-tune and evaluate our models on hypotheses and premises only in Appendix Table 12 (Gururangan et al., 2018; Poliak et al., 2018; Tsuchiya, 2018; Yin et al., 2021). Our labels show some modest predictability above random from our hypotheses at around 0.55 at the highest (in line with other NLI datasets⁷) and from

⁵Following this paper, we fine-tune on ANLI and then ConTRoL.

⁶See Appendix Section A.2 about converting LawngNLI to two labels.

⁷For hypotheses, similar to ANLI (Nie et al., 2020) A1 at 0.497 and MNLI at 0.55 (Williams et al., 2018; Poliak et al., 2018) and slightly above ANLI later rounds and ConTRoL (Liu et al., 2021) in the 0.40s.

our long premises at 0.45 at the highest (slightly higher than another imperfectly symmetric benchmark (ConTRoL (Liu et al., 2021)) at 38.56%) and short premises at around 0.52 at the highest.

Especially for short-sequence models, some difficulty on long premises may be due to distraction by less relevant text. We address this by adding a version of each model with a filtering step prepended, selecting the 5 paragraphs with highest BM25 (Robertson and Zaragoza, 2009) score across the premise when querying the hypothesis. This filtering step allows our models to rival their performance on LawngNLI’s short premises, when inputting only its unexcerpted long premises.

4.1.1 RQ1: Are NLI Models Competitive on LawngNLI’s Long Premises, Absent Fine-Tuning with Our Long Premises?

Among our NLI models fine-tuned with existing benchmarks only and/or LawngNLI with short premises, the top model trails the top model fine-tuned with our long premises in accuracy by 5 percentage points (bolded in columns (4) and (6) among all 21 three-label models in Appendix Table 13; the top 6 models among these on both our short and long premises plus all long-sequence models are shortlisted in Table 5).⁸

Thus these models fall substantially short of top performance on long-premise NLI, unless a large-scale long-premise NLI dataset is constructed and used for fine-tuning. Models have more to learn here from long contexts than short contexts alone teach.

Generalizing from short to long premises, we control for much example structure by keeping the same domain and examples. However, judge excerpting could create some long versus short premise differences besides length, and it is possible that these differences are impacting our results. Such differences may be somewhat unavoidable when excerpting NLI premises that can cover many arguments. Relevant passages depend on the given hypothesis and so may be variably imbalanced relative to the long premise. In contrast to multilevel summaries, our motivation is comparing nested premises that read like premises naturally occurring in the domain at their respective lengths.

We would thus argue that differences arising

⁸Appendix Tables 14 and 15 show minimal differences on examples with hypotheses with versus without pivotal negation and slightly lower performance on examples with above-median length premises.

from human excerpting need not confound drawing wider parallels to our results. Counterpart (if differently structured) differences between naturally occurring short contexts versus long contexts likely exist in various domains. These differences are arguably properly includable in our experiments since they impact model generalization from short to long contexts in such other domains as well. Still it remains to be studied if large-scale long-context datasets are needed to perform competitively on other long-context NLP tasks and domains.

4.1.2 RQ1A: Is RQ1 Robust to Any Error Rate in LawngNLI’s Automatic Labels?

On our human-validated subset, we check how the gap in balanced accuracies from RQ1 (i.e., between our top models fine-tuned using long premises versus those fine-tuned using short premises) is impacted by evaluating against our automatic labels, as compared to against our gold labels. While sampling variation leads to some divergence of this human-validated subset from the test set, this bias in balanced accuracies due to automatic labels (on our human-validated subset) provides a noisy estimate of the same bias for the test set.⁹

In Appendix Table 16, we can check this impact of the automatic labels on the performance gap between the top models fine-tuned using long premises (column 6) versus top models fine-tuned using short premises (column 4) from Table 5. For the single model (among these top 6) with the largest performance gap from Table 5, this gap is biased on the human-validated subset by the automatic labels by -0.066. Averaging across all top 6 models, this bias is -0.016. These non-positive biases provide evidence that the performance gap from RQ1 likewise likely does not arise from bias due to the automatic labels.

4.2 Implication-Based Legal Retrieval

Major legal retrieval systems (including leading commercial systems, according to publicly available information) rely on signals from lexical keywords, semantic similarity, and question answering (Section 5). But to our knowledge, such systems do not utilize signals from legal NLI. In contrast, LawngNLI’s hypotheses are derived from the actual implications from the premise case that judges write within new cases when applying the premise case as precedent.

⁹The modest gold subset size provides low power for statistical testing, but we can still calculate with the point values.

Suppose a user wished to retrieve documents that support or refute an argument: Can leading retrieval and NLI models perform well (zero shot or with fine-tuning) on retrieving the target document when given that argument as the query? We test this question by comparing models on implication-based retrieval in our domain.

4.2.1 RQ2: How Do Models Compare on Implication-Based Case Retrieval?

We proceed in three steps: cheap retrieval, bi-encoder ranking, and cross-encoder re-ranking.

(1) Building our test set, using cheap retrieval across all candidate cases: Our retrieval test set comprises the LawngNLI test set’s Entail and Contradict long premises (majority opinions, here before dataset processing) as positive examples, pooling each with 999 other majority opinions from the same state (or the federal level) selected by highest all-mpnet-base-v2 (Song et al., 2020) embedding dot-product similarity with the hypothesis as negative examples. An effective retrieval system would rank the premise case highly as correct, and we can compare models on Recall@k. So each non-Neutral LawngNLI test set hypothesis is a query paired with 1000 candidate documents.

This step corresponds to the initial retrieval step in a standard retrieve-and-re-rank approach to the retrieval problem. We control for this first step by uniformly applying this method and bracket a search for improvements. This allows our comparison to focus on which models rank target cases most highly in steps 2 and 3, with the aspiration of rankings high enough to be within users’ reach with minimal skimming beyond the top result.

(2) Zero-shot bi-encoder ranking of retrieved top 1000: Our zero-shot panel comprises a BM25 (Robertson and Zaragoza, 2009) baseline and four Sentence-Transformers (Reimers and Gurevych, 2021) lightweight bi-encoders: msmarco-distilroberta-base-v2, nli-distilroberta-base-v2, all-distilroberta-v1, and all-mpnet-base-v2 (Sanh et al., 2019; Liu et al., 2019; Bajaj et al., 2016; Song et al., 2020). The first three were chosen for comparable setups besides their training task, while the last is a top model across semantic search evaluations.¹⁰ They are respectively pre-trained for actual user queries from the Bing search engine, NLI, and a combined dataset of over 1 bil-

¹⁰https://www.sbert.net/docs/pretrained_models.html#sentence-embedding-models. No all-distilroberta-v2 was available on HuggingFace (Wolf et al., 2020).

lion pairs of related sentences covering many tasks. As in Section 4.1, we also evaluate each model version while prepending with a module that filters each candidate document to the 5 paragraphs with highest BM25 similarity to the query.

Appendix Table 17 shows that leading, lightweight bi-encoders and BM25 (Robertson and Zaragoza, 2009) can rank the target case reasonably well zero shot while processing a large number of candidate documents. Prepending with BM25 filtering improves performance. Since Entail and Contradict hypotheses are twinned, it is not so surprising that performance is similar between those two subsets. Still, these retrieval models seem capable of recognizing long-context inference (as a subtype of relevance) close to equally between entailed and contradicted queries. While Recall@10 of the target premise case reaches about 0.3, Recall@100 still only reaches about 0.5 (versus expected value of 0.1), showing that leading retrieval systems would miss important documents for legal retrieval zero shot.

(3) NLI cross-encoder re-ranking of ranked top 100: We re-rank our top Recall@100 rankings, achieved by all-distilroberta-v1 (Sanh et al., 2019; Liu et al., 2019) prepended with BM25 (Robertson and Zaragoza, 2009) filtering. Re-ranking is by (Entail or Contradict) label probability. All models are prepended with BM25 filtering, which improved performance in Appendix Table 17.

Our model panel starts with albert-xxlarge-v2 (Lan et al., 2019) fine-tuned on each of our three included previous NLI datasets (ANLI (Nie et al., 2020), ConTRoL (Liu et al., 2021), and DocNLI (Yin et al., 2021)). We also add cross-encoder/ms-marco-MiniLM-L-6-v2 (Wang et al., 2020; Reimers and Gurevych, 2021; Bajaj et al., 2016) as an additional zero-shot baseline. Among our dense baselines, fine-tuning on ANLI provides the top re-ranking performance. We then fine-tune the ANLI model on an adjusted retrieval version of LawngNLI.¹¹ We use the same setup as for our NLI intermediate fine-tuning on ANLI except with a learning rate of 1e-6.

Table 6 presents our comparison. Zero shot, re-ranking with these dense models underperforms

¹¹Here, premises are without prepending and without citation spans removed. And models are instead fine-tuned on LawngNLI’s Entail or Contradict examples, together with its Neutral hypotheses instead paired with a random incorrect all-mpnet-base-v2 (Song et al., 2020) top 100 candidate in order to align more closely with re-ranking task negatives.

Needs long premises for fine-tuning	No			Yes		
	Short premise			BM25 retrieval on short premise	Long premise	Long premise filtered by BM25 retrieval
Fine-tuning	No	No [512 [tokens]	Yes	Yes	No	Yes
Long premise filtered by BM25 retrieval at evaluation	(1)	(2)	(3)	(4)	(5)	(6)
[Entail/Neutral/Contradict. Chance=1/3]						
albert-xxlarge-v2_anli	0.742+/-0.014	0.742+/-0.014	0.81+/-0.013	0.817+/-0.012	0.789+/-0.013	0.853+/-0.011
albert-xxlarge-v2_ConTRoL-dataset	0.725+/-0.014	0.725+/-0.014	0.804+/-0.013	0.806+/-0.013	0.781+/-0.013	0.851+/-0.011
albert-xxlarge-v2_vanilla	0.743+/-0.014	0.743+/-0.014	0.816+/-0.012	0.825+/-0.012	0.781+/-0.013	0.852+/-0.011
allenai_longformer-base-4096_anli	0.562+/-0.016	0.652+/-0.015	0.706+/-0.014	0.707+/-0.014	0.691+/-0.015	0.784+/-0.013
allenai_longformer-base-4096_ConTRoL-dataset	0.56+/-0.016	0.641+/-0.015	0.698+/-0.015	0.691+/-0.015	0.693+/-0.015	0.781+/-0.013
allenai_longformer-base-4096_vanilla	0.541+/-0.016	0.604+/-0.015	0.691+/-0.015	0.701+/-0.015	0.51+/-0.016	0.791+/-0.013
facebook_bart-large_anli	0.664+/-0.015	0.675+/-0.015	0.808+/-0.013	0.809+/-0.013	0.76+/-0.014	0.868+/-0.011
facebook_bart-large_ConTRoL-dataset	0.665+/-0.015	0.666+/-0.015	0.809+/-0.013	0.814+/-0.012	0.758+/-0.014	0.868+/-0.011
facebook_bart-large_vanilla	0.667+/-0.015	0.699+/-0.015	0.813+/-0.012	0.807+/-0.013	0.782+/-0.013	0.875+/-0.011
google_bigbird-roberta-base_anli	0.613+/-0.015	0.666+/-0.015	0.764+/-0.014	0.774+/-0.013	0.77+/-0.013	0.829+/-0.012
google_bigbird-roberta-base_ConTRoL-dataset	0.601+/-0.015	0.656+/-0.015	0.761+/-0.014	0.769+/-0.013	0.757+/-0.014	0.819+/-0.012
google_bigbird-roberta-base_vanilla	0.648+/-0.015	0.657+/-0.015	0.767+/-0.013	0.777+/-0.013	0.763+/-0.014	0.827+/-0.012
Maximum of p-values versus (6)	0	0	0	0		
N				3966		

Table 5: Accuracy on LawngNLI’s “analysis” subset *with long premises*: Top 6 three-label models on both short and long premises plus long-sequence models (for all 28 models, see Appendix Table 13). The error provided is the larger deviation of the Clopper-Pearson (Clopper and Pearson, 1934) exact binomial 95% confidence bounds. All p-values round to zero (<0.0005) from an exact binomial McNemar’s (McNemar, 1947) test for a statistically significant difference in accuracies between each version fine-tuning using short premises as inputs (1-4) and the best version fine-tuning using long premises as inputs (6). For (2), 512 tokens is the overall sequence limit.

Evaluation (long premise filtered by BM25 retrieval)	[Entail hypotheses as queries]			[Contradict hypotheses as queries]		
	1	10	100	1	10	100
albert-xxlarge-v2_anli-LawngNLI-retrieval	0.138+/-0.018	0.321+/-0.019	0.505+/-0.003	0.1+/-0.016	0.271+/-0.019	0.486+/-0.003
albert-xxlarge-v2_anli	0.09+/-0.016	0.262+/-0.02	0.505+/-0.003	0.07+/-0.014	0.174+/-0.019	0.486+/-0.003
albert-xxlarge-v2_ConTRoL-dataset	0.034+/-0.011	0.144+/-0.018	0.505+/-0.003	0.048+/-0.012	0.177+/-0.019	0.486+/-0.003
albert-xxlarge-v2_DocNLI	0.067+/-0.014	0.226+/-0.02	0.505+/-0.003	0.008+/-0.006	0.083+/-0.015	0.486+/-0.003
all-mpnet-base-v2	0.09+/-0.016	0.228+/-0.02	0.505+/-0.003	0.076+/-0.015	0.219+/-0.019	0.486+/-0.003
cross-encoder_ms-marco-MiniLM-L-6-v2	0.019+/-0.009	0.116+/-0.017	0.505+/-0.003	0.013+/-0.008	0.123+/-0.017	0.486+/-0.003
BM25	0.016+/-0.008	0.061+/-0.014	0.505+/-0.003	0.008+/-0.006	0.052+/-0.013	0.486+/-0.003
N				1322		

Table 6: Recall@k of model panel, when re-ranking the bi-encoder top 100 (ranked by all-distilroberta-v1 (Sanh et al., 2019; Liu et al., 2019) prepended with BM25 (Robertson and Zaragoza, 2009) filtering from Appendix Table 17) for implication-based retrieval. The error provided is the larger of the two deviations of the Clopper-Pearson (Clopper and Pearson, 1934) exact binomial 95% confidence bounds from the point estimate.

BM25 (Robertson and Zaragoza, 2009) on Recall@1 and Recall@10 (compare Appendix Table 17). However, fine-tuning on our adjusted version of LawngNLI draws re-ranking performance toward BM25. Refining our approach to fine-tuning using LawngNLI may yield future improvements.

5 Related Work

Legal retrieval: Locke and Zuccon (2022) note that the most used commercial systems have not publicly released specific algorithms. Still, some providers have shared core feature types or modules. The retrieval module for a recent QA system (by authors including LexisNexis researchers) compares candidates by sparse (BM25-based) and dense representations (Khazaeli et al., 2021). West-Search Plus categorizes questions into frames and selects passages using QA-pair classifiers (McElvain et al., 2019). Modern commercial systems’ features include text, citation networks, annotations, and query logs.¹² Published retrieval systems such as Khazaeli et al. (2021) (by authors including LexisNexis researchers) and Tang and Clematide (2021) do leverage Transformer embedding similarity. To our knowledge, our application of NLI-based retrieval to legal case research is novel.

Long-context NLI benchmarks: SCROLLS (Shaham et al., 2022) analyzes long-context datasets across multiple tasks (ContractNLI (Koreeda and Manning, 2021) is the included NLI dataset). Schuster et al. (2022) investigated retrieval and aggregation methods to scale up their sentence-level NLI model and baselines for DocNLI (Yin et al., 2021) and ContractNLI (Koreeda and Manning, 2021). Without premise multigranularity, however, these long-context experiments do not evaluate models’ generalization to long contexts while holding constant the target domain and examples (e.g., no model-external NLI labels for evaluation data’s evidence spans).

ContractNLI (Koreeda and Manning, 2021) constructed an NLI dataset with similarly long premises, though containing a substantially smaller number of examples (607 contracts as premises, each paired with 17 shared hypotheses) in the quite linguistically different contracts domain. They study evidence identification and context segmentation within their premises, as compared to (in

the present paper) benchmarking models on multigranular premises. LawngNLI contains over a hundred thousand distinct hypotheses, providing models with domainwide supervision.

Other legal NLP: AutoLAW and CaseHOLD (Mahari, 2021; Zheng et al., 2021) construct datasets for a distinct task of predicting holdings from other cases that support the arguments in the nearby context in the *citing* case. These holdings exhibit an argument support relation with respect to their surrounding context, as opposed to necessarily any NLI relation. Recently, Shen et al. (2022) introduced a multi-document summarization dataset for documents around civil rights cases, with multiple granularities of summaries. The legal tasks closest to ours are from the annual COLIEE workshop.¹³ However, these tasks do not fully map to three-label NLI. For, e.g., relevant 2021 Tasks 2 and 4, their training corpora (in the hundreds of examples) are ballpark 1000 times smaller than usual single-sentence benchmarks, making supervised learning alone insufficient for reliably training models to generalize (Hudzina et al., 2020; Rabelo et al., 2021; Kim et al., 2021; Schilder et al., 2021).

6 Conclusion and Future Work

This work presents LawngNLI, a new NLI benchmark with multigranular long premises, each containing a shorter version. Experiments demonstrate some use cases. First, we show that leading NLI models fall substantially short of competitive performance when generalizing to LawngNLI with its long premises, even after fine-tuning using existing NLI benchmarks and/or LawngNLI with short premises (with the same domain and examples as the long-premise evaluation). Unconfounded by domain shift, these results show the need for a large-scale long-premise dataset like ours at fine-tuning time.

Second, we show that leading lightweight retrieval models can reasonably handle implication-based retrieval on LawngNLI zero shot with both entailed and contradicted arguments as queries. We then compare re-ranking by lexical overlap and models fine-tuned using a modified LawngNLI or several previous NLI datasets. Multiple other aspects of LawngNLI are left for future study.

¹²<https://blog.law.cornell.edu/voxpath/2013/03/28/next-generation-legal-search-its-already-here/>

¹³<https://sites.ualberta.ca/~rabelo/COLIEE2021/>

Limitations

LawngNLI contains automatic labels, derived from the construction process. Its Entail labels are effectively annotated by judges, who wrote Entail hypotheses as parentheticals asserted by the cited premise. Neutral and Contradict examples are derived from Entail examples by, respectively, repairing with a different non-adjacent premise in the citation network and by adding or removing pivotal negation (Appendix Section A.1). These steps could introduce some error rate, which we validate by human assessment (Section 3.2). And using a human-validated subset, Section 4.1.2 evinces that our conclusions from our generalization experiment likely do not arise from such differences between automatic and gold labels.

Both experiments test standard approaches when applied to distinct challenges: first, short-premise models on long-premise NLI and, second, semantic search models on implication-based retrieval. For the experiment on generalization to long contexts, we demonstrate that these standard approaches do not always suffice, but only in one (albeit important) counterexample domain: law. We have not established if these shortcomings extend to other domains more broadly.

Ethics Statement

Considerations for general NLI have been explored elsewhere (e.g., for gender bias by Sharma et al. (2021)). We discuss some considerations for the legal aspect. On the benefit side, NLI is a principal cognitive task in law, so progress here also stands to benefit the legal community: Building court cases and advising clients essentially is arguing for and against different natural language inferences from legal texts and facts. Implications may not be directly stated in the text or annotations (e.g., those at a different level of specificity or requiring compositional reasoning). Instead, holdings and rules inferable from case text must be extracted through costly human annotation and curation. All around the legal system, the pay grade and spare bandwidth of legal counsel is frequently starkly imbalanced between parties with adversarial interests: whether people in the courtroom or settlement conference, consumers or companies in a negotiation boardroom, or in everyday society where behavior is shaped by prospects of legal action. Anything that makes legal research and thus legal counsel cheaper, including more lightweight

or task-tailored case retrieval systems, can contribute toward fairer access to legal representation and justice regardless of financial means. Models that perform well on LawngNLI's retrieval setup could crosswalk between cases as premises and implications as hypotheses, performing implication-based retrieval automatically. We describe the state of legal retrieval systems (including limited public information about leading proprietary commercial algorithms) relative to an NLI-based approach in Section 5. Legal services overall comprise about 1.3% of U.S. GDP.¹⁴ The legal research industry's annual revenue meanwhile is in the multiple billions of dollars.¹⁵ And the full societal cost of suboptimal case retrieval should include the time and resources expended by human legal researchers in the loop (paralegals and lawyers) in unnecessary iterating with any suboptimal retrieval systems. Indeed, junior lawyers (less than 10 years of experience) spend almost a third (28%) of their working time on case research (Poje, 2014).¹⁶

On the risk side, while prospective human reliance for decision making on erroneous model predictions is an ever-present consideration in NLP, we do not view this as a practical risk for LawngNLI. Everyday people can turn to numerous simple articles online summarizing the law, without digging into complex case retrieval and jurisprudence. And regarding advising others, lawyers bound by professional duties are exclusively authorized to practice law in the U.S. and around the world.¹⁷ Nothing can even be done just knowing the most relevant cases or implications; they must be synthesized by human judgment into an argument sound enough to pass the muster of judges and juries. In other words, legal NLI models are in no way lawyers. Instead, they can work as screening tools for practitioners who then must apply their own judgment to make the results useful. In this way, legal NLI models could help save the resources of lawyers and clients and help improve the quality of legal representation.

Acknowledgements

This research is based upon work supported in part by the Office of the Director of National Intelli-

¹⁴<https://fred.stlouisfed.org/graph/?g=PfxD>

¹⁵As of 2020: e.g., Thomson Reuters.

¹⁶This percentage that falls to 16% for senior lawyers who are likely to take on a more managerial role.

¹⁷<https://www.ibanet.org/MediaHandler?id=199b20ec-b7ab-4ef4-99c4-cd45c7b6371b>

gence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via IARPA Contract No. 2019-19051600006 under the BETTER Program. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, the Department of Defense, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein. This research is supported by a Focused Award from Google.

References

- Pablo D. Arredondo. 2017. Harvesting and Utilizing Explanatory Parentheticals. *SCL Rev.*, 69:659. Publisher: HeinOnline.
- Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, and Tri Nguyen. 2016. Ms marco: A human generated machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268*.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Yonatan Bilu, Daniel Hershcovich, and Noam Slonim. 2015. Automatic claim negation: Why, how and when. In *Proceedings of the 2nd Workshop on Argumentation Mining*, pages 84–93.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python*. O’Reilly, Cambridge.
- Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, Nikolaos Aletras, and Ion Androutsopoulos. 2020. LEGAL-BERT: The Muppets straight out of Law School. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2898–2904, Online. Association for Computational Linguistics.
- Charles J. Clopper and Egon S. Pearson. 1934. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, pages 404–413. ISBN: 0006-3444 Publisher: JSTOR.
- Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, and Ellen M. Voorhees. 2020. Overview of the trec 2019 deep learning track. *arXiv preprint arXiv:2003.07820*.
- Jack Cushman, Matthew Dahl, and Michael Lissner. 2021. eyecite: A tool for parsing legal citations. *Journal of Open Source Software*, 6(66):3617. ISBN: 2475-9066.
- I. Dagan, O. Glickman, and B. Magnini. 2005. The pascal recognising textual entailment challenge. In *Proceedings of PASCAL first Workshop on Recognising Textual Entailment*.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. Annotation artifacts in natural language inference data. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112.
- Sebastian Hofstätter, Sheng-Chieh Lin, Jheng-Hong Yang, Jimmy Lin, and Allan Hanbury. 2021. Efficiently Teaching an Effective Dense Retriever with Balanced Topic Aware Sampling. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 113–122, Virtual Event Canada. ACM.
- John Hudzina, Kanika Madan, Dhivya Chinnappa, Jinane Harmouche, Hiroko Bretz, Andrew Vold, and Frank Schilder. 2020. Information Extraction/Entailment of Common Law and Civil Code. In *JSAI International Symposium on Artificial Intelligence*, pages 254–268. Springer.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*. ISBN: 2332-7790 Publisher: IEEE.
- Soha Khazaeli, Janardhana Punuru, Chad Morris, Sanjay Sharma, Bert Staub, Michael Cole, Sunny Chiu-Webster, and Dhruv Sakalley. 2021. A free format legal question answering system. In *Proceedings of the Natural Legal Language Processing Workshop 2021*, pages 107–113.
- Mi-Young Kim, Juliano Rabelo, R Goebel, and y. 2021. BM25 and Transformer-based Legal Information Extraction and Entailment. New York, NY, USA. ACM.
- Nikita Kitaev, Steven Cao, and Dan Klein. 2019. Multilingual Constituency Parsing with Self-Attention and Pre-Training. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3499–3505, Florence, Italy. Association for Computational Linguistics.
- Nikita Kitaev and Dan Klein. 2018. Constituency Parsing with a Self-Attentive Encoder. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2676–2686, Melbourne, Australia. Association for Computational Linguistics.
- Yuta Koreeda and Christopher D. Manning. 2021. ContractNLI: A Dataset for Document-level Natural Language Inference for Contracts. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1907–1919.

- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880.
- Hanmeng Liu, Leyang Cui, Jian Liu, and Yue Zhang. 2021. Natural Language Inference in Context- Investigating Contextual Reasoning over Long Texts. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13388–13396. Issue: 15.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).
- Daniel Locke and Guido Zuccon. 2022. Case law retrieval: problems, methods, challenges and evaluations in the last 20 years. *arXiv preprint arXiv:2202.07209*.
- Robert Zev Mahari. 2021. AutoLAW: Augmented Legal Reasoning through Legal Precedent Prediction. *arXiv preprint arXiv:2106.16034*.
- Gayle McElvain, George Sanchez, Sean Matthews, Don Teo, Filippo Pompili, and Tonya Custis. 2019. Westsearch plus: A non-factoid question-answering system for the legal domain. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1361–1364.
- Quinn McNemar. 1947. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*.
- Ines Montani, Matthew Honnibal, Matthew Honnibal, Van L, Sofie eghem, Adriane Boyd, Henning Peters, Paul O’Leary McCann, Maxim Samsonov, Jim Geovedi, Jim O’Regan, György Orosz, Duygu Altinok, Søren Lind Kristiansen, Roman, Explosion Bot, Le Fiedler, er, Grégory Howard, Wannaphong Phatthiyaphaibun, Yohei Tamura, Sam Bozek, Murat, Mark Amery, Björn Böing, Pradeep Kumar Tippa, Leif Uwe Vogelsang, Ramanan Balakrishnan, Vadim Mazaev, GregDubbin, Jeannefukumaru, and Walter Henry. 2021. [explosion/spaCy: v3.1.4: Python 3.10 wheels and support for AppleOps](#).
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. Adversarial NLI: A New Benchmark for Natural Language Understanding. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4885–4901.
- Joshua Poje. 2014. [Legal Research](#). Technical report, American Bar Association.
- Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. 2018. Hypothesis Only Baselines in Natural Language Inference. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 180–191.
- Juliano Rabelo, R Goebel, y, Yoshinobu Kano, Mi-Young Kim, Masaharu Yoshioka, and Ken Satoh. 2021. Summary of the Competition on Legal Information Extraction/Entailment (COLIEE) 2021. New York, NY, USA. ACM.
- Radim Rehurek and Petr Sojka. 2010. Software framework for topic modelling with large corpora. In *Proceedings of the LREC 2010 workshop on new challenges for NLP frameworks*. Citeseer.
- Nils Reimers and Iryna Gurevych. 2021. [The Curse of Dense Low-Dimensional Information Retrieval for Large Index Sizes](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 605–611, Online. Association for Computational Linguistics.
- Stephen Robertson and Hugo Zaragoza. 2009. *The probabilistic relevance framework: BM25 and beyond*. Now Publishers Inc.
- Victor Sanh, Lys Debut, re, Julien Chaumond, and Thomas Wolf. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Frank Schilder, Dhivya Chinnappa, Kanika Madan, Jinane Harmouche, Andrew Vold, Hiroko Bretz, and John Hudzina. 2021. A Pentapus Grapples with Legal Reasoning. In *Proceedings of the Eighth International Competition on Legal Information Extraction/Entailment (COLIEE 2021)*, New York, NY, USA. ACM.
- Tal Schuster, Sihao Chen, Senaka Buthpitiya, Alex Fabrikant, and Donald Metzler. 2022. Stretching Sentence-pair NLI Models to Reason over Long Documents and Clusters. *arXiv preprint arXiv:2204.07447*.
- Uri Shaham, Elad Segal, Maor Ivgi, Avia Efrat, Ori Yoran, Adi Haviv, Ankit Gupta, Wenhan Xiong, Mor Geva, and Jonathan Berant. 2022. SCROLLS: Standardized Comparison Over Long Language Sequences. *arXiv preprint arXiv:2201.03533*.
- Shanya Sharma, Manan Dey, and Koustuv Sinha. 2021. Evaluating Gender Bias in Natural Language Inference. *arXiv preprint arXiv:2105.05541*.

- Zejiang Shen, Kyle Lo, Lauren Yu, Nathan Dahlberg, Margo Schlanger, and Doug Downey. 2022. Multi-LexSum: Real-World Summaries of Civil Rights Lawsuits at Multiple Granularities. *arXiv preprint arXiv:2206.10883*.
- Tom De Smedt and Walter Daelemans. 2012. **Pattern for Python**. *Journal of Machine Learning Research*, 13(66):2063–2067.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2020. MpNet: Masked and permuted pre-training for language understanding. *Advances in Neural Information Processing Systems*, 33:16857–16867.
- Li Tang and Simon Clematide. 2021. Searching for legal documents at paragraph level: Automating label generation and use of an Extended Attention Mask for boosting neural models of semantic similarity. Association for Computational Linguistics.
- Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. 2020. Efficient transformers: A survey. *arXiv preprint arXiv:2009.06732*.
- The President and Fellows of Harvard University. 2018. **Caselaw Access Project**.
- Thomson Reuters. **Annual Report 2020**. Technical report.
- Masatoshi Tsuchiya. 2018. Performance Impact Caused by Hidden Bias of Training Data for Recognizing Textual Entailment. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.
- Chantal van Son, Emiel van Miltenburg, and Roser Morante. 2016. **Building a Dictionary of Affixal Negations**. In *Proceedings of the Workshop on Extra-Propositional Aspects of Meaning in Computational Linguistics (ExProM)*, pages 49–56, Osaka, Japan. The COLING 2016 Organizing Committee.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. **GLUE: A multi-task benchmark and analysis platform for natural language understanding**. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355.
- Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *Advances in Neural Information Processing Systems*, 33:5776–5788.
- 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.
- Thomas Wolf, Julien Chaumond, Lys Debut, re, Victor Sanh, Clement Delangue, Anthony Moi, Pierric Cistac, Morgan Funtowicz, Joe Davison, and Sam Shleifer. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45.
- Yunyang Xiong, Zhanpeng Zeng, Rudrasis Chakraborty, Mingxing Tan, Glenn Fung, Yin Li, and Vikas Singh. 2021. Nyströmformer: A Nyström-based Algorithm for Approximating Self-Attention. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 14138–14148. Issue: 16.
- Wenpeng Yin, Dragomir Radev, and Caiming Xiong. 2021. DocNLI: A Large-scale Dataset for Document-level Natural Language Inference. *arXiv preprint arXiv:2106.09449*.
- Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, and Li Yang. Big Bird: Transformers for Longer Sequences.
- Lucia Zheng, Neel Guha, Br Anderson, on R., Peter Henderson, and Daniel E. Ho. 2021. When Does Pretraining Help? Assessing Self-Supervised Learning for Law and the CaseHOLD Dataset. *arXiv preprint arXiv:2104.08671*.

A Appendix

A.1 Dataset Construction Procedure

A.1.1 Extraction from Caselaw Access Project

LawngNLI is constructed starting with all xml case files from the April 21, 2021 bulk export from the Caselaw Access Project (**The President and Fellows of Harvard University, 2018**). The word count of the full original corpus before processing at about 12 billion¹⁸ is around three times that of English Wikipedia¹⁹, though for our premises we limit to only the majority opinions.

Entail examples are pairs of citation parentheticals (hypotheses) and excerpts of majority opinions from cited cases with resolvable pincites (premises), extracted from case files using Eyecite (**Cushman et al., 2021**). Where Eyecite associates multiple consecutive citations resolving to the same case with the same citation parenthetical, only the first citation and its pincite, if any, is paired with the

¹⁸https://case.law/docs/site_features/trends

¹⁹About 4 billion as of December 1, 2021: <https://web.archive.org/web/20211201013917/https://en.wikipedia.org/wiki/Special:Statistics>

parenthetical and included as an example. In this paper, we only include examples from citations including a resolvable pincite (e.g., does not contain letters).

The short version of the premise consists of the resolvable cited pages within the cited case’s majority opinion, while the long version of the premise consists of the cited case’s full majority opinion.

A.1.2 Initial Filters

Examples are dropped or modified by simple “accuracy” filters. First, as an overbroad criterion to exclude examples where the (converted or unconverted) original Entail hypothesis was a parenthetical in a case that was later overturned, we drop all examples with hypotheses from cases where a later case shared the same party names in the same or reverse order.

Second, parentheticals with citations including a case history flag (e.g., “acq.”, “aff’d”) are excluded.

Third, we drop examples with hypotheses that contain certain regex keywords (‘quot-ing|len banclomitted|emphasis|applying|citing |concur|dissent|majority|, in chambers|per curiam|Lexis|opinion| v. |§|¶|[0-9]’) associated with parentheticals describing “metadata” about the cited case rather than its content.

Fourth, verbs ending with “ing” followed by “that” at the beginning of remaining hypotheses many times take a supporting stance toward the subsequent subordinate clause, so to adapt such hypotheses to be more similar to a standalone sentence, we remove such initial words and the subsequent “that” in hypotheses.

Finally, sentences are normalized with spaCy 3.1.1 (Montani et al., 2021) to, e.g., process contractions.

A.1.3 Identifying (Pivotal) Negation in Hypotheses

Next the Entail examples are automatically labeled by whether their hypotheses contain (pivotal) negation or not, depending on whether the contradiction algorithm described in Appendix Section A.1.5 removes or adds negation, respectively. Pairs with hypotheses rejected for processing by our contradiction algorithm are dropped from the dataset.

Since the absence versus presence of such negation in the hypothesis results in contradictory truth values (and thus also flips the NLI label between Entail and ‘Contradict’), such negation can be called “pivotal.” Negation is defined this way throughout

the paper except in Table 2 when comparing to other datasets, since our contradiction algorithm might exhibit a different error rate on those datasets and confound the comparison. For this reason, greater than 50% of LawngNLI’s hypotheses contain negation in Table 2, even though the dataset is constructed to contain 50% (pivotal) negation hypotheses.

A.1.4 NLI Label Split

Within examples from cases from each state (or federal) and pivotal negation or not, entail examples are randomly assigned to be 1/3 Entail, 1/3 converted to Neutral, and 1/3 converted to Contradict.

A.1.5 Converting Entail Examples to Contradict Examples: Contradiction Algorithm

For examples labeled Contradict in Appendix Section A.1.4, we use our contradiction algorithm to add or remove pivotal negation²⁰ from the hypothesis, toward aligning the NLI relation with the label.

Our contradiction algorithm builds on the negation algorithm outlined in Section 4.2 of Bilu et al. (2015), which in their paper was annotated by majority vote to have generated an opposing claim with probability 0.79.²¹

The algorithm chooses a random sentence for adding or removing negation and leaves the others unchanged. It finds a non-compound independent clause within the chosen sentence and then makes the first applicable change in the list below. If none of the changes’ conditions apply, the hypothesis is rejected for processing by the algorithm. This includes rejecting hypotheses consisting of verb phrases not nested within independent clauses; since these are rarely found in negated form in the original dataset, including them would leave an artifact of this contradiction algorithm. So for these hypotheses, we prioritize balance across labels over coverage of candidate examples.

²⁰“Pivotal” negation is negation the absence versus presence of which results in at least some contradictory truth values for the hypothesis, flipping its NLI label from Entail to Contradict.

²¹Hypotheses are parsed with the Berkeley Neural Parser 0.2.0 ‘benepar_en3’ with spaCy 3.1.1 ‘en_core_web_lg’ (Kitaev et al., 2019, Kitaev and Klein, 2018, Montani et al., 2021). Verb tense is modified using NLTK 3.6.2 WordNet Lemmatizer and Pattern 3.6 conjugate function (?; Bird et al., 2009; Smedt and Daelemans, 2012). We explored attempting to negate adjectives and verbs using the lexical negation dictionary compiled by van Son et al. (2016) but ultimately limited to just using direct negation.

1. If there are any contradictable indefinite pronouns in the first highest-level noun phrase, the first one is changed to a contradictory pronoun (e.g., “some” to “none” or “neither” to “either”).
2. If there are any verb phrases, the first highest-level verb phrase is contradicted using a modified version (e.g., also reversing negation by removing “do”/“does”/“did”+“not”) of the negation algorithm from [Bilu et al. \(2015\)](#) mentioned above.
3. If there are any adjective phrases, the first [‘no’, ‘not’, ‘never’] is removed from or else a ‘not’ is added to the first highest-level adjective phrase or past participle.

A.1.6 Filtering

Now we apply simple “difficulty” filters: examples with hypotheses containing quotation marks or fewer than four words or with at least 50% bigram overlap with their premise are dropped.

A.1.7 Converting Entail Examples to Neutral Examples: Neutralization Algorithm

For examples labeled Neutral in Appendix Section A.1.4, we use our neutralization algorithm to match the hypothesis with a different premise, toward aligning the NLI relation with the label. To balance attrition, the neutralization algorithm is applied to all examples regardless of NLI label, but only the hypotheses from Neutral examples are actually re-paired with the assigned premise.

The candidates for matching with each hypothesis are the premises from all examples that are from cases in the same state as the original premise (or from a federal case if the original premise is from a federal case). Excluded from candidacy are premises from cases citing or cited by the case containing the original hypothesis.

A hypothesis is paired with a candidate premise as follows. The short version of the premise is used for this step.

First, the top 30 (dot-product) nearest neighbors of the hypothesis among the candidates are retrieved using FAISS ([Johnson et al., 2019](#))²² on msmarco-distilbert-base-tas-b embeddings ([Hofstätter et al., 2021](#))²³ via

²²<https://github.com/facebookresearch/faiss>

²³https://huggingface.co/sebastian-hofstaetter/distilbert-dot-tas_b-b256-msmarco. <https://www.sbert.net/docs/pretrained-models/msmarco-v3.html>

Sentence-Transformers (<https://github.com/UKPLab/sentence-transformers>, [Reimers and Gurevych, 2021](#)).

Second, candidate premises with which a hypothesis has at least 50% bigram overlap are dropped. This step preserves the filter applied earlier to all examples through the re-pairing for the Neutral examples.

Finally, Neutral hypotheses only are paired with their remaining candidate premise with respect to which it has the highest BM25 ([Robertson and Zaragoza, 2009](#)) score via Gensim 3.8.3 ([Rehurek and Sojka, 2010](#)). For hypotheses of all labels, if no candidate premises remain, their example is dropped.

A.1.8 Balancing

We split the dataset into “analysis”/non-“analysis” subsets by the inclusion criterion for this paper’s experimental evaluation (Section 4): whether the sequence length of an example’s long premise is at most 4096 tokens, via a RoBERTa ([Liu et al., 2019](#)) tokenizer.

Within each of the “analysis”/non-“analysis” subsets, the dataset is then downsampled by randomly sampling each of the three label-plus-negation groups closed under the contradiction operation (Entail+negation plus Contradict+non-negation; Contradict+negation plus Entail+non-negation; Neutral+negation plus Neutral+non-negation) down to the minimum of their example counts. A 90/5/5 train/val/test split is stratified by “analysis”/non-“analysis” subset and these groups.

Each example is then complemented with its contradictory twin: the same premise paired with the hypothesis modified by adding or removing pivotal negation (so applying the contradiction algorithm). Neutral labels are unchanged from the original example, while Entail and Contradict labels are flipped. This twinning balances the dataset within the “analysis”/non-“analysis” subsets on NLI label by pivotal negation versus not.

A.1.9 Citation Removal Algorithm and Prepending

Our algorithm here attempts to remove as many in-line citations from premises as it can so that the premises are more customary English-language texts. The processed premises are studied in this

shows retrieval using dot-product similarity on this model’s embeddings to perform best among several models on TREC-DL 2019 ([Craswell et al., 2020](#)) and the MS Marco Passage Retrieval dataset ([Bajaj et al., 2016](#)).

paper. But the dataset obtainable from code to be released will include the pre-processing premises as well for future study. Finally, we copy and prepend at the beginning of the long premises the minimum number of paragraphs from the end that contain 512 tokens, to limit models from relying on cues for the NLI label near the start.

A.2 Implementation Details

External code is from GitHub repositories, with repository forking permitted under contemporaneous GitHub’s Terms of Service. External models are from HuggingFace Transformers (Wolf et al., 2020; contemporaneously governed by an Apache License 2.0 permitting modification, distribution, etc.) or from GitHub repositories. Cases from the Caselaw Access Project (The President and Fellows of Harvard University, 2018) are used to construct our datasets. Any dataset sharing will comply with Caselaw Access Project (The President and Fellows of Harvard University, 2018) terms of access or else any separate agreement with the licensor. In particular, if necessary to ensure this compliance, we will share code for constructing our datasets rather than the datasets themselves.

NVIDIA 12GB TITAN Xp, 11GB GeForce GTX 1080 Ti, 11GB GeForce RTX 2080 Ti, 24GB TITAN RTX GPUs, and NVIDIA 48GB RTX A6000 GPUs were used for all fine-tuning.

A.2.1 Evaluation 1

For our intermediate fine-tuning, we adapt the code and largely follow the respective model hyperparameters and fine-tuning settings of the three existing NLI benchmarks. The settings that we modify rather than follow are: attention gradient checkpointing, GPU setup while not changing accumulated batch size, and maximum sequence length (given our long sequence lengths, we also train for 3 epochs instead of 5 on DocNLI (Yin et al., 2021)). Maximum sequence lengths for intermediate fine-tuning are the lesser of the model maximum and 2048 (except for a maximum sequence length of 156 for pretrained short-sequence models fine-tuned on ANLI, consistent with Nie et al. (2020)²⁴).

After intermediate fine-tuning, the long-sequence models’ maximum sequence lengths are increased to 4096 for further fine-tuning on LawngNLI. We adapt the code from Xiong et al.

(2021).²⁵ We adapted this code in order to allow compatibility with their suite of efficient Transformers, but ultimately we did not pretrain them and did not further explore including them after several (initialized with copied RoBERTa-base (Liu et al., 2019) embeddings) did not rise far above random accuracy for LawngNLI fine-tuning under some initial hyperparameters explored. This reflects little on these models since we did not pretrain them.

For fine-tuning on LawngNLI in our NLI experiments, we use a batch size of 32, learning rate of $1e-5$, 4 epochs (2 epochs for DocNLI; see next paragraph), learning rate schedule adapted from Xiong et al. (2021) (Adam optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$, L2 weight decay of 0.01, warm-up over first 10,000 steps, and linear decay), and half precision. We explored hyperparameters among those explored by RoBERTa (Liu et al., 2019) for GLUE (Wang et al., 2018), along with batch size 128 so that all of our models in Appendix Section A.4 would start to converge during fine-tuning starting from their initial losses and accuracies.

To transfer learning from two-label DocNLI, the models intermediate-fine-tuned on DocNLI are further fine-tuned and evaluated on a two-label version of LawngNLI (where the Entail examples are duplicated and then (Entail, Neutral and Contradict) labels are mapped to (Entail, Not Entail)). This construction balances the two-label version between (Entail, Not Entail). For further fine-tuning these models on LawngNLI, the number of epochs is then halved. This is equivalent to splitting the Neutral and Contradict examples (now labeled Not Entail) in the original three-label dataset in half across pairs of consecutive original epochs (1 and 2, 3 and 4, and so on) so that the fine-tuning example count is $2/3$ of the original dataset times the original number of epochs, except that example shuffling also pools examples between these consecutive original epochs.

A.2.2 Existing NLI Datasets

For models in Appendix Section A.4 with fine-tuned checkpoints provided at <https://github.com/facebookresearch/anli> (ALBERT-xxlarge-v2 (Lan et al., 2019), BART-large (Lewis et al., 2020), and RoBERTa-large (Liu et al., 2019)), we used these model checkpoints. Otherwise we fine-tuned the models, aiming to

²⁴<https://github.com/facebookresearch/anli>

²⁵<https://github.com/mlpen/Nystromformer>

replicate the original hyperparameters.

A.3 Procedure for Human Assessment

Human assessment was limited to Amazon Mechanical Turk Master Workers based in the U.S.

Assessed accuracy of examples with long premises is lower than for with short premises, even though the former arguably should have a higher accuracy against the ground truth: they are a superset of the information in the short premise, thereby providing additional context while being written to be internally consistent. It may be then that the human-assessed error rates for the automatic labels are themselves imperfect against the ground truth, especially for more difficult examples.

Human assessment proceeded as follows:

- Examples were each reviewed by two workers in batches of 28 examples, which were drawn from a first and then second set of 504 examples with sequence length at most 4096. Each set consists of a stratified random sample of test examples. The stratification is as follows: First, balance over the Cartesian product of the automatic label and pivotal negation versus not. Then half using the short premise and half using the long premise.
- Workers provided NLI labels for batches effectively without a time limit (batches due 1 week after assignment). Batches were issued until there were 300 non-screening examples with their two worker labels in agreement. The accuracy of these examples' automatic labels was then evaluated against those agreed labels (as gold).
- Workers were advised that they were providing NLI labels to be used in an academic analysis evaluating a new dataset.
- A co-author provided NLI labels for a predetermined random sample as "screening" examples, *which did not enter the dataset with human-assessed gold labels or the error rate calculations and so do not directly impact them*. They were scattered throughout and not separately identified to workers. Performance formulae using the screening examples *only* were used to calculate worker bonuses and to exclude (ultimately two) workers who appeared to be guessing frequently.
- Workers were paid above the U.S. federal minimum wage on "reasonable" (as opposed to actual) time spent: 2 hours per batch, but workers may have spent more or less time on any batch up to 1 week. In addition, a performance bonus was provided for each label deemed correct on a screening example.
- Worker screening was as follows:
 - First, workers needed to qualify by answering at least 4 examples correct (credit was sometimes given for an incorrect label with defensible reasoning) on an initial pre-screen of six screening examples within a half hour. Several batches not meeting the minimum performance described in the instructions (which was itself below the qualification threshold) were rejected.
 - Because NLI is multiple choice, there is a risk that the initial screening may be insufficient or that workers may not consider examples thoroughly in selecting options (or even guess somewhat randomly). Though we saw evidence directly in the gold dataset, we included screening examples in the ongoing batches. We excluded two workers' examples for falling below a threshold.
 - * Each batch contains 3 screening examples and 25 non-screening examples.
 - * Labels on screening examples were selected by a co-author. Screening examples were not included in the 300 examples in the gold dataset.
 - * Workers could continue completing the batches of 28 unless at a time of audit their cumulative accuracy on screening examples fell below 50% (after at least 5 screening examples). If their cumulative accuracy fell below this threshold, they were still paid for all completed batches but the examples they labeled were not included in the gold dataset.
 - Workers provided labels via a six-option scale: 'definitely entail', 'probably entail', 'definitely neutral', 'probably neutral', 'definitely contradict', 'probably contradict'.

- For examples that workers labeled as entail or contradict, they also copied and pasted a portion of the premise relevant to determining the label they chose.
- We temporarily experimented with having a different version of the dataset assessed, but no workers labeled the same examples in that experiment and the current assessment set.

A.3.1 Interface for Main NLI Task

See Appendix Table 7.

A.3.2 Interface for Pre-Screen Task

See Appendix Table 8. An included illustrative example is also omitted here. Note that some earlier workers saw earlier versions.

A.4 Evaluation Panel: List of State-of-the-Art Pretrained Models

- Longformer-base (Beltagy et al., 2020)
- BigBird-RoBERTa-base (Zaheer et al.)
- ALBERT-xxlarge-v2 (Lan et al., 2019). It ranked highest on ANLI test A2 and A3²⁶. It also ranked highest besides T5 models and third overall on MNLI (Williams et al., 2018)²⁷.
- BART-large (Lewis et al., 2020). It ranked first on ConTRoL (Liu et al., 2021), after fine-tuning on ANLI (Nie et al., 2020).
- Custom Legal-BERT (Zheng et al., 2021). Pre-trained on the Caselaw Access Project (The President and Fellows of Harvard University, 2018) corpus.
- LEGAL-BERT-base-uncased, also known as LEGAL-BERT-SC (Chalkidis et al., 2020). It is pretrained on legal text from fields such as legislation, cases, and contracts.
- RoBERTa-large (Liu et al., 2019). It performed the better out of two models (over Longformer (Beltagy et al., 2020)) on DocNLI (Yin et al., 2021) and ranked second on ANLI test A1²⁸.

A.5 Appendix Tables

²⁶<https://paperswithcode.com/sota/natural-language-inference-on-anli-test>

²⁷<https://paperswithcode.com/sota/natural-language-inference-on-multinli>

²⁸<https://paperswithcode.com/sota/natural-language-inference-on-anli-test>

How is text1 related to text2: Entail, Neutral, or Contradict?

[If this task is visible to you, the batch of items should be new to you and you may complete the task.]

I. OVERVIEW

6 items limited to 30 minutes. The reward is \$3.75, with a bonus for 5 or 6 answers deemed correct for a total reward of \$4 or \$4.50, respectively.

For EACH item you answer, SELECT ONE LABEL AMONG THE THREE (there is exactly one reference answer for EACH item) and PASTE A SNIPPET FROM TEXT1 SUPPORTING THE LABEL. Any item with multiple or zero labels selected or not accompanied by relevant snippet from text1 (see II) will be marked incorrect.

Feel free to use the Find function to search the text. Some text1 portions standing alone could suggest an incorrect label and yet should be consistent with a correct label in the context of text1 overall. For example, if a text mentions a claim by a lower court or different case and then proceeds to reject that claim, then the text overall rejects that claim.

text2 may well refer to text1, and any of the three labels may apply. Note that in some items though, text2 may not be referring to text1 at all (e.g., it may be discussing a scenario from a text not included here, with different entities or conclusions neither entailed nor contradicted by text1). In the latter cases, the correct label would be Neutral.

Assignments with too few answers deemed correct (e.g., two or fewer) that are not accompanied by evidence of diligence in reading passages (relevant snippet from text1) may be rejected.

II. RESPONSES

(1 - Required)
Select the label that best describes how text1 is related to text2.

(2 - Required if Entail or Contradict label is provided, paste relevant snippet from text1)
This text can be short or incomplete. It is simply to help demonstrate reasonable diligence, only for if that is not already clear from the labels. A tiny fraction of the time should be spent on this step.

(3 - Optional - include explanation for your label)
Answers matching reference label will be given full credit regardless of if an explanation is provided. For other labels, it may depend on explanation.

III. LABELS (Entail, Contradict, and Neutral)

text1 in this task may be substantially longer than below. See "Instructions" for examples. text1 and text2 are sometimes known as premise and hypothesis, respectively.

Quoting from https://aclweb.org/aclwiki/index.php?title=Textual_Entailment_Portal and <http://u.cs.biu.ac.il/~dagan/publications/RTEchallenge.pdf>:

"An example of a positive TE (text entails hypothesis) is:
text: If you help the needy, God will reward you.
hypothesis: Giving money to a poor man has good consequences.

An example of a negative TE (text contradicts hypothesis) is:
text: If you help the needy, God will reward you.
hypothesis: Giving money to a poor man has no consequences.

An example of a non-TE (text does not entail nor contradict) is:
text: If you help the needy, God will reward you.
hypothesis: Giving money to a poor man will make you better person.

The entailment need not be pure logical - it has a more relaxed definition: "t entails h (t = h) if, typically, a human reading t would infer that h is most likely true." [1]"

Instructions
Shortcuts

Entail
 Neutral
 Contradict

text1: \${text1}
 text2: \${text2}

Relevant snippet from text1 (Required if Entail or Contradict label is provided.)

Lorem ipsum...

Explanation for your label (optional)

Lorem ipsum...

Table 7: Interface for the main NLI task.

How is text1 related to text2: Entail, Neutral, or Contradict?

****To qualify for this task, you must first perform sufficiently well on the [SCREENING TASK], listed under the same requester**.**

This task is natural language inference/RTE (same as under our [SCREENING TASK]).

Note: In completing this task, you are providing NLI labels to be used in an academic analysis evaluating a new dataset. These items were constructed from legal texts and probabilistically selected from a larger dataset without screening. They may include sensitive or derogatory language.

Items are in batches of 26. The time limit is one week, so that you can spend more time on certain difficult items if you wish.

I. REWARD

The reward per batch was calculated based on two hours of work. However, feel free to work at your own pace as the time limit is much longer. As a percent of the base reward, there is an overall bonus of 20% times your accuracy rate on "validation items" (see II), paid after you are finished completing all your batches.

II. EVALUATION

Assignments that are not accompanied by evidence of diligence in reading passages (snippets from text1 that are relevant) may be rejected. If your running average accuracy on "validation items" falls too low, you may no longer be able to access further batches.

"Validation items" are a random subset of the items that are separately labeled: correct labels are determined independent of your response, though these labels will not be visible to you.

III. INSTRUCTIONS

Feel free to use the Find function to search the text. Some text1 portions standing alone could suggest an incorrect label and yet should be consistent with a correct label in the context of text1 overall. For example, if a text mentions a claim by a lower court or different case and then proceeds to reject that claim, then the text overall rejects that claim.

text2 may well refer to text1, and any of the three labels may apply. Note that in some items though, text2 may not be referring to text1 at all (e.g., it may be discussing a scenario from a text not included here, with different entities or conclusions neither entailed nor contradicted by text1). In the latter cases, the correct label would be Neutral.

IV. RESPONSES

(1 - Required)
SELECT EXACTLY ONE LABEL AMONG THE THREE that best describes how text1 is related to text2. As part of the label, select a level of confidence: PROBABLY or DEFINITELY (this confidence level will in no way affect the evaluation or reward, so you can just be honest).

(2 - Required if Entail or Contradict label is provided, PASTE A SNIPPET FROM TEXT1 SUPPORTING THE LABEL)
This text can be short or incomplete. It is simply to help demonstrate reasonable diligence, only for if that is not already clear from the labels. A tiny fraction of the time should be spent on this step.

definitely entail definitely neutral definitely contradict
 probably entail probably neutral probably contradict

Relevant snippet from text1 (Required if Entail or Contradict label is provided.)

Lorem ipsum...

text1: \${text1}

text2: \${text2}

Table 8: Interface for the pre-screen task. An included illustrative example is also omitted here. Note that some earlier workers saw earlier versions.

Sample twin Entail/Contradict examples with same premise from LawngNLI	
Twin hypotheses with same premise, from “analysis” subset	<ul style="list-style-type: none"> • <i>Contradict</i>: city acted affirmatively to create or increase risk of harm on city street by ignoring residents’ requests to reduce speed limit or by taking down residents’ signs indicating drivers should adhere to a lower speed limit • <i>Entail</i>: city did not act affirmatively to create or increase risk of harm on city street by ignoring residents’ requests to reduce speed limit or by taking down residents’ signs indicating drivers should adhere to a lower speed limit
Some additional hypotheses with same premise	<ul style="list-style-type: none"> • <i>Entail</i>: failing to enforce or lower the speed limit on a residential street “did not create a ‘special danger’ to a discrete class of individuals..[ed.: excerpted]..as opposed to a general traffic risk to pedestrians and other automobiles” • <i>Contradict</i>: traffic laws and enforcement practices did not pose “a general traffic risk to pedestrians and other automobiles”
Relevant excerpts of shared premise	<ul style="list-style-type: none"> • [ed.: Plaintiffs] ...submit that the City of Fort Thomas..violated their son’s substantive due process rights by failing to act upon their request (and the requests of others) to lower the speed limit on the street..The police also removed signs posted by residents indicating that drivers should adhere to a 15 mile-per-hour speed limit.. • [ed.: Plaintiffs] ...alleged that the City’s failure to maintain safe conditions on Garrison Avenue violated their son’s substantive due process rights..established a “state-created danger” under DeShaney.. • ...DeShaney’s holding..precludes [ed.: Plaintiffs]’ argument that the Due Process Clause constitutionalizes a locality’s choices about what speed limit to adopt for a given street or how to enforce that speed limit.. • There are two exceptions to the DeShaney rule..Under the second exception..a plaintiff may bring a substantive due process claim by establishing (1) an affirmative act by the State that either created or increased the risk that the plaintiff would be exposed to private acts of violence.. • [ed.: Plaintiffs] fail to satisfy any of the three requirements for establishing our circuit’s “state-created danger” exception to DeShaney. First, the creation of a street and the management of traffic conditions on that street are too attenuated and indirect to count as an “affirmative act”..
Distractor excerpts of same premise	<ul style="list-style-type: none"> • ...After all, the City was told about the risks of not lowering the speed limit to 15 miles per hour (more accidents); it intentionally chose not to heed this warning (taking on the risk of more accidents); and the alleged risk came to pass when..was killed (an accident).. • ...For in one sense, it could be said that all governing bodies act with deliberate indifference when they consider and reject a traffic-safety proposal of this sort that comes with known risks..

Table 9: Sample twin Entail/Contradict examples with same premise from LawngNLI, also in the “analysis” subset analyzed in our experiments (Section 4): sequence length of long premise at most 4096. Each hypothesis pairs with the excerpted premise in a separate example. For those specific “additional hypotheses” above, the examples containing them are in unfiltered-LawngNLI2 (see GitHub link in first footnote) but not LawngNLI, the core dataset studied in this paper. See also Table 1.

Sample twin Neutral examples with same premise from LawngNLI	
Twin hypotheses with same premise, from “analysis” subset	<ul style="list-style-type: none"> • <i>Neutral</i>: a parade permit requirement did not violate the First Amendment • <i>Neutral</i>: a parade permit requirement violated the First Amendment
Distractor excerpts of same premise	<ul style="list-style-type: none"> • ...Section 13k prohibits two distinct activities: it is unlawful either “to parade, stand, or move in processions or assemblages in the Supreme Court Building or grounds,”.. • ...we shall address only whether the proscriptions of 13k are constitutional as applied to the public sidewalks..

Table 10: Sample twin Neutral examples from LawngNLI, but not in the “analysis” subset analyzed in our experiments (Section 4): sequence length of long premise at most 4096. Each hypothesis pairs with the excerpted premise in a separate example. See also Table 1.

Example 1 (human Neutral, automatic Entail)	<ul style="list-style-type: none"> • <i>Relevant premise excerpts</i>: Mebane’s sentence was clearly imposed by a court with jurisdiction, and his sentence is unambiguous. His sentence conforms to the statutory provisions regarding class A, B, and C felonies and the Habitual Criminal Act. • <i>Hypothesis</i>: claim that multiple sentences arose from single wrongful act and violated Double Jeopardy Clause does not establish that sentence is illegal
Example 2 (human Entail, automatic Contradict)	<ul style="list-style-type: none"> • <i>Relevant premise excerpts</i>: Moreover,(a)(6) states that “[a] report of an interdisciplinary team that contains the evaluation and signature of an acceptable medical source is also considered acceptable medical evidence,” while later in that section the statute designates nurse practitioners as an “other source.” 416.913(e)(3). While nowhere in the regulations is the term “interdisciplinary team” expressly defined, a plain reading of these sections taken together indicates that a nurse practitioner working in conjunction with a physician constitutes an acceptable medical source, while a nurse practitioner working on his or her own does not. • <i>Hypothesis</i>: non-medical source must not be integral to team, and the acceptable medical source must undersign her findings
Example 3 (human Contradict, automatic Neutral)	<ul style="list-style-type: none"> • <i>Relevant premise excerpts</i>: Where the BIA adopts the decision of the IJ, we review the IJ’s decision. The standard of review for factual findings made by the IJ is a deferential substantial evidence standard, and credibility findings will be upheld unless the evidence compels a contrary result. This deferential standard of review “precludes relief absent a conclusion that no reasonable factfinder could have reached the agency’s result.” <i>Thangaraja v. Gonzales</i>. • <i>Hypothesis</i>: the court may not make independent credibility determinations

Table 11: Several examples with differing automatic and gold labels.

Fine-tuning and evaluation	Hypothesis only	Long premise only (after BM25 filtering)	Long premise only	Short premise only (after BM25 filtering)	Short premise only
[Entail/Neutral/Contradict. Chance=1/3]					
albert-xxlarge-v2_anli	0.512+/-0.016	0.438+/-0.016	0.433+/-0.016	0.491+/-0.016	0.496+/-0.016
albert-xxlarge-v2_ConTRoL-dataset	0.525+/-0.016	0.437+/-0.016	0.436+/-0.016	0.496+/-0.016	0.489+/-0.016
albert-xxlarge-v2_vanilla	0.541+/-0.016	0.434+/-0.016	0.433+/-0.016	0.485+/-0.016	0.489+/-0.016
facebook_bart-large_anli	0.544+/-0.016	0.447+/-0.016	0.45+/-0.016	0.505+/-0.016	0.511+/-0.016
facebook_bart-large_ConTRoL-dataset	0.542+/-0.016	0.439+/-0.016	0.425+/-0.016	0.509+/-0.016	0.514+/-0.016
facebook_bart-large_vanilla	0.488+/-0.016	0.409+/-0.015	0.437+/-0.016	0.517+/-0.016	0.513+/-0.016
N	3966				

Table 12: Accuracy of baselines on LawngNLI’s “analysis” subset: Top 6 three-label models on both short and long premises from Appendix Table 13. The error provided is the larger deviation of the Clopper-Pearson (Clopper and Pearson, 1934) exact binomial 95% confidence bounds.

Needs long premises for fine-tuning	No				Yes	
Fine-tuning	Short premise		BM25 retrieval on short premise		Long premise	Long premise filtered by BM25 retrieval
Long premise filtered by BM25 retrieval at evaluation	No	No [512 [tokens]	Yes	Yes	No	Yes
	(1)	(2)	(3)	(4)	(5)	(6)
[Entail/Neutral/Contradict. Chance=1/3]						
albert-xxlarge-v2_anli	0.742+/-0.014	0.742+/-0.014	0.81+/-0.013	0.817+/-0.012	0.789+/-0.013	0.853+/-0.011
albert-xxlarge-v2_ConTRoL-dataset	0.725+/-0.014	0.725+/-0.014	0.804+/-0.013	0.806+/-0.013	0.781+/-0.013	0.851+/-0.011
albert-xxlarge-v2_vanilla	0.743+/-0.014	0.743+/-0.014	0.816+/-0.012	0.825+/-0.012	0.781+/-0.013	0.852+/-0.011
allenai_longformer-base-4096_anli	0.562+/-0.016	0.652+/-0.015	0.706+/-0.014	0.707+/-0.014	0.691+/-0.015	0.784+/-0.013
allenai_longformer-base-4096_ConTRoL-dataset	0.56+/-0.016	0.641+/-0.015	0.698+/-0.015	0.691+/-0.015	0.693+/-0.015	0.781+/-0.013
allenai_longformer-base-4096_vanilla	0.541+/-0.016	0.604+/-0.015	0.691+/-0.015	0.701+/-0.015	0.51+/-0.016	0.791+/-0.013
facebook_bart-large_anli	0.664+/-0.015	0.675+/-0.015	0.808+/-0.013	0.809+/-0.013	0.76+/-0.014	0.868+/-0.011
facebook_bart-large_ConTRoL-dataset	0.665+/-0.015	0.666+/-0.015	0.809+/-0.013	0.814+/-0.012	0.758+/-0.014	0.868+/-0.011
facebook_bart-large_vanilla	0.667+/-0.015	0.699+/-0.015	0.813+/-0.012	0.807+/-0.013	0.782+/-0.013	0.875+/-0.011
google_bigbird-roberta-base_anli	0.613+/-0.015	0.666+/-0.015	0.764+/-0.014	0.774+/-0.013	0.77+/-0.013	0.829+/-0.012
google_bigbird-roberta-base_ConTRoL-dataset	0.601+/-0.015	0.656+/-0.015	0.761+/-0.014	0.769+/-0.013	0.757+/-0.014	0.819+/-0.012
google_bigbird-roberta-base_vanilla	0.648+/-0.015	0.657+/-0.015	0.767+/-0.013	0.777+/-0.013	0.763+/-0.014	0.827+/-0.012
nlpaueb_legal-bert-base-uncased_anli	0.709+/-0.014	0.709+/-0.014	0.777+/-0.013	0.791+/-0.013	0.767+/-0.013	0.811+/-0.013
nlpaueb_legal-bert-base-uncased_ConTRoL-dataset	0.721+/-0.014	0.721+/-0.014	0.766+/-0.014	0.787+/-0.013	0.761+/-0.014	0.816+/-0.012
nlpaueb_legal-bert-base-uncased_vanilla	0.729+/-0.014	0.729+/-0.014	0.776+/-0.013	0.788+/-0.013	0.767+/-0.013	0.813+/-0.012
roberta-large_anli	0.716+/-0.014	0.716+/-0.014	0.778+/-0.013	0.799+/-0.013	0.778+/-0.013	0.842+/-0.012
roberta-large_ConTRoL-dataset	0.707+/-0.014	0.707+/-0.014	0.759+/-0.014	0.789+/-0.013	0.761+/-0.014	0.838+/-0.012
roberta-large_vanilla	0.718+/-0.014	0.718+/-0.014	0.777+/-0.013	0.801+/-0.013	0.766+/-0.013	0.843+/-0.012
zlucaia_custom-legalbert_anli	0.723+/-0.014	0.723+/-0.014	0.794+/-0.013	0.805+/-0.013	0.776+/-0.013	0.824+/-0.012
zlucaia_custom-legalbert_ConTRoL-dataset	0.727+/-0.014	0.727+/-0.014	0.792+/-0.013	0.802+/-0.013	0.782+/-0.013	0.824+/-0.012
zlucaia_custom-legalbert_vanilla	0.732+/-0.014	0.732+/-0.014	0.785+/-0.013	0.798+/-0.013	0.783+/-0.013	0.832+/-0.012
Maximum of p-values versus (6)	0	0	0	0		
N				3966		

Table 13: Accuracy on LawngNLI’s “analysis” subset *with long premises*: Full intermediate-fine-tuned model panel. Pretrained models are chosen for top performance on existing NLI benchmarks (see Appendix Section A.4 for versions). The error provided is the larger deviation of the Clopper-Pearson (Clopper and Pearson, 1934) exact binomial 95% confidence bounds. All p-values round to zero (<0.0005) from an exact binomial McNemar’s (McNemar, 1947) test for a statistically significant difference in accuracies between each version fine-tuning using short premises as inputs (1-4) and the best version fine-tuning using long premises as inputs (6). For (2), 512 tokens is the overall sequence limit.

Evaluation	Long Premise				
	Fine-tuning (on same)	No	Yes		
	Subset	All	All	Above-Median Long Premise Length	Hypothesis Contains Negation
[Entail/Neutral/Contradict. Chance=1/3]					
albert-xxlarge-v2_anli	0.501+/-0.016	0.789+/-0.013	0.738+/-0.02	0.789+/-0.019	
albert-xxlarge-v2_ConTRoL-dataset	0.434+/-0.016	0.781+/-0.013	0.732+/-0.02	0.787+/-0.019	
albert-xxlarge-v2_vanilla	0.345+/-0.015	0.781+/-0.013	0.733+/-0.02	0.784+/-0.019	
allenai_longformer-base-4096_anli	0.367+/-0.015	0.691+/-0.015	0.652+/-0.022	0.695+/-0.021	
allenai_longformer-base-4096_ConTRoL-dataset	0.355+/-0.015	0.693+/-0.015	0.664+/-0.022	0.692+/-0.021	
allenai_longformer-base-4096_vanilla	0.334+/-0.015	0.51+/-0.016	0.462+/-0.023	0.511+/-0.022	
facebook_bart-large_anli	0.345+/-0.015	0.76+/-0.014	0.695+/-0.021	0.767+/-0.019	
facebook_bart-large_ConTRoL-dataset	0.407+/-0.015	0.758+/-0.014	0.699+/-0.021	0.764+/-0.019	
facebook_bart-large_vanilla	0.336+/-0.015	0.782+/-0.013	0.723+/-0.021	0.789+/-0.019	
google_bigbird-roberta-base_anli	0.342+/-0.015	0.77+/-0.013	0.74+/-0.02	0.773+/-0.019	
google_bigbird-roberta-base_ConTRoL-dataset	0.354+/-0.015	0.757+/-0.014	0.735+/-0.02	0.757+/-0.019	
google_bigbird-roberta-base_vanilla	0.335+/-0.015	0.763+/-0.014	0.747+/-0.02	0.761+/-0.019	
nlpaueb_legal-bert-base-uncased_anli	0.478+/-0.016	0.767+/-0.013	0.725+/-0.02	0.775+/-0.019	
nlpaueb_legal-bert-base-uncased_ConTRoL-dataset	0.423+/-0.016	0.761+/-0.014	0.726+/-0.02	0.768+/-0.019	
nlpaueb_legal-bert-base-uncased_vanilla	0.336+/-0.015	0.767+/-0.013	0.725+/-0.02	0.769+/-0.019	
roberta-large_anli	0.353+/-0.015	0.778+/-0.013	0.727+/-0.02	0.781+/-0.019	
roberta-large_ConTRoL-dataset	0.429+/-0.016	0.761+/-0.014	0.717+/-0.021	0.763+/-0.019	
roberta-large_vanilla	0.333+/-0.015	0.766+/-0.013	0.712+/-0.021	0.772+/-0.019	
zlucaia_custom-legalbert_anli	0.499+/-0.016	0.776+/-0.013	0.736+/-0.02	0.779+/-0.019	
zlucaia_custom-legalbert_ConTRoL-dataset	0.445+/-0.016	0.782+/-0.013	0.753+/-0.02	0.783+/-0.019	
zlucaia_custom-legalbert_vanilla	0.49+/-0.016	0.783+/-0.013	0.753+/-0.02	0.787+/-0.019	
N	3966	3966	1936	1983	
[Entail/Not entail. Chance=1/2]					
albert-xxlarge-v2_DocNLI	0.422+/-0.013	0.847+/-0.01	0.804+/-0.016	0.872+/-0.013	
allenai_longformer-base-4096_DocNLI	0.5+/-0.014	0.777+/-0.011	0.76+/-0.017	0.814+/-0.015	
facebook_bart-large_DocNLI	0.497+/-0.014	0.641+/-0.013	0.602+/-0.019	0.64+/-0.019	
google_bigbird-roberta-base_DocNLI	0.513+/-0.014	0.817+/-0.011	0.798+/-0.016	0.843+/-0.014	
nlpaueb_legal-bert-base-uncased_DocNLI	0.48+/-0.014	0.833+/-0.01	0.814+/-0.015	0.85+/-0.014	
roberta-large_DocNLI	0.448+/-0.014	0.843+/-0.01	0.82+/-0.015	0.877+/-0.013	
zlucaia_custom-legalbert_DocNLI	0.557+/-0.014	0.822+/-0.011	0.801+/-0.016	0.842+/-0.014	
N	5288	5288	2637	2644	

Table 14: Accuracy on LawngNLI’s “analysis” subset: Full intermediate-fine-tuned model panel. If fine-tuned, fine-tuning on the LawngNLI subset is on premises with the same granularity as evaluation. The error provided is the larger of the two deviations of the Clopper-Pearson (Clopper and Pearson, 1934) exact binomial 95% confidence bounds from the point estimate.

Evaluation	Short Premise			
	Fine-tuning (on same)		Yes	
	Subset	No	All	Above-Median Long Premise Length
[Entail/Neutral/Contradict. Chance=1/3]				
albert-xxlarge-v2_anli	0.551+/-0.016	0.882+/-0.01	0.866+/-0.016	0.885+/-0.015
albert-xxlarge-v2_ConTRoL-dataset	0.478+/-0.016	0.878+/-0.011	0.858+/-0.016	0.881+/-0.015
albert-xxlarge-v2_vanilla	0.336+/-0.015	0.883+/-0.01	0.861+/-0.016	0.883+/-0.015
allenai_longformer-base-4096_anli	0.402+/-0.015	0.802+/-0.013	0.767+/-0.019	0.809+/-0.018
allenai_longformer-base-4096_ConTRoL-dataset	0.376+/-0.015	0.791+/-0.013	0.762+/-0.02	0.793+/-0.019
allenai_longformer-base-4096_vanilla	0.339+/-0.015	0.779+/-0.013	0.741+/-0.02	0.784+/-0.019
facebook_bart-large_anli	0.532+/-0.016	0.879+/-0.011	0.862+/-0.016	0.878+/-0.015
facebook_bart-large_ConTRoL-dataset	0.468+/-0.016	0.876+/-0.011	0.861+/-0.016	0.879+/-0.015
facebook_bart-large_vanilla	0.323+/-0.015	0.89+/-0.01	0.874+/-0.016	0.893+/-0.014
google_bigbird-roberta-base_anli	0.403+/-0.015	0.84+/-0.012	0.825+/-0.018	0.838+/-0.017
google_bigbird-roberta-base_ConTRoL-dataset	0.383+/-0.015	0.845+/-0.012	0.824+/-0.018	0.845+/-0.017
google_bigbird-roberta-base_vanilla	0.354+/-0.015	0.842+/-0.012	0.824+/-0.018	0.841+/-0.017
nlpaueb_legal-bert-base-uncased_anli	0.514+/-0.016	0.849+/-0.012	0.829+/-0.018	0.845+/-0.017
nlpaueb_legal-bert-base-uncased_ConTRoL-dataset	0.471+/-0.016	0.839+/-0.012	0.818+/-0.018	0.833+/-0.017
nlpaueb_legal-bert-base-uncased_vanilla	0.346+/-0.015	0.853+/-0.011	0.831+/-0.017	0.854+/-0.016
roberta-large_anli	0.374+/-0.015	0.884+/-0.01	0.867+/-0.016	0.886+/-0.015
roberta-large_ConTRoL-dataset	0.478+/-0.016	0.872+/-0.011	0.85+/-0.017	0.877+/-0.015
roberta-large_vanilla	0.4+/-0.015	0.887+/-0.01	0.868+/-0.016	0.89+/-0.015
zlucaia_custom-legalbert_anli	0.536+/-0.016	0.843+/-0.012	0.826+/-0.018	0.84+/-0.017
zlucaia_custom-legalbert_ConTRoL-dataset	0.462+/-0.016	0.839+/-0.012	0.821+/-0.018	0.837+/-0.017
zlucaia_custom-legalbert_vanilla	0.267+/-0.014	0.844+/-0.012	0.834+/-0.017	0.848+/-0.017
N	3966	3966	1936	1983
[Entail/Not entail. Chance=1/2]				
albert-xxlarge-v2_DocNLI	0.508+/-0.014	0.907+/-0.008	0.896+/-0.012	0.913+/-0.011
allenai_longformer-base-4096_DocNLI	0.496+/-0.014	0.834+/-0.01	0.813+/-0.015	0.855+/-0.014
facebook_bart-large_DocNLI	0.531+/-0.014	0.874+/-0.009	0.865+/-0.014	0.889+/-0.013
google_bigbird-roberta-base_DocNLI	0.508+/-0.014	0.863+/-0.01	0.849+/-0.014	0.875+/-0.013
nlpaueb_legal-bert-base-uncased_DocNLI	0.466+/-0.014	0.873+/-0.009	0.859+/-0.014	0.89+/-0.013
roberta-large_DocNLI	0.412+/-0.013	0.91+/-0.008	0.897+/-0.012	0.927+/-0.011
zlucaia_custom-legalbert_DocNLI	0.616+/-0.013	0.857+/-0.01	0.843+/-0.014	0.873+/-0.013
N	5288	5288	2637	2644

Table 15: Accuracy on LawngNLI’s “analysis” subset: Full intermediate-fine-tuned model panel. If fine-tuned, fine-tuning on the LawngNLI subset is on premises with the same granularity as evaluation. The error provided is the larger of the two deviations of the Clopper-Pearson (Clopper and Pearson, 1934) exact binomial 95% confidence bounds from the point estimate.

Needs long premises for fine-tuning	No		Yes	
Fine-tuning	BM25 retrieval on short premise		Long premise filtered by BM25 retrieval	
Long premise filtered by BM25 retrieval at evaluation	Yes		Yes	
Label	Automatic	Gold	Automatic	Gold
	(4)	(4)	(6)	(6)
[Entail/Neutral/Contradict. Chance=1/3]				
albert-xxlarge-v2_anli	0.668	0.688	0.684	0.689
albert-xxlarge-v2_ConTRoL-dataset	0.674	0.673	0.629	0.65
albert-xxlarge-v2_vanilla	0.685	0.677	0.646	0.636
facebook_bart-large_anli	0.558	0.537	0.683	0.68
facebook_bart-large_ConTRoL-dataset	0.614	0.604	0.672	0.67
facebook_bart-large_vanilla	0.622	0.583	0.672	0.699
N	140			

Table 16: Balanced accuracy/macro-averaged recall: Top 6 three-label models from RQ1 (Appendix Table 13), evaluated against the automatic labels versus gold labels on the human-validated subset. Column numbers are the same as in Appendix Table 13. As described in Section 4.1.2, for the gaps between the top model setups that need long premises for fine-tuning (column 6) versus those that do not (column 4), there are non-positive biases (automatic labels versus gold labels) here both for the model with the largest performance gap from Table 5 and on average across the 6 models. Thus this comparison provides evidence that the performance gap from RQ1 is not due to bias from error in the automatic labels.

Evaluation	Long premise			Long premise filtered by BM25 retrieval		
	1	10	100	1	10	100
[Entail hypotheses as queries]						
all-distilroberta-v1	0.039+/-0.012	0.108+/-0.018	0.262+/-0.025	0.075+/-0.016	0.219+/-0.023	0.505+/-0.027
all-mpnet-base-v2	0.042+/-0.012	0.115+/-0.018	0.224+/-0.023	0.083+/-0.016	0.219+/-0.023	0.433+/-0.027
nli-distilroberta-base-v2	0.005+/-0.006	0.035+/-0.011	0.132+/-0.019	0.014+/-0.008	0.063+/-0.014	0.243+/-0.024
msmarco-distilroberta-base-v2	0.02+/-0.009	0.053+/-0.013	0.152+/-0.021	0.045+/-0.013	0.14+/-0.02	0.359+/-0.027
BM25	0.138+/-0.02	0.281+/-0.025	0.447+/-0.027	0.158+/-0.021	0.308+/-0.026	0.458+/-0.027
[Contradict hypotheses as queries]						
all-distilroberta-v1	0.037+/-0.012	0.1+/-0.017	0.263+/-0.025	0.069+/-0.015	0.22+/-0.023	0.486+/-0.027
all-mpnet-base-v2	0.039+/-0.012	0.112+/-0.018	0.222+/-0.023	0.078+/-0.016	0.215+/-0.023	0.43+/-0.027
nli-distilroberta-base-v2	0.006+/-0.006	0.03+/-0.011	0.13+/-0.019	0.01+/-0.007	0.063+/-0.014	0.238+/-0.024
msmarco-distilroberta-base-v2	0.018+/-0.009	0.054+/-0.014	0.151+/-0.02	0.04+/-0.012	0.143+/-0.02	0.359+/-0.027
BM25	0.141+/-0.02	0.267+/-0.025	0.429+/-0.027	0.145+/-0.02	0.286+/-0.025	0.452+/-0.027
N	1322					

Table 17: Recall@k of SentenceTransformers (Reimers and Gurevych, 2021) bi-encoders and BM25 (Robertson and Zaragoza, 2009) baseline evaluated on implication-based retrieval (of each LawngNLI test set premise case using hypothesis as query, among 999 other cases similar to the query; see Section 4.2), using a dataset derived from LawngNLI. The error provided is the larger of the two deviations of the Clopper-Pearson (Clopper and Pearson, 1934) exact binomial 95% confidence bounds from the point estimate.