Health Insurance Coverage Rule Interpretation Corpus: Law, Policy, and Medical Guidance for Health Insurance Coverage Understanding

Mike Gartner mike@persius.org

ABSTRACT

U.S. health insurance is complex, and inadequate understanding and limited access to justice have dire implications for the most vulnerable. Advances in natural language processing present an opportunity to support efficient, case-specific understanding, and to improve access to justice and healthcare. Yet existing corpora lack context necessary for assessing even simple cases. We collect and release a corpus of reputable legal and medical text related to U.S. health insurance. We also introduce an outcome prediction task for health insurance appeals designed to support regulatory and patient self-help applications, and release a labeled benchmark for our task, and models trained on it. This is a work in progress, and the datasets, benchmark, annotations, and experimental results are all preliminary and being actively developed.

KEYWORDS

Machine Learning, Natural Language Processing, Health Insurance, Access to Justice, Outcome Prediction

1 INTRODUCTION

Health insurance coverage issues in the U.S. pose serious problems for patients. They cause delays, forgone care, and detrimental health outcomes. Sometimes patients receive services despite denials, incurring bills, debt, and collections lawsuits. For the most vulnerable, there are calamitous effects [9, 27, 39].

Evaluating coverage rules requires precise understanding of a complex web of case-specific laws, contracts, and medical literature. In cases where existing automation does not exist, fails, or is disputed, specialists such as medical coders, pharmacists, doctors, and attorneys do such work manually. The processes are time-consuming and error-prone, and successful navigation often involves tasks like synthesis of medical literature and contracts.

Advances in natural language understanding (NLU) present an opportunity to improve patient outcomes. In particular, extractive and generative models can increase the efficiency with which specialists understand cases. Tools that support efficient and accurate coverage understanding stand to serve patients, caseworkers, and regulators alike. They can lower unnecessary administrative costs, and improve both health and financial outcomes.

However, there are barriers preventing responsible and effective use of such tools. Most notably, there is a dearth of high-quality, curated corpora containing information sufficient to correctly answer coverage questions. We make four contributions. First, we collect and release a corpus of text that informs understanding of health insurance coverage rules. Second, we present an 'appeal adjudication' task. Third, we curate and release an annotated dataset to serve as a benchmark for the task. Finally, we train, evaluate, and release baseline models using our benchmark.

The datasets and code for all of our contributions have been made publicly available¹ under a permissive license. This work was previously accepted and presented as a workshop paper at the AI for Access to Justice workshop at JURIX 2024.

2 RELATED WORK

Pretraining Corpora. Domain adaptive pretraining can improve performance on downstream tasks [17, 50]. High quality pretraining corpora for both legal and medical applications abound.

In the legal domain, there is a rich history of pretraining models via masked language objectives [5, 6, 12, 19, 32, 34].

The medical domain shares a similar history [1, 8, 14, 22].

Recent years have seen an increasing focus on using language models to support arbitrary downstream tasks (e.g. [32, 35]). For tasks involving complex domain specific context, high quality pretraining corpora are critical. This is especially true for legal and medical question answering.

Our corpus shares some overlap with two existing datasets: Pile of Law [19] and the Guidance dataset [8]. The overlap can be understood in detail by consulting the documentation associated with the release.

To our knowledge, this work is among the first few efforts targeting curation of pretraining data in the intersection of legal and medical domains. The only other works of which we are aware are that of [21], which focused on tasks involving personal injury cases, and that of [2] and [23] which involve health insurance coverage tasks.

Supervised Finetuning Data. There has also been a large body of work focused on curating data for specific legal and medical tasks.

In the legal domain, people have developed datasets and models to support entity and clause selection [20, 48], multiple choice question answering [7, 51], span-based question answering [40, 41, 47], extractive summarization [3, 18], retrieval [49], and document and word level classification [4, 46], among other tasks.

In the medical domain, people have developed datasets and models to support multiple choice question answering

¹Code was released at https://github.com/TPAFS/hicric.

[38], span-based question answering [24, 25], document and clause classification [1, 22], and language modeling [8, 44], among other tasks.

In this work we construct a benchmark dataset for a document classification task. We also describe how to adapt our (pseudo)-annotation pipeline to yield related, but distinct benchmarks.

Case Outcome Prediction Tasks. Our task shares qualities with existing legal judgement prediction tasks [15, 28, 31, 33].

In particular, a few tasks presented in recent literature share many similarities.

For example, our task is at a high level the same as that presented in [2]. The task there is to predict health insurance coverage outcomes from Spanish language clinical notes. Our task differs in a few critical ways. It is a three class problem, focused on English, and concerned with both expert and layperson descriptions.

Similarly, [23] introduces a general purpose medical language model, and a health insurance denial task. Like [2], it focuses on making predictions from clinical notes. The authors did not release model weights nor training data, so the results are not reproducible, and the model is not usable, by anyone but the authors. Our work presents a distinct task, and a completely open source approach. We will release the training data and code, experimental logs, and model weights under permissive licenses.

Finally, we note that outcome prediction tasks in the literature are often actually retrospective outcome classification tasks [29]. The tasks, benchmarks, and models do not support real world forecasting. Instead, they 'predict' case outcomes from text which is not and could not be produced before adjudication. Our focus is on real world forecasting applications, and our methods reflect that focus.

3 A CORPUS

Our dataset consists of documents from diverse sources. It includes U.S. federal and state law, insurance contracts, official regulatory guidance, agency opinions and policy briefs, official coverage rules for Medicaid and Medicare, and summaries of appeal adjudications. In total, the dataset contains 8,311 documents, 419 million words, and 2.7 billion characters. Uncompressed, it occupies 2.8 gigabytes of disk space.

In the realm of pretraining corpora, this is a small dataset. Our focus in curation was to produce a corpus free from redundancies (which is not often a focus²), and to source text primarily from reputable, authoritative sources. An ultimate aim beyond the scope of this work is to produce a minimally sufficient, authoritative corpus for related adjudication tasks. The dataset is intended to support both pretraining and retrieval.

3.1 Composition

We describe the composition of our dataset via a partition into six categories. Each document in our corpus belongs to exactly one of the following categories:

- (1) Legal. Current or former U.S. law.
- (2) Regulatory Guidance.Guidance on U.S. law, released by agencies.
- (3) Coverage Rules, Contracts, and Medical Policies. Text outside formal law that describes binding coverage rules³. This includes text from contracts, and contract-referenced or proprietary medical policies.
- (4) Opinion, Policy, and Summary. Opinions, policy perspectives, or summaries of law, proposed law, executive actions, or compliance.
- (5) **Case Descriptions**. Reviews of individual health insurance coverage decisions.
- (6) Medical Guidelines and Literature. Clinical guidelines and medical literature, excluding contract-specific medical policy falling into category 3.

3.2 Document-Level Tags

We equip each document in our corpus with a set of plain text tags. For example, documents comprised of text from U.S. law carry a "legal" tag. These tags support use as a knowledge base. For example the tags generally support constrained retrieval. This applies in particular to Retrieval Augmented Generation (RAG) [26], where one can use the tags to limit the collection of document chunks considered for injection in a prompt. Exact tag-based filtering supports pipelines capable of providing certain types of performance guarantee.

3.2.1 Partitioning Tags. We formulated a privileged set of *partitioning tags*, i.e. tags with the property that each document in the dataset belongs to exactly one tag. The privileged set corresponds to the breakdown described in Section 3.1. The associated plain text tags are: "legal", "regulatory-guidance", "contract-coverage-rule-medical-policy", "opinion-policy-summary", "case-description", and "clinical-guidelines".

3.2.2 Knowledge Base Tag. We also make use of a knowledge base tag ("kb") to indicate that a document is authoritative. This is of course an ambiguous determination. Nonetheless, our goal was to formalize this association as a step toward supporting applications requiring authoritative retrieval.

Knowledge base tags were auto generated during dataset curation with a human in the loop. We tried to assign the tag only to documents which are binding in law, or which constitute legal guidance from government agencies. However, even this characterization is ambiguous. For example, a contract is binding in law for those who are parties to the contract, but not for others. Documents which are binding in law in some but not all contexts, such as contracts, were decorated with the 'kb' tag.

For example, we label all Medicare Coverage Determination documents with the knowledge base tag. This is because law

²For example, Pile of Law [19] contains U.S. state and territory code snapshots, for each year for which the data exists on justia.com. This data comprises a 6.7 GB subset of the 256 GB dataset. Removing verbatim redundancies from this subset, resulting from code that remains unchanged year to year, reduces its size by an order of magnitude.

³In some cases law references but does not contain such rules (e.g. state Medicaid handbooks).

stipulates the authority of these determinations. As a result they hold established and verifiable weight in individual case adjudications.

As a non-example, we omit the knowledge base tag for congressional testimony documents. While such testimony is often reputable, and accurate, it is neither binding nor is it formal administrative guidance.

3.2.3 Example Usage. Consider a caseworker seeking information on coverage rules for a Medicaid beneficiary in New York state. Here a caseworker could mean a pharmacist, nurse, physician, attorney, patient advocate, or other individual who is supporting the beneficiary in understanding and accessing their coverage. A pipeline can employ filtering to restrict attention to documents with "legal", "kb", "new-york" and "medicaid" tags. This provides a simple and efficient mechanism to deterministically avoid retrieval of irrelevant documents. For example, those corresponding to Ohio state law.

This capability is important for at least two reasons. First, it removes the necessity to filter documents using only their semantic content. This allows for more efficient retrieval, which makes deployment more accessible for low-resource efforts, and improves latency. Second, tag based filtering supports provable guarantees about the relevance of retrieved documents. While the end to end utility of such guarantees in typical *generative* pipelines is tenuous, the utility in extractive or otherwise guardrailed pipelines need not be.

3.2.4 Tag Distribution. Table 1 shows high level statistics about the tag distribution. In total there are 616 unique tags used across the 8,311 documents. Figure 1 shows the distribution of tags by associated character count.



Figure 1: Character counts for the twenty most common tags in our dataset. There is a long tail of less common tags.

4 AN APPEAL ADJUDICATION TASK

We now introduce an appeal adjudication task, and an associated dataset.

4.1 Task Definition

The task is to predict whether an external appeal of a health insurance coverage denial will result in a full or partial overturn, given a description of the denial context and two optional pieces of metadata. We formulate this as a three class classification problem, with classes corresponding to: full or partial overturns, upheld decisions, and insufficient descriptions. We do not require that the description be complete, consist of clinical notes, or be written by experts. The optional metadata specify insurance type and regulatory jurisdiction classifications⁴ - two factors that influence coverage rules and appeal outcomes.

We define an external appeal to be one submitted to an independent third party. This is in contrast to an internal appeal, which is one submitted to the insurance administrator that issues a denial. External appeal processes exist for most types of health insurance, but vary in nature across them.

4.2 Motivation

This task is important for patients and caseworkers engaged in appeals processes. The likelihood of appeal success informs strategy, expectations, and cost-benefit analyses for appellants. It can also help streamline costly administrative work, and inform oversight.

It is important for regulators, who typically want to ensure that external appeal overturn rates are low⁵. Low overturn rates reflect consistency between insurer and third party adjudication, which is some form of fairness.

We compiled a labeled dataset for this task by exploiting historical case outcomes. Those outcomes are present in the "case-description" subset of our corpus.

4.3 Prediction

Although our labeled benchmark dataset is generated from historical case outcomes, we note that we are singularly focused on supporting actual forecasting. That is, we aim to support predicting appeal adjudication outcomes in real world situations in which the outcomes have yet to occur. This means our task must be addressed using only information which would be known by appellants or their advocates at that time. While this goal may sound obvious or implied from our setup, there has recently been critical and justified examination of misuse of the term 'prediction' in legal language tasks that perform classification using information that is not realistically available in the real world settings they purportedly support [29, 30].

In the following sections we introduce a process for learning to extract subsets of case summaries to obtain context which is reflective of that known by patients prior to their appeals.

⁴In a given implementation, insurance types could take the form of high-level categorizations (e.g. "Medicaid"), or be more fine grained (e.g. "New York Medicaid Advantage"). Similarly one could use a high-level categorization of the primary regulatory jurisdiction, such as "New York", or "Employee Benefits Security Administration", or a more nuanced and accurate multi-label scheme.

 $^{^5\}mathrm{A}$ Medicare star rating measure reflects this goal. See e.g. measure C29 of the 2024 technical notes.

Category	Documents	Words (M)	Chars (M)	Size (GB)
legal	330	86	545	0.57
regulatory-guidance	1,110	5	38	0.04
contract-coverage-rule-medical-policy	7	196	1,228	1.31
opinion-policy-summary	2,093	19	131	0.14
case-description	2,630	214	1,351	1.44
clinical-guidelines	2,150	82	554	0.57
[kb]	[1,434]	[171]	[1,123]	[1.16]
Total	8.311	419	2,703	2.83

Table 1: Dataset statistics broken down by our privileged partitioning set. We also show statistics for the subset of documents tagged with the knowledge base tag ("kb"). The Total row indicates the sum of the rows excluding the [kb] row.

In this way, our approach to the task is truly one of of *outcome forecasting*, rather than post hoc *outcome-based judgement identification*, to use the language of [30]. This is because our process is designed to extract summaries that only contain statements which could be known before an appeal is submitted.

Our focus on this quality of our inputs is central to our formulation, as we are interested in real world use of our models to support patients in accessing coverage to which they are entitled, rather than retrospective classification models tailored to the qualities of review summaries.

4.4 Challenges

Our task is difficult for many reasons.

Most notable is that outcomes may depend on facts not present in our model inputs, such as jurisdiction, governing law, contracts, medical history, policy, and literature. These vary by state, insured, insurer, and insurance type, and change over time.

This renders our modeling task under specified in many contexts, with no correct answer for many inputs. For example, what should a model predict when an appeal would be overturned under Medicare rules but upheld under New York Medicaid rules, if insurance type isn't specified?

We introduce our task despite this difficulty for three reasons.

First, responsible deployment of a well-performing model would be valuable and pose few risks. Although such a model cannot be perfect, it can help improve access to justice. For some contexts, denials are overwhelmingly inappropriate, across jurisdictions, insurance types, and contracts. It is possible to identify such context as *likely* to be problematic. This is useful in the same way a prediction for 90% chance of rain (another theoretically under specified forecasting task) is useful. It would be irresponsible to bet one's wellbeing on the prediction, but carrying an umbrella is a relatively harmless precaution. Models that excel at our task can help promote access to justice, without posing risks, if deployed with caution and with limited scope.

Second, we hope this work promotes further research on related tasks. For example, one could extend our benchmark to additionally include optional input metadata encoding medical billing codes.

Finally, our task circumvents a common difficulty for patients: accessing important details. Patients face significant barriers in accessing important details. They may not have full access to contracts or medical records, and may not know who funds their plan. These details are critically relevant in helping patients navigate recourse available to them. Nonetheless, they are often inaccessible, even when law ostensibly protects access. Accurate prediction tools with lower barriers to entry can improve access to justice for those who might otherwise forgo recourse.

It is important to note that our task differs from that performed by adjudicators. Reviewers rely on charts, notes, diagnoses, and rationales, and can request more information. Our benchmark descriptions provide only brief explanations which are always insufficient for adjudication, and sometimes also insufficient for high quality outcome forecasting.

Some explanations in our benchmark are too general to hold any predictive power. For example, a summary might say "My hospital stay was denied". There are many reasons for hospital admission, and hospital admission is routinely rightfully *and* wrongfully denied.

At the other extreme, some summaries can be used to predict appeal adjudication outcomes effectively. This can be because they contain specific detail, or describe situations that typically yield consistent outcomes regardless of details. For example, among state regulated commercial plans in California, more than 80% of appealed denials for Harvoni for Hepatitis C treatment are overturned [13].

This variability is both challenging and reflects a real problem encountered in manual casework. Patients request help with varying specificity. Modeling whether a request is sufficient to predict an outcome is important for applications. We want to avoid making predictions when there is insufficient information, which is why the problem is cast as one of three class classification.

4.5 Raw Source Data

To construct annotated data, we used three sources whose unlabeled text are also present in our corpus. These sources are:

- (1) A New York DFS External Appeal Database.
- (2) A California CDI External Appeal Database.
- (3) A California DMHC External Appeal Database.

The sources contain plain text descriptions of external appeal cases, often including background, summaries of the outcomes, and rationales for decisions. Separately, structured metadata is reported. We extracted from the structured metadata a rudimentary insurance type and jurisdiction classification scheme: each case is labeled with an insurance type among {Commercial, Medicaid, Unspecified}, and a jurisdiction among {NY, CA, Unspecified}⁶.

The standards used to structure the plain text descriptions are inconsistent both between and within the sources. Typically, the adjudication rationale leaks information about the outcome, whereas the background does not.

Together, all these features make it nontrivial to extract useful (description, outcome) pairs, as meaningful training and evaluation of forecasting models requires descriptions which do not leak the outcome.

Figure 2 illustrates the issues. We reiterate that across our data sources, case descriptions leak case outcomes in varied ways. This example illustrates one of many case description formats and mechanisms of leaking. The high degree of variability makes immediate solutions like splitting on keywords ineffective.

4.6 Construction of Training Data

To construct a suitable training dataset, we first produced manual span annotations for a subset of case descriptions. We labeled subspans encoding background context one would typically know before an appeal submission, and scored the spans on a scale from 1 to 4, based on the degree to which they would be sufficient to make an informed prediction about the adjudication outcome in isolation.

We then used this annotated subset to train two 'bootstrapping' models. One is a span selection model responsible for extracting background from potentially leaking case descriptions. The other is a binary classifier for such background text, indicating whether the text is sufficient in isolation to predict a case outcome. We used the models together to perform extraction of non-leaking background from all 73,987 of our case descriptions, and to endow them with three-class pseudolabels. Finally, we used this larger collection of pseudo-annotated case summaries to train appeal outcome classifiers.

4.6.1 Manual Annotation. We manually annotated spans from 1,000 case descriptions, with one third coming from each of the three data sources. The descriptions from each data source were sampled randomly. We used argilla [10] to annotate non-leaking background spans and label those spans with sufficiency scores. The work described in this paper is ongoing, and we are currently working to streamline and iteratively improve our annotation guidelines, and to generate more manual annotations. Encouraged by our initial modeling results, described in more detail in what follows, we plan to enlist 5 to 10 professionals experienced in casework to each manually annotate the same set of 3,000 case descriptions. We will then measure inter-annotator agreement for spans and sufficiency scores via numerous measures, to evaluate the extent to which our manual annotation process yields reliable, and consistent results. At present, we have no such consistency results to share.

4.6.2 Span Selector Training. We trained a span selection model to extract non-leaking background context from our case descriptions. We finetuned Distilbert [43] for this task, and found subjectively⁷ that it gave results of sufficient quality. We did not experiment with alternative approaches to span selection once we successfully trained the Distilbert variant⁸.

4.6.3 Sufficiency Model Training. We trained a model to classify background as sufficient or insufficient to predict a case outcome. We also finetuned Distilbert for this task. We produced binary labels from our ranked sufficiency score annotations by labeling all scores less than 3 as "Insufficient", and all scores greater than or equal to 3 as "Sufficient". This choice was informed by our current annotation guidelines, which are documented in a previously released variant of this work. In short, scores less than 3 reflect descriptions which do not specify either the service, or the diagnosis or medical issue for which the service is being provided. Without both of these pieces of information, it is impossible to assess the likelihood of appeal overturn with any fidelity.

4.6.4 Bootstrap Model Application. To complete construction of the dataset, we applied our trained span selection model to 73,987 case descriptions, and then applied our sufficiency classifier. We partitioned the results into training and test splits; splits were stratified by source and appeal outcome and otherwise determined via uniform sampling.

An example of a non-leaking case summary extracted by our trained model artifact is displayed in Figure 3.

4.7 Baselines

We trained end to end outcome prediction models on our benchmark to serve as baselines for further modeling efforts. The choice of models we evaluated was, as is often the case, essentially arbitrary.

⁶Because the insurance types and jurisdictions do not vary much across our data, we do not expect them to be impactful for modeling *in our benchmark*. We do expect them to be useful for the general task.

 $^{^7\}mathrm{We}$ evaluated our span selector quantitatively, but our determination that the model was adequate was a subjective one based on qualitative human evaluation of many samples. The metrics for token-level classification on a held out test set were 89.8% recall, 92.2% precision, 91.0% F1, and 93.2% accuracy.

⁸Our first approach was to use a language model to perform the requisite span selection. We found initial results with gpt-40 and Claude 3 variants unsatisfactory. Given the success of this cheap approach which also affords perpetual, negligible-cost access to the model artifacts, we did not pursue that path further.

Source Text

The parents of a seventeen month-old male with respiratory distress requiring intubation, surfactant, and episodes of supraventricular tachycardia (SVT) are requesting reimbursement and prospective authorization for Synagis (pallivizumab). Reviewers Findings: The reviewing physician found that RSV infection has been associated with tachyarrhythmias including SVT that may be hemodynamically significant in infants such as this one. The reviewing physician found that the Health Plan denial should be overturned.

Figure 2: An example of a raw case description from our unlabeled corpus. Background context that summarizes the facts of the case knowable prior to its review by an external reviewer is indicated in red. The subsequent non-background text leaks the case outcome determination. We used argilla to manually annotate background context.

Extracted Case Summary

A 52-year-old female enrollee has requested reimbursement for the Decision Dx melanoma assay provided on 11/18/15. The Health Insurer has denied this request indicating that the testing at issue was considered investigational for evaluation of the enrollee's malignant melanoma.

Figure 3: An example of non-leaking background context extracted from a full, leaking case description, by our trained Distilbert-based span selector.

Table 2 shows the results of evaluating a set of finetuned models on the test set of our benchmark.

The models tested include pretrained bidirectional transformers (BERT and DistilBERT variants). One model, ClinicalBERT [22], was already pretrained on medical corpora. Another, LegalBERT [5], was already pretrained on legal corpora.

We also compared performance to a popular closed source language model, GPT-4o-mini. GPT-4o-mini was provided a system prompt explaining the task, and a required output format. It was then evaluated in a two shot context. The best performing model among those we tested was the Distilbert variant. Table 2 records the results on a per-metric basis.

Our training and evaluation methodologies are detailed further in our released code, and in previously published variants of this work.

5 DISCUSSION

We now turn to discussing potential applications of this work.

5.1 Intended Use

Our corpus, task, and benchmarks were all constructed with downstream applications in mind. We list a few.

Corpus as a Partial Knowledge Base Health insurance coverage rules, contracts, and medical policies are often complex. They are difficult to efficiently and accurately apply to the facts of particular situations, even for experts. Existing generative tools are capable of summarizing the complexities of individual provisions, but to do so they need access to authoritative ground truth. Many deployed solutions suffer from both insufficient access to such ground truth, and an abundance of access to non-authoritative or inaccurate text.

Our introduction of this corpus is a step towards a minimally sufficient knowledge base. It can support generative and extractive tasks related to understanding of U.S. health insurance. It is particularly well suited to legal question answering for health insurance related questions.

Models as Oversight Tools Regulators could use high quality appeal outcome predictors to promote improved outcomes and adjudication. For example, they could require that cases with high likelihood of overturn get reviewed in shortened time frames. This could improve mean time to access-to-care for overturned pre-service denials, which has a major impact on health outcomes [16, 39].

Models as Patient Self Help Tools Appeal outcome predictors can support patients who do not have access to expert support.

Blind reliance and automated use carries risk, and we do not recommend such use. Nonetheless, these models can play effective and safe supporting roles.

Blind reliance has problematic implications. Models could incorrectly convince a patient they can't win an appeal. If this led to a patient unnecessarily forgoing care, the consequences could be dire. This would be especially problematic when predictions differ from those that would be provided by human experts.

Risk mitigation is possible through worst case scenario assessment. By designing self-help applications for responsible use, acceptable risk levels are achievable. One mechanism is designing to promote 'qualified trust' - that is, granting trust qualified by explicit risk assessment.

For example, applications could instruct those desperately needing coverage to ignore predictions and seek the support of human experts. Many free and low-cost support systems exist across the U.S. In such cases, prediction inaccuracy risks outweigh benefits.

Model	Accuracy	Precision	Recall	F1	ROC AUC	Params (M)
gpt-40-mini-2024-07-18	.532	.499	.565	.397	N/A	unknown
(2-shot)						
legalbert-small-uncased	.707	.753	.669	.704	.857	35
distilbert-base-uncased	.733	.772	.695	.727	.878	67
clinical-bert	.740	.758	.688	.717	.889	135

Table 2: Classification metrics on the held out test set of our case adjudication benchmark. Each metric is represented by a macro average across the three classes, and the best performing model for each macro metric is in bold. Note that the macro F1 is not the harmonic mean of the macro recall and macro precision. In the same vein, it is not true that the value in the F1 column must fall between the values presented in the Precision and Recall columns, as seen in the gpt metrics. The test set is comprised of 9,745 records, of which 175 are labeled "Insufficient", 5,177 are labeled "Upheld", and 4,393 are labeled "Overturned".

When lack of coverage has important, but not dire, implications, applications could instruct patients to seek expert advice. For those without access to such help, the models can serve as fallbacks. The risk of prediction inaccuracy might or might not outweigh the risks of having no support at all. For example, pursuing an appeal that is ultimately unsuccessful costs time. Forgoing a post-service appeal and paying for care in a case that would have resulted in overturn costs money. Forgoing a pre-service appeal and forgoing care as a result in a case that would have resulted in overturn poses risks to one's health. The risk tradeoff is subjective, and case dependent.

When lack of coverage has minimal implications, our models can safely support patients most easily. Patients in such scenarios often forgo post-service appeals due to the mistaken belief that they are unlikely to win. The models provide a quick way for patients to update their belief in the likelihood of appeal success.

Self-help tools designed through a lens of qualified trust can improve patient outcomes.

Task as an Application Entry Point Our task has notable shortcomings and limitations. However, the more general conceptual context merits further exploration.

Our task illuminates the need for progress on a suite of related tasks. We hope that our task serves as an entry point for other researchers to contribute to this domain. Our adjudication data and annotation workflow provide a straightforward path to such contributions.

We list some alternate and complementary tasks of interest.

- It would be useful to predict whether an appeal merits overturn, according to law or policy. Such a model could help regulators and administrators triage appeal case loads. Effective triage could lower mean-time-to-case-resolution for *inappropriate* denials. A model of this type could also help regulators audit review outcomes. One could approach this task with our methods and alternate annotation guidelines.
- (2) It would be useful to cite evidence supporting predictions for the task in item 1. A simple formulation could use a language model and RAG pipeline to achieve this. It is straightforward to use our corpus as a partial knowledge base towards this end.

- (3) It would be useful to construct complete, coherent arguments supporting predictions for the task in item 1. This could take the form of generating 'proof' for an adjudication determination. Such proof could be primarily extractive in nature, relying on authoritative references. It would require a dynamic determination of a minimally sufficient set of source chunks, unlike a typical RAG pipeline, and coherent assimilation of those sources. The exact approach could take many forms (see e.g. [45] for one motivating example).
- (4) It would be useful to draft high quality appeal letters for patients and providers. Appeal processes are laborious, and present many barriers. One intensive step is writing a letter arguing for overturn. Some reviewers consider letters minimally, rendering generic or inaccurate letters harmless. However, letter quality and accuracy become critical in dire cases or when reviewers might miss important context. Deploying low quality letter generators without precautions harms the most vulnerable. Nonetheless, it is trivial to take this approach, and many companies are doing so with varied (monetary) success.

Responsibly deploying letter generators that present accurate and verifiable evidence is more difficult. Models with this capability offer many benefits. This task is of course deeply related to the previous task.

Benchmark as a Minimal Prototype Our appeal adjudication task formulation reflects a realistic application with broad potential impact. On the other hand, our preliminary *benchmark dataset* is a rudimentary formulation of the task. The corpus and task definition, rather than the task benchmark and model implementations, are the more central concerns of this work. The task data and models serve as minimal, first iteration prototypes for a more general research need. We hope the benchmarks will inspire further research and more refined formulations of benchmark data of the same flavor in this space.

5.2 Out of Scope Use

There are many ways to use this work that are not intended or advised. These include:

- Treating the unlabeled corpus as a *complete* knowledge base.
- Applying our benchmark models to alternate tasks (e.g. interpreting the outputs as appropriate, rather than expected, outcomes).
- Using our data to fully automate claims denial review and recourse processes.

Some out of scope uses of our work arise from inherent limitations, which we discuss in more detail in Section 7. Others arise from associated risks. For example, it is possible to use this work to deploy applications that would negatively affect patient outcomes, despite our intent. We discuss the risk landscape in Section 8.

6 CONCLUSION

In this work we introduced a corpus of reputable text pertaining to health insurance coverage rules, and discussed potential applications. We also introduced a case outcome prediction task, and a benchmark dataset for that task. We trained and evaluated model baselines on our introduced benchmark. The data and task show promise for improving efficient understanding of health insurance coverage rules.

7 LIMITATIONS

7.1 Task Shortcomings

Our appeal adjudication task has room for improvement; for example, it could support additional optional metadata. Appeal outcomes often depend on medical records, billing codes, and other missing factors.

7.2 Simplicity of the Benchmark

As detailed in Section 4.6, we bootstrapped our benchmark from 1,000 manual annotations. Furthermore, we extracted background descriptions from a set of relatively short case descriptions. It would be beneficial to draw descriptions from a more diverse distribution, and to perform more manual annotation.

7.3 Corpus Deficiencies

Our corpus lacks information necessary to make certain types of coverage determinations. This includes commercial insurance contracts, medical policies, Medicaid and CHIP state plans, Medicare managed care contracts, and the most recent state and U.S. code.

8 ETHICAL CONSIDERATIONS

8.1 Potential for Propagation of Bias

The risk of propagating harmful bias is central in machine learning. One hopes to learn about a target problem from training data, but risks instead learning about bias in that data. This is particularly problematic when tools are widely used before they are evaluated for such bias.

The proliferation of language models has worsened these problems [11]. It is easier than ever to develop, deploy, and gain adoption for tools with minimal effort. At the same time, language models have less built-in guarantees than many other ML models. Their evaluation also requires more subtle, subjective choices.

Healthcare applications are particularly fraught. There are already stark disparities and vast collections of data reflecting those disparities. Automated tools have consistently led to bias propagation [36, 37].

There are inequities in U.S. healthcare associated with race, sex, gender identity, age, disease states, disability status, and income, among many other things. Our work could propagate such biases through unintended use.

Particular concerns include:

- (1) Conflating historical trends with just or appropriate trends. This concern informed our task formulation. Modeling whether a case *merits* overturn with our data would be misguided because historical outcome data encodes biases we cannot sufficiently account for.
- (2) Demographic selection bias and inequitable benefit. Our case data has minimal demographic information, and mostly corresponds to individuals over the age of 50 with commercial plans in New York and California. It also reflects only externally appealed denials, rather than the complete distribution. Information about racial, ethnic, and income distributions is limited, and more robust evaluations are necessary. Models trained on our benchmark will perform best in contexts common in the data, potentially exacerbating disparities.

8.2 Potential for Misuse

There is also potential for our data or models to worsen patient outcomes in ways unrelated to bias. This could occur through reckless use, or intentionally (e.g. for financial gain).

Organizations quickly deploy prototypes to raise funding and acquire customers. A reckless entrepreneur could use our work to deploy immature patient support tools. Without guardrails or clear disclaimers about risks, such use could have disastrous consequences. Many people who need support are fighting for coverage of care critical to their well-being.

More concerning is the possibility of nefarious misuse. Corporations with extensive resources are deploying ML tools to lower their costs. When done legally, and responsibly, there are many benefits to society. But healthcare organizations have used ML based cost-cutting measures irresponsibly [42]. Such irresponsible misuse is a major risk. For example, one could repurpose our overturn prediction models as semiautomated appeal reviewers.

REFERENCES

- Emily Alsentzer, John R. Murphy, Willie Boag, Wei-Hung Weng, Di Jin, Tristan Naumann, and Matthew B. A. Mc-Dermott. 2019. Publicly Available Clinical BERT Embeddings. arXiv e-prints, Article arXiv:1904.03323 (April 2019), arXiv:1904.03323 pages. https://doi.org/10.48550/arXiv.1904. 03323 arXiv:cs.CL/1904.03323
- [2] Claudio Aracena, Nicolás Rodríguez, Victor Rocco, and Jocelyn Dunstan. 2023. Pre-trained language models in Spanish for health insurance coverage. In *Proceedings of the 5th Clinical Natural Language Processing Workshop*, Tristan Naumann, Asma Ben Abacha, Steven Bethard, Kirk Roberts, and Anna Rumshisky (Eds.). Association for Computational Linguistics, Toronto, Canada, 433–438. https://doi.org/10.18653/v1/2023. clinicalnlp-1.46
- [3] Emmanuel Bauer, Dominik Stammbach, Nianlong Gu, and Elliott Ash. 2023. Legal Extractive Summarization of U.S. Court Opinions. arXiv e-prints, Article arXiv:2305.08428 (May 2023), arXiv:2305.08428 pages. https://doi.org/10.48550/arXiv.2305. 08428 arXiv:cs.CL/2305.08428
- [4] Ilias Chalkidis, Ion Androutsopoulos, and Nikolaos Aletras. 2019. Neural Legal Judgment Prediction in English. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Anna Korhonen, David Traum, and Lluís Màrquez (Eds.). Association for Computational Linguistics, Florence, Italy, 4317–4323. https://doi.org/10.18653/v1/P19-1424
- [5] Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, Nikolaos Aletras, and Ion Androutsopoulos. 2020. LEGAL-BERT: The Muppets straight out of Law School. arXiv e-prints, Article arXiv:2010.02559 (Oct. 2020), arXiv:2010.02559 pages. https: //doi.org/10.48550/arXiv.2010.02559 arXiv:cs.CL/2010.02559
- [6] Ilias Chalkidis, Nicolas Garneau, Catalina Goanta, Daniel Martin Katz, and Anders Søgaard. 2023. LeXFiles and LegalLAMA: Facilitating English Multinational Legal Language Model Development. arXiv e-prints, Article arXiv:2305.07507 (May 2023), arXiv:2305.07507 pages. https://doi.org/10.48550/arXiv.2305. 07507 arXiv:cs.CL/2305.07507
- [7] Ilias Chalkidis, Abhik Jana, Dirk Hartung, Michael Bommarito, Ion Androutsopoulos, Daniel Martin Katz, and Nikolaos Aletras. 2021. LexGLUE: A Benchmark Dataset for Legal Language Understanding in English. arXiv e-prints, Article arXiv:2110.00976 (Oct. 2021), arXiv:2110.00976 pages. https: //doi.org/10.48550/arXiv.2110.00976 arXiv:cs.CL/2110.00976
- [8] Zeming Chen, Alejandro Hernández Cano, Angelika Romanou, Antoine Bonnet, Kyle Matoba, Francesco Salvi, Matteo Pagliardini, Simin Fan, Andreas Köpf, Amirkeivan Mohtashami, Alexandre Sallinen, Alireza Sakhaeirad, Vinitra Swamy, Igor Krawczuk, Deniz Bayazit, Axel Marmet, Syrielle Montariol, Mary-Anne Hartley, Martin Jaggi, and Antoine Bosselut. 2023. MEDITRON-70B: Scaling Medical Pretraining for Large Language Models. arXiv e-prints, Article arXiv:2311.16079 (Nov. 2023), arXiv:2311.16079 pages. https://doi.org/10.48550/arXiv.2311. 16079 arXiv:cs.CL/2311.16079
- [9] Sarah Collins, Shreya Roy, and Relebohile Masitha. 2023. Paying for It: How Health Care Costs and Medical Debt Are Making Americans Sicker and Poorer. https://www.commonwealthfund. org/publications/surveys/2023/oct/paying-for-it-costs-debtamericans-sicker-poorer-2023-affordability-survey Accessed on January 3, 2024.
- [10] Vila-Suero Daniel and Aranda Francisco. 2023. Argilla Opensource framework for data-centric NLP. https://github.com/ argilla-io/argilla
- [11] Jwala Dhamala, Tony Sun, Varun Kumar, Satyapriya Krishna, Yada Pruksachatkun, Kai-Wei Chang, and Rahul Gupta. 2021. Bold: Dataset and metrics for measuring biases in open-ended language generation. In Proceedings of the 2021 ACM conference on fairness, accountability, and transparency. 862–872.
- [12] Emad Elwany, Dave Moore, and Gaurav Oberoi. 2019. BERT Goes to Law School: Quantifying the Competitive Advantage of Access to Large Legal Corpora in Contract Understanding. arXiv e-prints, Article arXiv:1911.00473 (Nov. 2019), arXiv:1911.00473 pages. https://doi.org/10.48550/arXiv.1911. 00473 arXiv:cs.CL/1911.00473
- [13] Michael Gartner. 2023. California External Appeal Outcome Demographics. https://blog.persius.org/blog/ca-external-appealsdemographics Accessed on November 15, 2024.

- [14] Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2020. Domain-Specific Language Model Pretraining for Biomedical Natural Language Processing. arXiv e-prints, Article arXiv:2007.15779 (July 2020), arXiv:2007.15779 pages. https: //doi.org/10.48550/arXiv.2007.15779 arXiv:cs.CL/2007.15779
- [15] Neel Guha, Julian Nyarko, Daniel E. Ho, Christopher Ré, Adam Chilton, Aditya Narayana, Alex Chohlas-Wood, Austin Peters, Brandon Waldon, Daniel N. Rockmore, Diego Zambrano, Dmitry Talisman, Enam Hoque, Faiz Surani, Frank Fagan, Galit Sarfaty, Gregory M. Dickinson, Haggai Porat, Jason Hegland, Jessica Wu, Joe Nudell, Joel Niklaus, John Nay, Jonathan H. Choi, Kevin Tobia, Margaret Hagan, Megan Ma, Michael Livermore, Nikon Rasumov-Rahe, Nils Holzenberger, Noam Kolt, Peter Henderson, Sean Rehaag, Sharad Goel, Shang Gao, Spencer Williams, Sunny Gandhi, Tom Zur, Varun Iyer, and Zehua Li. 2023. Legal-Bench: A Collaboratively Built Benchmark for Measuring Legal Reasoning in Large Language Models. arXiv e-prints, Article arXiv:2308.11462 (Aug. 2023), arXiv:2308.11462 pages. https: //doi.org/10.48550/arXiv.2308.11462 arXiv:cSCL/2308.11462
- [16] Avni Gupta, Sara R. Collins, Shreya Roy, and Relebohile Masitha. 2024. Unforeseen Health Care Bills and Coverage Denials by Health Insurers in the U.S. https://www.commonwealthfund.org/publications/issuebriefs/2024/aug/unforeseen-health-care-bills-coverage-denialsby-insurers Accessed on April 27, 2025.
- [17] Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. Don't Stop Pretraining: Adapt Language Models to Domains and Tasks. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (Eds.). Association for Computational Linguistics, Online, 8342-8360. https://doi.org/10.18653/v1/ 2020.acl-main.740
- [18] Ivan Habernal, Daniel Faber, Nicola Recchia, Sebastian Bretthauer, Iryna Gurevych, Indra Spiecker genannt Döhmann, and Christoph Burchard. 2022. Mining Legal Arguments in Court Decisions. arXiv e-prints, Article arXiv:2208.06178 (Aug. 2022), arXiv:2208.06178 pages. https://doi.org/10.48550/arXiv.2208. 06178 arXiv:cs.CL/2208.06178
- [19] Peter Henderson, Mark S. Krass, Lucia Zheng, Neel Guha, Christopher D. Manning, Dan Jurafsky, and Daniel E. Ho. 2022. Pile of Law: Learning Responsible Data Filtering from the Law and a 256GB Open-Source Legal Dataset. arXiv e-prints, Article arXiv:2207.00220 (July 2022), arXiv:2207.00220 pages. https: //doi.org/10.48550/arXiv.2207.00220 arXiv:cs.CL/2207.00220
- [20] Dan Hendrycks, Collin Burns, Anya Chen, and Spencer Ball. 2021. CUAD: An Expert-Annotated NLP Dataset for Legal Contract Review. arXiv e-prints, Article arXiv:2103.06268 (March 2021), arXiv:2103.06268 pages. https://doi.org/10.48550/arXiv.2103. 06268 arXiv:cs.CL/2103.06268
- [21] Wenyue Hua, Yuchen Zhang, Zhe Chen, Josie Li, and Melanie Weber. 2022. LegalRelectra: Mixed-domain Language Modeling for Long-range Legal Text Comprehension. arXiv e-prints, Article arXiv:2212.08204 (Dec. 2022), arXiv:2212.08204 pages. https: //doi.org/10.48550/arXiv.2212.08204 arXiv:cs.CL/2212.08204
- [22] Kexin Huang, Jaan Altosaar, and Rajesh Ranganath. 2019. ClinicalBERT: Modeling Clinical Notes and Predicting Hospital Readmission. arXiv e-prints, Article arXiv:1904.05342 (April 2019), arXiv:1904.05342 pages. https://doi.org/10.48550/arXiv.1904. 05342 arXiv:cs.CL/1904.05342
- [23] Lavender Yao Jiang, Xujin Chris Liu, Nima Pour Nejatian, Mustafa Nasir-Moin, Duo Wang, Anas Abidin, Kevin Eaton, Howard Antony Riina, Ilya Laufer, Paawan Punjabi, et al. 2023. Health system-scale language models are all-purpose prediction engines. *Nature* 619, 7969 (2023), 357–362.
- [24] Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. 2020. What Disease does this Patient Have? A Large-scale Open Domain Question Answering Dataset from Medical Exams. arXiv e-prints, Article arXiv:2009.13081 (Sept. 2020), arXiv:2009.13081 pages. https://doi.org/10.48550/ arXiv.2009.13081 arXiv:cs.CL/2009.13081
- [25] Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William W. Cohen, and Xinghua Lu. 2019. PubMedQA: A Dataset for Biomedical Research Question Answering. arXiv e-prints, Article arXiv:1909.06146 (Sept. 2019), arXiv:1909.06146 pages. https: //doi.org/10.48550/arXiv.1909.06146 arXiv:cs.CL/1909.06146

- [26] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. arXiv e-prints, Article arXiv:2005.11401 (May 2020), arXiv:2005.11401 pages. https://doi.org/10.48550/ arXiv.2005.11401 arXiv:cs.CL/2005.11401
- [27] Lunna Lopes, Audrey Kearney, Alex Montero, Liz Hamel, and Mollyann Brodie. 2022. Health Care Debt In The U.S.: The Broad Consequences Of Medical And Dental Bills. https://www.kff.org/ report-section/kff-health-care-debt-survey-main-findings/ Accessed on January 3, 2024.
- [28] Vijit Malik, Rishabh Sanjay, Shubham Kumar Nigam, Kripabandhu Ghosh, Shouvik Kumar Guha, Arnab Bhattacharya, and Ashutosh Modi. 2021. ILDC for CJPE: Indian Legal Documents Corpus for Court Judgment Prediction and Explanation. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (Eds.). Association for Computational Linguistics, Online, 4046-4062. https://doi.org/10.18653/v1/2021.acl-long.313
- [29] Masha Medvedeva and Pauline Mcbride. 2023. Legal Judgment Prediction: If You Are Going to Do It, Do It Right. In Proceedings of the Natural Legal Language Processing Workshop 2023, Daniel Preotiuc-Pietro, Catalina Goanta, Ilias Chalkidis, Leslie Barrett, Gerasimos Spanakis, and Nikolaos Aletras (Eds.). Association for Computational Linguistics, Singapore, 73–84. https://doi.org/ 10.18653/v1/2023.nllp-1.9
- [30] Masha Medvedeva, Martijn Wieling, and Michel Vols. 2023. Rethinking the field of automatic prediction of court decisions. *Artificial Intelligence and Law* 31, 1 (01 Mar 2023), 195–212. https://doi.org/10.1007/s10506-021-09306-3
- [31] Joel Niklaus, Ilias Chalkidis, and Matthias Stürmer. 2021. Swiss-Judgment-Prediction: A Multilingual Legal Judgment Prediction Benchmark. In Proceedings of the Natural Legal Language Processing Workshop 2021, Nikolaos Aletras, Ion Androutsopoulos, Leslie Barrett, Catalina Goanta, and Daniel Preotiuc-Pietro (Eds.). Association for Computational Linguistics, Punta Cana, Dominican Republic, 19–35. https://doi.org/10.18653/v1/2021.nllp-1.3
 [32] Joel Niklaus and Daniele Giofré. 2022. BudgetLongformer:
- [32] Joel Niklaus and Daniele Giofré. 2022. BudgetLongformer: Can we Cheaply Pretrain a SotA Legal Language Model From Scratch? arXiv e-prints, Article arXiv:2211.17135 (Nov. 2022), arXiv:2211.17135 pages. https://doi.org/10.48550/arXiv.2211. 17135 arXiv:cs.CL/2211.17135
- [33] Joel Niklaus, Veton Matoshi, Pooja Rani, Andrea Galassi, Matthias Stürmer, and Ilias Chalkidis. 2023. LEXTREME: A Multi-Lingual and Multi-Task Benchmark for the Legal Domain. arXiv e-prints, Article arXiv:2301.13126 (Jan. 2023), arXiv:2301.13126 pages. https://doi.org/10.48550/arXiv.2301. 13126 arXiv:cs.CL/2301.13126
- [34] Joel Niklaus, Veton Matoshi, Matthias Stürmer, Ilias Chalkidis, and Daniel E. Ho. 2023. MultiLegalPile: A 689GB Multilingual Legal Corpus. arXiv e-prints, Article arXiv:2306.02069 (June 2023), arXiv:2306.02069 pages. https://doi.org/10.48550/arXiv. 2306.02069 arXiv:cs.CL/2306.02069
- [35] Joel Niklaus, Lucia Zheng, Arya D. McCarthy, Christopher Hahn, Brian M. Rosen, Peter Henderson, Daniel E. Ho, Garrett Honke, Percy Liang, and Christopher Manning. 2024. FLawN-T5: An Empirical Examination of Effective Instruction-Tuning Data Mixtures for Legal Reasoning. arXiv e-prints, Article arXiv:2404.02127 (April 2024), arXiv:2404.02127 pages. https: //doi.org/10.48550/arXiv.2404.02127 arXiv:cs.CL/2404.02127
- [36] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366, 6464 (2019), 447-453. https://doi.org/10.1126/science.aax2342
- [37] Jesutofunmi A Omiye, Jenna C Lester, Simon Spichak, Veronica Rotemberg, and Roxana Daneshjou. 2023. Large language models propagate race-based medicine. NPJ Digital Medicine 6, 1 (2023), 195.
- [38] Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. 2022. MedMCQA: A Large-scale Multi-Subject Multi-Choice Dataset for Medical domain Question Answering. In Proceedings of the Conference on Health, Inference, and Learning (Proceedings of Machine Learning Research), Gerardo Flores, George H Chen, Tom Pollard, Joyce C Ho, and Tristan Naumann

(Eds.), Vol. 174. PMLR, 248–260. https://proceedings.mlr.press/v174/pal22a.html

- [39] Karen Pollitz, Kaye Pestaina, Luna Lopes, Rayna Wallace, and Justin Lo. 2023. Consumer Survey Highlights Problems with Denied Health Insurance Claims. https://www.kff.org/healthreform/issue-brief/consumer-survey-highlights-problems-withdenied-health-insurance-claims/ Accessed on January 8, 2024.
- [40] Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know What You Don't Know: Unanswerable Questions for SQuAD. arXiv e-prints, Article arXiv:1806.03822 (June 2018), arXiv:1806.03822 pages. https://doi.org/10.48550/arXiv.1806. 03822 arXiv:cs.CL/1806.03822
- [41] Abhilasha Ravichander, Alan W Black, Shomir Wilson, Thomas Norton, and Norman Sadeh. 2019. Question Answering for Privacy Policies: Combining Computational and Legal Perspectives. arXiv e-prints, Article arXiv:1911.00841 (Nov. 2019), arXiv:1911.00841 pages. https://doi.org/10.48550/arXiv.1911. 00841 arXiv:cs.CL/1911.00841
- [42] Casey Ross and Bob Herman. 2023. Denied by AI: How Medicare Advantage plans use algorithms to cut off care for seniors in need. STAT News (2023). https://www.statnews.com/2023/03/13/ medicare-advantage-plans-denial-artificial-intelligence/
- [43] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. arXiv e-prints, Article arXiv:1910.01108 (Oct. 2019), arXiv:1910.01108 pages. https: //doi.org/10.48550/arXiv.1910.01108 arXiv:cs.CL/1910.01108
- [44] Karan Singhal, Shekoofeh Azizi, Tao Tu, S. Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, Perry Payne, Martin Seneviratne, Paul Gamble, Chris Kelly, Abubakr Babiker, Nathanael Schärli, Aakanksha Chowdhery, Philip Mansfield, Dina Demner-Fushman, Blaise Agüera y Arcas, Dale Webster, Greg S. Corrado, Yossi Matias, Katherine Chou, Juraj Gottweis, Nenad Tomasev, Yun Liu, Alvin Rajkomar, Joelle Barral, Christopher Semturs, Alan Karthikesalingam, and Vivek Natarajan. 2023. Publisher Correction: Large language models encode clinical knowledge. Nature 620, 7973 (Aug. 2023), E19–E19. https://doi.org/10.1038/s41586-023-06455-0 arXiv:cs.CL/2212.13138
- [45] Oyvind Tafjord, Bhavana Dalvi Mishra, and Peter Clark. 2020. ProofWriter: Generating Implications, Proofs, and Abductive Statements over Natural Language. arXiv e-prints, Article arXiv:2012.13048 (Dec. 2020), arXiv:2012.13048 pages. https: //doi.org/10.48550/arXiv.2012.13048 arXiv:cs.CL/2012.13048
- [46] Don Tuggener, Pius von Däniken, Thomas Peetz, and Mark Cieliebak. 2020. LEDGAR: A Large-Scale Multi-label Corpus for Text Classification of Legal Provisions in Contracts. In International Conference on Language Resources and Evaluation. https://api.semanticscholar.org/CorpusID:218974409
- [47] Wasi Uddin Ahmad, Jianfeng Chi, Yuan Tian, and Kai-Wei Chang. 2020. PolicyQA: A Reading Comprehension Dataset for Privacy Policies. arXiv e-prints, Article arXiv:2010.02557 (Oct. 2020), arXiv:2010.02557 pages. https://doi.org/10.48550/arXiv.2010. 02557 arXiv:cs.CL/2010.02557
- [48] Steven H. Wang, Antoine Scardigli, Leonard Tang, Wei Chen, Dimitry Levkin, Anya Chen, Spencer Ball, Thomas Woodside, Oliver Zhang, and Dan Hendrycks. 2023. MAUD: An Expert-Annotated Legal NLP Dataset for Merger Agreement Understanding. arXiv e-prints, Article arXiv:2301.00876 (Jan. 2023), arXiv:2301.00876 pages. https://doi.org/10.48550/arXiv.2301. 00876 arXiv:cs.CL/2301.00876
- [49] Weijie Yu, Zhongxiang Sun, Jun Xu, Zhenhua Dong, Xu Chen, Hongteng Xu, and Ji-Rong Wen. 2022. Explainable Legal Case Matching via Inverse Optimal Transport-based Rationale Extraction. arXiv e-prints, Article arXiv:2207.04182 (July 2022), arXiv:2207.04182 pages. https://doi.org/10.48550/arXiv.2207. 04182 arXiv:cs.IR/2207.04182
- [50] Lucia Zheng, Neel Guha, Brandon R. Anderson, Peter Henderson, and Daniel E. Ho. 2021. When Does Pretraining Help? Assessing Self-Supervised Learning for Law and the CaseHOLD Dataset. arXiv e-prints, Article arXiv:2104.08671 (April 2021), arXiv:2104.08671 pages. https://doi.org/10.48550/arXiv.2104. 08671 arXiv:cs.CL/2104.08671
- [51] Lucia Zheng, Neel Guha, Brandon R. Anderson, Peter Henderson, and Daniel E. Ho. 2021. When Does Pretraining Help? Assessing Self-Supervised Learning for Law and the CaseHOLD Dataset. arXiv e-prints, Article arXiv:2104.08671 (April 2021), arXiv:2104.08671 pages. https://doi.org/10.48550/arXiv.2104.

08671 arXiv:cs.CL/2104.08671