Building a Long Text Privacy Policy Corpus with Multi-Class Labels

Anonymous ACL submission

Abstract

This work introduces a new hand-coded dataset for the interpretation of privacy policies, containing the privacy policies from 149 firms, including documents incorporated by reference. The policies are coded across 64 dimensions that map onto commonly included terms and applicable legal rules. The coding method is designed to capture complexities inherent to legal interpretation and often absent from existing privacy policy datasets. These complexities include textual ambiguity, indeterminate meanings, interdependent clauses, contractual silence, and the implications of legal defaults.

1 Introduction

Privacy policies are important legal documents; they govern how firms collect, use, share, and secure personal information. They have become a prime target for automated interpretation. Despite their importance, they are rarely read. Privacy policies are long, complex, and require legal expertise to understand. At the same time, many privacy policies are publicly available and map to a consistent set of well-defined legal questions. This presents an opportunity to represent policy content in a reasonably consistent and somewhat objective manner against which automated interpreters can be tuned and measured.

Recent advancements in machine learning, especially the introduction of the large language model (LLM), increased interest in automating long-text interpretation. Practical legal use of NLP has expanded from term- and clause-level classification to nuanced interpretation of long bodies of legal text. Legal interpretation presents an especially challenging interpretative task. Privacy policies are a stereotypical legal document, as they are drafted by experts and include domain-specific vocabulary and interpretation (Zheng et al., 2021; Mellinkoff, 2004; Mertz, 2007); they can contain inconsistencies and be susceptible to multiple valid interpretations (Reidenberg et al., 2016); and they often contain interdependent clauses—sometimes spread across multiple documents—whose meaning is best understood when read wholesale. Like other legal texts, privacy policies must be interpreted in the context of applicable legal rules, which can define terms and provide guidance on issues not explicitly addressed in the text.

Current privacy policy datasets either offer high-granularity labels for short samples of policy text, or low-granularity classification of longer text. These approaches may not capture many domain-specific aspects of legal interpretation that are relevant to the expanding range of automated legal tasks. For example, neither approach accounts for how documents that are "incorporated by reference" may affect the way a policy restricts (or doesn't restrict) the ways in which a company can use user data. As legal interpretation increasingly becomes the target of automation, new datasets are needed. This paper aims to help address that need.

We provide a legal dataset of labeled online privacy policies coded by legally-trained experts. It contains 149 privacy policies and the documents they incorporate by reference, including Terms of Use, Cookie Policies, California Consumers Privacy Act (CCPA) disclosures, and documents pertaining to compliance with the European Union's General Data Protection Regulation (GDPR). Our coding accounts for the ways in which applicable legal rules and referenced documents can affect the meaning of terms. It also tracks relevant legal terms and reflects aspects of the legal interpretation task, including accounting for ambiguity and reasonable disagreements, and interpreting silence.

2 Related Work

Prior work building datasets for privacy policies mostly focuses on expert annotation or classifica-

[*] Equal contribution

tion of short text. (Lippi et al., 2019; Bui et al., 2021; Ahmad et al., 2021). Some investigation has also looked into crowd-sourcing annotation (Wilson et al., 2018). One privacy-policy-adjacent dataset involving classification of longer legal text labels the content of cookie banner disclosures with the stated purposes for data collection (Santos et al., 2021). In addition to annotated datasets, there are large-scale compilations of privacy policies scraped from the Internet and Internet Archive(Amos et al., 2021; Srinath et al., 2021).

Perhaps the most widely-used privacy policy dataset is the OPP-115 dataeset introduced by Wilson et. al. in 2016. OPP presented coders with paragraph-length excerpts from 115 privacy policies, to which coders added word-level annotations related to 36 data practices across 10 categories. The OPP dataset was used to train prominent tools used to pick out specific clauses from privacy policies (Harkous et al., 2018; Mousavi Nejad et al., 2020). It has also been used to generate related datasets, either by transforming its annotations for use in a new task like question-answering or GDPR compliance (Poplavska et al., 2020; Ahmad et al., 2020), or as an input into composite legal-task benchmarks like LEGALBENCH and PRIVACYGLUE (Guha et al., 2023; Chalkidis et al., 2022). The OPP taxonomy scheme has also been used to organize other privacy-related datasets (Ravichander et al., 2019). Another notable privacy policy dataset—the *unfair-TOS* dataset introduced by Lippi et. al.—annotates "potentially unfair" clauses in privacy practices and is also incorporated into some composite benchmarks, including the privacy-policy-specific PRIVACYGLUE benchmark (Shankar et al., 2023).

Benchmarking legal AIs goes beyond the traditional metrics-and-datasets approach. Alternative evaluation approaches include having NLP systems take the bar exam (Bommarito II and Katz, 2022; Katz et al., 2024) (though some have questioned the efficacy of that evaluation approach (Martínez, 2024)), grading LLM-generated law school exam answers (Choi et al., 2021), and measuring how law student performance is affected by LLM use (Choi and Schwarcz, 2023).

3 Dataset Preparation

3.1 Document Selection

Privacy policies vary significantly between services and across markets (Marotta-Wurgler, 2016b). To account for this, we evaluate policies from multiple markets and across a range of traffic rankings. We grouped US-based, english-language websites by Tranco traffic rating (Pochat et al., 2018). We selected 40 english-language websites from each of six ranking ranges: $[1, 10]$; $[10, 1000]$; $[1000, 10k]$; $[10k, 100k]$; $[100k, 1M]$; and $[1M, \infty]$. For ease of comparison, we first sampled from the websites that either were used in OPP-115 (Wilson et al., 2016) or in the legal empirical study that provides a basis for our coding scheme (Marotta-Wurgler, 2016b). When necessary, we fell back to random selection from english-language websites within that range. We then discarded any websites that had no privacy policy or terms of service linked from their landing page. We also discarded all but the highest-ranking from any set of urls sharing a single policy (e.g., facebook.com and instagram.com linked to the same policy). The resulting sample contained 174 policies. Our final dataset contains all "top ten" websites in the sample, along with 142 additional policies selected at random from the sample. Our dataset includes websites tagged with 23 of the 26 "tier 1" market categories in the IAB Content Taxonomy 2.0 (IAB Tech Lab, 2024). The remaining three categories are illegal content, religious content, and websites primarily about pets. A full list of urls, category tags, and Tranco ranks is included in appendix D.

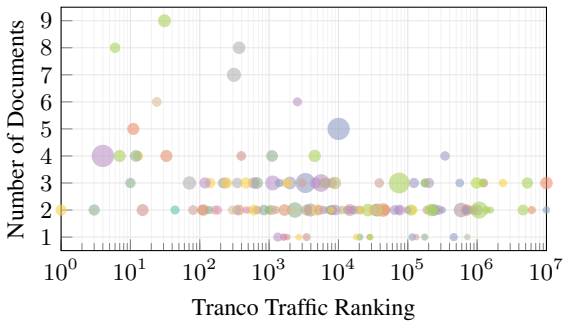

Figure 1: The documents in the dataset. The X axis shows the website's Tranco traffic rank. The Y axis reflects the number of documents collected for that site. Color and radius correspond to website industry (see Appendix D) and word count (ranging from 458 to $39,988$) respectively. An 18-document outlier is omitted from this chart.

Many privacy policies are spread across multiple linked documents. For example, Amazon's privacy policy's statement on cookie use reads: "*To enable our systems to recognize your browser or device and to provide and improve Amazon Services, we use cookies and other identifiers. For more*

information about cookies and how we use them, please read our Cookies Notice." The full terms regarding cookie use are found in the referenced "Cookies Notice." This is commonly referred to as "incorporation by reference." Terms incorporated by reference are enforceable by courts as long as the referenced documents are clearly identified and accessible.(ALI, 2024). Failing to include those referenced documents can result in incomplete (and potentially incorrect) interpretations of the policy. To compile the full set of included terms, we downloaded each website's privacy policy and terms of service, then iteratively added any referenced documents. Figure 1 shows the distribution of policies by rank and the number of distinct documents associated with each policy after applying this procedure.

3.2 Coding Process

Each policy was coded independently by two law students who had completed relevant coursework on contracts and received training regarding legal interpretation of privacy policies. All coders were compensated for their time at the appropriate rate for student employees at their institution. We presented each coder with all relevant legal documents corresponding to a firms' website and multiple-choice questions about those documents. For each question, coders were asked to highlight any and all text (including text in documents incorporated by reference) they found relevant to answering the question. They answered each question by selecting from a set of choices describing possible policy content, including, when applicable, the possibility that the policy was silent on a question. They also recorded their confidence in their answer on a Likert scale.

For the first ten weeks of coding, we held weekly overview meetings with each coder where we reviewed all codings. For each reviewed question where we identified a disagreement or error, we highlighted relevant text, and recorded our own answer, confidence, and notes. These reviews responses are included in the dataset for reference but are not used in this paper's summary statistics or benchmarks.

Coding, review, and project management were all performed using a suite of custom web tools we developed, as shown in figure 2. We provide a hosted version of the tools on our website. The tool includes several quality-of-life features, including offline mode, bookmarks and tabs in the document pane, and assignment and project management tools. We are releasing the tool as an open-source project in parallel with the dataset from this paper. The source code is available on our GitHub repository[1] and a hosted version of the tool is available at (omitted for ARR review).

Coders self-directed their approach to coding. Most coders spread their coding tasks across multiple sessions; using a six-week sample, we can estimate the length of each session by observing the time that the tool was open and did not lose input focus for more than ten minutes. Ignoring three multi-day outliers, the median contiguous coding session lasted 19 minutes, and total time to code a policy ranged from 51 to 283 minutes over 1 to 13 sessions, the median coding time was 89 minutes per policy, and the median policy took 4.5 sessions to code.

3.3 Coding Schema

We generated our coding schema following the procedure developed by Marotta-Wurgler (Marotta-Wurgler, 2016a). The approach has been used in legal empirical scholarship to make quantitative comparisons of privacy policy content and compliance between industries and over time (Marotta-Wurgler, 2016b,c; Davis and Marotta-Wurgler, 2019). The schema represents a policy content as a set of labels derived from significant and influential privacy guidelines and applicable rules that have shaped the content and structure of privacy policies. These are: the 1973 HEW Fair Information Practice Principles, the 2012 Federal Trade Commission's Information Privacy Guidelines, the 2012 White House Privacy Bill of Rights, the GDPR of 2018, and the CCPA of 2020, and contract law—the background rules courts have employed to enforce privacy policy. The resulting coding provides a granular representation of privacy policy content mentioned in relevant guidelines and laws. For example, there are three labels that encode the rights users and firms have with respect to changes to the policy (can the firm make changes, does a user have to assent to that change before it takes effect, and are changes retroactive). Another label marks whether the set of documents includes a class action waiver. The goal of a granular approach was to minimize ambiguity in representation of policy content and enhance consistency among coders. We translate these vari-

[1](omitted for ARR review)

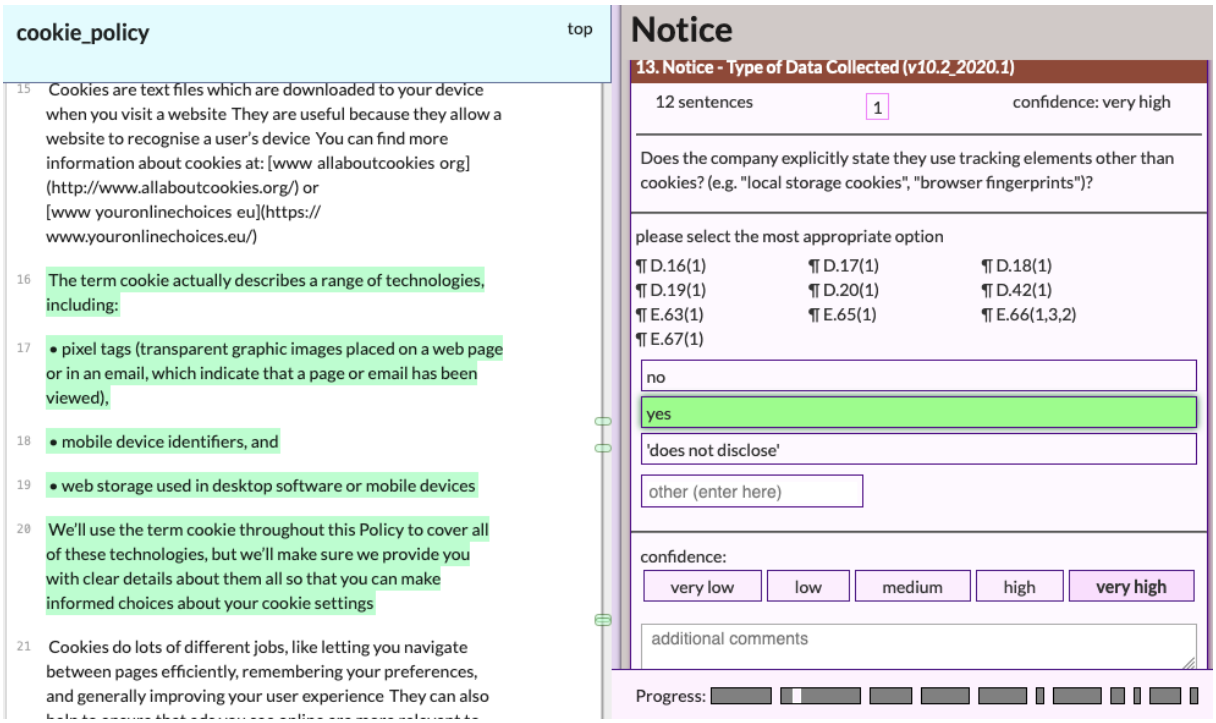

Figure 2: The highlight tool. Coders are shown all documents related to a policy on the left, and a list of questions on the right. They select a question, highlight any text from the documents they find relevant to answering the question.

ables into 64 multi-choice questions, which we group into 11 categories:

1. *CCPA* (10 labels): Tracks requirements unique to the California Consumer Protection Act, such as whether the subject can request that their personal information not be sold.
2. *GDPR* (5 labels): Tracks requirements unique to the GDPR, such as whether the entity has designated a Data Privacy Officer.
3. *Data Practices (DP)* (1 label): whether the firm has procedures to safely dispose of personal information.
4. *Enforcement (E)* (8 labels): Tracks mechanisms of legal redress.
5. *Notice (N)* (13 labels): Tracks notices pertaining to data collection and mandatory disclosures with state privacy laws.
6. *Contract (K)* (1 label): Tracks whether policy incorporates terms by reference.
7. *Privacy by Design (PBD)* (2 labels): Tracks general data practices in management and design.
8. *Security (SE)* (8 labels): Tracks information security practices.
9. *Sharing (SH)* (7 labels): Tracks sharing with third and other parties.
10. *User Control (UC)* (8 labels): Tracks user rights regarding personal information access and control.

For each question we drafted a set of answers designed to minimize ambiguity while providing granular representation of policy content. These include choice sets that are binary (*"Do the Terms of Use or Terms of Service incorporate the Privacy Policy by reference?" [yes/no]*), single-class (*"Does the Privacy Policy offer data requests by consumers explicitly free of charge?" [Not Applicable/Yes/No]*), and multi-class (*"Does the privacy policy provide means by which a user can contact the company with any privacy concerns or complaints? [select all that apply...]"*). Response options also include and distinguish between policies being silent regarding a term or the particular term being not applicable (e.g., a policy that states that no personal information is collected does not need to provide information about how such information is stored).

In contrast to other privacy policy datasets, which make efforts to maximize inter-coder agreement and often discard points of disagreement, we preserve disagreement and low-confidence coding. Because ambiguity is feature of many legal texts,

the ground truth is effectively probabilistic, meaning disagreement and low confidence are expected features of a legal interpretation classification task. To ensure that disagreements and reported low confidence correspond to ambiguity in the policy rather than unclear coding instructions, we engaged in an iterative process to reduce exogenous sources of ambiguity from our coding. Note that the goal is to ensure annotators are consistent in their coding approach, which may not always imply consistency in their coding outcomes. Our goal is to maximize the likelihood that inconsistencies between annotations reflect ambiguities in the legal text, not in the annotators' understanding of their instructions.

Once a week during the 10-week iterative revision period, we met with coders and discussed each of their coding choices. The recorded highlights for each question helped coders recall and explain their decision-making. We qualitatively assessed the source of each instance of inter-coder disagreement or or low reported confidence, choosing between five possible causes:

1. *Questions*: coders interpret the question in conflicting ways due to poor or confusing wording
2. *Choices*: the answer choice sets did not fully map onto the text and law
3. *Defaults*: the answer choice set does not properly account for the existence of default rules that alter the meaning of contractual silence
4. *Coder Error*: a coder made a mistake.
5. *Policy Text*: the text of the policy is ambiguous or susceptible to reasonable disagreements in interpretation

We addressed disagreements or confusion resulting from Categories 1, 2, and 3 by either adding clarifying details to the language of questions and answer choice sets, or adjusting the set of choices to better reflect the range of observed practices. When we detected coder error we corrected it and sent updated training guidance to other coders. We made no changes following disagreements and low confidence caused by unclear policy text. After any change to a question we removed any coding recorded using an outdated version of that question from the dataset. Coders relabeled those questions at the end of the iterative adjustment period.

Not every policy implicates the same questions or choices; some issues arose later in our revision window. After five rounds of iteration and revision, we stopped observing instances of the first three categories. The final set of coding instructions, including the history of changes corresponding to each variable, is included as an online appendix.

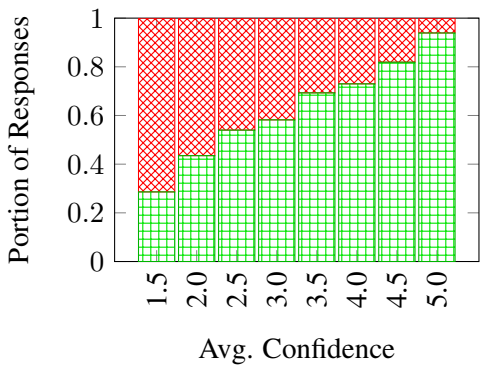

Figure 3: Pairwise agreement rate among coders by average self-reported confidence on a Likert scale.

As an initial sanity check on our data, we compare confidence and inter-coder agreement rates, as shown in figure 3. We observe that inter-coder agreement has a roughly linear relationship with self-reported confidence. This suggests that both disagreement and confidence correspond with situations where coders see multiple potentially appropriate answers. Our efforts to remove other sources of confusion and our use of expert coders mean that this ambiguity should largely come from the text of the policy.

3.4 Insights from Iterative Schema Refinement

To test whether our coding scheme and iterative refinements actually reduce measurement errors by reducing exogenous sources of ambiguity, we included the OPP annotation scheme for the first five weeks of our iterative process. OPP is a natural baseline to measure against: it is cited as the current "gold standard" privacy policy annotation scheme (Mousavi Nejad et al., 2020), and has been used to train numerous automated privacy policy interpreters and included in the LEGALBENCH and PRIVACYGLUE composite legal reasoning benchmarks. While the OPP taxonomy was originally designed to annotate short phrases by their associated data practices, the taxonomy and dataset has been adapted for several other contexts and tasks.

While initial inter-coder agreement rates were similar between our questions and the OPP schema, we found that clarity issues in the OPP scheme held steady week-over-week while our scheme's error

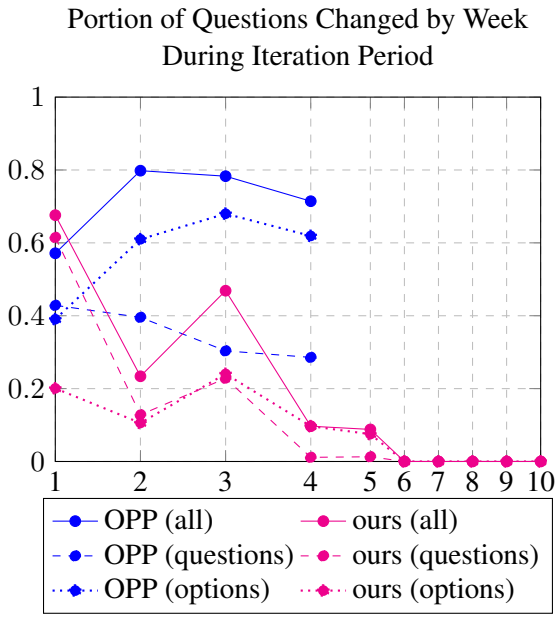

Figure 4: Portion of questions changed during iterative refinement, by coding scheme. OPP was removed from the scheme after week 5.

rate went down. This was true for both the classifications selected by coders, and the text coders marked as relevant to each question. Figure 4 shows the portion of questions changed after each week in our scheme and OPP. After four rounds of iterative adjustment, we removed the OPP annotations from our coding scheme. With enough iterative changes, both our scheme and OPP's would likely resolve; however, with each change we apply comparison with the original OPP-115 dataset and scheme becomes more difficult.

These results demonstrate how ambiguity and noise can come from the coding scheme or the way in which it is presented. While these results caution against uncritical reuse of the OPP taxonomy in new contexts, they do not apply directly to the OPP dataset. OPP was designed to pick out phrases disclosing data practices from short text samples; our scheme is designed to record the ways in which the full text of a privacy policy implicates the rights and responsibilities of firms and users.

We think that these differences in coder convergence between our scheme and OPP help justify our label selection method and iterative approach. There are two potential exogeneities worth addressing here. First, if training and regular meetings were only reason for convergence then we would expect all questions to converge. The OPP questions questions did not, which might suggest that our labels correspond more closely to the terms contained in a privacy policy. Second, if the difference between the two schemes were the result of our initial drafting choices, we would expect to see lower rates of confusion at the outset. Instead, we see similar rates of initial confusion that improve once we detect and correct our drafting issues, further supporting the notion that our labels are better-aligned with policy content.

Questions	64
Categories	11
Total Coders	18
Policies	149
Paragraphs	52,176
Words	1,524,570
Highlight Annotations	22,609
Policy Classifications	19,729
Confidence Scores	17,243

Table 1: Summary statistics on the corpus

4 Dataset Contents

The dataset comprises annotations for 149 documents, each coded by two or more coders. For each document and question, the dataset contains classifications selected by each coder, a list of sentences the coder marked as relevant to answering the question, and their self-reported confidence ranked on a Likert scale. The dataset includes 64 variables motivated by 11 legal categories. For a subset of documents and questions, the dataset also includes the amount of time each coder spent answering the question. Coders were all upper-level law students who had completed coursework covering the relevant topics in contract law. Table 1 provides additional descriptive statistics about the dataset.

Except for the CCPA category, correlation between question responses is low, as shown in Figure 5. Since answers to one question are not highly predictive of answers to another, we feel that this set of questions provides reasonable coverage over distinct legal interpretation tasks.

While achieving high intercoder agreement was explicitly not our goal, the dataset exhibits a moderate to high level of intercoder agreement, with Krippendorff's Alpha ranging from 0.40 to 0.96, averaging at 0.55. Absolute agreement rates across the dataset are 79%, though most disagreement is concentrated in 7 questions. This may suggest that there are a few areas where firms are more likely to use opaque language when describing their terms.

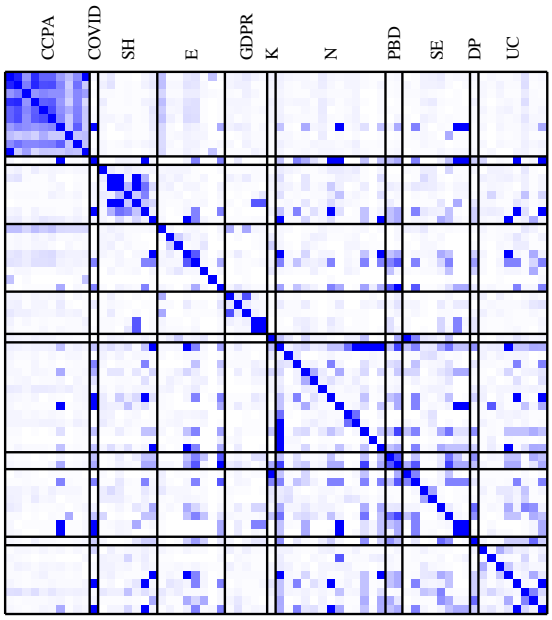

Figure 5: Correlation matrix showing correlation between question encodings. High correlation means that answers to questions covary by policy. For example, the high correlation between CCPA sections suggests that most policies are all-or-nothing along the CCPA-related dimensions we measure.

5 Results

Using our dataset, we evaluate current language models on their ability to perform two tasks. The first task ("holistic classification") is a multi-classification task that uses the entire policy as input: given our questions and a policy from our dataset, select the most likely answers for each question. The second task ("highlight prediction") is an annotation task that targets individual paragraphs: given a question from our dataset and a paragraph from one of the policies in our sample, predict whether a coder marked that paragraph as relevant to answering the question.

We evaluate the *holistic classification* task using batched cross-entropy loss. Each of the k policies in the dataset is associated with n sets of labels, $\{L_1, L_2, ..., L_n\}$, each corresponding to a question in our coding scheme. Each label L_i is associated with a set of options, $m_{i,j} \in M_i$. Given the ambiguity present in some privacy policies, the ground-truth value of L_i^k may not be a single value, but rather a probability distribution over M_i. Coder responses are therefore definitionally noisy. We compute the goal probability distribution $\mathbf{y}_i^k$ as $([c_{i1}^k, ..., c_{im}^k])$ normalized to sum to 1, where c_{ij}^k represents the number of coders who selected option j for question i on document k. We apply label smoothing to account for noise, as described in (Müller et al., 2020), setting α to .1. We use LLMs to generate a probability distributions p_i^k over the set of options for each label. When logprobs are available, we generate the distribution by crawling the response tree of each branch until an answer is selected or the net probability is negligible. We evaluate model responses by computing cross-entropy loss between the model's response and the reference distribution, $\frac{1}{m_i}\sum_{j=1}^{m_i}(y_{ij}^k \log(p_{ij}^k)+(1-y_{ij}^k)\log(1-p_{ij}^k))$. We report average loss by question, category, and across the entire dataset.

Because some policies contain more than 32 thousand tokens, we can only test LLMs with sufficiently large context windows without resorting to context-expanding techniques or alternative models, which are out of scope for this project. The performance of some commercial LLMs with sufficiently large context windows appears in table 2. For some of the questions in our dataset, some models performed worse than random guessing, a result we found surprising. The errors appear to be caused by the models incorrectly selecting "not applicable" and "does not disclose" options far too often.

We evaluate *highlight predictions*, a binary classification task, by concatenating individual paragraphs with question text and option descriptions. $y_{ij} = 1$ if at least one coder flagged paragraph j as relevant when answering question i, and 0 otherwise. We tested zero-shot labeling, prompting the model to answer whether the paragraph was relevant and computing the relative likelihood of an affirmative or negative response using the tree-crawling approach described above. The zero-shot accuracy, precision, and f_1 scores of several models are shown in table 3.

We note that performance varies significantly across categories and questions, including questions within the same category. While differences in performance between models may be an artifact of our prompt design, we found the variance between similar questions about similar topics striking. At least for the systems we tested, an LLM's ability to answer one legal question appears to not be predictive of that LLM's ability to answer other questions, even within extremely narrow domains like "properties of sharing practices described within a privacy policy."

Average Cross Entropy Loss

Category	Claude 3.5 "Sonnet"	Claude 3 "Haiku"	GPT-4o	GPT-4	random guesses
overall	0.174	0.309	0.175	0.217	0.458
CCPA	0.217	0.293	0.204	0.234	0.447
COVID	0.021	0.021	0.026	0.021	0.463
Data Production	0.084	0.293	0.214	0.266	0.466
Enforcement	0.092	0.215	0.118	0.158	0.469
GDPR	0.176	0.259	0.151	0.187	0.443
Contract	0.136	0.143	0.094	0.150	0.427
Notice	0.219	0.361	0.197	0.216	0.457
Privacy by Design	0.153	0.474	0.340	0.308	0.505
Security	0.154	0.295	0.134	0.233	0.466
Sharing	0.249	0.333	0.223	0.292	0.457
User Control	0.125	0.377	0.160	0.196	0.451

Table 2: Average cross-entropy on holistic classification task, by category.

model	**acc.**	**recall**	f_1
BERT-BASE	75.30	18.05	1.89
RoBERTa-BASE	45.41	32.72	1.48
LEGAL-BERT	53.43	29.19	1.57

Table 3: Average zero-shot performance on highlighting task, optimizing for f_1. Because highlighting is noisy and heavily skewed, we suspect a certain number of false positives are unavoidable. Current practical use-cases aim to identify relevant text from within policies (Lippi et al., 2019); setting a significantly higher β when computing thresholds (e.g., optimizing for f_{100} instead of f_1) can greatly increase recall.

6 Future Directions

This project is designed to contribute to the growing body of legal task corpora. We plan to add it to open-source legal benchmarks, such as the LEGALBENCH consolidated corpus.

One of the challenges of analyzing legal documents is how work-intensive it is, how ephemeral some documents are, and how difficult it can be to comparing documents across time, especially if they incorporate changing (external) legal contexts by reference. By releasing these tools and putting greater emphasis on reproducability, we plan to extend this dataset to observe how privacy policies change over time.

Our tools for classifying legal documents were intentionally designed to apply to other legal tasks, or to support future extensions of the question set and dataset. We hope to partner with other legal experts to expand this dataset to cover a broader range of legal questions and documents.

Finally, we have begun investigating the underlying cause of uneven rates of disagreement by question among coders (and, to a lesser extent, similarly uneven response rates between state-of-the-art LLMs). Future work may investigate whether firms are intentionally ambiguous to obscure practices or add flexibility, whether a mismatch between technology and law makes certain disclosure difficult or nonsensical, or whether some other factor is at play.

7 Conclusion

We have described a hand-coded dataset for the interpretation of privacy policies. It captures granular, multi-class data about 149 privacy policies and their associated documents along 64 dimensions, and is intended as a new resource for the development and benchmarking of NLP systems that interpret long legal text. Given the relatively poor performance of current state-of-the-art LLMs at task described in this dataset, we suspect that some aspects of the task that are not well represented by existing training data and benchmarks.

Our coding approach is designed to capture complexities inherent to the task of legal interpretation that are not present in current privacy policy datasets, such as addressing textual ambiguity, indeterminate meaning, interdependent clauses, contractual silence, and the effect of legal defaults. Along with our classification data, we include relevant-text annotations and confidence scores from each labeller. We supplement this dataset with our own coding of the questions where labellers disagree or report low confidence, which may provide additional insight into the textual ambiguities in the underlying policies. We include the tools we used to produce this dataset, including a hosted online tool that (non-technical) domain experts can use to produce similar classification datasets within their own areas of expertise.

Limitations

All but one of our population of coders learned consumer contracts (the relevant class for privacy policy interpretation) at the same law school from the same two law professors. They may have adopted some of those professors biases, or approach contract interpretation in similar ways. That overlap may have obscured lingering ambiguities in our coding scheme. It may also have biased them towards understanding a coding scheme designed by one of those professors. We think the latter possibility is fairly remote–privacy policies are a sufficiently esoteric corner of contract law; it receives very little dedicated class time.

Decision fatigue is always a concern with labeling. Though our method does not take any explicit steps to mitigate fatigue, there are three reasons we don't think that particular issue is likely to meaningfully to effect data quality. First, reading and interpreting the clauses of a contract is a common task for a lawyer; one that our coders were recently evaluated on. Second, our coders spread most of their coding across multiple sessions, reducing the risk of decision fatigue. Third, in the weekly sessions where coders discussed the reasoning behind their decisions with us, we did not observe any qualitative difference in the level of care or depth of discussion surrounding any particular subset of the coding questions.

As noted above, the "ground truth" meaning of a contract can be probabilistic. Our coders effectively took a noisy sample of each contract with $n = 2$. We think this is reasonable for two reasons. First, the two-coder approach matches the current state of the art for privacy policy datasets. Second, confidence seems to be a decent predictor of disagreement, which opens mitigation options. We didn't, but potentially could, explore mitigation options.

The noisiness of our measurements also means that our benchmarks in part 5 necessarily contain a (somewhat arbitrary) smoothing factor. We suspect that specific tasks might be better measured using other metrics, and that the smoothing factor could be tuned to reflect confidence.

Likert scores are notoriously messy, meaning our confidence measurements may not contain as much information as we'd ideally like to capture.

Our reported benchmark performance rely on the quality of our prompt design. We have more experience designing prompts for LEGALBERT and GPT-4; our measurement of Claude 3 and other BERT models may be influenced by a prompt that is better suited for GPT. (our prompt generation script is included in the associated github repo).

References

Wasi Uddin Ahmad, Jianfeng Chi, Tu Le, Thomas Norton, Yuan Tian, and Kai-Wei Chang. 2021. Intent classification and slot filling for privacy policies.

Wasi Uddin Ahmad, Jianfeng Chi, Yuan Tian, and Kai-Wei Chang. 2020. Policyqa: A reading comprehension dataset for privacy policies.

ALI. 2024. Adoption of standard contract terms. In *Restatement of Consumer Contracts*. American Law Institute.

Ryan Amos, Gunes Acar, Eli Lucherini, Mihir Kshirsagar, Arvind Narayanan, and Jonathan Mayer. 2021. Privacy policies over time: Curation and analysis of a million-document dataset. In *Proceedings of the Web Conference 2021*, pages 2165–2176.

Michael Bommarito II and Daniel Martin Katz. 2022. Gpt takes the bar exam. *arXiv preprint arXiv:2212.14402*.

Duc Bui, Kang G. Shin, Jong-Min Choi, and Junbum Shin. 2021. Automated extraction and presentation of data practices in privacy policies. *Proceedings on Privacy Enhancing Technologies*, 2021(2):88–110.

Ilias Chalkidis, Abhik Jana, Dirk Hartung, Michael Bommarito, Ion Androutsopoulos, Daniel Martin Katz, and Nikolaos Aletras. 2022. Lexglue: A benchmark dataset for legal language understanding in english.

Jonathan H Choi, Kristin E Hickman, Amy B Monahan, and Daniel Schwarcz. 2021. Chatgpt goes to law school. *J. Legal Educ.*, 71:387.

Jonathan H Choi and Daniel Schwarcz. 2023. Ai assistance in legal analysis: An empirical study. *Available at SSRN 4539836*.

Kevin E Davis and Florencia Marotta-Wurgler. 2019. Contracting for personal data. *NYUL Rev.*, 94:662.

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. Legalbench: A collaboratively built

benchmark for measuring legal reasoning in large language models.

Hamza Harkous, Kassem Fawaz, Rémi Lebret, Florian Schaub, Kang G. Shin, and Karl Aberer. 2018. Polisis: Automated analysis and presentation of privacy policies using deep learning.

IAB Tech Lab. 2024. IAB Content Taxonomy 3.0.

Daniel Martin Katz, Michael James Bommarito, Shang Gao, and Pablo Arredondo. 2024. Gpt-4 passes the bar exam. *Philosophical Transactions of the Royal Society A*, 382(2270):20230254.

Marco Lippi, Przemysław Pałka, Giuseppe Contissa, Francesca Lagioia, Hans-Wolfgang Micklitz, Giovanni Sartor, and Paolo Torroni. 2019. Claudette: an automated detector of potentially unfair clauses in online terms of service. *Artificial Intelligence and Law*, 27:117–139.

Florencia Marotta-Wurgler. 2016a. Self-regulation and competition in privacy policies. *The Journal of Legal Studies*, 45(S2):S13–S39.

Florencia Marotta-Wurgler. 2016b. Understanding Privacy Policies: Content, Self-Regulation, and Markets. SSRN Scholarly Paper ID 2736513, Rochester, NY.

Florencia Marotta-Wurgler. 2016c. Understanding privacy policies: Content, self-regulation, and markets. *NYU Law and Economics Research Paper*, (16-18).

Eric Martínez. 2024. Re-evaluating gpt-4's bar exam performance. *Artificial Intelligence and Law*, pages 1–24.

David Mellinkoff. 2004. *The language of the law*. Wipf and Stock Publishers.

Elizabeth Mertz. 2007. *The language of law school: learning to" think like a lawyer"*. Oxford University Press, USA.

Najmeh Mousavi Nejad, Pablo Jabat, Rostislav Nedelchev, Simon Scerri, and Damien Graux. 2020. Establishing a strong baseline for privacy policy classification. In *ICT Systems Security and Privacy Protection: 35th IFIP TC 11 International Conference, SEC 2020, Maribor, Slovenia, September 21–23, 2020, Proceedings 35*, pages 370–383. Springer.

Rafael Müller, Simon Kornblith, and Geoffrey Hinton. 2020. When does label smoothing help?

Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhoob, Maciej Korczyński, and Wouter Joosen. 2018. Tranco: A research-oriented top sites ranking hardened against manipulation. *arXiv preprint arXiv:1806.01156*.

Ellen Poplavska, Thomas B Norton, Shormir Wilson, and Norman Sadeh. 2020. From prescription to description: Mapping the gdpr to a privacy policy corpus annotation scheme. In *Legal Knowledge and Information Systems-JURIX 2020: 33rd Annual Conference*.

Abhilasha Ravichander, Alan W Black, Shomir Wilson, Thomas Norton, and Norman Sadeh. 2019. Question answering for privacy policies: Combining computational and legal perspectives.

Joel R Reidenberg, Jaspreet Bhatia, Travis D Breaux, and Thomas B Norton. 2016. Ambiguity in privacy policies and the impact of regulation. *The Journal of Legal Studies*, 45(S2):S163–S190.

Cristiana Santos, Arianna Rossi, Lorena Sanchez Chamorro, Kerstin Bongard-Blanchy, and Ruba Abu-Salma. 2021. Cookie banners, what's the purpose? analyzing cookie banner text through a legal lens. In *Proceedings of the 20th Workshop on Workshop on Privacy in the Electronic Society*, pages 187–194.

Atreya Shankar, Andreas Waldis, Christof Bless, Maria Andueza Rodriguez, and Luca Mazzola. 2023. Privacyglue: A benchmark dataset for general language understanding in privacy policies. *Applied Sciences*, 13(6):3701.

Mukund Srinath, Shomir Wilson, and C Lee Giles. 2021. Privacy at scale: Introducing the privaseer corpus of web privacy policies. 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)*. Association for Computational Linguistics.

Shomir Wilson, Florian Schaub, Aswarth Abhilash Dara, Frederick Liu, Sushain Cherivirala, Pedro Giovanni Leon, Mads Schaarup Andersen, Sebastian Zimmeck, Kanthashree Mysore Sathyendra, N Cameron Russell, et al. 2016. The creation and analysis of a website privacy policy corpus. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1330–1340.

Shomir Wilson, Florian Schaub, Frederick Liu, Kanthashree Mysore Sathyendra, Daniel Smullen, Sebastian Zimmeck, Rohan Ramanath, Peter Story, Fei Liu, Norman Sadeh, and Noah A. Smith. 2018. Analyzing privacy policies at scale: From crowdsourcing to automated annotations. *ACM Trans. Web*, 13(1).

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 of 53,000+ legal holdings. In *Proceedings of the eighteenth international conference on artificial intelligence and law*, pages 159–168.

A Question List

A.1 Privacy By Design (PBD)

Question ID	Question Text	Choices
PBD-1	Does the privacy policy require periodic compliance review of structural and technological data security measures?	• No • Yes
PBD-2	Does the privacy policy contain self-reporting measures in case the firm experiences a privacy violation to, for example, a privacy seal organization or third party consultant?	• No • Yes

A.2 Contract (K)

Question ID	Question Text	Choices
K-1	Do the Terms of Use or Terms of Service incorporate the Privacy Policy by reference?	• No • Yes

A.3 CCPA (CCPA)

Question ID	Question Text	Choices
CCPA-1	Does the privacy policy include a link to the CCPA section, as opposed to in the same privacy policy?	• N/A: there is no CCPA section or CCPA reference in the contract • No, it's part of the same privacy policy • Yes, it's on a separate link or separate document
CCPA-2	Does the privacy policy state that the firm's CCPA policy only applies to California residents? For example, does it inlcude a statement similar to the following one: "This California section supplements the Privacy Policy and applies solely to California consumers (excluding our personnel). The Table below describes how we process California consumers' personal information (excluding our personnel), based on definitions laid out in the California Consumer Privacy Act ("CCPA")."	• N/A: there is no CCPA section or CCPA reference in the contract • No, • Yes,
CCPA-3	Does the privacy policy include a California Privacy Rights Section that explains all rights afforded to users and consumers under the CCPA? For example, these include: the right to request disclosure of business' data collection and sales practices; the categories of personal information collected, the source of the information; use of the information and, if the information was disclosed or sold to third parties, the categories of personal information disclosed or sold to third parties and the categories of third parties to whom such information was disclosed or sold; the right to request a copy of the specific personal information collected about them during the 12 months before their request (together with right #1, a "personal information request"); the right to have such information deleted (with exceptions); the right to request that their personal information not be sold to third parties, if applicable; and the right not to be discriminated against because they exercised any of the new rights.	• N/A • No • Yes, fully compliant • It is partially compliant.
CCPA-4	Does the privacy policy directs California Residents to the CCPA section when describing general, non-California exclusive, data practices?	• N/A -there is no CCPA section or CCPA reference in the contract • No • Yes

CCPA-5	Does the privacy policy offer California residents an opportunity to request all information shared with third parties in the last year?	• N/A • No or does not disclose • Yes
CCPA-6	Does the privacy policy offer California residents a direct link via which to contact site and request information?	• N/A • The contact information appears in a different section • The contact information appears in the same section • The privacy policy includes a direct link in the same section.
CCPA-7	Does the privacy policy offer data requests by consumers or users explicitly free of charge?	• N/A • No • Yes
CCPA-8	Does the privacy policy list the categories of personal information sold in the past 12 months?	• N/A • No • Yes
CCPA-9	Does the privacy policy identify at least two methods for submitting a personally identifiable information or erasure request, in accordance with CCPA? These must include, at a minimum, a web page and a toll-free telephone number.	• No • The privacy policy identifies two methods • The privacy policy identifies one method • two methods identified, but they are different from web page and toll-free number.
CCPA-10	Does the privacy policy offer consumers or users the right to opt-out of selling personal information to third parties with a visible, direct link to "Do Not Sell My Personal Information"?	• N/A [this would be applicable if, for example, the company did not share information with third parties] • No mention • Yes • The privacy policy mentions the right but explains why it's not available or not applicable (e.g., firm does not sell information to third parties)

A.4 Data Practices (DP)

Question ID	Question Text	Choices
DP-1	Does company have a procedure for safely disposing unused or no longer needed data or personally identifiable information?	• No or doesn't say • Yes

A.5 Sharing (SH)

Question ID	Question Text	Choices
SH-1	Are affiliates or subsidiaries bound to this privacy policy, confidentiality agreements, or have contractual obligations outlining how the shared data will be used or secured?	• N/A • No • Yes • Does not disclose

SH-2	Are contractors, service providers, or processors (for example, payment process companies) bound by either the same privacy policy, confidentiality agreements, or are under contractual obligations outlining how data will be used and secured?	• N/A • No • Yes • Does not disclose
SH-3	Are third parties bound by the same privacy policy?	• N/A [this would be applicable if, for example, the company did not share information with third parties] • No • Yes • Does not disclose
SH-4	Does the company perform due diligence to ensure the legitimacy of third parties that may have access to personally identifiable information?	• No • Yes • Does not disclose
SH-5	Does the company have a contract with third parties, other than processors or service providers, establishing how the shared persinally identifiable data can be used?	• N/A • No • Yes • Does not disclose
SH-6	Does the privacy policy provide hyperlinks to the privacy policies of relevant third parties' PP's? For example, sometimes the privacy policy includes links to third party privacy policies when it states that any engagement with third parties will be governed by third party privacy policies	• N/A, for example in cases where there are no relationships with third parties of any sort • No • Yes
SH-7	What is consent mechanism for sharing or selling personally identifiable or sensitive information with entities that are not service providers? Please do not consider service providers whose function is to effect, administer, or enforce a transaction, send future correspondence to user, or perform research, internal database compilation, or servicing the website?	• N/A, in cases when the firm does not engage in this practice • Mandatory or does not disclose • The firm offers an opt-out choice. That is, the default is for the user to share personally identifiable information with third parties, but the privacy policy gives the user the opportunity to opt out • Opt-in. This is when the user must consent before personally identifiable information or data can be shared with or collected or used by a third party.

A.6 Enforcement (E)

Question ID	Question Text	Choices
E-1	Does the privacy policy provides means by which user can contact the company with any privacy concerns or complaints? Please select all that apply.	• No • Yes, for all users • Yes, and mentions it in accordance with CCPA [or listed under a CCPA section in PP] • Yes, and mentions it in accordance with GDPR [or listed under a GDPR section in PP]

E-2	Do the privacy policy or the Terms of Use have a forum selection clause? If so, which forum?	• No • Yes
E-3	Do the privacy policy or Terms of Use have choice of law clause? If so, which law?	• No • Yes
E-4	Do the privacy policy or Terms of Use have an arbitration clause?	• No • Yes • consumer may choose arbitration
E-5	Do the privacy policy or terms of use have a class action waiver?	• No • Yes
E-6	Do the privacy policy or terms of use disclaims liability for failure of security measures?	• No • Yes
E-7	Does the privacy policy provides a link to the Federal Trade Commision's Consumer Complaint Form or does it inlcude t he FTC telephone number?	• No • Yes
E-8	Does the privacy policy include a privacy seal, certification, or industry oversight organization, other than those mandated by international law, such as the Swiss Privacy Law? Privacy Seals are independent, third-party enforcement programs to monitor company practices and enforce privacy policies. They are designed to provide protection to consumers by allowing Web companies to standardize privacy policies. Privacy seal programs include, among others, TRUSTe, BBBOnline, and CPA Webtrust. These are different from regulatory compliance seals, such as those that the company complies with COPPA, the Children Online Privacy Protection Act).	• No • Yes

A.7 GDPR (GDPR)

Question ID	Question Text	Choices
GDPR-1	Does the privacy policy states that it complies with GDPR or it includes section on GDPR compliance?	• No or no mention • Yes
GDPR-2	Does the privacy policy state that it complies with EU-US Privacy Shield?	• No • Yes
GDPR-3	Does the privacy policy state that GDPR terms apply only and exclusively to EU residents?	• N/A [This is when there are no GDPR terms] • No • Yes
GDPR-4	Are users or consumers able to object to the processing or automated decision making that could impact them? This is only applicable if company does profiling or any other automated decision making, such as algorithmic decision making, or any automated decisions that don't involve a human.	• N/A • No • Yes • Does not disclose
GDPR-5	If the privacy policy state that the firm engages in automated decision making, does it provide meaningful information about the logic involved, or significance or effect of such decisions?	• N/A • No • Yes • Does not disclose

A.8 Notice (N)

Question ID	Question Text	Choices
N-1	Does the Privacy Policy include the company's cookie policy, such as an explanation on the text, or a hyperlink to a document with a cookie policy?	• No • Yes
N-2	Does the Privacy Policy note or explain that the company uses tracking elements other than cookies, such as local storage cookies, browser fingerprints, or other non-cookie tracking elements?	• No • Yes • Does not disclose
N-3	Does the privacy policy state that the company collects or stores biometric information, such as facial scans, fingerprints, facial patterns, voice or typing cadence?	• No • Yes • Does not disclose
N-4	Does the Privacy Policy include a statement noting that that personally identifiable nformation will be used internally only for business purposes, such as for effecting, administering, or enforcing a transaction, or for sending future correspondence to the user, or for research, internal database compilation, or for servicing the website? Note that using the data for advertising is not considered an internal business purpose.	• No • Yes • Yes, but there are other terms in the policy that conflict with such commitment.
N-5	Does the privacy policy include a commitment by the company to use personally identifiable information only for stated, context specific, purposes? These are purposes that a user would expect in the context of the service provided, such as users expecting their personal profiles made available to other users in a dating site?	• No • Yes • Yes, but there are other terms in the policy that conflict with such commitment.
N-6	Are third parties allowed to place advertisements that may track user behavior?	• No • Yes • Does not disclose
N-7	Does the privacy policy identifiy third party recipients of shared or sold data?	• Provides a general answer, such as "trusted 3rd parties" • The privacy policy includes a specific or named entity, or named category of third party • Does not disclose
N-8	Does the privacy policy define words such as "affiliates" or "third parties," if it uses them?	• N/A • No • Yes
N-9	Does the privacy policy include a "Change of Terms" or modifiacation provision that allows the firm to change the privacy policy?	• No • Yes
N-10	Does the privacy policy require the user/consumer to explicitly assent to any material changes?	• N/A • No, or does not say • Yes
N-11	Does the privacy policy states that material changes made to the policy will be retroactive or apply to previous data collection?	• No • Yes, it applies before data collected before the change.
N-12	Does the privacy policy summarize the key terms at the top of the policy? Just a table of contents doesn't count.	• No • Yes

N-13	Does the privacy policy explain any data procedures if company is sold or otherwise ceases to exist by, for example, filing for bankruptcy?	• No • Yes

A.9 COVID (COVID)

Question ID	Question Text	Choices
COVID-1	Does the privacy policy include any terms related to contact tracing, health tracking, or other terms in relationship to COVID?	• No • Yes

A.10 Security (SE)

Question ID	Question Text	Choices
SE-1	Is there a term in the Privacy Policy or Terms of Use guaranteeing data accuracy?	• No • Yes
SE-2	Does the privacy policy specify any reasonable procedures the company many have in place to ensure data accuracy?	• No • Yes
SE-3	Does the privacy policy note whether the firm reserves a right to disclose protected personally identifiable information to comply with law or prevent crime?	• No or doesn't say • Yes
SE-4	Does the firm preserve the right to disclose protected identifiable information to protect its own rights?	• No or doesn't say • Yes
SE-5	Will users or consumers be given notice of any government requests for information about the user?	• No or doesn't say • Yes
SE-6	Does the privacy policy state whether the user will be notified in case there is a data breach?	• No (or does not disclose) • Yes
SE-7	Does the privacy policy describe any substantive privacy and security protections incorporated into firm's managerial or structural procedures, such as limiting the number of employees who have access to personally identifiable data, allowing personally identifiable data access only for job-related functions, assigning employees to oversee privacy issues, employing Chief Privacy Officer, o rrequiring periodic audits?	• No • Yes
SE-8	Does the privacy policy identify which means of technological security it employs, such as encryption?	• No • Yes

A.11 User Control (UC)

Question ID	Question Text	Choices
UC-1	Does the privacy policy allow the user or consumer to request that incorrect data be either rectified, updated, or erased?	• No • Yes, within 30 days (in compliance with GDPR) • Yes, no time limit • Does not disclose
UC-2	Does the privacy policy allow users or consumers to adjust their privacy settings? Note that directing the user to control cookies through settings in the browser doesn't count. The ansswer to this question may also be found by exploring the privacy settings of the service.	• No • Yes

UC-3	Does the privacy policy allow users or consumers to access and correct or update any personally identifiable information collected by the company?	• No • Users or consumers can only acess data • Users or consumers can access and correct data • Users or consumers can access and correct data, and third parties will be notified of the correction.
UC-4	Does the privacy policy allow the user to request that personally identifiable information be deleted or anonymized?	• No • Yes, but partially. This is when the user or consumer can delete or anonymize the account, but the company may continue to keep some of the user's data • Yes. This is when the consumer or the user can delete account and all of the user's information is removed from company, its servers and databases.
UC-5	Do the Term of Use include a term explaining the ownership of the user's or consumer's personally identifiable information?	• The company owns this information • User owns the information, but licenses data to entity in a non-exclusive, royalty-free, form, with no right to compensation for company's use of the information • The user owns the information but the license is not as broad as in [the previous option], where the company can use, share, and commercialize data/proprietary media as it sees fit • Does not disclose
UC-6	Does the privacy policy state what happens to the data or personally identifiable information the company collects if the firm ceases to exist or is acquired?	• Sold with the company or otherwise distributed or disclosed [e.g., when the company liquidates and the assets are sold piecemal] • Sold but it is given continued protection under same privacy policy or transferred to acquiring entity • Destroyed, anonymized, etc. • The user is given choice as to what happens with data
UC-7	If company is sold or goes bankrupt, is the user or consumer given choice as to what happens to their data or personally identifiable information?	• No • Yes
UC-8	Does the privacy policy explain what happens to the personally identifialbe infomration of a user who quits the service or closes the account?	• Retained and treated as if the consumer or user is still using service • Retained but modified • Deleted or anonymized • Does not disclose

B LLM Performance on Holistic Reading Task Benchmark

The following tables shows the binary cross entropy between coding produced by several commercial LLMs and the coding provided in our dataset on a per-question and per-category basis.

We used the following prompt across all models, adjusting data format to match each LLM's API. As with any LLM, additional or model-specific prompt design might result in significantly different results.

```
{
  messages: {
    role: "system",
    content: [
      {
        type: "text",
        text: """You are an AI assistant tasked with answering multiple-choice
            questions about a collection of legal documents that comprise a website'
            s privacy policy.
        You respond with a JSON object containing two fields: "answer" (a character
            or list of characters) and "confidence" (an integer between 1-5, where 1
             is least confident, and 5 is most confident)."""
      }, {type: "text", text: (first document as markdown)}, {type: "text", text: (
          second document as markdown)}, ...
    ]
  }, {
    role: "user",
    content: [
      {
        type:"text",
        text: """<question>
        (content of the question)
         - "A": (first option)
         - "B": (second option)
         ...
        </question>
        <instructions>
        Select the best answer for the provided question
        </instructions>"""
      }
    ]
  },
  return_type: {
    type: "json_schema",
    json_schema: {
    schema: {
      answer: {type: "string", pattern: "^[A-Z]$"},
      "confidence": {"type": "integer", "enum": [1,2,3,4,5]}
    },
  },
  temperature: 0
}
```

B.1 Performance by Category

Category	Claude 3.5 "Sonnet"	Claude 3 "Haiku"	GPT-4o	GPT-4	random guesses
CCPA	0.217	0.293	0.204	0.234	0.447
COVID	0.021	0.021	0.026	0.021	0.463
DP	0.084	0.293	0.214	0.266	0.466
E	0.092	0.215	0.118	0.158	0.469
GDPR	0.176	0.259	0.151	0.187	0.443
K	0.136	0.143	0.094	0.150	0.427
N	0.219	0.361	0.197	0.216	0.457
PBD	0.153	0.474	0.340	0.308	0.505
SE	0.154	0.295	0.134	0.233	0.466
SH	0.249	0.333	0.223	0.292	0.457
UC	0.125	0.377	0.160	0.196	0.451

B.2 CCPA

Question ID	Claude 3.5 "Sonnet"	Claude 3 "Haiku"	GPT-4o	GPT-4	random guesses
CCPA-1	0.063	0.108	0.073	0.091	0.469
CCPA-2	0.191	0.357	0.098	0.103	0.527
CCPA-3	0.226	0.246	0.200	0.219	0.427
CCPA-4	0.179	0.289	0.157	0.275	0.455
CCPA-5	0.314	0.314	0.311	0.312	0.423
CCPA-6	0.157	0.344	0.150	0.199	0.444
CCPA-7	0.291	0.347	0.288	0.303	0.450
CCPA-8	0.324	0.316	0.321	0.337	0.404
CCPA-9	0.187	0.273	0.185	0.212	0.464
CCPA-10	0.212	0.324	0.224	0.257	0.420

B.3 Enforcement

Question ID	Claude 3.5 "Sonnet"	Claude 3 "Haiku"	GPT-4o	GPT-4	random guesses
E-1	0.120	0.209	0.147	0.160	0.436
E-2	0.184	0.177	0.255	0.218	0.471
E-3	0.055	0.089	0.221	0.075	0.532
E-4	0.054	0.178	0.044	0.105	0.478
E-5	0.082	0.236	0.017	0.207	0.476
E-6	0.140	0.219	0.142	0.161	0.401
E-7	0.006	0.021	0.008	0.008	0.472
E-8	0.092	0.584	0.113	0.332	0.489

B.4 GDPR

Question ID	Claude 3.5 "Sonnet"	Claude 3 "Haiku"	GPT-4o	GPT-4	random guesses
GDPR-1	0.101	0.156	0.085	0.085	0.462
GDPR-2	0.038	0.038	0.031	0.036	0.427
GDPR-3	0.274	0.478	0.234	0.244	0.423
GDPR-4	0.252	0.347	0.223	0.314	0.462
GDPR-5	0.208	0.275	0.179	0.244	0.436

B.5 Notice

Question ID	Claude 3.5 "Sonnet"	Claude 3 "Haiku"	GPT-4o	GPT-4	random guesses
N-1	0.094	0.097	0.097	0.090	0.458
N-2	0.134	0.138	0.110	0.262	0.444
N-3	0.249	0.326	0.252	0.337	0.423
N-4	0.481	0.438	0.473	0.325	0.486
N-5	0.500	0.446	0.459	0.433	0.447
N-6	0.102	0.138	0.091	0.106	0.464
N-7	0.274	0.345	0.279	0.268	0.405
N-8	0.105	0.405	0.056	0.074	0.466
N-9	0.035	0.049	0.028	0.049	0.499
N-10	0.172	0.453	0.104	0.198	0.489
N-11	0.063	0.685	0.057	0.057	0.464
N-12	0.198	0.562	0.113	0.136	0.447
N-13	0.444	0.592	0.444	0.481	0.444

B.6 Privacy By Design

Question ID	Claude 3.5 "Sonnet"	Claude 3 "Haiku"	GPT-4o	GPT-4	random guesses
PBD-1	0.199	0.556	0.603	0.543	0.522
PBD-2	0.107	0.393	0.078	0.073	0.489

B.7 Security

Question ID	Claude 3.5 "Sonnet"	Claude 3 "Haiku"	GPT-4o	GPT-4	random guesses
SE-1	0.072	0.058	0.052	0.045	0.481
SE-2	0.632	0.673	0.388	0.619	0.434
SE-3	0.057	0.077	0.024	0.044	0.470
SE-4	0.085	0.154	0.085	0.132	0.445
SE-5	0.020	0.223	0.041	0.264	0.520
SE-6	0.066	0.612	0.093	0.335	0.474
SE-7	0.215	0.409	0.308	0.315	0.440
SE-8	0.078	0.152	0.078	0.112	0.466

B.8 Contract

Question ID	Claude 3.5 "Sonnet"	Claude 3 "Haiku"	GPT-4o	GPT-4	random guesses
K-1	0.136	0.143	0.094	0.150	0.427

B.9 Sharing

Question ID	Claude 3.5 "Sonnet"	Claude 3 "Haiku"	GPT-4o	GPT-4	random guesses
SH-1	0.235	0.205	0.202	0.190	0.450
SH-2	0.153	0.212	0.163	0.197	0.454
SH-3	0.148	0.222	0.150	0.159	0.482
SH-4	0.173	0.561	0.176	0.433	0.487
SH-5	0.344	0.305	0.261	0.308	0.430
SH-6	0.374	0.485	0.294	0.428	0.436
SH-7	0.313	0.344	0.314	0.331	0.462

C Results for Highlighting Task

C.1 BERT

C.1.1 Performance by Category

Category	r	a	p	f_1
CCPA (CCPA)	11.90	75.52	1.07	1.36
COVID (COVID)	0.87	72.88	0.09	0.16
Contract (K)	7.16	96.45	1.31	2.04
Data Practices (DP)	0.68	98.95	0.68	0.68
Enforcement (E)	22.05	66.11	1.29	2.01
GDPR (GDPR)	5.52	88.37	0.84	1.00
Notice (N)	35.57	54.21	1.67	2.77
Privacy By Design (PBD)	2.18	95.92	0.30	0.50
Security (SE)	21.60	74.69	1.37	1.93
Sharing (SH)	6.16	95.36	1.71	2.23
User Control (UC)	17.63	82.95	1.19	1.95

C.1.2 CCPA (CCPA)

Question ID	r	a	p	f_1
CCPA-1	1.01	99.24	0.76	0.81
CCPA-2	17.49	68.96	0.37	0.72
CCPA-3	32.97	9.78	1.05	1.98
CCPA-4	34.98	8.96	0.23	0.45
CCPA-5	4.14	98.15	2.13	1.96
CCPA-6	6.62	95.17	1.33	1.80
CCPA-7	0.67	97.67	0.34	0.45
CCPA-8	13.81	84.77	1.04	1.80
CCPA-9	4.47	96.82	2.76	2.75
CCPA-10	2.82	95.71	0.65	0.86

C.1.3 COVID (COVID)

Question ID	r	a	p	f_1
COVID-1	0.87	72.88	0.09	0.16

C.1.4 Data Practices (DP)

Question ID	r	a	p	f_1
DP-1	0.68	98.95	0.68	0.68

C.1.5 Enforcement (E)

Question ID	r	a	p	f_1
E-1	41.56	69.30	2.73	4.86
E-2	4.79	96.54	1.47	2.17
E-3	12.90	90.14	1.03	1.79
E-4	54.11	0.56	0.56	1.09
E-5	39.04	0.25	0.25	0.50
E-6	18.10	89.53	3.26	4.39
E-7	1.34	96.77	0.75	0.81
E-8	4.52	85.82	0.26	0.46

C.1.6 GDPR (GDPR)

Question ID	r	a	p	f_1
GDPR-1	14.73	60.31	0.52	0.96
GDPR-2	2.53	97.94	1.36	1.36
GDPR-3	1.88	98.49	1.69	1.55
GDPR-4	4.76	89.32	0.33	0.59
GDPR-5	3.69	95.80	0.31	0.56

C.1.7 Contract (K)

Question ID	r	a	p	f_1
K-1	7.16	96.45	1.31	2.04

C.1.8 Notice (N)

Question ID	r	a	p	f_1
N-1	22.24	86.37	3.18	4.76
N-2	14.95	94.53	4.26	5.88
N-3	4.14	96.60	1.16	1.73
N-4	66.21	30.03	2.38	4.33
N-5	28.83	63.00	1.68	2.95
N-6	28.57	88.11	2.96	5.02
N-7	90.10	2.61	1.95	3.75
N-8	27.92	6.56	0.43	0.82
N-9	50.06	59.04	1.33	2.45
N-10	56.60	21.14	0.67	1.29
N-11	51.35	3.80	0.41	0.79
N-12	16.50	62.15	0.75	1.37
N-13	4.92	90.78	0.50	0.85

C.1.9 Privacy By Design (PBD)

Question ID	r	a	p	f_1
PBD-1	1.34	97.45	0.34	0.54
PBD-2	3.02	94.39	0.26	0.46

C.1.10 Security (SE)

Question ID	r	a	p	f_1
SE-1	2.04	95.83	1.04	1.25
SE-2	3.24	98.49	2.85	2.77
SE-3	88.59	0.98	0.98	1.92
SE-4	27.91	67.06	0.97	1.84
SE-5	16.38	55.57	0.31	0.60
SE-6	1.01	99.62	1.01	1.01
SE-7	12.11	94.21	2.39	3.62
SE-8	21.48	85.79	1.36	2.44

C.1.11 Sharing (SH)

Question ID	r	a	p	f_1
SH-1	4.82	96.45	1.09	1.59
SH-2	4.50	97.14	1.63	2.08
SH-3	8.02	95.04	2.13	3.18
SH-4	0.68	99.59	0.68	0.68
SH-5	12.93	86.74	0.84	1.53
SH-6	4.70	98.71	3.21	3.51
SH-7	7.44	93.86	2.40	3.03

C.1.12 User Control (UC)

Question ID	r	a	p	f_1
UC-1	19.30	80.65	1.04	1.92
UC-2	16.76	75.88	0.94	1.59
UC-3	10.99	92.91	1.76	2.76
UC-4	40.99	68.18	1.39	2.64
UC-5	5.36	94.70	2.14	2.53
UC-6	17.23	83.10	0.69	1.30
UC-7	14.14	81.27	0.58	1.06
UC-8	16.26	86.90	0.99	1.79

C.1.13 CCPA (CCPA)

Question ID	r	a	p	f_1
CCPA-1	1.01	99.24	0.76	0.81
CCPA-2	17.49	68.96	0.37	0.72
CCPA-3	32.97	9.78	1.05	1.98
CCPA-4	34.98	8.96	0.23	0.45
CCPA-5	4.14	98.15	2.13	1.96
CCPA-6	6.62	95.17	1.33	1.80
CCPA-7	0.67	97.67	0.34	0.45
CCPA-8	13.81	84.77	1.04	1.80
CCPA-9	4.47	96.82	2.76	2.75
CCPA-10	2.82	95.71	0.65	0.86

C.1.14 COVID (COVID)

Question ID	r	a	p	f_1
COVID-1	0.87	72.88	0.09	0.16

C.1.15 Data Practices (DP)

Question ID	r	a	p	f_1
DP-1	0.68	98.95	0.68	0.68

C.1.16 Enforcement (E)

Question ID	r	a	p	f_1
E-1	41.56	69.30	2.73	4.86
E-2	4.79	96.54	1.47	2.17
E-3	12.90	90.14	1.03	1.79
E-4	54.11	0.56	0.56	1.09
E-5	39.04	0.25	0.25	0.50
E-6	18.10	89.53	3.26	4.39
E-7	1.34	96.77	0.75	0.81
E-8	4.52	85.82	0.26	0.46

C.1.17 GDPR (GDPR)

Question ID	r	a	p	f_1
GDPR-1	14.73	60.31	0.52	0.96
GDPR-2	2.53	97.94	1.36	1.36
GDPR-3	1.88	98.49	1.69	1.55
GDPR-4	4.76	89.32	0.33	0.59
GDPR-5	3.69	95.80	0.31	0.56

C.1.18 Contract (K)

Question ID	r	a	p	f_1
K-1	7.16	96.45	1.31	2.04

C.1.19 Notice (N)

Question ID	r	a	p	f_1
N-1	22.24	86.37	3.18	4.76
N-2	14.95	94.53	4.26	5.88
N-3	4.14	96.60	1.16	1.73
N-4	66.21	30.03	2.38	4.33
N-5	28.83	63.00	1.68	2.95
N-6	28.57	88.11	2.96	5.02
N-7	90.10	2.61	1.95	3.75
N-8	27.92	6.56	0.43	0.82
N-9	50.06	59.04	1.33	2.45
N-10	56.60	21.14	0.67	1.29
N-11	51.35	3.80	0.41	0.79
N-12	16.50	62.15	0.75	1.37
N-13	4.92	90.78	0.50	0.85

C.1.20 Privacy By Design (PBD)

Question ID	r	a	p	f_1
PBD-1	1.34	97.45	0.34	0.54
PBD-2	3.02	94.39	0.26	0.46

C.1.21 Security (SE)

Question ID	r	a	p	f_1
SE-1	2.04	95.83	1.04	1.25
SE-2	3.24	98.49	2.85	2.77
SE-3	88.59	0.98	0.98	1.92
SE-4	27.91	67.06	0.97	1.84
SE-5	16.38	55.57	0.31	0.60
SE-6	1.01	99.62	1.01	1.01
SE-7	12.11	94.21	2.39	3.62
SE-8	21.48	85.79	1.36	2.44

C.1.22 Sharing (SH)

Question ID	r	a	p	f_1
SH-1	4.82	96.45	1.09	1.59
SH-2	4.50	97.14	1.63	2.08
SH-3	8.02	95.04	2.13	3.18
SH-4	0.68	99.59	0.68	0.68
SH-5	12.93	86.74	0.84	1.53
SH-6	4.70	98.71	3.21	3.51
SH-7	7.44	93.86	2.40	3.03

C.1.23 User Control (UC)

Question ID	r	a	p	f_1
UC-1	19.30	80.65	1.04	1.92
UC-2	16.76	75.88	0.94	1.59
UC-3	10.99	92.91	1.76	2.76
UC-4	40.99	68.18	1.39	2.64
UC-5	5.36	94.70	2.14	2.53
UC-6	17.23	83.10	0.69	1.30
UC-7	14.14	81.27	0.58	1.06
UC-8	16.26	86.90	0.99	1.79

C.1.24 CCPA (CCPA)

Question ID	r	a	p	f_1
CCPA-1	1.01	99.24	0.76	0.81
CCPA-2	17.49	68.96	0.37	0.72
CCPA-3	32.97	9.78	1.05	1.98
CCPA-4	34.98	8.96	0.23	0.45
CCPA-5	4.14	98.15	2.13	1.96
CCPA-6	6.62	95.17	1.33	1.80
CCPA-7	0.67	97.67	0.34	0.45
CCPA-8	13.81	84.77	1.04	1.80
CCPA-9	4.47	96.82	2.76	2.75
CCPA-10	2.82	95.71	0.65	0.86

C.1.25 COVID (COVID)

Question ID	r	a	p	f_1
COVID-1	0.87	72.88	0.09	0.16

C.1.26 Data Practices (DP)

Question ID	r	a	p	f_1
DP-1	0.68	98.95	0.68	0.68

C.1.27 Enforcement (E)

Question ID	r	a	p	f_1
E-1	41.56	69.30	2.73	4.86
E-2	4.79	96.54	1.47	2.17
E-3	12.90	90.14	1.03	1.79
E-4	54.11	0.56	0.56	1.09
E-5	39.04	0.25	0.25	0.50
E-6	18.10	89.53	3.26	4.39
E-7	1.34	96.77	0.75	0.81
E-8	4.52	85.82	0.26	0.46

C.1.28 GDPR (GDPR)

Question ID	r	a	p	f_1
GDPR-1	14.73	60.31	0.52	0.96
GDPR-2	2.53	97.94	1.36	1.36
GDPR-3	1.88	98.49	1.69	1.55
GDPR-4	4.76	89.32	0.33	0.59
GDPR-5	3.69	95.80	0.31	0.56

C.1.29 Contract (K)

Question ID	r	a	p	f_1
K-1	7.16	96.45	1.31	2.04

C.1.30 Notice (N)

Question ID	r	a	p	f_1
N-1	22.24	86.37	3.18	4.76
N-2	14.95	94.53	4.26	5.88
N-3	4.14	96.60	1.16	1.73
N-4	66.21	30.03	2.38	4.33
N-5	28.83	63.00	1.68	2.95
N-6	28.57	88.11	2.96	5.02
N-7	90.10	2.61	1.95	3.75
N-8	27.92	6.56	0.43	0.82
N-9	50.06	59.04	1.33	2.45
N-10	56.60	21.14	0.67	1.29
N-11	51.35	3.80	0.41	0.79
N-12	16.50	62.15	0.75	1.37
N-13	4.92	90.78	0.50	0.85

C.1.31 Privacy By Design (PBD)

Question ID	r	a	p	f_1
PBD-1	1.34	97.45	0.34	0.54
PBD-2	3.02	94.39	0.26	0.46

C.1.32 Security (SE)

Question ID	r	a	p	f_1
SE-1	2.04	95.83	1.04	1.25
SE-2	3.24	98.49	2.85	2.77
SE-3	88.59	0.98	0.98	1.92
SE-4	27.91	67.06	0.97	1.84
SE-5	16.38	55.57	0.31	0.60
SE-6	1.01	99.62	1.01	1.01
SE-7	12.11	94.21	2.39	3.62
SE-8	21.48	85.79	1.36	2.44

C.1.33 Sharing (SH)

Question ID	r	a	p	f_1
SH-1	4.82	96.45	1.09	1.59
SH-2	4.50	97.14	1.63	2.08
SH-3	8.02	95.04	2.13	3.18
SH-4	0.68	99.59	0.68	0.68
SH-5	12.93	86.74	0.84	1.53
SH-6	4.70	98.71	3.21	3.51
SH-7	7.44	93.86	2.40	3.03

C.1.34 User Control (UC)

Question ID	r	a	p	f_1
UC-1	19.30	80.65	1.04	1.92
UC-2	16.76	75.88	0.94	1.59
UC-3	10.99	92.91	1.76	2.76
UC-4	40.99	68.18	1.39	2.64
UC-5	5.36	94.70	2.14	2.53
UC-6	17.23	83.10	0.69	1.30
UC-7	14.14	81.27	0.58	1.06
UC-8	16.26	86.90	0.99	1.79

C.2 ROBERTA

C.2.1 Performance by Category

Category	r	a	p	f_1
CCPA (CCPA)	17.12	68.77	0.83	1.36
COVID (COVID)	1.30	91.32	0.32	0.50
Contract (K)	15.86	90.92	1.38	2.33
Data Practices (DP)	7.88	50.77	0.16	0.31
Enforcement (E)	42.60	36.84	0.95	1.71
GDPR (GDPR)	20.40	27.81	0.25	0.48
Notice (N)	47.60	33.23	1.14	2.00
Privacy By Design (PBD)	5.15	69.26	0.46	0.47
Security (SE)	37.01	42.46	0.71	1.21
Sharing (SH)	31.92	40.67	0.95	1.60
User Control (UC)	38.31	44.63	1.11	1.75

C.2.2 CCPA (CCPA)

Question ID	r	a	p	f_1
CCPA-1	7.20	89.23	0.87	1.38
CCPA-2	3.66	96.75	1.71	2.06
CCPA-3	20.93	54.38	1.36	2.40
CCPA-4	5.75	92.67	0.39	0.64
CCPA-5	13.57	73.27	0.54	0.94
CCPA-6	29.68	61.16	0.73	1.39
CCPA-7	4.36	81.90	0.35	0.60
CCPA-8	44.52	9.18	0.55	1.06
CCPA-9	29.69	56.67	1.02	1.85
CCPA-10	11.86	72.53	0.73	1.23

C.2.3 COVID (COVID)

Question ID	r	a	p	f_1
COVID-1	1.30	91.32	0.32	0.50

C.2.4 Data Practices (DP)

Question ID	r	a	p	f_1
DP-1	7.88	50.77	0.16	0.31

C.2.5 Enforcement (E)

Question ID	r	a	p	f_1
E-1	52.18	61.99	3.51	6.06
E-2	76.03	1.15	0.73	1.40
E-3	83.56	0.99	0.74	1.43
E-4	54.11	1.26	0.56	1.10
E-5	7.36	87.14	0.41	0.74
E-6	50.45	44.11	1.36	2.41
E-7	1.01	96.83	0.15	0.24
E-8	16.11	1.22	0.15	0.30

C.2.6 GDPR (GDPR)

Question ID	r	a	p	f_1
GDPR-1	39.12	1.81	0.48	0.93
GDPR-2	2.96	87.61	0.25	0.40
GDPR-3	27.88	21.21	0.29	0.57
GDPR-4	18.61	17.14	0.16	0.32
GDPR-5	13.42	11.28	0.08	0.15

C.2.7 Contract (K)

Question ID	r	a	p	f_1
K-1	15.86	90.92	1.38	2.33

C.2.8 Notice (N)

Question ID	r	a	p	f_1
N-1	46.96	48.99	1.57	2.82
N-2	64.54	15.30	1.07	2.04
N-3	20.47	37.93	0.70	1.13
N-4	79.62	18.31	2.44	4.52
N-5	10.83	86.05	2.17	2.93
N-6	68.84	24.89	1.14	2.18
N-7	90.60	1.93	1.93	3.72
N-8	25.94	11.50	0.42	0.80
N-9	87.56	9.94	0.99	1.92
N-10	56.04	22.79	0.67	1.28
N-11	36.49	24.14	0.43	0.82
N-12	2.48	89.71	0.92	1.07
N-13	28.41	40.48	0.39	0.77

C.2.9 Privacy By Design (PBD)

Question ID	r	a	p	f_1
PBD-1	8.39	46.11	0.12	0.24
PBD-2	1.90	92.42	0.81	0.69

C.2.10 Security (SE)

Question ID	r	a	p	f_1
SE-1	30.78	26.10	0.39	0.76
SE-2	3.91	96.60	0.88	1.33
SE-3	88.48	1.55	0.98	1.91
SE-4	66.26	26.34	0.87	1.69
SE-5	28.67	1.53	0.25	0.50
SE-6	2.48	89.38	0.35	0.57
SE-7	72.15	2.26	0.76	1.49
SE-8	3.36	95.90	1.20	1.46

C.2.11 Sharing (SH)

Question ID	r	a	p	f_1
SH-1	28.85	54.32	1.10	1.84
SH-2	7.77	95.28	1.70	2.36
SH-3	30.39	52.01	1.04	1.87
SH-4	22.30	1.15	0.16	0.31
SH-5	57.05	1.83	0.47	0.93
SH-6	15.17	72.32	1.26	2.05
SH-7	61.93	7.80	0.94	1.82

C.2.12 User Control (UC)

Question ID	r	a	p	f_1
UC-1	12.29	87.34	2.29	2.82
UC-2	29.20	60.40	1.01	1.87
UC-3	14.92	85.32	2.46	3.33
UC-4	25.14	61.20	1.09	1.99
UC-5	72.11	0.58	0.58	1.15
UC-6	61.45	17.31	0.49	0.96
UC-7	57.24	4.67	0.43	0.86
UC-8	34.12	40.21	0.51	0.99

C.2.13 CCPA (CCPA)

Question ID	r	a	p	f_1
CCPA-1	7.20	89.23	0.87	1.38
CCPA-2	3.66	96.75	1.71	2.06
CCPA-3	20.93	54.38	1.36	2.40
CCPA-4	5.75	92.67	0.39	0.64
CCPA-5	13.57	73.27	0.54	0.94
CCPA-6	29.68	61.16	0.73	1.39
CCPA-7	4.36	81.90	0.35	0.60
CCPA-8	44.52	9.18	0.55	1.06
CCPA-9	29.69	56.67	1.02	1.85
CCPA-10	11.86	72.53	0.73	1.23

C.2.14 COVID (COVID)

Question ID	r	a	p	f_1
COVID-1	1.30	91.32	0.32	0.50

C.2.15 Data Practices (DP)

Question ID	r	a	p	f_1
DP-1	7.88	50.77	0.16	0.31

C.2.16 Enforcement (E)

Question ID	r	a	p	f_1
E-1	52.18	61.99	3.51	6.06
E-2	76.03	1.15	0.73	1.40
E-3	83.56	0.99	0.74	1.43
E-4	54.11	1.26	0.56	1.10
E-5	7.36	87.14	0.41	0.74
E-6	50.45	44.11	1.36	2.41
E-7	1.01	96.83	0.15	0.24
E-8	16.11	1.22	0.15	0.30

C.2.17 GDPR (GDPR)

Question ID	r	a	p	f_1
GDPR-1	39.12	1.81	0.48	0.93
GDPR-2	2.96	87.61	0.25	0.40
GDPR-3	27.88	21.21	0.29	0.57
GDPR-4	18.61	17.14	0.16	0.32
GDPR-5	13.42	11.28	0.08	0.15

C.2.18 Contract (K)

Question ID	r	a	p	f_1
K-1	15.86	90.92	1.38	2.33

C.2.19 Notice (N)

Question ID	r	a	p	f_1
N-1	46.96	48.99	1.57	2.82
N-2	64.54	15.30	1.07	2.04
N-3	20.47	37.93	0.70	1.13
N-4	79.62	18.31	2.44	4.52
N-5	10.83	86.05	2.17	2.93
N-6	68.84	24.89	1.14	2.18
N-7	90.60	1.93	1.93	3.72
N-8	25.94	11.50	0.42	0.80
N-9	87.56	9.94	0.99	1.92
N-10	56.04	22.79	0.67	1.28
N-11	36.49	24.14	0.43	0.82
N-12	2.48	89.71	0.92	1.07
N-13	28.41	40.48	0.39	0.77

C.2.20 Privacy By Design (PBD)

Question ID	r	a	p	f_1
PBD-1	8.39	46.11	0.12	0.24
PBD-2	1.90	92.42	0.81	0.69

C.2.21 Security (SE)

Question ID	r	a	p	f_1
SE-1	30.78	26.10	0.39	0.76
SE-2	3.91	96.60	0.88	1.33
SE-3	88.48	1.55	0.98	1.91
SE-4	66.26	26.34	0.87	1.69
SE-5	28.67	1.53	0.25	0.50
SE-6	2.48	89.38	0.35	0.57
SE-7	72.15	2.26	0.76	1.49
SE-8	3.36	95.90	1.20	1.46

C.2.22 Sharing (SH)

Question ID	r	a	p	f_1
SH-1	28.85	54.32	1.10	1.84
SH-2	7.77	95.28	1.70	2.36
SH-3	30.39	52.01	1.04	1.87
SH-4	22.30	1.15	0.16	0.31
SH-5	57.05	1.83	0.47	0.93
SH-6	15.17	72.32	1.26	2.05
SH-7	61.93	7.80	0.94	1.82

C.2.23 User Control (UC)

Question ID	r	a	p	f_1
UC-1	12.29	87.34	2.29	2.82
UC-2	29.20	60.40	1.01	1.87
UC-3	14.92	85.32	2.46	3.33
UC-4	25.14	61.20	1.09	1.99
UC-5	72.11	0.58	0.58	1.15
UC-6	61.45	17.31	0.49	0.96
UC-7	57.24	4.67	0.43	0.86
UC-8	34.12	40.21	0.51	0.99

C.2.24 CCPA (CCPA)

Question ID	r	a	p	f_1
CCPA-1	7.20	89.23	0.87	1.38
CCPA-2	3.66	96.75	1.71	2.06
CCPA-3	20.93	54.38	1.36	2.40
CCPA-4	5.75	92.67	0.39	0.64
CCPA-5	13.57	73.27	0.54	0.94
CCPA-6	29.68	61.16	0.73	1.39
CCPA-7	4.36	81.90	0.35	0.60
CCPA-8	44.52	9.18	0.55	1.06
CCPA-9	29.69	56.67	1.02	1.85
CCPA-10	11.86	72.53	0.73	1.23

C.2.25 COVID (COVID)

Question ID	r	a	p	f_1
COVID-1	1.30	91.32	0.32	0.50

C.2.26 Data Practices (DP)

Question ID	r	a	p	f_1
DP-1	7.88	50.77	0.16	0.31

C.2.27 Enforcement (E)

Question ID	r	a	p	f_1
E-1	52.18	61.99	3.51	6.06
E-2	76.03	1.15	0.73	1.40
E-3	83.56	0.99	0.74	1.43
E-4	54.11	1.26	0.56	1.10
E-5	7.36	87.14	0.41	0.74
E-6	50.45	44.11	1.36	2.41
E-7	1.01	96.83	0.15	0.24
E-8	16.11	1.22	0.15	0.30

C.2.28 GDPR (GDPR)

Question ID	r	a	p	f_1
GDPR-1	39.12	1.81	0.48	0.93
GDPR-2	2.96	87.61	0.25	0.40
GDPR-3	27.88	21.21	0.29	0.57
GDPR-4	18.61	17.14	0.16	0.32
GDPR-5	13.42	11.28	0.08	0.15

C.2.29 Contract (K)

Question ID	r	a	p	f_1
K-1	15.86	90.92	1.38	2.33

C.2.30 Notice (N)

Question ID	r	a	p	f_1
N-1	46.96	48.99	1.57	2.82
N-2	64.54	15.30	1.07	2.04
N-3	20.47	37.93	0.70	1.13
N-4	79.62	18.31	2.44	4.52
N-5	10.83	86.05	2.17	2.93
N-6	68.84	24.89	1.14	2.18
N-7	90.60	1.93	1.93	3.72
N-8	25.94	11.50	0.42	0.80
N-9	87.56	9.94	0.99	1.92
N-10	56.04	22.79	0.67	1.28
N-11	36.49	24.14	0.43	0.82
N-12	2.48	89.71	0.92	1.07
N-13	28.41	40.48	0.39	0.77

C.2.31 Privacy By Design (PBD)

Question ID	r	a	p	f_1
PBD-1	8.39	46.11	0.12	0.24
PBD-2	1.90	92.42	0.81	0.69

C.2.32 Security (SE)

Question ID	r	a	p	f_1
SE-1	30.78	26.10	0.39	0.76
SE-2	3.91	96.60	0.88	1.33
SE-3	88.48	1.55	0.98	1.91
SE-4	66.26	26.34	0.87	1.69
SE-5	28.67	1.53	0.25	0.50
SE-6	2.48	89.38	0.35	0.57
SE-7	72.15	2.26	0.76	1.49
SE-8	3.36	95.90	1.20	1.46

C.2.33 Sharing (SH)

Question ID	r	a	p	f_1
SH-1	28.85	54.32	1.10	1.84
SH-2	7.77	95.28	1.70	2.36
SH-3	30.39	52.01	1.04	1.87
SH-4	22.30	1.15	0.16	0.31
SH-5	57.05	1.83	0.47	0.93
SH-6	15.17	72.32	1.26	2.05
SH-7	61.93	7.80	0.94	1.82

C.2.34 User Control (UC)

Question ID	r	a	p	f_1
UC-1	12.29	87.34	2.29	2.82
UC-2	29.20	60.40	1.01	1.87
UC-3	14.92	85.32	2.46	3.33
UC-4	25.14	61.20	1.09	1.99
UC-5	72.11	0.58	0.58	1.15
UC-6	61.45	17.31	0.49	0.96
UC-7	57.24	4.67	0.43	0.86
UC-8	34.12	40.21	0.51	0.99

C.3 LEGALBERT

C.3.1 Performance by Category

Category	r	a	p	f_1
CCPA (CCPA)	24.39	52.87	0.68	1.10
COVID (COVID)	4.03	1.83	0.04	0.09
Contract (K)	46.91	37.83	0.74	1.42
Data Practices (DP)	16.22	0.79	0.13	0.25
Enforcement (E)	37.92	50.50	0.98	1.67
GDPR (GDPR)	4.89	97.02	1.30	1.69
Notice (N)	37.91	41.06	1.20	1.99
Privacy By Design (PBD)	1.68	89.18	0.14	0.24
Security (SE)	30.77	54.70	0.94	1.45
Sharing (SH)	37.81	41.77	0.93	1.61
User Control (UC)	27.80	64.96	1.28	2.10

C.3.2 CCPA (CCPA)

Question ID	r	a	p	f_1
CCPA-1	10.70	72.81	0.38	0.70
CCPA-2	12.75	80.07	0.49	0.93
CCPA-3	28.24	17.43	1.03	1.92
CCPA-4	1.06	99.58	1.41	1.17
CCPA-5	4.27	94.59	0.67	1.01
CCPA-6	49.20	18.13	0.50	0.97
CCPA-7	3.36	95.95	0.68	1.05
CCPA-8	48.32	0.76	0.52	1.01
CCPA-9	53.46	33.27	0.89	1.71
CCPA-10	32.55	16.06	0.26	0.52

C.3.3 COVID (COVID)

Question ID	r	a	p	f_1
COVID-1	4.03	1.83	0.04	0.09

C.3.4 Data Practices (DP)

Question ID	r	a	p	f_1
DP-1	16.22	0.79	0.13	0.25

C.3.5 Enforcement (E)

Question ID	r	a	p	f_1
E-1	71.68	29.93	2.09	3.96
E-2	4.93	97.53	2.30	2.93
E-3	71.35	14.03	0.77	1.47
E-4	21.65	68.99	0.73	1.36
E-5	39.73	0.77	0.26	0.52
E-6	89.86	1.73	1.14	2.21
E-7	0.67	96.15	0.07	0.12
E-8	3.47	94.85	0.49	0.81

C.3.6 GDPR (GDPR)

Question ID	r	a	p	f_1
GDPR-1	10.94	95.94	3.37	3.92
GDPR-2	2.26	98.29	1.00	1.28
GDPR-3	7.14	94.29	1.20	1.92
GDPR-4	3.45	97.55	0.79	1.11
GDPR-5	0.67	99.03	0.13	0.22

C.3.7 Contract (K)

Question ID	r	a	p	f_1
K-1	46.91	37.83	0.74	1.42

C.3.8 Notice (N)

Question ID	r	a	p	f_1
N-1	17.31	83.89	2.14	3.10
N-2	73.28	1.14	0.93	1.80
N-3	3.69	90.39	0.42	0.74
N-4	90.66	3.06	2.24	4.22
N-5	64.36	5.09	1.38	2.59
N-6	82.55	1.17	1.02	1.98
N-7	32.08	70.54	2.25	3.95
N-8	28.86	0.60	0.40	0.77
N-9	3.38	96.32	1.97	2.23
N-10	14.43	84.85	1.18	1.99
N-11	4.84	84.78	0.87	1.03
N-12	24.41	11.48	0.41	0.80
N-13	53.02	0.45	0.35	0.69

C.3.9 Privacy By Design (PBD)

Question ID	r	a	p	f_1
PBD-1	2.35	79.79	0.17	0.30
PBD-2	1.01	98.57	0.11	0.19

C.3.10 Security (SE)

Question ID	r	a	p	f_1
SE-1	1.93	97.70	1.44	1.49
SE-2	27.99	54.67	0.54	1.04
SE-3	87.79	3.07	1.00	1.95
SE-4	83.22	1.05	0.76	1.49
SE-5	29.37	1.77	0.26	0.51
SE-6	2.36	95.75	0.50	0.80
SE-7	7.89	93.89	2.14	2.94
SE-8	5.59	89.71	0.94	1.38

C.3.11 Sharing (SH)

Question ID	r	a	p	f_1
SH-1	68.03	1.00	0.76	1.47
SH-2	72.97	0.89	0.67	1.32
SH-3	15.95	86.43	1.75	2.62
SH-4	15.54	30.35	0.17	0.33
SH-5	56.38	4.65	0.47	0.93
SH-6	5.26	95.62	1.13	1.77
SH-7	30.55	73.44	1.55	2.79

C.3.12 User Control (UC)

Question ID	r	a	p	f_1
UC-1	24.89	82.20	2.26	3.17
UC-2	38.86	29.30	0.70	1.34
UC-3	13.69	90.19	2.45	3.85
UC-4	13.48	92.16	1.87	3.01
UC-5	13.04	91.10	1.32	2.29
UC-6	53.85	40.74	0.59	1.15
UC-7	59.31	0.43	0.43	0.84
UC-8	5.26	93.57	0.64	1.11

C.3.13 CCPA (CCPA)

Question ID	r	a	p	f_1
CCPA-1	10.70	72.81	0.38	0.70
CCPA-2	12.75	80.07	0.49	0.93
CCPA-3	28.24	17.43	1.03	1.92
CCPA-4	1.06	99.58	1.41	1.17
CCPA-5	4.27	94.59	0.67	1.01
CCPA-6	49.20	18.13	0.50	0.97
CCPA-7	3.36	95.95	0.68	1.05
CCPA-8	48.32	0.76	0.52	1.01
CCPA-9	53.46	33.27	0.89	1.71
CCPA-10	32.55	16.06	0.26	0.52

C.3.14 COVID (COVID)

Question ID	r	a	p	f_1
COVID-1	4.03	1.83	0.04	0.09

C.3.15 Data Practices (DP)

Question ID	r	a	p	f_1
DP-1	16.22	0.79	0.13	0.25

C.3.16 Enforcement (E)

Question ID	r	a	p	f_1
E-1	71.68	29.93	2.09	3.96
E-2	4.93	97.53	2.30	2.93
E-3	71.35	14.03	0.77	1.47
E-4	21.65	68.99	0.73	1.36
E-5	39.73	0.77	0.26	0.52
E-6	89.86	1.73	1.14	2.21
E-7	0.67	96.15	0.07	0.12
E-8	3.47	94.85	0.49	0.81

C.3.17 GDPR (GDPR)

Question ID	r	a	p	f_1
GDPR-1	10.94	95.94	3.37	3.92
GDPR-2	2.26	98.29	1.00	1.28
GDPR-3	7.14	94.29	1.20	1.92
GDPR-4	3.45	97.55	0.79	1.11
GDPR-5	0.67	99.03	0.13	0.22

C.3.18 Contract (K)

Question ID	r	a	p	f_1
K-1	46.91	37.83	0.74	1.42

C.3.19 Notice (N)

Question ID	r	a	p	f_1
N-1	17.31	83.89	2.14	3.10
N-2	73.28	1.14	0.93	1.80
N-3	3.69	90.39	0.42	0.74
N-4	90.66	3.06	2.24	4.22
N-5	64.36	5.09	1.38	2.59
N-6	82.55	1.17	1.02	1.98
N-7	32.08	70.54	2.25	3.95
N-8	28.86	0.60	0.40	0.77
N-9	3.38	96.32	1.97	2.23
N-10	14.43	84.85	1.18	1.99
N-11	4.84	84.78	0.87	1.03
N-12	24.41	11.48	0.41	0.80
N-13	53.02	0.45	0.35	0.69

C.3.20 Privacy By Design (PBD)

Question ID	r	a	p	f_1
PBD-1	2.35	79.79	0.17	0.30
PBD-2	1.01	98.57	0.11	0.19

C.3.21 Security (SE)

Question ID	r	a	p	f_1
SE-1	1.93	97.70	1.44	1.49
SE-2	27.99	54.67	0.54	1.04
SE-3	87.79	3.07	1.00	1.95
SE-4	83.22	1.05	0.76	1.49
SE-5	29.37	1.77	0.26	0.51
SE-6	2.36	95.75	0.50	0.80
SE-7	7.89	93.89	2.14	2.94
SE-8	5.59	89.71	0.94	1.38

C.3.22 Sharing (SH)

Question ID	r	a	p	f_1
SH-1	68.03	1.00	0.76	1.47
SH-2	72.97	0.89	0.67	1.32
SH-3	15.95	86.43	1.75	2.62
SH-4	15.54	30.35	0.17	0.33
SH-5	56.38	4.65	0.47	0.93
SH-6	5.26	95.62	1.13	1.77
SH-7	30.55	73.44	1.55	2.79

C.3.23 User Control (UC)

Question ID	r	a	p	f_1
UC-1	24.89	82.20	2.26	3.17
UC-2	38.86	29.30	0.70	1.34
UC-3	13.69	90.19	2.45	3.85
UC-4	13.48	92.16	1.87	3.01
UC-5	13.04	91.10	1.32	2.29
UC-6	53.85	40.74	0.59	1.15
UC-7	59.31	0.43	0.43	0.84
UC-8	5.26	93.57	0.64	1.11

C.3.24 CCPA (CCPA)

Question ID	r	a	p	f_1
CCPA-1	10.70	72.81	0.38	0.70
CCPA-2	12.75	80.07	0.49	0.93
CCPA-3	28.24	17.43	1.03	1.92
CCPA-4	1.06	99.58	1.41	1.17
CCPA-5	4.27	94.59	0.67	1.01
CCPA-6	49.20	18.13	0.50	0.97
CCPA-7	3.36	95.95	0.68	1.05
CCPA-8	48.32	0.76	0.52	1.01
CCPA-9	53.46	33.27	0.89	1.71
CCPA-10	32.55	16.06	0.26	0.52

C.3.25 COVID (COVID)

Question ID	r	a	p	f_1
COVID-1	4.03	1.83	0.04	0.09

C.3.26 Data Practices (DP)

Question ID	r	a	p	f_1
DP-1	16.22	0.79	0.13	0.25

C.3.27 Enforcement (E)

Question ID	r	a	p	f_1
E-1	71.68	29.93	2.09	3.96
E-2	4.93	97.53	2.30	2.93
E-3	71.35	14.03	0.77	1.47
E-4	21.65	68.99	0.73	1.36
E-5	39.73	0.77	0.26	0.52
E-6	89.86	1.73	1.14	2.21
E-7	0.67	96.15	0.07	0.12
E-8	3.47	94.85	0.49	0.81

C.3.28 GDPR (GDPR)

Question ID	r	a	p	f_1
GDPR-1	10.94	95.94	3.37	3.92
GDPR-2	2.26	98.29	1.00	1.28
GDPR-3	7.14	94.29	1.20	1.92
GDPR-4	3.45	97.55	0.79	1.11
GDPR-5	0.67	99.03	0.13	0.22

C.3.29 Contract (K)

Question ID	r	a	p	f_1
K-1	46.91	37.83	0.74	1.42

C.3.30 Notice (N)

Question ID	r	a	p	f_1
N-1	17.31	83.89	2.14	3.10
N-2	73.28	1.14	0.93	1.80
N-3	3.69	90.39	0.42	0.74
N-4	90.66	3.06	2.24	4.22
N-5	64.36	5.09	1.38	2.59
N-6	82.55	1.17	1.02	1.98
N-7	32.08	70.54	2.25	3.95
N-8	28.86	0.60	0.40	0.77
N-9	3.38	96.32	1.97	2.23
N-10	14.43	84.85	1.18	1.99
N-11	4.84	84.78	0.87	1.03
N-12	24.41	11.48	0.41	0.80
N-13	53.02	0.45	0.35	0.69

C.3.31 Privacy By Design (PBD)

Question ID	r	a	p	f_1
PBD-1	2.35	79.79	0.17	0.30
PBD-2	1.01	98.57	0.11	0.19

C.3.32 Security (SE)

Question ID	r	a	p	f_1
SE-1	1.93	97.70	1.44	1.49
SE-2	27.99	54.67	0.54	1.04
SE-3	87.79	3.07	1.00	1.95
SE-4	83.22	1.05	0.76	1.49
SE-5	29.37	1.77	0.26	0.51
SE-6	2.36	95.75	0.50	0.80
SE-7	7.89	93.89	2.14	2.94
SE-8	5.59	89.71	0.94	1.38

C.3.33 Sharing (SH)

Question ID	r	a	p	f_1
SH-1	68.03	1.00	0.76	1.47
SH-2	72.97	0.89	0.67	1.32
SH-3	15.95	86.43	1.75	2.62
SH-4	15.54	30.35	0.17	0.33
SH-5	56.38	4.65	0.47	0.93
SH-6	5.26	95.62	1.13	1.77
SH-7	30.55	73.44	1.55	2.79

C.3.34 User Control (UC)

Question ID	r	a	p	f_1
UC-1	24.89	82.20	2.26	3.17
UC-2	38.86	29.30	0.70	1.34
UC-3	13.69	90.19	2.45	3.85
UC-4	13.48	92.16	1.87	3.01
UC-5	13.04	91.10	1.32	2.29
UC-6	53.85	40.74	0.59	1.15
UC-7	59.31	0.43	0.43	0.84
UC-8	5.26	93.57	0.64	1.11

D Websites in Dataset

Tranco Range	Tranco Rank	URL	IAB 2.0 Top-Level Categories	Document Count	Word Count
1-10	1	google.com	IAB19 (Technology & Computing)	2	12379
1-10	2	facebook.com	IAB14 (Society)	18	39915
1-10	3	youtube.com	IAB25 (Non-Standard Content) IAB1 (Arts & Entertainment)	2	12390
1-10	4	microsoft.com	IAB9 (Hobbies & Interests) IAB19 (Technology & Computing)	4	39988
1-10	6	twitter.com	IAB14 (Society)	8	10495
1-10	7	instagram.com	IAB19 (Technology & Computing) IAB14 (Society)	4	14897
1-10	10	netflix.com	IAB25 (Non-Standard Content) IAB1 (Arts & Entertainment)	3	10139
1-1000	11	linkedin.com	IAB4 (Careers) IAB14 (Society)	5	14495
1-1000	12	qq.com	IAB1 (Arts & Entertainment) IAB12 (News)	4	13692
1-1000	13	apple.com	IAB19 (Technology & Computing) IAB22 (Shopping)	4	9236
1-1000	15	wikipedia.org	IAB5 (Education)	2	14536
1-1000	24	amazon.com	IAB22 (Shopping)	6	9034
1-1000	31	pinterest.com	IAB18 (Style & Fashion) IAB14 (Society)	9	16726
1-1000	33	adobe.com	IAB19 (Technology & Computing) IAB3 (Business)	4	14684
1-1000	44	reddit.com		2	7229
1-1000	71	tumblr.com	IAB25 (Non-Standard Content)	3	17869
1-1000	80	msn.com	IAB25 (Non-Standard Content) IAB12 (News)	2	8635
1-1000	109	myshopify.com	IAB3 (Business) IAB22 (Shopping)	2	10987
1-1000	114	cnn.com	IAB12 (News)	2	13017
1-1000	119	twitch.tv	IAB25 (Non-Standard Content) IAB9 (Hobbies & Interests)	3	12413
1-1000	136	imdb.com	IAB25 (Non-Standard Content) IAB1 (Arts & Entertainment) IAB9 (Hobbies & Interests)	2	5318
1-1000	141	stackoverflow.com	IAB19 (Technology & Computing)	3	10484
1-1000	163	aliexpress.com	IAB22 (Shopping)	2	7345
1-1000	180	washingtonpost.com	IAB11 (Law, Gov't & Politics) IAB12 (News)	2	7565
1-1000	216	chaturbate.com	IAB25 (Non-Standard Content)	3	10693
1-1000	229	amazon.co.uk	IAB22 (Shopping)	3	12583
1-1000	245	researchgate.net	IAB19 (Technology & Computing)	3	11118
1-1000	295	walmart.com	IAB22 (Shopping)	2	7210
1-1000	311	pornhub.com	IAB25 (Non-Standard Content)	7	19749
1-1000	340	ted.com	IAB25 (Non-Standard Content)	2	7215
1-1000	347	livejasmin.com	IAB25 (Non-Standard Content)	3	10441
1-1000	365	okta.com	IAB3 (Business) IAB19 (Technology & Computing)	2	11675
1-1000	369	xvideos.com	IAB25 (Non-Standard Content)	8	15822
1-1000	397	instructure.com	IAB5 (Education)	4	8597
1-1000	464	theverge.com	IAB19 (Technology & Computing)	3	12697
1-1000	469	loc.gov	IAB11 (Law, Gov't & Politics)	2	2898
1-1000	524	craigslist.org	IAB25 (Non-Standard Content) IAB21 (Real Estate) IAB22 (Shopping)	2	2063
1-1000	598	homedepot.com	IAB10 (Home & Garden)	3	12885
1-1000	609	stumbleupon.com	IAB19 (Technology & Computing)	2	11417
1-1000	641	pbs.org	IAB25 (Non-Standard Content)	2	4196
1-1000	666	hulu.com	IAB25 (Non-Standard Content) IAB1 (Arts & Entertainment)	3	14169

1-1000	683	steampowered.com	IAB9 (Hobbies & Interests)	2	4574
1-1000	843	fortune.com	IAB1 (Arts & Entertainment) IAB13 (Personal Finance)	2	10844
1-1000	993	arstechnica.com	IAB1 (Arts & Entertainment)	2	6233
1000-10000	1090	fool.com	IAB13 (Personal Finance)	2	9979
1000-10000	1104	barnesandnoble.com	IAB1 (Arts & Entertainment)	4	13471
1000-10000	1128	corp.ign.com	IAB9 (Hobbies & Interests)	3	21645
1000-10000	1254	thehill.com	IAB12 (News)	2	5021
1000-10000	1337	dictionary.com	IAB9 (Hobbies & Interests)	1	7335
1000-10000	1396	usa.healthcare.siemens.com	IAB15 (Science) IAB7 (Health & Fitness)	3	5799
1000-10000	1621	usa.gov	IAB11 (Law, Gov't & Politics)	1	778
1000-10000	1628	archives.gov	IAB11 (Law, Gov't & Politics)	1	3153
1000-10000	1671	adweek.com	IAB3 (Business)	2	13495
1000-10000	1678	lonelyplanet.com	IAB20 (Travel)	3	5271
1000-10000	1831	wordreference.com	IAB19 (Technology & Computing) IAB5 (Education)	1	572
1000-10000	1889	amd.com	IAB19 (Technology & Computing)	3	8043
1000-10000	1999	allrecipes.com	IAB8 (Food & Drink)	3	16320
1000-10000	2146	slickdeals.net	IAB22 (Shopping)	2	6797
1000-10000	2363	reverbnation.com	IAB1 (Arts & Entertainment)	2	22859
1000-10000	2564	dpreview.com	IAB9 (Hobbies & Interests) IAB19 (Technology & Computing)	6	6980
1000-10000	2672	macrumors.com	IAB19 (Technology & Computing)	1	1272
1000-10000	2976	freep.com	IAB12 (News)	3	7884
1000-10000	3343	everydayhealth.com	IAB15 (Science) IAB7 (Health & Fitness)	3	34103
1000-10000	3474	uh.edu	IAB5 (Education)	1	4190
1000-10000	3510	edmunds.com	IAB2 (Automotive)	2	12616
1000-10000	3962	babycenter.com	IAB15 (Science) IAB6 (Family & Parenting) IAB7 (Health & Fitness)	2	16777
1000-10000	4541	match.com	IAB14 (Society)	4	15183
1000-10000	4683	ubi.com	IAB9 (Hobbies & Interests)	3	12588
1000-10000	4937	geocaching.com	IAB19 (Technology & Computing)	2	12973
1000-10000	5283	sltrib.com	IAB12 (News)	2	9042
1000-10000	5615	wizards.com	IAB9 (Hobbies & Interests)	3	28462
1000-10000	6002	thermofisher.com	IAB15 (Science)	3	9510
1000-10000	6145	newgrounds.com	IAB9 (Hobbies & Interests)	2	5409
1000-10000	6658	allstate.com	IAB13 (Personal Finance)	3	11136
1000-10000	7223	purevolume.com	IAB1 (Arts & Entertainment)	2	4547
1000-10000	7868	simplemachines.org	IAB19 (Technology & Computing)	2	2218
1000-10000	7951	basketball-reference.com	IAB17 (Sports)	2	3787
1000-10000	7973	gamepedia.com	IAB9 (Hobbies & Interests)	2	9944
1000-10000	8064	dailynews.com	IAB12 (News)	3	9532
1000-10000	8146	ebaumsworld.com	IAB1 (Arts & Entertainment)	2	8453
1000-10000	8851	moneysavingexpert.com	IAB13 (Personal Finance)	3	15539
1000-10000	9067	gaiaonline.com	IAB9 (Hobbies & Interests)	2	10800
10000-100000	10053	23andme.com	IAB7 (Health & Fitness) IAB15 (Science)	5	39584
10000-100000	11173	zacks.com	IAB13 (Personal Finance)	2	5880
10000-100000	13372	dailystrength.org	IAB7 (Health & Fitness)	2	6009
10000-100000	14512	valvesoftware.com	IAB10 (Home & Garden)	2	5320
10000-100000	14565	signonsandiego.com	IAB12 (News)	2	14777
10000-100000	15650	soundclick.com	IAB1 (Arts & Entertainment)	2	3699
10000-100000	17000	airliners.net	IAB9 (Hobbies & Interests) IAB19 (Technology & Computing) IAB20 (Travel) IAB25 (Non-Standard Content)	2	6880
10000-100000	18163	videohelp.com	IAB19 (Technology & Computing)	1	1963
10000-100000	20256	filefront.com	IAB9 (Hobbies & Interests)	2	6510
10000-100000	20440	somethingawful.com	IAB1 (Arts & Entertainment)	1	1612

10000-100000	20895	yardbarker.com	IAB17 (Sports)	2	5239
10000-100000	26711	eharmony.com	IAB14 (Society)	2	16146
10000-100000	28046	afterdawn.com	IAB12 (News)	1	847
10000-100000	28890	sci-news.com	IAB15 (Science)	1	569
10000-100000	30950	namepros.com	IAB19 (Technology & Computing)	2	8491
10000-100000	32720	us.mouthshut.com	IAB24 (Uncategorized)	2	7554
10000-100000	34840	drinksmixer.com	IAB8 (Food & Drink)	2	14662
10000-100000	36320	ashleymadison.com	IAB25 (Non-Standard Content) IAB14 (Society)	2	17552
10000-100000	37097	taylorswift.com	IAB19 (Technology & Computing)	2	12151
10000-100000	38498	kraftrecipes.com	IAB8 (Food & Drink)	2	7630
10000-100000	38793	cbsinteractive.com	IAB12 (News)	3	8446
10000-100000	44194	bolt.com	IAB3 (Business)	2	18194
10000-100000	47174	bio-rad.com	IAB3 (Business)	2	10793
10000-100000	64822	tgifridays.com	IAB8 (Food & Drink)	2	9020
10000-100000	74083	twoplustwo.com	IAB9 (Hobbies & Interests)	2	6218
10000-100000	75247	christianmingle.com	IAB14 (Society)	3	36186
10000-100000	76486	primagames.com	IAB9 (Hobbies & Interests)	2	11455
10000-100000	99193	dailyillini.com	IAB12 (News)	2	2275
100000-1000000	100437	listography.com	IAB25 (Non-Standard Content)	2	2925
100000-1000000	110999	friendfinder.com	IAB14 (Society)	2	12337
100000-1000000	116128	chicagomarathon.com	IAB7 (Health & Fitness)	1	3317
100000-1000000	118881	hardwarezone.com	IAB19 (Technology & Computing)	2	7314
100000-1000000	122817	cariboucoffee.com	IAB8 (Food & Drink)	3	10040
100000-1000000	132336	restaurantnews.com	IAB12 (News)	1	836
100000-1000000	175175	cincymuseum.org	IAB1 (Arts & Entertainment)	1	2323
100000-1000000	176865	myriad.com	IAB15 (Science) IAB7 (Health & Fitness)	3	6778
100000-1000000	178874	opendiary.com	IAB14 (Society)	3	6293
100000-1000000	187525	fightingillini.com	IAB12 (News)	2	5768
100000-1000000	201854	aq.com	IAB9 (Hobbies & Interests)	3	10208
100000-1000000	214643	mate1.com	IAB14 (Society)	2	15435
100000-1000000	230902	india-forums.com	IAB14 (Society)	2	8802
100000-1000000	237561	helix.com	IAB15 (Science) IAB7 (Health & Fitness)	2	14241
100000-1000000	239486	abita.com	IAB8 (Food & Drink)	2	2467
100000-1000000	277961	dnacenter.com	IAB15 (Science) IAB6 (Family & Parenting) IAB7 (Health & Fitness)	2	8662
100000-1000000	316184	coffeereview.com	IAB8 (Food & Drink)	2	2137
100000-1000000	344874	communitycoffee.com	IAB8 (Food & Drink)	4	8127
100000-1000000	459935	bgi.com	IAB7 (Health & Fitness)	1	4800
100000-1000000	562225	ambrygen.com	IAB7 (Health & Fitness)	3	6169
100000-1000000	585888	bigtent.com	IAB11 (Law, Gov't & Politics) IAB6 (Family & Parenting)	2	20840
100000-1000000	705374	orig3n.com	IAB15 (Science) IAB7 (Health & Fitness)	2	7119
100000-1000000	707387	enpnetwork.com	IAB4 (Careers)	2	13180
100000-1000000	722795	aboutus.disaboom.com	IAB24 (Uncategorized)	1	458
100000-1000000	814897	nygenome.org	IAB15 (Science) IAB7 (Health & Fitness)	2	1666
100000-1000000	879524	true.com	IAB13 (Personal Finance)	2	8363
100000-1000000	982740	sequencing.com	IAB15 (Science)	3	13097
1000000+	1015718	gays.com	IAB25 (Non-Standard Content)	2	11301
1000000+	1071352	spark.com	IAB14 (Society)	2	27220
1000000+	1201138	greensingles.com	IAB14 (Society)	3	6187
1000000+	1248563	completegenomics.com	IAB5 (Education)	3	6525
1000000+	1363985	veggiedate.org	IAB14 (Society)	2	3974
1000000+	1557614	metrodate.com	IAB14 (Society)	2	2316
1000000+	2364598	my.opera.com	IAB19 (Technology & Computing)	3	5731
1000000+	4599194	wealthymen.com	IAB14 (Society)	2	11487
1000000+	5301232	heremedia.com	IAB14 (Society)	3	12608

1000000+	6127774	epernicus.com	IAB3 (Business)	2	4917
1000000+	10000000	sediabio.com	IAB7 (Health & Fitness)	2	3323
1000000+	10000000	webmediabrands.com	IAB9 (Hobbies & Interests) IAB4 (Careers)	3	15587