

Bibliographic citation:

```
@inproceedings{ravichander-etal-2019-question,
  title = "Question Answering for Privacy Policies: Combining Computational and Legal Perspectives",
  author = "Ravichander, Abhilasha and
    Black, Alan W and
    Wilson, Shomir and
    Norton, Thomas and
    Sadeh, Norman",
  booktitle = "Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)",
  month = nov,
  year = "2019",
  address = "Hong Kong, China",
  publisher = "Association for Computational Linguistics",
  url = "https://www.aclweb.org/anthology/D19-1500",
  doi = "10.18653/v1/D19-1500",
  pages = "4947--4958",
  abstract = "Privacy policies are long and complex documents that are difficult for users to read and understand. Yet, they have legal effects on how user data can be collected, managed and used. Ideally, we would like to empower users to inform themselves about the issues that matter to them, and enable them to selectively explore these issues. We present PrivacyQA, a corpus consisting of 1750 questions about the privacy policies of mobile applications, and over 3500 expert annotations of relevant answers. We observe that a strong neural baseline underperforms human performance by almost 0.3 F1 on PrivacyQA, suggesting considerable room for improvement for future systems. Further, we use this dataset to categorically identify challenges to question answerability, with domain-general implications for any question answering system. The PrivacyQA corpus offers a challenging corpus for question answering, with genuine real world utility.",
}
```

Link to published version: <https://www.aclweb.org/anthology/D19-1500/>

Abstract:

Privacy policies are long and complex documents that are difficult for users to read and understand. Yet, they have legal effects on how user data can be collected, managed

and used. Ideally, we would like to empower users to inform themselves about the issues that matter to them, and enable them to selectively explore these issues. We present PrivacyQA, a corpus consisting of 1750 questions about the privacy policies of mobile applications, and over 3500 expert annotations of relevant answers. We observe that a strong neural baseline underperforms human performance by almost 0.3 F1 on PrivacyQA, suggesting considerable room for improvement for future systems. Further, we use this dataset to categorically identify challenges to question answerability, with domain-general implications for any question answering system. The PrivacyQA corpus offers a challenging corpus for question answering, with genuine real world utility.