# Learning Privacy Expectations by Crowdsourcing Contextual Informational Norms

Yan Shvartzshnaider, Schrasing Tong, Thomas Wies, Paula Klift, Helen Nissenbaum, Lakshminarayanan Subramanian, Prateek Mittal

*To appear in the proceedings of the Fourth AAAI Conference on Human Computation and Crowdsourcing (HCOMP), 2016*

Designing programmable privacy logic frameworks that correspond to social, ethical, and legal norms has been a fundamentally hard problem. Contextual integrity (CI) (Nissenbaum 2010) offers a model for conceptualizing privacy that is able to bridge technical design with ethical, legal, and policy approaches. While CI is capable of capturing the various components of contextual privacy in theory, it is challenging to discover and formally express these norms in operational terms.

In our work we propose a framework for crowdsourcing privacy norms based on the theory of contextual integrity. Our evaluation demonstrates that the combination of crowdsourcing and CI's privacy model allows us to distinguish between socially acceptable and contentious informational norms, on the one hand, and to identify those norms that make little sense, or require further examination, on the other hand.

To facilitate this effort, we introduced several indicators that can be used to produce an estimate of the overall community's privacy expectations towards a set of information flows within a given context. More specifically, the Norm Approval (NA) score allows us to measure the approval score for each norm; the User Approval (UA) score to reflect the number of norms approved by individual users; and the Divergent Score (DS) metric allows us to measure the overall satisfaction of users with regards to the collectively agreed upon privacy norms. These indicators can be used to fine-tune privacy policies on an ongoing basis to ensure that the majority of users actually support the enforced logic. Furthermore, we leverage formal verification techniques to detect inconsistencies in approved sets of norms that can then be eliminated by refining the survey questions.

To evaluate our metrics, we conducted an extensive survey on Amazon Mechanical Turk (AMT) with more than 450 participants and 1400 questions. Our methodology allowed us to discover norms that have majoritarian consensus among users, discard norms that do not make any sense in a given context as well as detect norms that require further examination. Our results leave us optimistic about the feasibility of a full-fledged information system that operates based on the design principles of crowdsourcing, formal verification, and contextual integrity.

Looking ahead, the mechanisms we developed for extracting, expressing, and validating a set of common rules could be integrated into a variety of online social platforms in the future. Discovering a set of common rules through crowdsourcing, e.g. as we have done through AMT, can be viewed as an efficient bootstrapping first move to populate a system at its budding stage. Incorporating these mechanisms into social information systems allows us to

elicit user feedback on an ongoing basis, which then ultimately enables our system to respond to evolving community norms and standards.

## References

Nissenbaum, H. 2010. Privacy in Context: Technology, Policy, and the Integrity of Social Life. Stanford Law Books.