Standardised Privacy Policies: A Post-mortem and Promising Developments

Reuben Binns, University of Southampton, r@reubenbinns.com

Introduction

Since the mid-1990's, frequent attempts have been made to standardise privacy policies, in order to help individuals make informed choices about whom to trust with their data (Egelman, Tsai, Cranor, & Acquisti, 2009). Standardisation has often been seen as a prerequisite for a functioning 'market for privacy', where organisations can compete on their privacy credentials in order to attract privacy-sensitive consumers. Recent examples come from the U.S. Department of Commerce (NTIA), 2013), elements of the proposed EU regulation[1], multiple initiatives from civil society, non-profits and consumer-oriented companies[2], as well as providers of privacy compliance and auditing services[3]. The common idea is that if organisations' disclosures of how they collect and use personal data could be standardised, this information could be aggregated, accessed, compared en mass using automated analysis tools – serving the interests of regulators, consumer advocacy groups, intermediaries and individuals themselves.

As Lorrie Cranor has noted, this idea has a long history (Cranor, 2012) beginning with

[1] See article 13(a) of the Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation).

[2] See for instance, Mozilla Icons project (https://wiki.mozilla.org/Drumbeat/Challenges/Privacy_Icons)

[3] See proposals from TRUSTe (http://www.truste.com/blog/2010/09/14/more-on-the-problem-with-p3p/) , and the Internet Advertising Bureau's CLEAR Ad Notice project (http://www.iab.net/clear).

the Platform for Privacy Preferences (P3P)[4] in the mid-1990's. The standard supported a system where privacy policies, encoded as structured data, could be 'understood' by web browsers which could automatically negotiate with websites according to their user's privacy preferences. So why didn't P3P or any of its many descendants achieve significant and sustained adoption?

A post-mortem

Incentives for adopters

While most organisations know that consumers don't read their policies, they may perceive that the current system isn't broken enough to be worth fixing. It doesn't prevent consumers using their services, and it generally doesn't result in significant penalties from regulators. Adopting a standardised policy might shake both the indifference of consumers and the relative complacency of regulators. Things may be different for those organisations whose privacy practices could hold up in the face of the increased public scrutiny that standardisation might bring. For such organisations, brand reputation and level of consumer trust would likely increase as a result of the standard. In which case, other organisations who would not see such gains have even greater reason to oppose adoption of the standard.

The Lacuna Between Legal, Human and Machine Languages

Translating a detailed, legalistic document into a rigid set of pre-determined criteria carries the usual risks of translation. It is possible that substantive content will be lost, added or changed in meaning. Essential nuances might be lost. If a standardised version of a privacy policy is to be treated with the same legal status as the original, these differences could have material effects on individual rights and organisations' liability. An alternative might be to have two policies, a

[4] http://www.w3.org/P3P/

traditional one written in legally robust, non-standardised legalese, and another standardised but with no legal power; but this would reduce the force of the latter. Even if high fidelity standardisation is possible, it requires a delicate blend of legal and technical skills which may be beyond the effective capabilities of many organisations. Expecting organisations to have comprehensive knowledge about how they and their myriad third parties use personal data may have been unrealistic.

Network effects

Finally, the adoption of any standard is likely to be strongly determined by network effects. In order to get off the ground, a voluntary standardisation effort faces a collective action problem; why should anyone adopt the standard unless a significant number of others have already adopted it? This goes for the organisations whose policies would be standardised, but also for the browser vendors and other intermediaries who would design the tools to parse and analyse the policies, and indeed individuals themselves who must invest time using those tools. The expected payoffs of a standard only materialise once a critical quota from each stakeholder adopt and use the standard. To make matters worse, the network effects problem faced by one standard can be multiplied by the number of competing standards.

Promising developments

Despite the problems noted above, there are several reasons why the idea may continue to be worth pursuing in future.

An Industry in Need of Efficient Compliance Tools

In recent years, many countries around the world have introduced comprehensive data protection and privacy regulations., containing provisions relating to the content of privacy notices. The European Union may soon adopt a new General Data Protection Regulation, with increased liability for organisations who flout the rules. An increasing number of organisations now feel the need for cheap and efficient means of compliance. At the same time, a new breed of legal services are arising which make greater use of automated tools and semantically structured information, for instance in services which generate legal agreements automatically through the use of templates.[5] Many of these services offer to generate privacy policies in this way, drawing from a standard set of options which the client can select as appropriate. If legal documents are written by default as structured and machine-readable data, then the machine-readable privacy policies may be a natural by-product.

Crowdsourcing

Terms of Service; Didn't Read (ToS;DR) is an initiative to crowdsource ratings of terms of service and privacy policies.[6] The project tracks several hundred of the most popular websites, and is gradually compiling a set of ratings. It uses a web scraper to keep track of changes, which are then rated by the community. Individual clauses are flagged, discussed, and rated as either good, bad or neutral. An overall rating for each website, and a summary of the most important clauses in its policies are provided to users to help them decide whether or not to use the site. Unlike many other standardisation efforts, this approach requires no effort (or permission) from the organisations whose practices are rated. This means it avoids at least part of the network effects problems and provides a degree of independence and trustworthiness over the ratings.

Automated analysis

[5] See, for instance, Docracy.com, NoLo.com, StandardLegal, Iubenda.org

[6] See www.tosdr.org.

There are some privacy-related practices which can be automatically discovered by the browser, such as types of first and third party cookies a website places, whether they use SSL, and other technical aspects of web privacy.[7] Some research projects are also developing means of reverse-engineering web tracking technologies[8]. These aspects can be surfaced to the user through the browser in ways which may help them make privacy choices.

Machine learning

Manual parsing of privacy policies is resource intensive and may not scale well. One way to extend the value of the such ratings from initiatives like ToS;DR is through the use of machine learning. This involves using data about human judgements of policies to train a classifier to recognise those features in new policies. The first step is categorisation of clauses in policies; i.e. the ability to distinguish which sentences relate to particular topics, which has been successfully demonstrated (Ammar et al 2012). The second step is automatically recognising features of particular clauses that would be rated positively or negatively by human raters, which has yet to be explored. Combined, these two techniques could allow the crowd-sourced approach to scale.

References:

Ammar, W., Wilson, S., Sadeh, N., & Smith, N. (2012). Automatic Categorization of Privacy Policies: A Pilot Study.

Cranor, L. F. (2012). NECESSARY BUT NOT SUFFICIENT : STANDARDIZED MECHANISMS FOR PRIVACY NOTICE AND CHOICE, 273–307.

Egelman, S., Tsai, J., Cranor, L. F., & Acquisti, A. (2009). Timing Is Everything ? The Effects of Timing and Placement of Online Privacy Indicators.

National Telecommunications and Information Administration (NTIA). (2013). SHORT FORM NOTICE CODE OF CONDUCT TO PROMOTE TRANSPARENCY.

[7] See for instance Ghostery.com, or the Electronic Frontier Foundation's PrivacyBadger eff.org/privacybadger

[8] See https://freedom-to-tinker.com/blog/randomwalker/web-measurement-for-fairness-and-transparency/