

Privacy technologies: An annotated syllabus

Arvind Narayanan

Princeton University

Abstract. The teaching of graduate courses is symbiotic with research — it helps systematize knowledge, often creates new knowledge, and influences the thinking of the next generation of scholars, indirectly shaping the future of the discipline. I therefore believe that this forum is well-suited to discussing the teaching of privacy technologies. As a starting point, I offer my annotated syllabus and other reflections on the graduate seminar I taught at Princeton during Fall 2012.

In developing the curriculum, I aimed to provide technical depth while integrating perspectives from economics, law, policy and other schools of thought on privacy. These worldviews, in my opinion, sometimes conflict with and often add much-needed nuance to the narrative on privacy technologies within computer science.

As a first-year faculty member, I'm keenly aware that my experience as an educator is relatively limited; my intent here is to provoke discussion rather than to be authoritative.

1 Introduction

During Fall 2012 I taught a graduate seminar course on privacy technologies at Princeton.¹ My conception of the term “privacy technologies” includes both privacy-enhancing and privacy-infringing technologies. Since this was a new course, I had freedom over both the content and the format.

While this was decidedly a computer science course, I aimed to integrate and reconcile the computer science literature the often divergent views on privacy and privacy technologies in the fields of economics, law, policy, philosophy, etc. This mirrors my view that privacy research as a whole stands to benefit from greater interdisciplinarity.

In terms of the format, my main departure from a traditional graduate seminar was to include an online Wiki discussion component, and in fact to make this the centerpiece, and to require students to participate in the online discussion of the day's readings before coming to class. The in-class discussion would use the Wiki discussion as a starting point.

The advantages of this approach are: 1. it gives the instructor a great degree of control in shaping the discussion of each paper, 2. the instructor can more closely monitor individual students' progress 3. class discussion can focus on particularly tricky and/or contentious points, instead of rehashing the obvious.

It was a small class of eleven students, of which ten completed it. Rigorously evaluating the effectiveness of teaching is hard (student evaluations are of limited use without baselines or controls), so I do not report any scientific results here. I only offer subjective thoughts.

Section 2 is about refuting privacy myths. Section 3, the bulk of this document, is an annotated syllabus. I have made available the entire set of Wiki discussion prompts for the class.² I consider this integral to the syllabus. Section 4 presents the broad thematic questions that I had in mind and my concluding thoughts are in 5.

¹ <http://randomwalker.info/teaching/fall-2012-privacy-technologies/>

² <http://randomwalker.info/teaching/fall-2012-privacy-technologies/discussion-prompts.html>
<http://randomwalker.info/teaching/fall-2012-privacy-technologies/discussion-prompts.pdf>

2 Refuting privacy myths via “expectation failure”

In this section I will discuss some major misconceptions about privacy, how to refute them, and why it is important to do this right at the beginning of the course.

Privacy’s primary pitfalls

Instructors are often confronted with breaking down faulty mental models that students bring into class before actual learning can happen. This is especially true of the topic at hand. Luckily, misconceptions about privacy are so pervasive in the media and among the general public that it wasn’t too hard to identify the most common ones before the start of the course. And it didn’t take much class discussion to confirm that my students weren’t somehow exempt from these beliefs.

One cluster of myths is about the supposed lack of importance of privacy. 1. “There is no privacy in the digital age.” This is the most common and perhaps the most grotesquely fallacious of the misconceptions; more on this below. 2. “No one cares about privacy any more” (variant: young people don’t care about privacy.) 3. “If you haven’t done anything wrong you have nothing to hide.”

A second cluster of fallacious beliefs is very common among computer scientists and comes from the tendency to reduce everything to a black-and-white technical problem. In this view, privacy maps directly to access control and cryptography is the main technical mechanism for achieving privacy. It’s a view in which the world is full of adversaries and there is no room for obscurity or nontechnical ways of improving privacy.³

The first step in learning is to unlearn

Why is it important to spend time confronting faulty mental models? Why not simply teach the “right” ones? In my case, there was a particularly acute reason — to the extent that students believe that privacy is dead and that learning about privacy technologies is unimportant, they are not going to be invested in the class, which would be really bad. But even in the case of misconceptions that don’t lead to students doubting the fundamental premise of the class, there is a surprising reason why unlearning is important.

A famous experiment in the ’80s demonstrated what we now know about the ineffectiveness of the “information transmission” model of teaching [5]. The researchers interviewed students after any of four introductory physics courses, and determined that they hadn’t actually learned what had been taught, such as Newton’s laws of motion; instead they just learned to pass the tests. When the researchers sat down with students to find out why, here’s what they found:

What they heard astonished them: many of the students still refused to give up their mistaken ideas about motion. Instead, they argued that the experiment they had just witnessed did not exactly apply to the law of motion in question; it was a special case, or it didn’t quite fit the mistaken theory or law that they held as true.

A special case! Ha. What’s going on here? Well, learning new facts is easy. On the other hand, updating mental models is so cognitively expensive that we go to absurd lengths to avoid doing so. The societal-scale analog of this extreme reluctance is well-illustrated by the history of science — we patched the Ptolemaic model of the Universe, with the Earth at the center, for over a millennium before we were forced to accept that the Copernican system fit observations better.

The instructor’s arsenal

³ For refutations of these myths, see [1], [2], [3], [4].

The good news is that the instructor can utilize many effective strategies that fall under the umbrella of active learning. Ken Bain's excellent book (which the preceding text describing the experiment is from) lays out a pattern in which the instructor creates an expectation failure, a situation in which existing mental models of reality will lead to faulty expectations. One of the prerequisites for this to work, according to the book, is to get students to care.

Bain argues that expectation failure, done right, can be so powerful that students might need emotional support to cope. Fortunately, this wasn't necessary in my class, but I have no doubt of it based on my personal experiences. For instance, back when I was in high school, learning how the Internet actually worked and realizing that my intuitions about the network had to be discarded entirely was such a disturbing experience that I remember my feelings to this day.

Let's look at an example of expectation failure in my privacy class. To refute the "privacy is dying" myth, I found it useful to talk about *Fifty Shades of Grey* — specifically, why it succeeded even though publishers initially passed on it. One answer seems to be that since it was first self-published as an e-book, it allowed readers to be discreet and avoid the stigma associated with the genre. (But following its runaway success in that form, the stigma disappeared, and it was released in paper form and flew off the shelves.)

The relative privacy of e-books from prying strangers is one of the many ways in which digital technology affords more privacy for specific activities. Confronting students with an observed phenomenon whose explanation involves a fact that seems starkly contrary to the popular narrative creates an expectation failure. Telling personal stories about how technology has either improved or eroded privacy, and eliciting such stories from students, gets them to care. Once this has been accomplished, it's productive to get into a nuanced discussion of how to reconcile the two views with each other, different meanings of privacy (e.g., tracking of reading habits), how the Internet has affected each, and how society is adjusting to the changing technological landscape.

3 Privacy technologies: An annotated syllabus

What should be taught in a class on privacy technologies? Before we answer that, let's take a step back and ask, **how does one go about figuring out what should be taught in any class?**

I've seen two approaches. The traditional, default, overwhelmingly common approach is to think of it in terms of "covering content" without much consideration to what students are getting out of it. The content that's deemed relevant is often determined by what the fashionable research areas happen to be, or historical accident, or some combination thereof.

A contrasting approach, promoted by authors like Bain, applies a laser focus on skills that students will acquire and how they will apply them later in life. On teaching orientation day at Princeton, our instructor, who clearly subscribed to this approach, had each professor describe what students would do in the class they are teaching, then wrote down only the verbs from these descriptions. The point was that our thinking had to be centered around skills that students would take home.

I prefer a middle ground. It should be apparent from my description of the traditional approach above that I'm not a fan. On the other hand, I have to wonder what skills our teaching coach would have suggested for a course on cosmology — avoiding falling into black holes? Alright, I'm exaggerating to make a point. The verbs in question are words like "synthesize" and "evaluate," so there would be no particular difficulty in applying them to cosmology. But my point is that in a cosmology course, I'm not sure the instructor should start from these verbs.

Sometimes we want students to be exposed to knowledge primarily because it is beautiful, and being able to perceive that beauty inspires us, instills us with a love of further learning, and I dare say satisfies a fundamental need. To me a lot of the crypto “magic” that goes into privacy technologies falls into that category (not that it doesn’t have practical applications).

With that caveat, however, I agree with the emphasis on skills and life impact. I thought of my students primarily as developers of privacy technologies (and more generally, of technological systems that incorporate privacy considerations), but also as users and scholars of privacy technologies.

I organized the course into sections, a short introductory section followed by five sections that alternated in the level of math/technical depth. Every time we studied a technology, we also discussed its social/economic/political aspects. I had a great deal of discretion in guiding where the conversation around the papers went by giving them questions/prompts on the class Wiki. Let us now jump in. The italicized text is from the course page, the rest is my annotation.

3.0 *Intro*

Goals of this section: Why are we here? Who cares about privacy? What might the future look like?

- *Dan Solove. Why Privacy Matters Even if You Have ‘Nothing to Hide’ (Chronicle) [1]*
- *David Brin. The Transparent Society (WIRED, circa 1996 [6], later expanded into a book [7])*

In addition to helping flesh out the foundational assumptions of this course that I discussed in the previous section, pairing these opposing views with each other helped make the point that there are few absolutes in this class, that privacy scholars may disagree with each other, and that the instructor doesn’t necessarily agree with the viewpoints in the assigned reading, much less expects students to.

3.1 *Cryptography: power and limitations*

Goals. Travel back in time to the 80s and early 90s, understand the often-euphoric vision that many crypto pioneers and hobbyists had for the impact it would have. Understand how cryptographic building blocks were thought to be able to support this restructuring of society. Reason about why it didn’t happen.

Understand the motivations and mathematical underpinnings of the modern research on privacy-preserving computations. Experiment with various encryption tools, discover usability problems and other limitations of crypto.

- *David Chaum. Security without Identification: Card Computers to make Big Brother Obsolete (1985)[8]*
- *Steven Levy. Crypto Rebels (WIRED, 1993 [9]; later a 2001 book [10])*
- *Eric Hughes. A cypherpunk’s manifesto. (short essay, 1993.)[11]*

I think the Chaum paper is a phenomenal and underutilized resource for teaching. My goal was to really immerse students in an alternate reality where the legal underpinnings of commerce were replaced by cryptography, much as Chaum envisioned (and even going beyond that). I created a couple of e-commerce scenarios for Wiki discussion and had them reason about how various functions would be accomplished.

My own views on this topic are set forth in the paper (and talk) “What Happened to the Crypto Dream?” [12]. In general I aimed to shield students from my viewpoints, and saw my role as helping them discover (and be able to defend) their own. At least in this instance I succeeded. Some students took the position that the cypherpunk dream is just around the corner.

- *The ‘Garbled Circuit Protocol’ (Yao’s theorem on secure two-party computation) and its implications (lecture)*

This is one of the topics that sadly suffers from a lack of good expository material, so I instead lectured on it.

- *Alma Whitten and Doug Tygar. Why Johnny Can’t Encrypt: A Usability Evaluation of PGP 5.0* [13]
- *Nikita Borisov, Ian Goldberg, Eric Brewer. Off-the-Record Communication, or, Why Not To Use PGP* [14]
- *Thomas Ptacek. Javascript Cryptography Considered Harmful* [15]

One of the exercises here was to install and use various crypto tools and rediscover the usability problems. The difficulties were even worse than I’d anticipated.

3.2 Data collection and data mining, economics of personal data, behavioral economics of privacy

Goals. Jump forward in time to the present day and immerse ourselves in the world of ubiquitous data collection and surveillance. Discover what kinds of data collection and data mining are going on, and why. Discuss how and why the conversation has shifted from Government surveillance to data collection by private companies in the last 20 years.

Theme: first-party data collection.

- *New York Times. How Companies Learn Your Secrets* [16]
- *Andrew Odlyzko. Privacy, Economics, and Price Discrimination on the Internet* [17]

Theme: third-party data collection.

- *Julia Angwin. The Web’s New Gold Mine: Your Secrets (First in the Wall Street Journal’s What They Know series)* [18]
- *Jonathan R. Mayer and John C. Mitchell. Third-Party Web Tracking: Policy and Technology* [19]

Theme: why companies act the way they do.

- *Joseph Bonneau and Sren Preibusch. The Privacy Jungle: On the Market for Data Protection in Social Networks* [20]
- *Bruce Schneier. How Security Companies Sucker Us With Lemons (WIRED)* [21]

Theme: why people act the way they do.

- *Alessandro Acquisti and Jens Grossklags. What Can Behavioral Economics Teach Us About Privacy?* [22]
- *Alessandro Acquisti. Privacy in Electronic Commerce and the Economics of Immediate Gratification* [23]

This section is rather self-explanatory. After the math-y flavor of the first section, this one has a good amount of economics, behavioral economics, and policy. One of the thought exercises was to project current trends into the future and imagine what ubiquitous tracking might lead to in five or ten years.

3.3 Anonymity and De-anonymization

Important note: communications anonymity (e.g., Tor) and data anonymity/de-anonymization (e.g., identifying people in digital databases) are technically very different, but we will discuss them together because they raise some of the same ethical questions. Also, Bitcoin lies somewhere in between the two.

- *Roger Dingledine, Nick Mathewson, Paul Syverson. Tor: The Second-Generation Onion Router* [24]
- *Satoshi Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System* [25]

Tor and Bitcoin (especially the latter) were the hardest but also the most rewarding parts of the class, both for them and for me. Together they took up 4 classes. Bitcoin is extremely challenging to teach because it is technically intricate, the ecosystem is rapidly changing, and a lot of the information is in random blog/forum posts.

In a way, I was betting on Bitcoin by deciding to teach it — if it had died with a whimper, their knowledge of it would be much less relevant. In general I think instructors should choose to make such bets more often; most curricula are very conservative. I'm glad I did.

- *Nils Homer et al. Resolving Individuals Contributing Trace Amounts of DNA to Highly Complex Mixtures Using High-Density SNP Genotyping Microarrays* [26]
- [Optional] *Arvind Narayanan, Elaine Shi, Benjamin I. P. Rubinstein. Link Prediction by De-anonymization: How We Won the Kaggle Social Network Challenge* [27]

It was a challenge to figure out which deanonymization paper to assign. I went with the DNA one because I wanted them to see that deanonymization isn't a fact about data, but a fact about the world. Another thing I liked about this paper is that they'd have to extract the not-too-complex statistical methodology in this paper from the bioinformatics discussion in which it is embedded. This didn't go as well as I'd hoped.

I've co-authored a few deanonymization papers, but they're not very well written and/or are poorly suited for pedagogical purposes. The Kaggle paper is one exception, which I made optional.

- *Paul Ohm. Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization* [28]
- [Optional] *Jane Yakowitz Bambauer. Tragedy of the Data Commons* [29]

This is another pair of papers with opposing views. Since the latter paper is optional, knowing that most of them wouldn't have read it, I used the Wiki prompts to raise many of the issues that the author raises.

3.4 *Lightweight Privacy Technologies and New Approaches to Information Privacy*

While cryptography is the mechanism of choice for cypherpunk privacy and anonymity tools like Tor, it is too heavy a weapon in other contexts like social networking. In the latter context, it's not so much users deploying privacy tools to protect themselves against all-powerful adversaries but rather a service provider attempting to cater to a more nuanced understanding of privacy that users bring to the system. The goal of this section is to consider a diverse spectrum of ideas applicable to this latter scenario that have been proposed in recent years in the fields of CS, HCI, law, and more. The technologies here are "lightweight" in comparison to cryptographic tools like Tor.

- *Scott Lederer, Jason Hong et al. Personal Privacy through Understanding and Action: Five Pitfalls for Designers* [30]
- *Franziska Roesner et al. User-Driven Access Control: Rethinking Permission Granting in Modern Operating Systems* [31]
- *Fred Stutzman and Woodrow Hartzog. Obscurity by Design: An Approach to Building Privacy into Social Media* [32]
- *Woodrow Hartzog and Fred Stutzman. The Case for Online Obscurity* [2]
- *Jerry Kang et al. Self-surveillance Privacy* [33]
- [Optional] *Ryan Calo. Against Notice Skepticism In Privacy (And Elsewhere)* [34]
- *Helen Nissenbaum. A Contextual Approach to Privacy Online* [35]

3.5 *Purely technological approaches revisited*

This final section doesn't have a coherent theme (and I admitted as much in class). My goal with the first two papers was to contrast a privacy problem which seems amenable to a purely or primarily technological formulation and solution (statistical queries over databases of sensitive personal information) with one where such attempts have been less successful (the decentralized, own-your-data approach to social networking and e-commerce).

- *Differential Privacy. (Lecture)*
 - *Cynthia Dwork. Differential Privacy.* [36]

Differential privacy is another topic that is sorely lacking in expository material, especially from the point of view of students who've never done crypto before. So this was again a lecture.

- *Arvind Narayanan et al. A Critical Look at Decentralized Personal Data Architectures* [37]
- *John Perry Barlow A Declaration of the Independence of Cyberspace (short essay, 1996)* [38]
- *James Grimmelmann. Sealand, HavenCo, and the Rule of Law* [39]

These two essays aren't directly related to privacy. One of the recurring threads in this course is the debate between purely technological and legal or other approaches to privacy; the theme here is to generalize it to a context broader than privacy. The Barlow essay asserts the exceptionalism of Cyberspace as an unregulable medium, whereas the Grimmelmann paper provides a much more nuanced view of the relationship between the law and new technological frontiers.

I have made available the entire set of Wiki discussion prompts for the class.⁴ I consider this integral to the syllabus, for it shapes the discussion very significantly. I really hope other instructors and students find this useful as a teaching/study guide. For reference, each set of prompts (one set per class) took me about three hours to write on average.

4 Themes

Here are the "big questions" and themes of the course, including those I discussed as part of the course contents:

- Who cares about privacy? Does privacy matter even if we have "nothing to hide?"
- How do values about privacy change over time and across cultures?
- What would a world of perpetual observation look like, technologically and socially?
- What's the difference, and the relationship, between security and privacy?
- Are privacy technologies about cryptography and access control, or is there more to them?
- What might an alternative world with ubiquitous crypto have looked like? Why are we not in that world?
- How does anonymity relate to privacy? How does crypto enable anonymity? What factors make it easier or harder to achieve?
- How can we classify privacy technologies in terms of the underlying techniques?
- How can we classify privacy technologies in terms of who they are intended for, and the incentives for usage?
- How can we go beyond the public/private dichotomy in privacy?
- What framework should we use for thinking about privacy technologies that have both good and bad uses?

⁴ <http://randomwalker.info/teaching/fall-2012-privacy-technologies/discussion-prompts.html>
<http://randomwalker.info/teaching/fall-2012-privacy-technologies/discussion-prompts.pdf>

- How much can be inferred or predicted about people’s actions, behavior and preferences via machine learning? Are there limits to these inferences and predictions?
- What are the economic incentives that make companies act the way they do?
- What are the biases and heuristics that make people act the way they do?
- How does better design and HCI contribute to improved privacy?
- What does the story of privacy technologies tell us about the relationship between technology and society?
- How do we find a balance, in various contexts, between regulating by technology, by law, by social norms, and by other means? What are the strengths of each approach?
- Should academics who study privacy have a normative stance on it? Is it acceptable for this stance to be reflected in scholarly work?
- What are some incentives that academics from different disciplines have that affect the type of privacy scholarship that they do and the type of conclusions they reach?

5 Discussion

Student feedback was predominantly positive. Essentially the only change that students wished for at the end of the course was more technical material. This was a surprise to me, since the evidence seemed to suggest that students were having difficulty with the more technical aspects — in the Wiki discussions that incorporated a mix of technical and less technical questions, the latter generally had more participation. One model of student behavior that might explain this is that they follow a work-minimization strategy during the course, but looking back at the end, wish it were structured so that they were forced to do more technical work.

Most of the students had only a passing familiarity with cryptography. I was pleasantly surprised by how far it is possible to go in teaching privacy technologies while building up the crypto background as necessary. While cryptography is obviously a key building block of many privacy technologies, perhaps it need not be a pedagogical prerequisite.

For several topics — Yao’s garbled circuit protocol, Bitcoin, and differential privacy — I did not find good expository material online. Developing these would be valuable to the community, in my opinion. Perhaps this could even be a part of the activity of the class. Partly motivated by my teaching experience, there is now a nascent effort in the security group at Princeton to write a survey paper on Bitcoin.

Finally, I’m happy to report that one of the class projects has led to novel research findings [40].

Acknowledgement. I’m very grateful to Vitaly Shmatikov for feedback on the syllabus. Joseph Lorenzo Hall’s privacy syllabus for his course at NYU was also very useful.⁵

A special thanks to the person on Elance who turned my blog posts on this topic into a LaTeX document (which became this paper), including hunting down the references. Seriously, outsourcing is amazing, you should try it.

References

1. D. J. Solove, “Why privacy matters even if you have ‘nothing to hide’,” *Chronicle of Higher Education*, vol. 15, 2011.
2. W. Hartzog and F. Stutzman, “The case for online obscurity.”
3. A. Narayanan, “Adversarial thinking considered harmful (sometimes).”
4. —, “The many ways in which the internet has given us more privacy.”

⁵ http://josephhall.org/papers/NYU-MCC-1303-S2012_privacy_syllabus.pdf

5. K. Bain, *What the Best College Teachers Do*. Harvard University Press, 2004.
6. D. Brin, "The transparent society. Wired," 1996.
7. —, *The transparent society: Will technology force us to choose between privacy and freedom?* Basic Books, 1999.
8. D. Chaum, "Security without identification: Transaction systems to make big brother obsolete," *Communications of the ACM*, vol. 28, no. 10, pp. 1030–1044, 1985.
9. S. Levy, "Crypto rebels. Wired," 1996.
10. —, *Crypto: How the Code Rebels Beat the Government—Saving Privacy in the Digital Age*. Penguin Books, 2001.
11. E. Hughes, "A cypherpunk's manifesto," 1993.
12. A. Narayanan, "What happened to the crypto dream?" *Security & Privacy, IEEE*, vol. 11, no. 2, pp. 75–76, 2013.
13. A. Whitten and J. D. Tygar, "Why johnny can't encrypt: A usability evaluation of PGP 5.0," in *Proceedings of the 8th USENIX Security Symposium*, vol. 99. McGraw-Hill, 1999.
14. N. Borisov, I. Goldberg, and E. Brewer, "Off-the-record communication, or, why not to use PGP," in *Proceedings of the 2004 ACM workshop on Privacy in the electronic society*. ACM, 2004, pp. 77–84.
15. T. Ptacek, "Javascript cryptography considered harmful," <http://www.matasano.com/articles/javascript-cryptography/>.
16. C. Duhigg, "How companies learn your secrets," *The New York Times*, vol. 2, p. 16, 2012.
17. A. Odlyzko, "Privacy, economics, and price discrimination on the internet," in *Proceedings of the 5th international conference on Electronic commerce*. ACM, 2003, pp. 355–366.
18. J. Angwin, "The web's new gold mine: Your secrets," *Wall Street Journal*, vol. 30, no. 07, p. 2010, 2010.
19. J. R. Mayer and J. C. Mitchell, "Third-party web tracking: Policy and technology," in *Security and Privacy (SP), 2012 IEEE Symposium on*. IEEE, 2012, pp. 413–427.
20. J. Bonneau and S. Preibusch, "The privacy jungle: On the market for data protection in social networks," in *Economics of information security and privacy*. Springer, 2010, pp. 121–167.
21. B. Schneier, "How security companies sucker us with lemons. Wired," 2007.
22. A. Acquisti and J. Grossklags, "What can behavioral economics teach us about privacy?" *DIGITAL PRIVACY*, p. 329, 2007.
23. A. Acquisti, "Privacy in electronic commerce and the economics of immediate gratification," in *Proceedings of the 5th ACM conference on Electronic commerce*. ACM, 2004, pp. 21–29.
24. R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The second-generation onion router," DTIC Document, Tech. Rep., 2004.
25. S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," *Consulted*, vol. 1, p. 2012, 2008.
26. N. Homer, S. Szlinger, M. Redman, D. Duggan, W. Tembe, J. Muehling, J. V. Pearson, D. A. Stephan, S. F. Nelson, and D. W. Craig, "Resolving individuals contributing trace amounts of DNA to highly complex mixtures using high-density SNP genotyping microarrays," *PLoS genetics*, vol. 4, no. 8, p. e1000167, 2008.
27. A. Narayanan, E. Shi, and B. I. Rubinstein, "Link prediction by de-anonymization: How we won the kaggle social network challenge," in *Neural Networks (IJCNN), The 2011 International Joint Conference on*. IEEE, 2011, pp. 1825–1834.
28. P. Ohm, "Broken promises of privacy: Responding to the surprising failure of anonymization," *UCLA Law Review*, vol. 57, p. 1701, 2010.
29. J. Yakowitz, "Tragedy of the data commons," 2011.
30. S. Lederer, I. Hong, K. Dey, and A. Landay, "Personal privacy through understanding and action: five pitfalls for designers," *Personal and Ubiquitous Computing*, vol. 8, no. 6, pp. 440–454, 2004.
31. F. Roesner, T. Kohno, A. Moshchuk, B. Parno, H. J. Wang, and C. Cowan, "User-driven access control: Rethinking permission granting in modern operating systems," in *Security and Privacy (SP), 2012 IEEE Symposium on*. IEEE, 2012, pp. 224–238.
32. F. Stutzman and W. Hartzog, "Obscurity by design: An approach to building privacy into social media."
33. J. Kang, K. Shilton, D. Estrin, and J. Burke, "Self-surveillance privacy," *Iowa L. Rev.*, vol. 97, p. 809, 2011.
34. R. Calo, "Against notice skepticism in privacy (and elsewhere)," *Notre Dame Law Review*, vol. 87, 2012.
35. H. Nissenbaum, "A contextual approach to privacy online," *Daedalus*, vol. 140, no. 4, pp. 32–48, 2011.
36. C. Dwork, "Differential privacy," in *Automata, languages and programming*. Springer, 2006, pp. 1–12.
37. A. Narayanan, V. Toubiana, S. Barocas, H. Nissenbaum, and D. Boneh, "A critical look at decentralized personal data architectures," *arXiv preprint arXiv:1202.4503*, 2012.
38. J. P. Barlow, "A declaration of the independence of cyberspace," 1996.
39. J. Grimmelmann, "Sealand, havenco, and the rule of law," *University of Illinois Law Review*, p. 405, 2012.
40. C. Eubank, M. Melara, D. Perez-Botero, and A. Narayanan, "Shining the floodlights on mobile web tracking — a privacy survey," in *Web 2.0 Security and Privacy Workshop*, 2013.