

Exploring Design and Governance Challenges in the Development of Privacy-Preserving Computation

Nitin Agrawal
nitin.agrawal@cs.ox.ac.uk
University of Oxford
Oxford, UK

Reuben Binns
reuben.binns@cs.ox.ac.uk
University of Oxford
Oxford, UK

Max Van Kleek
max.van.kleek@cs.ox.ac.uk
University of Oxford
Oxford, UK

Kim Laine
kim.laine@microsoft.com
Microsoft Research
Redmond, CA, USA

Nigel Shadbolt
nigel.shadbolt@cs.ox.ac.uk
University of Oxford
Oxford, UK

ABSTRACT

Homomorphic encryption, secure multi-party computation, and differential privacy are part of an emerging class of Privacy Enhancing Technologies which share a common promise: to preserve privacy whilst also obtaining the benefits of computational analysis. Due to their relative novelty, complexity, and opacity, these technologies provoke a variety of novel questions for design and governance. We interviewed researchers, developers, industry leaders, policymakers, and designers involved in their deployment to explore motivations, expectations, perceived opportunities and barriers to adoption. This provided insight into several pertinent challenges facing the adoption of these technologies, including: how they might make a nebulous concept like privacy computationally tractable; how to make them more usable by developers; and how they could be explained and made accountable to stakeholders and wider society. We conclude with implications for the development, deployment, and responsible governance of these privacy-preserving computation techniques.

CCS CONCEPTS

• **Security and privacy** → **Usability in security and privacy**; **Social aspects of security and privacy**; **Privacy protections**.

KEYWORDS

privacy-enhancing technologies, expert interview, cryptography, policy

ACM Reference Format:

Nitin Agrawal, Reuben Binns, Max Van Kleek, Kim Laine, and Nigel Shadbolt. 2021. Exploring Design and Governance Challenges in the Development of Privacy-Preserving Computation. In *CHI Conference on Human Factors in Computing Systems (CHI '21)*, May 8–13, 2021, Yokohama, Japan. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3411764.3445677>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI '21, May 8–13, 2021, Yokohama, Japan

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8096-6/21/05...\$15.00

<https://doi.org/10.1145/3411764.3445677>

1 INTRODUCTION

HCI research on privacy has traditionally focused on end-users: understanding their privacy attitudes and mental models, studying their privacy-related behaviours, and designing tools to help them manage data disclosure according to their preferences. While important, this paradigm of end-user privacy also has limitations. First, individuals may have their data processed in remote and opaque ways by dint of being taxpayers, credit risks, or suspected terrorists — not ‘end users’ as traditionally conceived. In such cases we still need to understand how privacy as a human right and public good can be reflected and governed in such systems. Second, the end-user privacy paradigm neglects the many other entities who play an important role in articulating, navigating, and embedding privacy in a range of contexts. If HCI is about reflecting human values in computer systems more broadly [99], it is equally important to study those people, whether they are developers [10], designers, risk managers, or policy makers. Finally, addressing privacy as a problem of end-user interaction often yields depressing results due to the sheer complexity of personal data processing making it difficult for end-users to comprehend the choices and tools available. This is perhaps especially true when such complexity is the result of modern cryptographic techniques designed to *protect* privacy [115].

These three limitations are particularly salient in the context of this paper, which addresses technologies for *privacy-preserving computation*. These are a subset of Privacy Enhancing Technologies (PETs) which have emerged in recent years. These include Homomorphic encryption (HE), secure multi-party computation (SMPC), and differential privacy (DP). These foundational technologies share a common promise: to preserve privacy while also obtaining the benefits of computational analysis. HE enables computation on encrypted data, making it possible to outsource computation to another entity without them ever having access to the input data in the clear. SMPC allows multiple parties to jointly perform a computation based on multiple respective inputs without revealing those inputs to each other. DP refers to a way of measuring the extent to which the output of a computation reveals information about an individual, and a range of associated techniques for reducing it. While these technologies have been available in some

form for years, and to some extent already are in deployment, recent progress in their foundational techniques and computational tractability has led some to anticipate their imminent adoption.¹

Compared to systems and contexts typically studied in privacy-related HCI research, these privacy-preserving computation techniques may be far removed – conceptually, operationally and experientially – from the entities whose privacy they purport to protect. While HE, SMPC, and DP are sometimes touted in the marketing campaigns of some device makers,² for the most part these technologies are deployed as invisible infrastructure rather than being positioned as features which end-users are expected to value, let alone understand or control themselves. In many other (actual or envisioned) deployment contexts, the data being kept private may relate to individual data subjects who are not informed or engaged with its processing; and even if they were aware, they may have no ability to distinguish between whether such processing was genuinely ‘privacy-enhancing’ or not. Furthermore, the mathematical and computational complexity underpinning these techniques raises particular challenges to explaining them to various stakeholders; not only end users and/or data subjects, but also developers, investors, product managers, and policymakers.

These differences make privacy-preserving computation technologies a prime case study for an expanded understanding of privacy within HCI beyond traditional paradigms of user attitudes and behaviours[41], to consider developers[3], managers, policymakers and others [86], and the roles they play in defining and operationalising goals like security and privacy. For better or worse, the development and adoption of these technologies, and the political values and consequences they reflect may ultimately have relatively little to do with ‘end users’ as traditionally conceived. With these considerations in mind, this paper aims to explore the following:

- (1) What challenges are associated with the adoption of privacy-preserving computation techniques for different stakeholders?
- (2) What are the motivations for adopting them?
- (3) Why and how should privacy-preserving computation technologies be explained, governed, and made accountable to data subjects and wider society?

To gain insight into these questions, we undertook a series of interviews with a variety of stakeholders involved in various ways in the development and adoption of privacy-preserving computation technologies (PPCTs). These included cryptographers and theoretical computer scientists working on foundational PPC techniques, developers of practical tools and libraries for non-expert developers, senior managers and policymakers assessing and identifying real-world use cases, practitioners building PPC products, and designers working with PPCs as a design material. Our aim was to draw out implications for HCI and design raised by this new class of technologies.

¹Recent industry analyst reports have suggested that PETs are ‘experiencing a renaissance’[25], and that 2020 was the ‘the year of PETs’[29]

²See e.g. Apple https://www.apple.com/privacy/docs/Differential_Privacy_Overview.pdf and Google’s [38] DP initiatives

2 BACKGROUND

We begin by briefly introducing emerging privacy-preserving computation techniques. We then situate our approach to studying them in relation to prior related work in HCI.

2.1 Overview of Privacy-Preserving Computation Technologies

Privacy-preserving computation is a subset of Privacy-Enhancing Technologies (PETs). PETs are a broad category which could include everything from a sticker placed over a webcam [78] to advanced cryptographic techniques [82]. Existing and well-established examples include encryption schemes used to secure data at rest, end-to-end encryption protecting data over the network, and anonymous routing protocols to prevent interactions between identities from being revealed. Such technologies are already widespread, embedded in products and as part of the global internet infrastructure. While they each have different underlying approaches and motivations, these technologies are primarily concerned with the protection of data, at rest and in transit. They generally assume that once data is safely transferred to a secure endpoint, it can be decrypted and computed on in the clear; that a single entity performs the computation; and that whether or not the result of the computation is ‘private’ has a binary answer.

A more recent wave of PETs – including homomorphic encryption, secure multi-party computation, and differential privacy – allow these assumptions to be relaxed or even abandoned altogether. We briefly introduce them here.

2.1.1 Homomorphic Encryption. Informally, homomorphic encryption (HE) enables computation over encrypted data without ever ‘seeing’ the input or the output. This is realized through a specific encryption and decryption scheme. In effect, a user could send their encrypted data to a service provider who could then perform the desired computation and send back the output to the user, while remaining oblivious to both the input and the output. More formally, homomorphic encryption is an encryption primitive that enables secure evaluation of an arbitrary circuit f on an encryption $C(x)$ of a plaintext x , without decrypting $C(x)$ in the process, and without requiring any information about the private key. Such an encrypted evaluation results in an encryption $C(f(x))$ [45, 94], which can at a later point be decrypted by the owner of the private key, to reveal the result $f(x)$, as if f had been evaluated on the plaintext data. In principle, homomorphic encryption can be used to evaluate any circuit on encrypted data, but often a weaker functionality called leveled homomorphic encryption is used instead, which allows only circuits of a predetermined (but arbitrarily high) depth to be evaluated on encrypted data. In practice, the encryption scheme must be parameterised according to a desired depth bound of some interesting class of circuits. Homomorphic encryption, and often leveled homomorphic encryption, has found its application in problems such as secure data retrieval [6, 7, 21, 119], outsourced computation [12, 66] and secure machine learning as a service for sensitive data [46, 53, 97], amongst others.

2.1.2 Secure Multi-Party Computation. Secure Multi-Party Computation (SMPC) is a class of cryptographic primitives which enables secure evaluation of a function over data shared across multiple parties. It was formally introduced in 1982 as a 2 party protocol for the Millionaire’s problem [118]. Informally, SMPC primitives allow multiple parties to come together and jointly compute a function on their combined inputs while remaining oblivious to each other’s inputs; the Millionaire’s problem involves two parties learning which has greater wealth without revealing their respective fortunes. Formally, in an n -party setting, party P_i possess an input x_i and gets an output y_i upon computation of function f over the combined set of x_i s ($i \in \{1..n\}$). The secure computation guarantees privacy of the individual inputs x_i s. Most SMPC protocols could be defined by the choice of the circuit for computing a particular function and the type of secret sharing scheme. Use cases for SMPC include secure operations over distributed sensitive data such machine learning [5, 44, 59, 84, 93], genomic comparison [39, 61] and private set operations [55, 56].

2.1.3 Differential Privacy. Differential Privacy (DP) [33] is a framework for sharing information based on a dataset while statistically limiting information exposure about the individuals in the dataset. More broadly, the idea of differential privacy is to deploy a mechanism where the effect of a single substitution in a dataset is very small. In effect, a query on a dataset with such a mechanism in place does not reveal anything substantial about a single individual. Differential privacy may not always be considered a privacy-enhancing technology *per se*, but rather a theory for *measuring* privacy in a particular way. However there are several techniques which are closely associated with differential privacy, all of which involve adding noise to results according to differentially private constraints; we therefore refer to this family of techniques loosely as differential privacy technologies. Formally, a randomized function f gives (ϵ, δ) -differential privacy for all databases D and D' , non-negative values ϵ, δ and $\forall S \subseteq \text{range of } f$, where D and D' differs by at most one record iff,

$$\Pr[f(D) \in S] \leq \delta + e^\epsilon \Pr[f(D') \in S]$$

Here δ and ϵ are the privacy parameters. Differential privacy is one of the more widely deployed privacy-preserving computation technologies. It can be applied to querying databases[62], building differentially private machine learning models [2, 114] and performing statistical analysis [34, 35] with privacy guarantees. More recently, the US census used DP in 2020³, Apple has deployed local DP for a number of features⁴ and Google has been using DP for collecting data over its Chrome browser [38] in a privacy preserving manner.

2.1.4 ‘Privacy-preserving computation’. Despite their differences, these technologies have all been classed as ‘tools for privacy-preserving computation’[108], which enable ‘the derivation of useful results from data without giving other people access’ to such data [102]. These privacy-preserving computation technologies are still an emerging technology. While significant theoretical progress has been made, this has yet to be translated into widespread adoption.

However, numerous libraries exist for HE, SMPC and DP⁵, the number of government-funded PPC projects is increasing [108], and various industry and policymaking forums have publicly heralded their potential. Recent reports and working papers have catalogued actual or potential use-cases, as well as noting possible usability barriers including programming complexity, computational overhead, and parameter selection [20, 108].

2.2 Related work on PETs in HCI

Much research on privacy in HCI is concerned with how end-users value, negotiate, and manage privacy in the context of their interactions with computers. Work in this vein involves: understanding the attitudes [67, 69, 101, 109, 110], expectations [9, 75] and mental models [64] of end-users regarding how their data is collected and used; studying privacy-related behaviours such as willingness to share data [4, 74] and use of protective measures [14]; and evaluating and designing tools for privacy management such as permission settings [73, 76], privacy notices [11, 98], and privacy assistants [77]. Related research in usable privacy and security addresses the usability of various end-user PETs tools. These include privacy and security aspects of ubiquitous tools e.g. web browsers [26, 57], as well as more advanced specialist tools, such as end-to-end encryption [115] and anonymous communication and routing tools [22]. Such work is highly relevant to contexts in which end users directly interact with systems in ways that may affect their privacy, and where there are opportunities to (re)design tools and interfaces to give them more control. Such work is premised on the ideal of individual users being able to understand at least some aspects of how their data is processed, and having the potential to exert some meaningful choices over it.

In some cases, privacy-preserving computation technologies might be usefully studied from this end-user perspective. Bullek et al. [19] studied people’s comprehension of the randomized response method for local differential privacy [113]. Participants were asked a series of questions, the answers to which were perturbed with noise to provide privacy. In response to a final question about a particularly sensitive topic, they were able to choose how much perturbation to add (i.e. the value of ϵ). While most participants selected the lowest (most privacy-preserving) value for ϵ , surprisingly, 20% chose the highest (least privacy-preserving) value for ϵ . Some participants explained this was because adding more noise felt like lying. Xiong et al. [117] also studied participants’ willingness to share data with a hypothetical differentially private system. They examined the effect of different descriptions of differential privacy (including real descriptions provided by technology companies and the U.S. Census bureau) on willingness to share, and their findings suggest that certain descriptions (in particular, implication descriptions) are more understandable and increase willingness to share data as a result. Finally, Qin et al. explored usability and understanding in the context of privacy-preserving data aggregation initiatives based on MPC, finding that using various analogies to

³https://www.census.gov/newsroom/blogs/research-matters/2018/08/protecting_the_conf.html

⁴https://www.apple.com/privacy/docs/Differential_Privacy_Overview.pdf

⁵For HE: Microsoft SEAL [1], HELib (<https://github.com/homenc/HElib>), PALISADE (<https://palisade-crypto.org>); for SMPC: Crypten (<https://crypten.ai>), emp-toolkit [112], SPDZ (<https://github.com/bristolcrypto/SPDZ-2>); for DP: Google’s DP framework (<https://github.com/google/differential-privacy>), Diffprivlib (<https://github.com/IBM/differential-privacy-library>), Pysift [96].

explain the process of additive secret sharing increased participants' confidence in the scheme [92].

However, in many contexts, the 'data subjects' are not co-extensive with the 'users'. In the case of the PPCTs mentioned above, there may be several primary users (which may include developers and others) and many wider 'stakeholders' (e.g. commercial and government partners, the wider public). Rather than studying end users who are also data subjects, then, we might instead follow previous HCI research on privacy and security which focuses instead on other actors, such as developers (e.g. [3, 8, 10, 48]). Balebako et al. note that while users may be concerned about privacy, they are generally not 'empowered to protect themselves'; by contrast 'the decisions made by app developers have great impact' [10]. Studying developers, designers, and others can reveal both practical and organisational challenges hindering the deployment of privacy and security technologies [3, 42], highlight discrepancies between privacy research and privacy engineering [51, 68], as well as elucidate the moral dimensions of design. While some studies of software developers suggest that they may 'not have sufficient knowledge and understanding of the concept of informational privacy' [52], others do explicitly engage with the ethical and political ramifications of their work; e.g. Rogaway [95] who acknowledges how the field of cryptographic privacy technologies has 'an intrinsically moral dimension'.

This kind of reflexivity on the part of developers and designers is something acknowledged and addressed in approaches like Value Sensitive Design (VSD), which aim to 'illuminate the ethical and moral responsibility on the part of the designer rather than the user' [41]. To understand how particular technologies are imagined as solutions to problems [60], we may need to study a wide variety of actors involved in their development, not only engineers but also those involved in the business of marketing them [86]. By encompassing the full breadth of different actors involved in creating and deploying these systems, we are also able to grapple in different ways with the trade-offs and tensions inherent in the field of privacy-preserving computation, and ask questions like "Who is making the design decision?", "Who is paying for it?", "What is this saying about the user?" [54].

Finally, there is also work which critically addresses Privacy-Enhancing Technologies from a philosophical and conceptual perspective. For instance, Tavani and Moor [107] assess how earlier PETs such as PGP and anonymity tools may address privacy as individual control, but do not provide 'external' control beyond the user, which they argue is necessary to protect privacy in the round. Gurses and Berendt point to the limitations of PETs that stem from understanding privacy solely in terms of confidentiality [50]. Stalder points to the ways that PETs designed for individual use may occlude broader social meanings of privacy [104], while Phillips notes how PETs designed to assist businesses with automating compliance with privacy laws reinforce a restricted notion of privacy as unwanted intrusion [89].

3 RESEARCH APPROACH

Given that the technologies being addressed here are still emerging, and the broad and exploratory nature of our research questions, we chose to undertake in-depth semi-structured interviews with a

select range of experts from a range of backgrounds and roles [15]. All had direct experience of working on projects relating to privacy-preserving computation, and occupied different strategic positions in the developing ecosystem. They included: researchers working across HE, SMPC and DP research; industry practitioners and designers with experience delivering practical applications of these technologies; as well policy experts with experience in PETs. We deliberately selected some experts whose careers and roles bridged between the domains of research, industry, or policy, some having moved from one to the other over the course of a career, while others maintaining feet in multiple domains simultaneously. These participants can be seen as 'boundary workers', working between the boundaries of science and policy to facilitate the co-production of knowledge and innovation [58] and 'knowledge brokers' who facilitate connections between scientific and other audiences [83]. Including a variety of different roles also reflects the nature of these technologies as 'use-inspired basic research' [87] operating between 'basic' and 'applied' research paradigms [106]. This enabled us to not only understand how the knowledge surrounding these technologies is made in specific places (e.g. research labs, technology companies, government) but also 'how transactions occur between places' [100].

Because these technologies are still emerging, we inevitably could only draw from a small class of professionals, whose roles in the production of these technologies are to some extent ill-defined. As is typical with expert interviews, there was no comprehensive list of relevant experts to sample from; we therefore built a 'sample frame' based on publicly available materials from a wide variety of sources including research papers, industry and policymaking fora, and press [47], to identify potentially relevant experts, and also used snowball sampling. As a result of the variety of roles and experiences of our 9 experts (see Table 1), we used a semi-structured interview format slightly tailored to four different roles (research, industry, policy, design). We invited participants to discuss their experiences, motivations, perceived opportunities and barriers, relating to this space. Open ended questions allowed us to have relatively free reign to explore these issues [88]. The interviews were conducted over video chat during Spring and Summer 2020. The 9 interviews varied in length between 35-75 minutes, with the average taking 55 minutes, producing 8.3 hours of audio recordings in total, which were transcribed. All parts of the study were approved by our institution's ethics review committee. We used thematic analysis [17] to identify key themes and ideas discussed by experts. Two of the authors independently developed a set of codes based on close reading of disjoint subsets of the interview transcripts, using an open coding process. The two authors then discussed and consolidated their codes to derive a common set, which was then applied by both authors to all the interview transcripts, and memo notes were taken to record observations about the codes and their relation to one another [72]. A final round of discussion based on this data resulted in a set of themes and sub-themes, presented in the following section.

4 FINDINGS

We divide the findings from our interviews into two main areas. The first addresses the technical challenges and opportunities around

Participants	Description (& years experience in PETs)
P1[R]	Cryptography Researcher and an industrial PETs library developer (20-30y experience)
P2[R]	Cryptography Researcher at company specialising in privacy preserving computation (5-10y experience)
P3[R]	Cryptography Researcher and an industrial PETs library developer (5-10y experience)
P4[P]	Law and Policy professional working on industrial adoption of privacy-preserving ML (4-6y experience)
P5[P]	Senior government adviser on technology, with strong interest in PETs and their applications (5-10y experience)
P6[I]	Security and Privacy Researcher working in executive role at large tech company (2-5y experience)
P7[I]	Data Scientist working on privacy at a 'big four' accounting firm (5-10y experience)
P8[I]	Researcher at consultancy specialising in privacy preserving computation as a service (2-5y experience)
P9[D]	Designer at a tech and design agency specialising in ethical use of data & AI (2-5y experience)

Table 1: A summary table of the total participant sample from Research (R), Policy (P), Industry (I) and Design (D).

the adoption of privacy-preserving computation techniques - primarily concerning their transition from theoretical research into practical application. The second concerns the motivations and goals for deploying these techniques to address commercial and societal goals, and how the institutions that deploy them might explain and be held accountable for their use.

4.1 Technical challenges and opportunities

4.1.1 From Theory to Practice. Many participants brought up how advances in the theoretical grounding on which privacy-preserving computation techniques are based may not translate straightforwardly into specific real-world applications. This was acknowledged by both research scientists working on those foundations, and practitioners attempting to deploy the technologies in particular contexts. Many participants expressed confidence that those theoretical advances would translate into practice in time. P3, a researcher, argued that while it is *'early stages, from the business development perspective of homomorphic encryption'*, he was nevertheless *'confident that [the] technology is useful, practical'* (P3[R]).

It was generally accepted that much of the research had only reached the stage of proofs-of-concept rather than deployments; while *'doable in principle'*, *'there is a lot of work to do before that bigger picture potential can be realized'* (P1[R]). Similarly, P9[D] explained that as a designer she was *'trying to make them more widely understood within the design and tech community ... There's lots of research in academia at the moment, but not very many examples of them being used in practice'*.

The policy experts we interviewed were optimistic that the technology was already nearly ready for practical deployment. For P4[P] *'the technology has scaled to the point that ... it's definitely commercially deployable'*. For P5[P], while practical deployment would require a series of *'reasonable engineering and architectural compromises'*, he was still optimistic that *'existing approaches to homomorphic encryption are tractable'*.

While both research scientists and policy experts were optimistic about the big picture, those trying to bridge theory and practice on the ground expressed frustration that much of the research was not directly relevant. In some cases, this was because the scientific work made simplifying assumptions that were rarely satisfied in real use cases. In the context of trying to apply differential privacy techniques to a project involving time series data, P7[I] admitted that he *'really struggled, you know, seeing the value in all those techniques that academia likes to talk about ... how am I going to use that with time series?'*

Similarly, several participants pointed to the variety of messy underlying data and software issues that exist on the ground that hamper deployment. P7[I] explained that the initial challenges are around *'how do we list data assets, and manage access, at an enterprise scale?'*; while P8[I] spoke of clients with various custom systems and data formats, so *'while we're solving the security aspect of the communication between counterparties ... we still haven't fixed this engineering problem - it's not going away'*.

P2[R], a research scientist who had worked on both foundational theory and engineering, explained how applying techniques in practice involved moving carefully between theory and engineering:

'Engineers do need to learn a lot to deploy these kind of technologies and what makes the whole thing complicated is that they need to acquire a kind of knowledge that is not something that the university professor knows ... a lot of low level optimizations make a huge difference, and yeah, from theory to practice they need to somehow be invented.'

In P2's opinion, such work would not come from a *'linear transmission of knowledge'*, but rather through continuous iteration and *'course correction'* between theoreticians and engineers.

4.1.2 Interdisciplinarity and translation between roles. As well as theoreticians and engineers to bridge the gap between theory and practice, several participants discussed the need for people with different skills, backgrounds and motivations to work together. This included not only the combinations of different expertise involved in foundational privacy-preserving computation research such as math, statistics, and cryptography, but also specialists in specific application domains. As P2[R] noted, successful deployment depends on a *'component of multidisciplinary'*:

'In my experience this is very hard to get the right people and crucially, with the right incentives in the same room ... You not only need data scientists, but also security engineers, mathematicians and then experts in the application domain.'

The differing motivations and cultures of these different communities was seen by some as a problem as it leads to certain

important problems being neglected. Commenting on the misalignment between incentives of academic researchers and industry, P3[R], who had worked in both sectors, lamented how research is *'driven by the need to get published; it favors more ... performance breakthroughs, functional breakthroughs'*, meanwhile, topics like usability are *'not so interesting for the basic core research community'*.

The need for an even broader range of disciplinary expertise and professional skills was articulated by P4[P], who described her role as *'to bridge the lexical gap between technologists, lawyers, and policymakers to defragment the current initiatives in PETs'*. Drawing from previous experience working on AI in government, where *'insulated development'* led by technologists failed to account for the *'constitutional implications'* of these technologies, she warned that *'the same could happen for PETs without this sort of ... interdisciplinary discourse'*.

4.1.3 Usability for Developers. A common theme among both researchers and industry practitioners was the complexity of applying privacy-preserving PETs from a software engineering perspective. They discussed a set of inherent challenges facing developers around flexibility, performance, and specifying appropriate parameters.

Several participants described how PPCTs, in particular homomorphic encryption, can be very *'brittle'* (P1[R]): small changes in parameters can result in drastic reductions in performance, security, or privacy guarantees. Such sensitivity can be hard for developers to anticipate and manage, especially as there are many different parameters to tune. This was contrasted against other PETs, like public key cryptoschemes, where there is one main parameter — key size— which has a fairly predictable relationship with security guarantees and computational overheads:

'I mean RSA, you have the bit length and that's pretty much it. These things [homomorphic encryption applications] you have a ton of decisions to make when it comes to how to instantiate it, and they have implications for both speed and for the actual function that you will need to compute.' (P1[R])

This results in problems for developers not just in the initial implementation of a privacy-preserving technique, but also as they inevitably need to update a system to *'evolve when you need to change various details ... performance and tractability can be so highly dependent on small details'* (P1[R]).

Participants articulated this as a trade off between approaches to building PPCTs that either work out-of-the-box but have poor and unpredictable performance, *or* that have reasonable performance, but require fine-tuning by engineers. For instance, while the invention of fully homomorphic encryption enables both addition and multiplication and therefore arbitrary computation, specific applications still need to be converted into those arithmetic operations and may incur great computational costs depending on how that is implemented. While it may be possible to *'come up with a system with adequate performance for your application'*, this often requires having an application which is *'fully specified and well defined in mind, and you have a team of experts working for you'* (P1[R]).

Several participants spoke about the development of privacy-preserving computation software libraries for developers (see 2.1.4),

often contrasting two approaches which reflect the trade-off articulated above: on the one hand, libraries which create an abstraction layer which obscures the underlying complexity; and on the other, libraries which expose all of that complexity so that the developers still need to create bespoke solutions for their application context. P1[R] noted that many developers expect a library to provide *'abstractions that are convenient'*; otherwise *'It's like telling people: OK, I'll give you transistors and you'll build from them ... people don't think this way and for good reason'*. P2[R] made a comparison to machine learning frameworks (e.g. Tensorflow and Scipy):

They have a very nice abstraction layer that allows them to say "OK, here's my function over the reals: optimize it, under the hood"... We will have worried about implementing matrix multiplication super quickly over floating point numbers so that the data scientists can assume [it's] like doing math on their on their notebook, right? This level of abstraction doesn't exist yet [for privacy-preserving computation]

However, P2[R] cautioned against such an approach for privacy-preserving computation libraries, because it would preclude *'a lot of optimizations that come from understanding the underlying protocol'*. P1 echoed this, stating that:

'The only way that we know now of making the computation go reasonably fast is to use a lot of tricks and the developer needs to know about those tricks'

As a result, P2[R] felt that *'general purpose tools'* would inevitably fail to meet developer's performance expectations and thus give them a mistaken impression about the true potential of PETs.

While many participants were in favor of some form of standardisation via libraries, P8[I] explained that the prospects for a standard platform depends on where the technologies are being deployed to. In the context of SMPC, because smartphone operating system providers *'control the platform, they can decide ... this is how it's going to work'*; whereas P8[I]'s work involved deploying SMPC into a wide range of different clients' environments where *'we can't really dictate to them how they store their data'*; as a result the possibility for standardisation was small.

4.2 Motivating, explaining and governing privacy-preserving computation

Our participants also raised various insights relating to the motivations for adopting PPCTs, and challenges relating to explanation and accountability.

4.2.1 Motivations. Unsurprisingly, 'privacy' was often cited as the motivation for developing and deploying privacy-preserving computation technologies. However, some subtly different articulations and understandings of privacy emerged from our interviews, as well as some other motivations which went beyond privacy altogether.

Some very directly motivated the adoption of privacy-preserving computation by reference to the interests of individuals in privacy and the protection of their personal data: *'it's individual privacy - it's human rights'* (P6[I]). In comparison to the push for similar technologies in other markets, privacy-preserving computation was more a response to individual privacy:

‘People do understand when their privacy is violated. So ... the push for these technologies is very different to the push the semiconductor industries have had ... So I give a lot of time on the examples that are user-centric’ (P6[I]).

However, appeals to individual privacy were often mediated via other pressures. First, organisations deploying PETs may not have a direct relationship with those individuals, but are instead concerned with third-parties affected through business-to-business relations:

‘You have customers: these may be business to business customers, but that also extends to customers of customers and therefore it boils down to individuals’ (P6[I])

Second, some cited the existence of privacy and data protection regulation as an incentive to provide and deploy PETs: ‘*because of GDPR [the E.U. General Data Protection Regulation]... all the regulatory environment is ... very favorable for providers*’ (P7[I]). This regulatory pressure meant that investment in PPC could be accounted for in terms of corporate risk management: ‘*to have compliance at least formally speaking with GDPR ... it’s really protecting assets of a company*’ (P6[I]).⁶

Third, P9[D] argued that rather than just enabling existing data processing to be done in a more privacy-preserving way, these technologies could enable new insights which ‘*you might not have been able to gain before because of the sensitivities around the data that you are using*’. P4[P] highlighted a range of ‘*missed opportunities*’ for privacy-preserving computation ‘*for a good purpose*’. These included cases such as the Boston Women’s Workforce Council who ‘*used secure multiparty computation to confidentially analyze gender wage gaps without ... disclosing who the salary belonged to*’.⁷ P5[P] noted the opportunities for government national security services to use HE techniques like private set intersection to identify suspects without combining certain databases in the clear, something that might not otherwise be undertaken due to the ‘*intrusiveness*’ of sharing data of large numbers of innocent citizens between departments.⁸

While individual privacy was cited by all participants as an important motivator, it was often an indirect motivator, and in some cases perhaps insufficient on its own (e.g. without being coupled with new opportunities to extract value from data). Other participants articulated motivations for pursuing privacy-preserving computation which had nothing to do with individual privacy as such. For example, for some researchers (e.g. P1, P2), it was basic intellectual curiosity (‘*somebody thinks of something that ... looks interesting to them*’ (P1)). Other cases included where competing businesses would have a mutual interest in the output of some computation on their respective data, but would not otherwise share that data out of ‘*fear of losing a competitive edge*’ (P6[I]). Intellectual property protection was also frequently cited as a key motivation for many business applications.

Privacy-preserving computation techniques were also seen by some as offering the possibility to navigate regulatory obligations and trade-offs in different ways. First, they have the potential to

fulfill obligations to protect data in new, more ‘technological’ ways, offering ‘*technological safeguards that can’t be easily overridden*’, the kind of protection that ‘*paper safeguards, like contractual guarantees and policies, just can’t provide*’ (P4[P]). They were seen as especially promising in cases where different regulatory obligations might appear to be in conflict, as P4[P] explained:

‘Anti money laundering regulations are very data maximalist; they want you to collect more data [to prevent] financial crimes. But in the meantime the GDPR is quite the opposite; it wants you to minimize data, ... and this really conflicts with the regime of AML. I think that PETs could actually cut through these legal conflicts and really provide a practical solution ... it’s not actually transferring PII, but it still allows for banks to prepare for AML protocols’

Similarly, for P5[P], privacy-preserving methods had the capacity to change what is possible without sharing data and thereby shift the scales in legal balancing tests [18] that might otherwise make certain data analysis unlawful:

‘UK law ... sets out a test for those of us in national security which is *necessity* and *proportionality*. So if you can shift the proportionality, then you’re in a better position so you can avoid intruding, you can avoid privacy risk’.

In these ways, such techniques were envisioned by P4[P] and P5[P] as enabling organisations in the public and private sector to break free of what P4[P] called ‘*legal gridlocks*’ that currently are (or are perceived to) exist around data use and enable new kinds of analysis.⁹

4.2.2 Explanation. Our participants discussed various facets relating to *explaining* privacy-preserving computation, including *how* they go about explaining it to different audiences (and in some cases, why they don’t even try).

The researchers described a variety of contexts in which they had had to explain underlying techniques and their strategies for doing so. For a general audience, P2’s strategy was to explain simplified versions of protocols, such as simulating a secure multi-party computation for dating using playing cards (see [80]). While these were ‘*fun to explain*’, P2[R] was unsure about the effectiveness of such explanations:

‘Then in the future, [the audience] will be like: “Oh yeah, multi-party computation, the thing with the cards.” That doesn’t mean that my explanation was effective... My feeling is that people tend to end up amused and satisfied.’ (P2[R])

Such explanations were offered as a starting point to encourage people ‘*who are attracted by that kind of magic*’ and would ‘*go into Wikipedia immediately after*’ (P2[R]). However, P3[I] felt that there was a lack of accessible educational material: ‘*there is certainly not enough material and the classical crypto papers are essentially useless for someone who is not an encryption expert*’; they suggested that

⁶A sentiment echoed in [20]

⁷<https://thebwwc.org/mpc>

⁸A use case discussed in [27].

⁹While a legal analysis of potential conflicts between these two areas of law is beyond the scope of this paper, we note that data protection and financial services regulators have, at least in the UK, affirmed their compatibility in general terms: see <https://www.fca.org.uk/news/statements/fca-and-ico-publish-joint-update-gdpr>.

explanations of core concepts might be more effective if tackled as part of a standardisation process and included within libraries.

Several participants also cautioned that the kind of explanations offered (if any) need to be tailored to the audience. On the one hand, explanations could be too technical: *'If you start with equations ... you lose 99 percent of the audience right away'* (P7[I]). On the other hand, short intuitive explanations might be too simplistic for informing executive decisions:

'So one thing is getting people interested, and the other one is informing, like, executive decisions. I don't think they should be informed ... by two minute stories... I don't think decisions about encryption are made based on an intuitive understanding of crypto'

For P9[D], designers have a role to play in explaining privacy-preserving techniques through prototyping their use in specific contexts. This included explanations to end users, but also *'a different language to explain it to those designers as well'*. Previously, their design agency hadn't *'seen much demand for them on the industry side'*; however, that changed after publishing a blog post explaining visually how differential privacy could work in the context of a project on identifying inequalities in urban mobility:

'Each step of the randomized response process ... we had an image to go with it, so that you could see ... the noise that you are adding to data. Visually seeing it was really helpful for me as a designer and then tying it to sort of real life stories so that I could see how you wouldn't be able to re-identify someone. Imagining what that makes possible forces you to think about the qualities of that technique, what it now enables you to do'

When it came to explaining these systems to end-users, however, some participants questioned whether this was a worthwhile goal. P9[D] couldn't imagine *'many scenarios where it's necessary to explain what privacy preserving techniques are being used to an end user who is trying to do something with their phone'*. P7[I] asked himself whether end-users understood these techniques, and answered: *'Well in general, not. Is it a problem? I'm not sure it's a problem'*. In such cases, it was seen as sufficient that end-users *'trust the provider of the solution that they do a good job'* (P7[I]).

4.2.3 Governance and Accountability. A final theme was around the challenges of *governance* and *accountability* of privacy-preserving computation. These topics often followed organically from discussions of explanation; attempts to explain these systems were often made in the course of trying to *justify* their use to affected stakeholders, and justification is a key element of accountability [16]). But even if explanations don't lead to real understanding on an individual level, it might still be possible to justify them to the public. P5[P] put it this way:

'These technologies are extremely difficult to understand... Do they meaningfully address genuine privacy issues? Yes they do. Do they address public concern? That's not to do with the technologies *per se*, [but] how the technologies are explicated and made available. If you told the public: "As a result of using these technologies, we are able to limit the amount of your

personal information that's shared, and are still able to offer you valuable services", they would be enthusiasts.'

Other participants expressed scepticism that the public would take such guarantees at face value. In the context of proposals for privacy-preserving facial recognition in border control, P6 asked:

'if someone publicises this new system ... just by saying: "and by the way the privacy of the information is very well handled because we use the state of the art cryptography", what does that mean to a citizen?'

Both P6 and P7 suggested that certifications and trust marks applied to services which use these techniques could enable individuals to seek out more trustworthy systems. However, expecting individuals to exercise meaningfully informed choices in relation to different services involving privacy-preserving computation was seen by some as adding to the burden of responsibility unhelpfully placed on individuals. P9 reflected on how *'constantly making decisions about data in the technology that we use is just not sustainable'*; instead, they suggested that *'collective consent models and other governance mechanisms ... that can make decisions on behalf of people'* might be a better approach. Similarly, P2[R] felt decisions about the technical details of the adoption of these technologies ought to be made by *'using experts or authorities'* who can act as *'proxies ... [who] understand their communities'*.

While most of our participants pointed to the positive potential of privacy-preserving computation techniques, a few were also concerned about the power imbalances they might reinforce. When the stakeholders are individuals, they are *'by definition, the weaker party'*, and *'lack the resources ... to induce changes; every time we talk about privacy there is some asymmetry that is implied by it.'* (P1[R]). For P9[D], it is important to recognise the limitations of PPCs as they are just:

'a technical solution to protecting people's privacy ... you have to think about the wider system that they sit within and what other kind of power dynamics are in that system.' (P9[D])

5 DISCUSSION

The findings from our interviews raised several important implications for the design and governance of privacy-preserving computation. They reveal how these techniques are being not only technically but also socio-technically constructed and constituted by a variety of actors, each pursuing overlapping and sometimes diverging agendas. Clearly, privacy-preserving computation techniques entail a variety of human-centric challenges which HCI research could seek to address. These challenges are multifaceted and will require diverse approaches; something that HCI as a methodologically diverse field is well-positioned to reflect. Furthermore, these challenges are inter-related: for instance, the way in which these technologies are translated from theory to practice may well affect how they can be explained and held accountable; while closer inspection of how 'privacy' and other motivations are unpacked might reconfigure what kinds of interdisciplinary collaborations are required in a particular context. Our aim in this section is to reflect on these, to understand both the design problems facing these techniques, and the challenges they raise in relation to the

interests of a variety of users and wider society. This discussion is not intended as direct ‘implications for design’; rather, we hope to draw attention to issues which require further research, as well as interdisciplinary discussion.

5.1 Whither the end user?

While our experts generally acknowledged the individuals whose personal data is being privately computed on as an important stakeholder group, few seemed to prioritise seeking their understanding and acceptance. This is in contrast to the small number of existing HCI studies that investigate ‘user acceptability’ of particular privacy-preserving computation techniques such as differential privacy [19, 117] and MPC [92]. User acceptability could and should be further examined in particular contexts; for instance, Colnago et al. suggest further work is needed to explore whether such techniques embedded in Internet-of-Things privacy assistants might ‘help mitigate people’s reservations about data collection practices and reduce the chance they opt out’ [24]. There is clearly great scope for important research within this paradigm of user acceptability.

However, our experts spoke about privacy-preserving computation technologies more as tools enabling *organisations* to achieve a variety of goals (including managing privacy risks, but also protecting corporate assets and secrets), rather than as a means of directly serving users’ interests. While user acceptance was not entirely disregarded, it did not appear to be a primary concern; even P9, a designer well-versed in user-centred design, doubted that people could or should be expected to understand and make decisions about privacy-preserving computation. Privacy may be important, and these techniques may have the potential to meaningfully embed it, but whether or not individuals understand and accept them seemed to be almost a secondary issue. In many of the use cases they mentioned, individuals whose data is being computed may not have any direct interaction with the system, nor any choice about whether to use it. In expressing such doubts, our interviewees might appear to be denying a sacrosanct tenet of HCI as a *human-centred* discipline. However, rather than denying the importance of user acceptance, we believe that these doubts should in fact point us towards alternative human-centric approaches to the development of privacy-preserving computation, in addition to solely looking at end users as data subjects.

First, our findings point towards studying the needs of different kinds of end users; specifically, those developers and designers who attempt to apply foundational privacy-preserving computation techniques in real-world applications. This echoes recent calls to acknowledge that ‘developers are users too’, as Green and Smith argue in relation to crypto and security libraries [49]. Similarly, P9[D] pointed to the relative lack of awareness and understanding of these techniques among designers. As with the application of other complex methods in computer science, such as machine learning, it may be difficult for designers to use privacy-preserving computation techniques in design practice due to unfamiliarity with how they work and awareness of what they can achieve [32]. P9[D] made the case for technical specialists and designers to work together to translate these technologies into ‘design material’ which design practitioners can use to explore real use cases.

In addition to understanding developers, designers, and others as *users* of privacy-preserving computation techniques, studying them also allows us to explore how a human value like privacy shapes the construction of complex computational systems. This perspective accords with ‘third wave’ approaches to HCI which orient attention towards the ethical obligations and values of designers [41], and incorporate different disciplinary perspectives which examine how social and political dimensions are embedded and reflected in systems [13]. As such, rather than just considering whether end users or laypeople understand, trust, and accept privacy-preserving computation technologies, we might also benefit from considering the perspectives of the various people involved in constructing their technical, commercial and regulatory foundations. Assessing whether an innovative technology will be acceptable to users through lab and field studies may be valuable, but such approaches often neglect the ways in which such technologies are interpreted, shaped, and mutually constructed over time through their designers, users, and broader political, economic and regulatory forces [60, 79]. As a result, it is equally important to consider the plurality of different actors and broader contexts through which values like privacy will be understood, traded-off, and embedded in these systems (or not).

5.2 The Limits of Abstraction

If, as suggested above, we are to consider the needs of developers and designers as users of underlying privacy-preserving computation techniques, then how might those needs be met? Many of the interviewees identified the need to create building blocks for privacy-preserving computation. In an ideal world, these building blocks would allow developers to *abstract away* the technical details and apply them to applications in different contexts. Creating such abstractions is fundamental to progress in computer science and programming; in Edsger Dijkstra’s words, it is ‘our only mental aid ... to organize and master complexity’ [31]. However, many of the experts expressed uncertainties about the form such abstractions should take and the extent to which they could reasonably be made in the domains of privacy-preserving computation. Especially with homomorphic encryption, abstracting away the details of implementation could mean losing the ability to optimise performance through engineering ‘tricks’ (P1[R]).

Attempts to create tools for developers to enable them to integrate privacy-preserving computation techniques into their products may therefore need to grapple with this need to balance abstraction and engagement with the implementation details. Specific applications will always require some ‘intimacy with the details’ [105] that might otherwise be abstracted away. Some of our interviewees argued that the necessary education required for developers could potentially be integrated into standardised APIs. This suggests that broader adoption of privacy-preserving computation may benefit from work in HCI which considers APIs and libraries as ‘first class design objects’ [85, 120], with the goal of ‘driving adoption of software components’ [81]. This could involve (re)designing them around the typical ways programmers learn, e.g. on-the-fly, via information foraging, and trial and error [65, 71].

However, the nature and extent to which developers need to become intimate with the details, and how they might do so, will

clearly depend on the particular technique in question. For instance, a DP library might implement a variety of noise sampling and injection techniques, but this is relatively simple compared to the much more complex mathematics and reasoning involved in deciding on and managing an appropriate privacy budget, which requires case-by-case human consideration. For SMPC, libraries might take care of some of the networking details, but leave difficult decisions regarding the protocol up to the developer. The nature and value of these standardised building blocks will therefore vary greatly between approaches.

Ultimately, the design and adoption of these privacy-preserving computation building blocks may need to reckon with the messy realities of underlying enterprise IT infrastructure, agile and iterative approaches to software development [51], and service-oriented architectures [68]. Given these practical considerations, the full complexity of these technologies might instead need to be mediated via a two-step process: general-purpose libraries which expose all of the complexity of a domain (e.g. homomorphic encryption) that enable specialist privacy engineers to create particular privacy-preserving computation components for common operations or use cases (e.g. private set intersection for contact discovery); those components could then be adapted and deployed with minimal configuration by non-specialist developers as microservices.

5.3 Privacy-Enhanced Technocracy

The way privacy-enhancing technologies are sometimes described can make them seem esoteric, exotic, and mysterious. For instance, in industry press they have been described as ‘black magic’¹⁰ and a ‘holy grail’¹¹. Such language suggests that their development is entirely in the hands of a small and specialised cabal of cryptographers and engineers, much like the early programmers who regarded themselves as ‘high priests’ of assembly code.¹² It is possible to imagine how in these respects, they might end up sharing the same ‘rampant hyperbole and political envisioning’ [37] of a higher-profile cryptographic technology — blockchain.

While our interviewees avoided such language, and even criticised the perceived ‘hype’ around PETs, they did reflect the highly specialised knowledge required to make use of the underlying mathematics, and drew parallels with magic. P6[I] described feeling like ‘Gandalf the wizard’ upon telling people that computation on encrypted data was possible, while P1[R] described the need for engineering ‘tricks’ to optimise performance within reasonable levels. From this perspective, the technical work of applying privacy-preserving computation seems more like craft than science, which the guild of cryptographers and engineers are uniquely capable of performing [91].

However, the mystery of their inner workings could easily serve as an excuse for not making these systems accountable to affected stakeholders. When reflecting on the challenges laypeople face in trying to understand PPCTs on any meaningful level, both P5[P] and P7[I] expressed some doubts about the possibility that individuals could ever be expected to really understand *how* they work. However, without some form of explanation, and absent any other

mechanism for meaningfully communicating their risks and opportunities, there is a risk that privacy-preserving computation becomes not just a *technical* but a *technocratic* solution imposed on populations without popular consent by grey eminences operating behind the scenes.

However, several of the experts did acknowledge the need for mechanisms of accountability and governance to developed as these technologies are rolled out. P6[I] and P7[I] suggested this could involve certification schemes. Similarly, while P1[R] and P9[D] were doubtful about individuals being able to meaningfully consent to these technologies, they proposed alternative forms of collective governance, where the interests of affected individuals could be represented by relevant representatives and experts who can make informed choices and demands on their behalf. These and other democratic mechanisms will need to be explored in order to counter a privacy-enhanced technocracy, and methods from HCI — such as participatory design [36], futures workshops [63], and other governance approaches — may have much to offer.

5.4 Secrets, Assets, Human Rights: Unpacking ‘Privacy’

Our findings attest to the many varied interpretations and uses of the term privacy. As previous work has explored, and as discussed above, the notion of privacy in Privacy-Enhancing Technologies is often a narrow interpretation of what is a multi-faceted and contested concept [50, 89, 104, 107]. This is certainly the case for the subset of privacy-preserving computation PETs studied here. They turn privacy into something mathematically formalisable, e.g. in terms of entropy in cryptographic approaches, or indistinguishability in statistical approaches, which can all be understood as variations of ‘confidentiality’, a pillar of the security triad [30]. This means that other ways of understanding privacy may be de-emphasised and de-prioritised.

There are continuities here with earlier PETs, such as de-identification techniques based on hashing personal data. Phillips argues that that these techniques embody privacy as protection ‘from unwanted intrusion’ [89]. However, they leave in place the ability of powerful observers to produce ‘panoptic’ knowledge which can be used to sort and discipline populations [23, 43]. Similarly, if we understand privacy as confidentiality, this can be engineered through architectures of data minimisation [103]; but this can lead to design choices which preclude alternative understandings of privacy (e.g. privacy-as-control), and hinder the exercise of related rights afforded by data protection law [111]. In our expert’s discussions, these alternative understandings of privacy were conspicuous by their absence.

Our findings also demonstrate that even while discourse around privacy-preserving computation restricts certain interpretations of privacy, it also stretches the meaning of privacy to incorporate unorthodox meanings, such as competitive secrecy, corporate asset protection, and government security. These are clearly significant and important use cases for the technology, but they arguably bear only a family resemblance to privacy as it relates to individuals and society. Indeed, intellectual traditions which value privacy as an individual right and public good have often been associated with opposition to corporate and government secrecy; according to

¹⁰<https://dualitytech.com/tag/homomorphic-encryption/>

¹¹<https://www.fastcompany.com/90314942/duality-homomorphic-encryption>

¹²In the words of Rear Admiral Grace Hopper [116]

them, privacy should be reserved for the weak, while transparency should be an obligation required of the strong [28]. In referring to all of these things as ‘privacy’, privacy-preserving computation technologies may elide significant political tensions between them. This is not to deny that they may have a powerful role to play in supporting privacy as an individual right and as a public good [40, 70]; but this confluence of quite different values under one banner complicates the narrative around whose interests they serve.

As well as tending to address narrow and perhaps unorthodox conceptualisations of privacy, it is important to recognise that these technologies do not protect other important values and interests. If our aim is to build and shape systems encompassing multiple social goals, where privacy is just one such goal, then privacy-preserving computation techniques have to be considered in relation to the whole system and the social context. The danger is that the societal problems of data processing technologies — such as the ways they create distinctions and hierarchies that reinforce power, shape politics, or facilitate abuse — are sidelined, redefined, or collapsed under the banner of ‘privacy’, so that privacy-preserving computation techniques can be positioned as *the* solution (what Pinch and Bijker term ‘closure by problem redefinition’ [90]). This danger was alluded to in P6[I]’s example of privacy-preserving border control (where people are still ultimately at the mercy of a powerful state), and in P9[D]’s concern about considering the wider power dynamics in the context of deployment.

6 CONCLUSIONS

Like other technologies which have been touted as potentially revolutionary in recent years, the concrete impact of these privacy-preserving computation techniques remains to be seen. New technologies often emerge in unexpected ways, at unpredictable times from niches of computer science: hypertext, Merkle trees, and neural networks were once confined to their respective research sub-fields before they became known more widely as the world wide web, blockchain, and ‘AI’ (in its latest guise of deep learning). Prior to their take-up in wider society, these specialised areas of research were conceived as laying the groundwork for purely technical pieces of invisible infrastructure, whose implications for human-computer interaction were remote and unclear.

However, we believe it is worth HCI researchers studying such technologies prior to their widespread adoption. Whatever technical and institutional forms they take, the journey of privacy-preserving computation techniques from the annals of cryptography into production code will be shaped in substantial part by the approach they take to a variety of human and societal challenges. Indeed, these challenges directly implicate some fundamental concerns of HCI, including: multifaceted (re)conceptualisations of the notion of ‘the user’; helping people navigate and manage computational complexity and its consequences; exploring how values like privacy can be reflected in the systems we build; and examining how different political agendas, economic rationale, and user groups shape and are shaped by those systems. These concerns all cohere and overlap in the emerging space of privacy-preserving computation.

This paper has aimed to provide a preliminary and partial outline of those challenges, laying some of the groundwork for substantial

further exploratory and in-depth work to be done. In addition to several recent studies which focused on people’s understanding of these techniques and their willingness to disclose personal data in the presence of them, we have outlined a broader set of research questions prompted for HCI by PPCs. These include understanding specific application contexts; usability of PPC libraries and tools from a non-specialist developer’s perspective; and understanding the explanation and governance challenges associated with these techniques.

ACKNOWLEDGMENTS

This work was funded by EPSRC grant EP/S035362/1 and Callsign Inc.

REFERENCES

- [1] 2020. Microsoft SEAL (release 3.5). <https://github.com/Microsoft/SEAL>. Microsoft Research, Redmond, WA.
- [2] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. 308–318.
- [3] Yasemin Acar, Sascha Fahl, and Michelle L Mazurek. 2016. You are not your developer, either: A research agenda for usable security and privacy research beyond end users. In *2016 IEEE Cybersecurity Development (SecDev)*. IEEE, 3–8.
- [4] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015. Privacy and human behavior in the age of information. *Science* 347, 6221 (2015), 509–514.
- [5] Nitin Agrawal, Ali Shahin Shamsabadi, Matt J Kusner, and Adrià Gascón. 2019. QUOTIENT: two-party secure neural network training and prediction. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. 1231–1247.
- [6] Adi Akavia, Dan Feldman, and Hayim Shaul. 2019. Secure Data Retrieval on the Cloud: Homomorphic Encryption meets Coresets. *IACR Transactions on Cryptographic Hardware and Embedded Systems* (2019), 80–106.
- [7] Sebastian Angel, Hao Chen, Kim Laine, and Srinath Setty. 2018. PIR with compressed queries and amortized query processing. In *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 962–979.
- [8] Hala Assal and Sonia Chiasson. 2019. ‘Think secure from the beginning’ A Survey with Software Developers. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [9] Rebecca Balebako, Jaeyeon Jung, Wei Lu, Lorrie Faith Cranor, and Carolyn Nguyen. 2013. Little brothers watching you: Raising awareness of data leaks on smartphones. In *Proceedings of the Symposium on Usable Privacy and Security*. ACM, 12.
- [10] Rebecca Balebako, Abigail Marsh, Jialiu Lin, Jason I Hong, and Lorrie Faith Cranor. 2014. The privacy and security behaviors of smartphone app developers. (2014).
- [11] Rebecca Balebako, Florian Schaub, Idris Adjerid, Alessandro Acquisti, and Lorrie Cranor. 2015. The Impact of Timing on the Salience of Smartphone App Privacy Notices. In *Proceedings of the ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices*. ACM, 63–74.
- [12] Manuel Barbosa and Pooya Farshim. 2012. Delegatable homomorphic encryption with applications to secure outsourcing of computation. In *Cryptographers’ Track at the RSA Conference*. Springer, 296–312.
- [13] Jeffrey Bardzell and Shaowen Bardzell. 2015. Humanistic hci. *Synthesis Lectures on Human-Centered Informatics* 8, 4 (2015), 1–185.
- [14] Lemi Baruh, Ekin Secinti, and Zeynep Cemalcilar. 2017. Online Privacy Concerns and Privacy Management: A Meta-Analytical Review. *Journal of Communication* 67, 1 (2017), 26–53.
- [15] Alexander Bogner, Beate Littig, and Wolfgang Menz. 2009. *Interviewing experts*. Springer.
- [16] Mark Bovens, Thomas Schillema, and Robert E Goodin. 2014. Public accountability. *The Oxford handbook of public accountability* 1, 1 (2014), 1–22.
- [17] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.
- [18] Ian Brown and Douwe Korff. 2009. Terrorism and the proportionality of internet surveillance. *European Journal of Criminology* 6, 2 (2009), 119–134.
- [19] Brooke Bullek, Stephanie Garboski, Darakhshan J Mir, and Evan M Peck. 2017. Towards Understanding Differential Privacy: When Do People Trust Randomized Response Technique?. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 3833–3837.
- [20] Rosario Cammarota, Matthias Schunter, Anand Rajan, Fabian Boemer, Ágnes Kiss, Amos Treiber, Christian Weinert, Thomas Schneider, Emmanuel Stapf,

- Ahmad-Reza Sadeghi, et al. 2020. Trustworthy AI Inference Systems: An Industry Research View. *arXiv preprint arXiv:2008.04449* (2020).
- [21] Hao Chen, Zhicong Huang, Kim Laine, and Peter Rindal. 2018. Labeled PSI from fully homomorphic encryption with malicious security. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. 1223–1237.
 - [22] Jeremy Clark, Paul C Van Oorschot, and Carlisle Adams. 2007. Usability of anonymous web browsing: an examination of tor interfaces and deployability. In *Proceedings of the 3rd symposium on Usable privacy and security*. 41–51.
 - [23] Julie E Cohen. 2012. *Configuring the networked self: Law, code, and the play of everyday practice*. Yale University Press.
 - [24] Jessica Colnago, Yuanyuan Feng, Tharangini Palanivel, Sarah Pearman, Megan Ung, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. 2020. Informing the design of a personalized privacy assistant for the internet of things. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
 - [25] CPOmagazine.Com. 2020. What's Old Is New Again: Examining Privacy Enhancing Technologies. <https://www2.deloitte.com/nl/nl/pages/risk/articles/is-2020-finally-the-year-of-the-pets.html>
 - [26] Lorrie Faith Cranor, Manjula Arjula, and Praveen Guduru. 2002. Use of a P3P user agent by early adopters. In *Proceedings of the 2002 ACM workshop on Privacy in the Electronic Society*. 1–10.
 - [27] Emiliano De Cristofaro and Gene Tsudik. 2010. Practical private set intersection protocols with linear complexity. In *International Conference on Financial Cryptography and Data Security*. Springer, 143–159.
 - [28] Paul De Hert and Serge Gutwirth. 2006. Privacy, data protection and law enforcement. Opacity of the individual and transparency of power. *Privacy and the criminal law* (2006), 61–104.
 - [29] Deloitte.com. 2020. Is 2020 finally the year of the PETs (Privacy Enhancing Technologies)? <https://www2.deloitte.com/nl/nl/pages/risk/articles/is-2020-finally-the-year-of-the-pets.html>
 - [30] Gurpreet Dhillon and James Backhouse. 2001. Current directions in IS security research: towards socio-organizational perspectives. *Information systems journal* 11, 2 (2001), 127–153.
 - [31] Edsger W Dijkstra. 1982. Selected writings on computing-a personal perspective. Texts and monographs in computer science. Springer, doi 10 (1982), 978–1.
 - [32] Graham Dove, Kim Halskov, Jodi Forlizzi, and John Zimmerman. 2017. UX design innovation: Challenges for working with machine learning as a design material. In *Proceedings of the 2017 chi conference on human factors in computing systems*. 278–288.
 - [33] Cynthia Dwork. 2008. Differential privacy: A survey of results. In *International conference on theory and applications of models of computation*. Springer, 1–19.
 - [34] Cynthia Dwork and Jing Lei. 2009. Differential privacy and robust statistics. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*. 371–380.
 - [35] Cynthia Dwork and Adam Smith. 2010. Differential privacy for statistics: What we know and what we want to learn. *Journal of Privacy and Confidentiality* 1, 2 (2010).
 - [36] Pelle Ehn. 1988. *Work-oriented design of computer artifacts*. Ph.D. Dissertation. Arbetslivscentrum.
 - [37] Chris Elsdén, Arthi Manohar, Jo Briggs, Mike Harding, Chris Speed, and John Vines. 2018. Making sense of blockchain applications: A typology for HCI. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–14.
 - [38] Úlfar Erlingsson, Vasily Pihur, and Aleksandra Korolova. 2014. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*. 1054–1067.
 - [39] David Evans, Jonathan Katz, Yan Huang, and Lior Malka. 2011. Faster secure two-party computation using garbled circuits. (2011).
 - [40] Joshua AT Fairfield and Christoph Engel. 2015. Privacy as a public good. *Duke LJ* 65 (2015), 385.
 - [41] Daniel Fallman. 2011. The new good: exploring the potential of philosophy of technology to contribute to human-computer interaction. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1051–1060.
 - [42] Steven M Furnell, Nathan Clarke, Rodrigo Werlinger, Kirstie Hawkey, and Konstantin Beznosov. 2009. An integrated view of human, organizational, and technological challenges of IT security management. *Information Management & Computer Security* (2009).
 - [43] Oscar H Gandy Jr. 1993. *The Panoptic Sort: A Political Economy of Personal Information*. Critical Studies in Communication and in the Cultural Industries. ERIC.
 - [44] Chong-zhi Gao, Qiong Cheng, Pei He, Willy Susilo, and Jin Li. 2018. Privacy-preserving Naive Bayes classifiers secure against the substitution-then-comparison attack. *Information Sciences* 444 (2018), 72–88.
 - [45] C Gentry. 2009. A fully homomorphic encryption scheme. PhD thesis, Stanford University, 2009. (2009).
 - [46] Ran Gilad-Bachrach, Nathan Dowlin, Kim Laine, Kristin Lauter, Michael Naehrig, and John Wernsing. 2016. Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy. In *International Conference on Machine Learning*. 201–210.
 - [47] Kenneth Goldstein. 2002. Getting in the door: Sampling and completing elite interviews. *PS: Political Science and Politics* 35, 4 (2002), 669–672.
 - [48] Peter Leo Gorski, Yasemin Acar, Luigi Lo Iacono, and Sascha Fahl. 2020. Listen to Developers! A Participatory Design Study on Security Warnings for Cryptographic APIs. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
 - [49] Matthew Green and Matthew Smith. 2015. Developers are users too: designing crypto and security APIs that busy engineers and sysadmins can use securely. (2015).
 - [50] Seda Gürses and Bettina Berendt. 2010. PETs in the surveillance society: a critical review of the potentials and limitations of the privacy as confidentiality paradigm. In *Data Protection in a Profiled World*. Springer, 301–321.
 - [51] Seda Gürses and Joris Van Hoboken. [n.d.]. Privacy after the agile turn. ([n.d.]).
 - [52] Irit Hadar, Tomer Hasson, Oshrat Ayalon, Eran Toch, Michael Birnhack, Sofia Sherman, and Arod Balissa. 2018. Privacy by designers: software developers' privacy mindset. *Empirical Software Engineering* 23, 1 (2018), 259–289.
 - [53] Rob Hall, Stephen E Fienberg, and Yuval Nardi. 2011. Secure multiple linear regression based on homomorphic encryption. *Journal of Official Statistics* 27, 4 (2011), 669.
 - [54] S Harrison, D Tatar, and P Sengers. 2007. The three paradigms of HCI. alt. In *CHI'07*.
 - [55] Carmit Hazay. 2018. Oblivious polynomial evaluation and secure set-intersection from algebraic PRFs. *Journal of Cryptology* 31, 2 (2018), 537–586.
 - [56] Carmit Hazay and Yehuda Lindell. 2008. Efficient protocols for set intersection and pattern matching with security against malicious and covert adversaries. In *Theory of Cryptography Conference*. Springer, 155–175.
 - [57] Amir Herzberg. 2009. Why Johnny can't surf (safely)? Attacks and defenses for web users. *computers & security* 28, 1-2 (2009), 63–71.
 - [58] Robert Hoppe. 2009. Scientific advice and public policy: expert advisers' and policymakers' discourses on boundary work. *Poiesis & Praxis* 6, 3-4 (2009), 235–263.
 - [59] Siam Hussain, Baiyu Li, Farinaz Koushanfar, and Rosario Cammarota. 2020. TinyGarble2: Smart, Efficient, and Scalable Yao's Garble Circuit. In *Proceedings of the 2020 Workshop on Privacy-Preserving Machine Learning in Practice*. 65–67.
 - [60] Sheila Jasanoff and Sang-Hyun Kim. 2009. Containing the atom: Sociotechnical imaginaries and nuclear power in the United States and South Korea. *Minerva* 47, 2 (2009), 119.
 - [61] Somesh Jha, Louis Kruger, and Vitaly Shmatikov. 2008. Towards practical privacy for genomic computation. In *2008 IEEE Symposium on Security and Privacy (sp 2008)*. IEEE, 216–230.
 - [62] Noah Johnson, Joseph P Near, and Dawn Song. 2018. Towards practical differential privacy for SQL queries. *Proceedings of the VLDB Endowment* 11, 5 (2018), 526–539.
 - [63] Robert Jungk and Norbert Müllert. 1987. *Future Workshops: How to create desirable futures*. Inst. for Social Inventions.
 - [64] Ruogu Kang, Laura Dabbish, Nathaniel Fruchter, and Sara Kiesler. 2015. "My Data Just Goes Everywhere:" User Mental Models of the Internet and Implications for Privacy and Security. In *Proceedings of Symposium On Usable Privacy and Security*. 39–52.
 - [65] Caitlin Kelleher and Michelle Ichinco. 2019. Towards a model of API learning. In *2019 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE, 163–168.
 - [66] Florian Kerschbaum. 2012. Outsourced private set intersection using homomorphic encryption. In *Proceedings of the 7th ACM Symposium on Information, Computer and Communications Security*. 85–86.
 - [67] Jennifer King. 2013. "How Come I'm Allowing Strangers To Go Through My Phone?"—Smartphones and Privacy Expectations. In *Symposium on Usable Privacy and Security (SOUPS)*.
 - [68] Blagovesta Kostova, Seda Gürses, and Carmela Troncoso. 2020. Privacy Engineering Meets Software Engineering. On the Challenges of Engineering Privacy ByDesign. *arXiv preprint arXiv:2007.08613* (2020).
 - [69] Ponnurangam Kumaraguru and Lorrie Faith Cranor. 2005. Privacy indexes: a survey of Westin's studies. (2005).
 - [70] Zbigniew Kwecka, William Buchanan, Burkhard Schafer, and Judith Rauhofer. 2014. "I am Spartacus": privacy enhancing technologies, collaborative obfuscation and privacy as a public good. *Artificial intelligence and law* 22, 2 (2014), 113–139.
 - [71] Joseph Lawrance, Christopher Bogart, Margaret Burnett, Rachel Bellamy, Kyle Rector, and Scott D Fleming. 2010. How programmers debug, revisited: An information foraging theory perspective. *IEEE Transactions on Software Engineering* 39, 2 (2010), 197–215.
 - [72] Lora Bex Lempert. 2007. Asking questions of the data: Memo writing in the grounded. *The Sage handbook of grounded theory* (2007), 245–264.
 - [73] Pedro Leon, Blase Ur, Richard Shay, Yang Wang, Rebecca Balebako, and Lorrie Cranor. 2012. Why Johnny can't opt out: a usability evaluation of tools to limit online behavioral advertising. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1–13.

- Factors in Computing Systems*. ACM, 589–598.
- [74] Pedro Giovanni Leon, Blase Ur, Yang Wang, Manya Sleeper, Rebecca Balebako, Richard Shay, Lujo Bauer, Mihai Christodorescu, and Lorrie Faith Cranor. 2013. What matters to users?: factors that affect users' willingness to share information with online advertisers. In *Proceedings of Symposium on Usable Privacy and Security*. ACM, 1–7.
 - [75] Jialiu Lin, Shahriyar Amini, Jason I Hong, Norman Sadeh, Janne Lindqvist, and Joy Zhang. 2012. Expectation and purpose: understanding users' mental models of mobile app privacy through crowdsourcing. In *Proceedings of Conference on Ubiquitous Computing*. ACM, 501–510.
 - [76] Jialiu Lin, Bin Liu, Norman Sadeh, and Jason I Hong. 2014. Modeling users' mobile app privacy preferences: Restoring usability in a sea of permission settings. In *Symposium On Usable Privacy and Security*. 199–212.
 - [77] Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhamidi, Shikun Aerin Zhang, Norman Sadeh, Yuvraj Agarwal, and Alessandro Acquisti. 2016. Follow My Recommendations: A Personalized Assistant for Mobile App Permissions. In *Proceedings of the Symposium on Usable Privacy and Security*.
 - [78] Dominique Machuletz, Stefan Laube, and Rainer Böhme. 2018. Webcam covering as planned behavior. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13.
 - [79] Donald MacKenzie and Judy Wajcman. 1999. *The social shaping of technology*. Open university press.
 - [80] Antonio Marcedone, Zikai Wen, and Elaine Shi. 2015. Secure Dating with Four or Fewer Cards. *IACR Cryptol. ePrint Arch.* 2015 (2015), 1031.
 - [81] Adam Mechtley. 2020. API as Curriculum: Designing High-Level API Affordances as Instructional Scaffolds. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–9.
 - [82] Alfred J Menezes, Jonathan Katz, Paul C Van Oorschot, and Scott A Vanstone. 1996. *Handbook of applied cryptography*. CRC press.
 - [83] Morgan Meyer. 2010. The rise of the knowledge broker. *Science communication* 32, 1 (2010), 118–127.
 - [84] Payman Mohassel and Yupeng Zhang. 2017. Secureml: A system for scalable privacy-preserving machine learning. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 19–38.
 - [85] Brad A Myers and Jeffrey Stylos. 2016. Improving API usability. *Commun. ACM* 59, 6 (2016), 62–69.
 - [86] Nelly EJ Oudshoorn and Trevor Pinch. 2003. *How users matter: The co-construction of users and technologies*. MIT press.
 - [87] Antti Oulasvirta and Kasper Hornbæk. 2016. Hci research as problem-solving. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 4956–4967.
 - [88] Edward C Page, Bill Jenkins, William Ieuan Jenkins, et al. 2005. *Policy bureaucracy: Government with a cast of thousands*. Oxford University Press on Demand.
 - [89] David J Phillips. 2004. Privacy policy and PETs: The influence of policy regimes on the development and social implications of privacy enhancing technologies. *New Media & Society* 6, 6 (2004), 691–706.
 - [90] Trevor J Pinch and Wiebe E Bijker. 1984. The social construction of facts and artefacts: Or how the sociology of science and the sociology of technology might benefit each other. *Social studies of science* 14, 3 (1984), 399–441.
 - [91] Maarten Roy Prak. 2006. *Craft guilds in the early modern low countries: Work, power and representation*. Ashgate Publishing, Ltd.
 - [92] Lucy Qin, Andrei Lapets, Frederick Jansen, Peter Flockhart, Kinan Dak Albab, Ira Globus-Harris, Shannon Roberts, and Mayank Varia. 2019. From usability to secure computing and back again. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS)*. 2019.
 - [93] M Sadegh Riazi, Mohammad Samragh, Hao Chen, Kim Laine, Kristin Lauter, and Farinaz Koushanfar. 2019. {XONN}: XNOR-based Oblivious Deep Neural Network Inference. In *28th {USENIX} Security Symposium ({USENIX} Security 19)*. 1501–1518.
 - [94] Ronald L Rivest, Len Adleman, Michael L Dertouzos, et al. 1978. On data banks and privacy homomorphisms. *Foundations of secure computation* 4, 11 (1978), 169–180.
 - [95] Phillip Rogaway. 2015. The Moral Character of Cryptographic Work. *IACR Cryptol. ePrint Arch.* 2015 (2015), 1162.
 - [96] Theo Ryffel, Andrew Trask, Morten Dahl, Bobby Wagner, Jason Mancuso, Daniel Rueckert, and Jonathan Passerat-Palmbach. 2018. A generic framework for privacy preserving deep learning. *arXiv preprint arXiv:1811.04017* (2018).
 - [97] Amartya Sanyal, Matt J Kusner, Adria Gascon, and Varun Kanade. 2018. Tapas: Tricks to accelerate (encrypted) prediction as a service. *arXiv preprint arXiv:1806.03461* (2018).
 - [98] Florian Schaub, Rebecca Balebako, Adam L Durity, and Lorrie Faith Cranor. 2015. A design space for effective privacy notices. In *Proceedings of the Symposium On Usable Privacy and Security*. 1–17.
 - [99] Abigail Sellen, Yvonne Rogers, Richard Harper, and Tom Rodden. 2009. Reflecting human values in the digital age. *Commun. ACM* 52, 3 (2009), 58–66.
 - [100] Steven Shapin. 1998. Placing the view from nowhere: historical and sociological problems in the location of science. *Transactions of the Institute of British Geographers* 23, 1 (1998), 5–12.
 - [101] Irina Shklovski, Scott D Mainwaring, Halla Hrunnd Skúladóttir, and Höskuldur Borgthorsson. 2014. Leakiness and creepiness in app space: Perceptions of privacy and mobile app use. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2347–2356.
 - [102] The Royal Society. 2019. *Protecting Privacy in Practice: The Current Use, Development and Limits of Privacy Enhancing Technologies in Data Analysis*. Technical Report. The Royal Society.
 - [103] Sarah Spiekermann and Lorrie Faith Cranor. 2008. Engineering privacy. *IEEE Transactions on software engineering* 35, 1 (2008), 67–82.
 - [104] Felix Stalder. 2002. The failure of privacy enhancing technologies (PETs) and the voiding of privacy. *Sociological Research Online* 7, 2 (2002), 25–39.
 - [105] Friedrich Steimann. 2018. Fatal abstraction. In *Proceedings of the 2018 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*. 125–130.
 - [106] Donald E Stokes. 2011. *Pasteur's quadrant: Basic science and technological innovation*. Brookings Institution Press.
 - [107] Herman T Tavani and James H Moor. 2001. Privacy protection, control of information, and privacy-enhancing technologies. *ACM Sigcas Computers and Society* 31, 1 (2001), 6–11.
 - [108] UN Privacy Preserving Techniques Task Team. 2020. *UN Handbook on Privacy-Preserving Computation Techniques*. Technical Report. <http://publications.officialstatistics.org/handbooks/privacy-preserving-techniques-handbook/UN%20Handbook%20for%20Privacy-Preserving%20Techniques.pdf>.
 - [109] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. Smart, useful, scary, creepy: perceptions of online behavioral advertising. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*. ACM, 4.
 - [110] Max Van Kleek, Ilaria Liscardi, Reuben Binns, Jun Zhao, Daniel J Weitzner, and Nigel Shadbolt. 2017. Better the devil you know: Exposing the data sharing practices of smartphone apps. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 5208–5220.
 - [111] Michael Veale, Reuben Binns, and Jef Ausloos. 2018. When data protection by design and data subject rights clash. *International Data Privacy Law* 8, 2 (2018), 105–123.
 - [112] Xiao Wang, Alex J. Malozemoff, and Jonathan Katz. 2016. EMP-toolkit: Efficient MultiParty computation toolkit. <https://github.com/emp-toolkit>.
 - [113] Stanley L Warner. 1965. Randomized response: A survey technique for eliminating evasive answer bias. *J. Amer. Statist. Assoc.* 60, 309 (1965), 63–69.
 - [114] Kang Wei, Jun Li, Ming Ding, Chuan Ma, Howard H Yang, Farhad Farokhi, Shi Jin, Tony QS Quek, and H Vincent Poor. 2020. Federated learning with differential privacy: Algorithms and performance analysis. *IEEE Transactions on Information Forensics and Security* (2020).
 - [115] Alma Whitten and J Doug Tygar. 1999. Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0.. In *USENIX Security Symposium*, Vol. 348. 169–184.
 - [116] Kathleen Broome Williams. 2012. *Grace Hopper: Admiral of the cyber sea*. Naval Institute Press.
 - [117] Aiping Xiong, Tianhao Wang, Ninghui Li, and Somesh Jha. 2020. Towards Effective Differential Privacy Communication for Users' Data Sharing Decision and Comprehension. *arXiv preprint arXiv:2003.13922* (2020).
 - [118] Andrew C Yao. 1982. Protocols for secure computations. In *23rd annual symposium on foundations of computer science (sfcs 1982)*. IEEE, 160–164.
 - [119] Xun Yi, Mohammed Golam Kaosar, Russell Paulet, and Elisa Bertino. 2012. Single-database private information retrieval from fully homomorphic encryption. *IEEE Transactions on Knowledge and Data Engineering* 25, 5 (2012), 1125–1134.
 - [120] Minhaz Zibran. 2008. What makes APIs difficult to use. *International Journal of Computer Science and Network Security (IJCSNS)* 8, 4 (2008), 255–261.