A Taxonomy of Privacy Salience Research

Meredydd Williams* and Jason R. C. Nurse

Department of Computer Science,
University of Oxford,
Oxford, UK
*meredydd.williams@cs.ox.ac.uk

Abstract Privacy is an intuitive concept in the physical world, with all of us needing some escape from the public gaze. However, while individuals might recognise a locked door as protecting one's privacy, they have difficulty practising equivalent actions online. Privacy salience considers the tangibility of this important principle, one which is often obscured in digital environments. By categorising existing work, we can both develop a deeper understanding of privacy and identify areas rich in research. Therefore, through surveying a range of prior studies, we iteratively construct a taxonomy of privacy salience. By coding research and identifying commonalities, we converge on several methodological categories, including surveys, tools and observations. These groups are further subdivided, presenting the range of techniques employed in privacy salience research. Exploring gaps in this taxonomy, we consider opportunities for further work, such as psychological analyses of privacy perceptions. We finish by discussing potential extensions to the taxonomy, based on growing trends in the field. It is through refining our understanding of this important topic that we can highlight the subject of privacy.

Keywords: Privacy salience · Taxonomy · Tangibility · Privacy paradox

1 Introduction

Privacy is a well-understood concept in the physical world. We all need some respite from the public gaze to enjoy our lives; indeed, it is essential to natural human development [8]. However, whereas individuals might intuitively consider a locked door as protecting one's privacy, they have difficulty practising equivalent actions online [17]. This can create a number of risks as users might be unaware of the digital dangers they face. 'Privacy salience' refers to how conspicuous the topic is in a specific environment. Risk in cyberspace is often intangible [23] and research [3] suggests reduced salience can lead to unwise decisions. Schneier [42] explained this intangibility could even contribute to the 'Privacy Paradox' [6], the disparity between what individuals claim about privacy and how they act. As technology permeates our society and we begin to live our lives 'online', privacy salience gains critical importance.

Previous research has considered the topic from a number of angles. Rao *et al.* [40] used an online survey to study privacy expectations across a range of websites. After analysing the portals' data management practices, they concluded identifying unexpected policies can increase privacy salience. Tsai *et al.* [45] modified search engine interfaces to highlight the results which respected privacy. Their analysis of 15,000 queries found that their alterations increased salience and led to more-private behaviour. Adjerid *et al.* [4] studied how the provision of additional privacy information could influence user actions. They discovered that a delay of only 15 seconds between a privacy notice and a decision could lead to less-private behaviour. While these analyses all concern salience, their methodological approaches differ in several respects. It was this fact which prompted our development of an extensive taxonomy, enabling the identification of further research opportunities.

The Oxford English Dictionary defines 'salience' to be "*[t]he quality or fact of being more prominent in a person's awareness or in his memory of past experience*" [36]. While privacy has been described as complex and ill-defined [1], we scope our definition to encompass informational privacy. Clarke [13] described this concept as "*the interest an individual has in controlling, or at least significantly influencing, the handling of data about themselves*". Therefore, we refer to 'privacy salience' as whether an individual is currently considering the topic of informational privacy. This differs slightly from 'privacy awareness', which we take to reflect long-term awareness of privacy, such as that which can be improved through educational campaigns. While this topic has been studied through a wide range of works [27,40], an extensive taxonomy has not yet been constructed.

We survey existing privacy salience literature, classifying studies based on their scope, methodology and underlying purpose. We select a range of varied research originating from computer science, cybersecurity and Human-Computer Interaction (HCI). Through identifying both commonalities and disparities, we iteratively formulate a number of methodological categories. These are further subdivided during taxonomy construction, before we explore those areas not yet thoroughly considered. For example, the increasing proliferation of smart devices could mask privacy issues, leading to the topic being further misinterpreted in the future. To mitigate this risk, privacy-protective Internet-of-Things (IoT) technologies could highlight the subject. To the best of our knowledge, we are the first to either survey privacy salience or construct a taxonomy of the subject, and our work looks to further understanding of online risk.

The remainder of our paper is structured as follows. Section 2 discusses our methodology, selection criteria and taxonomy design. Section 3 then explores our taxonomy in detail, highlighting previous literature of importance. In Section 4 we consider the research gaps identified through this process and discuss required future work. Finally, we conclude the paper in Section 5 and reflect on possible taxonomy extensions.

2 Methodology

Before detailed discussion of our methodology, we should define what we consider to be a 'taxonomy'. Oxford English Dictionary [37] defines a taxonomy as a "*particular system of classification*", and in this work we classify privacy salience research. De Hoog [19] explains how a taxonomy consists of three fundamental parts: representation, ordering and nomenclature. In terms of representation, the elements should be "maximally simple" and atomic in character. Categories should be arranged in a certain order and this order expressed through "character correlation". Finally, the product of logical ordering should be named formally, with minimal ambiguity. This aims to ensure that the structure is both accessible and usable. We incorporated these key principles into the development of our own taxonomy.

Constructing our taxonomy consisted of several stages. Firstly, previous research was surveyed to identify those works which directly concerned the topic. This was undertaken through Google Scholar, Scopus and Web of Science searches for the synonymous terms 'privacy salience' and 'privacy saliency'. While inspecting these articles, we identified further works of interest by considering references and citations. Continuing in this fashion, we studied other papers from surveyed authors to locate related articles. We also searched for papers through associated terms, such as 'privacy tangibility', before verifying these works represented the topic of salience. This stage was undertaken on a periodic basis to ensure that new papers would not be omitted from our analysis. While many works were selected through this process, several papers were rejected due to their purely-tangential relevance. Research originated from a number of fields, including computer science, cybersecurity and HCI, and ranged from 1999 to 2016. We believe this enables our taxonomy to comprehensively explore the important topic of privacy salience.

Secondly, these studies were coded based on their similarities, whether these were in domain, methodology or findings. Domains included web-based, mobile and stand-alone platforms, while methodologies ranged from tool development to system design recommendations. Although most findings agreed that salience encouraged private behaviour, some studies found that additional information increased data disclosure [22]. We proceeded through coding in an iterative fashion, gradually converging on several categories. This was crucial for the aforementioned 'ordering' principle [19], and we ensured similar elements were correctly grouped. These groups were eventually based on methodology, with some works studying the topic, while others looked to remediate the situation. This division was achieved organically, as method emerged as the most appropriate means of classification. While differences did exist in terms of domain, we preferred methodology to allow identification of further research opportunities.

Thirdly, where these groupings contained otherwise differing elements, further subdivisions were made. For example, although protective tools were distinct from surveys or observations, these applications used a range of techniques. While some looked to directly influence an individual, others took a neutral stance and presented relevant privacy information. For example, interfaces could

be annotated with warning lights [21], or details simply summarised into safety ratings [11]. Still others contextualised the topic by customising the tool with user data. This was seen in the case of the PrivAware application [7], which highlighted the information disclosed through web browsing. Subdivision continued until all papers collected in a single section exhibited similar methodologies. This was essential for the principle of 'representation' [19], as categories should be as atomic as possible.

Fourthly, we created representative formal names for our sections and subsections. The selection of nomenclature was an iterative process, refining definitions as categories evolved. Titles aimed to encapsulate commonalities in a group; for example, desk analyses, surveys and observations could all be considered types of Analysis. Using the hierarchical technique described by de Hoog [19], we appended superclass titles to subclass titles. For example, with preventative tools successively originating from the Tools and Remediation categories, they are titled 'Remediation:Tools:Preventative'. This contextualises an individual group, clarifies its position within the taxonomy, and complies with the important 'nomenclature' principle [19].

Finally, once the taxonomy was completed, we used the structure to identify research gaps in the field. This was undertaken by judging the number of articles in each subcategory and considering which approaches might have been overlooked. For example, while many tools promoted privacy salience, they seldom considered domains outside of browsers and smartphones. This presents opportunities to explore technologies which integrate with novel devices, such as wearables (e.g. smartwatches) and autonomous vehicles. Further details regarding research opportunities are discussed in Section 5.

3 Privacy Salience Taxonomy

We now discuss our privacy salience taxonomy, shown below in Figure 1, in detail. The figure omits full nomenclature for the purpose of brevity, and therefore 'Remediation:Tools:Illustrative' is presented as 'Illustrative'. We describe each category and subcategory in turn, our rationale behind such divisions, and the literature which inhabits each group.

At the highest level, we divided research between that which analyses the topic (**Analysis**), and that which aims to increase privacy salience (**Remediate**). Whereas articles in the former explore the prevalence and influence of salience, works in the latter attempt to improve the situation. Studies might be undertaken through a wide range of approaches, and, similarly, methods to increase salience are diverse. We first discuss the Analysis section and how its subcategories are composed.

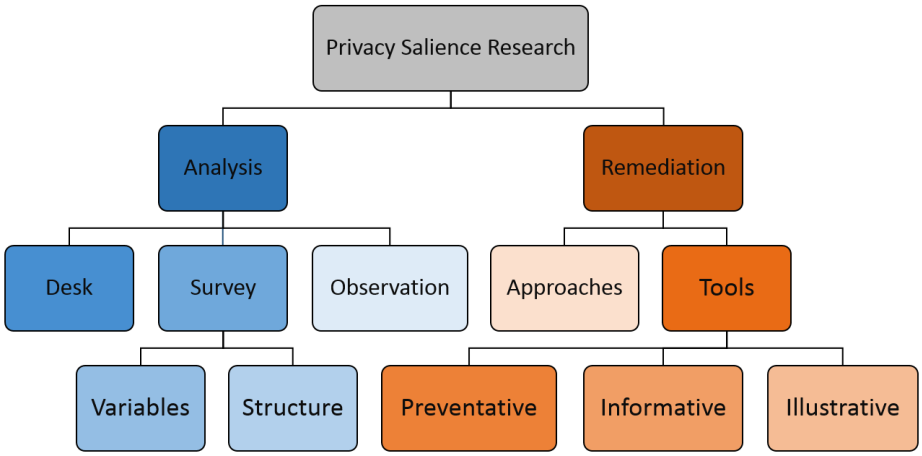

Fig. 1 Privacy salience research taxonomy

3.1 Analysis

The Analysis section concerns works which look to improve understanding of privacy salience. Rather than altering user behaviour or theorising new approaches, these articles explore the current situation. Although they do not aim to make privacy more tangible, findings from these papers inform the development of new tools and techniques. Through iterative construction of the taxonomy, we found analyses differed in several respects. For example, while some involved in-depth user experiments, others used online surveys to gauge privacy attitudes. Therefore, this category was further subdivided into three groups: Desk, Survey and Observation.

Analysis: Desk. This category concerns syntheses of existing research, including both systematic reviews and analyses of historical datasets. Rather than soliciting the opinions of new users, desk studies involve analysing previous work and deriving novel findings. This might be an academic pursuit, performing secondary research and identifying commonalities, or more practical by reinterpreting public domain data. These works do not directly engage with end-users, but might reinterpret a topic and promote it to a new audience. The taxonomy which we describe will exist within such a category, as it surveys the current state of privacy salience without conducting novel data collection.

Acquisti [2] discussed how 'nudges' can move privacy settings closer to individuals' expectations. Drawing on literature from behavioural economics and psychology, he discussed how contextualised decisions can reduce the disparity between concern and action. Rather than collecting new data, the article discusses prior work which might be of interest to privacy academics. While Schneier [42] also did not conduct novel experiments, his article communicated privacy salience to a wider audience. He surveyed a number of studies, all considered in

this taxonomy, before calling on social networking sites to improve their practices. By discussing privacy salience in a broader context, Schneier is able to provide a novel interpretation of existing research. Leech [29] briefly considered the topic through the lens of user interface design. He discussed the research of John *et al.* [24], and concluded that privacy messages should not be so overt as to discourage customers. While this might be a consideration for website designers, priming the topic of privacy should help protect ordinary users. Cichy and Salge [12] studied 35 years of privacy discourse in The New York Times. Through considering social norms and topic salience, they found perceptions to be susceptible to myopia and manipulation.

Analysis: Survey. Surveys question individual behaviour in a controlled and structured environment. These works can solicit participant opinions, collect sample demographics, or record self-reported actions. While surveys are often more expensive than desk studies, online questionnaires can canvass the perceptions of large numbers of individuals. These articles can themselves be further subdivided based on whether the structure of the survey itself is adjusted.

John *et al.* [24] conducted three conventional studies, with the former increasing salience and the latter two distracting participants from the topic. The first experiment included a consent warning, the second used a 'How BAD are U??' light-hearted survey, and the third questioned participants on ethics. In this case, the independent variable adjusted was that of question focus: playful in study two and serious in study three. They found that even if the risks of disclosure are low, people still refuse to disclose information when privacy concerns are primed. Statements which were intended to reassure participants had the opposite effect, reminding individuals of privacy risks. Tuunainen *et al.* [47] surveyed 210 Facebook users to understand the relationship between privacy awareness and data disclosure. Their study was undertaken through a conventional online questionnaire, with respondents self-reporting their social network behaviour. The survey structure was not adjusted at any stage, with participants simply providing their opinions. Individuals were found to be frequently unaware of post visibility and privacy policies, indicating that this topic was far from salient. Rao *et al.* [40] also used an online survey, studying privacy expectations across a range of websites. Site popularity and type were adjusted as independent variables, as the authors analysed standard practices and disclosure behaviour. Through their study of 240 participants, they found highlighting unexpected policies could increase privacy salience.

Adjusting the survey structure, Joinson *et al.* [26] increased salience by adding a "prefer not to say" option to web forms. They also analysed the effect of 'blurring', in which individuals could report their income in vague increments. They discovered that the inclusion of an opt-out decreased disclosure quantity, while the approximation of responses protected sensitive data. Brandimarte *et al.* [10] also varied the structure of their survey, with some individuals given greater control of information disclosure. They found that when participants perceived themselves as having greater agency, they were more willing to divulge

their personal data. Identified as the 'control paradox', the authors discussed how this might have parallels with aversions to air travel. Johnson *et al.* [25] conducted two online surveys to ascertain the opt-in/opt-out effect on privacy decisions. In both experiments, participants were randomly assigned questions in differing formats: some receiving opt-in questions while others faced opt-out. Through varying their survey structure, they found that participants agree 30% more often when opt-in is default. In this manner, the salience of privacy was strongly influenced by default framing. In a 2016 survey, Marreiros *et al.* [33] showed participants privacy documents before data disclosure. These documents either highlighted the protections of tech companies or discussed their negative activities. The authors found that individuals disclosed less even when positive behaviour was discussed, suggesting that salience affected action.

Analysis: Observation. Some studies are more empirical, collecting new data through live observations. Rather than analysing or reinterpreting existing data, these works conduct research through novel experiments. In contrast to surveys, observations involve greater flexibility and interaction between researchers and participants. These experiments might comprise a simulated environment, exploration of a novel system, or monitoring of user behaviour. Observations need not require direct surveillance in the sense of user observation, but involve the conduction of new experiments. While Observations might occasionally require application development, they differ from Tools as the system is used to analyse salience.

Houghton and Joinson [20] conducted a thematic analysis of their interviews concerning online friendship. They questioned 8 students on their digital communications, with results transcribed and divided into 18 privacy violation categories. They found that a lack of privacy salient information on social networks contributed to unwise disclosure decisions. In their analysis of Facebook behaviour, Stuzman and Kramer-Duffield [44] discovered that privacy discussions can encourage the tightening of online settings. By comparing a variety of opinions on the subject of privacy, the topic is made salient. Through questioning 444 college students and using logistic regression, they also found individuals do not appreciate how much information they present. Bonneau and Preibusch [9] evaluated 45 social networking sites to study how online policy and practice differed. They registered a number of online accounts and compared the information they received with how the platform functioned. The authors saw that when attention to privacy was increased, users disclosed less personal information.

Schaub *et al.* [41] conducted a study of privacy-protective extensions and how they influence user awareness. They recruited 24 participants for their qualitative lab study, who were tasked to think aloud while undertaking a series of online actions. Through their evaluation of the privacy extensions Ghostery, DoNotTrackMe and Disconnect, they found the tools increased the salience of tracking. However, overall privacy concerns were mitigated as participants felt protected by these browser plug-ins. Hui *et al.* [22] took an alternative approach, analysing how privacy statements and TRUSTe seals affected user behaviour. Through

their observation of 109 participants, they discovered that statements actually increased levels of data disclosure. Privacy seals were seen to have little effect, presenting that increases in salience do not always translate into improved behaviour. In this case, individuals were reassured by the presence of privacy policies, whether or not their content was actually taken into consideration.

3.2 Remediation

While some research aims improve our understanding, other works look to increase privacy salience. Although exploration of subtleties can be interesting, it is through remediation that individual privacy is protected. In this manner, analyses inform remediation and remedial efforts encourage a further degree of analysis. These approaches look to create a novel artefact, whether this is technological, socio-technical or academic. During taxonomy construction, we recognised that while some products are practical, others inform the development of future applications. Therefore, we divided Remediation into *Approaches* and *Tools*, with the former less direct than the latter.

Remediation: Approaches. Some research works are prescriptive, providing advice for future technologies. Rather than directly protecting individuals' privacy, these proposals inform technical developments. Recommendations use a number of techniques, ranging from psychological considerations to product design guidelines. While less practical than prototype tools, these articles ensure that future developments continue in appropriate directions.

Creese and Lamberts [17] considered how to best present online risk, examining visual and numerical techniques. Through their discussion of usability, visualisation and cognitive psychology, they provided advice for highlighting risk salience. They described how individuals might not appreciate numerical warnings and that graphical representations can be preferable. This work was later extended by Nurse *et al.* [35], who assessed the methods in which online risks may be communicated. The authors explored literature pertaining to the message content, the sender and the receiver, before presenting a number of recommendations. These included careful interface design, appropriate risk framing and timely message communication. Pötzsch [38] discussed the novelty of online platforms and the contrast between human cognition and machine resources. She considered that existing mental models might not be appropriate for digital interactions and this could place privacy in danger. To increase salience she advised that tools should be understandable to their target audiences and consider the cognitive boundaries that individuals possess.

Further design advice came from Lederer *et al.* [28], who presented five pitfalls for technical systems. Among their recommendations, they advised implementers to avoid obscuring information flows. As they wisely explained, "users can make informed use of a system only when they understand the scope of its privacy implications". Their guidelines should increase the salience and clarity of privacy. Cranor *et al.* [16] gave design advice for privacy agents: tools which alert users to potential risks. They described their 'Privacy Bird' system, which compares

website behaviour with pre-defined P3P preferences [15]. Their recommendations included providing access to privacy settings and increasing transparency to the end-user. It is through these approaches that privacy should become more accessible.

Remediation: Tools. Much research has concerned the development of technical tools. These works might themselves contain analyses or observations, but are conducted primarily to evaluate a system's performance. While aforementioned papers either study privacy salience or offer remedial advice, practical implementations can actually protect end-users. These tools are informed by Approaches and their performance evaluated through analytical works. Although Tools differ as a category from other research efforts, they themselves can be subdivided based on their technique.

Tools might be preventative, insomuch as they attempt to deter individuals from privacy threats. These works explicitly aim to protect end-users when privacy is perhaps not a salient topic. LaRose and Rison [27] conducted an assessment of user behaviour, concluding that explicit privacy warnings increase risk salience and reduce information disclosure. Through their considerations of social cognitive theory, they also found that the presence of privacy seals reduced participant concerns. Hughes-Roberts [21] annotated social media interfaces to highlight privacy, using coloured icons to emphasise disclosure risk. He found participants revealed less data when the topic was made salient, this having further implications for usability design. Ackerman and Cranor [1] discussed 'privacy critics', tools which present warnings in response to online actions. They described two prototype implementations: the first consulting a complaint database, and the second highlighting identifiability. Both products displayed large warning messages to deter unsafe actions, an approach which could increase the salience of privacy.

Alternatively, tools might be informative by simply presenting additional information. Not explicitly warning of online dangers, these technologies guide individuals by making privacy a salient topic. While these tools are more neutral in stance than preventative applications, their underlying purpose is still privacy protection. Tsai *et al.* [45] found that privacy indicators on search engines can lead to more-private behaviour. They conducted a 10-month study over 460 participants who used the Privacy Finder tool. The discovered that even when sites appeared near the bottom of results, they were more likely to be viewed if annotated as privacy-respecting. A shopping search engine from the same authors [46] annotated results with concise privacy policies. They analysed the activities of over 200 participants to understand whether privacy concerns influenced e-commerce. The researchers found that individuals preferred protective sites when privacy was salient, with some willing to pay a premium. Choe *et al.* [11] created visual representations of smartphone app privacy. Through conducting two studies with Amazon Mechanical Turk, they evaluated whether risk could be communicated effectively. They found that this visualised rating indeed influenced user perceptions, suggesting additional information increased the salience of privacy. Rajivan and Camp [39] conducted a similar study in 2016, comparing

app choices in the presence or absence of privacy icons. They discovered that priming the topic indeed increased participant concerns, with positive framing encouraging improved selections.

To further emphasise the risks of online environments, some platforms illustrate the potential results of user actions. By highlighting the consequences of unwise disclosures and lax software configurations, the topic of privacy becomes tangible. Although individuals might not find privacy salient in the abstract, they are more likely to show concern when it is contextualised [43]. Wang *et al.* [48] used 'nudges' to remind individuals of their Facebook audience. They conducted a 6-week study of 28 users, analysing their public messages and collecting their opinions. They found that illustrations of potential consequences could reduce unintended disclosure, likely due to increased privacy salience. Lipford *et al.* [31] developed a similar 'audience view', allowing users to observe their profiles as others do. Their HTML approach added tabs to the Facebook user interface, with 16 participants completing a number of tasks. They found that individuals better-understood the results of their online actions when the privacy was highlighted.

Almuhimedi *et al.* [5] implemented nudges on smartphones, developing an app which underlined data sharing between applications. They found that when the consequences of lax privacy settings were illustrated, 95% of participants reassessed their permissions. Over half of the sample went on to restrict their settings, suggesting salience influenced user behaviour. Melandrino *et al.* [32] developed a client-side tool to analyse data leakage through web browsing. They evaluated 180 popular sites, logging HTTP headers and saving HTML pages. By showcasing the range of information collected through third-party services, the salience of online privacy was increased. Becker and Chen [7] developed PrivAware, a tool aiming to reduce unintended social media disclosure. This application interacted with Facebook and was able to infer user attributes on almost 60% of occasions. The tool also gave recommendations for social networking privacy, increasing salience by contextualising the topic.

4 Research Gaps and Opportunities

While constructing our taxonomy, we identified several gaps in existing research. This was performed by considering which categories appear adequately addressed, and which could benefit from further work. As we iteratively refined our groups and differentiated between studies, we also saw where additional subcategories could form. We discuss several gaps and how these might be addressed through further research.

While salience has been considered through observations and surveys, there have been few comparative studies. Although research suggests risk tangibility is reduced by many platforms [17,23], the effects of these environments have not been compared. Individual works study the topic on social networking sites [44], web browsers [32] and mobile phones [11], but none conduct cross-platform analyses. This is particularly important with the increased development of the

Internet-of-Things (IoT), which can regard security and privacy as an afterthought [14]. As these novel smart devices begin to proliferate, privacy salience could be further diminished.

Although some effort has been invested in psychological research, we should further analyse decision-making processes. While understanding the prevalence of salience is important, this does not explain individuals' perceptions. Furthermore, even if technical tools can make the topic tangible, we require a theoretical understanding for progress to be repeatable. We believe studies should analyse users' thought processes, rather than make inferences from statistically-significant metrics. For example, cognitive science [17] might allow investigation of effective privacy communication. Li *et al.* [30] studied the role of affect in influencing disclosure decisions. They found that the initial appearance of a website can invoke emotions that condition later privacy behaviour. Mohammed and Tejay [34] took this approach even further, proposing the field of cognitive neuroscience. Through analysing brain activity while individuals interact online, we might be able to isolate privacy decisions. Through a variety of approaches, we could investigate how salience interacts with other factors, such as usability, security and familiarity.

We believe privacy salience surveys should incorporate empirical findings into their analyses. While surveys consolidate the views of many people, responses are generally self-reported. For this reason, it is often challenging to ascertain the veracity of participant claims. Individuals might not be consciously deceptive, but might have little awareness of what constitutes private behaviour. Williams and Nurse [49] conducted a street survey across the UK, questioning the public on their privacy opinions. To mitigate self-reporting biases, they included several optional demographic questions, using these as a proxy for data disclosure. They found that despite 92% of participants claiming to value their privacy, 99% revealed at least one piece of data needlessly. We encourage the conduction of similar surveys to ascertain whether privacy is salient or reported in response to social norms.

More comprehensive studies should be conducted of online activity. Thus far, much research has concerned social networking sites, both due to their popularity and the quantity of data disclosed. Facebook and Twitter have become ubiquitous in many people's lives, and APIs allow researchers to extract interesting information. While these platforms can show that users forget their privacy concerns, people visit a wide range of other sites. Whether purchasing goods or streaming media, individuals generate a large digital footprint. Privacy is unlikely to be a salient thought when using Amazon or Skype, and therefore significant risks might exist. For example, Amazon tracks all searches on their site, even if a customer neglects to purchase a product. Future work should consider privacy salience across a number of online environments, and develop tools which highlight disclosure regardless of platform.

Finally, while remedial approaches provide useful guidelines, many might not be appropriate for novel technologies. The popularity of mobile devices and wearable gadgets has led many away from traditional browser-based websites. Privacy salience might be further obscured by small screens, limited user interfaces and inaccurate mental models. Research has suggested the Privacy Paradox might be exacerbated by these new devices. Williams *et al.* [50] discussed how the unfamiliarity and heterogeneity of the IoT could contribute to unwise privacy decisions. They considered a number of Internet-of-Things idiosyncrasies, including resource constraints and ubiquitous data collection. In 2013, the UK government commissioned a usability study on several IoT heating devices, concluding none of the five market leaders offered sufficient user interfaces [18]. For these reasons, we believe guidelines should be developed to ensure privacy is tangible on novel devices. As the Internet-of-Things begins to proliferate, privacy salience takes on even greater importance.

5 Conclusions

In this paper we have surveyed privacy salience, a topic of great importance as we begin to live our lives 'online'. Through identifying commonalities in existing research, we iteratively constructed a taxonomy on the subject. While several papers analysed the current situation, others discussed the development of novel technologies. Categories were successively subdivided until we converged on a stable division of existing work. Considering gaps highlighted by this structure, we then discussed several research opportunities for further study. We believe that greater attention should be given to new platforms, particularly smart devices, and how these could impact our privacy perceptions. We also encourage the research community to explore the psychology behind privacy decisions, as a means of increasing topic salience. Risk can be obscured by online environments, and as technology pervades our lives, further work is required to protect our privacy.

There are several possible extensions to our taxonomy, depending on the content of future publications. While Approaches differ from Tools and Analyses, their recommendations originate from a variety of fields. Within our current structure, we have research from cognitive science [17], psychology [38] and user interface design [16]. Although present disparities do not warrant division, due to the sparsity of work in this area, distinct subcategories could be introduced in the future. We also found that Tools, whether Preventative, Informative or Illustrative, operate in a number of different domains. Some nudge smartphone users [5], others highlight social network privacy [48], and still others display web browser tracking [32]. As Privacy-Enhancing Technologies (PETs) develop over the coming years, the Tools section could be further subdivided based on domain. Finally, the Observation subsection could be split between simulations and lab-based studies. Currently we saw little need for division considering paper sparsity, but we predict privacy salience research will become increasingly popular as we live our lives 'online'.

References

1. Ackerman, M., Cranor, L.: Privacy critics: UI components to safeguard users' privacy. In: CHI'99 Extended Abstracts on Human Factors in Computing Systems. pp. 258–259 (1999)
2. Acquisti, A.: Nudging privacy: The behavioral economics of personal information. IEEE Security & Privacy 7(6), 82–85 (2009)
3. Adjerid, I., Acquisti, A., Loewenstein, G.: Framing and the malleability of privacy choices. In: Proceedings of the 13th Workshop on the Economics of Information Security (2014)
4. Adjerid, I., Acquisti, A., Brandimarte, L., Loewenstein, G.: Sleights of privacy: Framing, disclosures, and the limits of transparency. In: Proceedings of the Ninth Symposium on Usable Privacy and Security (2013)
5. Almuhimedi, H., Schaub, F., Sadeh, N., Adjerid, I., Acquisti, A., Gluck, J., Cranor, L.F., Agarwal, Y.: Your location has been shared 5,398 times!: A field study on mobile app privacy nudging. In: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems. pp. 787–796 (2015)
6. Barnes, S.: A privacy paradox: Social networking in the United States. First Monday 11(9) (2006)
7. Becker, J., Chen, H.: Measuring privacy risk in online social networks. In: Proceedings of the 2009 Workshop on Web (2009)
8. Berscheid, E.: Privacy: A hidden variable in experimental social psychology. Journal of Social Issues 33(3), 85–101 (1977)
9. Bonneau, J., Preibusch, S.: The privacy jungle: On the market for data protection in social networks. Economics of Information Security and Privacy (2010)
10. Brandimarte, L., Acquisti, A., Loewenstein, G.: Misplaced confidences privacy and the control paradox. In: The Ninth Workshop on the Economics of Information Security (2010)
11. Choe, E., Jung, J., Lee, B., Fisher, K.: Nudging people away from privacy-invasive mobile apps through visual framing. INTERACT 2013 in Lecture Notes in Computer Science 8119, 74–91 (2013)
12. Cichy, P., Salge, T.: The evolution of privacy norms: Mapping 35 years of technology-related privacy discourse, 1980-2014. In: 2015 International Conference on Information Systems (2015)
13. Clarke, R.: Introduction to dataveillance and information privacy, and definitions of terms. Tech. rep. (1999), `http://www.qatar.cmu.edu/iliano/courses/10F-CMU-CS349/slides/privacy.pdf`
14. Coates, M.: The Internet of Things will be vulnerable for years, and no one is incentivized to fix it. VentureBeat (2014), `http://venturebeat.com/2014/08/23/the-internet-of-things-will-be-vulnerable-for-years-and-no-one-is-incentivized-to-fix-it/`
15. Cranor, L.: P3P: Making privacy policies more useful. IEEE Security & Privacy 6, 50–55 (2003)
16. Cranor, L.F., Guduru, P., Arjula, M.: User interfaces for privacy agents. ACM Transactions on Computer-Human Interaction 13(2), 135–178 (2006)
17. Creese, S., Lamberts, K.: Can cognitive science help us make information risk more tangible online? In: Proceedings of the WebSci'09 (2009)
18. Department of Energy and Climate Change: Usability testing of smarter heating controls. Tech. rep. (2013), `https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/266220/usability_testing_smarter_heating_controls.pdf`

19. de Hoog, G.: Methodology of taxonomy. Taxon 30(4), 779–783 (1981)
20. Houghton, D., Joinson, A.: Privacy, social network sites, and social relations. Journal of Technology in Human Services 28(1-2), 74–94 (2010)
21. Hughes-Roberts, T.: Reminding users of their privacy at the point of interaction: The effect of privacy salience on disclosure behaviour. Human Aspects of Information Security, Privacy, and Trust in Lecture Notes in Computer Science 9190, 347–356 (2015)
22. Hui, K., Teo, H., Lee, S.: The value of privacy assurance: an exploratory field experiment. MIS Quarterly 31(1), 19–33 (2007)
23. Jackson, J., Allum, N., Gaskell, G.: Perceptions of risk in cyberspace. Edward Elgar (2005)
24. John, L., Acquisti, A., Loewenstein, G.: The best of strangers: Context dependent willingness to divulge personal information. In: The Best of Strangers. Pittsburgh, Carnegie Mellon-University (2009)
25. Johnson, E., Bellman, S., Lohse, G.: Defaults, framing and privacy: Why opting in-opting out. Marketing Letters 13(1), 5–15 (2002)
26. Joinson, A., Paine, C., Buchanan, T., Reips, U.: Measuring self-disclosure online: Blurring and non-response to sensitive items in web-based surveys. Computers in Human Behavior 24(5), 2158–2171 (2008)
27. LaRose, R., Rifon, N.: Promoting i-safety: Effects of privacy warnings and privacy seals on risk assessment and online privacy behavior. Journal of Consumer Affairs 41(1), 127–149 (2007)
28. Lederer, S., Hong, J., Dey, A., Landay, J.: Personal privacy through understanding and action: Five pitfalls for designers. Personal and Ubiquitous Computing 8(6), 440–454 (2004)
29. Leech, J.: Reassuring people about privacy is counter productive (2014), `https://mrjoe.uk/reassuring-people-privacy-counter-productive/`
30. Li, H., Sarathy, R., Xu, H.: The role of affect and cognition on online consumers' decision to disclose personal information to unfamiliar online vendors. Decision Support Systems 51(3), 434–445 (2011)
31. Lipford, H., Besmer, A., Watson, J.: Understanding privacy settings in Facebook with an audience view. In: Proceedings of the First Conference on Usability, Psychology, and Security. pp. 1–8 (2008)
32. Malandrino, D., Petta, A., Scarano, V., Serra, L., Spinelli, R., Krishnamurthy, B.: Privacy awareness about information leakage. In: Proceedings of the 12th ACM Workshop on Privacy in the Electronic Society. pp. 279–284 (2013)
33. Marreiros, H., Tonin, M., Vlassopoulos, M., Schraefel, M.: 'Now that you mention it': A survey experiment on information, salience and online privacy. In: CESifo Working Paper, Ludwig Maximilian University of Munich (2016)
34. Mohammed, Z., Tejay, G.: Re-examining the privacy paradox through cognitive neuroscience perspective. In: 21st Americas Conference on Information Systems (2015)
35. Nurse, J.R.C., Creese, S., Goldsmith, M., Lamberts, K.: Trustworthy and effective communication of cybersecurity risks: A review. In: First Workshop on Socio-Technical Aspects in Security and Trust. pp. 60–68. IEEE (2011)
36. Oxford English Dictionary: Salience (2016), `http://www.oed.com/view/Entry/170000`
37. Oxford English Dictionary: Taxonomy (2016), `http://www.oed.com/view/Entry/198305`
38. Pötzsch, S.: Privacy awareness: A means to solve the privacy paradox? In: The Future of Identity in the Information Society, pp. 226–236. Springer (2009)

39. Rajivan, P., Camp, J.: Influence of privacy attitude and privacy cue framing on android app choices. In: Twelfth Symposium on Usable Privacy and Security (2016)
40. Rao, A., Schaub, F., Sadeh, N., Acquisti, A., Kang, R.: Expecting the unexpected: Understanding mismatched privacy expectations online. In: Federal Trade Commission PrivacyCon Conference (2016)
41. Schaub, F., Marella, A., Kalvani, P., Ur, B., Pan, C., Forney, E., Cranor, L.: Watching them watching me: Browser extensions' impact on user privacy awareness and concern. In: NDSS Workshop on Usable Security (2016)
42. Schneier, B.: Facebook should compete on privacy, not hide it away (2009), `http://www.theguardian.com/technology/2009/jul/15/privacy-internet-facebook`
43. Shirazi, F., Volkamer, M.: What deters Jane from preventing identification and tracking on the web? In: Proceedings of the 13th Workshop on Privacy in the Electronic Society. pp. 107–116 (2014)
44. Stutzman, F., Kramer-Duffield, J.: Friends only: Examining a privacy-enhancing behavior in Facebook. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. pp. 1553–1562 (2010)
45. Tsai, J., Egelman, S., Cranor, L., Acquisti, A.: The impact of privacy indicators on search engine browsing patterns. In: Proceedings of the Fifth Symposium on Usable Privacy and Security (2009)
46. Tsai, J., Egelman, S., Cranor, L.F., Acquisti, A.: The effect of online privacy information on purchasing behavior: An experimental study. Information Systems Research 22(2), 254–268 (2011)
47. Tuunainen, V., Pitkänen, O., Hovi, M.: Users' awareness of privacy on online social networking sites - Case Facebook. In: BLED 2009 Proceedings (2009)
48. Wang, Y., Leon, P.G., Acquisti, A., Cranor, L.F., Forget, A., Sadeh, N.: A field trial of privacy nudges for Facebook. In: Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems. pp. 2367–2376 (2014)
49. Williams, M., Nurse, J.R.C.: Optional data disclosure and the online privacy paradox: A UK perspective. In: Fourth International Conference on Human Aspects of Information Security, Privacy and Trust at the 18th International Conference on Human Computer Interaction. pp. 186–197. Springer (2016)
50. Williams, M., Nurse, J.R.C., Creese, S.: The perfect storm: The privacy paradox and the Internet-of-Things. In: Workshop on Challenges in Information Security and Privacy Management at the 11th International Conference on Availability Reliability and Security (ARES). IEEE (2016)