IoT Refine: Making Smart Home Devices Accountable for Their Data Harvesting Practices

Max Van Kleek, William Seymour, Reuben Binns, Jun Zhao, Daniel Karandikar, Nigel Shadbolt

Department of Computer Science, University of Oxford, Oxford, UK, OX1 3QD
{*max.van.kleek, william.seymour, reuben.binns, jun.zhao, daniel.karandikar, nigel.shadbolt*} @ *cs.ox.ac.uk*

Keywords: smart home, privacy, data disclosure, visualisation, ambient display

Abstract

While smart home devices have the potential to improve people's lives by providing increased safety, security, and comfort, they also pose unprecedented privacy risks by having access to highly privileged aspects of people's lives. Already a complex concept, privacy is made more challenging in the smart home because devices are often designed to channel data to ad networks and other third parties unbeknownst to their users. In this paper, we propose a way to start to make smart home IoT devices accountable for their data collection, disclosure, and use practices by introducing the concept of a *privacy-empowering network disaggregator*. This disaggregator actively monitors and analyses all network traffic passing into and out of the home, helping to build a visual atlas that helps end-users understand such practices. We then describe the design and implementation of the first such privacy disaggregator, *IoT Refine*, demonstrating the feasibility and the potential for this approach towards addressing the privacy problem.

1 Introduction

Privacy has been widely recognised as one of the biggest challenges for smart home IoT technology. Embedded within the highly private context of the home, these devices gain access to some of the most sensitive aspects of people's lives and are capable of capturing large volumes of data from which behaviours and preferences can be inferred, stored, and exploited by manufacturers for purposes such as advertising and profiling. Unfortunately, these threats are not merely hypothetical; the prevailing model of surveillance capitalism [1] adopted by the technology industry has meant that most digital services, apps, and devices are designed out-of-the-box to maximise the potential for personal data exploitation at the cost of user privacy [2].

Many within the privacy community see this as a problem requiring an immediate solution. The most ardent privacy advocates propose eschewing smart devices entirely in favour of 'traditional', disconnected devices. However, such a strategy precludes use of many devices that can significantly improve people's lives by making their homes more safe, secure, and energy-efficient. Thus, other strategies are needed: ones which permit the use of home IoT devices without forfeiting one's own privacy.

One such approach has been to only allow IoT devices designed to respect their users' privacy, or at least to follow principles and guidelines such as Privacy by Design [3] or, more recently, Data Protection by Design and Default [4]. These guidelines specify at a high-level how data should be handled and used by such devices, including specific recommendations such as *data minimisation*[1]. When followed, these principles go far to curtail the ancillary data capture activities that pose the greatest threats to users' privacy. Unfortunately, few (if any) devices have emerged that adhere these principles, and so the impact thus far has been negligible. Another has been the creation of consumer advice guides, in which experts review data handling practices of devices to assess the relative risk(s) that each might pose. An example of such a guide is Mozilla's *Privacy Not Included* [6], which, like earlier projects such as *ToSDR: Terms of Service Didn't Read* [7], interprets privacy policies and provides an easy-to-understand, visual breakdown of the practices described in each device's privacy policy.

While a guide-based approach provides users with important information when choosing a device for their home, the potential value they provide is somewhat limited. Firstly, these guides are primarily based on what devices are said to do—that is, how they are described by the manufacturer—rather than on an actual analysis (e.g. through reverse engineering or operational analysis) of the device in question. This means that discrepancies, intentional or not, between the actual and reported operation of the device will not be known to potential users. The second problem is that these characterisations are static and not made with reference to specific configurations or firmware versions of the device—changes to which can lead to considerably different levels of privacy risk. Finally, many

[1] The requirement for devices to only capture, store, and use data in ways that directly support core functionality [5].

smart home devices function as "platforms" for other software, such as apps for smart TVs and skills for voice assistants, but the risks that third party software could introduce to these devices are not reflected in privacy guides as they often review unmodified version of the underlying system.

Thus, these four strategies highlight a need to suitably address the privacy risks posed by smart devices, especially those that change dynamically between configurations and over time. So how can users become better informed about how their devices operate in order to become empowered to make 'good' privacy-related decisions?

In this paper, we describe *IoT Refine*, a system that aims to give smart home users a high-level understanding of how their devices are sharing data with companies, as well as to start to understand for what purposes their data are being used. IoT Refine is adapted from earlier work on *X-Ray Refine*, a system designed to help people understand and "refine" their exposure to third party data collection via smartphone apps [8]. Unlike X-Ray Refine, IoT Refine is as an ambient situational awareness display, designed to operate continuously within a home setting. It is designed to work with a user's broadband router, and requires no configuration, operating purely by inspecting network traffic passing through the router from the home to the Internet and vice versa. In the following sections we describe IoT Refine, including ongoing and future work.

2 Background

Individuals' concerns about the collection and use of their personal information predate the era of the smart home. The notion of privacy as autonomy over what information is communicated, by whom, and when, was articulated by Alan Westin at the dawn of the digital computing era [9]. This work has exerted a strong and lasting influence over the design of privacy tools and interfaces which seek to set and enforce end-user preferences and rules. Alternative conceptions of privacy as a dynamic process of boundary negotiation [10, 11, 12], have also inspired contextually-aware design patterns for HCI [13] and cautioned against treating privacy preferences as persistent and universal [14, 15].

Recent studies in understanding the factors that influence individuals' willingness to share information have found a multitude of factors that relate not just to facts about the data and their handling, but to many personal factors around individuals and their situations. While the former include what data are disclosed, how long they are retained, and how they are used, the latter include personal experiences, cultural norms, understanding of risks, and even exposure to recent news articles about high-profile privacy violations (such as data breaches) [16, 17]. Other studies that have found the seemingly contradictory or paradoxical nature of privacy preferences have concluded that such incongruities may at least be partially explained by competing personal and situational factors [15, 8].

In light of the complex and personal nature of privacy preferences, many privacy researchers have turned away from systems that enforce normative privacy policies, towards those that seek to empower people to form and act upon their own priorities. Towards this end, *Privacy Leaks* by Balebako et al. revealed how personally identifying information (PII) was shared between apps and third parties [18], finding that users were particularly interested in surprising and unfamiliar destinations. Inspired by this work, we designed a visualisation called *Data Controller Indicators*, and ran a study that found that showing end-users information flows from apps to first and third parties allowed them to make more considered, informed, and consistent choices between smartphone apps [19].

Both Privacy Leaks and Data Controller Indicators provided highly granular and detailed information about individual apps and data sharing activities. While this detail might be useful for making a specific decision—such as deciding between similar apps or devices—it made it difficult for users to understand, at a high level, how their data were being processed, the purposes it was being used for, and with whom it was being shared. With this in mind, we designed *X-Ray Refine*, a visualisation aimed at allowing end-users to *refine* their existing data exposure to both first and third parties by giving them a high level overview of their data exposure. X-Ray Refine was made possible by creating data disclosure models of over 1 million Android apps using a static code analysis process [8].

2.1 Energy disaggregators

Another kind of system with similar goals to IoT Refine—to empower inhabitants of smart homes with useful information for making better-informed decisions—are energy meters. Specifically, visualisations of electricity consumption known as *energy disaggregators* have sought to provide end-user inhabitants with more relevant detail than standard smart meter energy-consumption displays that depict consumption totals in aggregate. Energy disaggregators, as their name suggests, instead provide a per-device breakdown of energy consumption, often accompanied trend analysis such as how specific devices are used throughout the day.

This relatively simple conceptual change has been seen as immensely empowering for end-users, allowing people to more easily identify opportunities for energy savings, such as by identifying the worst energy 'offenders' and providing opportunities to reschedule certain activities to off-peak times [20].

3 IoT Refine in Detail

The following section describes the design of IoT Refine, including its architectural and visualisation components.

3.1 Data Architecture: Monitoring Data Flows To and From the Home

At a high level, IoT Refine could be considered a *network traffic analyser for non-experts*. It works by performing shallow

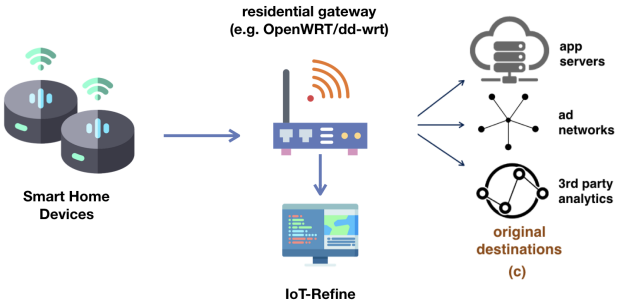

Fig. 1: Use within the smart home - IoT Refine works with the home's residential gateway to intercept traffic passing form the home to the external Internet.

packet inspection [21] on all traffic passing into and out of the home via the residential home gateway. As can be seen in Figure 1, all traffic through the router passes on to its usual destinations; therefore IoT Refine is a passive observer rather than a firewall. As it observes packets, IoT Refine intercepts and records the data in a relational database, as shown in Figure 2.

This database keeps an indefinite history of all packets that have passed through the router. Since this history can be extremely large, only a small amount of data is stored per packet: source and destination IP addresses, MAC address of the local device, timestamp, packet length, and protocol.

As packet signatures are captured, a separate process performs data enrichment, identifying any new sources and destinations in incoming packet signatures. For sources and destinations representing local devices, the process attempts to identify what kind of device it is (using MAC address vendor resolution), adding it to the *devices* table in the database. (This table can be updated later by the user to assign user-friendly names to devices.)

For external sources and destinations, the enrichment process does two things: it attempts to identify the data controller (e.g. organisation, company, or owner) using WHOIS lookups [22] of the IP address, and, secondly, its geographic location/jurisdiction, through GeoIP resolution [23]. The main goal of this is to abstract the data to a level more closely with the logical entities that people reason with; that is, users are likely to care much more about the *companies* receiving their data than the specific endpoints that their traffic is going to. Geographic destinations and jurisdictions may help end-users to understand if they have jurisdictional data rights under regional data protection schemes.

3.2 Front-end Visualisation

The front-end of IoT Refine comprises three sub-displays containing the following visualisations:

1. *Exposure per company* - indicates how much data has been sent by devices to each company in a coloured stacked bar chart sorted by companies from those receiving the most data to the least, with each stack representing a distinct device (see Figure 3).
2. *Geographic destinations* - comprises a world map with all geo-resolved destinations of data and indicators sized proportional to how much data each has received (see Figure 4).
3. *Jurisdictions* - consists of a stacked bar chart indicating the jurisdictions within which the destination of traffic falls, sorted by from most to least traffic volume (see Figure 4).

Together, these displays collectively display how much data has been exchanged between specific devices and particular companies, and the location(s) within which the endpoints fall. All visualisations are updated from a single data model, which is, in turn, updated in real-time through an event stream posted by the back-end.

The interface is designed, by default, to show the total data exchanged over the past 24 hours; however, this can be interactively changed to shorter intervals (e.g. past hour, or past week). This will be extended to support arbitrary time-based queries and interactive scrubbing (described in Future Work).

4 Evaluation and Future Work

The IoT Refine prototype was developed and iteratively usability tested with discount usability heuristic evaluation methods with researchers in the lab. It was then installed in a prototype smart home testbed known as *Barratt House* at the UK's Building Research Establishment, as shown in Figure 5.

There are plans to extend IoT Refine in three key areas. The first pertains to greatly enhancing the front-end to support additional sensemaking and exploratory data analysis. The second planned set of work pertains to extending the analytical capabilities of the back-end, while the third consists of a field study of IoT Refine.

4.1 Sensemaking support

In its current state, IoT Refine is designed as a largely non-interactive ambient display. However, preliminary user testing showed that people who saw the display were often interested in finding out more about specific aspects of what they saw in the visualisation.

Thus, we intend to add functionality to the front end to support this exploratory data analysis [24] and sensemaking [25]. We have written previously on how the formation of privacy preferences by individuals could be seen as a sensemaking process, suggesting the potential for sensemaking interfaces for supporting the formation of privacy preferences [26]. To this end we think it would be useful to support exploration by porting the company view-facilities from X-Ray Refine, making it possible to select the name of a company to find out associated jurisdictions, endpoints, and disclosures. Additionally, we wish to

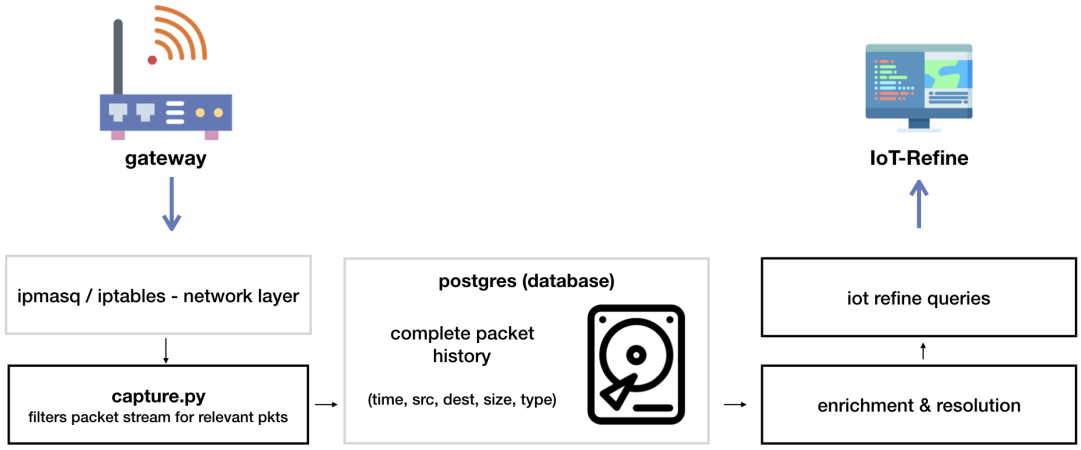

Fig. 2: Detailed Architecture and Data Flow - IoT Refine follows a *capture* ⇒ *categorise* ⇒ *process* flow. Metadata about network traffic passing through the built in WiFi hotspot are captured using Tshark (the command line component of Wireshark), before being stored in a PostgreSQL database. In the categorisation phase, company ownership and geographic location is retrieved and stored for each observed IP address. As a final step, standing queries/views on the database and notification change triggers provide the IoT Refine front-end (an Angular and D3 Javascript app) with real time updates to be rendered.

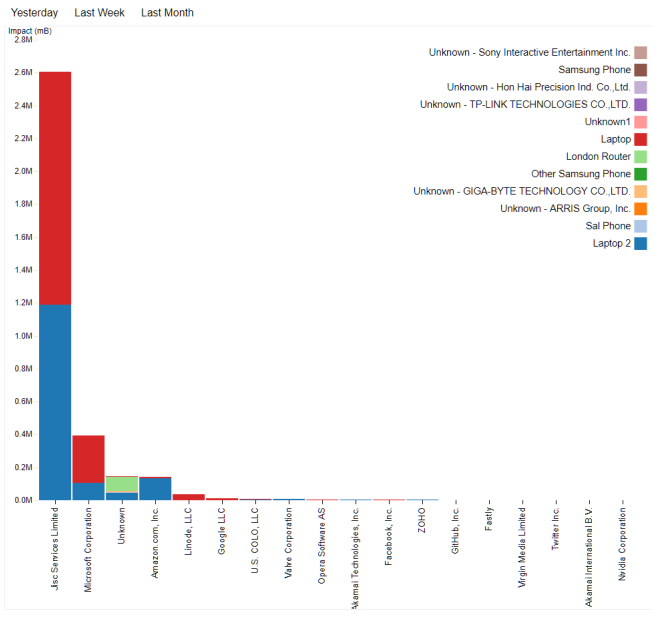

Fig. 3: Traffic Flows in IoT Refine, sorted by company

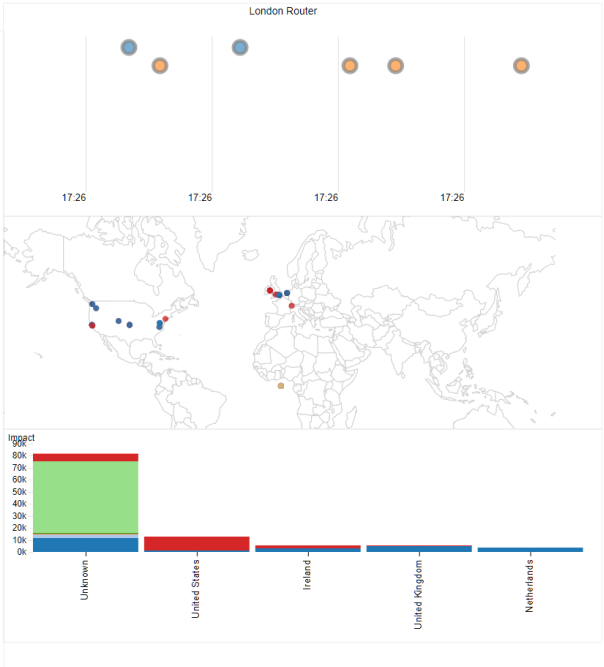

Fig. 4: Traffic Flows in IoT Refine, sorted by country

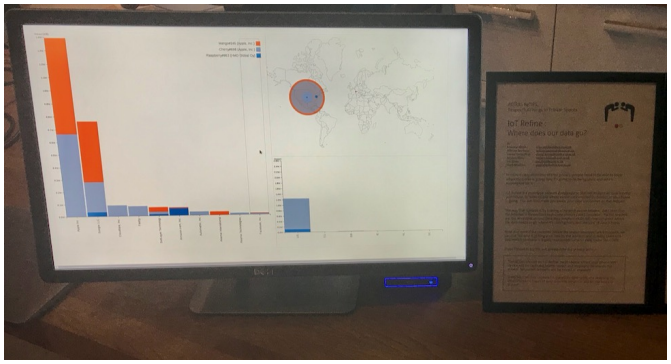

Fig. 5: IoT Refine installed at a Smart Home testbed - IoT Refine was installed as the world's first smart home privacy network disaggregator in Barratt House, a smart home testbed at the Building Research Establishment (BRE) in Watford, UK.

explore the use of statistical trends at various temporal scales such as to bring greater interpretability to disclosure patterns, such as being able to easily see the effects of software updates or the use of specific features.

4.2 Back-end analytical capability

To support sensemaking on the front end, one would ideally be able to allow users to "see" *what* data was disclosed by each device to each destination. However, due to the fact that many devices (fortunately) send their traffic end-to-end encrypted, it is impossible to simply capture and read the contents of the packet. We have demonstrated, however, that pattern-analysis of packet metadata alone can reveal more about about the data being disclosed. For instance, informal experiments we ran revealed that supervised machine learning techniques can be used to identify device behaviours, from voice assistants retrieving weather reports to light bulbs being switched on.

4.3 Evaluation and use as experimental testbed

Ultimately, the test of whether the approach taken in IoT Refine does help to alleviate smart home privacy concerns for users will require evaluating it in the field with real users and homes. We plan on using IoT Refine as a *technology probe* [27] with the aim of understanding the potential for user-empowering interfaces like it, as well as inherent challenges, through field-testing with families starting the second quarter of 2019.

5 Conclusions

The challenge of helping end-users stay in control of their privacy in the smart home is one the greatest, due to both the ever-expanding capabilities of these devices as well as the sensitive nature of the home as a private space. In this paper we have described IoT Refine, a prototype privacy network disaggregator for smart homes that aims to help people become more aware of how the devices in their home are sending and receiving data. Whilst some devices are designed to talk only to their manufacturers, many others contain third party libraries which disclose data directly from devices to third parties. It is for this latter kind of device that IoT Refine will provide the greatest benefit, by identifying such disclosure activities directly in an easy-to-read display.

We have demonstrated the potential uses for this type of basic disaggregator, as well as how one can feasibly be built today. Although informal, the user tests conducted to date indicate that the information displayed by IoT Refine is both interesting and much-needed—users feel woefully uninformed about the data handling practices of their devices. We have briefly outlined plans on how we wish to continue this work, developing IoT Refine into a fully-fledged sensemaking tool, as well as completing a formal evaluation of its use in the field. Finally, the full source code and documentation for IoT Refine is made available under an open source license, with the aim of helping privacy researchers and enthusiasts alike[2].

Acknowledgements

This work was supported by *ReTiPS: Respectful Things in Private Spaces*, a project funded through the PETRAS IoT Hub Strategic Fund, which, in turn, was funded by the UK Engineering and Physical Sciences Research Council (EPSRC) under grant number N02334X/1.

References

[1] S. Zuboff, "Big other: surveillance capitalism and the prospects of an information civilization," *Journal of Information Technology*, vol. 30, no. 1, pp. 75–89, 2015.

[2] A. Hern, "UK homes vulnerable to 'staggering' level of corporate surveillance." The Guardian. `https://www.theguardian.com/technology/2018/jun/01/uk-homes-vulnerable-to-staggering-level-of-corporate-surveillance`, June 2018. [Online; accessed 2019-01-04].

[3] A. Cavoukian, *Privacy by design in law, policy and practice: a white paper for regulators, decision-makers and policy-makers*. Information and Privacy Commissioner of Ontario, Canada, 2011.

[4] ICO, "Data protection by design and default." `https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/accountability-and-governance/data-protection-by-design-and-default/`, 2018. [Online; accessed 2019-02-02].

[2] Available at `https://github.com/OxfordHCC/IoT-Refine`

[5] ICO, "Principle (c): Data minimisation." https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/principles/data-minimisation/, 2018. [Online; accessed 2019-02-02].

[6] Mozilla, "Privacy not included." https://foundation.mozilla.org/en/privacynotincluded/, 2018. [Online; accessed 2019-02-02].

[7] OpenCollective, "Terms of service; didn't read." https://tosdr.org/about.html, 2012. [Online; accessed 2019-02-02].

[8] M. Van Kleek, R. Binns, J. Zhao, A. Slack, S. Lee, D. Ottewell, and N. Shadbolt, "X-ray refine: Supporting the exploration and refinement of information exposure resulting from smartphone apps," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, p. 393, ACM, 2018.

[9] A. F. Westin, "Privacy and freedom," *Washington and Lee Law Review*, vol. 25, no. 1, p. 166, 1968.

[10] I. Altman, "The environment and social behavior: Privacy, personal space, territory, and crowding.," 1975.

[11] L. Palen and P. Dourish, "Unpacking "privacy" for a networked world," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '03, (New York, NY, USA), pp. 129–136, ACM, 2003.

[12] H. Nissenbaum, "Privacy as contextual integrity," *Wash. L. Rev.*, vol. 79, p. 119, 2004.

[13] S. Lederer, J. I. Hong, A. K. Dey, and J. A. Landay, "Personal privacy through understanding and action: five pitfalls for designers," *Personal and Ubiquitous Computing*, vol. 8, no. 6, pp. 440–454, 2004.

[14] L. Barkhuus, "The mismeasurement of privacy: using contextual integrity to reconsider privacy in hci," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 367–376, ACM, 2012.

[15] K. Martin and K. Shilton, "Putting mobile application privacy in context: An empirical study of user privacy expectations for mobile devices," *The Information Society*, vol. 32, no. 3, pp. 200–216, 2016.

[16] P. G. Leon, A. Rao, F. Schaub, A. Marsh, L. F. Cranor, and N. Sadeh, "Privacy and behavioral advertising: Towards meeting users preferences," in *Proceedings of the Symposium on Usable Privacy and Security*, 2015.

[17] P. G. Leon, B. Ur, Y. Wang, M. Sleeper, R. Balebako, R. Shay, L. Bauer, M. Christodorescu, and L. F. Cranor, "What matters to users?: factors that affect users' willingness to share information with online advertisers," in *Proceedings of Symposium on Usable Privacy and Security*, pp. 1–7, ACM, 2013.

[18] R. Balebako, J. Jung, W. Lu, L. F. Cranor, and C. Nguyen, "Little brothers watching you: Raising awareness of data leaks on smartphones," in *Proceedings of the Symposium on Usable Privacy and Security*, p. 12, ACM, 2013.

[19] M. Van Kleek, I. Liccardi, R. Binns, J. Zhao, D. J. Weitzner, and N. Shadbolt, "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*, CHI '17, (New York, NY, USA), pp. 5208–5220, ACM, 2017.

[20] A. Schoofs, A. Sintoni, G. M. O'Hare, and A. G. Ruzzelli, "Demo abstract: appliance load monitoring by power load disaggregation," in *2010 7th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, pp. 1–2, IEEE, 2010.

[21] C. Parsons, *Deep Packet Inspection in Perspective: Tracing its lineage and surveillance potentials*. Citeseer, 2008.

[22] J. Gargano and K. Weiss, "Whois and network information lookup service, whois++," tech. rep., 1995.

[23] S. Siwpersad, B. Gueye, and S. Uhlig, "Assessing the geographic resolution of exhaustive tabulation for geolocating internet hosts," in *International Conference on Passive and Active Network Measurement*, pp. 11–20, Springer, 2008.

[24] J. W. Tukey, *Exploratory data analysis*, vol. 2. Reading, Mass., 1977.

[25] P. Pirolli and S. Card, "The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis," in *Proceedings of international conference on intelligence analysis*, vol. 5, pp. 2–4, McLean, VA, USA, 2005.

[26] M. Van Kleek, W. Seymour, M. Veale, R. Binns, and N. Shadbolt, "The need for sensemaking in networked privacy and algorithmic responsibility," in *Sensemaking in a Senseless World: Workshop at ACM CHI18, 22 April 2018, Montréal, Canada*, ACM Conference on Human Factors in Computing Systems (CHI'18), 2018.

[27] H. Hutchinson, W. Mackay, B. Westerlund, B. B. Bederson, A. Druin, C. Plaisant, M. Beaudouin-Lafon, S. Conversy, H. Evans, H. Hansen, *et al.*, "Technology probes: inspiring design for and with families," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 17–24, ACM, 2003.