Poli-see: An Interactive Tool for Visualizing Privacy Policies

Wentao Guo
Pomona College
wgtc2015@mymail.pomona.edu

Jay Rodolitz
Pomona College
jbra2017@mymail.pomona.edu

Eleanor Birrell
Pomona College
eleanor.birrell@pomona.edu

ABSTRACT

Prior work has shown that current privacy policies fail to effectively implement informed consent. This work investigates how data use practices might be conveyed by a graphical representation. We present Poli-see, an interactive tool for visualizing privacy policies. We then describe the results of an in-person user study ($n = 24$) and an online study ($n = 600$) that evaluate how well Poli-see conveys information about data use practices. In our in-person study, we found that participants answered factual questions about privacy policies more accurately when shown a Poli-see representation than when shown an annotated text representation. In our online study, we found that participants who were shown a Poli-see representation reported higher levels of enjoyment and higher likelihood of looking at the policy than participants who were shown a conventional text representation or an annotated text representation. These results suggest that graphical representations might be useful for conveying data use practices to users, but that further research and refinement will be required before graphical representations can be effectively deployed in real-world systems. We conclude by identifying key advantages and challenges for graphical representations of privacy policies drawn from our experience.

CCS CONCEPTS

• Social and professional topics → Privacy policies; **• Security and privacy → Usability in security and privacy**; **• Human-centered computing → Information visualization**; *User studies*; *Graph drawings*; *Empirical studies in visualization*; *Empirical studies in interaction design.*

KEYWORDS

privacy policies, visualization, usable privacy

ACM Reference Format:
Wentao Guo, Jay Rodolitz, and Eleanor Birrell. 2020. Poli-see: An Interactive Tool for Visualizing Privacy Policies. In *19th Workshop on Privacy in the Electronic Society (WPES '20), November 9, 2020, Virtual Event, USA.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3411497.3420221

WPES '20, November 9, 2020, Virtual Event, USA

ACM ISBN 978-1-4503-8086-7/20/11…$15.00
https://doi.org/10.1145/3411497.3420221

1 INTRODUCTION

Recent trends towards increased personalization have created economies of personal data, in which data about individual users are collected, used, exchanged, and sold. The only legal limitations on how these economies operate on user data come from data protection regulations (e.g., the European Union's General Data Protection Regulation [8]) and standards (e.g., the Federal Trade Commission's Fair Information Practice Principles [9]); these regulations authorize data use only under certain conditions, including when individuals provide informed consent to the use of their data. Today, informed consent is implemented through the practice of notice and consent: companies publish a privacy policy (typically a long text document) describing their data use practices, and people who use their service are considered to have consented to the terms of the privacy policy.

However, fifteen years of critiques have provided compelling evidence that current privacy policies fail to adequately implement informed consent. People do not read privacy policies [2, 14, 23], the number of policies and frequency of updates make it infeasible to keep up to date [4, 17], and policies are difficult to understand [18, 22, 25]. To address these problems, alternate representations of privacy policies have been designed [15, 19, 27, 29]. However, despite evidence that visualizations can effectively convey actual data collection and dissemination [5, 10], there is limited work investigating tools for visualizing privacy policies.

Our aim in this work is to address this gap through the design and implementation of Poli-see, the first tool for generating interactive, graphical visualizations of privacy policies. Poli-see uses a graph-like representation to depict data collection by a company and to show how this data and derived data flow to third parties. Collected data types are denoted by icons—with detailed information about data use provided when a user hovers over an icon—and data inference, sharing, and sale are depicted with arrows.

To evaluate the usability of Poli-see, we conducted an in-person user study with 24 participants and an online study with 600 participants. In each study, we quantitatively compared Poli-see to alternative, non-graphical representations using four different metrics: user accuracy, user confidence, user speed, and quality of user experience. We consistently found that Poli-see either matched or outperformed alternatives on all four metrics. In our in-person study, we found that participants answered factual questions about privacy policies more accurately and with more confidence when shown a Poli-see representation than when shown an annotated text representation. After revising Poli-see's design in response to qualitative feedback, we conducted an online study in which we evaluated the graphical representations generated by Poli-see compared to conventional text representations and to annotated text representations. We found that participants answered simple questions more accurately with Poli-see than with other representations, although the opposite was true for complex questions. Participants also reported greater enjoyment from interacting with Poli-see and a higher likelihood of looking at policies if real-world policies were represented in this way.

We conclude that the small quantitative improvements demonstrated by Poli-see, counterbalanced by the ongoing challenges presented by graphical policy representations, are not yet sufficient to justify widespread deployment of this tool. Nonetheless, we believe that this work constitutes a critical first step towards the development and deployment of graphical policy representations.

We view the primary contributions of this work to be:

(1) the development of the first interactive, graphical representation of privacy policies,
(2) an evaluation of Poli-see that demonstrates the potential for graphical representations to accurately convey data use practices to users in a manner that encourages user engagement, and
(3) the discovery of insights into the advantages and challenges of graphical policy representations.

We believe that this work validates the idea of using graphical representations to improve usability of privacy policies, and that it can serve as a foundation for future research on graphical privacy policy representations.

2 RELATED WORK

To address the shortcomings of privacy policies, a large body of work has been devoted to designing and implementing *transparency-enhancing tools* that might better convey an organization's data use practices to data subjects.

Privacy Policy Representations. One approach to enhancing the transparency of data use practices is to develop an alternative privacy policy representation that conveys information to users more effectively than conventional text policies. This approach, which is adopted in this work, has previously led to several alternative policy representations.

Annotated text privacy policies [27, 30] are text policies augmented with interactive annotations designed to make it easier for a user to find particular information about an organization's data use practices. Segments of text are highlighted in different colors based on the type of information that segment contains (e.g., "First Party Collection/Use," "Third Party Sharing/Collection," or "User Choice/Control"); buttons to the left of the text indicate the meaning of different colors and allow users to filter policies on a specific type of information. A color-coded scrollbar provides an overview of how types of information are dispersed through the policy. Annotations are automatically generated for a text policy using a machine learning model that was trained on privacy policies annotated by domain experts. This annotated privacy policy representation has not yet been evaluated in a published user study.

Polisis [13] is a privacy policy analysis framework comprised of algorithms that divide policies into segments and label each segment with relevant data use practices. Users can interact with the policy representation by sending natural-language questions to a chatbot named PriBot. The authors found that PriBot was able to answer privacy-related questions more accurately than prior automated alternatives. In a large-scale user study, PriBot's responses were rated as relevant more often than responses generated by the baseline tools. As part of the Polisis project, researchers developed an online exploratory tool for privacy policies that visualizes types of data collected, reasons for data collection, third-party sharing, and more.

The Privacy Policy Visualization Model [11] is a framework for graphical policy representations. Policy representations in this model are diagrams that connect personal data and relevant entities, including the user and the data collector. Solid and dashed lines indicate different kinds of relationships between elements. A heuristic evaluation by experts suggests that the model is adequate for providing notice and raising awareness of privacy policy information, but no tool has been implemented to generate policy representations, and this representation has not been evaluated in a user study.

Privacy Policy Summaries. Rather than representing a full policy, some work attempts to convey information about a privacy policy by summarizing key data use practices.

Privacy Nutrition Labels [15, 16] were an early effort to address the length and complexity of text privacy policies by providing a summary, called a privacy label, which conveyed key data use practices. These labels were automatically generated from machine-readable P3P descriptions of the privacy policy. In the initial user study and in a large-scale follow-up study, participants were able to answer questions about privacy policies more accurately and more quickly using privacy labels than using text privacy policies. Participants rated the privacy label as easier to use and more enjoyable to interact with.

Privacy Icons [19] are comprised of four icons that convey an organization's key data use practices: how long it retains personal data, whether the organization shares data with advertisers, whether third-parties can use data generally or only to help the organization fulfill the intended transaction, and whether the organization complies with voluntary requests for data from law enforcement.

Privacy Bird [29] is a browser plug-in that compares a user's stated preferences to the P3P-defined data use terms of a privacy policy. An icon in the browser informs the user whether the policy matches their preferences: a green bird indicates a match, a red bird indicates a conflict between the organization's policy and the user's preferences, and a yellow bird indicates that the organization has not published a P3P policy.

Ex-Post Visualization Tools. Instead of representing privacy policies, some work focuses on showing users what data have been collected and how those data have been shared since a user started using an application; tools that adopt this approach are termed *ex-post* tools. To date, the majority of work on privacy-related visualizations has been conducted in the context of ex-post tools.

Privacy Panel [7] is a mobile app that tracks data disclosures by other apps and visually represents how much information each app shares about the user's location data, content (such as pictures or emails), and contacts. Privacy Panel also provides detailed visualizations for each app, including a chart depicting the frequency of data accesses over time and a heatmap displaying specific locations collected.

The Data Track [1, 10] is an ex-post visualization tool designed to inform data subjects about data that have been collected about them by different organizations across the web. The Data Track presents a network-like visualization that depicts which types of data each organization has collected. Users can also view details regarding specific pieces of their data—including which organizations have collected that data—and can request corrections or deletions. In a user study, researchers found that all 17 participants were able to use the Data Track successfully to identify which of their data had been collected by a particular organization, as well as which organizations had collected a particular piece of data.

PrivacyInsight [5] is an ex-post tool that uses a graph-based visualization to depict the flow of personal data into, through, and out of an organization. Nodes in the graph represent different organizational units: data sources (such as the user), the data collector, and data sinks (such as third parties with whom personal data have been shared). Next to each node is a group of visual icons indicating what kinds of personal data are held by the organizational unit. Links in the form of arrows between nodes represent data flow. Users can also view the flow of a single piece of data in isolation and can request that data be deleted or corrected. In a user study, researchers found that participants completing tasks in a fictional scenario were more able to correctly find information using PrivacyInsight than using the Data Track or a JSON document. Eye-tracking data collected during the same study demonstrated that the graph-based visualization was effective at directing users' attention to the intended spots in the interface.

3 DESIGN

We identified three goals that guided the design of our privacy policy visualization tool:

(1) **Expressive:** The graphical representations generated by the tool should be able to depict complete data use practices as described in current conventional text privacy policies.
(2) **Usable:** Users should be able to accurately and confidently understand the information conveyed by a graphical representation quickly, even if they have never seen this type of representation before. Their experience using the tool should be positive.
(3) **Scalable:** The tool should make the process of generating a privacy policy visualization as simple and scalable as possible. The process for creating a graphical representation from an existing text privacy policy should require minimal human effort.

Following these goals, we designed Poli-see, a JavaScript application for visualizing privacy policies. We used D3.js, an open-source JavaScript library that can apply data-driven transformations to elements on a webpage. This allows Poli-see to generate graph-based visualizations of privacy policies from JSON files encoding the contents of a privacy policy. In Subsection 3.1, we describe our process of iteratively developing Poli-see and seeking user feedback. In Subsection 3.2, we describe the current design of Poli-see.

3.1 Design Process

We began the process of designing Poli-see by closely reviewing ten current privacy policies published by companies in various sectors. During this review we observed several patterns in how data use practices are described:

(1) Current privacy policies describe both what data are collected and how these data are used. An expressive graphical policy representation will therefore need to feature both collection and use.
(2) Current privacy policies describe data use practices both for data collected from the data subject and for derived data inferred by the data collector. An expressive graphical policy representation will therefore need to convey both the inference of additional data types and the data use practices for these derived data.
(3) Current privacy policies use a variety of different terms to describe the same types of data or the same types of uses. This variation in language has been previously noted as a likely barrier to comprehension of privacy policies [12]. A usable graphical policy representation will therefore need to employ standardized terminology for depicting data types and data uses.
(4) Current privacy policies typically describe many data use practices non-specifically, referring to "all data we collect." Where necessary, policies then provide details and exceptions to those practices for specific data types, such as location data. A usable graphical policy representation will therefore need to clearly convey both non-specific and specific data use practices.

We used these observations, along with insights drawn from existing ex-post visualization tools, to create a preliminary design for Poli-see. We then built a functional prototype as a JavaScript web application.[1] We conducted three rounds of one-on-one design pilot studies at Pomona College to gather feedback on the preliminary design of Poli-see and to refine our user study design. After each round of pilots, we redesigned Poli-see based on feedback. The first two rounds consisted of three pilots each; the third consisted of two. We recruited students at our institution and nearby institutions through Facebook groups. Pilots lasted 30 minutes, and participants were compensated $5. During the pilots, participants were introduced to Poli-see and shown a visualization for Strava, a fitness-tracking and social network app, which we anonymized using the name buildUp. We asked participants to complete some information-finding tasks and then interviewed them to gather open-ended feedback.

We then conducted a 24-participant in-person user study (described in Section 4). Following this in-person study, we again revised the design of Poli-see to address qualitative feedback from in-person study participants, resulting in the current Poli-see design.

3.2 Current Design

The current design of Poli-see draws inspiration from the design of current ex-post visualization tools (the Data Track and PrivacyInsight), but we focused on incorporating design elements informed

[1] We chose to use an implemented prototype for our preliminary design pilots rather than a low-fidelity prototype for two reasons. First, due to the large quantity of information expressed in privacy policies, we decided that an interactive design, in which much of the information is visible only when a user takes particular actions, was necessary for a usable visualization. Second, our graphical representation diverges significantly from representations that participants were likely familiar with. For both of these reasons, we believed that low-fidelity prototypes would not reliably convey the feeling of this interface.

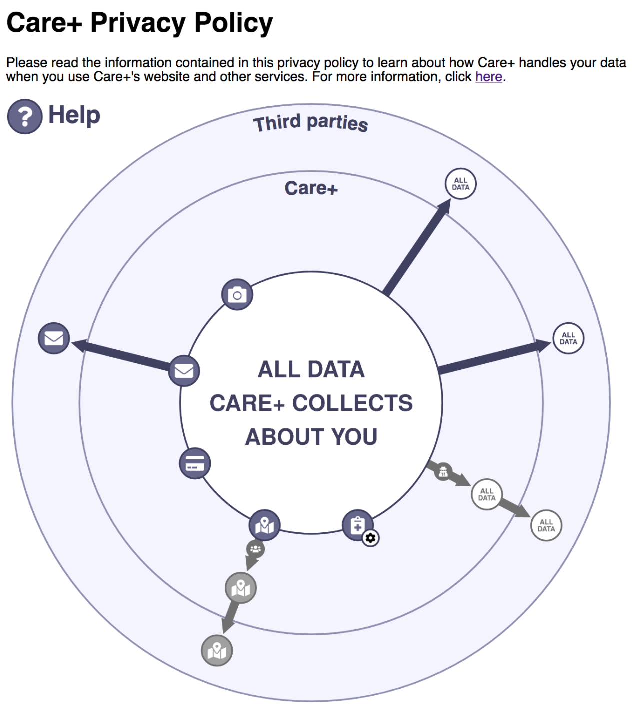

Figure 1: The Poli-see visualization of a fictional company named Care+.

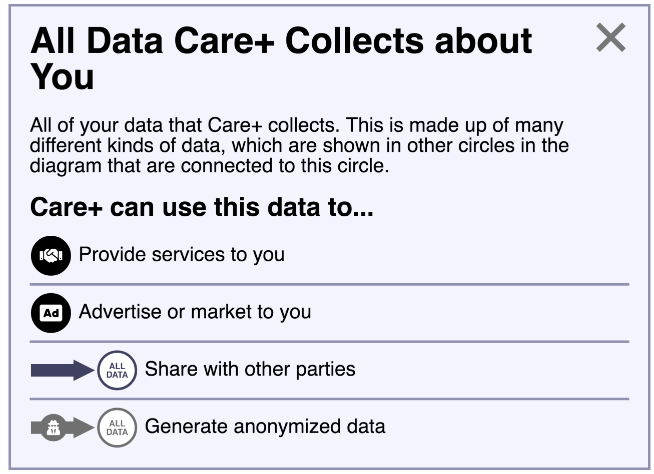

Figure 2: An example sidebar displayed when a user hovers over a node.

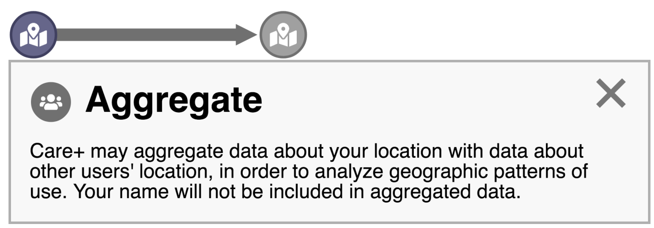

Figure 3: An example sidebar displayed when a user hovers over a link.

by our review of current privacy policies in order to support an expressive, usable graphical policy representation. A screenshot of the Poli-see visualization for a fictional company named Care+ is shown in Figure 1.

Like the ex-post visualizations, we generated a standardized list of data types and selected an icon[2] to depict each data type. The list of data types was generated during our review of current privacy policies, but it is intended to be extensible. The Poli-see representation includes a circle with an icon, or a *node*, for each type of data collected. Prior work has found that users' privacy preferences change significantly depending on whether data is personally identifiable [6], so Poli-see representations denote identifiable data nodes in blue and non-identifiable data nodes in gray. This distinction, alongside an icon legend, is noted in a help bar that appears when a user hovers over the help button. Like PrivacyInsight, we use arrows, or *links*, to represent the flow of information between principals.

We observed in our review that privacy policies focus primarily on three classes of principals: the data collector, its affiliates and business partners, and other third parties. However, early feedback suggested that the distinction between business partners and other third parties was not meaningful to users, so the current design of our policy representation depicts just two classes of principals: the data collector and third parties. These principals are depicted as concentric circles, with the inner circle representing the data collector and the outer circle representing third parties.

We observed that how data are used is central to many current privacy policies. To convey data use practices authorized by a privacy policy, we added an informational sidebar that appears when a user hovers over a node; the sidebar includes a list of ways in which the principal in possession of the data may use that type of data. We generated a list of standard data uses during our review of privacy policies, and we selected icons to depict each type of use. An example node sidebar is shown in Figure 2.

Prior work has found that users care about controlling how their data are used, but they struggle to locate options in privacy policies [21]. To signal to users when they have the option to control data use, we annotated nodes for which options are available with a small gear icon, and we added details about the available options (e.g., opting out of data collection, or deleting personal data) to the corresponding sidebar.

We also observed that privacy policies contain information about data transformations that derive new data: for example, automatically tagging a photo with the names of the people in it, or analyzing accelerometer data to infer the number of steps a user takes in a day. To depict these types of data use practices, we added additional links that denote data transformations. These links are labeled with a transformation-specific icon; a list of standardized transformation

[2]The icons that we used in the application are attributed under the Creative Commons license to Font Awesome and to Adam Stevenson from the Noun Project.

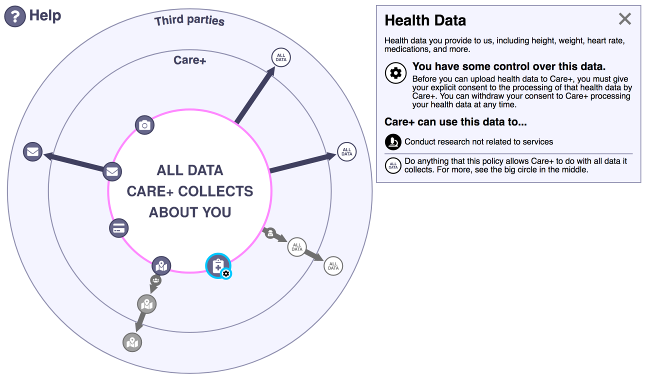

Figure 4: Poli-see when an individual data node (in this case, the Health Data node) has been clicked. The health data node and the all-data node are highlighted, and the node sidebar is shown.

types was generated during our review of current policies. Additional information about a transformation or disclosure is presented in the sidebar when a user hovers over the corresponding link; an example of a link sidebar is shown in Figure 3. However, in our design pilots we found that participants focused primarily on nodes and ignored links, causing them to miss information about authorized transformations. To address this issue, in each node sidebar we listed the ways in which the associated data type may be shared or transformed.

The final observation that emerged during our review of privacy policies was that current privacy policies often describe data use practices for non-specific data. For example, rather than describing how individual data types may be used, many policies describe how all collected data may be used. In our preliminary designs, we repeated this information in all nodes. However, this presented a usability problem: pilot participants found this repetition confusing, and they had trouble noticing when the data use practices for a specific data type differed from those for other data types. Our solution was to develop a visually distinct type of node called an *all-data node*, which represents non-specific data use practices for all data collected about the data subject. Conceptually, all-data nodes are composed of the data contained in nodes representing specific types of data, or *specific-data nodes*. Because participants in our design pilots had difficulty understanding this relationship, specific-data nodes are displayed embedded in the large all-data node in the center; hovering over one also causes the other to be outlined, as shown in Figure 4, and an item in the sidebar for specific-data nodes directs users to examine the central all-data node to see more information. Thus, information about how all data may be used or shared is displayed just once, by the all-data nodes.

4 IN-PERSON USER STUDY

To evaluate usability, we conducted an in-person user study with 24 participants comparing Poli-see representations to the Usable Privacy Policy Project's annotated privacy policy representations [28]. This study was approved by the Pomona College Instutional Review Board. A description of the annotated policy tool is in Section 2, and a screenshot is shown in Figure 5b. We evaluated each representation using four key metrics: how accurately participants could answer factual questions about the privacy policy after interacting with the representation, how confident participants felt about their answers to those factual questions, how quickly participants could answer those questions, and how participants rated the quality of their experience interacting with the representation.

Table 1: Reading metrics for the ASICS and Strava privacy policies.

	ASICS	Strava
Word Count	3851	4091
Flesch Reading Ease	36.8	36.1

We recruited students at our institution through Facebook groups, email lists, and paper flyers. Of our participants, 17 identified as female and 7 as male, and all were between the ages of 18 and 22. Each study lasted 30 minutes, and participants were compensated $5.

This study used a beta version of Poli-see that differed from the current design in two key ways. First, the beta version conveyed information about how "all data" were handled differently. In the beta version, the central node represented the data subject, rather than all data collected. The generic "all data collected" was represented by a separate, smaller all-data node, and light-colored *all-data links* pointing from specific-data nodes to this all-data node conveyed that specific data were also subject to the "all data" policies. Second, the beta version lacked the help bar, which contains the legend in the current design; instead, a tutorial button linked to a webpage briefly explaining how to interpret Poli-see representations. These design points were changed post-study in response to qualitative feedback received from study participants who found the number of links overwhelming and the icons confusing. A screenshot of the beta version of Poli-see used in this study is shown in Figure 5a.

4.1 In-person Study Design

During the study, participants interacted sequentially with the privacy policies of two fitness-tracking apps: ASICS Runkeeper [3] and Strava [24].[3] We selected privacy policies for fitness-tracking platforms because these services collect sensitive personal data, such as location and health information, and because there is a large selection of such services for users to choose from, making privacy-driven decision-making more feasible. ASICS and Strava were chosen in particular because their policies have similar lengths and reading levels (Table 1).

One policy was presented as a Poli-see visualization, and the other was presented as an annotated privacy policy. Both representations of the ASICS policy are shown in Figure 5. For each participant, we randomly determined which representation was used to depict which policy, as well as which policy was shown

[3] ASICS Runkeeper is an app that uses GPS to track users' runs. In the app, users can analyze their activity and follow personalized fitness plans. Strava is a social fitness network that can be integrated with smartphones and wearable devices to track users' activity through GPS. It enables users to analyze their activity, and it has many features for sharing and social interaction.

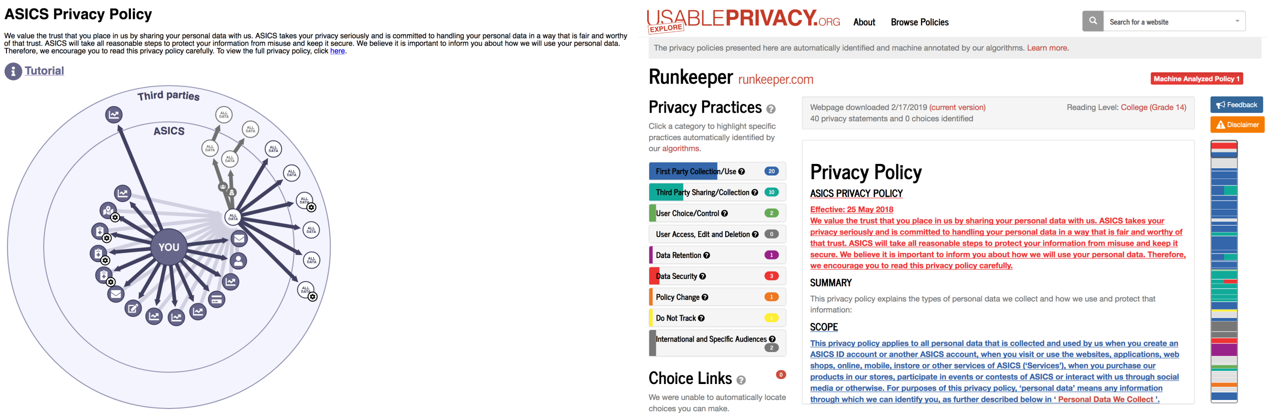

(a) Poli-see

(b) Annotated privacy policy

Figure 5: Screenshots of the two privacy policy representations for ASICS used in the in-person study. A beta version of Poli-see is depicted.

Table 2: Factual questions asked during the in-person study. The correct answer to Q5 under the ASICS policy was determined to be ambiguous and thus was excluded from our analysis.

	Question	Correct Answer ASICS	Correct Answer Strava
Q1	Does the policy allow [company] to collect information about your location?	Yes	Yes
Q2	Does the policy allow [company] to use your location information to conduct research?	Yes	Yes
Q3	Does the policy allow [company] to use your health information (such as your height and weight) to conduct research?	Yes	Yes
Q4	According to the policy, does [company] provide you with specific options for controlling how they collect or use your health information?	Yes	Yes
Q5	Does the policy allow [company] to use your health information to infer additional information about you?	–	Yes
Q6	According to the policy, does [company] provide you with specific options for controlling how they collect or use your contact information?	No	No
Q7	Does the policy allow [company] to share your contact information with third parties?	Yes	Yes
Q8	Does the policy allow [company] to sell aggregated information about its users?	No	Yes

first. When interacting with a privacy policy, participants were first given a short time to familiarize themselves with the representation used; the tutorial button was pointed out in Poli-see. Participants then answered a series of factual questions about the policy, recording their responses in an online survey form, which was open in a browser window side by side with the policy representation so that they could refer back and forth.

Participants answered the same series of eight factual questions, given in Table 2, for each policy. Answers were chosen on a five-point rating scale: "Definitely no," "Probably no," "Unsure," "Probably yes," and "Definitely yes." Participants were informed that their responses would be timed, and they were instructed to prioritize answering accurately first and answering quickly second. Unlike in our pilot studies, we did not anonymize the privacy policies, so participants were also instructed to answer based on the information contained in the privacy policy and not on pre-existing knowledge of the company or its practices.

After answering the factual questions for each policy representation, participants were asked three additional questions about their user experience; these questions are given in Table A.2. Answers were chosen on a five-point rating scale: "Strongly disagree," "Somewhat disagree," "Neither agree nor disagree," "Somewhat agree," and "Strongly agree." The full set of questions is given in Appendix A.

The study concluded with a five-minute interview to gather qualitative feedback from the participant on the two privacy policy representations they interacted with.

4.2 Quantitative Results

We evaluated the usability of Poli-see using four key metrics: user accuracy, user confidence, user speed, and quality of user experience. To evaluate user accuracy, we counted the number of participants who correctly answered the eight factual questions (Q1–Q8) in each set of responses. Answers were coded as correct if the participant selected "Definitely yes" or "Probably yes" when the correct answer

Table 3: Mean accuracy rates by question for the in-person study, broken down by policy and by representation. Mean confident accuracy rates are given in parentheses.

	ASICS		Strava	
	Poli-see	Annotated	Poli-see	Annotated
Q1	1.00 *(.75)*	0.92 *(.75)*	1.00 *(.92)*	0.92 *(.75)*
Q2	0.67 *(.58)*	0.83 *(.50)*	0.92 *(.58)*	0.67 *(.25)*
Q3	0.92 *(.50)*	0.83 *(.42)*	1.00 *(.83)*	0.58 *(.25)*
Q4	1.00 *(.67)*	0.75 *(.33)*	0.67 *(.50)*	0.58 *(.42)*
Q5	–	–	0.75 *(.58)*	0.75 *(.42)*
Q6	0.42 *(.25)*	0.25 *(.00)*	0.33 *(.08)*	0.42 *(.00)*
Q7	0.75 *(.42)*	0.92 *(.67)*	0.92 *(.75)*	0.75 *(.42)*
Q8	0.08 *(.00)*	0.08 *(.00)*	0.92 *(.75)*	0.75 *(.58)*
Avg.	0.69 *(.45)*	0.66 *(.38)*	0.82 *(.63)*	0.67 *(.38)*

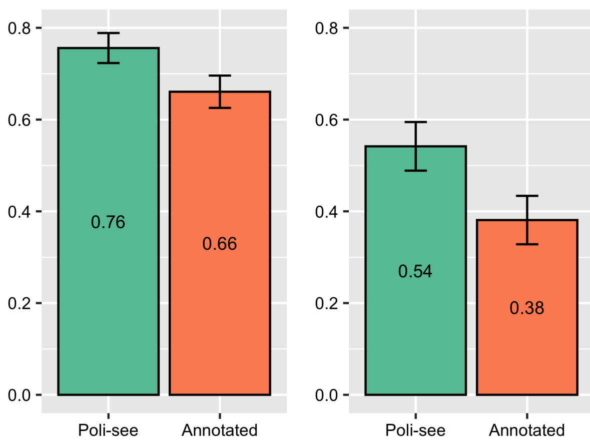

(a) Accuracy **(b) Confident accuracy**

Figure 6: Mean rates of accuracy and confident accuracy on the factual questions for the in-person study, shown with standard error bars.

was yes, and "Definitely no" or "Probably no" when the correct answer was no. To evaluate user confidence, we counted the number of participants who answered each of the eight factual questions both confidently and correctly (for example, answered "Definitely yes" if the correct answer was yes). These results by question are shown in Table 3.[4] Mean rates of accuracy and confidence are show in Figure 6. To evaluate user speed, we timed how long it took participants to answer each factual question. To evaluate quality of user experience, we analyzed participants' responses to the user experience questions as ordinal data.

To analyze our results on accuracy, confidence, and timing, we used ANOVA statistical tests with two independent variables: representation (annotated or Poli-see) and company (ASICS or Strava). For the purposes of analysis, we treated each of a participant's two sets of responses as independent observations.[5]

We found that our participants demonstrated improved accuracy using Poli-see compared to using the annotated policy representation ($F = 4.08$, $p = 0.0495$), while rates of inaccuracy (answers that were neither correct nor unsure) did not change ($F = 1.08$, $p = 0.305$). Moreover, participants also demonstrated higher rates of confident correct responses ($F = 4.71$, $p = 0.036$), while there was no significant difference between rates of confident incorrect responses ($F = 1.32$, $p = 0.258$). There were no significant interaction effects between representation and company.

We found that the mean time participants spent per question was less using Poli-see: 24.2 seconds for Poli-see and 30.5 seconds for the annotated representation. However, when we conducted an ANOVA on the log-normalized mean time, the difference was not statistically significant ($F = 1.87$, $p = 0.178$). We observe that the large variance (from under 10 seconds to over 100 seconds) and the small sample size preclude statistically significant timing differences in this study.

We conducted a Kruskal-Wallis test to detect differences between the two representations in the distribution of responses to our questions about participants' user experiences. However, there were no statistically significant differences.

4.3 Qualitative Results

We deemed the sample size of our in-person study to be too small to merit formal analysis of the qualitative comments; nonetheless, we found many of these comments to be informative.

Several participants expressed a preference for the visual presentation of Poli-see. P12 said, "I like that like all of it was within view . . . The icons are the icons that you would usually see, so it was pretty clear—versus the language of [the annotated representation], a lot of them seemed intermixed." Some participants said that Poli-see displayed information more concisely, which made information easier to find. P24 said, "I found [the annotated representation] really hard to read. It was just like a wall of text, so it was really hard to answer the questions." They stated that even locating information they had previously seen was difficult because of the scrolling involved. With Poli-see, they said, "I think why it's easier is they kind of substitute like huge chunks of words that describe, say, how you can control your data, for instance, or how you can control how the company uses your data, like in [the annotated representation]—but in [Poli-see], they simplify that block of text into like a gear icon." Others preferred Poli-see because they found it fun or interesting to use. P21 said that Poli-see was easier to navigate because "the picture really helped me, and you could just go round in a circle; it was kind of fun." P10 contrasted it to the annotated representation, which "was designed in a way that's just like a long essay that you just don't want to read. So, it's something that I'd rather just scroll over and it's just not as interactive or fun."

A common reason expressed for preferring Poli-see over the annotated representation was that Poli-see made it easy for users to follow sequences of nodes and links to see what can happen to

[4] In the process of validating the correct answers to the factual questions after conducting the study, we determined that the correct answer to Q5 for the ASICS policy was not a matter of settled law, due to the vague language in the original text policy. We therefore excluded Q5 from our analysis.

[5] We believe this is justified because we found that the order had no significant effect on any of the outcomes we analyzed, and because a given participant's two sets of responses deal with two different companies' privacy policies and two different privacy policy representations.

Table 4: **Factual questions asked during the online study. The correct answer to Q4 under the ASICS policy was determined to be ambiguous and thus was excluded from our analysis.**

	Question	Correct Answer	
		ASICS	Strava
Q1	Does the policy allow [company] to collect information about your location?	Yes	Yes
Q2	Does the policy allow [company] to use your location data to conduct research?	Yes	Yes
Q3	According to the policy, does [company] provide you with additional options for controlling how they collect or use your location data in particular?	Yes	Yes
Q4	Does the policy allow [company] to use your health information (such as your heart rate) to infer additional information about you (such as calories burned)?	–	Yes
Q5	According to the policy, does [company] provide you with additional options for controlling how they collect or use your contact information in particular?	No	No
Q6	Does the policy allow [company] to share your contact information with third parties?	Yes	Yes
Q7	Does the policy allow [company] to share all data they collect about you with third-party service providers?	Yes	Yes
Q8	Does the policy allow [company] to sell aggregated information about its users?	No	Yes

data. P6 said they liked Poli-see because "you could see where the information flowed." P5 contrasted it with the difficulty of tracing information in the annotated representation, saying that in the annotated representation "you're kind of like bouncing all over the place, looking for like keywords that might be relevant, and then when you do find a keyword, then you have to go to a different section to see if they're sharing it, for example. So, like, you can see that they collect it, but you're not clear if they share it until you check another section. Whereas, in [Poli-see], it was much easier because it was just like a flowchart basically." Similarly, P10 said that answering questions using the annotated representation required a lot of scrolling and felt like putting together a "puzzle," while they could follow arrows in Poli-see to easily see what might happen to data.

However, several other participants expressed a preference for the annotated privacy policies. P15 found the number of links overwhelming, saying, "all the arrows were like pointing every way, so it was actually hard to see." P22 expressed confusion about the meaning of icons, especially repeated icons, and the difference in color between identifiable links and unidentifiable links; they said, "maybe it's too much information at once."

Several participants described Poli-see's all-data nodes in particular as confusing or vague. P4 expressed concern about the nonspecificity of all-data nodes, specifically pointing out the use "Not limited by this privacy policy," which appears when the privacy policy does not limit how a third party may use shared data; they expressed concern that information was being hidden. They contrasted this with the annotated privacy policy, which they felt had more "upfront" language; after pointing out a passage in the annotated policy that listed examples of how shared data might be used by the third party, they said, "that's like a concrete action that they're doing. Like, oh, we may share your information with these people. And so I guess I see exactly what's going on, rather than like, oh, they have access to your data . . . and so I felt like I had a better sense of what was going on with the data and where it was going." Similarly, P24 said, about Poli-see, "I wish that they had made the all-data icons in third parties more specific to the kinds of data they were sharing out to all parties, because then it just seems like, oh, like they're still sharing all of my data to like an outside company . . . It was hard to tell the differences between the different all-data icons."

Some participants expressed concern that Poli-see might not convey all the details or the original wording of the privacy policy. P2 said, "I also feel like more things might have been hidden, because it was more complex, and it wasn't like—the details didn't seem as fleshed out." P22 said that the annotated representation was more straightforward because "it helped like being able to see the entire wording of the policy."

This qualitative feedback was incorporated into a final round of design revisions, culminating in the current Poli-see design described in Section 3.

5 ONLINE USER STUDY

Acknowledging the small size ($n = 24$) and unrepresentative nature (entirely students at our institution) of our in-person study, we determined that a larger-scale study with a more representative population was necessary to enable us to draw general conclusions about the usability of Poli-see. We therefore conducted a pilot study with 36 participants, followed by a full user study with 658 participants, on Amazon Mechanical Turk. This study was approved by the Pomona College Institutional Review Board. The survey was restricted to people located in the United States who had previously completed at least 50 HITs with an approval rate of 95%. On average, participants spent 8 minutes and 5 seconds to complete the study, and participants were compensated $2. As an attention check question, we asked participants to enter the URL of the tool they used into a text response box; we rejected 58 responses for failing this.

Of the 600 accepted participants in the full study, 62.5% identified as male, and 37.2% identified as female. 72.5% of our participants identified as white, 13.8% as black, 7.3% as Asian, and 5% as Hispanic. 10.7% were ages 18–25, 48.2% ages 26–35, 29.3% ages 36–50, 9.5% ages 50–65, and 2.3% ages 66 and older. Although the demographics of our study population do not match the demographics

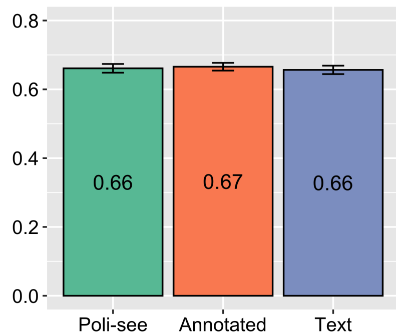

(a) Overall

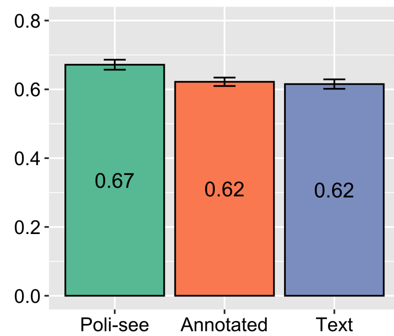

(b) On simple questions

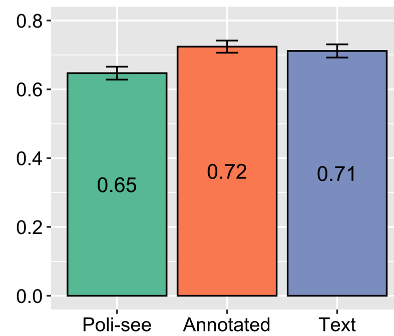

(c) On complex questions

Figure 7: Mean rates of accuracy on the factual questions for the online study, shown with standard error bars.

of the United States (as reported by the American Community Survey [26]), prior work has found that Mechanical Turk survey responses on the topics of security and privacy are generally representative of the U.S. population for people aged 18-49 who have completed some college [20].

5.1 Online Study Design

Our Mechanical Turk study compared Poli-see representations, annotated text representations, and conventional text policies. Each participant was shown one of the two fitness-tracking app privacy policies used in the in-person study (Strava or ASICS Runkeeper), depicted using one of the three representations; the privacy policy and the representation were both randomly determined.

As with the in-person user study, the goal of our online study was to evaluate Poli-see along four key metrics: user accuracy, user confidence, user speed, and quality of user experience. To evaluate the first three criteria, participants answered eight factual questions about their assigned privacy policy (these questions are given in Table 4)), choosing responses on a five-point rating scale: "Definitely no," "Probably no," "Unsure," "Probably yes," and "Definitely yes." The amount of time it took each participant to answer each question was recorded. To evaluate the quality of participants' user experience, participants were then asked six questions about their experience with the policy representation and asked to select a response on a different five-point rating scale: "Strongly disagree," "Somewhat disagree," "Neither agree nor disagree," "Somewhat agree," and "Strongly agree". The full list of questions used in the online study are given in Appendix B.

The factual questions were slightly revised from the in-person study to improve clarity. In order to better gauge participants' understanding of the all-data node concept, we also replaced one question about how a specific type of data could be used (which was very similar to another existing question) with a new question about whether the policy allowed all data collected to be shared with third parties.

In order to access their assigned policy representation, participants were given a link that opened the representation in a new browser window. Participants were instructed to keep their policy representation open, but the link was provided with each factual question in case the participant accidentally closed the window.

Table 5: Mean accuracy rates by question for the online study, broken down by policy and by representation.

	ASICS			Strava		
	Poli-see	Anno.	Text	Poli-see	Anno.	Text
Q1	0.90	0.91	0.90	0.89	0.91	0.96
Q2	0.73	0.84	0.81	0.75	0.88	0.83
Q3	0.85	0.76	0.77	0.88	0.66	0.68
Q4	–	–	–	0.91	0.88	0.87
Q5	0.26	0.23	0.17	0.18	0.19	0.12
Q6	0.71	0.80	0.86	0.72	0.83	0.76
Q7	0.69	0.64	0.67	0.73	0.69	0.66
Q8	0.16	0.25	0.27	0.80	0.78	0.74
Avg.	0.61	0.63	0.64	0.71	0.70	0.68

After the last factual question, when the link was no longer available, participants were asked to provide the link; if the link did not match the assigned condition, then the response was rejected.

5.2 Accuracy Results

To evaluate the usability of Poli-see, we counted the number of participants who answered the eight factual questions correctly (as well as confidently and correctly). We determined that the correct answer to Q4 for the ASICS policy was ambiguous, so we excluded Q4 from our analysis. The average accuracy results for each policy representation are shown in Figure 7a, and a breakdown of the results by question is given in Table 5.

As in the in-person study, we conducted two-way ANOVA tests used to analyze the in-person study. As before, there were no significant interaction effects between representation and company. On average, participants who were shown Poli-see visualizations answered 66.1% of the questions correctly; this is lower than the the average accuracy observed in the in-person study (75.6%), and it was not significantly different from the average accuracy of participants in the online study who were shown an annotated privacy policy or a text privacy policy (p = 0.862). Similarly, the representation used had no significant effects on rates of inaccuracy (p = 0.544), confident accuracy (p = 0.581), and confident inaccuracy (p = 0.140).

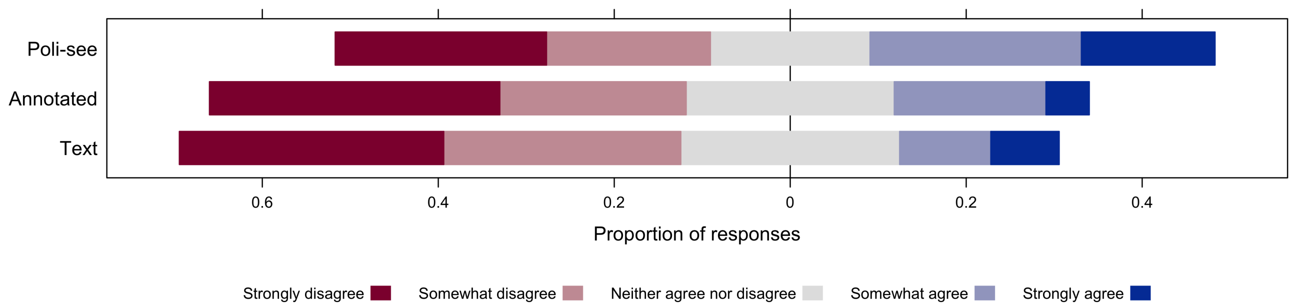

(a) "Finding information in this policy was an enjoyable experience."

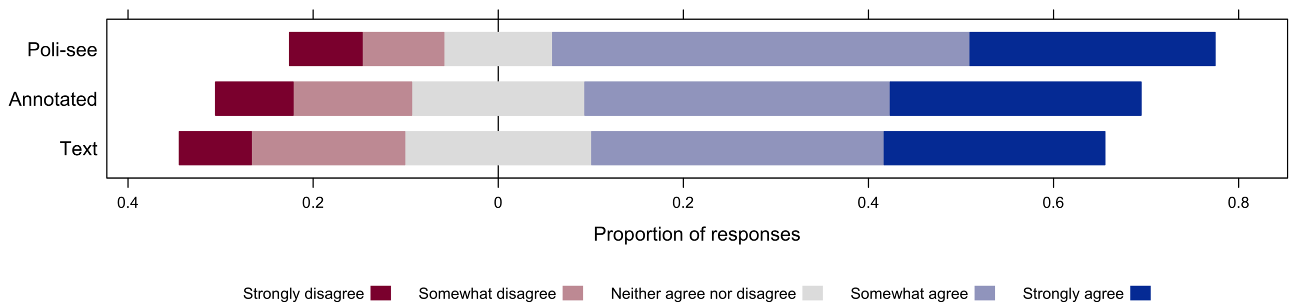

(b) "If this were how [company] presented its privacy policy, I would be likely to look at this policy before deciding whether to use this service."

Figure 8: Response breakdowns for two of the user experience questions for the online study.

There are several possible reasons for the discrepancy between accuracy rates for Poli-see in the in-person study and in the online study, including different sample populations, different sample sizes, different length of exposure to the Poli-see visualization, or different levels of attention to detail.

To further investigate the discrepancy between the in-person results and the online results, we categorized each question as simple or complex based on how the answer could be found in Poli-see. Simple questions required interacting with a single element of the interface (e.g., looking at a single specific-data node). Complex questions required interacting with multiple elements (e.g., looking at both a specific-data node and an all-data node). For both ASICS and Strava, four questions were categorized as simple (Q1, Q3, Q5, Q7), and three were categorized as complex (Q2, Q6, Q8). The mean accuracy results for simple and for complex questions are shown in Figures 7b and 7c, broken down by policy representation.

We found that participants who used Poli-see had significantly higher accuracy on simple questions and significantly lower accuracy on complex questions. For simple questions, an ANOVA test showed significant effects of representation on accuracy rates ($F = 5.47, p = 0.004$). A post-hoc Tukey's HSD test found statistically higher pairwise accuracy for participants who used Poli-see compared to participants who saw annotated policies ($p = 0.025$) and text policies ($p = 0.010$). For complex questions, ANOVA tests also found significant effects ($F = 5.8, p = 0.003$); the HSD test showed lower accuracy with Poli-see visualizations than with annotated policies ($p = 0.004$) and text policies ($p = 0.024$). An HSD test did not find any statistically significant differences between annotated policies and text policies for either simple questions ($p = 0.937$) or complex questions ($p = 0.866$).

5.3 Timing Results

As in the in-person study, we used an ANOVA test to analyze the log-normalized mean time per question. We excluded timing data from Q1 because many participants spent a far longer time on Q1 than on any other question (we believe they used that time to familiarize themselves with their policy and representation). The mean time per question was 26.9 seconds for participants who were shown a Poli-see visualization, which was shorter than for participants who were shown an annotated representation (29.3 seconds) or a text representation (31.9 seconds). However, the difference was not statistically significant ($p = 0.152$). As before, we believe the large variance precluded statistically significant results.

5.4 User Experience Results

We also asked participants to rate how strongly they agreed with a series of subjective statements about their experience interacting with the policy representation they were shown (Poli-see, annotated, or text); the complete set of statements is given in Appendix B.2, and responses to selected statements are graphically represented in Figure 8.

We analyzed the results using a Kruskal-Wallis test and found that representation had a significant effect on the distribution of participants' responses for whether they found their experience enjoyable ($p = 0.0003$). Pairwise tests showed that the distribution was significantly different between Poli-see and text representations ($p = 0.0005$) and between Poli-see and annotated representations ($p = 0.0006$), while there was no significant difference between the text and the annotated representations ($p = 0.904$). These significant results hold up under a Bonferroni correction for multiple testing.

We also partitioned the responses into positive ("Strongly agree" or "Somewhat agree") and not positive ("Neither agree nor disagree", "Somewhat disagree", or "Strongly disagree") and analyzed the results with a chi-squared contingency test. Participants who used Poli-see rated their experience as more enjoyable than participants who used an annotated representation ($p = 0.0001$) or a text representation ($p < .0001$). Participants also stated that they were more likely to consult the policy if it were represented with Poli-see than if the company used an annotated representation ($p = 0.004$) or a text representation ($p = 0.0002$). There were no significant results for other statements.

We interpret these results as a positive indication that users might be more likely to look at privacy policies if presented with a graphical representation such as Poli-see.

6 DISCUSSION

In this work, we investigated how data use practices might be conveyed by a graphical representation. Our goal was to develop a representation that would be expressive, usable, and scalable. Through the design and evaluation process of Poli-see, we discovered insights into key advantages and challenges that may guide future design and development of graphical privacy policy representations.

6.1 Advantages of Graphical Representations

Visualizations have previously been shown to effectively convey actual data collection and dissemination [5, 10]; in the course of this work, we discovered that visualizations also have two key advantages when it comes to representing privacy policies.

Expressiveness. In our review of privacy policies, we identified key patterns that appear across current policies. In particular, we observed that data use practices are an important part of privacy policies. The specified uses include how data are transformed, including what derived data are inferred. As we iterated on the design of Poli-see through design pilots and user studies, we found that a significant advantage of graphical privacy policy representations is their ability to express the flow of information through such data processing.

Engagement. A key challenge posed by current text privacy policies is that users rarely bother to read them at all [2, 14, 23]. In our in-person study, many participants thought that the visual representation of Poli-see made it easier to use and more engaging. In our online study with the current version of Poli-see, participants were significantly more likely to respond that they enjoyed interacting with Poli-see and that they would consult a Poli-see representation outside of a research study. These answers suggest that graphical representations might be able to increase user engagement, thereby offering an effective alternative to text representations.

6.2 Challenges of Graphical Representations

Despite the advantages offered by graphical representations and despite an extensive design process, Poli-see did not demonstrate significant improvements in user accuracy, user confidence, or user speed compared to alternative text representations in our large-scale online study. Although improved engagement—combined with comparable usability based on other metrics—might be sufficient to justify the adoption of graphical privacy policy representations, we believe this approach has the potential to outperform text policies along all four usability metrics we considered. However, designing such a representation will require further research to develop better solutions to the challenges we identified while designing Poli-see.

Intuitiveness. Although current text policies have many limitations, users are (by necessity) reasonably familiar with the format. Any graphical representation will face the challenge that users are unfamiliar with the format and will need to figure out how to extract the information they want from the graphical representation. Both the qualitative feedback and the quantitative results observed in this work suggest that Poli-see has a steep learning curve. In our in-person study, some users found links, all-data nodes, and other features confusing. And comparing the accuracy results of our in-person study (which included an initial exploratory phase before users were asked the series of factual questions, and in which participants answered questions significantly more accurately and more confidently using Poli-see than using an annotated text policy) to the results of our online study (which did not include the initial exploratory phase, and in which there was no statistically significant improvement in accuracy or confidence) suggests that this learning curve might affect the usability of Poli-see. Future graphical representations would benefit from explicitly considering the intuitiveness and the learning curve of their tool during their design phase.

Policy Specificity. One of the key challenges we identified in designing a usable graphical policy representation was the non-specificity in conventional privacy policies. That is, many privacy policies specify the types of data the company may collect and then define a set of data uses for "all data we collect"; specific uses of specific data types are typically defined only in special cases. A key challenge for graphical policy representations is therefore how to convey both the normal case and the exceptions.

In early versions of our design, we attached each non-specific data use to each specific element it applied to. For example, if a policy stated that a company may collect the data subject's name, address, and location information, and that all information it collects may be used for marketing, then Poli-see would show, separately in each respective node, that the data subject's name, email address, and location information could be used for marketing. However, participants in our design pilots found that the level of clutter in the resulting graphical representation rendered it unusable.

In later designs, including the current Poli-see design, we separated this non-specific information into the more abstracted design of all-data nodes. However, this design required users to consult both the specific "email" node and an "all-data" node in order to understand fully how their email might be used—a requirement that was not readily apparent to many of our study participants. The discrepancy in accuracy rates for simple questions (which did not require understanding this feature) and for complex questions (which did) observed in our online study[6] might be explained by insufficient comprehension of the all-data nodes feature. We believe that future graphical privacy policy representations could benefit from further research into how to express non-specific policy rules.

Scalability. In an effort to meet our goal of scalability, we implemented Poli-see as a graphical representation that is automatically generated from a JSON encoding of the privacy policy. Unfortunately, generating the initial JSON encoding is a time-consuming process requiring a domain expert to interpret and encode each privacy policy by hand. Although this is an improvement over manually generating a similar interactive representation of each policy without Poli-see, it falls short of achieving useful scalability; this is a challenge that should be addressed by future graphical policy representations. However, we note that orthogonal work has explored how machine learning and crowdsourcing might be leveraged to reliably parse existing text representations of privacy policies into discrete data practices [30]; we believe this work could eventually be incorporated into future graphical representations to achieve the level of scalability necessary for widespread deployment in the real world.

Accessibility. Any visualization tool or graphical representation inherently comes with the challenge of how to make it accessible to colorblind or visually-impaired users. We ensured that the color scheme adopted by Poli-see was accessible to colorblind users by incorporating feedback from colorblind individuals throughout the design process. However, accessibility for visually-impaired users more broadly was beyond the scope of this work. Addressing the issue of accessibility is a challenge that should be addressed by future graphical policy representations.

7 CONCLUSION

We view the primary contributions of this work to be (1) the development of Poli-see, the first interactive, graphical representation of privacy policies, (2) an evaluation of Poli-see that demonstrates the potential for graphical representations to accurately convey data use practices to users in a manner that encourages user engagement, and (3) the discovery of insights into the advantages and challenges of graphical policy representations. Based on the results of this work and the ongoing challenges we identified, we do not believe that Poli-see is ready for large-scale deployment in production systems. Instead, we view this work as validation of the idea that graphical representations have the potential to improve usability compared to text-based representations, and we believe that this work can serve as a foundation for future research on graphical privacy policy representations.

[6] We did not observe this discrepancy in our in-person study. It is possible that participants gained a better understanding of this feature while initially exploring the policy representation. Alternatively, it is possible that the discrepancy is due to sampling bias or small sample size in our in-person study.

REFERENCES

[1] Julio Angulo, Simone Fischer-Hübner, Tobias Pulls, and Erik Wästlund. 2015. Usable Transparency with the Data Track: A Tool for Visualizing Data Disclosures. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 1803–1808. https://doi.org/10.1145/2702613.2732701

[2] Manon Arcand, Jacques Nantel, Mathieu Arles-Dufour, and Anne Vincent. 2007. The Impact of Reading a Web Site's Privacy Statement on Perceived Control over Privacy and Perceived Trust. *Online Information Review* 31, 5 (2007), 661–681. https://doi.org/10.1108/14684520710832342

[3] ASICS 2018. *ASICS Privacy Policy*. Retrieved October 30, 2019 from https://www.asics.com/us/en-us/privacy-policy.html

[4] Solon Barocas and Helen Nissenbaum. 2009. On Notice: The Trouble with Notice and Consent. In *Proceedings of the Engaging Data Forum: The First International Forum on the Application and Management of Personal Electronic Information*. https://ssrn.com/abstract=2567409

[5] Cristoph Bier, Kay Kühne, and Jürgen Beyerer. 2016. PrivacyInsight: The Next Generation Privacy Dashboard. In *Privacy Technologies and Policy*. Springer, 135–152. https://doi.org/10.1007/978-3-319-44760-5_9

[6] Eleanor Birrell and Fred B. Schneider. 2017. *A Reactive Approach for Use-Based Privacy*. Technical Report. Cornell University. https://hdl.handle.net/1813/54843

[7] Debmalya Biswas, Imad Aad, and Gian Paolo Perrucci. 2013. Privacy Panel: Usable and Quantifiable Mobile Privacy. In *2013 International Conference on Availability, Reliability and Security*. IEEE, 218–223. https://doi.org/10.1109/ARES.2013.29

[8] European Parliament and Council of the European Union. 2016. General Data Protection Regulation (GDPR). https://eur-lex.europa.eu/eli/reg/2016/679/oj

[9] Federal Trade Commission et al. Last modified June 25, 2007. Fair Information Practice Principles. (Last modified June 25, 2007).

[10] Simone Fischer-Hübner, Julio Angulo, Farzaneh Karegar, and Tobias Pulls. 2016. Transparency, Privacy and Trust – Technology for Tracking and Controlling My Data Disclosures: Does This Work?. In *Trust Management X*. Springer, 3–14. https://doi.org/10.1007/978-3-319-41354-9_1

[11] Kambiz Ghazinour and Tahani Albalawi. 2016. A Usability Study on the Privacy Policy Visualization Model. In *2016 IEEE 14th Intl Conf on Dependable, Autonomic and Secure Computing, 14th Intl Conf on Pervasive Intelligence and Computing, 2nd Intl Conf on Big Data Intelligence and Computing and Cyber Science and Technology Congress(DASC/PiCom/DataCom/CyberSciTech)*. IEEE, 578–585. https://doi.org/10.1109/DASC-PICom-DataCom-CyberSciTec.2016.109

[12] Hana Habib, Yixin Zou, Aditi Jannu, Neha Sridhar, Chelse Swoopes, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. 2019. An Empirical Analysis of Data Deletion and Opt-Out Choices on 150 Websites. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. USENIX Association, Santa Clara, CA. https://www.usenix.org/conference/soups2019/presentation/habib

[13] Hamza Harkous, Kassem Fawaz, Rémi Lebret, Florian Schaub, Kang G. Shin, and Karl Aberer. 2018. Polisis: Automated Analysis and Presentation of Privacy Policies Using Deep Learning. In *27th USENIX Security Symposium (USENIX Security 18)*. USENIX Association, Baltimore, MD, 531–548. https://www.usenix.org/conference/usenixsecurity18/presentation/harkous

[14] Carlos Jensen, Colin Potts, and Christian Jensen. 2005. Privacy Practices of Internet Users: Self-Reports versus Observed Behavior. *International Journal of Human-Computer Studies* 63, 1 (2005), 203–227. https://doi.org/10.1016/j.ijhcs.2005.04.019

[15] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W. Reeder. 2009. A "Nutrition Label" for Privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security*. ACM. https://doi.org/10.1145/1572532.1572538

[16] Patrick Gage Kelley, Lucian Cesca, Joanna Bresee, and Lorrie Faith Cranor. 2010. Standardizing Privacy Notices: An Online Study of the Nutrition Label Approach. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1573–1582. https://doi.org/10.1145/1753326.1753561

[17] Aleecia M. McDonald and Lorrie Faith Cranor. 2008. The Cost of Reading Privacy Policies. *I/S: A Journal of Law and Policy for the Information Society* 4, 3 (2008), 543–568. http://hdl.handle.net/1811/72839

[18] Aleecia M. McDonald, Robert W. Reeder, Patrick Gage Kelley, and Lorrie Faith Cranor. 2009. A Comparative Study of Online Privacy Policies and Formats. In *Privacy Enhancing Technologies*. Springer, 37–55. https://doi.org/10.1007/978-3-642-03168-7_3

[19] Aza Raskin et al. Last modified June 28, 2011. Privacy Icons. https://wiki.mozilla.org/Privacy_Icons

[20] Elissa M. Redmiles, Sean Kross, and Michelle L. Mazurek. 2019. How Well Do My Results Generalize? Comparing Security and Privacy Survey Results from MTurk, Web, and Telephone Samples. In *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1326–1343. https://doi.org/10.1109/SP.2019.00014

[21] Joel R. Reidenberg, Travis Breaux, Lorrie Faith Cranor, Brian French, Amanda Grannis, James T. Graves, Fei Liu, Aleecia McDonald, Thomas B. Norton, Rohan Ramanath, N. Cameron Russell, Norman Sadeh, and Florian Schaub. 2015.

Disagreeable Privacy Policies: Mismatches between Meaning and Users' Understanding. *Berkeley Technology Law Journal* 30, 1 (2015), 39–88. https://doi.org/10.2139/ssrn.2418297

[22] Ravi Inder Singh, Manasa Sumeeth, and James Miller. 2011. A User-Centric Evaluation of the Readability of Privacy Policies in Popular Web Sites. *Information Systems Frontiers* 13, 4 (2011), 501–514. https://doi.org/10.1007/s10796-010-9228-2

[23] H. Jeff Smith, Tamara Dinev, and Heng Xu. 2011. Information Privacy Research: An Interdisciplinary Review. *MIS Quarterly* 35, 4 (2011), 989–1015. https://misq.org/catalog/product/view/id/1518

[24] Strava 2018. Strava Privacy Policy. Retrieved October 30, 2019 from https://www.strava.com/legal/privacy

[25] Joseph Turow, Lauren Feldman, and Kimberly Meltzer. 2005. *Open to Exploitation: American Shoppers Online and Offline.* Technical Report. Annenberg Public Policy Center of the University of Pennsylvania. https://repository.upenn.edu/asc_papers/35

[26] United States Census Bureau [n.d.]. American Community Survey. https://www.census.gov/programs-surveys/acs/

[27] Usable Privacy Policy Project 2017. *The Usable Privacy Policy Project.* Retrieved October 30, 2019 from https://usableprivacy.org

[28] Usable Privacy Policy Project 2018. *Machine-Annotated Privacy Policies.* Retrieved October 30, 2019 from https://explore.usableprivacy.org/?view=machine

[29] Kim-Phuong L. Vu, Vanessa Chambers, Beth Creekmur, Dongbin Cho, and Robert W. Proctor. 2010. Influence of the Privacy Bird® User Agent on User Trust of Different Web Sites. *Computers in Industry* 61, 4 (2010), 311–317. https://doi.org/10.1016/j.compind.2009.12.001

[30] Shomir Wilson, Florian Schaub, Frederick Liu, Kanthashree Mysore Sathyendra, Daniel Smullen, Sebastian Zimmeck, Rohan Ramanath, Peter Story, Fei Liu, Norman Sadeh, and Noah A. Smith. 2018. Analyzing Privacy Policies at Scale: From Crowdsourcing to Automated Annotations. *ACM Transactions on the Web* (2018), 1:1–1:29. https://doi.org/10.1145/3230665

Appendices

Appendix A IN-PERSON USER STUDY QUESTIONS

A.1 Factual Questions

		Correct Answer	
	Question	ASICS	Strava
Q1	Does the policy allow [company] to collect information about your location?	Yes	Yes
Q2	Does the policy allow [company] to use your location information to conduct research?	Yes	Yes
Q3	Does the policy allow [company] to use your health information (such as your height and weight) to conduct research?	Yes	Yes
Q4	According to the policy, does [company] provide you with specific options for controlling how they collect or use your health information?	Yes	Yes
Q5	Does the policy allow [company] to use your health information to infer additional information about you?	–	Yes
Q6	According to the policy, does [company] provide you with specific options for controlling how they collect or use your contact information?	No	No
Q7	Does the policy allow [company] to share your contact information with third parties?	Yes	Yes
Q8	Does the policy allow [company] to sell aggregated information about its users?	No	Yes

A.2 User Experience Questions

	Question Text
$Q_{UX}1$	I understand how [company] uses my data.
$Q_{UX}2$	I feel comfortable with the ways [company] would use my data.
$Q_{UX}3$	If I wanted a fitness app, I would be willing to use this service, knowing how [company] would use my data.

A.3 Demographic and Background Questions

	Question Text
Q_D1	What is your gender?
Q_D2	What is your age?
Q_D3	What is your nationality?
Q_D4	What is your ethnicity?
Q_D5	What is your highest level of formal education?
Q_D6	How frequently do you read privacy policies?
Q_D7	If you've read privacy policies recently, what was your goal in doing so?

Appendix B ONLINE USER STUDY QUESTIONS

B.1 Factual Questions

	Question	Correct Answer	
		ASICS	Strava
Q1	Does the policy allow [company] to collect information about your location?	Yes	Yes
Q2	Does the policy allow [company] to use your location data to conduct research?	Yes	Yes
Q3	According to the policy, does [company] provide you with additional options for controlling how they collect or use your location data in particular?	Yes	Yes
Q4	Does the policy allow [company] to use your health information (such as your heart rate) to infer additional information about you (such as calories burned)?	–	Yes
Q5	According to the policy, does [company] provide you with additional options for controlling how they collect or use your contact information in particular?	No	No
Q6	Does the policy allow [company] to share your contact information with third parties?	Yes	Yes
Q7	Does the policy allow [company] to share all data they collect about you with third-party service providers?	Yes	Yes
Q8	Does the policy allow [company] to sell aggregated information about its users?	No	Yes

B.2 User Experience Questions

	Question Text
$Q_{UX}1$	I feel that I understand how [company] uses my data.
$Q_{UX}2$	I would feel comfortable with the ways [company] would use my data.
$Q_{UX}3$	If I wanted a fitness app, I would be willing to use this service, knowing how [company] would use my data.
$Q_{UX}4$	If this were how [company] presented its privacy policy, I would be likely to look at this policy before deciding whether to use this service.
$Q_{UX}5$	Finding information in this policy was easy to do.
$Q_{UX}6$	Finding information in this was policy was an enjoyable experience.

B.3 Demographic and Background Questions

	Question Text
Q_D1	What is your gender?
Q_D2	What is your age?
Q_D3	What is your nationality?
Q_D4	What is your ethnicity?
Q_D5	What is your highest level of formal education?
Q_D6	Have you read a privacy policy in the last three months?
Q_D7	How frequently do you read privacy policies?