Cookie Slayer: Challenging Learned Helplessness Through Instrumental Interaction Design

BENCE SZABO*, Aalborg University in Copenhagen, Denmark
LOUISE FOLDØY STEFFENS*, Aalborg University in Copenhagen, Denmark
SARA SELMAN*, Aalborg University in Copenhagen, Denmark

CCS Concepts: • **Human-centered computing → Human computer interaction (HCI)**; **Web-based interaction**; • **Security and privacy → Usability in security and privacy**.

Additional Key Words and Phrases: Human-Computer Interaction, Instrumental Interaction Design, Digital Agency, Learned Helplessness, Online Consent, Consent Management Platforms (CMPs), User Appropriation, Tool Ownership, AI-Powered Automation, Seamfulness

ACM Reference Format:
Bence Szabo, Louise Foldøy Steffens, and Sara Selman. 2026. Cookie Slayer: Challenging Learned Helplessness Through Instrumental Interaction Design. 1, 1 (May 2026), 67 pages. https://doi.org/XXXXXXX.XXXXXXX

Abstract

Current online consent mechanisms frequently utilise dark patterns and manipulative architectures, inducing widespread privacy fatigue and a state of resignation known as the "Okay, whatever" effect. This behaviour reflects deeply entrenched learned helplessness, where users feel their actions have no real effect on tracking outcomes. These manipulative frameworks are typically manifested through complex consent management platforms (CMPs), which hide tracking intentions within dense, inaccessible interface layers and underlying system metadata. Moving beyond traditional data-processing interaction models, this paper introduces a browser-level privacy protection tool built on third-wave Human Computer interaction (HCI) and instrumental interaction design principles. By reifying abstract data policies into manipulatable objects and providing polymorphic reusable instruments, the Cookie Slayer extension lowers the degree of indirectness for querying, evaluating, and enforcing privacy personal privacy preferences. To facilitate this interaction, the extension integrates two key components: a large language model (LLM) that parses dense privacy interfaces to deliver contextual explanations, and a transparent, dynamically evolving recommendation algorithm that suggests personally relevant privacy actions directly at the objects of interest. A qualitative technology probe study ($N = 9$) across users with high, medium, and low levels of privacy literacy suggests that direct interaction with CMP metadata fosters a distinct sense of digital ownership and active agency. Our findings indicate that while instrumental mediation successfully fosters user appropriation and heightened privacy care, true preference alignment remains structurally limited by language model reliability and recommendation algorithm flexibility. This underscores critical design trade-offs between seamless automation and active reflection.

*All authors contributed equally to this research.

Authors' Contact Information: Bence Szabo, bszabo21@student.aau.dk, Aalborg University in Copenhagen, Copenhagen, Denmark; Louise Foldøy Steffens, Aalborg University in Copenhagen, Copenhagen, Denmark, lfst21@student.aau.dk; Sara Selman, Aalborg University in Copenhagen, Copenhagen, Denmark, sselma24@student.aau.dk.

ACM XXXX-XXXX/2026/5-ART
https://doi.org/XXXXXXX.XXXXXXX

1 Introduction and motivation

Current online consent mechanisms are fundamentally broken. Although the GDPR mandates "freely given, specific, and informed" [12] consent, dark patterns and unethical click-maximisation designs frequently manipulate users. This systematic failure has led to widespread privacy fatigue and the "Okay, whatever" effect, where people reflexively click through pop-ups rather than truly enforcing their privacy preferences [11]. Such actions reflect a state of behavioural, emotional, and cognitive apathy rooted in learned helplessness. This phenomena is not merely a result of insufficient comprehension, but of a perceived lack of contingency; the belief that no matter what action a user takes, the outcome, such as tracking, remains the same [30].

Overcoming this futility requires a step beyond second-wave HCI, which treats users as data processors and making their jobs easier with static automation or information pre-processing. Today's users are just as conscious in the online space as in the real world, bearing unique experiences as well as cultural and personal backgrounds. We argue that a more promising avenue is a shift toward interactive instruments that provide users with clear evidence of their influence through appropriation and mediation. By promoting digital ownership through these principles, it becomes a possible prospect to transform the user from resigning to the system to actively shaping it. As users appropriate a tool to fit their unique life contexts, they move beyond mere compliance toward active agency, turning what once felt like a lost cause into a manageable domain of personal control.

This paper makes two primary contributions to the field of HCI. First, we introduce an empirical and artifact-driven contribution through the design, implementation, and in-the-wild deployment of an artificial intelligence (AI) powered browser extension that instantiates generative instrumental interaction design principles [5]. By providing direct manipulation and localised visual cues over adversarial CMPs, this artifact serves as an architectural blueprint for dismantling the "Okay, whatever" effect and countering learned

helplessness at the user interface level. Second, we provide an empirical understanding of how different privacy-literacy user groups appropriate automated recommender systems and language models in adversarial environments. Through a mixed-methods evaluation, we uncover critical interaction dynamics regarding how algorithmic friction and system disagreements can paradoxically act as a cognitive safety net that stimulates user agency, while identifying the systemic limitations that black-box models impose on sustained tool ownership.

2 Related work and theoretical foundations

Moving from information processing to meaning-making requires a shift in how privacy tools mediate the user's relationship with their data. The process begins with an evaluation of current browser-level solutions.

Extensive research has utilised black-box automation alongside rich data management and abstraction to reduce user burden and combat dark patterns through protocols like Generalisable Active Privacy Choice or tools like Consent-O-Matic [9] [10] [32]. Building on these foundations, recent studies leverage AI-driven systems for automated scraping and personalised explanations to bridge the transparency gap in black-box solutions [27] [17] [1] [9]. However, despite these advancements, interactions remain trapped in rigid popup conventions or become entirely decoupled from the browsing context. From an instrumental interaction design perspective, these tools function as disconnected utilities rather than true mediators.

While these approaches are frequently validated through high performance in traditional usability metrics, one may argue that such measurements are insufficient indicators of a solution's true quality. These browser-level tools often score well on second-wave metrics such as task-completion speed and error reduction, precisely because they automate the user out of the loop. However, by prioritising functional efficiency over the human experience, they fail to address the underlying erosion of agency. Consequently, high usability in this context does not signify user empowerment, but rather a streamlined path to the "Okay, whatever" effect, leaving the user's situated experience unaddressed [11]. To reinforce these observations, Duarte et al. [18] further elaborates on these points, highlighting the fact that although addressing the burden of privacy fatigue, such solutions often introduce complacency and foster a placebo-effect of security.

In our previous study, *From Learned Helplessness to Digital Agency: Evaluating Seamful Design Interventions in Consent Management Platforms* [29], we evaluated the short-term effects of a plethora of design approaches on user-perceived quality in CMPs. While our work produced valuable findings regarding the necessity of "seams" in automation to establish psychological safety nets and increase perceived performance, both the intervention and its evaluation remained confined to second-wave methodologies. By developing these solutions in a vacuum, we inadvertently overlooked the broader, solidified impact of everyday CMP interactions on the user experience.

Nevertheless, our findings, alongside numerous related studies [11] [26] [9], reinforce the notion that trust is not monolithic; it is built differently across individuals, making "one-size-fits-all" privacy tools fundamentally inadequate. Because privacy choices are deeply situated, they require a level of flexibility that current, rigid tools fail to provide. This is a limitation that prevents users from appropriating technology into the complexities of their daily lives; a challenge that is directly addressed by instrumental interaction design theory. As proposed by Beaudouin-Lafon, this paradigm suggests that digital tools should function as flexible instruments that users can manipulate and adapt to their context, rather than rigid containers that restrict user agency [3] [3].

Currently, CMPs rely heavily on declarative transparency, expecting users to read and synthesise abstract information. However, as Gibson's theory of affordances suggests, humans learn how the world works primarily through interaction [13]. By adopting a generative theory of interaction and reifying abstract privacy constraints into manipulatable objects, we could allow users to explore the system's boundaries through epistemic actions. Rather than treating consent as a static declaration, this instrumental approach transforms privacy preferences into tangible tools that users can actively probe and adapt. This direct interaction approach within a privacy instrument could, in theory, provide immediate, tangible feedback that builds a more robust sense of data ownership than static, pre-determined text ever would, as the user discovers system properties by probing it. This shift from passive observation to active exploration is not only a learning advantage; it is the fundamental mechanism through which user agency is reclaimed.

As explored in the paper, *Understanding Agency in Human-Computer Interaction Design* [21], the concepts of instrumental interaction design and agency are inherently linked; agency being defined as *the capacity of an actor to take intentional action and exert control within a digital environment* [21]. In the context of online privacy, this capacity is defined by a user's ability to make autonomous decisions that align with their personal preferences. Designing privacy tools as instruments rather than automated barriers may potentially transition users from passive objects of data processing to active agents. This shift suggests the possibility of fostering a stronger sense of digital ownership, possibly moving individuals past a state of passive compliance and toward active appropriation, where they might more freely shape the technology to fit their specific life contexts. Transforming the interaction from a passive encounter into a deliberate exercise of control and ownership is central to moving beyond purely functional interactions and addressing the broader socio-technical shifts that define the third-wave of HCI [21].

However, the transition from the second- to the third-wave of HCI introduces significant issues, as outlined by Susanne Bødker in *When Second Wave HCI Meets Third Wave Challenges* [8]. Bødker suggests that designers should not abandon second-wave theories, but rather refine them through two key principles; appropriation and mediation. By focusing on appropriation, designers can study how users integrate tools into their daily lives in ways that often

diverge from their original intent. Furthermore, building on the principles of instrumental interaction design, designers should not only view technology as a mediator of human activity but acknowledge that what is being mediated is often social connection or personal expression rather than just a job. Ultimately, Bødker argues for a situated understanding where context is dynamically created through user action, shifting the evaluative focus from functional usability to the experience.

3 Instrumental interaction design inspirations

While instrumental interaction design has yet to gain traction in the domain of online privacy, several state-of-the-art applications demonstrate its potential. For our system's design, we primarily draw inspiration from innovative interfaces such as DirectGPT [33], Memolet [35], StickyLines [4], and CPN2000 [2]. By examining how these tools allow users to bypass rigid, predefined workflows through direct interaction and appropriation, we can establish a blueprint for moving automated privacy management away from black-box automation toward a model of user-led, situated control. In this section, we provide a summary of the most notable ideas that inform our design approach. If you are interested in a rundown of all examples of reification, polymorphism, and reuse we found in other systems, see appendix A.

DirectGPT is a primary source of inspiration for us. This paper documents a new approach to LLM interactions, replacing traditional linear conversations with physical actions like clicking, dragging, and dropping. Our key takeaway from this paper is how the authors chose to reshape their users' hard-wired interaction models with physical actions like clicking, dragging, and dropping. Importantly, any object of interest is interacted with directly, and queries and prompts are turned into actionable items for later reuse. Furthermore, DirectGPT provides an outstanding example of reification through their prompt objects; instead of a static chat history, previous queries become actionable buttons that can also be executed on outputs, providing improved support for follow-up questions. This point also ties into reuse, as successful prompts are not lost. DirectGPT also uses polymorphism, enabling users to apply actions to different media, which provides more flexibility to their system.

The second example, Memolet, is an interface designed to address the linear memory problem in AI chat-bots. Our key takeaway from this paper is how it applies the principle of reification to turn conversational snippets into tangible, reusable memory-objects. Memolet applies polymorphism by having memos that are applicable in a wide range of contexts and promotes reuse by allowing users to manually curate and reuse specific memories across different sessions.

StickyLines gives a good example for shaping a workplace to accommodate direct interactions and increase autonomy. Instead of hiding alignment options below a sub-menu disconnected from the canvas, StickyLines reifies this solution by placing their tools right where the object of interest is located. StickyLines also supports polymorphism in their design by enabling their lines to adapt their behaviour based on how the user interacts with them or what objects are attached, and applies the reuse-principle through configurations.

Lastly, CPN2000 presents an innovative approach to colour palettes where, instead of moving the cursor to a menu, the user can move the instrument directly over the object of interest. CPN2000 promotes reuse by enabling macros, turning a complex series of edits into a single reusable instrument. A major takeaway from this product is the ability to stack functions and perform otherwise complex operations in a single action.

Other interesting examples include the reification in dragging and dropping files available in most operating systems, translating the semi-complex action of relocating resources into a more tangible action. Additionally, as the cursor changes, although the tool remains the same, the appearance is adjusted to the target interaction - a prime example of polymorphism. One gesture leads to multiple object-specific outcomes.Microsoft Excel's fill handle also gives a good example of a mix of principles, enabling the user to drag the corner of a cell to continue a pattern (1, 2, 3...), reusing the logic and data from the previous cells to generate new ones.

While some modern interfaces leverage tangible shortcuts, they often lack true polymorphism. For example, Instagram maps a long-press interaction to entirely different commands - previewing a photo on a profile, pausing a video on a story, or opening a menu in a messaging list. Because the underlying command changes rather than the object type, this represents a context-dependent mode mapping rather than polymorphism, which strictly requires a singular command to be applicable across diverse object types.

Regarding the LLM-driven integration of our approach, PRISMe (from the paper, *Helping Johnny Make Sense of Privacy Policies with LLMs* [7]) offers a significant technological precedent. This study demonstrates that while LLM-based tools can effectively increase privacy awareness and emotional engagement, noted by participants who remarked that visual cues like "a sad face does something to me emotionally". Such tools primarily function as advisory dashboards rather than active instruments. While PRISMe primarily assists users in the meaning-making process, our project extends this paradigm by reifying abstract privacy criteria and model outputs into tangible, directly interactive workspace objects. By translating data into manipulable artifacts such as sticky notes and reusable action items, our probe should enable users to actively negotiate tracking architectures and retain complete ownership over the final decision-making process.

PRISMe's findings indicate that users often struggle with inconsistent privacy criteria across websites; this provides empirical justification for our focus on reusable instruments, which should allow users to apply a consistent logic across diverse digital contexts. At the same time, these issues with criteria are strongly linked to leaving the user out of the loop when producing the assessments. In contrast, we aim to keep the users in control of their algorithms. Another major issue with this product was the lack of context-awareness. PRISMe is not aware of which document object model

(DOM) element the users are looking at and which website they accessed. Because of these issues, interactions are reportedly cumbersome and confusing at times.

As these papers and corresponding systems illustrate, the transition from abstract data management to tangible, instrumental interaction can yield higher degrees of user flexibility, control, and ownership. Combining this approach with an appropriate LLM inclusion, which is aware of the user in the loop, seems like a promising yet unproven approach that we aim to test. By applying the principles of reification, polymorphism, and reuse found in these precedents, we can begin to dismantle the rigid, black-box nature of current privacy automation.

4 Project aims and research question

This section details how we aim to move beyond usability benchmarks to investigate how such an instrumental interaction design approach can foster appropriation, ownership and care, providing an empirical path to reclaiming user agency in the context of online privacy.

In this project, we seek to adopt a third-wave lens, moving beyond functional metrics to explore how privacy technology fits into everyday lives, emotion, and social context. Through an innovative approach driven by instrumental interaction, we aim to move from information processing to a model of meaningful appropriation and user ownership.

Through an exploratory study, we aim to investigate how applying instrumental interaction design principles to AI-powered, browser-level privacy tools influences user agency. Specifically, we will examine the process of meaningful appropriation and mediation, exploring how users move beyond passive compliance to actively integrate and adapt these tools into the unique messiness of their everyday lives and complex network of personal values. By reifying complex CMP architectures and privacy policies into manipulatable interface instruments, specifically through spatial, colour-coded post-it notes and draggable action tokens, we seek to understand if this direct interaction fosters a sense of digital ownership and care that can effectively counter the "Okay, whatever" effect of learned helplessness. To implement the theoretical principles of instrumental interaction, this research focuses on three core terminisms - reification, polymorphism, and reuse [6].

Furthermore, highlighting the exploratory nature of this research, we aim to gain insights on how users' experience changes with this unique approach and what value they find in it.

While our study involves a functional browser extension, we do not frame this phase as a traditional usability evaluation of a finished product. Instead, following the methodology of Hutchinson et al. [31], we deploy our tool as a technology probe. According to this framework, a probe serves distinct, but overlapping, goals which define the core aims of this research.

We aim to use the tool as an exploratory technology probe to understand the effects of instrumental interaction design on users in real-world online privacy contexts, primarily on CPMs encountered during everyday browsing. We seek to uncover how users would navigate the tension between convenience and privacy when provided with an open-ended instrumental interface, and how it can be appropriated to serve as a mediator of personal goals. Secondly, by implementing the tool in-the-wild, we aim to field-test the technical viability of combining LLMs with instrumental interaction for supporting informed consent decisions. This is not to prove the readiness of the software, but to observe how the technology performs and fails under the unpredictable conditions of diverse web architectures. Importantly, we aim to move beyond existing design constraints and gather "inspirational data" that will inform the requirements for future third- or maybe even fourth-wave privacy solutions. By adopting the technology probe approach, we acknowledge that our current design is a non-commercial prototype. Our primary aim is not to validate the usability of the extension's current features, but to use it as a medium for co-discovery - learning alongside the user what it means to appropriate a privacy tool and what specific affordances are required to successfully reclaim digital agency.

4.1 Problem statement

Using an AI-powered, browser-level privacy management tool built on instrumental interaction design principles, can users make active decisions that actually align with their preferences? How does direct interaction with CMP interface data influence the way users appropriate these technologies, and does this appropriation foster a sense of ownership, agency and care that counters learned helplessness?

5 Design process

This section details the central features and interface design of our proposed probe, Cookie Slayer. For a full list of features and system components along with their completion statuses, see appendix B.

5.1 Feature design

Drawing inspiration and reflecting on our previous research into online privacy protection software, the diverse landscape of state-of-the-art solutions, and the practical applications of instrumental interaction design, we have implemented our learning in practice and developed our proposed probe.

Through a high-level rundown of the core features of our probe, this section outlines the measures taken during our design process to implement instrumental interaction design principles that support appropriation and mediation, and most importantly aid us in addressing the fundamental issue at hand; empowering users to actually make decisions that align with their preferences through ownership and care through their personal perception and use of our probe.

In light of the current trends in privacy protection tools and the extensive opportunities presented by LLMs, we have established our technical baseline as a browser extension providing AI-powered,

browser-level automation. By equipping users with such advanced, yet familiar technology, we expect the potential for meaningful appropriation and mediation increases significantly.

Based on our key takeaways from previous research on CMPs, we emphasise the importance of implementing a framework of psychological safety nets alongside seamfulness [16] and operational flexibility. By making automated systems modifiable, we aim to ensure that automation acts as a supportive mediator, allowing users to actively regain digital ownership and agency.

1. Promoting ownership through concern calibration and dynamic privacy choice recommendations
Upon installation, following the optional viewing of the provided introductory material, the user engages in a recommender algorithm calibration phase to establish their general privacy profile. This design choice is grounded in the behavioural clustering research of Dupree et al. [19], which categorises privacy personas into three major groups, namely marginally concerned, pragmatists, and fundamentalists. Rather than imposing an initial one-size-fits-all automated policy, our system uses this calibration to gauge the user's self-reported behavioural archetype. These preferences are then systematically applied through transparent automation, directly indicating recommended choices with icons placed on the objects of interest (e.g. cookie popups buttons). Going forward, the system functions dynamically by adapting to the user's choices by listening to their subsequent choices, while the underlying algorithm remains directly manipulable by the user itself at any time.

By approaching automation through the lens of these personas, the system employs goal estimation even from a cold-start scenario to provide dynamic, context-aware, automated recommendations that align with the user's likely intent, while remaining easily modifiable. This dynamic goal estimation model is loosely inspired by the intent prediction frameworks of Qu et al. [14]. Our goal is to enable the logic-based algorithm to provide increasingly accurate, context-aware recommendations that evolve from the initial semi-cold-start persona. This transition from broad archetypes to situated, real-time intent prediction ensures that the automation remains a responsive mediator of users' unique goals rather than a rigid, predefined script. Importantly, our system is designed with the objective to ensure that we do not act on the user's behalf in a black-box manner; instead, we provide a flexible psychological safety net. Users retain ownership of the automation, possessing the agency to manually override both the persona-based defaults and subsequent fine-tuning of the algorithm through everyday use. Essentially, both the underlying algorithm and the final decision remains in the users' hands, and they are free to make exceptions at any time.

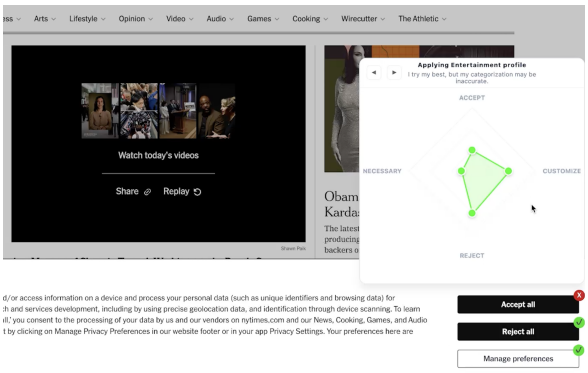

Fig. 1. Recommender interface

2. Direct manipulation through DOM object selection
Any website and internal system element can be interacted with directly through manual selection, by opening an LLM prompt field and using a list of actions, both suggested by us and created by the user. This feature is primarily intended for CMPs, but can in fact be used on any website elements on any website, including notes created by our system. This gives the users a big wiggle-room for appropriation, and lets them express any concerns they might have regarding a specific website, or following up on a response they might have more questions for. DOM elements and potential previous outputs are also used in the prompt-engineering process, so the responses are context-aware.

3. Actions powered by instrumental interaction design principles
When executing a prompt, user input is translated into reusable action items, which is executable in numerous ways - by dragging them on a website element, clicking them, stacking them in the prompt field, or following up on previous system output. This feature, inspired by the reuse principle, turns a static prompt history into actions that can be reused in future queries - similarly to DirectGPT and Memolet.

Additionally, all previous queries are accessible at all times, allowing users to shape their workspaces to their individual preferences, always having their most used queries at hand. The principle of reification plays an important role in this feature with respects to how prompts can be dragged or selected to appear in the input field, so the user can rework their previous actions for a new context.

Lastly, inspired by Memolet and StickyLines, polymorphism and appropriation come to fruition in how actions applied using multiple tangible interactions (e.g. clicking and dragging) can be performed on any website elements and notes created by our system. This feature also promotes directness and tangibility. Actions are performed directly on - and are aware of - the selected object of interest, and can be drag-and-dropped if they are in the taskbar. As a small extra detail, we decided to stylise saved actions with emojis to capture a complex query and simplify it into a symbol that the user can associate with their action, making further reuse easier.

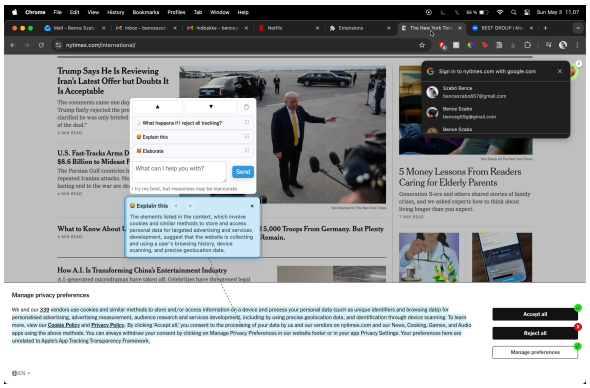

Fig. 2. Inspector tool

4. Post-it notes containing flexible answers combating current LLM interface shortcomings e.g. command-line-esque setup of current models

When an action or prompt is executed, a post-it note is created containing the answer linked to the selected object. This means that the tool-object relation is clearly established through visible links between notes and selected elements - similarly to StickyLines. Furthermore, implementing one of the key ideas behind Direct GPT, even though notes represent the output of the system, they remain flexible for further follow-ups and edits. The contents of notes which have been interacted with change based on the user's query, but interaction is reversible by the undo safety net. Importantly, notes can linger as long as the user considers them useful, or potentially for future reference - otherwise they can be closed.

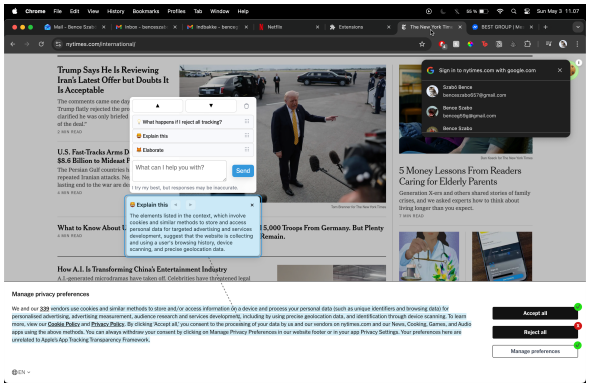

Fig. 3. Notes

A comprehensive breakdown of the client-side UI layouts, including the user onboarding flows, the recommender system dashboard, the inspector tool overlay, and the spatial tracking of interactive post-it notes mapped against usability heuristics, can be found in appendix C.

While these core components define the structural capabilities of Cookie Slayer, the system is flexible enough to support open-ended individual appropriation. To preserve the focus on our primary design choices, a detailed breakdown of the system's intended workflows, including baseline use cases and an exploration of the unintended yet possible examples of user appropriation across political, contextual, and linguistic boundaries is provided in appendix D.

6 System architecture

The technical implementation is split into a client-side browser extension and a server-side backend. This is done to ensure a clear separation of concerns while balancing user-side interactive performance with the significant computational requirements of our Gemma 4 LLM ran through Ollama, in a remote environment. The setup of the system can be seen in figure 4.

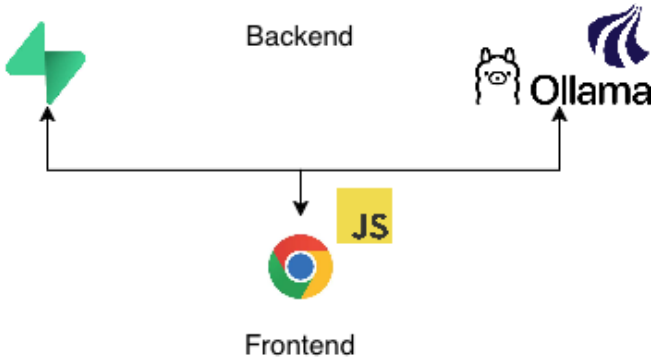

Fig. 4. Overview of the system design

6.1 Frontend (Browser Extension)

The frontend serves as the primary mediator for the user's interaction with the web environment. Its primary technical responsibility is the extraction and manipulation of DOM elements. To identify potential CMPs, the system employs a scraping heuristic that targets elements with a $z - index > 1$, operating on the assumption that privacy banners typically occupy the topmost visual layer of the interface.

Once these high-level containers are identified, the system utilises regular expressions (regex) to parse the inner text of button elements, mapping them to specific privacy decisions/categories (shopping, social, entertainment, productivity, health, finance and government). Beyond automated scraping, the frontend implements the custom inspector logic required for direct manipulation, allowing the user to select and reify any DOM element into a manipulatable object. This inspector element handles the direct manipulation aspects of the interaction, including the drag-and-drop interactions for prompts placed in the toolbar, and the dynamic injection of UI overlays onto native CMP buttons to display personalised recommendations based on the user's calibrated persona.

Additionally, the frontend interacts with the Supabase database and the Gemma 4 LLM by sending asynchronous HTTP requests

to their respective API endpoints via GET and POST methods. We have structured our application to be loosely coupled, rendering a potential shift to another service provider easy. During our development work, this approach saved us from a lot of headache when switching from Google AI Studio to Ollama due to Google's rate limits and costly service.

6.2 Backend (Ollama & Supabase)

The system's backend is divided into two functional areas running in parallel; an LLM serving service for real-time reasoning, Ollama, and a remote database for data persistence, Supabase. This infrastructure is designed to maintain high availability for the user study. Hosting these services separately also ensures that database-related processes are not subjects to bottlenecks by, for instance, a high number of concurrent LLM request processes.

6.2.1 Ollama LLM. Our LLM serving service, Ollama [23], is hosted on AAU Strato through OpenStack [24], which provides the necessary bandwidth and computational power to support real-time AI mediation. The model utilises specific prompt engineering to ensure that regardless of the user's input, the output is systematically steered toward privacy-centric evaluations and situated data ownership. On the other hand, AAU Strato comes with a couple of known limitations. For one, setting up an autonomous development pipeline is largely blocked by the numerous security measures, and more crucially, students are limited to CPU usage. The lack of GPU access is unfortunately a massive limitation regarding the potential complexity of the LLMs we are able to host with a reasonable response time.

6.2.2 Supabase. We utilise Supabase [15] as a lightweight, off-the-shelf remote solution to handle the logging of user interactions during the duration of the study. This includes capturing prompts and system actions, which are essential for evaluating how users appropriate the tool over time. The structure of the database is described in more detail in figure 5.

7 Methodology

With the design and development phases complete, the research shifts toward an exploratory user study designed to capture insights regarding our problem statement through real-world use done by real people. This methodology is centered on understanding how users weave instrumental interaction into the messiness of their daily lives.

7.1 Study type

This technology probe user study is designed to adopt a third-wave HCI perspective, with respect to our wish to elevate the evaluation of privacy tools to more current HCI concerns. Our key source of inspiration for the technology probe approach is Hutchinson et al.'s paper, *Technology probes: inspiring design for and with families* [31], originally published in the proceedings of CHI '03. Furthermore, this approach is most suitable to investigate deeply-rooted cultural issues regarding a loss of agency and learned helplessness. Rather than relying on traditional second-wave metrics like efficiency or task completion, the study uses a longitudinal, in-the-wild approach that captures how users engage with and appropriate the tool into their everyday lives and capture the contextual nature of privacy interactions. By avoiding predefined tasks and allowing interaction to emerge from users' own contexts and goals, the study enables observation of whether users make meaningful, preference-aligned decisions and how they appropriate the system over time. This approach is essential for examining whether and in what ways direct interaction with privacy mechanisms fosters a sense of ownership and control, and whether it can counter the passive "Okay, whatever" behaviour.

By combining these perspectives with the probe's development's instrumental interaction design-driven generative theory approach, the study aims to move beyond evaluating whether users can complete tasks, and instead investigates how they engage with, shape, and take ownership of an unconventional system. Additionally, we are interested in learning about how a flexible system can be a mediator of a plethora of personal goals and beliefs, which ultimately fuel a decision-making process.

The longitudinal user study is conducted over a period of one week. This design allows us to observe how user experiences evolve over time and how they naturally occur in everyday contexts. By moving beyond controlled laboratory settings, the study prioritises ecological validity and embraces the inherent complexity and messiness of everyday life. This is particularly important in the context of privacy, where decisions are shaped by situational factors, habits, and emotions rather than isolated tasks. Since our probe is vastly different from other available solutions, seeing how habituation can take effect as the novelty-effect wears off over time is another major factor in opting for a longitudinal study.

While this approach introduces variability and limits direct comparability across participants, it provides deeper insight into how agency and appropriation emerge in practice, on a personal level.

7.2 Practical information

The study includes 9 participants recruited to represent a diverse range of technical proficiency, privacy awareness, and browsing habits. Participants are required to use a browser compatible with the developed extension.

Before participation, all users provide informed consent and are briefed on the purpose of the study. They are informed that their interaction data will be anonymised and used solely for research purposes. Participants are also made aware that they can withdraw from the study at any time without consequences.

Participants are instructed to use the tool as part of their normal browsing activities. No artificial tasks are imposed, ensuring that interaction emerges from their own needs and contexts. This decision should reflect both third-wave HCI and instrumental interaction design principles, as it allows interaction to emerge from users' own goals and contexts, supporting meaningful appropriation rather than prescribed use.

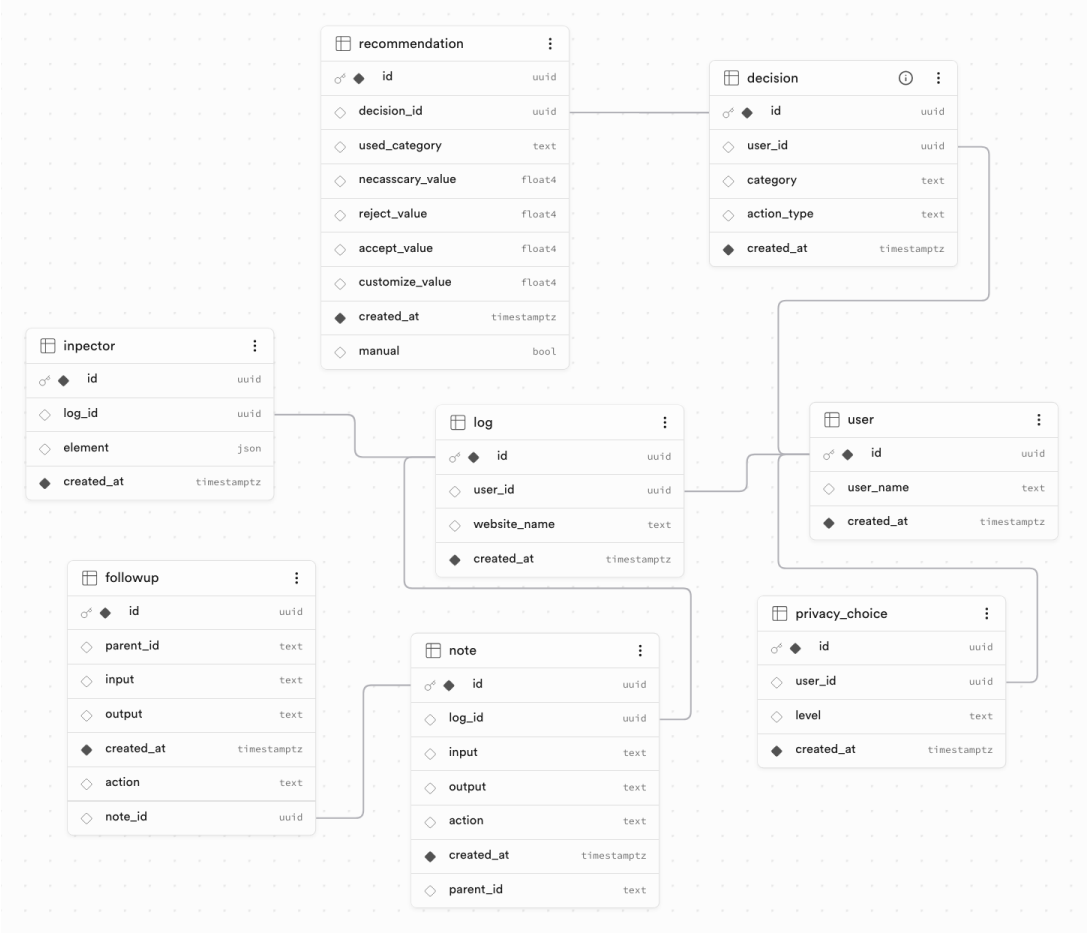

Fig. 5. Overview of the database structure

To ensure a problem-free installation process, users are provided with a document detailing the installation procedure. This document can be found in appendix I.

7.3 Procedure

The user study consists of three main phases.

Phase 1: Pre-study survey Participants take part in a survey aimed at understanding current privacy practices, attitudes toward consent mechanisms, and perceived level of control over their data. Some of the questions are inspired by the *Affinity for Technology Interaction Scale* [28] to assist our understanding of our users' likeliness to actively engage in or avoid intensive technology interactions. This phase establishes a baseline of users' experiences and expectations.

Our pre-study questions are designed to give us valuable insights on a personal-level regarding our participants' primary privacy concerns, contextual dependence, level of resignation, and perceived control among other things. Other than being able to map these characteristics to different concern-level groups, it also aids us in crafting more personal inquiries for the post-study survey in alongside our with the logs. For a complete list of questions, see appendix F.

Phase 2: Probe introduction Participants are introduced to the browser extension. Instructions are intentionally minimal to encourage exploration and support epistemic interaction, allowing users to discover how the system works through direct engagement. In case of any doubts that the revisitable tutorial does not answer, facilitators can be reached out to during the duration of the study.

Phase 3: In-the-wild use Participants use the tool during their normal browsing activities over a period of one week. No specific tasks are assigned, allowing the tool to be integrated into participants' everyday routines and contexts. Participants encounter the tool across different contexts such as shopping, social, or entertainment, providing insight into how privacy practices vary across situations. During this phase, interactions, prompts, outputs, and general usage metrics are logged. This data will aid us in proposing

new discussion points and observing tendencies among participants.

Throughout the study, we analyse the logs and do regular check-ins through private messaging channels, calls, or in-person to gather feedback and learn about participants' experiences on an individual level - similarly to a study diary, but with the possibility of immediate followups. The data gathered during these check-ins culminate in the the final question of the closing survey, primarily focusing on gaining further insights about personally interesting use-cases, tendencies, feedback, or other data points that could give us valuable insights regarding our problem statement.

Phase 4: Post-study survey At the end of the study, participants engage in a second survey reflecting on their experiences. Here, we aim to explore if the participants' use of the tool has evolved, whether they appropriated it in unexpected ways, and how it influenced their sense of agency and control. Additionally, our objective is to learn about the effects instrumental interaction design has had during their time with the probe. It is worth mentioning, that we adjust our initial questions to account for the feedback we get along the way, data we gather, and the users' discoveries shortcomings of our probe. For a complete list of questions, including personalised questions based on a mid-study analysis of the logs and previous survey responses, see appendix E.

7.4 Data collection methods

To evaluate the psychological shifts between baseline learned helplessness and meaningful appropriation, this study utilises a mixed-methods data collection approach. This combination allows us to cross-examine subjective user reflections (qualitative) alongside objective behavioural patterns (quantitative) across the one-week in-the-wild deployment.

The data collection architecture is explicitly structured to support our core research goals along three dimensions: verifying whether an AI-driven, browser-level instrument can enable active, preference-aligned decisions; examining how direct interaction with CMP metadata influences technology appropriation; and observing whether this appropriation fosters a sense of personal ownership capable of dismantling the passive "Okay, whatever" effect. By avoiding artificial, predefined laboratory tasks, this experimental setup prioritises ecological validity. Consequently, it allows us to gather empirical evidence that directly addresses our problem statement within the natural complexity of daily browsing routines.

7.4.1 System Logging. During the third phase of our study (in-the-wild use), the extension automatically logs user interactions to record how participants engage with the interface elements. Rather than collecting sensitive user data, the logging architecture is strictly scoped to record structural system events. The objective is to track the transition from passive automation bias to active, instrumental interaction, capturing the reality of user engagement without self-report bias. Each logged category addresses our core research aims:

System override frequency (countering learned helplessness) Reversals or mismatches of privacy recommendations. This verifies active decision-making, indicating whether users treat the tool as an instrumental extension of intent rather than defaulting to recommended passivity.
Epistemic clicks (fostering understanding) Tracks how often users explore or inspect system layers. This measures epistemic interaction, evaluating whether exposing abstract tracking mechanisms triggers the curiosity and care required to dismantle the effect of complex CMP wordings.
Prompt generation mechanics and reuse (appropriation) Prompt executions and the use of reusable interaction items such as action buttons. Tracking how these features are customised over time verifies the velocity of appropriation, showing whether users actively tame the technology for context-dependent routines.

7.4.2 Semi-structured questionnaires. The qualitative inquiries are structured around specific theoretical benchmarks to explicitly trace how users engage with the tool over time:

Mapping baseline resignation (pre-study): Probes emotional stressors, fatigue, and passivity caused by standard website CMPs, documenting where users felt forced into tracking defaults or trapped on a predetermined route.
Assessing the tactile shift and ownership (post-study): Evaluates how physically probing elements or interacting with recommender status symbols (ticks and crosses) shifted data ownership feelings, moving users closer to a domain they actively influence.
Evaluating system confrontation and overrides (post-study): Explores moments where users disagreed with automated recommendations and used the interface to "argue" with the system, serving as a direct qualitative test of active agency.
Reactive care and technical imperfections (post-study): Investigates user reactions to lackluster LLM outputs to evaluate if system imperfections trigger a protective form of reactive care, prompting users to take ultimate responsibility for outcomes.
Velocity of appropriation (post-study): Prompts users to identify the exact browsing session during which the tool altered how they navigate a site, pinpointing the moment the extension felt like it truly belonged to them.

7.5 Analysis approach

To interpret the potential psychological and behavioural shifts from baseline learned helplessness to active digital agency, we apply an iterative, mixed-methods analytical approach. First, qualitative text from the pre- and post-study surveys and mid-study diary checks are evaluated using an open- and axial-coding process. These codes were grouped into macro-themes - such as data proximity, cognitive safety nets, and systemic power imbalances; also accounting for the participants' technical backgrounds (high, medium, and low privacy literacy groups). This qualitative sweeping allows us to uncover conceptual thresholds and behavioural tendencies unique to each literacy tier.

Second, we triangulate these subjective insights with quantitative, objective interaction logs automatically captured in our database during the in-the-wild deployment phase. This log data serves as an empirical lens to identify interaction anomalies; instances where

a user's self-reported perception was disproven or validated by their actual system usage, mitigating self-report and recall bias. Finally, we synthesise these behavioural trajectories into high-level, colour-coded cross-comparison matrices representing user stances before and after deployment. Conducting a semi-direct comparison between these tabular artifacts should allow us to map individual empowerment trajectories and isolate precisely where systemic algorithmic boundaries restricted true tool ownership.

8 Results

This section presents the empirical findings gathered throughout our study, pairing baseline data with the behavioural logs and experiential insights observed during the deployment period. We first establish a user-typology baseline through an analysis of our pre-study surveys, subsequently evaluating how these distinct user groups appropriated the system, managed systemic errors, and negotiated tracking architectures across natural browsing contexts.

8.1 Pre-study insights

Based on the results of our shortened OPLIS questionnaire results, we have grouped our participants into three privacy literacy groups. This was done to gain a more nuanced view of our results and prerequisites with regards to individual concern levels. In this new grouping, there are two high-level (P6, P8), five medium-level (P2, P3, P4, P7, P9), and two low-level participants (P1, P5).

To give a quick overview of the most crucial differences between literacy groups, we have created a table seen in figure 6. The colour-coding of the field should represent the degree to which a behaviour, feeling, or status aligns with the user's true preferences or preferred actions. To present a high-level overview of our participant insights, we conducted an analysis on our qualitative data to identify, analyse, and interpret key patterns within the dataset (for the full anonymised dataset, see appendix F). To provide deeper context for the resulting user archetypes, this section details the distinct behaviours, emotional responses, and privacy attitudes of high-, medium-, and low-level participants. See appendix J for the full breakdown of the differences in results.

8.2 User study results

We concluded the study with a closing questionnaire. Below, we walk through the most significant insights we have gathered from our analysis of the responses. To see all individual responses, see appendix G and H.

Direct manipulation can aid in closing the data proximity gap

6 out of 9 participants reported that direct, tangible interactions, both with the cookie popups and the underlying algorithm itself, increased their sense of control and made their preferences more tangible and systematically enforceable. For instance, P6 noted that through these interactions, they *"felt more in control"* and found that their data-sharing preferences were *"a lot easier to maintain and not compromise"*. The specific interaction, described by participants, included being able to double-check the meaning of complex terms directly where they were presented, thereby implying a circumvention of the deceptive nudging tactics of unique CMPs through AI-driven content condensing. Furthermore, several users attributed their closeness to the data to the active critical thinking enabled by the tools at their disposal; as P7 summarised, the system *"changed in that way that I felt closer to my data, and who I let get to my data"*.

On the other hand, as some participants pinpoint, although they felt closer to their data, the known shortcomings of our LLM made it hard for them to verify the true outcomes, even when they made a choice that aligned with their preferences. The main pain points and drawbacks were the generic answers and the inherent issues that arise with only guiding users to customise their privacy options, then letting go of their hand. These downsides resulted in participants feeling less agency over their decisions. One participant saw the possibility of them exercising more caution, but without complete confidence, and two participants did not feel that direct, tangible interactions influenced their behaviour in any major way.

Overlaid privacy recommendations can inspire cautiousness toward consent choices

6 out of 9 participants indicated that the recommender system directly placed on the objects of interest made them more cautious, or felt that the mediation of their intentions was more straightforward. As P2 reflected on this shift: *"I did become more cautious, as to what website should have access to my data, and more comfortable, because i got honor my personal preferences"*. Interestingly, this participant also noted becoming *"more aware of the times, where I wasn't able to simply reject the cookies and I [had] to compromise with my chosen preferences"*. This highlights that users felt extra cautious in cases where none of the interface options perfectly matched their privacy profiles. However, when the system flagged conflicts across all available choices, it could also induce friction; for example, P5 shared that *"in the cases where all buttons were marked with red I felt somewhat lost and demotivated in terms of understanding"*. In these select instances, users frequently treated the visual indicators as a cognitive shortcut, remaining more mentally present in complicated scenarios where optimal choices were buried behind sub-pages or where websites proposed functional restrictions based on enforced preferences.

Elaborating on this point, another participant pinpoints that when the choice is straightforward, they would not oppose complete automation either - inspired by our previous study, a seamful implementation of this, while retaining some safety net features from the extension, could yield interesting results. Conserving user effort for highly problematic consent encounters where the enforcement of preferences is actively threatened, while automating straightforward scenarios could dramatically lower overall privacy fatigue. A comparative A/B study between these two interaction paradigms would clarify the exact behavioural and attitudinal differences that emerge. Such an evaluation could demonstrate whether maintaining a visible safety net during automated tasks prevents user resignation while still alleviating the friction of repetitive compliance actions.

Some participants also mention that although the prompt-driven interaction could feel tedious at times, if they felt like taking a

	High-Level	Medium-Level	Low-Level
Feelings about CMPs	Understanding to a certain extent, then hurriedness and annoyance	Frustration, waste of time	Impatience, irritation
Physical reaction to CMPs	Reading and selective choice, mostly rejection	Sometimes brief reading, somewhat selective choice, mostly rejection.	Choose the fastest option, mostly acceptance as that is the easiest option
Level of compromise	No compromise, avoidance	Mostly compromising	Total resignation
Distinction between site types	Systematic. Okay with internal-use of their data, but not with profiting off of them. More careful on shopping and finance-related sites.	Nuanced. Giving a pass to websites where they believe their data could be used for valid purposes. General dislike towards tracking done by social media companies.	Generally uncaring. Sites where potential material loss is perceived possible may recieve increased attention. Loose about social media tracking in hopes of more personalized ads and content.
Trust in the systems	They feel punished and unsure if their choices are truly respected	To a certain extent, they believe that their choices align with their preferences	Minimal
Perception of control	Somewhat in control	Mostly not in control	Total lack of control
Response to forced acceptance	Avoidance or frustration during time-pressure situations	Steering away, workarounds or acceptance	Acceptance
Degree of learned helplessness	Medium	Medium	High

Fig. 6. An overview of our pre-study insights on the different study-participant types

backseat, they could just rely on the simple visual cues. Although checkmarks and crosses were reported to be highly intuitive and aided users in unique ways, this approach introduces a potential danger of complacency. When interface elements become "invisible", users risk falling into passive automation bias, relying on simple visual shorthands rather than evaluating the underlying privacy data. For example, P3 highlighted this effortless reliance, noting: *"It was very easy, because you intuitively know to click on the green tick"*. This suggests that while colourful icons may effectively reduce interaction friction, they can inadvertently trigger peripheral cognitive processing. If users blindly trust a green checkmark as an absolute stamp of safety, they may bypass the critical evaluation needed to verify if the recommendation actually aligns with their specific privacy boundaries.

Based on the findings in table 7, users semi-successfully taught the algorithm to reflect their personal preferences, resulting in an overall 63.6% alignment rate between user decisions and system recommendations. Despite the system's inconsistent success rate, the high level of overall trust expressed in the post-study questionnaire suggests that participants successfully adapted their strategies to leverage the recommender for their specific needs. Interestingly, all accept decisions from users were in the start of the study, which suggests a display of learned behaviour of instinctively going for the easiest choice, seen in table 6. The recommended action for reject only had one mismatch the first day of the study, and otherwise all the reject user decisions matched the system-advised decision. Although more data would be needed to confirm this point, the logs suggest that the users felt it was easier to identify their true choice as they truly inhabited the system.

System inaccuracies dismantle passive compliance through reflective friction

8 out of 9 participants had examples for disagreeing with the AI assistant or recommender algorithm. Although most answers outline clear frustrations, it is clear that all users were ready to "argue" for their case, expressing high ownership over their decisions or resorting to sharing the least amount of data possible. This tendency can be linked to the concept of "rage-baiting". [25] Taking inspiration from Oxford University Press' definition (REF), our privacy-focused adaptation of this would be "vague or incorrect content provoking users to exercise heightened caution within a privacy decision". As Cunningham's law states, *"the best way to get the right answer on the internet is not to ask a question; it's to post the wrong answer"* [34].

This provocative friction actively broke users' passive scrolling habits; as P4 remarked, *"It definitely made me question myself (...)"*. Instead of yielding to the system, this confusion instigated deeper investigation. For instance, P5 noted that *"in the case of all cookie choices having red x marks I felt the need of understanding further exactly why this was,"* while P3 bypassed the automated suggestion entirely: *"It didn't understand what to do, so i just clicked to show details (...)"*. When the algorithm failed or was overly restrictive, users instinctively fell back on defensive privacy defaults, exemplified by P1 who shared, *"(...) so then I went with only necessary (...)"*.

The critical design implication extracted from this behaviour is that privacy instruments must be inherently dialectic, supporting counter-prompting, and low-friction correction mechanics. When users encounter an incorrect AI interpretation, their immediate behavioural impulse is to override and correct it. Therefore, interfaces should avoid rigid binary choice architectures and instead provide visible, physical avenues for disagreement, such as immediate override shortcuts, dynamic feedback loops, and polymorphic tools that let users explicitly reshape the recommendation algorithm on the fly. Designing for disagreement ensures that when the system inevitably fails, the failure mode actively supports user agency rather than causing frustration or abandonment.

Objective interaction logs expose user blind spots and self-report biases

Interestingly, when users were asked specifically about the above-mentioned behavioural shift, which was explicitly confirmed by our system logs through an increased number of prompts, choices, and self-reported behavioural data, the majority (6 out of 9) reported that they did not believe they had behaved any differently regarding the care or effort put into their decisions. This blind spot is neatly illustrated by P4, who noted: *"I would tend to move away from what the AI told me and return back to my usual routine"*, despite explicitly describing a behavioural shift in their previous response. This inconsistency between clear-cut interaction logs and user perception brings the classic issue of self-report bias into question, pointing

toward either a lack of introspective awareness or a heavy recall bias where users anchor their memory to prior web habits rather than their actual actions within the tool.

Another example of one such data point hinting towards the previously mentioned behavioural shift can be found in our logs; the number of followup questions clearly out-weighing the creation or reuse of previous queries, see table 1.

	Note		Follow-up	
Action Type	**Count**	**%**	**Count**	**%**
TYPED	47	70.1%	95	93.1%
CLICKED	17	25.4%	4	3.9%
DRAGGED	3	4.5%	3	2.9%
Total Valid Actions	**67**	**100%**	**102**	**100%**

Table 1. Comparison of Reuse Actions on Note vs. Follow-up post-its

Direct interaction shortcuts can shift website navigation and prompt broader privacy curiosity
5 out of 9 participants described that the extension changed their way of navigating websites. The listed examples include asking the AI assistant about certain websites' data policies and political agendas before accessing them; as P7 noted, *"Before even going into the website I would use the tool to see what they used my information for and if they had any political agenda behind them"*. Other participants used the interface to ask questions that they would otherwise have to open new pages for, or to request concise summaries of long legal texts. Remarkably, this structural directness also shifted some users out of lifelong passive compliance habits. P4 shared, *"For the first time I actually read the text inside of cookie banners instead of just straight out rejecting it"*. Generally, it seems like the directness of the tool, driven by instrumental interaction design principles, acted as a form of shortcut, enabling users to perform tedious compliance and auditing actions with significantly less effort.

An interesting remark by P3 is also worth mentioning here: "*I found out that I [encounter] less cookies than I thought I do*". As we have learned, looking out for your privacy is not always directly related to CMPs. As shown in the table below, only 46% of the prompts were targeted at cookie-related elements.

Interaction Category	**Click Count**
Cookie-Related Elements	170
Non-Related Elements	201
Total Interactions	**371**

Table 2. Inspector Tool Click Distribution

Personal value aligned customisation can foster ownership and serve as a congnitive safety net
8 out of 9 participants could describe a moment where the extension truly felt like it belonged to them, suggesting a transition toward ownership. A primary catalyst for this feeling was knowing the system was explicitly tailored to their values; as P3 noted, *"(...) the fact that [the recommendation] was based on my decision to protect my data was a point when I felt ownership over it"*. This tailored foundation allowed users to feel significantly safer, with the extension acting as a protective baseline. P6 illustrated this comforting presence, sharing: *"I felt a lot safer and that my data wasn't being unnecessarily shared. It felt like a friend trying to look out for me, in a way"*.

Additionally, participants embraced the tool as an active assistant for navigating cluttered CMPs and avoiding deceptive layouts, especially when optimal choices were intentionally obscured by the website. Within these complex architectures, the extension served as a cognitive safety net; P4 recalled that *"(...) the tool helped me reevaluate and click the right things so my choice aligned with what I actually intended"*, while P2 added that it *"does some of the hard lifting when it comes to the additional decision making that cookies impose"*. In contrast, the single user who did not experience this sense of natural interaction explained that it was entirely due to their pre-existing confidence in their personal privacy workflows.

Algorithmic transparency can help in demystifying online surveillance but highlights barriers to complete user agency
4 out of 9 participants believed they discovered something interesting during their time with the extension, noting that it was the first time they truly understood what online tracking meant. However, while the tool successfully heightened their structural awareness, it simultaneously exposed the deep systemic imbalances of the web, often leaving users without direct, actionable pathways. For instance, P1 reflected on this tension, noting that while the system *"(...) made me aware that I have a bit more power, it still felt too cumbersome to actually do something about it"*. This friction quickly translated into a sense of resignation regarding the current state of internet surveillance; as P4 states, the experience *"(...) opened my eyes to how alarmingly little there truly is I can do to protect myself when using the internet, and how regardless of my best efforts, some of me and my activity will always be monitored, analysed, and used for somebody's advantage"*.

To give a quick overview of the most crucial differences within and between literacy groups, with some aspect being compared to the pre-study questionnaire, we have created a table 7. The colour-coding of the field should represent the degree to which a behaviour, feeling, or status aligns with the user's true preferences or preferred actions. An interesting insight that can be observed on this table is how users naturally prefer and put more trust into the fully flexible and user-driven recommender algorithm in comparison to the black-box LLM. Although, to draw definitive conclusions, a more advanced language model should be used.

8.3 Personal followups

Throughout our study, we monitored our users closely and did regular check-ins with them to gather insights without subjecting them to having to keep a study-diary. For the final question of our study,

	High-Level	Medium-Level	Low-Level
Feelings about direct interaction	Mostly comfortable, frustrating at times	Easy, time-saving, at times inconvenient (LLM)	Promising, efficient, helpful
Level of compromise	Low, unchanged	Lower than before	Lower than before
Main pain points	LLM not helping with “custom settings”	Generic LLM responses	Generic LLM responses
Trust in the recommender algo.	High	High	Medium
Trust in the LLM	Low	Low	Low
Perception of control	Increased, more easily maintained	Felt more ownership regarding their data	More aware about their options
Response to conflicts (with the extension)	More caution	Extra care, arguing with the LLM or resignation to their own tools	Extra care or slight demotivation
Degree of learned helplessness	Lower than before	Mixed, on average lower than before	Mixed, showing slightly more concern
Examples of appropriation	Questions about why a particular ad is shown to them, and cross-site tracking related queries.	Preemptive questions before entering a website, simplifying all sort of legal texts (one privacy-related example being an EULA)	Learning about trackers and their general behavior

Fig. 7. An overview of our post-study insights on the different study-participant types

we chose to dive deeper into a single aspect of each person’s data (see appendix H), for instance unique use-cases or potential changes in opinions.

Overall, participants were eager to share their personal insights from the study and offer concrete suggestions on how to adapt the probe to better serve their needs.

A major takeaway from these insights is the participants’ newfound perception of cross-site tracking as creepy; a mechanism often considered far more invasive than sharing sensitive information directly with a primary service provider due to its stalker-like surveillance behaviour. For instance, P6 highlighted the unsettling nature of this boundary-crossing tracking, noting that because social media is so *"closely connected to one’s personal life,"* also highlighting that, *"random websites you browse should not know much about exactly what types of reels you watch the most"*. This systemic lack of transparency suggests a massive disconnect between corporate data practices and public awareness; as P8 argued, the public *"would be completely outraged if they actually learned what happens to their data, where it goes, and who really knows everything about them"*.

Furthermore, we got a clearer picture about how some users would like to be more knowledgeable about who their data is shared with, and not what data is shared, with the majority of the arguments being linked to political and cultural ties. Elaborating on this point, multiple users believe that a key feature that would give them more agency would be proper citations of up-to-date sources and website elements - P4: *"(...) I’d feel better about the AI guiding my hand in picking which privacy options I want, as it would understand up-to-date literature as well as my preferences"*. Given that state of the art AI is already able to do this, its implementation would not pose a massive technical hurdle.

Interestingly, two of the participants outlined an idea similar to an eventually scrapped feature of the probe, namely the seamful automatic application of batch actions. One of the proposed use cases was preemptively informing the user whether their preferences could be truly enforceable on a website (before entering the site), and if yes, by which option. Another idea was making seamful, automatic, override-able choices on the user’s behalf.

9 Discussion

This section contextualises our mixed-methods findings against the theoretical backdrop of online consent architectures, privacy fatigue, and generative instrumental interaction design.

9.1 Evaluation of our findings in light of the problem statement

The empirical findings from our closing questionnaire, when combined with our interaction logs, offer a nuanced answer to our core problem statement: whether an instrumental, AI-powered privacy tool can foster active, preference-aligned user decisions to counter learned helplessness. By implementing a probe grounded in generative instrumental interaction design theory, utilising direct, tangible interactions, and contextual visual overlays, the data suggests that our tool successfully lowered the degree of interface indirectness.

This structural shift allowed the majority of participants to migrate from a state of passive resignation to one of active agency. As demonstrated in Q1, 6 out of 9 participants reported an increased sense of control, explicitly exemplified by P6 feeling that their data-sharing preferences became *"a lot easier to maintain and not compromise"*, and P7 noting that they felt physically *"closer to my data, and who I let get to my data"*. This qualitative shift is quantitatively supported by our system logs, which suggest a a temporally-evolving 63.6% overall alignment rate between user actions and system recommendations in table 7, alongside purposeful personal adjustments made beyond the baseline dynamicity of the recommender algorithm. This interaction architecture enabled users to actively utilise the system for the mediation of their personal goals, transforming the otherwise overwhelming, opaque task of privacy management into a tangible, personally adaptable domain.

Crucially, the emergence of user-driven ownership and care was most evident when the system encountered operational friction or delivered inaccurate recommendations. Rather than reverting to a state of learned helplessness when faced with vague or incorrect AI behaviour, users leveraged the flexible interface to "argue" with the system and assert their unique privacy profiles. This behavioural phenomenon represents a privacy-focused application of Cunningham’s Law [34], where 8 out of 9 participants (Q3) successfully navigated algorithmic disagreements.

Instead of yielding to the interface, this provocative friction acted as a cognitive speed bump that broke passive scrolling habits, leading P4 to note that it *"definitely made me question myself"* and prompting P5 to seek a deeper understanding when confronted with a wall of red indicators. Users actively deployed the tool's instruments to override errors (e.g., P3 clicking to *"show details"* when the AI faltered), frequently falling back on defensive privacy defaults, such as P1 choosing to go with *"only necessary"*.

However, our results also expose a critical tension between automation and active reflection. While localised visual cues (checkmarks and crosses) directly mediated user intentions during straightforward choices (Q2), they introduced a clear risk of user complacency and automation bias. This vulnerability is perfectly captured by P3, who admitted: *"It was very easy, because you intuitively know to click on the green tick"*.

Consequently, the logs documenting that the vast majority of user choices aligned with system recommendations cannot exclusively be attributed to the system's flexibility; it also highlights a potential pattern of uninformed trust and peripheral cognitive processing. If users blindly accept a visual shorthand as an absolute stamp of safety, the instrument risks replacing legacy "cookie-clicking" habits with a new form of learned compliance, undermining the deep reflection required for genuine digital agency.

Furthermore, as highlighted by structural drawbacks like the LLM's common vagueness, simply guiding users to customise options before "letting go of their hand" is insufficient to dismantle systemic web power imbalances. While the directness of the tool acted as an empowering shortcut, true preference alignment remains tightly bound to the reliability and transparency of the underlying model. When the LLM delivered generic answers, participants experienced a localised collapse in agency, realising, as P4 states, *"how alarmingly little there truly is I can do to protect myself... some of me and my activity will always be monitored"*. This underscores that for an interaction instrument to fully conquer learned helplessness, it must couple actionable, direct manipulation at the UI layer with verifiable, transparent system outcomes from the underlying language model.

9.2 Evaluation of our key takeaways for individual privacy-concern groups

A granular evaluation of our user-type matrix (see figure 7) reveals that the capacity to make active, preference-aligned decisions varies structurally depending on the user's technical baseline (privacy literacy). For high-level users, who entered the study with pre-existing privacy paradigms, the tool served as an efficiency-maximising instrument that effectively lowered their level of daily compromise. Their appropriation manifested in advanced, proactive ways during browsing sessions, such as utilising the inspector tool to explicitly query cross-site tracking and hidden ad mechanics. By routing their intentions through a customisable recommender algorithm, these users experienced a pronounced reduction in learned helplessness, translating system friction into heightened caution rather than defeat.

Conversely, medium-level and low-level users experienced the most transformative shifts in agency. Previously paralysed by the opaque, multi-layered architectures of standard CMPs, these groups used the extension to simplify complex legal texts (such as EULAs) and demystify tracking behaviours. This behavioural shift is empirically substantiated by our follow-up logs (see table 1), which show that users generated 102 follow-up actions compared to 67 initial notes, with 93.1% of those follow-ups being explicitly typed inquiries. This high volume of text-driven interaction highlights a transition from passive compliance to an actively articulated search for system understanding; suggesting a shift from passive compliance to expressing more care and exerting heightened ownership over their data.

This interaction layer enabled users to break passive compliance loops, leading P4 to reflect that, *"for the first time I actually read the text inside of cookie banners instead of just straight out rejecting it"*. Furthermore, the tool's instrumental directness allowed users to expand their privacy auditing behaviour beyond traditional cookie dialogs; as shown in table 2, 54% of user interactions (201 clicks) targeted non-cookie elements, with users like P7 proactively deploying the tool as a gateway to check a website's underlying data policies and *"political agenda"* before even engaging with the page.

However, figure 7 also highlights a shared architectural ceiling to this empowerment: across all literacy groups, trust in the black-box LLM remained universally low. When the AI agent's guidance failed or turned vague at highly customisable configurations, users were abruptly reminded of their systemic lack of power. When confronted with these system conflicts, medium and low-level users occasionally exhibited a localised resurgence of demotivation or a resignation back to legacy tools. For instance, P5 reported feeling *"somewhat lost and 'demotivated'"* when confronted with conflicting red indicators, while P1 noted that despite an increased awareness of their own power, executing advanced choices still *"felt too cumbersome to actually do something about it"*.

This cross-group disparity suggests that while direct interaction with CMP elements may trigger the initial stages of user appropriation, the complete erasure of learned helplessness is entirely dependent on sustained system reliability. True tool ownership requires an unbroken chain of transparency; if the underlying model remains an unverifiable black box, users will ultimately drift back toward legacy tools or default to historical routines, as summarised by P4: *"I would tend to move away from what the AI told me and return back to my usual routine"*.

9.3 Advancing the state of the art in privacy fatigue and learned helplessness

Our findings advance the HCI literature on privacy fatigue and learned helplessness by moving beyond passive, educational notification systems and automated black-box tools. Prior research establishes that manipulative architecture on CMPs induces a chronic "Okay, whatever" effect, leaving users feeling structurally powerless against tracking. Traditional technical interventions attempt to counter this resignation either by shielding the user entirely through absolute automation or by bombarding them with dense privacy notifications. However, our study demonstrates that absolute automation merely replaces old habits of passive compliance with new forms of unreflective trust, while notification-heavy approaches aggravate cognitive fatigue. By contrast, Cookie Slayer advances the state of the art by demonstrating that embedding an active, instrumental layer directly on top of adversarial interface elements can systematically disrupt these deeply ingrained habit loops.

Crucially, this work introduces a paradoxically productive design paradigm to the privacy domain: the deliberate leveraging of algorithmic friction as a tool for cognitive empowerment. While classical usability heuristics prioritise seamlessness and the minimisation of user disruption, our empirical results show that when an AI recommendation tool exhibits localised inaccuracies or ambiguous states, it acts as a constructive cognitive speed bump. This friction effectively triggers a privacy-centric application of Cunningham's Law. Rather than inducing a resurgence of learned helplessness, system errors and disagreements actively provoke users to re-assert their agency, investigate underlying policy text, and defensively override the system to demand data minimisation. Consequently, we push the boundaries of current privacy design by showing that an "invisible" system can be counterproductive to digital literacy; instead, a dialectic interface that leaves room for user disagreement is what scaffolds genuine tool ownership.

Finally, this study refines our understanding of user-tailored privacy mediation across distinct technical baselines. Existing literature frequently treats privacy-fatigued users as a monolithic group suffering from identical modes of resignation. Our stratification of high, medium, and low privacy literacy groups reveals that the path out of learned helplessness varies structurally. While higher-literacy individuals leverage instrumental tools as efficient shortcuts to minimise data compromise, low- and medium-literacy users utilise the exact same tools to completely alter their web navigation workflows, transforming passive resignation into active domain pre-vetting. By mapping these distinct trajectories, this work proves that overcoming privacy helplessness is not dependent on achieving absolute corporate transparency, but rather on providing users with an adjustable, value-aligned UI safety net that restores their perceived behavioural control.

9.4 Study limitations

While our findings provide interesting insights into user appropriation, several methodological limitations must be acknowledged regarding the study's validity and long-term generalisability. First, the evaluation faces constraints in external validity due to the relatively small sample size ($N = 9$). Although this setup is not uncommon for qualitative HCI evaluations, the distinct behavioural splits between high, medium, and low-Level users suggest that a wider, more diverse demographic might yield entirely different appropriation patterns.

Second, while the deployment of a functional browser extension directly within the participants' personal browsing routines granted the study high ecological validity, it simultaneously introduced threats to internal validity. Because users navigated uncontrolled, real-world web environments, confounding variables, such as varying website complexity, pre-existing ad-blockers, or distinct structural layouts of unexpected CMPs, unavoidably impacted user frustration levels independent of our tool's design.

Furthermore, the study's brief timeline leaves the critical question of true habituation unanswered. The documented increase in user care, prompt frequency, and willingness to "argue" with the AI might partially reflect a temporary novelty effect or an artificial evaluation bias, where users feel a heightened sense of duty simply by participating in a study. Over an extended deployment, this active agency might erode into a different form of learned helplessness or complacency, especially given the low universal trust in the generic LLM responses.

Finally, our reliance on a mixed-methods approach exposed a stark self-report bias (as seen in Q4); users frequently claimed their baseline behaviour remained unchanged despite clear back-end logs indicating highly active follow-up questioning and customised actions. This mismatch underscores the difficulty of accurately mapping a user's internal cognitive shift away from helplessness using retrospective metrics alone, highlighting the need for future long-term longitudinal studies to observe how true habituation affects digital ownership over time.

9.5 Technical limitations

The most significant challenge encountered in our implementation is the inherent context-window and memory limitations of current LLMs [20]. While even the smaller models we work with are mostly decent at point-in-time analysis, their short-term memory restricts our ability to perform longitudinal goal estimation; a process where the system reduces uncertainty about a user's intent over time through repeated interactions.

In the context of third-wave HCI, where the focus shifts from functional information processing to the "felt life" and situated experience of the user, this technical bottleneck also becomes a theoretical one. A truly third-wave privacy instrument should evolve alongside the user, learning from their experiences alongside their environmental stressors and emotional contexts to move beyond rigid, rule-based automation. However, because the LLM often "forgets" the nuances of past appropriations, the interaction remains trapped in a series of disconnected, second-wave encounters. An approach we have experimented with, where we fed information

into our LLM from a database also proposed massive challenges; for example what we consider relevant information, and how do we compress it to not overload the model's context-window. Ultimately, this lack of persistent, experience-driven UX prevents the system from becoming a true sustainable and dynamic mediator of human activity. Instead of a tool that matures into a personalised partner, the system risks resetting the user's agency at the start of every new session, potentially undermining the very sense of digital ownership and care we aim to foster.

Another major technical bottleneck was the lack of standard naming conventions for cookies and trackers, combined with a structural inability to trace data due to the absence of fingerprinting. Additionally, the unpredictable DOM structures and element classifications found across various CMPs presented engineering hurdles that exceeded the scope of this project. The fact that these issues remain unresolved even in production-ready state-of-the-art privacy tools underscores the immense systemic complexity of modern web tracking and highlights the technical gap that current solutions have yet to bridge.

9.6 Future work

Based on the findings from the user study and the closing questionnaire responses, several new design alternatives were developed to explore how the extension could better support users in understanding and navigating cookie consent interactions.

Dark Pattern Detection Interface

One of the findings from the study suggested that participants became more aware of cookie consent interactions once the extension made direct visual changes to cookie banners. Rather than automatically ignoring popups, several users described how the extension encouraged them to pay attention to privacy choices and the structure of consent interfaces. One participant explained that the extension helped them quickly assess whether websites actually offered a visible "Reject all" option: *"I became more aware if there was a reject all option or not"*.

The study results also suggest that participants already perceived many CMPs as visually manipulative before using the extension, but that the checkmark and cross overlay aided in highlighting these dark patterns. One participant stated: *"I've always had a sense of being forced into a choice"*.

These findings from the user study motivated the development of the Dark Pattern Detection interface. The proposed design introduces contextual warnings that explicitly identify dark patterns directly within the browsing interface. Rather than exclusively relying on conversational AI explanations, the users are met with explicit warnings of these dark patterns.

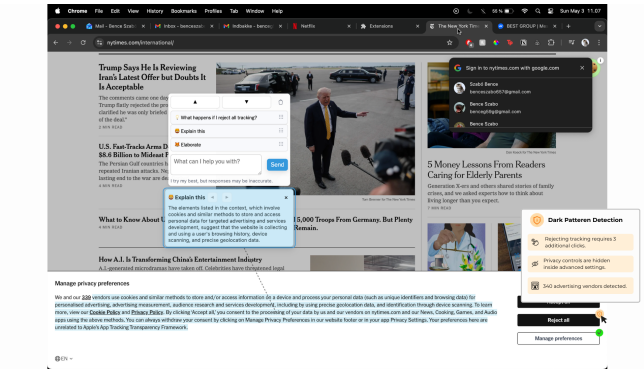

Fig. 8. Dark Pattern Detection interface

By exposing dark patterns directly at the moment of interaction, this design directly supports the problem statement by aiming aims to increase user awareness of how websites may influence privacy-related decisions and encourage more reflective consent behaviour.

"Where Your Data May Go" Transparency Preview

Another recurring theme in the study results suggested that participants often lacked, and expressed a demand for a concrete understanding of what happens after accepting cookies. Before the study, participants reported to be mainly focused on completing the consent interaction quickly rather than understanding the consequences of their choices. However, the extension displaying contextual privacy information on top of the cookie banner may result in users becoming more reflective about their decisions and spending more time examining consent options before interacting with them. One participant described how the experience changed their browsing behaviour: *"For the first time I actually read the text in cookie popups instead of just clicking accept"*.

To explore this opportunity, an transparency preview was designed. Before consent is given, hovering over the button reveals a pop-up describing the potential destinations and uses of personal data.

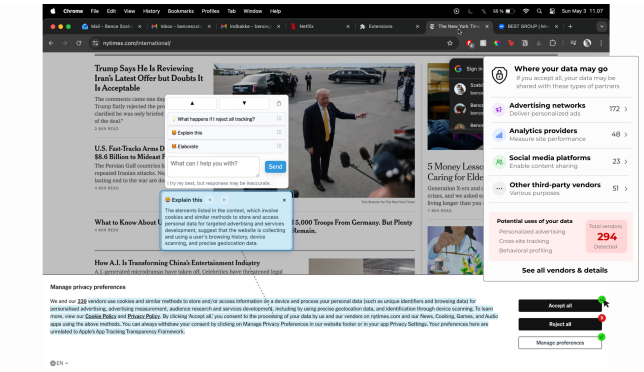

Fig. 9. "Where your data may go" Preview

"Who's Tracking You Right Now" Awareness Layer

A third design alternative was developed to further explore how tracking activity could be made more visible and understandable during browsing activities. While the previous mock-ups focused on identifying manipulative consent flows and visualising the potential consequences of accepting cookies, this concept aims to show relevant privacy implications directly at the moment of decision making.

This proposed design alternative therefore is aimed at gathering insights regarding the problem statement by improving transparency around data sharing practices and helping users better understand the implications of accepting tracking.

To address these shortcomings, the "Who's Tracking You Right Now" awareness layer shows a overview of active trackers directly within the browsing interface. The interface continuously displays how many trackers are currently active on the website and categorises them into groups such as advertising networks, analytics services, and social media integrations.

Fig. 10. "Who's Tracking You Right Now" Awareness Layer

This concept was primarily motivated by insights suggesting that participants responded positively when tracking related information became more transparent during browsing. The design focuses on lightweight visual transparency that remains accessible while users interact with website content. By integrating tracking awareness directly into the browsing environment, the design attempts to reduce the invisibility of online tracking infrastructures and support a stronger sense of situational awareness.

Overall, this proposed design alternative contributes toward the problem statement by aiming to increase transparency around online tracking practices and helping users better understand how their browsing activity may be monitored across websites.

10 Conclusion

To challenge the learned helplessness in the context of online privacy decision making, we designed and deployed an interactive browser extension probe that instantiates the instrumental interaction design principles of reification, polymorphism, and reuse. By transforming abstract privacy policies into manipulatable spatial objects (post-it notes), allowing queries to act as polymorphic, draggable instruments, and enabling the iterative reuse of contextually aware actions, the system succesfully altered how users negotiated their digital data boundaries.

Our empirical evaluation suggests that moving from an information-processing model toward an interactive, instrumental mediator can trigger meaningful appropriation and foster a distinct sense of digital ownership.

Furthermore, as our user-type matrix highlights, while high-level users appropriated the instrument to optimise existing privacy paradigms and audit advanced background mechanisms like cross-site tracking, the tool's impact was most profound for low- and medium-level users. For these groups, the extension acted as a critical cognitive shortcut and linguistic bridge, decoding dense legal prose and rendering tracking behaviours immediately visible. However, our findings also expose a delicate design tension. While localised visual indicators (checkmarks and crosses) streamlined straightforward choices, they simultaneously introduced risk of complacency.

Conclusively, our technical and empirical findings reveal that while direct interface interaction can combat hostile CMP design, true preference alignment cannot be achieved without verifiable trust and an overarching knowledge of complex network structures going beyond the accessed website itself. Future iterations of instrumental interaction design driven privacy-preserving systems must therefore move beyond black-box automation or isolated advisory dashboards. They must provide an unbroken chain of transparency, combining actionable, direct manipulation with verifiable data outcomes to permanently transform online consent into an enduring domain of personal control.

Acknowledgments

We would like to thank our supervisors Carla Florencia Griggio for her support and input throughout the project.

References

[1] Holger Szüsz Ala Sarah Alaqra. 2025. User Trust and AI-Enhanced Transparency: A Study on Cookie Banner UI Design. In *HCII 2025*. HCII 2025, 128–147. doi:10.1007/978-3-031-92840-6_8 Accessed: 2026-05-18.

[2] Michel Beaudouin-Lafon. 2000. Instrumental Interaction: An Interaction Model for Designing Post-WIMP User Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '00)*. doi:10.1145/332040.332473 Accessed: 2026-05-20.

[3] Michel Beaudouin-Lafon. 2020. *Towards Unified Principles of Interaction*. Youtube. https://www.youtube.com/watch?v=NvB_2vL1UmA Accessed: 2026-05-18.

[4] Marianela Ciolfi Felice Nolwenn Maudet Wendy E Mackay Michel Beaudouin-Lafon. 2016. Beyond Snapping: Persistent, Tweakable Alignment and Distribution with StickyLines. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. UIST 2016, 133–144. doi:10.1145/2984511.2984577 Accessed: 2026-05-18.

[5] Susanne Bødker Wendy Mackay Michel Beaudouin-Lafon. 2021. Generative Theories of Interaction. In *Transactions on Computer-Human Interaction (TOCHI)*. ACM 2021, 1–54. doi:10.1145/3468505 Accessed: 2026-05-18.

[6] Wendy Mackay Michel Beaudouin-Lafon. 2000. Reification, polymorphism and reuse: three principles for designing visual interfaces. In *Proceedings of the working*

conference on Advanced visual interfaces. AVI 2000, 102–109. doi:10.1145/345513.345267 Accessed: 2026-05-18.

[7] Vincent Freiberger Arthur Fleig Erik Buchmann. 2026. Helping Johnny Make Sense of Privacy Policies with LLMs. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. CHI 2026, 1–21. doi:10.1145/3772318.3791465 Accessed: 2026-05-18.

[8] Susanne Bødker. 2021. When second wave HCI meets third wave challenges. In *NordiCHI 2006*. NordiCHI, Oslo, Norway, 1–8. doi:10.1145/1182475.1182476 Accessed: 2026-05-18.

[9] Sebastian Zimmeck Eliza Kuller Chunyue Ma Bella Tassone Joe Champeau. 2024. Generalizable Active Privacy Choice: Designing a Graphical User Interface for Global Privacy Control. *Proceedings on Privacy Enhancing Technologies* 2024, 1 (1 2024), 258–279. doi:10.56553/popets-2024-0015 Accessed: 2026-05-18.

[10] Consent-O-Matic. [n. d.]. *Consent-O-Matic*. https://chromewebstore.google.com/detail/consent-o-matic/mdjildafknihdffpkfmmpnpoiajfjnjd?hl=en Accessed: 2026-05-18.

[11] Hana Habib Megan Li Ellie Young Lorrie Faith Cranor. 2022. "Okay, whatever": An Evaluation of Cookie Consent Interfaces. In *CHI 2022*. Association for Computing Machinery, New York, NY, USA, 1–16. doi:10.1145/3491102.3501985 Accessed: 2026-05-18.

[12] EDPU. 2020. *Guidelines*. EDPB. https://www.edpb.europa.eu/sites/default/files/files/file1/edpb_guidelines_202005_consent_en.pdf Accessed: 2026-05-18.

[13] James J. Gibson. 1979. *The Ecological Approach to Visual Perception*. Psychology Press Ltd. https://cs.brown.edu/courses/cs137/2017/readings/Gibson-AFF.pdf Accessed: 2026-05-18.

[14] Colin M. Gray, Jingle Chen, Shruthi Sai Chivukula, and Liyang Qu. 2021. End User Accounts of Dark Patterns as Felt Manipulation. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2, Article 372 (10 2021), 25 pages. doi:10.1145/3479516 Accessed: 2026-05-18.

[15] Supabase Inc. 2026. *Supabase*. Supabase Inc. https://supabase.com Accessed: 2026-05-18.

[16] Sarah Inman and David Ribes. 2019. "Beautiful Seams": Strategic Revelations and Concealments. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland, UK) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, Article 285, 14 pages. doi:10.1145/3290605.3300508 Accessed: 2026-05-18.

[17] David Martínez Aniol Molero Eusebi Calle Dolors Canals Ametller Albert Jové. 2025. Large-scale web tracking and cookie compliance: Evaluating one million websites under GDPR with AI categorization. In *Journal of Network and Computer Applications*. Journal of Network and Computer Applications. doi:10.1016/j.jnca.2025.104222 Accessed: 2026-05-18.

[18] Maximiliane Windl Roman Amberg Thomas Kosch. 2025. The Illusion of Privacy: Investigating User Misperceptions in Browser Tracking Protection. In *CHI 2025*. CHI 2025, 1–10. doi:10.1145/3706598.3713912 Accessed: 2026-05-18.

[19] Janna Lynn Dupree Richard Devriesg Daniel M. Berry Edward H Lank. 2016. Privacy Personas: Clustering Users via Attitudes and Behaviors toward Security Practices. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. CHI 2016, 5228–5239. doi:110.1145/2858036.2858214 Accessed: 2026-05-18.

[20] Gianluca Mondillo. 2023. Context Window: The memory limits of LLMs. Medium. https://medium.com/@gianluca.mondillo/context-window-the-memory-limits-of-llms-f11887390490 Accessed: 2026-05-20.

[21] Romualdo Gondomar Enric Mor. 2021. Understanding Agency in Human-Computer Interaction Design. In *HCII 2021*. Springer, Switzerland, 137–149. doi:10.1145/3491102.3501985 Accessed: 2026-05-18.

[22] Jakob Nielsen. 1994. *10 Usability Heuristics for User Interface Design*. Nielsen Norman Group. https://www.nngroup.com/articles/ten-usability-heuristics/ Accessed: 2026-05-18.

[23] Ollama. 2026. *Ollama*. Ollama. https://ollama.com Accessed: 2026-05-18.

[24] OpenStack Foundation. 2026. OpenStack: The Open Source Cloud Operating System. Official Website. https://www.openstack.org/ Accessed: 2026-05-20.

[25] Oxford University Press. 2025. The Oxford Word of the Year 2025 is 'rage-bait'. OUP Corporate News. https://corp.oup.com/news/the-oxford-word-of-the-year-2025-is-rage-bait/ Accessed: 2026-05-20.

[26] Tom Biselli Laura Utz Christian Reuter. 2024. Supporting Informed Choices about Browser Cookies: The Impact of Personalised Cookie Banners. *Proceedings on Privacy Enhancing Technologies* 2024, 1 (1 2024), 171–191. doi:10.56553/popets-2024-0011 Accessed: 2026-05-18.

[27] Marija Slavkovik Than Htut Soe Cristiana Teixeira Santos. 2022. Automated detection of dark patterns in cookie banners: how to do it poorly and why it is hard to do it any other way. In *Automated detection of dark patterns in cookie banners: how to do it poorly and why it is hard to do it any other way*. ResearchGate. doi:10.48550/arXiv.2204.11836 Accessed: 2026-05-18.

[28] ATI Scale. 2020. *Affinity for Technology Interaction Scale*. ATI Scale. https://ati-scale.org/ Accessed: 2026-05-18.

[29] Bence Szabó Louise Foldøy Steffens Sara Selman. 2020. *From Learned Helplessness to Digital Agency: Evaluating Seamful Design Interventions in Consent Management Platforms*. AAU. https://kbdk-aub.primo.exlibrisgroup.com/permalink/45KBDK_AUB/3b147k/alma9922389677905762 Accessed: 2026-05-18.

[30] Irina Shklovski, Scott D. Mainwaring, Halla Hrund Skúladóttir, and Höskuldur Borgthorsson. 2014. Leakiness and Creepiness in App Space: User Perceptions of Privacy and Mobile App Use. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) *(CHI '14)*. ACM, Association for Computing Machinery, New York, NY, USA, 2347–2356. doi:10.1145/2556288.2557421 Accessed: 2026-05-18.

[31] Hilary Hutchinson Wendy Mackay Bosse Westerlund Benjamin B. Bederson Allison Druin Catherine Plaisant Michel Beaudouin-Lafon Stéphane Conversy Helen Evans Heiko Hansen Nicolas Roussel Björn Eiderbäck Sinna Lindquist Yngve Sundblad. 2003. Technology Probes: Inspiring Design for and with Families. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI 2003, 17–24. doi:10.1145/642611.642616 Accessed: 2026-05-18.

[32] Holger Szüsz. 2025. *How transparency and UI/UX optimization influences user perspective: User perspective of cookie banner and cookie consent forms*. Karlstad Business School. https://www.diva-portal.org/smash/record.jsf?dswid=5748&pid=diva2%3A1992200 Accessed: 2026-05-18.

[33] Damien Masson Sylvain Malacria Géry Casiez Daniel Vogel. 2024. DirectGPT: A Direct Manipulation Interface to Interact with Large Language Models. In *Conference on Human Factors in Computing Systems*. CHI 2024. doi:10.48550/arXiv.2310.03691 Accessed: 2026-05-18.

[34] Wikimedia Contributors. 2026. Cunningham's Law. Meta-Wiki. https://meta.wikimedia.org/wiki/Cunningham%27s_Law Accessed: 2026-05-20.

[35] Ryan Yen Jian Zhao. 2024. Memolet: Reifying the Reuse of User-AI Conversational Memories. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. UIST 2024, 1–22. doi:10.1145/3654777.3676388 Accessed: 2026-05-18.

Appendix

This appendix provides supplementary data to support the findings discussed in the main text.

A Examples of instrumental interaction design

Examples of reification

- **DirectGPT (Prompt-Objects)**: instead of a chat history, DirectGPT reifies specific prompt syntax (e.g., "Add texture") into a toolbar of commands.
- **DirectGPT (Manipulable Outputs)**: the generated result (code, image, or text) is treated as a first-class object that stays at a fixed position, allowing for direct, incremental edits rather than full regenerations.
- **Memolet (Snippetization)**: reifies abstract "chat history" into tangible Memos (interactive cards). Instead of a long, monolithic scroll, specific pieces of information become first-class objects that can be moved and grouped in an active grid map.
- **StickyLines (Magnetic Guides)**: standard WIMPs (like PowerPoint), "alignment" is a command hidden in a menu. StickyLines reifies alignment into a physical-like "stick" that objects snap to. The relationship itself becomes an object you can click, move, and delete.
- **CPN2000 (Toolglasses)**: floating translucent palettes that reify commands. Instead of moving the cursor to a menu, you move the "instrument" (toolglass) over the object of interest.
- **Adobe Illustrator (Graphic Styles)**: after styling an object (shadows, gradients, strokes), you can drag that finished object into the "Graphic Styles" panel.
- **File Management (Drag and Drop)**: when you pick up an object and drag to a desired placement to drop the object.

Examples of polymorphism

- **DirectGPT (Localized Prompts)**: a single "prompt instrument" can be applied to different parts of an output. For example, a "Style" tool could be applied to a specific layer of a vector image or a specific function in a code block.
- **DirectGPT (Universal Undo)**: The "undo" mechanism is polymorphic, working across different types of media (text, image, code) produced by the LLM.
- **Memolet (Instrumental Usage)**: a single "Memo" can act as an instrument for multiple tasks—it can be used to inform a code generation task, then repurposed to provide context for a summary task without re-typing.
- **StickyLines (Flexible Lines)**: a single StickyLine can act as a vertical axis, a horizontal boundary, or a distribution spacer. It adapts its behavior based on how the user interacts with it or what objects are attached.
- **CPN2000 (Generic Tools)**: a single "style picker" instrument can be used on arcs, nodes, or text.
- **Instagram (Long Press)**: on a profile page, a long press on a photo previews it. On a story, it pauses it. In your DM list, it opens a menu. One gesture, multiple object-specific outcomes.
- **Cursor (Appearance)**: It clicks a button, drags a file, selects text, or resizes a window. The tool is the same, but the depicted interaction changes based on the target.

Examples of reuse

- **DirectGPT (Prompt Syntax Reuse)**: successful prompts are not lost in a scrolling chat; they are saved into a toolbar, allowing the user to reuse "engineered prompts" as rapid, repeatable tools for different projects.
- **Memolet (Persistent Context)**: solves the "short-term memory" problem of LLMs by allowing users to manually curate and reuse specific memories across different sessions, acting as a persistent instrument for long-term AI collaboration.
- **StickyLines (Sticky Lines)**: once you have "stuck" five objects to a line, that configuration can be reused. You can move the entire group by dragging the line, or copy the "sticky" property to another set of objects. It turns a one-time layout task into a persistent instrument.
- **CPN2000 (Input Reuse)**: command sequences can be captured and re-applied (macros), turning a complex series of edits into a reusable instrument.
- **Adobe Illustrator (Graphic Styles)**: the object captured in the styles panel is now a tool that can be applied to any new object.
- **Excel (Fill Handle)**: when you drag the corner of a cell to continue a pattern (1, 2, 3...), you are reusing the logic and data from the previous cells to generate new ones.

B System requirements and their completion statuses

B.1 Backend

B.1.1 Ollama LLM API [DONE].

(1) API endpoint for receiving prompt and returning LLM response [DONE]
(2) API endpoint for website classification [DONE]
(3) Appropriate prompt engineering for optimal answers [DONE]
(4) 24-7 availability hosting without bandwidth or other limitation bottlenecks (AAU Strato) [DONE]

B.1.2 Supabase Database [DONE].

(1) Save prompts for each user (input, translation, action)[DONE]
(2) Save outputs for each user [DONE]

B.2 Frontend

B.2.1 Extension Must-haves [DONE].

(1) Direct DOM element selection [DONE]
(2) Prompt input [DONE]
(3) Prompt output in form of notes [DONE]
(4) Translating inputs into reusable actions [DONE]
(5) List reusable actions [DONE]
(6) Object-selection aware prompt engineering [DONE]
(7) Scraping DOM for cookie buttons [DONE]
(8) Indicating personal privacy recommendations directly on the CMP buttons [DONE]

B.2.2 Extension Should-haves (Done).

(1) LLM responses should be directly manipulatable (follow-ups) [DONE]
(2) Actions should be polymorphic (draggable, clickable, revisitable, and reusable for any website element) [DONE]
(3) Recommended privacy choice should be based on dynamic user preferences [DONE]
(4) Recommended privacy algorithm should be easily manipulatable by the user [DONE]
(5) Recommended privacy choice should be grouped into website-type categories [DONE]
(6) Brief startup pop-up should be displayed on first use and should be revisitable [DONE]
(7) Introductory videos/gifs should play in the startup pop-up showing examples of use [DONE]
(8) Use scraped tracking data when performing prompt engineering [DONE]

B.2.3 Extension Nice-to-haves (Acceptable).

(1) Scraping the DOM for trackers (Scrapped) [Scrapped]
(2) Undo button (one button applicable for everything; partially implemented) [PARTIALLY IMPLEMENTED]
(3) Preset examples of actions [DONE]
(4) Batch actions (Partially implemented) [PARTIALLY IMPLEMENTED]
(5) Auto-apply action(s) (single or batch; scrapped) [SCRAPPED]

C User interface design

The UI for Cookie Slayer is designed to be approachable, lightweight, and highly intuitive, showing all relevant UI elements when/where they are intended for utilisation. Given that cookie banners often overwhelm users with complex choices, the interface prioritises simplicity, transparency, and immediate access. The design is guided by established usability principles, particularly Nielsen's 10 usability heuristics [22], ensuring that interactions feel intuitive.

C.1 Onboarding

Upon installing the extension, users are guided through a onboarding flow consisting of short, auto-playing videos that introduce the key features of the product. The onboarding flow emphasises clarity through the use of visual cues such as check-marks for highlighting recommendations, and the allowance of diverse physical interaction types (e.g. dragging, clicking, and hovering). This hopefully aids in reducing ambiguity and addresses any potential discouragement by allowing otherwise unconventional physical interaction types. This directly reflects principles such as visibility of system status and match between system and the real world, as the system is constantly communicating its internal logic through clear visual indicators, and users have the opportunity to apply familiar physical metaphors instead of navigating complex menus.

The playful cookie mascot visible during the introductory phase, along with the simple, conversational language and emojis, aim make the experience feel less technical and more engaging. Rather than presenting dense explanations, key benefits are broken down into short bullet points, ensuring that users can quickly understand the value of the tool without feeling overwhelmed. The overall structure also follows aesthetic and minimalist design, presenting only the information necessary at each step.

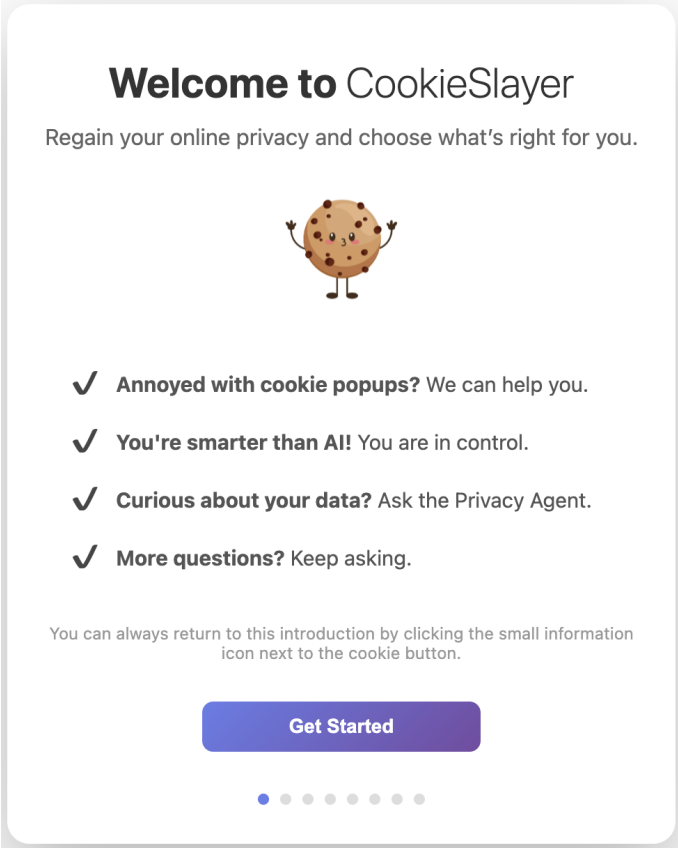

Fig. 11. Welcome screen

After completing onboarding, a small, draggable floating cookie icon appears within the browser interface. This serves as a entry point to Cookie Slayer's controls without disrupting the browsing experience. The icon is minimal yet easily recognisable, ensuring accessibility while maintaining low visual noise. This reflects user control and freedom, allowing users to access or ignore the tool at any time without interruption. By clicking the small "info-icon" attached to the floating cookie, the users are free to explore the introductory slides whenever they want to do so.

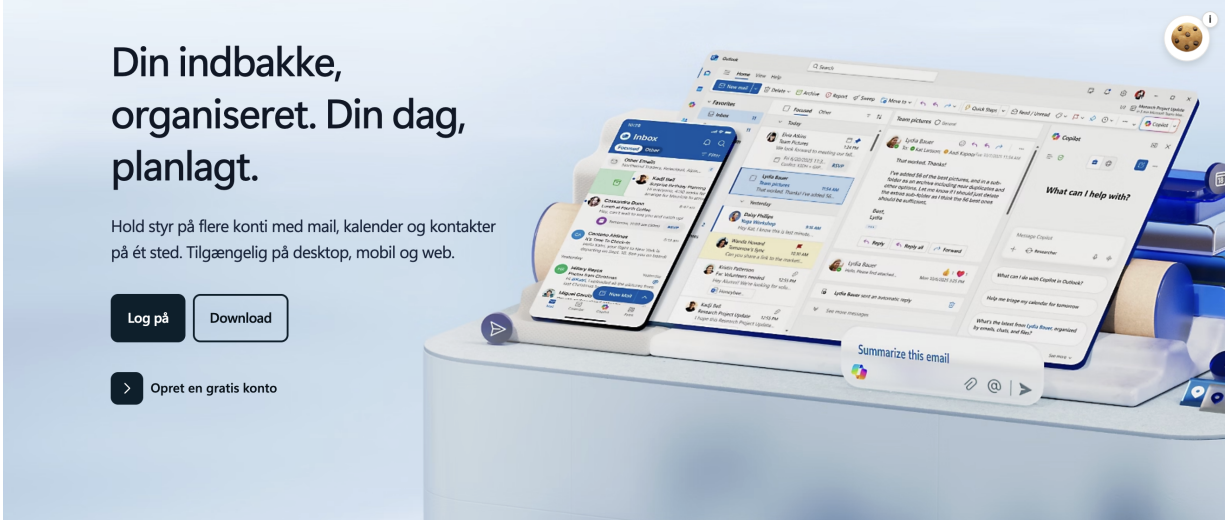

Fig. 12. Cookie Icon

C.1.1 Recommender interface. The main concern with the recommender interface was making it obvious to the user that their algorithm is tweakable without going into some messy sub-menu system. Therefore, we decided to make the inner-workings of the system directly visible when hovering over a recommendation indication (check-mark or cross). We have also attempted to translate the logic-based algorithm into a visual space providing a clearer overview upon first glance. Algorithms are abstract and math-heavy, so to reduce the degree of indirectness, we implementing feature with the key inspirations being the principles of visibility of system status, easy error recovery, and minimalist design.

C.2 Inspector tool

The inspector tool is arguably the most complex tool to navigate, so have therefore chosen to make it look as simple as possible. When selecting a website element, a box containing it is drawn for the client, and the input popup appears right next to it in a safe area of the window. Users are provided with some pre-determined example-questions, and system status is clearly communicated by a loading icon when waiting for a response, and a green flash upon successful response. Sending the prompt also accounts for multiple possible user inputs, such as clicking the send button or pressing enter on their keyboard, to ensure that any realistic use-style is viable. Our main considerations with this are flexibility and efficiency of use and error-prevention.

C.3 Notes

Notes, along with the corresponding text-highlights, are color-coded to make the distinction between multiple responses clearly visible. Each notes is also linked to the corresponding context using dotted lines, which disappear upon scroll, and appear again when stationary. This design decision aligns with the recognition rather than recall principle, By using dotted lines to physically link a note to its corresponding text, we are removing the need for the user to remember which note belongs to which highlight. The interface shows the relationship rather than making the user recall it.

D Intended use and appropriation

The system is designed to support users in making privacy choices align with their personal values and preferences with the help of safety nets placed right where point where those decisions occur. Its intended use centers on assisting users with active learning and acting on CMP user interface elements through a combination of automated recommendations, contextual explanations, and direct manipulation tools.

In a typical use case, a user encounters a CMP, receives a context-aware recommendation aligned with their calibrated privacy profile, and can either act on the suggested action or explore further by selecting UI elements on the CMP, asking questions, and optionally modifying the outcome. The system operates as both a decision aid and an interactive layer that augments existing web interfaces with transparency and control.

A more advanced intended use-case we consider would be asking the AI agent to make a recommendation based on some specific criteria, like *"I do not want to share my computer's MAC address with the owners of this website - is there an option that guarantees this?"* The reuse of such a prompt would be a reasonable action for fundamentalists.

At the same time, the design also moves beyond restrictive use and instead enables appropriation as a core quality of interaction. Users are not confined to predefined workflows but can reinterpret system features to fit their own goals.

A good example of appropriation we would expect is the exploration of political or financial connections of companies, and making privacy decisions based on these grounds rather than the facts regarding data collection. A concrete use case could entail accessing Spotify's website and using our tool to ask a question regarding whether Daniel Ek, who notoriously invested 694 million dollars into military tech, is still the CEO of the company or not. Establishing such connections could be an additional unintended safety net for those who are more likely to make their choices based on a case-by-case, pragmatic point of view. For such users, we would expect a higher degree of agency by their appropriation. Alternatively, questions regarding previous data-leaks would not be far-fetched either. Although this is not intended use, as this data is not directly related to the data available by scraping we do, it could be very relevant for some, and it would account for potential previous personal experiences regarding data-leaks.

Additionally, although revisitabilty of decisions is not supported by our tool, we could imagine that users would ask questions about similarities between different websites' data collection methods, likely ones they have made decisions on previously. We expect the AI agent to give somewhat helpful answers, but if this behavior is recognised as a valid tendency, this feature may we worth looking deeper into.

A user might also encounter a website in a foreign language with a complex CMP. Instead of using a generic page translator, users could ask the AI to give a simplified translation of the banner's text. The appropriation here is using the instrument as a linguistic and legal bridge to ensure their agency is not limited by their language proficiency, ensuring informed choice.

Another unintended but feasible use would be similar to PRISMe, where a user might create a reusable prompt along the lines of: *"Summarise this policy using only red flags and green flags emojis."* By reifying the complex output into simple emotional symbols, the user is appropriating the tool to manage their cognitive load. This would be most appropriate for the marginally concerned users, who are not seeking deep information; they are seeking a quick gut feeling to decide if the site is trustworthy enough to spend time on, effectively using the tool as an emotional gatekeeper.

A user might also use the tool to generate a one-time "mask" for a specific session. For example: *"I only want to be on this site for 5 minutes to read one article. Which settings will leave the smallest digital footprint for this specific short-term visit?"* This shows appropriation through temporal context, where the user adapts the tool's logic based on how long they intend to stay on a site, rather than a permanent preference.

Lastly, since we are not limiting the tool for cookie popups, we could imagine that users will also apply our AI agent to consent forms or other lengthy legal text. A concrete example could be looking at an EULA, marking the text, and asking our agent to *"explain this like they were five years old"*.

Not limiting the use of our AI agent to the appearance of cookie popups also means that even after a decision was made, users could ask subsequent questions about the choice they have just made, and how they could recover from a potential error. The system is in no way designed for this, but could function as a helping hand for regaining agency even after a decision has been "finalised".

It is also worth mentioning, that the AI agent could also function outside of the domain of online privacy. In theory, it could just be used as a floating, context-aware AI-assistant, providing a shortcut from copy-pasting a recipe to ChatGPT or Google Gemini, to simply selecting the text and asking it questions directly at the object(s) of interest.

A comprehensive breakdown of the client-side UI layouts, including the user onboarding flows, the recommender system dashboard, the inspector tool overlay, and the spatial tracking of interactive post-it notes mapped against usability heuristics, can be found in appendix C.

E Pre-study and post-study question lists

E.1 Pre-study questions

Could you walk me through the inner dialogue you have when a cookie banner interrupts your browsing? What is the immediate vibe or emotional reaction you feel toward the site in that moment? (Emotional stressors, learned helplessness and the "Okay, whatever" effect)

We often hear people say 'Okay, whatever' to these popups. In what specific ways does your current approach to these banners feel like a compromise rather than a choice that follows your preferences? (Learned helplessness and the "Okay, whatever" effect)

How does your care for your data shift when you move from a social media to a shopping or a bank's site? What makes one feel worth the effort of protection while the other does not? (Agency and importance of context)

When you make a choice on a cookie popup, how far do you go to understand what each option (e.g. "Accept", "Necessary Only", "Reject") actually does versus just picking one and moving on? (Affinity for technology and the "Okay, whatever" effect)

When interacting with privacy settings, some people feel like they are driving the car, while others feel like they are just passengers on a pre-determined route. Where do you see yourself in that analogy, and why? (Agency, ownership, and perceived pre-determinism)

If you were to spend five to ten minutes carefully selecting your preferences, for example hand-picking what to allow on a site, how much do you believe that effort would actually change how the site treats you? (Learned helplessness and affinity for technology)

Describe a situation where you felt "forced" to accept cookies even though you were uncomfortable with it. What was the specific pressure or friction that made you abandon your own preferences? ("Okay, whatever" effect, perseverance and importance of context)

Pre-study questions that stayed in the drafts
Do the political or cultural impact of specific sites and their personnel have an effect on your privacy choices? (e.g. Spotify CEO investing $600.000.000 in military defense company Helsing) (Importance of context - already addressed by other question)

Do you find yourself wanting to understand how tracking or cookies technically work, or is that something you prefer not to think about? (Affinity for technology - too on the nose, not super relevant)

E.2 Post-study questions

In our product, we focused on providing some examples of "direct interactions" - meaning that you could see recommendations, create notes and write prompts directly on top of cookie popups. (Intro)

In what ways did physically interacting with the website elements change how close or far away you felt from your data and potential outcomes of your privacy decisions? (Learned helplessness)

We have designed our cookie choice algorithm to be completely transparent and modifiable by you. (Intro)

Despite its numerous flaws, in what ways do you think this overlay (ticks and crosses) influenced the way you think about enforcing your preferences? (Ownership, care, learned helplessness)

Describe a time the tool made a recommendation that you disagreed with or surprised you. How did the interface allow you to argue with the system or assert your own preference? How did it feel to override the AI in that moment? (Agency, care)

A big part of our extension was the ability to ask questions. (Intro)

As you may have noticed, many of the answers our extension gave you were super vague. Interestingly, we have noticed a tendency in users expressing increased care regarding their privacy, exactly because they doubt the validity of the responses. If you have experienced something similar, can you describe it? (Being reactive to lackluster outputs and adapting to the new setting, care and ownership).

Throughout the week-long user study, we have noticed some cool ways you have used the extension we have not thought of. (Intro)

Can you describe a specific browsing session where the tool actually changed the way you navigated a website? What did you do differently? (Mediation)

Can you describe a specific moment when the tool felt like it truly belonged to you, rather than just completing manual tasks to contribute to the user study? In what moment did the extension become genuinely useful? (Appropriation)

When you used the tool to probe or interact with specific website elements, did you discover something interesting that you previously did not realise? How did that moment change your sense of who holds the power on that page? (Agency and ownership)

Do you have any comments? (Optional)

Post-study questions that stayed in the drafts Think back to the first day versus the last day. How did your relationship with the tool change as the newness wore off? (Novelty effect - interesting but not super important)

In what ways did you find yourself adapting the tool to fit a specific need we did not mention in the introductory material? (Appropriation - too on the nose, we can just look at data instead)

Many privacy tools work by hiding everything in the background. In contrast, this tool aims to keep the "seams" visible. How did this visibility affect your sense of being in charge versus just being a passenger to an algorithm? (Automation - already addressed in our previous project)

Instead of just clicking buttons on a banner, you were interacting with a mediated instrument. How did this change your confidence that your actual preferences - not just the site's defaults—were being respected? (Instrumental interaction design - too on the nose, we have a revised, better question in place)

If this tool were to become your "privacy companion," what physical or interactive features would it need to grow with you? How could it become more of a personal instrument for your values and less of a standard utility? (Not too important, address this with a comments part)

How would you describe your ability to be confident about following your privacy preferences, now compared to before using these tools? (we have similar question addressing this)

E.3 Personal questions

P1: Looking at your data, we have noted that the way you interact with cookie banners is highly dependent on your mood. Super relatable! In a world, where our extension was more feature-complete and smarter, one of the potential advantages would be that it would be much better at getting your vibe and attitude toward cookies - being much more personal. In this case, how would this more understanding, human-like partner in privacy change the way you treated and cared about your data?

P2: Looking at your data, we have recognized a pattern in how one of your goals in preserving your privacy is also putting a stop to large data-brokers making money off of your personal information. If they make money, then you should get a cut at least! In a world, where our extension was more feature-complete and smarter, one of the potential advantages would be that it could provide you with a verifiable trace of your data, and tell you where your data actually ends up, and who is profiting off of it. In this case, would you act to stop it, and if it proved feasible, how would it affect the way you think about your ownership of personal information in the online space?

P3: Looking at your data, we have noticed that you might agree with the statement, that "most online tracking is probably predetermined, so one's choices would not make a big difference". In a world, where our extension was more feature-complete and smarter, one of the potential advantages could be that it could give you highly accurate answers to whether your choice would actually make a difference - using proper citations to the website's inner workings (code), legal documents, and other verifiable sources. In this case, would you feel like keeping your data to yourself would be less of a lost cause?

P4: Looking at your data, you were one of the participants that asked the most questions. Curious one! In a world, where our extension was more feature-complete and smarter, one of the potential advantages would be that it would cite sources and search the web for accurate and up-to-date answers. In this case, how would the extension affect your confidence in the choices you make?

P5: Looking at your data, you really want to get through cookie popups as fast as possible. Speedrunner type! In a world, where our extension was more feature-complete and smarter, one of the potential advantages would be that it could make a choice on your behalf and let you know what it meant for you - for example automatic reject-all, and a text saying that this site cannot track you in any way. In this case, would you feel more in control about your data?

P6: Looking at your data, we have noted that you are concerned with cross-site tracking, especially when it comes to other sites seeing your social media activities. Super valid concern! In a world, where our extension was more feature-complete and smarter, one of the potential advantages of this product could be that it would be able to let you know exactly how each website would be able to track your activities on other services. In this case, how would that change your sense of control and/or feelings towards the way you manage online tracking?

P7: Looking at your data, we have noted that you were curious about the political and financial connections, potential controversies

(e.g. data leaks), and other cultural ties of websites even before you visited them. Very woke! In a world, where our extension was more feature-complete and smarter, one of the potential advantages of this extension could be that it would actually give you a good overview of all of these things. In this case, how would this influence your decisions? Would it have an effect on your sense of control and care for your personal information?

P8: Looking at your data, we recognize that you have a tendency to avoid websites where you cannot enforce your preferences. But finding out about this can take a looong time! In a world, where our extension was more feature-complete and smarter, one of the potential advantages would be that the extension could automatically apply a prompt, and tell you from the get-go whether a site is compatible with you. In this case, how would this increasingly automated solution influence your daily browsing processes, or your sense of control over your time spent on enforcing your privacy?

P9: Looking at your feedback, we have noted that you use many of the same websites during your daily browsing sessions, and that you tend to retroactively clean up your traces (trackers, cookies, etc.). Smart move! In a world, where our extension was more feature-complete and smarter, one of the potential advantages could be being able to learn about where your data actually washes up at the end of the endless data-broker operations. In this case, would you feel empowered by learning exactly where to step in to stop the leakage of your data instead of doing routine tasks like tracker cleanups? Would this give you more ownership over your data?

F Pre-study questions and OPLIS questionnaire results (Anonymised)

Study about online consent

10 responses

Publish analytics

Information about the study and informed consent form

What is your name? (first name is enough)

10 responses

Certificate of Consent

Copy

10 responses

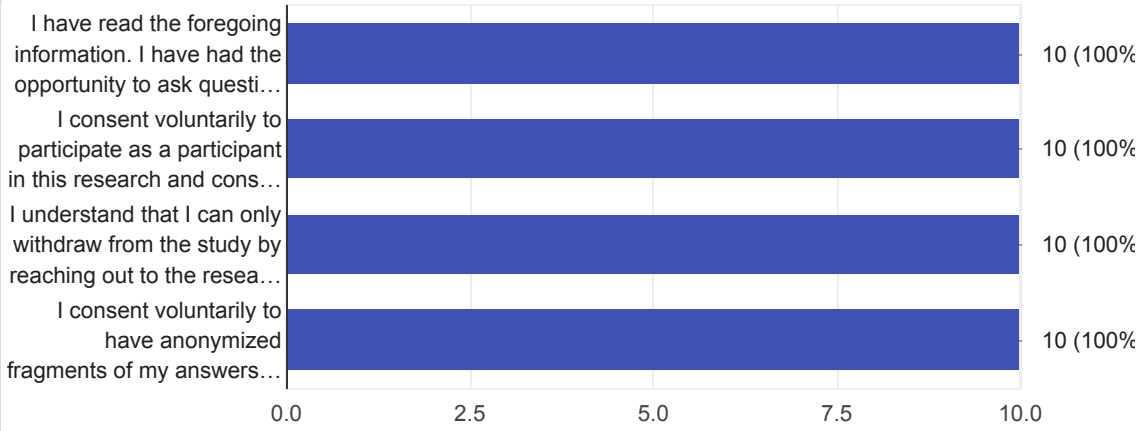

Questions about your experiences with cookies

Could you walk me through the inner dialogue you have when a cookie banner interrupts your browsing? What is the immediate vibe or emotional reaction you feel toward the site in that moment?

10 responses

Negative emotions. I dislike the interruption.

My initial reaction is always annoyance, even though I understand and appreciate the function of them. But I get so annoyed every single time because most cookie banners make it so painfully difficult to switch off non-necessary cookies, and even when it's simple, it's forcing me to read things I don't care about and do extra clicking around.

I think it depends on what mood I am in and what website I am on. Before, maybe a year + ago, I always clicked accept all on autopilot, not really thinking about it. Now, I often click deny all if I see the button. Sometimes they have the preferance alternative instead of deny all, then I look for the quickest way to close the banner without accepting all, which is often a button like "accept only essentials" or something like that. If I am in a very lazy mood I just click accept autopilot, just wanting to close the banner as quick as possible. In general, my inner dialog is "how can I close it the fastest with accepting the least amount of stuff" and sometimes "I just want to close this, and i do not care how".

Also I think I have a thought that it does not really matter if I accept/reject, like technically I do not want to accept because I know that teqnically it is bad to share data etc. but at the same time I think "how bad can it be" "does it really matter" and "whatever, I have already accept 10000x before" (this is not something I think while the popup show up, but just a general reflecting on my approach to cookies).

Jeg vil helst sige nej, det wack når det ikk er en mulighed

I feel annoyed - because 99% of the time i see no benefit for me that some goofy aah company needs to know my browsing history, age, name, mothersmaiden name, location etc. I understand that its good for business but me as a consumer does not like it.

I mean, I get that it's necessary. Depending on how intrusive it is, I don't mind it so much. If it's a banner or something similar, it's okay. But if the whole screen is taken up by a pop-up and it's made difficult to navigate and to reject it, then that annoyes me.

it is annoying, i want to make sure that i get none or only neccesary cookies, but it is time consuming

I feel somewhat irritated but I immediately look for the "reject" button. If I have to unclick all tags for "sending data to a third party" etc. I will unclick and make sure these settings are saved correctly. I take my time doing this even though I am hurrying. Except for e.g. when I have been in queue for tickets to a concert that is almost selling out. Then I honestly will allow all cookies and hope for the best.

I always click only share the necessary information, its irritating but is a need for the website to work. It kind of throw me of from what i was doing.

I find it irritating and i tend to become quite impatient. When that happens, i quickly look for any way to make it disappear. Often without really thinking through which option i click on.

We often hear people say 'Okay, whatever' to these popups. In what specific ways does your current approach to these banners feel like a compromise rather than a choice that follows your preferences?

10 responses

I generally dislike accepting cookies on websites I don't trust, so it's quite the compromise to me.

I compromise my time and patience because I tend to go to great lengths to click through everything and turn off all non-necessary cookies.

I totally compromise. I technically do not want to accept the cookies, but I am also lazy and just want to visit websites without interruption or delay, therefore, I accept a lot more than I would like. So it is a priority thing where I prioritise a more efficent/efforless web experience over thinking about the data I share.

Vil helst bare undgå dem

a lot of the time it is a compromise as a lot of websites will not let u continue if u do not accept their cookies. if i can i always reject - though i will say i have a hard time believing that that actually truly rejects all cookies. (conspiracy coded)

I don't feel like it's a compromise, I usually take the time and I get why they need them. Unless the site is purely just out for my data or for marketing purposes, then I don't mind it. Though I have been to a few site where I have to pay to disable cookies, and those get blocked super fast on my end.

i try to get none, but sometimes i can be lazy and think whatever

Even though I unclick all tags for the third party services, I still feel uncertain about if these choices actually make any difference.

It feels like a compromise because you have to share some data with the site for it to work appropriatly, where i would rather not

Well, i would like to be able to reject most of them, but it’s rarely an option anymore.

How does your care for your data shift when you move from a social media to a shopping or a bank's site? What makes one feel worth the effort of protection while the other does not?

10 responses

I would be inclined to think that government agencies (skat.dk) uses my data for something productive compared to a random shopping site I stumble upon.

I understand that I have non consensually and unknowingly sold so much of my personal information to Meta and who knows who else at this point, so I no longer feel protective over these. Obviously my bank details and similar feel more sensitive because I don't want people to have access to my money, but in principal, basic personal information should have the same level of care as bank info. I only feel like it's not worth protecting is because it's a lost cause now.

Not sure what you mean by the question. Of course I care more of how my data at skat / bank is protected than at social media, because if someone access my bank data they can steal money from me etc. but that does not relate to whether I accept cookies or not or how I use the website. I trust that the bank / goverment has sufficent security measures, and if not, i expect they would compensate if any harm caused.

In theory I would love if no company / people had access to my data, however, by using social media etc. I have accepted that some people (or an algorithnm) have some data about me, and can use that to ads, recommended content etc. however, it does not harm me directly. I care more for my card information much more than my search log history. However, if my search log history was publically available, I would care, but if a random guy in the US at Meta has it, I do not care.

Har det ok med at sige ja til feks online shopping algo, men alt andet er et pænt nej tak

same - no one needs to know what i do and i don't think cookies are good, unless its like "we are saving ur language preferences for said website so u dont have to switch from FR to DK" but anything in the ballpark of looking at browsing history, seeing how i interact with their website, how long i stay and worst of all "optimizing advertisments" needs to get removed IMO.

Because the sites need the data for different reasons.
The banks primary income doesn't rely on my data and they loose a lot by me not becoming a customer because they use my data incorrectly. The bank also has a reason to need my data other than just to sell it.
Where social media needs my data because they need to sell it and benefit from it.
One (social media) has hidden intentions with my data while the other (the bank) usually states quite clearly why they need the data.

In any case, I usually don't give much of my data unless it's specifically stated why they need it and I agree to that reason in either case.

when you are doing more serious tasks and dealing with more sensitive information, you want to feel more safe, i will give more effort

Først og fremst, wow sikke en summe hahaha!!!

I feel that e.g. using online shopping pages (which I actually don't use so much of) I always decline. I protect my personal data as much as I can e.g. by not always writing my full name in the hope that the carrier of the package does not see my full name. In these cases however my address is almost always on the package so it is often not worth the hassle. Regarding cookies, whicever type of site I am visiting I always try my best to disregard the cookies. Even though that officials (myndighederne) have all my personal information, like personal id or how much money I own, I still know that they don't own/control the browser so I have the same approach as to any other webpage. Decline cookies. But honestly I still don't know what the cookies indebærer, and I also haven't read up on it. I just know/hope that it is smart to decline.

I dont know if i understand the question correctly, but i would rather have that my banking site is better protected than other browsing websites.

I think it's mostly because, one site contains more of my personal information than the other. One of them gains/have access to my creditcard information (in a way?) the other has, at most, my email.

When you make a choice on a cookie popup, how far do you go to understand what each option (e.g. “Accept”, “Necessary Only”, “Reject”) actually does versus just picking one and moving on?

10 responses

I generally don't go to any lengths to understand. I think my cookie preferences are largely trust based. I usually accept on websites I frequent, and for others I reject.

I read up on everything when cookie banners first became mandatory on websites to understand what I'm reading and clicking on. I always reject everything that I can and only give the absolute bare minimum when visiting any site.

I do not understand. I have never looked into what the cookies actually does. I have a programming background, so I probably should know more than I do, but I do not. Other than it decide what data I will share. But I have no idea of what data is shared and not, and what is in each of the options. I never used time to look into it. I think that if it was really bad, it would be illigeal, so it cannot be that bad.

Not far

always reject, but as mentioned earlier i don't believe in companies and i try to delete all cookies at least once a week. ATM it is not feasibly to actually read what theyre tracking or not as they have too much text, and i need to look at this assorted website right now.

Not super far. I feel like a I have a pretty good understanding of them from working with cookies on the other end as a developer. But I also know that not every company has setup their cookies completely correct and I should use more time on it if I really cared. But I just press reject without reading further most of the time anyway.

i try to click reject most often, but sometimes it is not there in the forst popup, so i just press neccesary

Not far. If I see "reject" I choose that. I choose the most strict option. So if "reject" is not an option I choose "necessary only". And if "necessary only" is not an option I might choose to leave the page.

Earlier i just clicked accept all, now my choice of study has made me smarter and now i choose necessary. But before i didnt do anything to understand what info i was sharing with teh site.

Think when cookie popups first began, i tried to read through some of them - i gave up over time. Haven’t done it in years.

Quiz time!

Social network site operators (e.g. Facebook) also collect and process information about non-users of the social network site.

Copy

10 responses

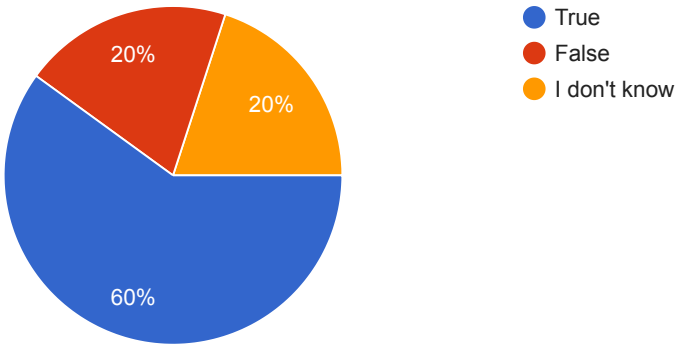

User data that is collected by social network site operators (e.g. Facebook) is deleted after five years.

Copy

10 responses

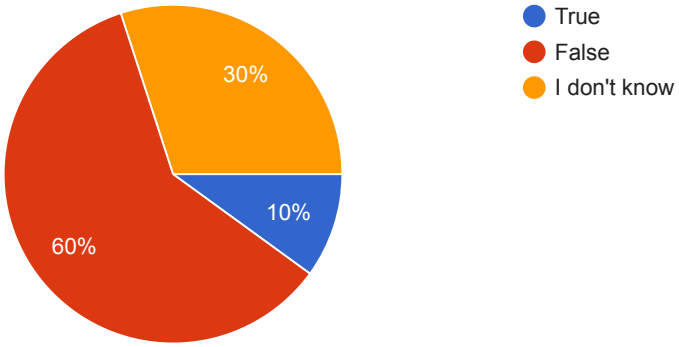

Companies combine users' data traces collected from different websites to create user profiles.

Copy

10 responses

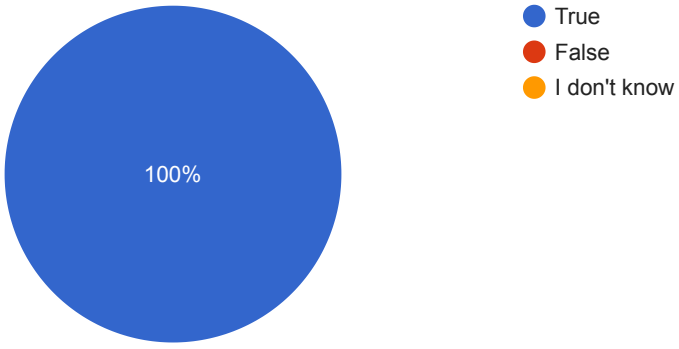

Copy

What does the term „browsing history“ stand for?

In the browsing history…

10 responses

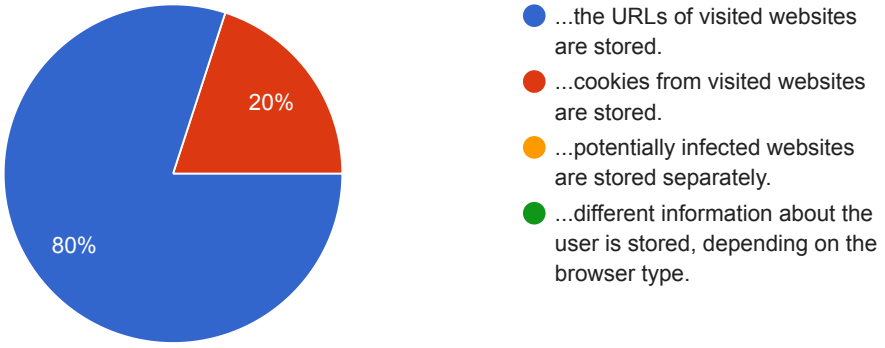

Copy

What is a “cookie”?

10 responses

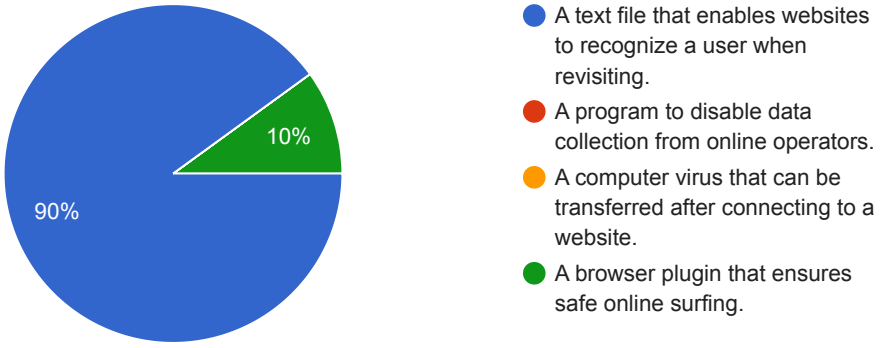

Copy

What does the term “cache” mean?

10 responses

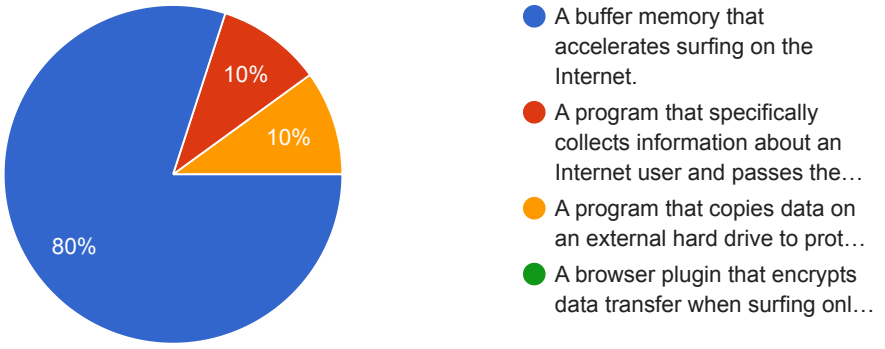

Copy

Forwarding anonymous user data for the purpose of market research is legal in the European Union.

10 responses

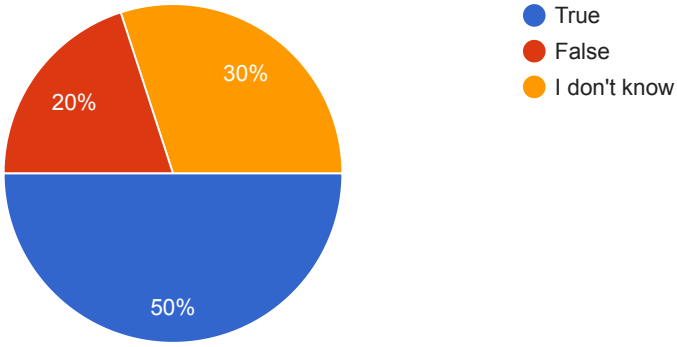

Copy

The EU-Directive on data protection…

10 responses

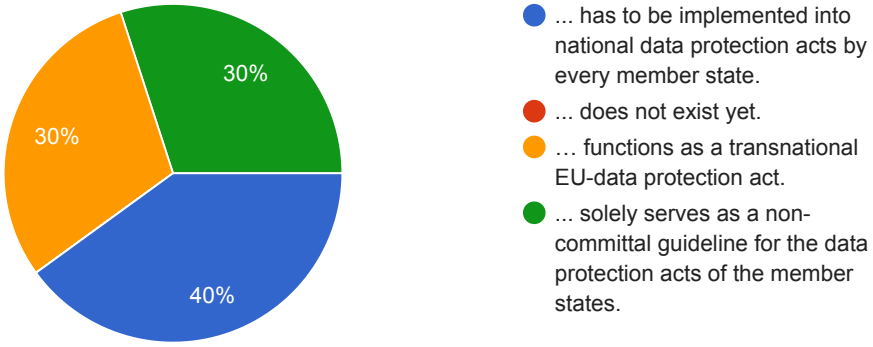

Copy

Tracking one's own internet usage is made more difficult if one deletes browser information (e.g. cookies, cache, browser history) regularly.

10 responses

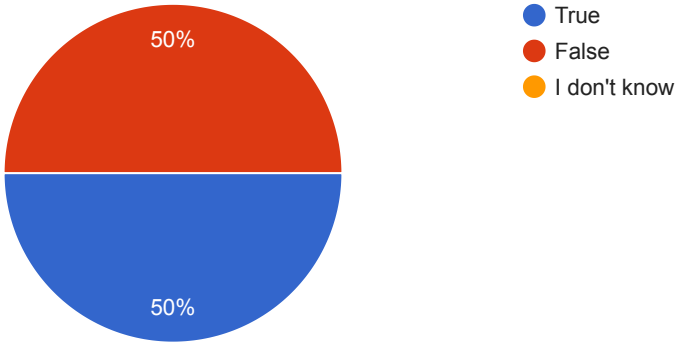

Copy

Surfing in the private browsing mode can prevent the reconstruction of your surfing behavior, because no browser information is stored.

10 responses

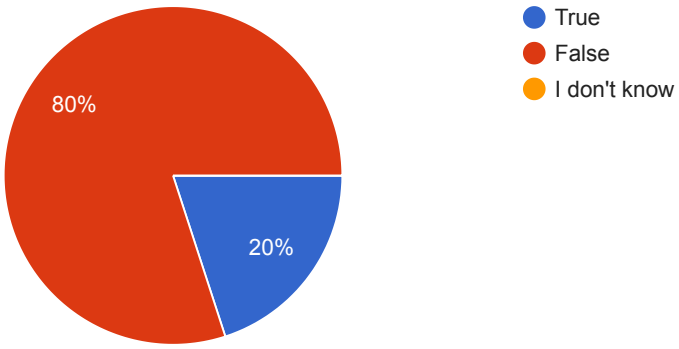

Copy

Using false names or pseudonyms can make it difficult to identify someone on the Internet.

10 responses

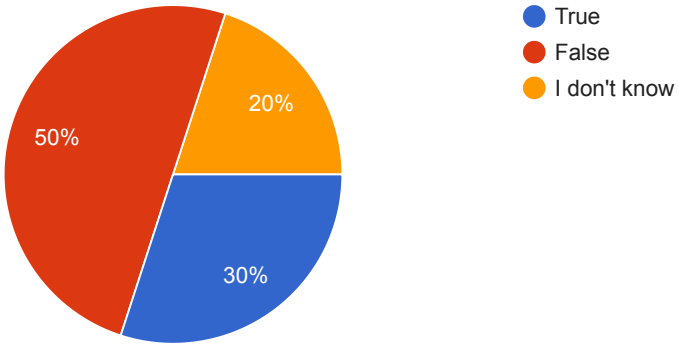

Some more questions about your experiences with cookies

When interacting with privacy settings, some people feel like they are driving the car, while others feel like they are just passengers on a pre-determined route. Where do you see yourself in that analogy, and why?

10 responses

definitely the driver like im ryan gosling

Definitely more like a passenger, as I am aware that much of my data has already been stolen and traded off a million times. It does not make a huge difference for me to reject cookies or play with different privacy settings, because I feel like all my data has been stolen from me a long time ago. "Secure" apps and browsers and whatnot else have also time and time again been exposed for selling our data the same way. I am aware I have little to no control over who does what with my data and that I pay with my information to use pretty much anything on the internet. I still try my best to protect what I can but there is only so much EU regulation can do, plus I feel like data trading isn't regulated the way it's supposed to be. Plus websites and apps can just straight up lie and steal my data anyway.

passenger :p I create secure passwords, protect myself from phising, not sharing my bank details, avoid scams etc. but when it comes to cookies and sharing data that does not cause hacking or could cause economic loss etc. I do not use my time on trying to influense what is shared.

Jeg tror ikk jeg tænker så meget over det, udover at jeg ikke Sys ddt nice at man bliver tracket på alt hvad men laver

i feel like a passenger, as mentioned earlier i feel like cookies have been so normalized to the point where u have to accept that its a part of the web experience. Therefore the only thing to do is reject as much as possible and then clean out often.

A bit of both. I feel like I have a pretty good understanding on what I can do to keep private, but it is made harder and harder on more and more sites.
Privacy is definitly a constant war of trying to one up the other with new methods of getting or protecting data.

i never thought about it. i think i feel more like the driver, but i also dont know a lot about privacy, so im propably most likely a passanger. most things are propably predetermined and done to use my info.

I feel like I could be sitting in the drivers seat of a Tesla that drives itself most of the time but I monitor closely and interfere if I feel the need to.

Hmm in between maybe a little torwards driving the car - one hundred percent because i work and study with cyber security, but before that definetlly more a passenger

More like the passenger, i have no idea what i'm doing most of the time, i just follow along.

If you were to spend five to ten minutes carefully selecting your preferences, for example hand-picking what to allow on a site, how much do you believe that effort would actually change how the site treats you?

10 responses

i would guess it wont change much

Realistically I am not even sure, I understand what the settings are supposed to mean, but I have little to no trust to anything online in regards to stealing or dealing with my data. I still go to the lengths so reject everything but I genuinely don't know if they won't still steal my data anyway.

i would never do that haha. But I do not believe it would make that big of a difference tbh.

Forhåbentlig treater den mig ikke anderledes

I would assume very little. i've seen some website say "oh accept terms with our 1128 partners and cookies", i feel like if i were to go in and reject every single one manually, they would most likely still track at least "necessary cookies" which i assume is up for interpretation.

I think most social media platforms would not want me as a customer anymore.
But usually, if it's a site I use often and a lot, I take the time to at least get myself acquainted with their privacy policy.

i would hope i protected my information more. im not sure i understand the question

I would believe that I see less personalised ads (which I am also very very afraid of.... how is that information about me stored #scary). I would also maybe see some of the text on the website I am trying to visit somewhat blurry. Using 5-10 minutes selecting the preferences would I say is too much time. I almost never visit websites for that long time either way, that is if I am browsing. If I were to read e.g. research papers online I would think about investing more time in it, but then again I think "noooo these companies are in the academic field so they must be kind towards users".... whiiiich is maybe not so smart way to think about it. But as I mentioned earlier in this questionnaire, I am unsure of how far these cookie preferences go. What actually does it mean to "reject". This is a question I am missing the oversight of.

It would not change that much, but it can prevent that they collect extra data that is not needed, and could change my "personalized" experience

I honestly have no idea. I'm not sure i even know half of what the cookies do or how much it would change.

Describe a situation where you felt "forced" to accept cookies even though you were uncomfortable with it. What was the specific pressure or friction that made you abandon your own preferences?

10 responses

some online news sites make you unable to view the article unless u accept their cookies.

A couple times when something major happened and I wanted a recent article, some online newspaper magazines force you to either pay money or accept cookies. I have accepted cookies a few times but I began to just use a paywall remover tool when I need to instead.

In general, when I do not see the deny all alternative I feel forced to click accept all / only neccessary.

Ville gerne ind og tjekke en opskrift på en hjemmeside, måtte ikke komme derind hvis ikk jeg sagde ja til cookies (eller betalte) så jeg lod vær (:

i think last time was on an american news site - where for me to access the textual version of the news story i had to partially accept cookies, and if i wanted to watch the video of the news article and get additional information i had to accept all cookies.

There have never been a time where I have accepted cookies that I didn't want to.
I'm very stubborn when it comes to things I don't want to do, so any time that happens, I just don't use the site.

There have diffinitely been attempts, like having to pay to disable cookies, but I just block the site and move on to another site.

when there is no imediate chance to reject cookies. when i have to click more and deselect multiple types of cookies, then i feel forced to give up and just accept

Often if I am in a group setting, maybe showing my friends something on my phone, I might feel for a little while that I need to be quick with making the cookie banner go away. Then I would be going against my own beliefs and previously (in high school for example) I would always just accept them, but now (maybe as I have gained more knowledge about computers in general) I feel I can take my time in declining and maybe even educate my friends about it at the same time. I have also sometimes when I am, as I also said before, trying to hurry to e.g. buy something online that is selling out fast just accepted the cookies, or in the setting where I hope the website has all functionality available as made from the designers. This could be e.g. on pages with many graphics or something where the interaction on the webpages matter. I don't remember exactly now what this could have been but if I feel it would affect my user experience visiting the page I would accept the cookies.

I cant remember anytime, but maybe sites i dont use that much normally, i would prefer they didnt have info on me

I can rarely reject the cookies, the only options i see is “Manage options” or “Accept all”. It’s very frustrating to not have a choice, and have it basically taken away from you.

Thank you so much for completing this questionnaire.

G Post-study questions (anonymised)

Closing questionnaire

9 responses

Publish analytics

Page 1/4

In what ways did directly interacting with the website elements change how close or far away you felt from your data and potential outcomes of your privacy decisions?

9 responses

It changed in that way that I felt closer to my data, and who I let get to my data. Ive always been a bit paranoid about what cookies I say yes and no to, so being able to use ur ai to dive deeper and get a quick answer was super helpfull.

I felt more in control and i actually felt that my preferences, when it comes to sharing my data, it was a lot easier to maintain and not compromise, when it comes to the way the website presents the cookies.

it definitely made it feel more comprehensible in a sense. By being able to click on an element and then getting ur questions answered was pretty nice - i could definitely see my self using this type of extension more. Tho i will say because i made it reject everything then from time to time i couldn't really check what was happening under the hood as it went too fast lol. But i think its a great extension and very handy to be able to double check how and what is happening with your data.

it was honestly super nice being able to condense it down, making it easier for me to make a decision based on what the AI said. I felt closer to the decisions and data, and I felt like I could make more qualified decisions

In theory I think it would make me feel closer, however, because of how the AI/tool responded, it felt too general and not specific to the selected text/website. I would love if the tool could do some "research" on its own, looking at the context of the website, interpet the text, look into the website policies (that are linked in the cookie banner) etc. so it feels like it actually has knowledge and that I can trust the answers it provide. Unfortunately it did not feel like that, and I did not trust the answers it provided, it just felt like it talked about cookies in general, not specific for the website.

I found out that I meet less cookies than I thought I do. The extension saved me a little time, when the first pop up only contains reject or accept. But when I had to go into settings and choose my cookie settings, the extension didn’t register anything. Maybe it is unnecessary for it to register, but it would help, to make the process easier, by for example showing where to click.

Not really, but I think that is a me situation. I know cookie setup pretty well, so there wasn't really anything new in that regard.

For me the direct interaction didn't do much, but mostly because I felt that it was inconvenient at times. I would sometimes open up the extension by accident and then click something on the website and instead of interacting with the website, the prompt window came up. I felt that while it can be handy for some users who may want to investigate individual parts of the banners or texts, for me this made it more difficult to interact with the cookie banners and sites in general.

It felt like I could understand better what data was being collected, and it made me stop and think critically about my choices.

Page 2/4

Despite its numerous flaws, in what ways do you think this overlay (ticks and crosses) influenced the way you think about enforcing your preferences?

9 responses

it was super helpful when it worked, it did often give a weird result tho. I did like that you could change your preferences in ur little model, but I hated that there was 8 different models you had to customize.

I did become more cautious, as to what website should have access to my data, and more comfortable, because i got honor my personal preferences. I became more aware of the times, where i wasn't able to simple just reject the cookies and i to compromise with my chosen preferences.

I think im a bit more boring in that sense because for me its always just reject all. So i didnt even really get that as the extension would just insta reject all all the time. (which is nice!)

It was a nice addition and it simplified the decision making process if I didn't want to sit down and read (or make the AI) the text.
I do think it would be nice(r) to have the AI make a "decision" behind the choice, aside from just the distribution profile.

I quickly let me see if there is a "reject all" button (because I had made strictly that I only wanted reject all). And it maybe made me a bit more aware of the text on the buttons in general, when I just want to click something quick on autopilot.

It was very easy, because you intuitively know to click on the green tick.

I don't think it influenced my opinion, but it difinitely made me think an extra time before mindlessly clicking something.

It didn't make an influence on my decisions, since I have been heavy on rejecting every single possible thing I can while browsing anyway.

I think it was very helpful to be able to see the ticks and crosses reflect directly my choices toward privacy. In the cases where all buttons were marked with red I felt somewhat lost and "demotivated" in terms of understanding. I wanted to visit the webpage but didn't want to confirm cookie choices that didn't align with my preferences.

Describe a time the tool made a recommendation that you disagreed with or surprised you. How did the interface allow you to argue with the system or assert your own preference? How did it feel to override the AI in that moment?

9 responses

I didn't really use this tool a lot, but I did argue a little with the ai - I did make some weird proposals

That did not happen for me.

I did not have that happen to me, though i did have it happen once when i asked about what data was being collected on a supplement site and then it just started talking about buying things online and what magnesium pills were good so that was pretty funny.

The AI was barely able to understand that I was trying to correct it. I think the idea behind the system is excellent, but the answers it gave me were pretty lackluster for the most part.

I had a strict policy that I only wanted reject-all, however, many websites only have "all" or "only necessary", so then I went with "only necessary", even tough it had a red cross. I did not think much about it, I just wanted to see the website content. Sometimes there were bugs so it made green click on "only necessary" or cross on "reject all" because the cookie banner may have used another wording or a different language, so I used a couple of more seconds myself to find the best alternative, ignoring the cross/tick.

I think when opening the website nakedcph.com i had to choose accept or show details it didn’t understand what to do, so i just clicked to show details and here there were red crosses on all options. I don’t remember if it’s true, please check for yourselves, but I remember it as if it was wrong and I had to correct it.

I don't think there was a time where it's recommendation was wrong compared to my preference settings, but there was a few times where the cookie UI was too complicated for the extension to figure out and just not give a recommendation.

I used it on a website where the tool was telling me that rejecting everything and then clicking "save preferences" was incorrect. It definitely made me question myself, and a little bit confused, even though I made sure this was the only way I could reject cookies and continue while rejecting them.

In the case of all cookie choices having red x marks I felt the need of understanding further exactly why this was. The answers from the AI didn’t always help and were often repetitive. So I felt I couldn’t trust the output of the AI 100%.

As you may have noticed, many of the answers our extension gave you were super vague. Interestingly, we have noticed a tendency in users expressing increased care regarding their privacy, exactly because they doubt the validity of the responses. If you have experienced something similar, can you describe it?

9 responses

If you kept on asking the ai the same question on repeat, it would crash a little and now it would give you the same question that you gave to it- like im n´the one knowing the answer. Not a very clever ai </3

I haven't experienced this actually, because i fully trusted that the extension knew my preferences.

It made me think a bit more when i was browsing just out of pure curiosity, per the last questionaire i raised questions about when u reject cookies if they actually do it - so sometimes i would ask it was is being collected after rejecting, and it was pretty vague about it, though i wouldnt say i doubted the extension, but as it was a bit vague i would then further research said website and cookies.

The AI was hallucinating half the time so I honestly don't think it made me care more. I would tend to move away from what the AI told me and return back to my usual routine of declining cookies if I don't trust the website. Maybe the first time I would've experienced what you are describing, but not the times.

This is what I said on the first question I think. So I did not trust the answers the extension gave, and did not believe it had the knowledge to answer the questions. Maybe if the tool had been like "alarm this website will actually use your data for x, y, z and it can harm you like this and this" maybe it could influence my choices, but it just gave a super vague explaination of cookies so I did not learn anything and it did not influence my preference or approach to cookie banners.

I only prompted once I think, but the extension didn't answer the question, I didn't think to use it again to be honest.

Yeah. I definitely think about the validity if the AI is confident. If the AI just concludes something without asking follow up questions on something that could be ambigous, I will almost instantly disregard the answer and try to rephrase my question.

I mainly experienced that the extension AI didn't always directly answer my question or give a fully truthful response. Like once I asked if a site is selling my data, to which the answer was no, but then, I asked it again and it said yes. It definitely added some confusion to the mix at times.

Right now I can't think of any situation but I utilise AI tools with very general information rather than inputting very personal and private info that could possibly 'corrupt' the answer. Like for example gender discrimination.

Page 4/4

Can you describe a specific browsing session where the tool actually changed the way you navigated a website? What did you do differently?

9 responses

The tool gave mi insight in whether I wanted to support the website, and if I wished to give them my data. Before even going into the website I would use the tool to see what they used my information for and if they had any political agenda behind them

I felt more safe, in general. I felt like i browse more freely. I did also try to ask the tool if it knew more about the website i was on, after i had chosen to reject the cookies. It didn't necessarily change the way i navigated, because regardless i probably would have found a different way to answer the question i had, like i usually would.

overall just faster as it would reject popups fast which is nice, gonna keep that. (idk if u have to devops it so it stays up to date tho)

When I was browsing news with the extension: I spent more than 1 second just clicking the reject button. It was quite interesting to prompt the AI to explain what the data was used for, etc.
Another time I prompted the AI with a large chunk of text about what the particular website was using my data for, and it nicely explained it concisely.

Not specific, but maybe that I use a couple of more seconds looking at the buttons to see if there is a green checkmark on one of the alternatives.

I don't recall doing that.

No not really. Every time the extension had influence on my browsing was mostly to reinforce my already existing choice, but it was sometimes nice to have something help me confirm my choice.

For the first time I actually read the text inside of cookie banners instead of just straight out rejecting it. I used the extension to gain a sense of understanding of the cookie banner texts and the different ways websites handle my data. It helped me learn about how differently some websites cookie policies are.

Not really.

Can you describe a specific moment when the tool felt like it truly belonged to you, rather than just completing manual tasks to contribute to the user study? In what moment did the extension become genuinely useful?

9 responses

At some point I started using the tool as a quick way to understand and search for stuff. Example if I read an article and I didn't understand the word, or there was a happening I didn't know about, I would use your tool to do a quick search and a briefing on the subject for me

Everytime i saw it recommended that i should “reject” a cookie i felt a lot safer and that my data wasn’t being unnecessarily shared. It felt like a friend trying to look out for me, in a way.

i completely forgot about it and just used it when i was unsure about stuff, which is perfect - that means that the extensions is just a part of my normal day2day browsing. It became a well oiled machine in my paranoia of the internet.

I was browsing a website for shopping something and it was nice that it reduced the amount of decision making I had to perform, when I just wanted to browse. It's nice that the tool can be used to actively understand cookies/privacy on websites, and also that it does some of the hard lifting when it comes to the additional decision making that cookies impose.

I became more aware if there was a reject all option or not, when I visted new websites. However, I browse mostly on safari on my phone, so it would probably be more useful there, because I actually notice how little aware of cookie banners I am on my phone versus my laptop (I probably click accept all often, without even knowing it). Probably because I am generally more rushed and impatient when I browse on my phone. I think on my phone seeing if there is a red/green cross/checkmark, would make be autopilot to the green checkmark, rather than just whatever close the banner the quickest.

It made the decision slightly easier because I didn’t have to read the pop up it think about my answer, the fact that it was based on my decision to protect my data was a point when I felt ownership over it.

Unfortunatly not. It was very much just confirming my existing knowlegde, so most of the usefulness for me comes from having me not doubt my choice.

I was close to misclicking or accepting cookies against my will when I was clicking too fast to get rid of the banners and the tool helped me reevaluate and click the right things so my choice aligned with what I actually intended.

When visiting new websites I know nothing about and might not trust the authors. Like non-governmental or non-official websites.

When you used the tool to probe or interact with specific website elements, did you discover something interesting that you previously did not realize? How did that moment change your sense of who holds the power on that page?

9 responses

I did indeed get surprised in what websites do send ur data to other platforms, its very unfortunate but im happy to be informed so I in the future can skip those sites (:

When i wanted the AI tool to help me manage the “mange preferences” and it couldn’t help me, i truly felt annoyed and frustrated that the website “won” and i just felt the urge to press “accept” just to move on with my browsing.

i don't think i had that happen to me - though it was fun to just double check and get more info about said website other than just "BUY BUY BUY" or " LOOK DANGER" to then be like "yo so whats the deal with this website" and its like yo man this is a political right leaning website etc. I had one website were it was saying it was trying to portray UK as a rogue nation and USA as an Ally of EU, so overall pretty cool!

Honestly not really, and it didn't change my sense of who holds the power

Not really, again I may have defaulted more to the reject all button because I could quickly identify it, but if it was not available I still choosed "only necessary" etc. even though it was crossed out, and in that moment I did not really think about what data it would get etc. Also asking the tool what data that it would collect was not useful, because it just gave a vague untrustworthy answer. To say something, I may have became more aware that I let this websites collect my data, but not what data it actually collects or what it does with it. So maybe it made my aware that I have a bit more power, but it still felt to cumbersome to actually do something about it.

I didn’t interact much with the tool, other than to show me what button to click

No, unfortunatly not. I didn't use the tool much to probe websites, but that was mostly because there wasn't many things I was confused about. I tried it a few times, just to see what it did, but it just gave me a response that I already knew the answer to.

I've always had a sense of being forced into a choice but it opened my eyes to how alarmingly little there truly is I can do to protect myself when using the internet, and how regardless of my best efforts, some of me and my activity will always be monitored, analyzed, and used for somebody's advantage.

Maybe the fact that cookies are always there. According to the tool (answers from it) the cookies are always the same website from websites. Never a new cookie. So all websites know everything. Very creepy.

You are done!

Do you have any comments? (Optional)

6 responses

no

the different memes did make the survey more entertaining. Thanx<333:p

Jeg indløser min gratis øl snart. Tak fordi jeg måtte deltage, og håber I kan bruge feedbacken til noget :') Held og lykke med projektet!

very cool extension and great job u guys <333333 even though it oftentimes did not recognize the cookie banner buttons. I think in general there were more times it could not recognize buttons versus the times it did, therefore, I did probably not interact with the extensions as much as I potentially could + I browse a lot on safari on my phone (and this is a chrome extension) so maybe my experience would be different if I had it on my phone as well.

I don't have any comments ☺️

Only that the cookie in the corner was a bit in the way quite a few times, so had to move it around depending on the website. And once I accidently clicked it without noticing it, and it took me an embarrassing long time to figure out why the extension was preventing me from interacting with anything on the website...

H Personal questions (anonymised)

P1

Looking at your data, we have noted that the way you interact with cookie banners is highly dependent on your mood. Super relatable!

In a world, where our extension was more feature-complete and smarter, one of the potential advantages would be that it would be much better at getting your vibe and attitude toward cookies - being much more personal.

In this case, how would this more understanding, human-like partner in privacy change the way you treated and cared about your data?

Answer:

I think for me it's not the personality of the tool/AI that matters, but its trustworthiness. So if the tool were actually reading the websites policies and I knew it had some validation checks etc. I would trust it more (even though LLMs are always prone to hallucinate, you can do a lot of refining to reduce that risk). So yeah personality does not matter, but it's legitimacy and trustworthiness does. So I had to be convinced that it actually based it answers on actual sources/data, before I would consider to take anything it said seriously.

P2

Looking at your data, we have recognized a pattern in how one of your goals in preserving your privacy is also putting a stop to large data-brokers making money off of your personal information. If they make money, then you should get a cut at least!

In a world, where our extension was more feature-complete and smarter, one of the potential advantages would be that it could provide you with a verifiable trace of your data, and tell you where your data actually ends up, and who is profiting off of it.

In this case, would you act to stop it, and if it proved feasible, how would it affect the way you think about your ownership of personal information in the online space?

Answer:

I would immediately stop it as I never ethically/willingly consented to selling my data. I would feel much safer, much more in control, and it would give me a sense of ownership that I currently think is inaccessible.

P3

Looking at your data, we have noticed that you might agree with the statement, that "*most online tracking is probably predetermined, so one's choices would not make a big difference*".

In a world, where our extension was more feature-complete and smarter, one of the potential advantages could be that it could give you highly accurate answers to whether your choice would actually make a difference - using proper citations to the website's inner workings (code), legal documents, and other verifiable sources.

In this case, would you feel like keeping your data to yourself would be less of a lost cause?

Answer:

Yes, since I would have more of an understanding of the inner workings, I would feel more in control and, yes, like it would be less of a lost cause.

P4

Looking at your data, you were one of the participants that asked the most questions. Curious one!

In a world, where our extension was more feature-complete and smarter, one of the potential advantages would be that it would cite sources and search the web for accurate and up-to-date answers.

In this case, how would the extension affect your confidence in the choices you make?

Answer:

It would make me much more confident and make me feel informed, without having to do the research myself. I can imagine that I'd feel better about the AI guiding my hand in picking which privacy options I want, as it would understand up-to-date literature as well as my preferences. If it could create a tailor-made cookie preference profile for me, it would no doubt make me much more confident.

P5

Looking at your data, you really want to get through cookie popups as fast as possible. Speedrunner type!

In a world, where our extension was more feature-complete and smarter, one of the potential advantages would be that it could make a choice on your behalf and let you know what it meant

for you - for example automatic reject-all, and a text saying that this site cannot track you in any way.

In this case, would you feel more in control about your data?

Answer:

Oh yeah, it definitely would. If i didn't have to click on anything and have the extension do it for me, i would remove a lot of the cognitive load when i'm browsing. I would feel more in control, because i would still get notifications from the extension, telling me what it chose for me and i would therefore still feel informed about what is happening <33

P6

Looking at your data, we have noted that you are concerned with cross-site tracking, especially when it comes to other sites seeing your social media activities. Super valid concern!

In a world, where our extension was more feature-complete and smarter, one of the potential advantages of this product could be that it would be able to let you know exactly how each website would be able to track your activities on other services.

In this case, how would that change your sense of control and/or feelings towards the way you manage online tracking?

Answer:

I feel like because social media is such closely connected to one's personal life that other random websites you browse should not know much about exactly what types of reels you watch the most. I would like to visit newspaper websites and see ads meant toward the general public. If I could see exactly how each website manages my data I believe I would feel more in control of my personal information. I could visit pages knowing what they know about me to make well informed choices/views/opinions about what content/ads they are showing me. And now I am referring maybe mostly about data I see on social media. E.g. how do they know I like this specific makeup brand.

P7

Looking at your data, we have noted that you were curious about the political and financial connections, potential controversies (e.g. data leaks), and other cultural ties of websites even before you visited them. Very woke!

In a world, where our extension was more feature-complete and smarter, one of the potential advantages of this extension could be that it would actually give you a good overview of all of these things.

In this case, how would this influence your decisions? Would it have an effect on your sense of control and care for your personal information?

Answer:

I would definitely use your tool if you made it more like this description. And I think most people would like to know what they are supporting in terms of if the websites has a political agenda. For sure a lot of stuff is hidden and it would be nice to know where not to go <33

P8

Looking at your data, we recognize that you have a tendency to avoid websites where you cannot enforce your preferences. But finding out about this can take a looong time!

In a world, where our extension was more feature-complete and smarter, one of the potential advantages would be that the extension could automatically apply a prompt, and tell you from the get-go whether a site is compatible with you.

In this case, how would this increasingly automated solution influence your daily browsing processes, or your sense of control over your time spent on enforcing your privacy?

Answer:

Yeah, I think it would. I don't often come across websites where I can't enforce my preference, but when I do, it's super frustrating.

I would greatly value if the extension was able to tell me preemptively if a website had paid cookie removal or the likes.

P9

Looking at your feedback, we have noted that you use many of the same websites during your daily browsing sessions, and that you tend to retroactively clean up your traces (trackers, cookies, etc.). Smart move!

In a world, where our extension was more feature-complete and smarter, one of the potential advantages could be being able to learn about where your data actually washes up at the end of the endless data-broker operations.

In this case, would you feel empowered by learning exactly where to step in to stop the leakage of your data instead of doing routine tasks like tracker cleanups? Would this give you more ownership over your data?

Answer:

I'm generally super cautious about my data as a whole, and I've always been interested in the idea of having actual control over your own data. The business of data brokering and advertisements is inherently shady and hidden away from the public, because companies make way too much money from it, and the public (in my personal belief) would be completely outraged if they actually learned what happens to their data, where it goes, and who really knows everything about them.

If we look at examples like in DK, where politicians wanted to add more cameras to bigger cities to help prevent and solve crime, there was huge public outrage about it. That's the government doing it openly, but you're telling me assorted evil corps are completely allowed to know everything about me, where I am, how healthy I am, my interests, hobbies, etc. just because I committed the cardinal sin of using the internet. (Slight rant.)

It would 100% make me feel much better if I could see what happens to my data and where it actually ends up. It would also make it easier to potentially raise political questions and push for change, because right now it's pretty difficult to do that when ain't nobody knows where your data goes or who you're even supposed to call out.

I also think it would raise awareness for the average individual / 50+ year olds that don't really know much about the internet, giving them a better understanding of what actually happens and making them more cautious in general.

And as an added bonus, I'd absolutely love the ability to request all the data assorted evil corps have on me and force them to print out 1000+ pages and deliver them to my door just to fuck with them a little.

I Extension setup guide

1. Udpak din fil

Først skal du klargøre filerne, så Google Chrome kan læse dem.

- Når du modtager din **extension.zip** , skal du **extract** (udpakke) den.
  Mac: **Double click extension.zip**
  PC: **Right click → Extract All**
- Efter udpakningen vil du have en almindelig mappe (f.eks. navngivet "extension")

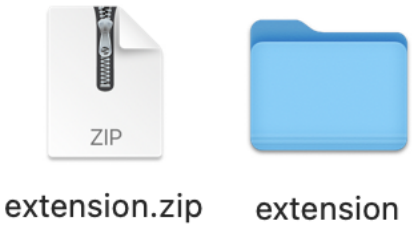

2. Åbn udvidelses-siden i Google Chrome

- Åbn Google Chrome, copy-paste det følgende i søgefeltet, og tryk på enter:
  `chrome://extensions`

- Alternativt kan du finde siden via Chrome-menuen under
  Extensions → Manage Extensions.

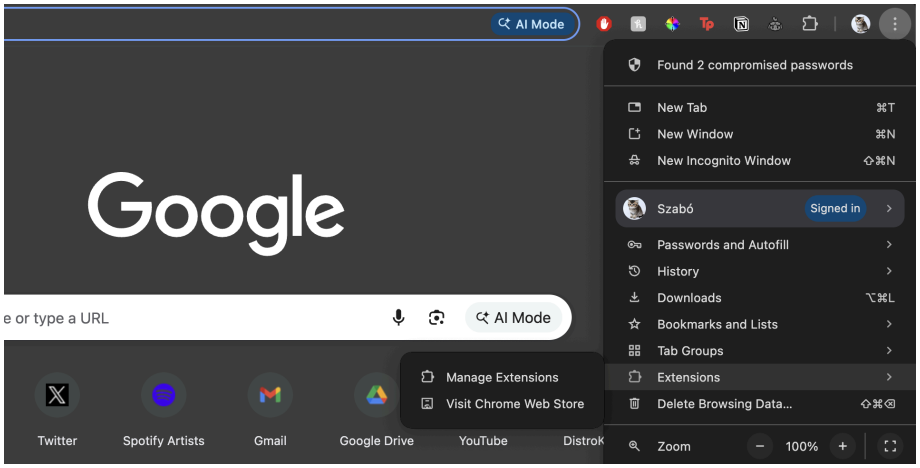

3. Aktivér Developer mode

For at kunne uploade dine egne filer skal Chrome være i udviklertilstand.

- Find kontakten **Developer mode** øverst i højre hjørne.
- Slå **Developer mode til**, så siden ændrer udseende og viser flere valgmuligheder.

4. Upload mappen

- Klik på knappen **Load unpacked**, som nu er kommet frem til venstre.
- Et vindue åbner nu på din computer. Find og **vælg den mappe**, du pakkede ud i trin 1. Den hedder “extension”.
- Klik på **Select**.

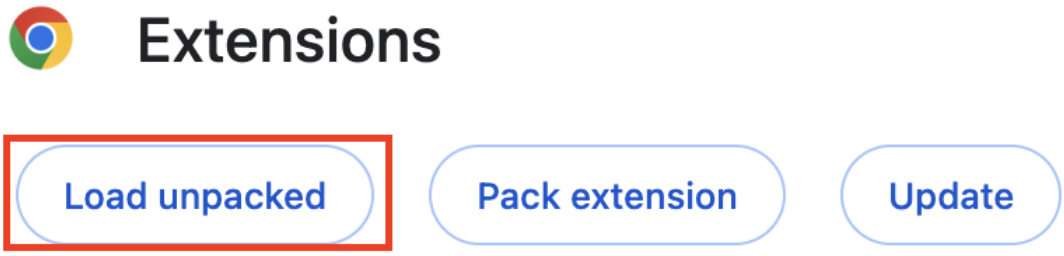

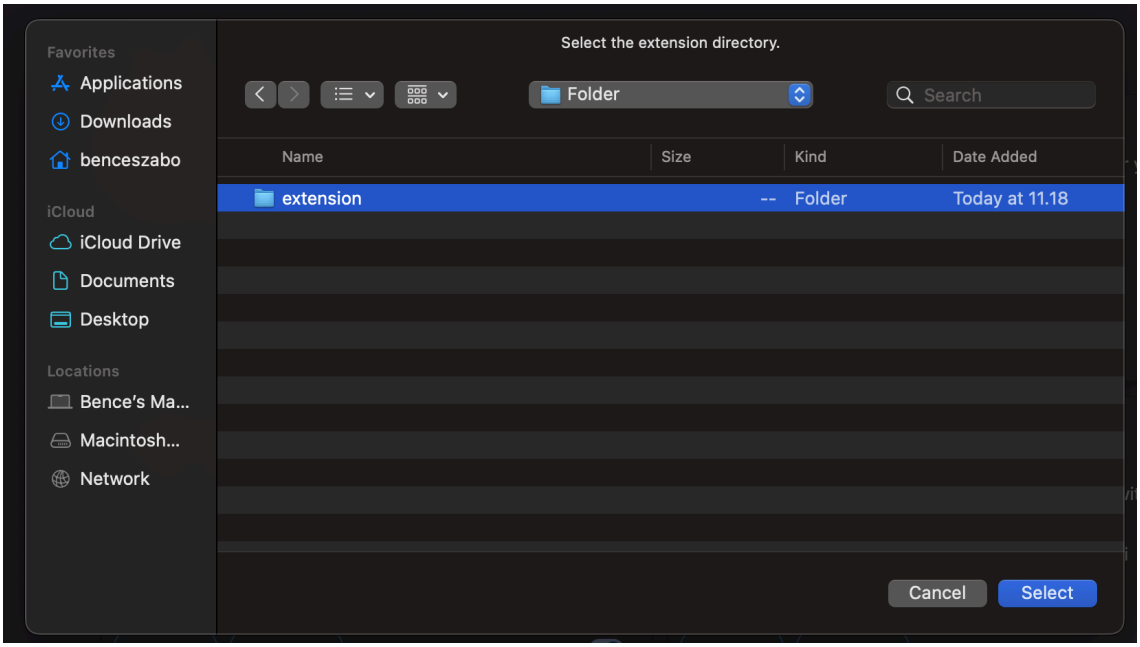

5. Bekræft installationen

- Hvis alt er gjort korrekt, vil udvidelsen nu fremgå på listen over "All extensions".
- Nu burde du kunne se udvidelsen med navnet **"Cookie Slayer"**.

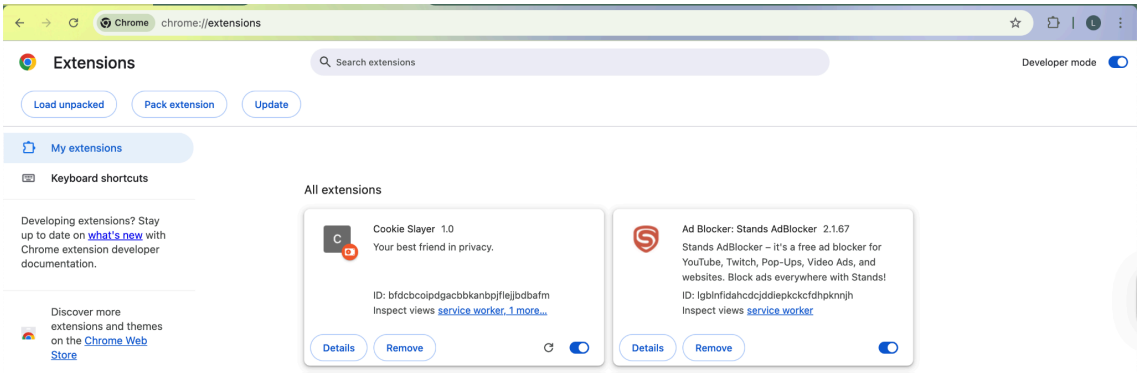

6. Eventuelle spørgsmål

Hvis du er usikker på noget som helst eller har andre spørgsmål angående vores brugerstudie, så kan du altid komme i kontakt med os, enten via privat besked eller email.

bszabo21@student.aau.dk

lfst21@student.aau.dk

sselma24@student.aau.dk

J Results from the user study

J.1 Analysis of pre-study results based on privacy literacy level

High-level participants have an understanding attitude towards cookie banner as long as there is a reasonable way to navigate them. However, when enforcing their preferences is uncertain, the first non-neutral feelings they mention of are irritation, annoyance and hurriedness. These users tend to avoid sites that force them into potential compromises. Otherwise, they give the popups a chance and try to understand them. High-level participants' main concern is what the site would gain from their data - if it is used for internal purposes, then they are mostly okay with them, but when it comes to advertisement and making money from their data, they tend to steer away. Site-specifically, there is a clear tendency to decline all cookies on shopping pages, more careful choices regarding finance, and a general tendency to withhold sensitive personal data. On the other hand, these users do not go far to understand what each privacy option (e.g. Accept, Decline, etc.) encompasses on a site-by-site basis, as they believe to behold a solid understanding already.

When questioned about being "drivers" or "passengers" in an online privacy analogy, these users tend to take a fence-sitting stance. In general, they feel somewhat in control, but limited by the effort it takes to find to navigate to the nitty-gritty of tracking. Regarding their feeling about how they are treated regarding the time spent on cookie popups, high-level users feel punished when spending more time to learn about the potential outcomes of their decisions. Furthermore, they are unsure whether their choices really make a big difference, as the only tangible evidence is the lack of personalised ads. In general, high-level users avoid sites where they feel forced to accept cookies. However, one of the participants mentions that they feel pressured to get rid of popups as fast as possible in group settings or in a time-pressure situation, for instance buying fast-selling products like tickets on a website like Ticketmaster, they feel pressured into "just getting on with it".

Medium-level participants have a tendency to feel frustration, annoyance and wasteful regarding their use of time when confronted with cookie banners. Everyone in this group have reported that they feel like their choice is mostly a compromise due to their lack of willingness to spend more time making meaningful choices. At times, some report a feeling of laziness. Medium-level participants have a nuanced take on how their care for data shifts in different browsing contexts. Some are selective with their acceptance, making exceptions for shopping sites they like, or on pages where they believe that their data is used for valid purposes. However, they seem to dislike tracking by social media companies. Although these users are not as willing to go deep into cookie-sub-menus to find out what their choices truly represent as high-level participants, some of them do tend to go to a certain extent to understand the outcomes of their choices.

However, in general, most users in this group do not feel in control, as they are aware of unethical data collection and monetisation, and have been betrayed by privacy tools and secure browser alternatives beforehand. Some attempt to retroactively "clean up" their traces, but believe that the real driver should be the EU regulations. Although these users show some hope regarding how their choices would result in receiving a treatment they would expect, they mention that in reality they do not believe that spending more time making intricate choices would not change much. At times, they go to great lengths, but are cynical about the true effects of it. When feeling forced by a website to accept cookies, medium-level users have a mixed reaction. They tend to either steer away from such sites, find potential workarounds like paywall removers, or at times give up on their preferences and accept all cookies in resignation.

Low-level participants report feelings of impatience and irritation when confronted with cookie popups. They tend to look for the quickest way out without thinking which option they click on. Their reasoning stems from a place of resignation, describing that they believe that if they have already accepted tracking on a bunch of sites beforehand, it does not matter to press accept on one more site. When it comes to personal preferences, low-level users totally compromise. They would like to reject everything, but most likely click accept all as it is the first fast-track option that is immediately available.

When making decisions, these users' concern is not based on the category of the site, but perception of what a site has access to. A site having access to their credit card would be a bigger concern to them than most other personal information - according to their answers, it could result in material loss. Sometimes, these users they are okay with social media tracking them to give better ads and more personalised content. Showing clear resignation, low-level users have self-reportedly accepted not knowing the outcomes of their decision, and have given up on reading about any of the options. They feel like passengers and are unclear about what they are doing. These users also show a tendency to feel forced more promptly than other groups - as soon as they cannot see an immediate "manage-options"-button, they feel forced to accept all cookies.

J.2 Logged data analysis results

K The project's Burndown Chart

Name	Privacy	Ent.	Shop.	Prod.	Fin.	Health	Total
P1	HIGH		R, R, R, N	N		R	6
P2	HIGH	R, C, R	R, R, R, R, R, C	R			10
P3	MEDIUM		N				1
P4	HIGH	R, N, R, R		A, C, A			7
P5	HIGH		R, R, R		N		4
P6	HIGH	A, C					2
P7	MEDIUM	R, R, A					3
P8	LOW		N				1
Total		**12**	**15**	**5**	**1**	**1**	**34**

Table 4. Breakdown of Decisions by Category (Accept, Reject, Customise, Necessary only

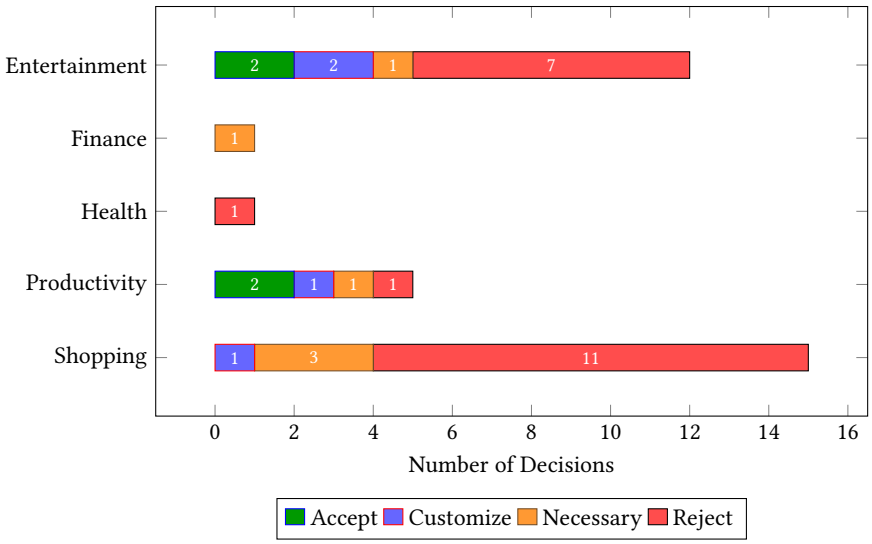

Fig. 13. Number of decisions grouped by type and category

User Action	Is Match?	System Recommended Choices (≥ 0.5)	Created At
ACCEPT	True	ACCEPT	2026-05-04 17:34
	True	ACCEPT, REJECT, CUSTOMIZE, NECESSARY	2026-05-04 17:16
	False	CUSTOMIZE	2026-05-04 11:36
	False	REJECT	22026-05-03 20:42
CUSTOMIZE	False	REJECT	2026-05-11 09:34
	False	REJECT	2026-05-08 08:47
	False	ACCEPT	2026-05-07 20:47
	False	REJECT, NECESSARY	2026-05-05 12:08
NECESSARY	False	REJECT	2026-05-07 20:38
	False	REJECT	2026-05-07 17:03
	True	NECESSARY, REJECT	2026-05-07 08:18
	True	NECESSARY, CUSTOMIZE	2026-05-05 16:02
	False	REJECT	2026-05-05 09:46
	False	REJECT	2026-05-05 09:46
REJECT	True	REJECT	2026-05-11 09:34
	True	REJECT	2026-05-11 09:33
	True	REJECT	2026-05-11 09:33
	True	REJECT, CUSTOMIZE	2026-05-08 10:15
	True	REJECT	2026-05-08 07:50
	True	REJECT	2026-05-08 07:50
	True	REJECT	2026-05-07 20:43
	True	REJECT	2026-05-07 17:22
	True	REJECT	2026-05-06 20:44
	True	REJECT	2026-05-06 10:41
	True	REJECT	2026-05-05 19:56
	True	REJECT	2026-05-05 17:22
	True	REJECT, NECESSARY	2026-05-05 10:32
	True	REJECT	2026-05-05 10:31
	True	REJECT	2026-05-05 07:33
	True	REJECT	2026-05-04 17:25
	True	REJECT	2026-05-04 17:13
	True	REJECT	2026-05-04 11:04
	False	CUSTOMIZE, NECESSARY	2026-05-03 20:40

Table 6. Comparison of User Actions and System Recommendations (Grouped)

Total Matches	Total Mismatches
21	12

Table 7. Summary of Recommendation Alignment Counts

Match Status	Category Alignment	Count
True	Recommended Category matches Website Category	33
False	Recommended Category differs from Website Category	1
Total Joined		**34**

Table 8. Recommendation Accuracy Summary for Website Contexts

	A	B	C	D	E	F	G	H	I	J	K	L
3	Write intro	Week 2	Done	5%			**Week 1**	06-03-2026	97%	95%	95%	
4	Write related work	Week 2	Done	5%			**Week 2**	13-03-2026	80%	85%	85%	
5	Design interface	Week 3	Done	10%			**Week 3**	20-03-2026	70%	75%	75%	
6	Design interface	Week 4	Done	5%			**Week 4**	27-03-2026	65%	65%	65%	
7	Write design	Week 4	Done	5%			**Week 5**	03-04-2026	60%	50%	50%	
8	Write methods	Week 5	Done	5%			**Week 6**	10-04-2026	55%	45%	45%	
9	Design intrastructure	Week 5	Done	5%			**Week 7**	17-04-2026	50%	40%	40%	
10	Design intrastructure	Week 5	Done	5%			**Week 8**	24-04-2026	40%	35%	35%	
11	Build extension	Week 6	Done	5%			**Week 9**	01-05-2026	30%	30%	30%	
12	Build extension	Week 7	Done	5%			**Week 10**	08-05-2026	20%	15%	15%	
13	Build extension	Week 8	Done	5%			**Week 11**	15-05-2026	10%	5%	5%	
14	Build extension	Week 9	Done	5%			**Week 12**	22-05-2026	0%	0%	0%	
15	User study	Week 10	Done	15%								
16	Evaluate results	Week 11	Done	10%								
17	Finish paper	Week 12	Done	5%								
18												
19												
20												
21												
22												
23												
24												
25												
26												
27												
28												
29												
30												
31												

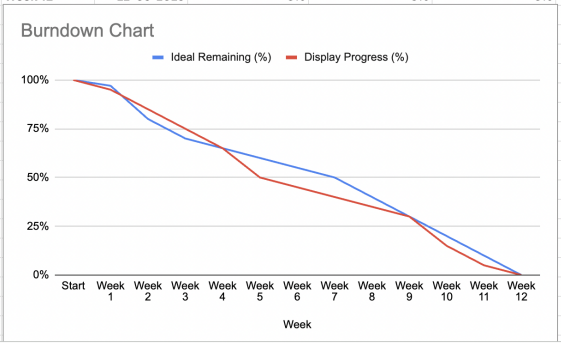

Fig. 14. The project's Burndown Chart