When Designers Meet GenAI: Understanding the Role of Prompt-to-Design Generators in Privacy Dark Patterns

Jingzhou Ye, Zhaojie Hu, Yao Li, Xueqiang Wang
University of Central Florida, USA
{*jingzhou.ye, zhaojie.hu, yao.li, xueqiang.wang*}*@ucf.edu*

***Abstract*—User interfaces (UIs) are the primary gateway for users to access digital systems, ranging from healthcare and finance applications to AI-powered platforms. Yet many of these interfaces embed dark patterns – manipulative design strategies that deceive, coerce, or otherwise influence users into making unintended decisions that benefit services while threatening users' privacy, safety, and autonomy. Recently, the UI design process has evolved rapidly with the widespread adoption of generative AI. In particular, designers increasingly use prompt-to-design generators (PRODEGENs) to create UI prototypes directly from natural language prompts. This emerging paradigm of designer–PRODEGEN collaboration may reshape the landscape of dark patterns, yet its effects remain largely unexplored. In this study, we investigate how such collaborations influence dark patterns that harm users' privacy rights. Specifically, We conducted designer-centric studies with 17 professional UI designers to examine their awareness, attitudes, and practices, and complemented this with a technical evaluation of real-world prompts across multiple PRODEGENs. Our analysis identifies both designer practices related to dark pattern occurrence and mitigation, as well as system-level factors in PRODEGENs that facilitate dark patterns, providing insights for mitigating privacy risks from the perspectives of designers, PRODEGEN developers, and end users.**

1. Introduction

In recent years, deceptive user interface (UI) designs that manipulate, deceive, or coerce users into making choices that benefit service providers at the expense of users' interests, also known as dark patterns (DPs) [1], [2], have become increasingly pervasive, posing a significant threat to the privacy, safety and autonomy of online users [3], [4], [5]. Examples include permission request interfaces that mislead users into granting data access for tracking [6], hard-to-find privacy settings [7], consent dialogs that make rejection difficult or impossible [8], and dialogs with harmful defaults [4] or preselected options [2], [9]. Prior research has demonstrated that UI designers play a pivotal role as key stakeholders in the creation and, at times, the perpetuation of DPs [5], [10], [11], [12]. However, little research has examined UI designers' perceptions and practices related to DPs. Understanding their DP awareness, attitudes, strategies, practices, and constraints is crucial for preventing the continued spread and reinforcement of DPs in real-world applications.

Yet, obtaining such an understanding around UI designers is becoming complicated due to the continuous evolution of today's UI design processes. In particular, with the rapid advancement of generative AI (GenAI), many designers have begun integrating GenAI tools into their everyday design workflows [13]. A recent survey indicated that 89% of designers use some form of GenAI tools in routine UI design tasks [14]. One of the most commonly used types of tools is the *text-to-design generators* (PRODEGENs), such as Figma Make [15], Uizard [16], and Vercel's V0 [17]. With these tools, UI designers can generate UI design prototypes directly from natural language descriptions of UI design requirements and usage scenarios. For instance, a designer can receive both the code and visual rendering of a UI, such as a consent dialog, by simply providing a descriptive prompt (e.g., *"create a consent dialog with options to accept or reject data collection"*) to Figma Make.

The adoption of PRODEGENs has effectively enabled a new UI design paradigm, where designers collaborate with PRODEGENs, instead of following the traditional UI research, ideation, and prototyping process. Such a design paradigm introduces previously unobserved factors that could shape the development of DPs, such as DPs generated by PRODEGENs (versus those created from scratch by designers), the attitudes of UI designers toward them, and the various ways in which they might handle and integrate the generated UIs. However, while recent research has shown that GenAI may generate DPs [18], *there is still limited understanding of how designer–PRODEGEN collaboration might reshape the landscape of DPs in real-world UI design*. In this study, we address this knowledge gap by conducting designer-centric studies on their use of PRODEGENs in privacy-related design tasks, and technical evaluations of real-world PRODEGENs. Specifically, we aim to answer the following research questions:

RQ1: *What are UI designers' knowledge of DPs and their capability to identify DPs in the UIs generated by* PRODEGEN*s?*

RQ2: *What are UI designers' attitudes toward DPs in the UIs generated by* PRODEGEN*s?*

RQ3: *What strategies do UI designers adopt to mitigate DPs in the use of* PRODEGEN*s?*

RQ4: *What UI designer practices are associated with the presence or mitigation of DPs in AI-generated UIs?*

RQ5: *How do different real-world* PRODEGEN*s produce (or are able to avoid) DPs in real-world UI design tasks?*

To answer the first three questions (RQ1-RQ3), we recruited 17 professional UI designers and engaged them in privacy-related UI design tasks using a representative, leading PRODEGEN, i.e., Figma Make [15]. Following this, we conducted semi-structured interviews to explore their knowledge, awareness, attitudes, and mitigation strategies related to DPs in UIs generated by PRODEGENs. The prompts that participants used in the design tasks form a dataset representing real-world UI designers' practices. To answer the latter two questions (RQ4 and RQ5), we perform a technical evaluation by running the prompts in the dataset on several of the most widely used PRODEGENs, with attribution analysis to identify designers' potential practices related to DP occurrence (pro-DP) or DP mitigation (anti-DP) when creating prompts for PRODEGENs (RQ4), and empirical comparative analysis to characterize the capabilities of different PRODEGENs (RQ5). Together, the designer-centric studies and technical evaluations help address the lack of understanding regarding designer–PRODEGEN collaborations, offering further implications and suggestions for reducing the harms of such patterns in the new era of GenAI (e.g., from the perspectives of PRODEGEN developers, interdisciplinary DP researchers, and end users).

More specifically, the study presents the following noteworthy insights:

RQ1: *UI designers have concerningly limited knowledge of and ability to identify DPs in* PRODEGEN*-generated UIs.* For instance, fewer than a quarter of the recruited UI designers had previously heard of DPs; even after being introduced to the concept, most of them fail to identify such patterns in the generated UIs.

RQ2: *In contrast to current research and user perceptions of DPs, many UI designers viewed DPs in* PRODEGEN*-generated UIs as beneficial.* When encountering UIs containing DPs, many UI designers perceived them as beneficial for both businesses (e.g., helping collect more user data, promote products, and retain users) and end users (e.g., avoiding UI interaction errors, receiving more information with better clarity), and aligned with prevailing industry norms. Consequently, they reported that they would adopt such generated UIs in their real-world UI designs.

RQ3: *UI designers are often unsure about the causes and mitigation strategies of DPs.* For instance, UI designers reported uncertainty about whether their prompts or the inherent behavior of PRODEGENs contributed to the emergence of DPs. They either experiment with different prompts in an attempt to reduce such patterns, or rely on subsequent usability testing to identify and address them.

RQ4: *UI designer's pro-DP practices in creating prompts may lead to the generation of DPs.* Our attribution analysis shows that some designers adopt anti-DP practices, such as adding usability, ethical, or privacy requirements to the prompts, which can reduce the generation of DPs. However, we also noticed that some prompts used by UI designers are highly likely to induce the generation of specific types of DPs.

RQ5: *DPs are commonly found in UIs generated by different* PRODEGEN*s, and some* PRODEGEN*s have design and implementation failures that allow DPs to persist.* We found that all three PRODEGENs investigated in the study generated UIs with DPs, and none of them have built-in mitigations. Moreover, even when prompts are created using anti-DP practices to avoid DPs, some PRODEGENs still generate them due to various factors, such as flaws in reasoning and UI code analysis, and failures to properly compose UI components (or templates) from different sources.

Suggestions. Addressing these questions enables us to offer several suggestions to different stakeholders, including PRODEGEN developers, regulatory authorities, and educators. First, the lack of built-in mitigation for DPs in PRODEGENs leaves UI designers entirely responsible for avoiding DPs in UI design. Our study shows that it is technically feasible for PRODEGENs to follow anti-DP practices in prompts and generate fewer DPs. Hence, to reduce the burden on the large population of designers, we encourage the development of DP mitigations built into PRODEGENs. Second, although regulatory resources for DPs are limited, existing regulations largely rely on abstract judgments of user perceptions to address the prevalence of DPs, which undermines enforcement effectiveness and scalability. To address this, we suggest that new regulatory frameworks define clear harm–benefit tradeoffs across multiple stakeholders (e.g., users and designers) and use these tradeoffs to characterize DPs, thereby enabling more actionable and scalable regulatory enforcement in practice. Third, our study finds that the majority of recruited UI designers are still unaware of the concept of DPs and have limited ability to manage them in design tasks. To tackle this, we believe it is important to develop systematic education and training that provides UI designers with both knowledge about DPs and the skills to create UIs free of DPs using PRODEGENs, especially before PRODEGENs offer built-in mitigations.

The study involving UI designers was reviewed and approved by the IRB of our institution. We have also made the prompts gathered during the study, along with other materials, available on GitHub [19].

2. Background

Dark patterns. In 2010 [1], Harry Brignull introduced the term "*dark patterns*" (DPs) to describe deceptive UI designs that manipulate users into actions they did not intend, often to the advantage of online businesses. An example of a DP is a preselected option that opts users in by default, or a misleading button that is either visually deceptive or not noticeable to users. Such DPs have been reported to appear frequently in privacy-sensitive UIs, such as consent management [8], [20], [21], [22], [23], [24], [25], permission requests [6], privacy settings [7], [26], and account deletion [27]. In 2018, Gray et al. further classified DPs into five broad categories in their taxonomy [2]. Specifically, (1) *Nagging*: Repeatedly prompting users to take action, often in a persistent or intrusive manner. (2) *Obstruction*: Deliberately making a process more complex to discourage or block certain actions. (3) *Interface Interference*: Modifying the user interface to prioritize certain actions over others. (4) *Sneaking*: Hiding, disguising, or delaying the display of important information from users. and (5) *Forced Action*: Requiring users to complete a specific task to access or maintain certain functionality. We will use this taxonomy to guide DPs reporting in the PRODEGEN-generated UIs.

UI design workflow in the pre-GenAI era. UI design workflow [28], [29] is an iterative process guiding designers from understanding user needs to delivering and refining user-friendly interfaces. It typically starts with research and discovery to gather insights on users, business objectives, and context. Next, designers define user personas, information architecture, and user flows to set the design direction. Based on this, they create wireframes, visual designs, and prototypes. Usability testing is then conducted to identify issues and refine the design before handing it over to developers for implementation. Throughout this workflow, designers use various UI/UX tools, such as Miro [30] for mapping user flows, and Figma [31] for designing and prototyping UIs. However, they still rely on their expertise and creative judgment to translate abstract requirements into tangible user interfaces. DPs can be introduced at any stage of the process, from business objectives that overshadow user-centric design to the prototype stage that translate design ideas to UIs.

UI design with PRODEGENs. With PRODEGENs, UI designers only need to provide prompts specifying the requirements for UI design and receive interactive prototypes, typically in the form of front-end code and rendered UIs. Although the implementations of PRODEGENs may vary, they generally share several core components, as illustrated in Figure 1. PRODEGENs often feature a *designer conversational GUI* where designers can enter prompts that describe the expected UI. They then uses a *design requirement analysis module* to analyze the prompt and extract key semantic elements, such as the UI's purpose, required features, and stylistic preferences. Subsequently, PRODEGENs typically uses LLMs to generate the front-end code that produces a UI prototype. This process often includes selecting and composing external UI resources (either from the PRODEGEN's built-in knowledge or retrieved through online searches), such as component libraries, pre-built templates, and image repositories. Additionally, PRODEGENs may adapt these resources on demand based on the prompt, e.g., modifying visual styles to align with the extracted stylistic preferences. Once the front-end code is generated, PRODEGENs leverage a *UI rendering module* to display the design for the designer's review. Based on this review, UI designers can perform iterative refinement by entering new prompts in the designer conversational GUI. Our study shows that both the prompts and the internal implementation of PRODEGENs (e.g., the reasoning of LLMs and their adaptation of UI code) can impact the generation of UIs with DPs.

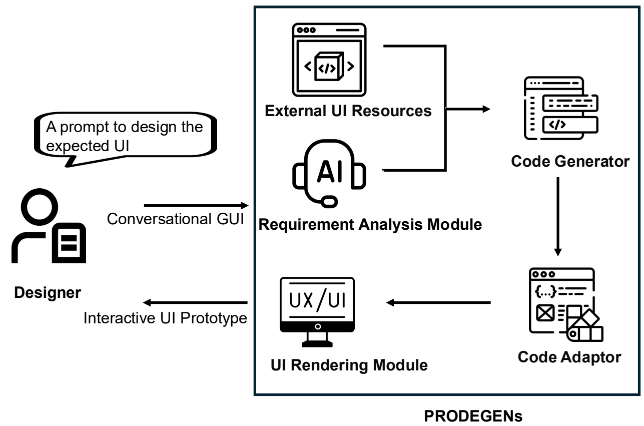

Figure 1: An overview of a PRODEGEN and its components

3. Designer-Centric Study

We conducted a designer-centric study to address three research questions (RQ1–RQ3) concerning UI designers' awareness of DPs, and their attitudes toward and approaches to mitigating these patterns. As part of the study, UI designers completed design tasks based on UIs that are related to privacy and have previously been identified as affected by DPs, alongside pre- and post-task interviews.

3.1. Recruitment of Participants

We recruited participants in the United States, targeting UI designers aged 18 years or older who had prior experience using PRODEGENs for UI design. This criterion ensured that participants' input reflected authentic, real-world use of AI design tools, rather than hypothetical or experimental engagement. To reach a broad audience, we distributed recruitment information through several online communities for UI designers, such as Facebook group "UI & UX Designers" [32] and Reddit community "UI_Design" [33]. We also leveraged participant referrals to reach additional designers beyond these online communities. We observed thematic saturation (when no new codes emerged during the interview data analysis) and stopped recruiting more designers after 17 participants (P1-P17). In total, 17 participants were recruited, 16 from the social media platforms and one through referral.

Potential participants first completed a screening survey administered via Qualtrics [34]. Before presenting any questions, the survey provides an informed consent form that outlines the purpose and procedure of the study. Since DPs are the central construct being examined, participants need a definition to understand what they are agreeing to. However, common DP definitions in the literature may carry evaluative framing that can prime participants to respond in socially desirable ways rather than reflecting their actual reasoning and practices. To avoid this, we use a neutral definition adapted from prior literature [1], with morally charged terms (e.g., "*deceptive*," "*unethical*," or "*tricks*") intentionally removed to avoid evaluative framing. Specifically, the form states: "*The study examines UI designers' perceptions and practices in using AI-powered UI design tools (e.g., Figma Make) to design privacy-related UIs. We aim to understand designers' reasoning and strategies regarding 'dark patterns' (DPs) in AI-powered UI design. Dark patterns refer to interface designs that may influence users' decisions in ways they did not intend.*" The form further informs participants that their responses and use of PRODEGENs during the study will be analyzed for research purposes, and presents explicit options: "*Yes, I agree to participate*" and "*No, I do not agree to participate*."

For those who agree to participate, the survey presents a set of eligibility questions, such as age and prior experience using PRODEGENs for UI design. Then, the survey asked about demographic information such as gender, education level, and years of experience with UI design to ensure a diverse sample. Finally, eligible participants were asked to provide their email for scheduling the study session. Each session was around 60 minutes and conducted via Zoom. Each participant received a $40 Amazon gift card as compensation for completing the study. The study was approved by the IRB of our institution and adhered to the procedures established by the IRB. The demographic information of participants is presented in Table 4 in Appendix.

3.2. Gathering of UI Design Tasks

We developed eight UI design tasks for participants to complete. This approach ensured that their responses were contextualized within DP-related scenarios and captured their spontaneous design attitudes and behaviors toward DPs, grounded in their prior UI design experience. To ensure the real-world representativeness of the tasks, we first conducted a systematic literature review of prior studies reporting DPs. Specifically, we gathered papers published over the past 10 years from several leading privacy, software engineering, and HCI venues, including CCS, USENIX, NDSS, IEEE S&P, FSE, ICSE, ASE, SOUPS, CHI, PETS, UIST, and CSCW. By searching for the direct keywords "`[dark|deceptive] pattern`", as well as the co-occurrence of "`[deception|deceptive]`" and "`user interface`", we identified 60 papers that report on DPs. We extracted eight UI design scenarios reported in at least five papers as having privacy implications, covering a variety of UIs such as *cookie banners*, *ad personalization settings*, *data access requests*, and *account deletion*. Three researchers with expertise in privacy and HCI collaborated to create UI design task descriptions for each of the identified scenarios. This was followed by an iterative discussion and refinement process, in which three students with UI

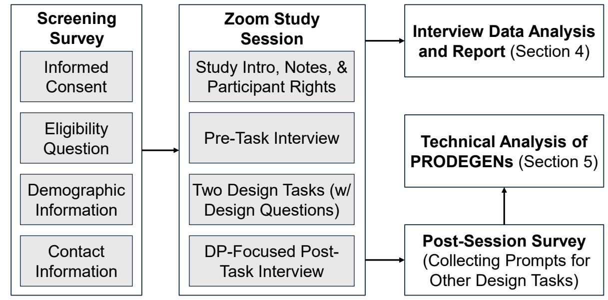

Figure 2: Study Overview

design or frontend development experience reviewed the descriptions (T1-T8) and provided feedback to enhance their readability and comprehensibility. The full list of UI design scenarios and task descriptions are presented in Table 1.

3.3. Study Procedure

As shown in Figure 2 and detailed at [19], we invited each UI designer participant to a Zoom study session. Upon joining a session, we welcomed each participant, introduced ourselves, and explained the study procedure. After that, we reminded participants that they could stop the interview at any time if they had concerns. We also emphasized that our goal was to understand their perceptions and practices, and that there were no right or wrong answers to the interview questions. We then asked pre-task questions about their UI design background, including the size of their team or company, their role within it, the tools (both traditional and AI-powered) they used for UI design, and how frequently they integrated AI into their design workflow.

Afterwards, each participant was provided with a testing account for Figma Make [15], the PRODEGEN built into Figma (a leading UI design platform with over 13 million monthly active users [80]), followed by a brief introduction to the tool's features and a demonstration of its use to familiarize them with the platform. This procedure ensured that participants without prior experience with Figma Make (e.g., those who used other tools) could familiarize themselves with it before completing the design tasks. Each participant was randomly assigned two tasks from Table 1 and completed them one by one in Figma Make. No guidance was provided during their design to preserve the authenticity of their design approach. During the design phase, participants began by creating prompts and using Figma Make to generate UIs based on their prompts. They were free to iteratively modify their prompts and interact with the tool until they were satisfied that the tasks had been completed. Upon completion of each task, we asked participants open-ended, design-related questions, such as their feelings regarding the generated UIs, their overall user experience, and the considerations they had in mind when designing the prompts.

After completing the two assigned tasks, we conducted a semi-structured interview focusing on DPs in the UIs participants created in Figma Make. We first asked participants

TABLE 1: Privacy-Related UI Design Tasks

UI Design Scenario	UI Design Task Description
T1: Cookie banners [10], [20], [21], [22], [24], [35], [36], [37], [38], [39], [40], [41], [42], [43], [44], [45], [46], [47]	Please design a cookie banner for a website. When users first visit this website, the banner will ask whether they agree to the storage and use of cookies.
T2: Ads personalization settings [2], [7], [8], [46], [48], [49], [50], [51], [52], [53]	Please design an ads personalization settings interface for a website (or app) that allows users to adjust whether their personal data could be used for personalized advertising or not.
T3: Account deletion flow [4], [9], [27], [36], [41], [46], [51], [54], [55], [56], [57], [58], [59], [60], [61]	Please design an account deletion function for a website (or app). When users choose this function, the flow will allow them to permanently remove their account and all associated personal information from the service.
T4: Account sign up with newsletter subscription [7], [36], [44], [46], [57], [62]	Please design a sign-up with email subscription option for a website (or app). When users first register, they should be able to create an account while opting in/out to receive marketing and promotional emails (such as newsletters and discount offers).
T5: Newsletter unsubscription flow [3], [6], [26], [36], [41], [44], [46], [57], [63], [64], [65], [66], [67], [68], [69], [70], [71], [72]	Please design an unsubscribe page for a website that allows users to opt in/out of receiving newsletters or marketing emails.
T6: Privacy policy consent [7], [8], [23], [27], [36], [46], [57], [66], [73], [74], [75], [76], [77], [78]	Please design a privacy policy consent interface for a website (or app). When users first access the website (or app), the interface will require them to read and agree to its privacy policy before continuing to use it.
T7: User data access request flow [27], [41], [46], [55], [66]	Please design a user data access request function for a website (or app). When users initiate this function, it will let them request and receive a copy of the personal data the website (or app) has collected about them.
T8: Permission redirect UI [46], [57], [58], [76], [79]	Please design a permission redirect UI for an app. The interface will guide users to the system settings for them to grant system permissions such as location and notification.

whether they had heard of the term dark pattern (DP). If participants were unfamiliar with the concept, we briefly explained what DPs are using the neutral definition used in the consent form, along with an illustrative example. Participants were asked whether they noticed any DPs in the generated UIs. At the same time, the interviewer, who has extensive DP research experience, identified DPs present in the participants' designs based on the taxonomy and examples outlined in Gray et al. [2] (these DPs were later confirmed by the full research team). If participants were unaware of these DPs or misidentified them, the interviewer pointed out the actual DPs in their designs and probed their thoughts about them. We then asked about participants' attitudes toward the DPs, including how they perceive the DPs (e.g., *"What are your thoughts on the presence of such dark patterns?"*), how the DPs impact user experience, and how participants' teams or companies might view the DPs. We also asked why they think such patterns might have appeared and which parts of their prompts could have contributed to them. Finally, we asked whether these DPs should be kept or removed, how the prompts could be improved to avoid them, and what other technical or non-technical approaches could help prevent similar issues in the future. Note that all interview questions were designed to be open-ended and neutrally framed, in order to avoid leading language and minimize researcher bias to ensure that participant perspectives were elicited and presented respectfully. The list of interview questions are available at [19].

At the end of the Zoom session, we thanked the participants and concluded the meeting. Afterwards, we followed up with them via a survey and asked them to create prompts for the six tasks not covered during the session due to time constraints. In this study, we used this set of prompts to facilitate the technical assessment of designers' use of PRODEGENs in Section 5. Overall, we collected participants' data, including recordings of the Zoom sessions (capturing both screen sharing during the tasks and audio during the interviews), as well as their prompts and the UIs they generated in Figma Make. Each participant completed two UI design tasks, except for P4, who completed only one due to internet issues, resulting in a total of 33 tasks. Our analysis shows that 27 of the tasks resulted in UIs containing DPs, with the most prevalent types of DPs being *Interface Interference (Visual)*, *Interface Interference (Preselection)*, and *Obstruction* (examples of these DPs will be discussed in Section 4.2). Table 2 presents the distribution of DPs across different UI design tasks and participants.

3.4. Data Analysis

We used inductive thematic analysis [81] to analyze the data. We first transcribed the audio-recorded interviews into text and anonymized all sensitive information. We then familiarized ourselves with the data, including audio transcripts, screen-sharing recordings, participants' prompts, and the UIs generated in Figma Make. Each of the three researchers from the study team independently noted potential codes related to participants' awareness, understanding, attitudes, rationale, practices, approaches, and challenges regarding DPs in their UI design. An initial set of candidate codes was generated through independent coding of the first five transcripts and subsequently consolidated into a codebook comprising 61 codes. The remaining transcripts were then analyzed using this codebook, with additional codes incorporated as needed until data saturation was reached after P14. To ensure consistency and rigor, the researchers iteratively compared codes, revisited the original data, and refined interpretations through multiple discussions. Finally, similar codes were collated into subthemes, resulting in 14

TABLE 2: Dark Patterns Appearing in PRODEGEN-generated UIs for In-Session Tasks

Task	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12	P13	P14	P15	P16	P17
T1	▲										▲ ◆	▲				▲	
T2			▲ ◆			■						▮	■				
T3	★						▲			★			▲ ★		▲		
T4		◆ ◗								▮							
T5				▲ ★		■		◆	▮		■				▲	▮	
T6		■			■		▲ ★							▮			▮
T7								▲						★			
T8			▲		◗				▲								▲ ◗

* ■: *Interface Interference (Preselection)*, ▲: "*Interface Interference (Visual)*," ◆: "*Sneaking*," ★: "*Obstruction*,", ◗: "*Forced Action*," and ▮: "*None of DP*."

subthemes. For example, P8 noted that using different colors makes elements more visually salient, while P12 mentioned that color differences help indicate distinct button functions; both were categorized under the subtheme "DPs provide users with visual clarity". These subthemes were then grouped into overarching themes by identifying conceptual relationships among them. For instance, subthemes such as "Benefit businesses" and "Benefit end users" were grouped into overarching themes like "Perceived Benefits of DPs" based on their co-occurrence across participants. In total, we identified three overarching themes. A thematic map was created to illustrate the relationships among themes, subthemes, and codes. Finally, we reviewed and refined this map to ensure that the themes and subthemes accurately captured the meanings within the coded data and formed a coherent pattern.

4. Findings from the Designer-Centric Study

4.1. Designers' Prior Knowledge about DPs

Awareness of DPs among participants is generally low. Only four participants (P7, P9, P10, P14) reported having heard of DPs, while the majority (13 participants) were unfamiliar with the concept. Among those who were aware of DPs, the sources of their knowledge varied. P7, a freelancer, mentioned learning about the concept from a YouTube video. Participants P9, P10, and P14, on the other hand, were all working in large companies and learned about DP through formal training, such as school training and company's bootcamp. They later gained a deeper understanding through work experience and social media. For example, P14 noted:

> *"I actually learned about it from one of my professors, and then I went on YouTube to learn what it truly is."*

These findings suggest that the awareness of DPs among designers is relatively scattered and largely dependent on formal education and informal learning via social media, lacking systematic training or a unified framework. This uneven understanding may affect their ability to identify and address DPs effectively during user experience design.

4.2. Designers' Attitude toward DPs

Most UI designers could not identify DPs in UIs and, even when they do, tend to keep them. After we introduced the concept of DPs, four participants were able to identify DPs in the UIs they had generated, whereas 13 participants did not recognize any. When we subsequently pointed out the DPs in their UIs, only five participants agreed that their UIs indeed contained a DP, six participants did not consider their UIs to contain any DPs. The remaining two participants agreed that some of the DPs we identified are indeed DPs (e.g., preselected choices in T4), while arguing that others are not (e.g., color differences between the reject and agree buttons in T1). Even among the five participants who agreed that the UIs contain DPs, three of them preferred to keep the DPs and use them in real-world UI design.

4.2.1 Perceived Benefits of DPs

Participants consider DPs to benefit businesses. Three participants (P1, P10, P11) believed that DPs could benefit companies and their business interests in several aspects:

(1) Designers viewed DPs as a strategic design choice for collecting more data from users. Some DPs observed in the generated UIs are related to the processing of user personal data. For example, making the "*Accept all*" option in the cookie banner more visually prominent or framing consent requests as "*Help us improve your experience*" can nudge users into permitting the collection and sharing of more data than they initially intended. We classify these DPs as belonging to *Interface Interference*, where manipulative visual presentations or interaction flows steer user choices about data collection. However, participants indicated that they would like to retain such DPs, viewing them as a strategic design choice that serves business interests. For example, P11 described intentionally designing a cookie banner to hide the "*Reject all*" option, making it less likely that users would quickly decline data collection:

> *I hid the "Reject all" option underneath the interface, because I more or less did not want users to make decisions faster or more transparently. From the perspective of a shopping website designer, I want to obtain users' data and understand their journey on the site.*

While noting that the use of DPs could be beneficial, participants also noted that some applications of DPs constitute abuse and are malicious, e.g., deliberately blocking users from unsubscribing or preventing them from exercising control over their data. For example, the same participant (P11) mentioned that:

> *As an employee of the company, what I'm doing is for the company's benefit. However, if I use malicious tactics, such as preventing you from unsubscribing, or other more extreme measures, I consider that an abuse of dark patterns.*

(2) DPs help companies promote their products and services. Participants also reported that they would like to keep some DPs because they serve marketing purposes. An example is the use of pre-checked options for newsletter subscriptions during account sign-up (T4). We interpret this DP as *Interface Interference* [2] (or more specifically, preselection [1]), as it automatically opts users into a service they may not want, making it harder to decline while increasing the number of users receiving promotional newsletters. Besides the DPs identified in the UI design tasks, some participants also cited UIs from their own companies, which leverage DPs to promote slow-selling products by disguising them as personalized recommendations – practices that were previously less known. For instance, P10 described:

> *"I found that our company's "For You" section, which is supposed to be a recommendation panel, actually displayed slow-selling products. This was created by the marketing team. After reporting it, it was changed so that the "For You" section truly reflected personalized recommendations."*

(3) DPs help companies retain customers. Participants mentioned that some DPs in the generated UIs would help companies retain customers, and they would not remove them in the company's interest. The most prominent examples of such DPs are *Obstruction* [2], which are designed to make a process more difficult than necessary, with the goal of discouraging users from taking actions they intend to perform, e.g., making the unsubscription process cumbersome, obscuring exit options, and complicating the delete account feature. When probed about their motivations, some participants revealed the mindset that users intending to leave are no longer customers and, therefore, do not warrant customer obsession. As a result, they felt no obligation to consider the negative impacts that these DPs might have on users who unsubscribe or delete their accounts. For instance, P1 explained his rationale for retaining the DPs:

> *"We don't need to worry about the user experience for those who use the delete option, because once they do, they're no longer our users.*

Participants consider DPs to benefit end users. Participants believed that not all DPs necessarily harm users. Instead, some DPs can benefit the end users in their user experience, such as helping users avoid accidental errors, get more visual clarity, learn more about a service, and increase their satisfaction.

(1) Participants found that DPs help users avoid accidental errors in their UI interactions. Four participants (P7, P10, P13, and P14) argued that certain DPs can help users avoid errors. For example, P7, P10, and P13 preferred to keep DPs that add extra steps before account deletion, such as asking users to confirm multiple times, select reasons for deletion, and re-enter their passwords (even though they were already signed in when attempting to delete their account). We consider that such DPs fall under the category of *Obstruction* [2], as they introduce friction into the deletion process, potentially trapping and frustrating users, which may ultimately prevent them from deleting their accounts. P14 noted that they would use DPs that include alarming visual elements (e.g., large, big red button) to warn users about the consequences of deleting an account. However, such DPs, as reported in [82], would evoke an emotional response in the user and create the impression that deleting an account is a dangerous action, which discourages the user from proceeding. However, from the designers' perspective, participants argued that all these DPs in the account deletion UIs are rooted in Nielsen's usability heuristic [83], which asserts that interfaces should prevent unintentional errors. Thus, they perceived these DPs as user-centered design practices that could help users rather than cause any harm. For instance, P14 mentioned that:

> *"There are certain instances probably your might just, you know, mistakenly type on your keyboard or probably tab on your phone. So it might be that you don't necessarily want to proceed to this page, probably by accident."*

(2) Participants note that DPs provide users with visual clarity. Participants (P8, P12, P15, P16, and P17) did not consider using different colors for buttons to be a DP, even when the "*Accept*" button appeared more prominent than "*Opt Out*." They argued that color differentiation is a good design practice, as different colors help users understand that the buttons perform different functions, with visual contrast clarifying meaning and guiding decision-making. However, we interpret such a design as constituting *Interface Interference*, since a visually dominant "*Accept*" button and a less noticeable "*Opt Out*" button are common DP tactics used to highlight options that benefit the company while de-emphasizing those it disfavors. This subtly nudges users toward clicking "*Accept*," even though both options should be equally accessible. For example, P12 noted

> *"It is a good design because if the two buttons had the same color, users might be tricked into thinking they perform the same function."*

(3) Participants believe that DPs help users learn more about a service. Some UIs generated by PRODEGEN offer a one-month free trial during the registration process. Participants considered that such free trials help users by allowing them to explore or understand product features they might not initially be familiar with. However, based on the taxonomy [2], we classify this type of design as *Sneaking* and *Forced Action*, as the services conceal the true cost or intent behind the "*free trial*". Specifically, free trials often require users to enter payment details upfront, leading to automatic charges if users forget to cancel, thereby forcing them into a transaction they may not have intended. This design benefits the company by converting trial users into paying subscribers through inertia or forgetfulness. For instance, P2 emphasized the value of providing limited-time free trials as a means of exposing users to additional product functionality:

> *"No, I won't modify it because there should be a free one-month premium trial to test the app and know more about its features."*

(4) Participants found that DPs can increase user satisfaction. We found that some participant-created UIs adopt linguistic framing that influences users' perceptions of value. For instance, participants use selective clarity and strategic wording (e.g., "*5 GB data transfer*") to evoke a sense of worth or exclusivity, encouraging users to choose higher-priced options, while employing vague or less informative language (e.g., "*high-speed data*") for cheaper plans. We consider such UI designs to constitute the *Interface Interference* DP, as they manipulate how information is presented to influence user choices. This design creates a sense of emotional satisfaction associated with higher spending, subtly nudging users toward premium plans and reinforcing the perception that paying more equates to greater value and quality. By leveraging ambiguity and presentation cues, the design makes users feel that the premium option reflects better personal worth, thereby improving both conversion rates and perceived customer satisfaction, as P11 described his work experience:

> *"I want users who are willing to spend only $5 a month to feel my product is great value for money, and I also want users who spend $25 to feel what they are buying is even more valuable."*

Participants consider DPs as aligned with industry norms or following best practices. For example, participants noted that the use of color contrast to make one option (typically "*Accept*" or "*Continue*") more prominent than others (such as "*Decline*" or "*Cancel*") is widespread in real-world interfaces. However, we believe that this design falls under *Interface Interference*, as it steers users toward a particular choice through visual emphasis. Participants noted that such practices have become normalized within the industry, and therefore they did not view them as deceptive tactics but rather as conventional design standards that align with users' expectations of interface behavior. This perspective reflects a utilitarian UX mindset, designing for the majority outcome to reduce friction and improve flow, even if it compromises the neutrality of choice. For instance, P5 mentioned:

> *"I think that visual is okay because we see it everywhere... I don't think all options are equal and UX tend to favor the one most people choose."*

Participants also believed that as long as a UI design complies with legal norms, including pre-selected options that align with the company's preferred outcomes, the design is acceptable. Yet, based on the taxonomy [2], such a design represents a DP of *Interface Interference*, as it subtly steers users toward a particular choice unless they actively change the default setting. Although some participants acknowledged the potential bias in these designs, they emphasized that the practice is legally justified and reflects the institutional normalization of DPs, provided users retain the ability to opt out later. This represents a compliance-minimalist approach, where the interface technically meets legal requirements (by including an opt-out mechanism) but prioritizes regulatory compliance over user experience. For instance, P5 explained:

> *"Legal eventually gave us the answer: you don't have to show it to them on the flow, but you do have to give them an option afterwards to opt out."*

4.2.2 Perceived Harms of DPs

Although most participants found the use of DPs is beneficial, a small number of participants expressed reluctance to implement specific DPs due to reputational and usability concerns. First, two participants (P10, P12) recognized the reputational risks of using *Obstruction* DPs in design. They noted that overtly manipulative practices, such as making unsubscribing or opting out cumbersome and time-consuming, create negative user emotions (e.g., anger, frustration, and betrayal) and ultimately damage the company's reputation and brand perception. These participants viewed such DPs as ethically unacceptable and counterproductive, as such DPs prioritize short-term user retention over long-term credibility and loyalty. For instance, P10 illustrated this issue using the example of The New York Times:

> *"I feel like they (The New York Times) are serious about journalism, and they are a good news company. Their dark pattern is that you must call to unsubscribe, which is very annoying for users."*

Only one participant raised concerns about the usability issues caused by DPs. He identified *Obstruction* DPs, such as requiring users to download their data as a JSON file, a machine-readable but not human-readable format, as creating confusion and barriers for users. For most non-technical users, this format is inaccessible, requiring additional technical skills or file conversion to interpret the content. Although the company technically provides "data access," this usability barrier effectively discourages users from viewing or managing their own information. Participants' reasoning reflected a user-centered ethic, as they criticized this practice for being unfriendly and exclusionary. They argued that companies should provide data in accessible, human-readable formats to promote genuine transparency and user empowerment, as P14 described:

> *"This part is stating that the download will be in JSON file, which I believe some users might not be tech-friendly in the sense that they might not know what a JSON file is. "*

4.3. Approaches to Removing DPs

When asked about the approaches the UI designers believe could remove the DPs in the generated UIs, they suggested several potential solutions.

Through prompt modification. Participants reported three strategies they could adopt to modify the prompts and remove DPs in the UIs. The first strategy focuses on prompt-level control (P3, P14, P15), explicitly instructing the AI not to generate UI elements that contain DPs by using corrective language and intentional guidance grounded in design ethics. For example, after we explained that using buttons with different colors could be considered as DPs, P3 began providing more specific instructions in the prompt, such as adjusting the prompt to say "*buttons have similar color*" to remove the DPs. Similarly, when the AI initially generated a feature that allowed users to download their data in JSON format, P14 revised the prompt to say: "*Instead of the data being downloaded in JSON format, make it plain text*." P15 also used this strategy, adding prompts such as "*don't use default choice*" to avoid misleading defaults.

The second strategy focused on element-level prompt editing. Rather than concentrating on the entire prompt, this strategy emphasized micro-level control, allowing participants to guide the AI in making context-specific, localized modifications instead of broad, automatic changes. Such control helps prevent unbounded AI edits that could introduce new DPs, inconsistencies, or unintended global changes. For instance, P10 would select specific UI components and give targeted instructions such as "no pre-selected" to eliminate deceptive defaults.

> *"If you don't specifically select the element and only give Figma Make instructions verbally, it often ends up modifying other parts of the interface as well, which you actually don't want to change. So I usually select the specific element before asking it to make changes."*

The third strategy focused on code-level control, representing a more technically intensive approach to removing or correcting DPs. In this strategy, participants directly intervened at the code level, allowing for complete control over design behavior and visual implementation. However, this method required substantial technical knowledge and was time-consuming, as participants needed to interpret and modify complex PRODEGEN-generated code structures. For instance, P17 mentioned that they sometimes modified the UI by directly editing the generated code but noted that this approach was challenging because they first had to spend considerable time understanding the structure of the code before making any meaningful changes.

Through usability testing. Another approach highlighted by participants (P9, P13, P14) for removing DPs was internal usability testing. This method enables designers to identify potential DPs early in the design process and address them before products are released to the public. The advantage of this approach lies not only in the proactive detection of DPs but also in serving as an ethical audit, embedding ethical principles into the design process. However, the effectiveness of this strategy depends on whether usability testing teams are trained to recognize and avoid DPs and whether the organization fosters a shared ethical standard. In companies lacking a strong ethical culture or explicit guidelines, usability testing may focus solely on technical functionality rather than DPs. For example, P13 noted:

> *"Our company focuses on transparency and user trust, reviewing designs through usability testing, ensuring options like Accept and Decline are equally clear and follow ethical guidelines. We also train the design team to support and always avoid manipulative elements, early in the design process."*

Compromise through rationalization. Three participants (P6, P9, P13) expressed uncertainty about why DPs appeared in PRODEGEN-generated designs. They reflected on both the behavior of PRODEGENs and their own instructions, attempting to diagnose the cause of the DP rather than immediately correcting it. These participants were unsure whether the issue stemmed from their own prompts or from the AI system itself. This ambiguity led to a form of compromise, in which they did not directly remove the DPs but instead attributed their presence to AI limitations or misinterpretations rather than intentional design choices. For example, P6 reflected on the AI's output and thought it was not caused by the prompt: *"Well, I don't think those words are used... using those prompts in AI can cause a dark pattern."* Yet, another participant (P9) suggested that AI might have misinterpreted their instructions: *"I'm not pretty sure as to why, because probably the AI misread my prompts in a way. Although I added that, you know, like 'user friendly'."*

5. Technical Analysis of PRODEGENs

This section explores answers to RQ4 (anti-DP and pro-DP UI designer practices) and RQ5 (how different PRODEGENs produce or can avoid DPs) by running prompts used by UI designers on real-world PRODEGENs, including Figma Make [15], Lovable [84], and Vercel V0 [17], all of which are among the top AI prototyping tools [85], [86].

Prompt dataset for UI design. In the post-session survey, we received responses from 16 study participants, including prompts created by the designers for a variety of design tasks (listed in Table 1). P4 chose not to complete the survey, and P11 did not provide a prompt for T4. Therefore, our dataset includes a total of 127 prompts for UI design with PRODEGENs, with each task having 16 prompts, except for T4 that has 15. We expect this dataset to reflect how designers with preliminary DP knowledge (who participated in our task-based and interview studies) would leverage PRODEGENs in real-life UI design.

5.1. RQ4: Understanding UI Designer Practices

To answer RQ4, we perform (1) an *attribution study* to understand the key texts in the designer prompts that contributed most to the generation or disappearing of DPs, followed by (2) a *thematic analysis* of the texts to form a tentative categorizations of designer anti-DP and pro-DP practices. Note that, due to the limited number of designer prompts we obtained, we explore a preliminary understanding of UI designer practices (which we found helpful in informing useful suggestions to the community), rather than aiming for a comprehensive, large-scale analysis.

5.1.1 Analysis Methodologies

Perturbation-based text attribution. In machine learning (ML), attribution helps identify which parts of the input contributed most to the output of ML models. Various attribution methods have been proposed, with two major categories being *gradient-based* [87], [88], [89] and *permutation-based* [90], [91], [92], [93], [94] attribution. In this study, we adopt permutation-based attribution due to the noise in gradient-based methods [95] and the difficulty of inspecting gradients in black-box PRODEGENs.

Specifically, as shown in Figure 8 (in the Appendix), each prompt ($\mathcal{P}$) consists of a main UI design *purpose* sentence (q) that describes the overall design task (such as "*design a cookie banner...*"), along with a series of relatively independent sentences that detail the design *expectations* (e_i) (such as "*keep it clean and easy to read*"), i.e., $\mathcal{P} = q, e_1, e_2, \ldots, e_n$. With multiple runs of PRODEGENs, each prompt will be associated with a probability score indicating the likelihood that the prompt leads to a DP, $S_{\mathcal{P}} = f(O = DP|\mathcal{P})$. Informed by the *Leave-One-Out* permutation method [90] (a classical method suited for isolating the effect of individual inputs), we evaluate the contribution of each e_i to the generation of a DP by removing each e_i and calculating the conditional probability change. In other words, for each e_i, we compute the probability that the modified prompt $\mathcal{P}' = \mathcal{P} \backslash e_i$ leads to a DP, i.e., $S_{\mathcal{P}'}$. The change in the probability score, $\Delta = S_{\mathcal{P}'} - S_{\mathcal{P}}$, indicates how much e_i affects the occurrence of DPs. For example, $\Delta > 0$ suggests that removing e_i increases the occurrence of DPs (in other words, using e_i indicates a potential anti-DP practice for UI designers). Conversely, $\Delta < 0$ suggests a potential pro-DP practice.

In this study, we manually identified the purpose sentences in the 127 prompts and then break each prompt into smaller sentences using spaCy [96]. By permuting these prompts and removing each expectation sentence, we generate a total of 335 unique prompts. For each prompt, we run Figma Make five times to calculate the probability that it leads to a DP, resulting in a total of 1,675 runs. We choose five runs, as we observe that, although Figma Make is probabilistic, the generated UI across different attempts is largely consistent when the same prompt is used.

Labeling DPs. To determine whether the PRODEGEN-generated UIs contain DPs, two authors familiar with the concept of DPs independently reviewed the 1,675 generated UIs and labeled potential DPs. We leveraged Krippendorff's alpha coefficient [97], a widely used measure to evaluate agreement between the professionals. Specifically, the two professionals achieved a Krippendorff's alpha of 0.91, indicating high reliability in DP labeling. The two professionals then compared their results and addressed disagreements through discussions. In total, among the 1,675 UIs, 1120

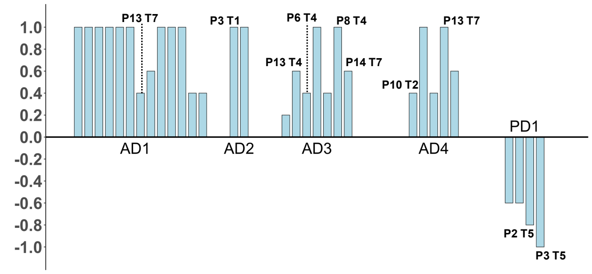

Figure 3: Probability changes caused by different sentences in the prompts regarding DP generation.

(66.9%) were found to contain at least one DP. Further comparison of the probability scores revealed that the permutation of 26 expectation sentences resulted in $\Delta > 0$ (potential anti-DP practice), while 4 expectation sentences resulted in $\Delta < 0$ (potential pro-DP practice). Figure 3 shows the distribution of the probability change in DP generation. Later, we use these sentences as a basis to approach a preliminary understanding of designer practices.

Inductive thematic analysis. We used a similar thematic analysis approach [81] (as described in Section 3.4) to analyze the 30 sentences. The two cybersecurity professionals read all the sentences and independently developed initial codes related to DP-specific details and designer considerations. They then compared their initial codes and refined and consolidated them through an iterative process. This process resulted in a total of 32 codes, with two sentences yielding more than one code. After grouping these codes, we identified five themes, with four suggesting anti-DP designer practices and one corresponding to potential pro-DP practices. We released the dataset of prompts and the sentences identified as affecting the appearance of DPs on an anonymous website [19].

5.1.2 Potential Anti-DP Practices of UI Designers

AD1: Provide detailed UI design specifications. Fourteen of the 26 sentences with $\Delta > 0$ provide detailed UI design specifications. For instance, P13's prompt for T7 specified that "*a user data request flow with a confirmation step and a download option*". With this sentence, the generated "data access request" UI uses *a single page* for users to confirm and submit their data access request (as shown in Figure 4a). However, without this text, the generated UI will require users to navigate through four different pages to request, verify identity, process, then download (as shown in Figure 4b). This UI design represents an Obstruction DP that deliberately makes it difficult for users to exercise their right of data access, by requiring extra navigation through the pages. Although it is difficult to draw a conclusive conclusion on the cause, we believe that providing detailed UI design specifications places constraints on the probabilistic UI generation of PRODEGENs, reducing the space for unintended manipulative features to emerge.

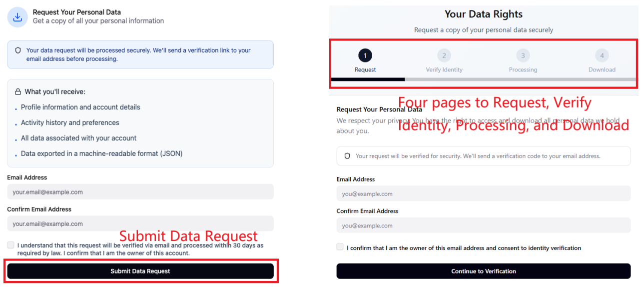

(a) UI without DP (b) UI with Obstruction

Figure 4: Two UIs for data access request

AD2: Include explicit guidelines to avoid DPs. Two of the 26 sentences provide clear guidelines for avoiding specific DPs. For example, P3's prompt for T1 explicitly asks PRODEGENs to "*use the same color for buttons*." When this guideline is not used, the prompt leads to DPs in all runs of Figma Make, where the "*accept*" button in cookie banner UIs is visually more prominent than the others – an instance of *Interface Interference*. This attribution result confirms the effectiveness of the strategies reported by participants in the interview study (Section 4.3), who noted that they could use targeted changes to the prompts to eliminate DPs.

AD3: Add usability requirements to UI design. Seven of the 26 sentences highlight the requirement that the generated UIs should be usable and simple. For example, P6's prompt for T4 (Account sign-up with newsletter subscription) instructs to "*keep it clear, minimal, and easy to complete*." With this requirement, the generated UI has only a 40.0% (2/5) chance of including a preselection for the newsletter subscription option (again, an instance of *Interface Interference*). Without this requirement, however, the UI always comes with the option preselected, influencing users' decision on the potentially unwanted subscription. We notice that similar requirements also helped reduce DPs in the prompts of three other designers, such as "*make the UI user-friendly*" (P8), "*keep it clear and polite*" (P13), and "*keep the page clean and simple*" (P14), confirming the cross-pollination of usability and DPs in UI designs.

AD4: Asks for UI designs that prioritize ethics and privacy. Five of the 26 sentences instruct PRODEGENs to follow ethical practices and respect users' privacy. For example, P13 created a prompt asking PRODEGENs to "*make it feel secure and respectful*" for the UI design task of user data access request (T7). No DPs were observed in the generated UI when this instruction was present, while all five runs of Figma Make without this instruction contained a DP with Obstruction, identical to the one presented in Figure 4b. Similarly, P10's prompt for T2 (Ads personalization settings) includes an instruction to "*respect users' privacy preferences and increase transparency*." With this instruction, the UI has a 60.0% lower chance of containing a DP that opts users into ads personalization by default (which allows businesses to use personal data, such as browsing history and search queries, for targeted advertising). These examples highlight the fact that Figma Make, specifically its underlying GenAI model (i.e., Claude Sonnet 4), was capable, to some extent, of assessing whether the UI design is ethical and respects privacy. However, in the PRODEGEN, this capability is not enabled in its design and implementation. As a result, *all UI designers have to take on the responsibility* of carefully crafting prompts, identifying, and preventing the emergence of DPs when using PRODEGENs for UI design tasks.

5.1.3 Potential Pro-DP Practices of UI Designers

PD1: Use of DP-inducing UI prompts. All four sentences leading to $\Delta < 0$ are found to include information that induces the appearance of certain DPs. An example is P3's prompt for T5 (newsletter unsubscription), which contains an instruction to "*use a hyperlink with the unsubscribe title*". This sentence precisely describes an Obstruction that makes confirming unsubscription difficult, as shown in Figure 5a, where a less noticeable link is used for users to confirm unsubscription, compared to the more prominent button to cancel the unsubscription (i.e., stay subscribed). We were unable to conclude the cause behind the designer using the instruction (after the study has been done). However, we suspect that the UI designer either used or observed such a design in their own practice but did not recognize it as a DP due to a lack of awareness.

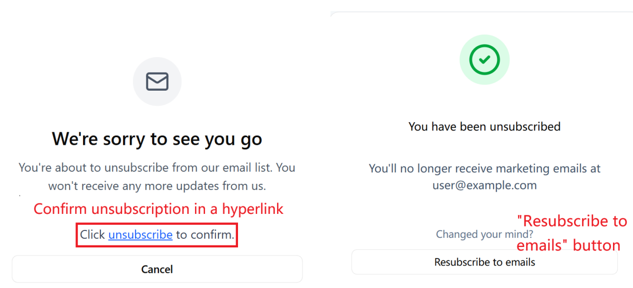

(a) Unsubscription UI with obstruction. (b) Unsubscription UI with a sneaking resubscription.

Figure 5: Two UIs for newsletter unsubscription

Another prompt created by P2 for T5 also induces DPs, but the cause is different. Specifically, P2 describes part of the unsubscription flow as "*clicks 'Confirm Unsubscribe' → success message appears → 'Resubscribe' link displayed.*" As a result, the unsubscription UI immediately displays a resubscribe button upon unsubscription, as shown in Figure 5b, similar to the *Sneaking* DP reported in [98]. This DP creates the impression that the unsubscription was not fully completed when it actually was, leading the user to accidentally click the button and resubscribe to what they intended to unsubscribe from. Upon closer analysis of this prompt, we found that P2 used a structured, lengthy prompt template, as shown in Figure 9. In other words, not only can PRODEGENs generate UIs with DPs, but the everyday use of general-purpose LLMs may also produce UI design instructions that induce DPs, leading to *the propagation of DP-inducing instructions from such LLMs to* PRODEGENs.

5.2. RQ5: Comparative Analysis of PRODEGENS

5.2.1 Overall Results of Different PRODEGENS

TABLE 3: DPs Caused by Different PRODEGENS

	Figma Make	Lovable	Vercel V0
PRODEGEN Runs w/ DPs	409 (64.4%)	420 (66.1%)	403 (63.5%)
Prompts w/ DPs	88 (69.3%)	97 (76.3%)	95 (74.8%)
Design Tasks w/ DPs	8	8	8

Besides Figma Make, we analyze two other PRODEGENS: Lovable [84] and Vercel V0 [17], which UI designers reported using most often that others. Similar to Section 5.1, for each PRODEGEN, we run the prompts in the dataset five times and report the probability score indicating whether the resulting UIs contain DPs. We followed the same labeling process to determine DPs. In total, we ran PRODEGENS 1,270 times (127 prompts, each run five times on two PRODEGENS). We reused the Figma Make results from Section 5.1.

DPs are commonly found in UIs generated by PRODEGENS, with no substantial differences across PRODEGENS. Specifically, the two cybersecurity professionals reported that, in total, two PRODEGEN runs led to UIs containing DPs, with a Krippendorff's alpha [97] of 0.93. Table 3 shows the distribution of DPs across different PRODEGENS. Notably, Lovable generates 420 (66.1%) UIs with DPs out of 635 runs, while Vercel V0 generates 403 (63.5%) UIs with DPs, compared to Figma Make which generated 409 (64.4%) UIs with DPs. Given the similar results, we believe that none of these PRODEGENS are significantly better than the others at mitigating DPs in UI design generation. This finding is confirmed by the *Repeated Measures ANOVA* [99], which showed no significant difference among the conditions (F(2, 30) = 0.34 , $p = 0.713 > 0.05$).

5.2.2 Potential Design and Implementation Failures

Figure 6 uses a heatmap to visualize how the three PRODEGENS produce UIs with DPs across different prompts and UI design tasks. We hypothesize that the differential aspects in the figure, where the probability of one PRODEGEN is significantly higher than that of the other PRODEGENS in terms of producing UIs with DPs, reveal potential design and implementation failures in the PRODEGEN. To test this hypothesis, we leverage an empirical rule to narrow down the scope of manual analysis. Specifically, for each prompt, given the probabilities that the three PRODEGENS generate UIs with DPs, i.e., (S_1, S_2, S_3), we analyze PRODEGEN i for the prompt if the z-score of S_i is greater than 1. With this rule, we prioritize the analysis of 12 prompts for a specific PRODEGEN, which are highlighted with rectangles in Figure 6.

F1: Reliance on LLM-based understanding of UI code. We found that in the UI code generation step, PRODEGENS generally relied on LLMs' capabilities to evaluate whether the UI code met the design expectations specified in the prompt. However, LLMs are limited in analyzing the complex UI code, which causes DPs to appear in the UIs generated for five prompts (highlighted in red in Figure 6). For example, for P3's prompt in T1, *Lovable* is more likely to cause DPs compared to the other two PRODEGENS. Specifically, P3 used a prompt that contains an anti-DP practice, i.e., "*use the same color for buttons*." for generating cookie banners, which led Figma Make to avoid the *Interface Interference* DP (see Section 5.1.2). However, Lovable still creates UIs with the DP, where the "*Accept*" button appears in a filled, vivid blue, while the "*Decline*" button is white and less noticeable. From the UI design generation logs, we found that Lovable mistakenly considers both buttons to be the same blue color, noting that "*Currently both buttons use the same blue color from the design system*." Closer analysis of the UI code reveals that the "*Decline*" button appears blue only when the website is in dark mode, while the UI code sets the mode to light by default. In other words, Lovable failed to analyze the dependency between the website mode and the button color, resulting in an incorrect assessment that the designer's expectation had been met (but actually not). To address the failure, one potential suggestion for Lovable could be to base the evaluation of UIs against designer expectations not only on the UI code (which may become unreliable as the code complexity increases), but also on more reliable methods, such as analyzing the UIs (e.g., screenshots) that are actually rendered.

F2: Failure to compose incompatible UI components and templates. For the other seven prompts (highlighted in blue in Figure 6), we found that DPs largely stem from the failed composition of incompatible UI components and templates from different sources. An example is P16's prompt in T6, where the designer asks for a design of "*privacy policy consent interface for website or app*." We noticed that, Figma Make appears to search for and compose pre-built UI components from an existing dataset, such as Tailwind [100]. However, while composing the application UI layout (with a white background) and the individual button (with white text on a transparent background), incompatibility issues arise. This incompatibility leads to an *Obstruction* DP where the "*Accept and Continue*" button in the privacy policy consent UI is highlighted where "*Decline*" is invisible, as shown in Figure 7. We note that such failures in composing different UI components are not accidental errors, but rather a design gap that appears in other PRODEGENS as well, such as Lovable when handling P17's prompt for T6.

6. Discussion

6.1. Suggestions to Different Stakeholders

Built-in DP mitigation in PRODEGENS. Our study reveals that it is common for PRODEGENS to generate UI designs containing DPs. Currently, UI designers need to craft proper designs or to review the generated UIs to ensure that they are free of DPs. Unfortunately, many UI designers have low awareness of DPs (Section 4.1) and are likely unable to identify or avoid DPs in UIs, leaving a chance that DPs

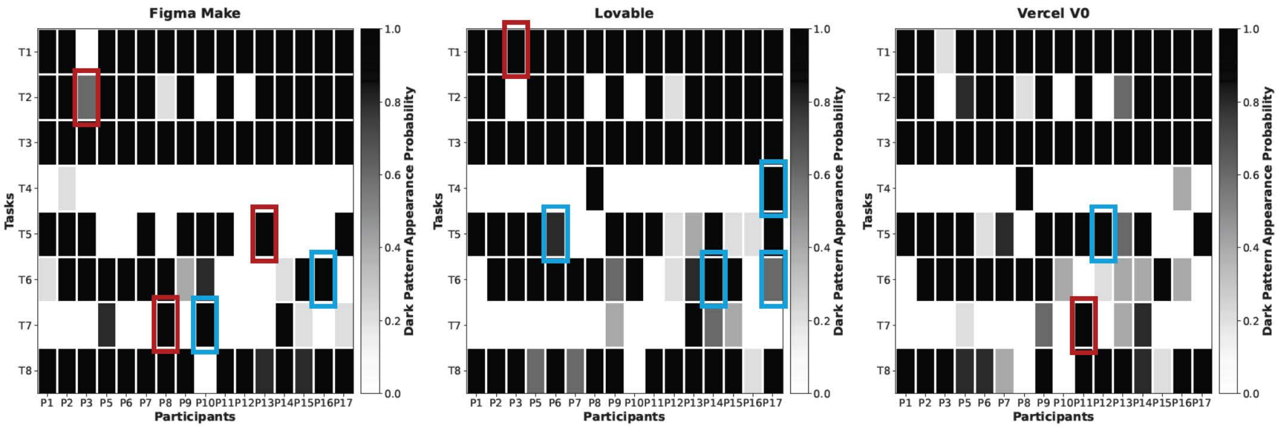

Figure 6: The comparison of three PRODEGENs in terms of generating DPs. (: F1, : F2)

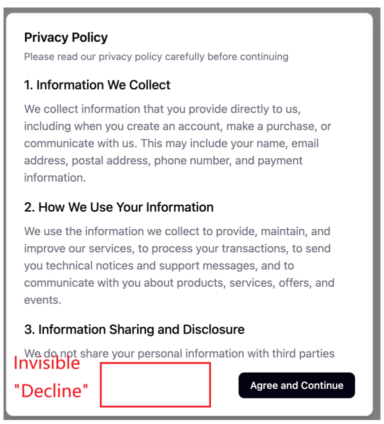

Figure 7: An obstruction DP for privacy policy consent (T6).

will continue to be generated and adopted in the real world. Through the analysis of the prompts, we found that the LLMs used by PRODEGENs are often capable of reasoning about DPs and avoiding them, for example, when prompts emphasize usability, ethics, and privacy. Therefore, we think that a technically feasible option could be for PRODEGENs to support built-in DP mitigation, which would reduce the burden on all UI designers (whether senior or novice).

DP regulation based on more concrete harm–benefit tradeoffs of multiple stakeholders. A key finding of this study is that, although many DPs are reported as harmful from end-user perspectives [25], [43], [57], [66], UI designers often perceive them as beneficial to both businesses and end users and consider them aligned with current industry norms (Section 4.2.1). As a result, even though DP regulations exist and base enforcement on user perceptions (i.e., the "*net impression*" of users, as noted in the FTC staff report [101]), UI designers (or services) may continue to adopt such practices in real-world UIs, in part due to their perceived benefits. This situation creates a practical challenge for existing DP regulation, which must operate under limited resources while addressing the widespread presence of DPs. Currently, regulatory agencies like the FTC provide high-level guidance that "*businesses should look not just at the effect their design choices have on sales,..., but also on how those choices affect consumers* [101]." However, due to the lack of concrete and operationalizable criteria, such guidance is difficult for designers to follow in real-world design processes. Hence, for DP regulations to be more effective and scalable, we argue that more concrete harm–benefit tradeoffs should be defined and discussed across multiple stakeholders, including end users and designers (or services), and that enforcement should be grounded in these tradeoffs rather than relying solely on "net impressions" or high-level design judgment.

Systematic DP knowledge and skills training for UI designers. Our study finds that the majority of UI design professionals (76.5% of participants) are still unaware of the concept of DPs, limiting their ability to address harmful patterns in design tasks. Therefore, we advocate for increased efforts to provide UI designers with systematic training, through formal education or workforce development, on DPs and their harmful impacts before they engage in real-world UI design. Moreover, we found that some UI designers use prompts containing anti-DP practices, while others either do not or even adopt pro-DP practices that lead to the appearance of DPs. Therefore, in addition to introducing DP concepts, we suggest that education and training also cover how to compose prompts with anti-DP practices, so that designers can develop skills for effectively interacting with PRODEGENs.

6.2. Limitations and Future Work

This study has several notable limitations. First, the findings are based on a small sample of participants and a limited set of design tasks, which may restrict the generalizability of the results. In this study, we relied on thematic saturation to determine the number of participants and developed design tasks that were reported to have privacy implications. Therefore, we consider the findings likely to capture common types of DPs that pose privacy risks to end users, rather than to provide a statistically robust and comprehensive evaluation of designers' practices (which would require a significantly larger sample size and a wider variety of design tasks). Second, we recruited US participants because the US has extensive research and legal enforcement related to DPs, such as regulations under the FTC Act [102], [103], that combat deceptive practices. While we believe the findings from these US-only participants provide valuable insights, they may not generalize to other regions where different legislation shapes designers' attitudes and practices. Hence, future research is needed to examine how designers' perceptions and practices about DPs vary across countries with different regulatory frameworks. Third, due to the lack of transparency in commercial PRODEGENs, we were only able to assess them as largely black boxes and make best-effort attributions and analyses of their responses to prompts. This has limited our ability to comprehensively identify design and implementation failures and offer more technical recommendations to improve PRODEGENs. To address this gap, our future work plans to engage the developers of PRODEGENs to gain a better understanding of how they are designed, implemented, and the considerations surrounding the generation of DPs, among other factors.

We also note that our goal was to understand how designers notice, evaluate, and respond to DPs generated by GenAI tools, rather than to compare AI-generated DPs with human-designed ones, which would require controlled comparative experiments. Given that GenAI-powered UI design tools are still emerging, this exploratory study is a necessary first step to capture the reasoning processes that influence whether AI-generated DPs are resisted or normalized in practice. While controlled comparative studies are an important next step, our work provides the behavioral context needed to design and interpret such causal comparisons. Additionally, we did not examine the relationships between participants' awareness of, or training in, DPs and the prompts they used. Investigating this relationship would require a larger participant sample and carefully designed educational interventions and assessments, which we leave for future work.

7. Related Work

Research on Dark Patterns. Recently, dark patterns (DPs) have attracted broad attention across multiple research domains, such as Human–Computer Interaction (HCI) [2], [36], [41], [46], [49], [51], [68], [69], [79], security and privacy [4], [6], [7], [8], [21], [23], [24], [25], [47], [59], [61], [71], [74], [104], [105], and software engineering [50], [70]. Some of the studies investigated DPs by developing taxonomies [2], [4], [105], or unveiling the presence of DPs across various platforms and services, such as social media [51], shopping websites [68], digital games [79], and smart home [49], [74]. Several other studies reported the impacts and perceptions of end users [25], [43], [57], [66], or proposed DP mitigation strategies like education [46] and adding technical interventions on UIs [41], [44]. The study most closely related to this research is Zhang et al. [3], which reported on the perceptions of UI designers regarding the implementation of privacy dark patterns. However, while both this study and our study engage UI designers, they differ in that Zhang et al. examine how designers interpret other designers' DPs (attribution), whereas we study how designers justify their own willingness to use DPs (self-justification). This is conceptually and methodologically distinct: our participants articulate internal decision frameworks that legitimize DP adoption, which Zhang et al. were not designed to capture. Although some high-level themes overlap (e.g., business goals), our findings extend beyond perception to endorsement, e.g., designers reframe DPs as "user-centered" practices. We also examine a distinct context, i.e., DPs generated by GenAI, and in that context, we reveal designers' awareness of DPs in AI-generated UIs, perceived harms, GenAI-specific mitigation strategies, and technical assessments of anti-/pro-DP strategies in prompts and across different GenAI tools.

GenAI used in UI Design and DPs. AI-powered tools have been applied at various stages of UI design and development, including creating UI mockups from high-level text descriptions [106], [107], [108], [109], translating UI mockups into working code [110], [111], [112], and evaluating UI designs to provide feedback for improvement [113], [114], etc. In recent years, a key factor accelerating the development of such tools has been generative AI, particularly large language models (LLMs), which have proven useful in most of these tasks [107], [108], [109], [111], [112], [113], [115]. The PRODEGENs analyzed in this study share similar goals with some of the tools discussed in previous studies [106], [108], [111], but they are primarily commercialized and non-transparent. In addition to assisting UI design, LLMs have also been found to generate DPs when creating UI designs (e.g., in e-commerce design tasks [18]). On the other hand, recent studies highlight that LLMs can be used to detect [116], address [117], and prevent [118] the creation of DPs in UI designs. These studies are related to our research as both involve LLMs. However, instead of general-purpose LLMs, we provide new insights into how UI designers interact with, perceive, and collaborate with PRODEGENs that are specifically built for UI design. This leads to a deeper understanding of how this emerging type of design tool may affect the landscape of DPs.

8. Conclusion

This study explores the impact of PRODEGENs on the generation of dark patterns (DPs) in UI design, particularly in privacy-related tasks, through designer-centric studies and technical analysis of PRODEGENs. We found that UI designers often lack awareness of DPs and struggle to identify them in PRODEGEN-generated UIs; many designers perceive DPs as beneficial, which could exacerbate their spread. Our findings also reveal that PRODEGENs lack built-in DP mitigations, and place the burden on designers to avoid DPs. We suggest that better designer training, built-in mitigations in PRODEGENs, and a reevaluation of the harm-benefit tradeoff of DPs are necessary to reduce their impact in real-world applications.

Acknowledgments

The authors would like to thank the anonymous reviewers and shepherd for their insightful and constructive feedback. The authors are supported in part by the UCF Seed Funding Program and NSF Grant No. 2520321.

Ethics considerations

This study uses several measures to ensure that our research practices adhere to ethical standards: **(1)** The designer-centric study (e.g., design tasks, interviews, and surveys) was reviewed and approved by the IRB of our institution. **(2)** We followed informed consent practices to ensure that all recruited participants are fully aware of the study's purpose, procedures, goals, and potential risks. We also clearly communicate that participation is voluntary, and participants may withdraw at any time. **(3)** All participant data are anonymized and securely stored on OneDrive, managed by our institution, and shared exclusively within the study group. **(4)** We interacted with publicly accessible PRODEGENs through manual analysis and official APIs (as normal UI designers would). Also, we used prompts created by real UI designers to test the PRODEGENs, which we believe poses no risk to the PRODEGENs themselves. These prompts were not designed to exploit, bypass, or circumvent any protections or mechanisms built into the PRODEGENs (e.g., authentication, authorization). **(5)** We recognize the importance of timely reporting and communication with involved stakeholders. Therefore, we have shared the prompts that result in DPs with all three PRODEGENs, along with our suggestions for implementing built-in DP mitigation in the PRODEGENs. We chose not to anonymize these PRODEGENs because, although UIs with DPs have privacy implications for end users, the generation of UIs with DPs does not constitute a vulnerability in the PRODEGENs that could allow attackers to benefit. Additionally, we consider that this study will cause low reputational harm to any specific PRODEGEN, as the lack of mitigation for DPs represents a systematic gap across all evaluated PRODEGENs. To date, we have not received concrete feedback from the PRODEGENs regarding their next steps for mitigating DPs. **(6)** We have also made the anonymized prompts, the prompt sentences causing (or reducing) DPs, and the interview questions used in the study publicly available on GitHub [19]. We remain open to any additional practices or suggestions that could further enhance the ethical integrity of this research.

LLM usage considerations

LLMs were used for editorial purposes in this manuscript, and all outputs were inspected by the authors to ensure accuracy and originality. The authors did not use LLMs to generate any ideas or content for this paper.

References

[1] H. Brignull, M. Leiser, C. Santos, and K. Doshi, "Deceptive patterns – user interfaces designed to trick you," 4 2010. [Online]. Available: https://www.deceptive.design/

[2] C. M. Gray, Y. Kou, B. Battles, J. Hoggatt, and A. L. Toombs, "The dark (patterns) side of ux design," in *Proceedings of the 2018 CHI conference on human factors in computing systems*, 2018, pp. 1–14.

[3] L. Zhang-Kennedy, M. Keleher, and M. Valiquette, "Navigating the gray: Design practitioners' perceptions toward the implementation of privacy dark patterns," *Proceedings of the ACM on human-computer interaction*, vol. 8, no. CSCW1, pp. 1–26, 2024.

[4] C. Bösch, B. Erb, F. Kargl, H. Kopp, and S. Pfattheicher, "Tales from the dark side: Privacy dark strategies and privacy dark patterns," *Proceedings on Privacy Enhancing Technologies*, 2016.

[5] A. Beattie, C. Lacey, and C. Caudwell, ""it's like the wild west": User experience (ux) designers on ethics and privacy in aotearoa new zealand," *Design and Culture*, vol. 16, no. 1, pp. 63–82, 2024.

[6] R. Mohamed, A. Arunasalam, H. Farrukh, J. Tong, A. Bianchi, and Z. B. Celik, "{ATTention} please! an investigation of the app tracking transparency permission," in *33rd USENIX Security Symposium (USENIX Security 24)*, 2024, pp. 5017–5034.

[7] Y. Chen, M. Zha, N. Zhang, D. Xu, Q. Zhao, X. Feng, K. Yuan, F. Suya, Y. Tian, K. Chen *et al.*, "Demystifying hidden privacy settings in mobile apps," in *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2019, pp. 570–586.

[8] S. Koch, B. Altpeter, and M. Johns, "The {OK} is not enough: A large scale study of consent dialogs in smartphone applications," in *32nd USENIX Security Symposium (USENIX Security 23)*, 2023, pp. 5467–5484.

[9] L. Di Geronimo, L. Braz, E. Fregnan, F. Palomba, and A. Bacchelli, "Ui dark patterns and where to find them: a study on mobile applications and user perception," in *Proceedings of the 2020 CHI conference on human factors in computing systems*, 2020, pp. 1–14.

[10] C. M. Gray, C. Santos, N. Bielova, M. Toth, and D. Clifford, "Dark patterns and the legal requirements of consent banners: An interaction criticism perspective," in *Proceedings of the 2021 CHI conference on human factors in computing systems*, 2021, pp. 1–18.

[11] J. Luguri and L. J. Strahilevitz, "Shining a light on dark patterns," *Journal of Legal Analysis*, vol. 13, no. 1, pp. 43–109, 2021.

[12] L. Nelissen and M. Funk, "Rationalizing dark patterns: Examining the process of designing privacy ux through speculative enactments," *International Journal of Design*, vol. 16, no. 1, pp. 75–92, 2022.

[13] J. Li, H. Cao, L. Lin, Y. Hou, R. Zhu, and A. El Ali, "User experience design professionals' perceptions of generative artificial intelligence," in *Proceedings of the 2024 CHI conference on human factors in computing systems*, 2024, pp. 1–18.

[14] W. AI, "What ai tools do designers use the most in 2024?" https://www.weavely.ai/blog/what-ai-tools-do-designers-use-the-most, 2024, accessed: 2025-09-30.

[15] Figma, "Figma make," https://www.figma.com/make/, 2024, accessed: 2025-11-01.

[16] U. Technologies, "Uizard – ui design made easy," https://uizard.io/, 2024, accessed: 2025-11-01.

[17] I. Vercel, "v0 by vercel," https://v0.app/, 2025, accessed: 2025-11-01.

[18] Z. Chen, J. Shen, K. Vaccaro *et al.*, "Hidden darkness in llm-generated designs: Exploring dark patterns in ecommerce web components generated by llms," *arXiv preprint arXiv:2502.13499*, 2025.

[19] SPIRIT-security, "PRODEGENs material," https://github.com/SPIRIT-security/PRODEGENs_material, 2026.

[20] H. Habib, M. Li, E. Young, and L. Cranor, "“okay, whatever”: An evaluation of cookie consent interfaces," in *Proceedings of the 2022 CHI conference on human factors in computing systems*, 2022, pp. 1–27.

[21] C. Utz, M. Degeling, S. Fahl, F. Schaub, and T. Holz, "(un) informed consent: Studying gdpr consent notices in the field," in *Proceedings of the 2019 acm sigsac conference on computer and communications security*, 2019, pp. 973–990.

[22] D. Bollinger, K. Kubicek, C. Cotrini, and D. Basin, "Automating cookie consent and {GDPR} violation detection," in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 2893–2910.

[23] C. Matte, N. Bielova, and C. Santos, "Do cookie banners respect my choice?: Measuring legal compliance of banners from iab europe's transparency and consent framework," in *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2020, pp. 791–809.

[24] M. Toth, N. Bielova, and V. Roca, "On dark patterns and manipulation of website publishers by cmps," *Proceedings on Privacy Enhancing Technologies*, vol. 2022, no. 3, pp. 478–497, 2022.

[25] D. Machuletz and R. Böhme, "Multiple purposes, multiple problems: A user study of consent dialogs after gdpr," *Proceedings on Privacy Enhancing Technologies*, vol. 2, pp. 481–498, 2020.

[26] V. H. Tran, A. Mehrotra, R. Sharma, M. Chetty, N. Feamster, J. Frankenreiter, and L. Strahilevitz, "Dark patterns in the opt-out process and compliance with the california consumer privacy act (ccpa)," in *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 2025, pp. 1–25.

[27] B. Schaffner, N. A. Lingareddy, and M. Chetty, "Understanding account deletion and relevant dark patterns on social media," *Proceedings of the ACM on Human-Computer Interaction*, vol. 6, no. CSCW2, pp. 1–43, 2022.

[28] J. Hannah, "A complete guide to the ui design process," https://www.uxdesigninstitute.com/blog/guide-to-the-ui-design-process/, 2023.

[29] M. Myre, "The ux design process: The ultimate 8-step guide," https://designlab.com/blog/what-is-the-ux-design-process, 2023.

[30] I. Miro, "Miro — ai innovation workspace," https://miro.com/, 2025, accessed: 2025-11-01.

[31] I. Figma, "Figma — the collaborative interface design tool," https://www.figma.com/, 2025, accessed: 2025-11-01.

[32] Meta. Ui & ux designers. https://www.facebook.com/groups/461999237551673.

[33] Reddit. r/ui_design. https://www.reddit.com/r/UI_Design/.

[34] Qulatrics, "Free online survey maker tool - qualtrics," https://www.qualtrics.com/.

[35] C. Bermejo Fernandez, D. Chatzopoulos, D. Papadopoulos, and P. Hui, "This website uses nudging: Mturk workers' behaviour on cookie consent notices," *Proceedings of the ACM on human-computer interaction*, vol. 5, no. CSCW2, pp. 1–22, 2021.

[36] J. Chen, J. Sun, S. Feng, Z. Xing, Q. Lu, X. Xu, and C. Chen, "Unveiling the tricks: Automated detection of dark patterns in mobile applications," in *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, 2023, pp. 1–20.

[37] C. M. Gray, C. T. Santos, N. Bielova, and T. Mildner, "An ontology of dark patterns knowledge: Foundations, definitions, and a pathway for shared knowledge-building," in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 2024, pp. 1–22.

[38] H. Habib and L. F. Cranor, "Evaluating the usability of privacy choice mechanisms," in *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*, 2022, pp. 273–289.

[39] S. Hidaka, S. Kobuki, M. Watanabe, and K. Seaborn, "Linguistic dead-ends and alphabet soup: Finding dark patterns in japanese apps," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 2023, pp. 1–13.

[40] R. Khandelwal, A. Nayak, H. Harkous, and K. Fawaz, "Automated cookie notice analysis and enforcement," in *32nd USENIX Security Symposium (USENIX Security 23)*, 2023, pp. 1109–1126.

[41] Y. Lu, C. Zhang, Y. Yang, Y. Yao, and T. J.-J. Li, "From awareness to action: Exploring end-user empowerment interventions for dark patterns in ux," *Proceedings of the ACM on Human-Computer Interaction*, vol. 8, no. CSCW1, pp. 1–41, 2024.

[42] D. Machuletz and R. Böhme, "Multiple purposes, multiple problems: A user study of consent dialogs after gdpr," *arXiv preprint arXiv:1908.10048*, 2019.

[43] M. Nouwens, I. Liccardi, M. Veale, D. Karger, and L. Kagal, "Dark patterns after the gdpr: Scraping consent pop-ups and demonstrating their influence," in *Proceedings of the 2020 CHI conference on human factors in computing systems*, 2020, pp. 1–13.

[44] R. Schäfer, P. M. Preuschoff, R. Röpke, S. Sahabi, and J. Borchers, "Fighting malicious designs: towards visual countermeasures against dark patterns," in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 2024, pp. 1–13.

[45] L. Schöni, K. Kubicek, and V. Zimmermann, "Block cookies, not websites: Analysing mental models and usability of the privacy-preserving browser extension cookieblock," *Proceedings on Privacy Enhancing Technologies*, vol. 2024, no. 1, pp. 192–216, 2023.

[46] J. Ye, Y. Li, W. Zou, and X. Wang, "From awareness to action: The effects of experiential learning on educating users about dark patterns," in *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 2025, pp. 1–22.

[47] V. Zimmermann, A. Toth, H. Sievers, L. Fanconi, Y. Isenring, M. Henz, A. Stöver, and N. Gerber, "Let's get visual-testing visual analogies and metaphors for conveying privacy policies and data handling information," in *2025 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2025, pp. 2303–2321.

[48] S. Hsu, V. Koshy, K. Vaccaro, C. Sandvig, and K. Karahalios, "Placebo effect of control settings in feeds are not always strong," in *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 2025, pp. 1–16.

[49] M. Kowalczyk, J. T. Gunawan, D. Choffnes, D. J. Dubois, W. Hartzog, and C. Wilson, "Understanding dark patterns in home iot devices," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 2023, pp. 1–27.

[50] S. H. Mansur, S. Salma, D. Awofisayo, and K. Moran, "Aidui: Toward automated recognition of dark patterns in user interfaces," in *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, 2023, pp. 1958–1970.

[51] T. Mildner, G.-L. Savino, P. R. Doyle, B. R. Cowan, and R. Malaka, "About engaging and governing strategies: A thematic analysis of dark patterns in social networking services," in *Proceedings of the 2023 CHI conference on human factors in computing systems*, 2023, pp. 1–15.

[52] M. Tahaei, A. Frik, and K. Vaniea, "Deciding on personalized ads: Nudging developers about user privacy," in *Seventeenth Symposium on Usable Privacy and Security (SOUPS 2021)*, 2021, pp. 573–596.

[53] B. Zhang and H. Xu, "Privacy nudges for mobile applications: Effects on the creepiness emotion and privacy attitudes," in *Proceedings of the 19th ACM conference on computer-supported cooperative work & social computing*, 2016, pp. 1676–1690.

[54] C. M. Gray, T. Mildner, and R. Gairola, "Getting trapped in amazon's" iliad flow": A foundation for the temporal analysis of dark patterns," in *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 2025, pp. 1–10.

[55] J. Gunawan, A. Pradeep, D. Choffnes, W. Hartzog, and C. Wilson, "A comparative study of dark patterns across web and mobile modalities," *Proceedings of the ACM on Human-Computer Interaction*, vol. 5, no. CSCW2, pp. 1–29, 2021.

[56] K. Hausladen, O. Wang, S. Eng, J. Wang, F. Wijaya, M. May, and S. Zimmeck, "Websites' global privacy control compliance at scale and over time," in *34th USENIX Security Symposium (USENIX Security 25)*, 2025, pp. 5837–5856.

[57] T. Mildner, D. Fidel, E. Stefanidi, P. W. Woźniak, R. Malaka, and J. Niess, "A comparative study of how people with and without adhd recognise and avoid dark patterns on social media," in *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 2025, pp. 1–17.

[58] A. Monge Roffarello, K. Lukoff, and L. De Russis, "Defining and identifying attention capture deceptive designs in digital interfaces," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 2023, pp. 1–19.

[59] A. Pradeep, J. Gunawan, Á. Feal, W. Hartzog, and D. Choffnes, "Gig work at what cost? exploring privacy risks of gig work platform participation in the us," *Proceedings on Privacy Enhancing Technologies*, 2025.

[60] K. M. Ramokapane, J. Such, and A. Rashid, "What users want from cloud deletion and the information they need: A participatory action study," *ACM Transactions on Privacy and Security*, vol. 26, no. 1, pp. 1–34, 2022.

[61] D. Smullen, Y. Yao, Y. Feng, N. Sadeh, A. Edelstein, and R. Weiss, "Managing potentially intrusive practices in the browser: A user-centered perspective," *Proceedings on Privacy Enhancing Technologies*, 2021.

[62] A. Sheil, G. Acar, H. Schraffenberger, R. Gellert, and D. Malone, "Staying at the roach motel: Cross-country analysis of manipulative subscription and cancellation flows," in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 2024, pp. 1–24.

[63] L. Alberts, U. Lyngs, and M. Van Kleek, "Computers as bad social actors: Dark patterns and anti-patterns in interfaces that act socially," *Proceedings of the ACM on Human-Computer Interaction*, vol. 8, no. CSCW1, pp. 1–25, 2024.

[64] E. Caragay, K. Xiong, J. Zong, and D. Jackson, "Beyond dark patterns: A concept-based framework for ethical software design," in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 2024, pp. 1–16.

[65] J. Ceci, H. Khan, U. Hengartner, and D. Vogel, "Concerned but ineffective: User perceptions, methods, and challenges when sanitizing old devices for disposal," in *Seventeenth Symposium on Usable Privacy and Security (SOUPS 2021)*, 2021, pp. 455–474.

[66] C. M. Gray, J. Chen, S. S. Chivukula, and L. Qu, "End user accounts of dark patterns as felt manipulation," *Proceedings of the ACM on Human-Computer Interaction*, vol. 5, no. CSCW2, pp. 1–25, 2021.

[67] A. Lewis, J. J. Martinez, M. Das, and J. Fogarty, "Inaccessible and deceptive: Examining experiences of deceptive design with people who use visual accessibility technology," in *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 2025, pp. 1–18.

[68] A. Mathur, G. Acar, M. J. Friedman, E. Lucherini, J. Mayer, M. Chetty, and A. Narayanan, "Dark patterns at scale: Findings from a crawl of 11k shopping websites," *Proceedings of the ACM on human-computer interaction*, vol. 3, no. CSCW, pp. 1–32, 2019.

[69] A. Mathur, M. Kshirsagar, and J. Mayer, "What makes a dark pattern... dark? design attributes, normative considerations, and measurement methods," in *Proceedings of the 2021 CHI conference on human factors in computing systems*, 2021, pp. 1–18.

[70] L. Nie, Y. Zhao, C. Li, X. Luo, and Y. Liu, "Shadows in the interface: A comprehensive study on dark patterns," *Proceedings of the ACM on Software Engineering*, vol. 1, no. FSE, pp. 204–225, 2024.

[71] C. Yue, C. Zhong, K. Chen, Z. Zhang, and Y. Lee, "{DARKFLEECE}: Probing the dark side of android subscription apps," in *33rd USENIX Security Symposium (USENIX Security 24)*, 2024, pp. 1543–1560.

[72] Z. Zhang, M. Jia, H.-P. Lee, B. Yao, S. Das, A. Lerner, D. Wang, and T. Li, ""it's a fair game", or is it? examining how users navigate disclosure risks and benefits when using llm-based conversational agents," in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 2024, pp. 1–26.

[73] P. Bahirat, M. Willemsen, Y. He, Q. Sun, and B. Knijnenburg, "Overlooking context: How do defaults and framing reduce deliberation in smart home privacy decision-making?" in *Proceedings of the 2021 chi conference on human factors in computing systems*, 2021, pp. 1–18.

[74] G. Chalhoub, I. Flechais, N. Nthala, and R. Abu-Salma, "Innovation inaction or in action? the role of user experience in the security and privacy design of smart home cameras," in *Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*, 2020, pp. 185–204.

[75] I. Chordia, L.-P. Tran, T. J. Tayebi, E. Parrish, S. Erete, J. Yip, and A. Hiniker, "Deceptive design patterns in safety technologies: A case study of the citizen app," in *Proceedings of the 2023 CHI conference on human factors in computing systems*, 2023, pp. 1–18.

[76] L. Kyi, S. Ammanaghatta Shivakumar, C. T. Santos, F. Roesner, F. Zufall, and A. J. Biega, "Investigating deceptive design in gdpr's legitimate interest," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 2023, pp. 1–16.

[77] T. Mildner, O. Cooney, A.-M. Meck, M. Bartl, G.-L. Savino, P. R. Doyle, D. Garaialde, L. Clark, J. Sloan, N. Wenig *et al.*, "Listening to the voices: Describing ethical caveats of conversational user interfaces according to experts and frequent users," in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 2024, pp. 1–18.

[78] S. Samat and A. Acquisti, "Format vs. content: the impact of risk and presentation on disclosure decisions," in *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*, 2017, pp. 377–384.

[79] T. Hardwick, M. Carter, S. Harkin, T. Zhangshao, and B. Egliston, "" they're scamming me": How children experience and conceptualize harm in game monetization," in *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 2025, pp. 1–14.

[80] R. A. Lee, "Figma statistics 2025: How users, revenue & ai are rising," https://sqmagazine.co.uk/figma-statistics/, 2025.

[81] V. Braun and V. Clarke, "Using thematic analysis in psychology," *Qualitative research in psychology*, vol. 3, no. 2, pp. 77–101, 2006.

[82] D. Kelly and V. L. Rubin, "Identifying dark patterns in user account disabling interfaces: Content analysis results," *Social Media+ Society*, vol. 10, no. 1, p. 20563051231224269, 2024.

[83] J. Nielsen. (2024, Jan.) 10 usability heuristics for user interface design. Nielsen Norman Group. Accessed: 2025-11-10. [Online]. Available: https://www.nngroup.com/articles/ten-usability-heuristics/

[84] Lovable, "Build something lovable," https://lovable.dev/, 2025, accessed: 2025-11-01.

[85] A. Gupta, "Tutorial of top 5 ai prototyping tools: Bolt, lovable, v0, replit, and cursor," https://www.news.aakashg.com/p/ai-prototyping-for-pms/, 2025.

[86] V. Solomakha. (2025) Best ai prototyping tools in 2025, ultimate review. https://www.banani.co/blog/best-ai-prototyping-tools. Accessed: 2025-11-01.

[87] A. Shrikumar, P. Greenside, and A. Kundaje, "Learning important features through propagating activation differences," in *International conference on machine learning*. PMlR, 2017, pp. 3145–3153.

[88] K. Simonyan, A. Vedaldi, and A. Zisserman, "Deep inside convolutional networks: Visualising image classification models and saliency maps," *arXiv preprint arXiv:1312.6034*, 2013.

[89] M. Sundararajan, A. Taly, and Q. Yan, "Axiomatic attribution for deep networks," in *International conference on machine learning*. PMLR, 2017, pp. 3319–3328.

[90] R. D. Cook and S. Weisberg, "Characterizations of an empirical influence function for detecting influential cases in regression," *Technometrics*, vol. 22, no. 4, pp. 495–508, 1980.

[91] B. Cohen-Wang, H. Shah, K. Georgiev, and A. Madry, "Contextcite: Attributing model generation to context," *Advances in Neural Information Processing Systems*, vol. 37, pp. 95 764–95 807, 2024.

[92] J. Enouen, H. Nakhost, S. Ebrahimi, S. Arik, Y. Liu, and T. Pfister, "Textgenshap: Scalable post-hoc explanations in text generation with long documents," in *Findings of the Association for Computational Linguistics: ACL 2024*, 2024, pp. 13 984–14 011.

[93] V. Petsiuk, A. Das, and K. Saenko, "Rise: Randomized input sampling for explanation of black-box models," *arXiv preprint arXiv:1806.07421*, 2018.

[94] M. T. Ribeiro, S. Singh, and C. Guestrin, "" why should i trust you?" explaining the predictions of any classifier," in *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 2016, pp. 1135–1144.

[95] Y. Wang, T. Zhang, X. Guo, and Z. Shen, "Gradient based feature attribution in explainable ai: A technical review," *arXiv preprint arXiv:2403.10415*, 2024.

[96] Spacy. Industrial-strength natural language processing. https://spacy.io/.

[97] K. Krippendorff, "Computing krippendorff's alpha-reliability," 2011.

[98] UXP2, "Pitcher & piano: Resubscribe button," https://darkpatterns.uxp2.com/pattern/22/, 2025, accessed: 2025-11-10.

[99] E. R. Girden, *ANOVA: Repeated measures*. sage, 1992, no. 84.

[100] Tailwind Labs Inc., "Tailwind css," 2023, accessed: 2025-11-01. [Online]. Available: https://tailwindcss.com/

[101] Federal Trade Commission, "Bringing dark patterns to light: Staff report," https://www.ftc.gov/system/files/ftc_gov/pdf/P214800%20Dark%20Patterns%20Report%209.14.2022%20-%20FINAL.pdf, 2022, accessed: 2024-08-26.

[102] F. R. Board, "Consumer compliance handbook: Section 5 of the ftc act," https://www.federalreserve.gov/boarddocs/supmanual/cch/200806/ftca.pdf, 2008, accessed: 2024-08-26.

[103] Federal Trade Commission, "Bringing dark patterns to light: Ftc workshop," https://www.ftc.gov/news-events/events/2021/04/bringing-dark-patterns-light-ftc-workshop, 2021, accessed: 2024-08-26.

[104] Z. Cai, Y. Nan, X. Wang, M. Long, Q. Ou, M. Yang, and Z. Zheng, "Darpa: combating asymmetric dark ui patterns on android with run-time view decorator," in *2023 53rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 2023, pp. 480–493.

[105] G. Conti and E. Sobiesk, "Malicious interface design: exploiting the user," in *Proceedings of the 19th international conference on World wide web*, 2010, pp. 271–280.

[106] F. Huang, G. Li, X. Zhou, J. F. Canny, and Y. Li, "Creating user interface mock-ups from high-level text descriptions with deep-learning models," *arXiv preprint arXiv:2110.07775*, 2021.

[107] J. Ma, K. Sreedhar, V. Liu, S. Wang, P. A. Perez, and L. B. Chilton, "Didup: Dynamic iterative development for ui prototyping," *arXiv preprint arXiv:2407.08474*, 2024.

[108] M. Yuan, J. Chen, Y. Hu, S. Feng, M. Xie, G. Mohammadi, Z. Xing, and A. J. Quigley, "Towards human-ai synergy in ui design: Supporting iterative generation with llms," *ACM Transactions on Computer-Human Interaction*, 2025.

[109] W. Feng, W. Zhu, T.-j. Fu, V. Jampani, A. Akula, X. He, S. Basu, X. E. Wang, and W. Y. Wang, "Layoutgpt: Compositional visual planning and generation with large language models," *Advances in Neural Information Processing Systems*, vol. 36, pp. 18 225–18 250, 2023.

[110] M. Bajammal, D. Mazinanian, and A. Mesbah, "Generating reusable web components from mockups," in *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*, 2018, pp. 601–611.

[111] T. Calò and L. De Russis, "Leveraging large language models for end-user website generation," in *International Symposium on End User Development*. Springer, 2023, pp. 52–61.

[112] Y. Wan, C. Wang, Y. Dong, W. Wang, S. Li, Y. Huo, and M. Lyu, "Divide-and-conquer: Generating ui code from screenshots," *Proceedings of the ACM on Software Engineering*, vol. 2, no. FSE, pp. 2099–2122, 2025.

[113] P. Duan, J. Warner, Y. Li, and B. Hartmann, "Generating automatic feedback on ui mockups with large language models," in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 2024, pp. 1–20.

[114] J. Wu, Y.-H. Peng, X. Y. A. Li, A. Swearngin, J. P. Bigham, and J. Nichols, "Uiclip: a data-driven model for assessing user interface design," in *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*, 2024, pp. 1–16.

[115] A. Ahmed and A. S. Imran, "The role of large language models in ui/ux design: A systematic literature review," *arXiv preprint arXiv:2507.04469*, 2025.

[116] S. Mills and R. Whittle, "Detecting dark patterns using generative ai: Some preliminary results," *Available at SSRN 4614907*, 2023.

[117] R. Schäfer, P. M. Preuschoff, R. Niewianda, S. Hahn, K. Fiedler, and J. Borchers, "Don't detect, just correct: Can llms defuse deceptive patterns directly?" in *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, 2025, pp. 1–11.

[118] O. Quijada and T. Serezlic, "Can ai avoid deceptive design?: Investigating the effectiveness of prompting ai to prevent dark patterns in web design," 2025.

Appendix A. Additional Tables and Figures

```
Design a cookie banner that appears when users first
    visit the site. Include a short message, a link to
    learn more, and clear "Accept" and "Decline" buttons
    . Keep it clean and easy to read.
```

Figure 8: An example of prompt used by P6.

TABLE 4: Participants' Demographic Information

No.	Company Scale	Age	Years of UI/UX Design	UI Platform	Design Domain
P1	Small Team	18–25	3–5	Web	Small Business (e.g., restaurant)
P2	Small Team	26–35	3–5	Web	Branding/Logo
P3	Small Team	18–25	3–5	Web	†
P4	Small Team	26–35	5+	Mobile	Digital Product
P5	Large Company	46–60	5+	Web	†
P6	Small Team	18–25	3–5	Web/Mobile	†
P7	Freelancer	18–25	3–5	Mobile	Game
P8	Freelancer	26–35	3–5	Web	Graphic/Visual
P9	Large Company	18–25	1–2	Web	†
P10	Large Company	26–35	5+	Web	Enterprise Software
P11	Large Company	26–35	3–5	In-Vehicle HMI	Automotive HMI
P12	Small Team	18–25	1–2	Web	†
P13	Large Company	18–25	1–2	Web	Graphic/Visual
P14	Large Company	18–25	5+	Web	†
P15	Small Team	26–35	3–5	Web	Delivery Product
P16	Freelancer	26–35	5+	Web/Mobile	Graphic/Visual
P17	Small Team	36–45	5+	Web	Entertainment/ Media/ Gallery

† Prefer not to disclose.

```
1. Figma Make Features
    Auto Layout: Structure text, radio buttons, and
         confirmation buttons neatly.
    Components & Variants: Use for radio buttons or
         checkboxes (selected/unselected).
    Prototyping: Simulate the unsubscribe confirmation
         flow.
    Text & Color Styles: Maintain clear hierarchy and
         branding consistency.
2. Prompts to Use
    "Design a clean unsubscribe page for a website with
         an option to stop receiving marketing emails."
    "Include a short thank-you or feedback message, and a
          visible 'Unsubscribe' button."
    "Use a simple layout with friendly wording and brand
         colors."
3. User Interaction Flow
    User clicks the "Unsubscribe" link in an email.
    Page opens, showing their email address and message:
    "We're sorry to see you go. Would you like to
         unsubscribe from all newsletters?"
    User selects: "Yes, unsubscribe me."
    User clicks: "Confirm Unsubscribe."
    Success message appears: "You have been unsubscribed.
          You'll no longer receive marketing emails."
    A "Resubscribe" link is displayed.
```

Figure 9: An example of prompt template created by P2.

Appendix B. Meta-Review

The following meta-review was prepared by the program committee for the 2026 IEEE Symposium on Security and Privacy (S&P) as part of the review process as detailed in the call for papers.

B.1. Summary

This paper studies privacy dark patterns in prompt-to-design generators (PRODEGENs) through a designer-centered user study and a follow-up analysis of prompts and outputs across multiple tools. The paper provides empirical insight into how designers identify, justify, and attempt to mitigate dark patterns in AI-assisted UI design workflows.

B.2. Scientific Contributions

- Provides a New Data Set For Public Use
- Provides a Valuable Step Forward in an Established Field

B.3. Reasons for Acceptance

1) This paper provides a valuable step forward in an established field. It addresses a timely security and privacy problem at the intersection of dark patterns and GenAI-assisted UI design, and extends prior work on manipulative interface design to AI-assisted design tools.
2) The paper combines a designer-centered user study with a follow-up technical analysis on multiple PRODEGEN tools. This gives the paper useful empirical scope and provides new evidence about designers' awareness of dark patterns, the justifications they offer, and the prompting strategies they use in response.
3) The paper provides a new data set for public use. The collected prompts, outputs, and qualitative findings have the potential to support future work on privacy dark patterns, AI-assisted design, and the security and privacy implications of generative design systems.

B.4. Noteworthy Concerns

1) The paper does not include a baseline comparison to traditional UI design workflows. Therefore, it does not establish how much GenAI changes the prevalence of dark patterns.