AdsDP: A Video Dataset for Recognizing and Examining Dark Patterns in iOS In-App Advertisements

YUXUAN SHANG, The State Key Laboratory of Blockchain and Data Security, Zhejiang University, China
GUANXIAO WANG, The State Key Laboratory of Blockchain and Data Security, Zhejiang University, China
MENGXIA REN, The State Key Laboratory of Blockchain and Data Security, Zhejiang University, China
HAOMIN ZHANG, The State Key Laboratory of Blockchain and Data Security, Zhejiang University, China
XINGMING CHEN, The State Key Laboratory of Blockchain and Data Security, Zhejiang University, China
HAITAO XU*, The State Key Laboratory of Blockchain and Data Security, Zhejiang University, China
CHUAN YUE, Department of Computer Science, Colorado School of Mines, US
SHUAI HAO, Department of Computer Science, Old Dominion University, US
BO ZHOU, The State Key Laboratory of Blockchain and Data Security, Zhejiang University, China
WENRUI MA, School of Computer Science and Technology, Zhejiang Gongshang University, China
FAN ZHANG, The State Key Laboratory of Blockchain and Data Security, Zhejiang University, China
ZHAO LI, Hangzhou Yugu Technology and Zhejiang University, China

Ads are widely deployed in mobile apps, significantly affecting user experience. In recent years, malicious or deceptive user interfaces (UI) designs, known as dark patterns, have increasingly been employed in in-app ads to manipulate users into unintended actions. While previous studies have identified a limited set of dark patterns in in-app ads and highlighted user concerns, a thorough understanding remains lacking, partly due to the absence of publicly available, contextual datasets on dark patterns in in-app ads. In this study, we systematically investigate dark patterns in iOS in-app ads, identifying 15 types, 11 of which were previously unreported. We also introduce AdsDP, an annotated video dataset documenting in-app ads appearing during normal usage of iOS apps, along with any dark patterns these ads may exhibit. AdsDP includes 718 videos totaling 60 hours and features 5,782 instances of ad-related dark patterns across 485 apps. Furthermore, we evaluate the

*Haitao Xu is the corresponding author.

Authors' Contact Information: Yuxuan Shang, shangyuxuan@zju.edu.cn, The State Key Laboratory of Blockchain and Data Security, Zhejiang University, Hangzhou, Zhejiang, China; Guanxiao Wang, 22421327@zju.edu.cn, The State Key Laboratory of Blockchain and Data Security, Zhejiang University, Hangzhou, Zhejiang, China; Mengxia Ren, mengxiaren@hotmail.com, The State Key Laboratory of Blockchain and Data Security, Zhejiang University, Hangzhou, Zhejiang, China; Haomin Zhang, 3230102472@zju.edu.cn, The State Key Laboratory of Blockchain and Data Security, Zhejiang University, Hangzhou, Zhejiang, China; Xingming Chen, 3230102275@zju.edu.cn, The State Key Laboratory of Blockchain and Data Security, Zhejiang University, Hangzhou, Zhejiang, China; Haitao Xu, haitaoxu@zju.edu.cn, The State Key Laboratory of Blockchain and Data Security, Zhejiang University, Hangzhou, Zhejiang, China; Chuan Yue, chuanyue@mines.edu, Department of Computer Science, Colorado School of Mines, Golden, Colorado, US; Shuai Hao, shao@odu.edu, Department of Computer Science, Old Dominion University, Norfolk, Virginia, US; Bo Zhou, zbxjdh@foxmail.com, The State Key Laboratory of Blockchain and Data Security, Zhejiang University, Hangzhou, Zhejiang, China; Wenrui Ma, mawenrui@zjgsu.edu.cn, School of Computer Science and Technology, Zhejiang Gongshang University, Hangzhou, Zhejiang, China; Fan Zhang, fanzhang@zju.edu.cn, The State Key Laboratory of Blockchain and Data Security, Zhejiang University, Hangzhou, Zhejiang, China; Zhao Li, lzjoey@gmail.com, Hangzhou Yugu Technology, Hangzhou, Zhejiang and Zhejiang University, Hangzhou, Zhejiang, China.

ACM 2474-9567/2025/9-ART127
https://doi.org/10.1145/3749515

performance of state-of-the-art dark pattern detection solutions using AdsDP, revealing a significant decline in performance on our novel dataset, underscoring the need for new detection methods. Finally, we demonstrate AdsDP's potential to enhance future detection efforts and increase user awareness of ad-related dark patterns. The AdsDP dataset is available at https://zenodo.org/records/15316373.

CCS Concepts: • **Human-centered computing → Empirical studies in ubiquitous and mobile computing**.

Additional Key Words and Phrases: Dark Pattern, Mobile Advertising, Video Dataset, Dark Pattern Detection

ACM Reference Format:
Yuxuan Shang, Guanxiao Wang, Mengxia Ren, Haomin Zhang, Xingming Chen, Haitao Xu, Chuan Yue, Shuai Hao, Bo Zhou, Wenrui Ma, Fan Zhang, and Zhao Li. 2025. AdsDP: A Video Dataset for Recognizing and Examining Dark Patterns in iOS In-App Advertisements. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 9, 3, Article 127 (September 2025), 31 pages. https://doi.org/10.1145/3749515

1 Introduction

Mobile ads offer substantial benefits to app developers, including increased revenue and enhanced product visibility. In pursuit of maximizing profits, some unethical developers have increasingly resorted to dark patterns to display as many in-app ads as possible and manipulate users into greater ad engagement through deceptive tactics, which lead to users' frustration and annoyance [42]. Dark patterns refer to design strategies that deceive or manipulate users into making choices they would not have otherwise made, such as disclosing more personal information than intended or subscribing to unwanted services [16].

Prior research [13, 20, 31, 40] has investigated general dark patterns in domains like e-commerce and mobile applications, yet a comprehensive study specifically targeting in-app advertising dark patterns (hereafter *ads dark patterns*) remains absent. Current dark pattern datasets [15, 20, 39, 40] exhibit notable limitations: they either document patterns through isolated screenshots [15, 39, 40] or impose access restrictions [20], while being exclusively limited to Android platforms. Although detection tools have emerged [15, 39], their reliance on individual screenshots prevents identification of dynamic patterns unfolding across multiple interfaces. This fundamental limitation originates from the lack of publicly available, contextual datasets that systematically capture the full progression of dark patterns in mobile in-app ads.

Constructing a contextual dataset of dark patterns in mobile in-app ads presents several challenges. (i) The process is both time-consuming and labor-intensive, encompassing extensive app usage, the identification of ad-related video clips, and the annotation of many instances of ads dark patterns. (ii) There may be divergent interpretations of what qualifies as a dark pattern in UI design, especially given the lack of consistent definitions of ads dark patterns (*e.g.*, "Disguised Ads" in [28] and [43]) in prior studies. (iii) Recording user interactions with an app, such as tapping or swiping, presents a nontrivial challenge, as these actions may not leave visible traces on the screen, unlike explicit mouse actions on a computer.

To address these challenges, we present AdsDP, an annotated video dataset of dark patterns within iOS in-app ads, generated through a multi-collaborator workflow. We begin by selecting 485 distinct apps from both Chinese and English language markets, each featuring extensive user reviews regarding in-app ads. From this selection, we randomly choose 100 apps and assign four researchers to independently explore them, identify ad-related UI designs, hold regular meetings to define newly identified dark patterns, resolve disagreements through majority voting, and ultimately produce a codebook with fine-grained dark pattern definitions, detailing user interactions and UI transitions. Concurrently, we compile a manual outlining common UI characteristics for each dark pattern, providing interaction guidelines for researchers.

Using the generated codebook and manual, the researchers thoroughly operate each of the 485 apps on an iPhone for at least five minutes, recording the entire app usage process on video. This ensures that potential ad occurrences, along with the user interactions before and after the ad's appearance, are clearly documented.

To track finger-based interactions (e.g., tapping or swiping), we connect a mouse to the iPhone, allowing us to monitor all user actions via the cursor's movement.

Finally, the four researchers annotate the dark patterns appearing in the recorded videos according to the codebook and manual. For instances involving unknown dark patterns, consensus is reached to determine whether the occurrence represents a new dark pattern. Additionally, we develop an online annotation-aiding tool to further streamline the annotation process. In total, the four researchers complete the entire workflow over the course of five months, producing AdsDP, which consists of 718 videos with a total duration of 60 hours and 5,782 instances of ads dark patterns collected from 485 apps.

Furthermore, we analyze the prevalence of ads dark patterns in AdsDP. On average, Chinese apps contain 4.7 ads dark patterns per app, compared to 3.5 in English apps. Our findings show distinct regional trends: *unexpected full-screen ads* and *shake-to-open ads* dominate in Chinese apps, while *ads with closure failure* and *auto-redirect ads* are more frequent in English apps, suggesting divergent ad design strategies across markets. Moreover, we validate these patterns across 60 popular Android and 60 additional iOS apps and assess state-of-the-art detection methods using AdsDP. Performance drops significantly, with micro F1 falling from 0.77 to 0.11, macro F1 from 0.82 to 0.05, and 94.2% (97/103) of sampled dark pattern instances undetected. Finally, by aligning user reviews from a larger iOS app set with our ads dark pattern definitions, we confirm their widespread use in app advertising and show our codebook's utility in identifying them.

In summary, this study makes the following contributions:

- We conduct the first systematic investigation of ads dark patterns across iOS apps in Chinese and English markets, identifying 15 distinct types - including 11 novel patterns not previously documented.
- We introduce AdsDP, a comprehensive annotated video dataset capturing complete app usage sessions with detailed interaction context.
- Our analysis reveals that ads dark patterns are prevalent in iOS apps, with significant differences in their distribution between distinct app markets, with user reviews of a broader iOS app set confirming these patterns' widespread adoption in production apps.
- We assess the performance of state-of-the-art dark pattern detection solutions using AdsDP, revealing a substantial decline in performance, with 94.2% of the sampled dark pattern instances undetected, demonstrating the AdsDP's value and the urgent need for improved detection approaches.

This study advances dark pattern research in iOS app monetization by introducing an enhanced taxonomy of ad-specific dark patterns and examining their revenue-driven behavioral exploitation, while providing practical resources for detecting manipulative UI designs and insights to improve user awareness of in-app ad practices.

2 Related Work

2.1 User Perception of Ads

Advertisements are frequently disliked by users [63], often perceived as detrimental to user experience and consumer welfare. This distrust partly stems from deceptive ads [19], which may contain false claims [7] or obscure the distinction between ads and organic content—a phenomenon termed "Disguised Ads" [20]. Even non-malicious ads can frustrate users due to intrusive timing and placement [10, 26].

To mitigate these issues, researchers have proposed improved ad placement strategies. Nguyen et al. [42] introduced a balanced approach considering advertisers, developers, and users, while Häglund et al. [32] developed an AI-driven system avoiding personal data tracking to alleviate privacy concerns. Advertisers also employ refined content strategies, such as material incentives [17], personalized ads [47], and interactive formats [2]. Other efforts focus on transparency: Silberstein et al. [54] suggested penalizing frequently dismissed ads, Gao et al. [23] proposed disclosing ad-related resource usage, and Lee et al. [35] found explanations helpful when ads feel intrusive. While these approaches reduce frustration and improve perception, most neglect UI design as a source of user dissatisfaction.

2.2 *Generic* Dark Patterns

UI design plays a crucial role in shaping user behavior in modern mobile apps. Through deliberate UI design, developers can obscure beneficial information, manipulate emotions, and even influence user decisions [41]. In 2010, Harry Brignull coined the term "dark pattern" to describe UI designs that prioritize developers' interests at the expense of users, potentially deceiving or harming them, and identified 12 types of such practices [12].

Since then, HCI researchers have explored and classified dark patterns across various platforms (e.g., web [31, 40] and mobile [20, 31]) and domains (e.g., e-commerce [40], games [1], and social media [27]). Gray *et al.* [28] expanded Brignull's original framework by adding new examples and categorizing dark patterns into five themes based on developer intentions: "Nagging," "Forced Action," "Sneaking" and etc. Di Geronimo *et al.* [20] analyzed 240 popular mobile apps, revealing that 95% employed dark patterns, with most users unaware of their existence. Mathur *et al.* [40] conducted a large-scale study across 11,000 e-commerce websites, identifying 1,818 instances of dark patterns. Gunawan *et al.* [31] compared dark patterns across mobile apps, mobile browsers, and web versions of 105 online services, finding differences in patterns across platforms. Building upon prior work, we empirically investigate UI instances to systematically identify and characterize dark patterns in iOS apps.

Several tools have been proposed to detect dark patterns [15, 39, 40]. Mathur *et al.* [40] used clustering to group UI screenshots from e-commerce sites and manually evaluated each group for dark patterns. Mansur *et al.* [39] introduced AidUI, a tool that identifies dark patterns by analyzing visual and textual cues from UI screenshots. Chen *et al.* [15] developed UIGuard, another detection tool, which compares UI screenshots against common patterns to detect dark patterns. However, as demonstrated in §5.3, these tools, relying on single-screenshot analysis, are ineffective at detecting ad-related dark patterns, which often span multiple frames.

2.3 *Ad-specific* Dark Patterns

While HCI research has explored various dark patterns, studies specifically targeting ad-related dark patterns remain limited. Existing work on mobile ads [24, 30] has identified such patterns through user reviews. Brignull *et al.* [13] introduced the dark pattern "Disguised Ads," which refers to ads designed to mimic regular content, making it challenging for users to recognize them. Di Geronimo *et al.* [20] expanded on this concept with additional examples and introduced three new ad-related dark patterns. Long *et al.* [38] conducted a case study on mobile apps in China, identifying five types of dark patterns in ad design, including fake close buttons that are non-functional. Nie *et al.* [43] reviewed existing research on dark patterns, categorizing 64 types, of which only five were ad-related. Gao *et al.* [24] analyzed user feedback on ads and highlighted negative phenomena such as excessive, overly long, or oversized ads. In this work, we focus on investigating dark patterns in iOS in-app ads.

2.4 Datasets Pertaining to Dark Patterns

To our knowledge, no prior research or dataset has specifically focused on ads-related dark patterns. In this context, we compare AdsDP with existing datasets dedicated to studying dark patterns in other domains. Table 1 presents four publicly available datasets related to dark patterns, each of which has several limitations.

Data Format: The three datasets—ContextDP [39], Rico' [15], and AGM [40]—record dark patterns using single screenshots, which lack sufficient contextual information and fail to capture dynamic interactions over time, a limitation noted by Nie *et al.* [43]. In contrast, AdsDP uses video recordings to capture the entire process of UI changes and user interactions with ads, providing rich contextual data.

Data Scale: Previous datasets annotated fewer than 2,000 instances of dark patterns, while AdsDP offers a much larger dataset with 5,782 instances. Specifically, compared to the LLE [20] dataset, which has a total video

Table 1. **List of Dark Pattern Datasets.** In cases where the datasets were not explicitly named by their authors, we utilize the initials of the first three authors as identifiers. The quantities are precisely calculated based on the descriptions provided in each respective paper. The term “# Instances” refers to the number of dark pattern instances documented in a dataset. “# Dark Pattern Types” and “# Ads Dark Pattern Types” indicate the number of general dark pattern types and ad-specific dark pattern types in a dataset, respectively.

	ContextDP [39]	Rico’ [15]	AGM [40]	LLE [20]	AdsDP (**Ours**)
Publishing Year	2023	2023	2019	2020	2025
Domain	E-commerce websites & Mobile Apps	Mobile Apps	E-commerce websites	Mobile Apps	Ads in Mobile Apps
Data Format	Screenshot	Screenshot	Screenshot	Video	Video
# Instances	301	1,660	1,818	1,787	5,782
# Dark Pattern Types	10	15	15	16	15
# Ads Dark Pattern Types	1	4	0	4	15
Data Availability	Publicly Accessible	Publicly Accessible	Publicly Accessible	Permission Required	Publicly Accessible

duration of 40 hours and 1,787 dark pattern instances, AdsDP provides 60 hours of video and a significantly higher number (5,782) of dark pattern instances.

Accessibility: The only video-based dataset, LLE, suffers from accessibility issues. While 15 videos, each lasting 10 minutes, are publicly available, access to the full dataset requires contacting the authors, limiting its usability and leading to inadequate analysis in [43]. In contrast, AdsDP has committed to making all 718 five-minute videos publicly available to the research community, ensuring greater accessibility for future studies.

Moreover, all three mobile dark pattern datasets—ContextDP [39], Rico’ [15], and LLE [20]—focus exclusively on Android apps, while our dataset, AdsDP, focuses on iOS apps, thus complementing the existing datasets. For all the reasons mentioned, AdsDP represents a significant advancement over current datasets and is better suited for research on ads-related dark patterns.

3 Motivation

Table 2. **Categories of Topics Addressed in User Reviews Concerning In-App Ads**

Topic Category	**Description**
Excessive Ads (36.9%)	Ads appear excessively during app usage, impeding functionality and reducing usability.
Ads Disrupting App Functionality (11.4%)	Ads disrupt usability or cause technical issues such as lagging, app crashing, or making certain app features unusable.
Incentivized Ad Viewing (7.5%)	Users are offered incentives or rewards for watching ads to earn in-game currency or features, which they find frustrating and manipulative, as it restricts progress without engaging with ads.
Deceptive Claims of Ad-Free Experience (6.6%)	Users felt tricked by ads, such as when an app was advertised as free of ads, but users found themselves inundated with ads after download.
Unexpected Ads in Paid Apps (4.7%)	Users encounter ads even after purchasing an ad-free version, leading to dissatisfaction.
Non-Skippable or Mandatory Ads (3.7%)	Users express their frustration with ads that cannot be skipped or are mandatory to watch for progression.
Inappropriate Ad Content (2.1%)	Ads contain offensive or adult content deemed unsuitable for the app’s audience, or ads are irrelevant and do not reflect the app’s functionality or features.
Repetitive Ads (1.3%)	Users reported annoyance overseeing the same ads repeatedly in short periods.
Excessive Ad Duration (0.7%)	Users perceive certain ads as excessively lengthy.
Total (81.1%)	n/a

Dark patterns have been employed by sophisticated UI designers [16], which has raised significant concerns among app users, as dark patterns can severely degrade the user experience. This study was prompted by the prevalence of negative comments regarding in-app ads in the user review sections of various apps.

To thoroughly investigate these user concerns, we first identified 231,421 ad-related comments across 24,415 applications through keyword matching (e.g., "ad(s)," "advertise," "advertisement(s)"). From this collection, we randomly selected 5,000 comments spanning 4,351 apps for detailed analysis. Using GPT-4, we extracted ad-relevant segments from these comments and generated an initial topic taxonomy, which was subsequently refined by us to produce the final classification scheme presented in Table 2. Each comment was then individually categorized by GPT-4 according to this taxonomy. To assess the accuracy of GPT-4's classifications, we performed manual annotation on 200 randomly selected comments, achieving substantial agreement with the model's classifications (Cohen's Kappa coefficient $k = 0.79$). The complete set of prompts employed for topic modeling and classification tasks is provided in Figures 6 and 7 (Appendix A).

Table 2 categorizes user comments regarding in-app ads, highlighting various issues that negatively impact user experience. The most prevalent concern, accounting for 36.9% of feedback, is the *excessive frequency of ads*, which hinders app functionality and usability. Collectively, these issues represent 81.1% of user complaints, highlighting widespread dissatisfaction with in-app advertising practices. *These negative sentiments are partly attributed to the use of dark patterns, which prompted us to systematically investigate the ads dark patterns that lead to user dissatisfaction.*

4 Dataset Acquisition and Description

Figure 1 illustrates the workflow for dataset acquisition aimed at investigating the prevalence of dark patterns in in-app ads. The workflow encompasses four sequential steps: app selection, codebook generation, operation and video recording, and dark pattern annotation. During the app selection phase, we established specific criteria and curated a list of apps from the iOS App Store (§4.1). In the codebook generation step (§4.2), four researchers downloaded and interacted with a subset of these apps (§4.2.1) to analyze possible dark patterns employed in their ads. They subsequently compiled a codebook (detailed in §4.2.2) to categorize the identified dark patterns. In the operation and video recording phase (§4.3), the researchers downloaded and installed all selected apps from §4.1, operating each app for five minutes to simulate typical user behavior and explore its functionalities. The iPhone's built-in screen-recording feature was utilized to capture the entire interaction as a video. In the dark pattern annotation step (§4.4), the researchers visually analyzed the recorded videos to identify and annotate ads dark patterns, guided by the generated codebook. Through this systematic process, we developed a comprehensive codebook that defines and describes various ads dark patterns, along with the AdsDP dataset, which includes screen recording videos and their corresponding annotations.

4.1 App Selection

To compile an app list for analyzing dark patterns and producing the AdsDP dataset, we implemented the following selection criteria: (1) apps that have garnered significant criticism from users regarding their advertising practices and maintain popularity within the user base; (2) apps that demonstrate content diversity and span various categories. We concentrated on apps from the U.S. and China App Stores, as these regions represent the two largest mobile app markets globally [49].

The app list was derived from 853,600 iOS apps launched worldwide since 2023, through a systematic process: (1) Within each category, we selected the top 20 apps with the most ad-related user reviews using the keyword matching method described in §3. (2) From the remaining apps, we included popular ones meeting minimum rating thresholds (10,000 for English apps and 1,000 for Chinese apps) to ensure a balanced representation of apps from the U.S. and China App Stores. Note that the quantity of ad-related user reviews serves as a reliable

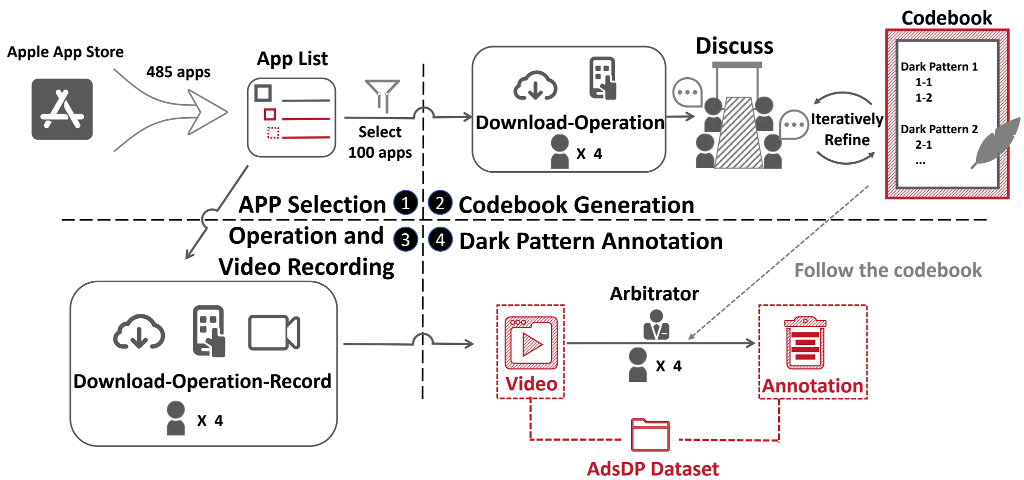

Fig. 1. Workflow of Dataset Acquisition

indicator of the prevalence of dark patterns, as evidenced by the observation that over 80% of such reviews are negative, as discussed in §3. Furthermore, prior research has demonstrated the effectiveness of inferring an app's advertising behavior from user reviews [6, 50]. This process ultimately resulted in a final selection of 485 apps, consisting of 245 apps from the U.S. App Store (referred to as the English app dataset) and 240 apps from the China App Store (referred to as the Chinese app dataset).

4.2 Codebook Generation

4.2.1 Generation of Codebook and Manual. We selected 100 apps from the curated app list to conduct a qualitative analysis using an iterative inductive thematic analysis approach, as outlined by Braun et al. [11], while adhering to the trustworthiness criteria established by Nowell et al. [45]. During this thematic analysis, four researchers were assigned to independently explore the functionalities of the apps, identify ad-related user interface designs, and hold regular meetings to define newly identified dark patterns. These patterns were supported by screen recordings captured during their interactions as evidence. Disagreements were resolved using the "arguing to consensus" approach described by Johnstone et al. [34], culminating in the creation of a codebook documenting the definitions of the dark patterns observed in in-app ads. Concurrently, the researchers collaboratively developed a manual providing guidelines for the actions to be taken and the information to be recorded in various scenarios for each dark pattern outlined in the codebook.

The codebook and manual underwent continuous refinement as researchers encountered cases that the initial guidelines could not adequately address. This iterative refinement process persisted until a saturation criterion [22] was achieved, at which point no further questions or revision requests were proposed by the researchers. Throughout this process, we consistently alternated between the stages of app usage, meetings, revisions, and documentation, as recommended by Creswell [18].

During the refinement of the codebook and manual, we encountered several challenges (e.g., discrepancies in how researchers described dark patterns) and established principles to address them:

- **Challenge 1**: Subjective perceptions among researchers led to disagreements over certain design patterns.
  Exemplar 1: One researcher argued that displaying a full-screen ad upon opening the app significantly degrades the user experience, while others deemed it acceptable.

> **Principle 1**: Researchers discussed each design until reaching a consensus; if consensus was unattainable, decisions were made based on majority rule. In Exemplar 1, after deliberation, they agreed that the design of "Ads upon opening the app" does not constitute a dark pattern.

- **Challenge 2**: Researchers, influenced by differing perspectives and interpretations, provided varying descriptions of the same dark pattern.
  Exemplar 2: An ad automatically appears after a game round ends. One researcher focused on its consequences, proposing the dark pattern "Ad hindering players from continuing to the next round," while another emphasized the ad's appearance, suggesting "Automatically popping up without user consent."

> **Principle 2**: Researchers should primarily describe dark patterns based on observable phenomena at the time they occur. Thus, in exemplar 2, we favored the description "Automatically popping up without user consent" over the alternative.

- **Challenge 3**: In real-world scenarios, UI design patterns are diverse and complex, with some designs closely approaching the boundaries of our dark pattern definitions, complicating the annotation process.
  Exemplar 3: We initially defined the dark pattern "Ad Closure Failure" as "Clicking the close button does not close the ad." However, during recording and annotation, we observed cases where clicking the close button immediately triggered another ad, forcing users to continue watching ads—a behavior closely aligned with "Ad Closure Failure."

> **Principle 3**: For recurring cases, researchers identified common characteristics and incorporated new explanations into the codebook and manual. In exemplar 3, researchers unanimously concluded that this design met the criteria for "Ad Closure Failure" and added corresponding explanations to the manual.

Following refinement, we finalized the codebook and manual. The final versions consolidated the dark patterns proposed by each researcher and referenced relevant works on dark pattern classification [15, 20, 39]. The codebook provides concise descriptions of the identified dark patterns and includes illustrative cases for each. These patterns were categorized into several themes, as discussed in §4.2.2. The final manual (provided in the online dataset repository [5]) supplements the codebook with specific interaction guidelines for various app scenarios as well as resolution approaches for borderline cases of dark patterns.

4.2.2 Codes & Descriptions. Our codebook encompasses 15 ads dark patterns identified in this study, as detailed in Table 3. Notably, 11 of these dark patterns have not been reported in prior works, while the remaining 4 patterns were previously identified by [15, 28, 38]. Following the classification framework proposed by Gray et al. [28], 14 dark patterns were grouped into 4 established themes: "Nagging," "Forced Action," "Sneaking," and "Interface Interference." One pattern, "Increased Ads with Use," did not align with these existing themes and was categorized under a newly introduced theme, "Deception of Acquaintance."

Building upon established classification frameworks [15, 24], we distinguish these dark patterns along two dimensions: ***static/dynamic*** and ***app/ad-level***. *Static* patterns are identifiable through single UI screenshots without additional context, while *dynamic* patterns require analysis of video-recorded interactions. Similarly, *ad-level* patterns occur specifically within ad displays, whereas *app-level* patterns persist throughout general app usage. Among the 15 identified dark patterns (Table 3), 7 are static and 8 dynamic. Furthermore, 8 patterns operate at ad level and 7 at app level, as indicated in the respective columns of Table 3. Dynamic patterns particularly benefit from AdsDP's video recordings, which provide necessary contextual information for identification.

Table 3. **Our Codebook on the Definitions, Descriptions, and Attributions of Ads Dark Patterns in Apps.** The "Novel" column distinguishes newly identified patterns ("Yes") from those previously documented in literature ("No"). The "Static or Dynamic" classification differentiates between patterns identifiable through single screenshots ("Static") versus those requiring sequential data like videos ("Dynamic"). The "Level" column specifies whether patterns occur exclusively during ad displays ("Ad-level") or persist throughout the application interface ("App-level").

Theme	Dark Pattern	Description	Example Figures (Figures 8 and 9)	Novel	Static or Dynamic	Level
Nagging	App Resumption Ads	Users return to the home screen temporarily. Upon switching back to the app, they are forced to watch an ad	Figure 8(a)	Yes	Dynamic	App-level
	Unexpected Full-Screen Ads	A full-screen ad unexpectedly appears when clicking a button, or appears unprompted without any user interaction	Figure 8(b)	Yes	Dynamic	App-level
	Auto-Redirect Ads	Ad redirects to its landing page without the user clicking a "Skip" or "Close" button	Figure 8(c)	Yes	Dynamic	Ad-level
	Long Ad/Many Ads	Excessive ad exposure during app usage	Figure 8(d)	Yes	Dynamic	App-level
Forced Action	Barter for Ad-Free Privilege	App offers ad-free options through methods such as watching videos, viewing ads, or leaving a five-star review	Figure 8(e)	Yes	Static	App-level
	Paid Ad Removal	App offers paid options to remove ads	Figure 8(f)	No	Static	App-level
	Reward-Based Ad	App requires user to watch ads for benefits	Figure 8(g)	No	Static	App-level
	Ad Without Exit Option	Ad does not provide a close button or delays its appearance during display	Figure 8(h)	Yes	Static	Ad-level
	Ad Closure Failure	Ad does not close after clicking the close button	Figure 8(i)	Yes	Dynamic	Ad-level
Deception of Acquaintance	Increased Ads with Use	An app displays an increased number of ads to users who repeatedly use it	Figure 9(a)	Yes	Dynamic	App-level
Sneaking	Gesture-Induced Ad Redirection	Ad interprets easily triggered action signals (such as shaking the phone or hovering a finger over the ad) as clicking and be activated inadvertently	Figure 9(c)	Yes	Dynamic	Ad-level
Interface Interference	Button-Covering Ads	Ad is inappropriately positioned, obscuring users' interaction with the system buttons or in-app buttons	Figure 9(d)	Yes	Dynamic	Ad-level
	Multiple Close Buttons	Ad shows multiple close buttons simultaneously	Figure 9(e)	Yes	Static	Ad-level
	Bias-Driven UI Ads	Ad presents options to users and prominently highlights the option benefiting the advertiser	Figure 9(b)	No	Static	Ad-level
	Disguised Ads	Ad visually resembles normal content in apps or system UI of iOS	Figure 9(f)	No	Static	Ad-level

We provide detailed descriptions of all these 15 ads dark patterns below. Illustrative examples from the AdsDP dataset are provided in Figures 8 and 9 (Appendix B). Note that each dark pattern may exhibit multiple distinct manifestations; these variations are explicitly identified and labeled for clarity.

- **App Resumption Ads**: When using an app, users may temporarily exit the app by accessing the iPhone's Control Center or swiping up to return to the Home Screen. Upon returning to the app, they may be forced to view an advertisement before resuming their activity, disrupting their original experience.

- **Unexpected Full-Screen Ads**: These ads may manifest in two distinct forms: either triggered by user interaction with a button (denoted as *"Button-Triggered Unexpected Ads"*, a phenomenon also identified by Chen et al. [15]), or appearing spontaneously without any user input (classified as *"Unprompted Intrusive Ads"*).
- **Auto-Redirect Ads**: An Auto-Redirect Ad automatically redirects to its landing page after the ad concludes, without the user clicking a "Skip" or "Close" button. Redirecting users to a landing page generates more revenue for developers than merely displaying an ad. However, users are required to exert extra effort to exit the landing page and return to the app. This dark pattern is similar to "Automating the User Away" identified by Gray et al. [27].
- **Long Ad/Many Ads**: Excessive exposure to ads—whether due to an overabundance of ads or excessively long ad durations—negatively impacts the user experience.
- **Barter for Ad-Free Privilege**: Some apps offer users the option to remove ads (or access an ad-free experience) through methods such as watching videos, viewing ads, or leaving a five-star review. Users are thus required to invest time and effort (e.g., *watching ads* or *completing specific actions*) to gain the privilege of avoiding further ads. In practice, developers may deploy excessive ads to boost ad revenue while simultaneously pressuring users into taking actions that benefit the developers in exchange for the ad-free experience.
- **Paid Ad Removal**: Some apps offer a paid option to remove ads, as observed by Chen et al. [15]. Even if users are not interested in purchasing an ad-free service, developers may repeatedly prompt users to make this purchase through intrusive pop-up windows. Some apps further encourage spending by bundling the ad-free option with a more expensive "premium" package.
- **Reward-Based Ads**: Users may also be required to watch ads in exchange for other benefits, such as *"earning game items"* or *"unlocking additional features,"* a practice also identified by Chen et al. [15]. While some developers and users consider it reasonable to invest effort to access optional content, many users indicate a preference for paying or completing more meaningful tasks (such as in-game challenges) to unlock content instead of watching ads.
- **Ad Without Exit Option**: Some ads do not provide a close button or delay their appearance to force users to watch them, a behavior identified in mobile apps by Long et al. [38]. This approach extends the time during which users are exposed to ads, but may result in users wasting time on content they are not interested in.
- **Ad Closure Failure**: After clicking the close button, an ad may fail to close as expected. Despite users' clear intention to "exit the ad," the developers may not execute this action and may attempt to retain users by displaying another ad interface or redirecting them to a landing page. Some developers even provide fake close buttons that do not respond or display another ad immediately after clicking.
- **Increased Ads with Use**: Initially, an app may not display startup ads or show very few ads. However, after repeated use, users may begin to see more ads. This gradual increase in ad frequency can lead to users either abandoning the app or paying for an ad-free experience.
- **Gesture-Induced Ad Redirection**: Some ads are triggered by easily detected user actions (e.g., shaking the phone or hovering a finger over the ad), which can result in inadvertent activation, as observed in [38]. *"Shake-to-open"* ads, for example, are triggered by movement sensors, while *"Hover-to-open"* ads are activated when a finger hovers over the ad. Developers may set low action thresholds to capture slight user movements.
- **Button-Covering Ads**: Ads may obscure system or functional buttons within the app, preventing users from interacting with them. Developers position ads in a way that overlaps other buttons to increase the click-through rate, causing users to inadvertently click on the ads instead of the intended button.
- **Multiple Close Buttons**: Some ads display multiple close buttons simultaneously, making it difficult for users to choose the correct one. This tactic increases the time users spend on the ad interface. Some buttons may be fake, leading to an ad redirect rather than closing the ad.

- **Bias-Driven UI Ads**: In ads where users are presented with options, the visual design often emphasizes choices that benefit the advertiser, such as directing users to a landing page or prompting them to watch another ad. This strategy creates a false impression that these options are essential for users to select.
- **Disguised Ads**: Ads may be designed to closely resemble regular content within apps or the system UI, making them difficult for users to recognize as ads. This dark pattern, previously identified by Brignull et al. [13], often manifests in ads that imitate iOS notifications or display alarming messages.

4.2.3 Comparison of Our Ads Dark Patterns with Existing Taxonomies. Our taxonomy of ads dark patterns substantially advances existing classification frameworks through both extension and refinement. Among the 11 novel patterns we identified, three patterns—namely "Auto-Redirect Ads," "Ad Without Exit Option," and "Button-Covering Ads"—represent adaptations of behaviors previously documented in other domains, including social media and e-commerce platforms. The "Auto-Redirect Ads" pattern, for instance, corresponds conceptually to the "Automating the User Away" strategy identified by Gray et al. [27] in their study of Reddit interfaces. Two additional patterns ("Unexpected Full-Screen Ads" and "Gesture-Induced Ad Redirection") exhibit partial alignment with existing literature but are substantially refined through empirical evidence from the AdsDP dataset. Specifically, the former pattern provides more precise characterization than the generic "Pop-up Ads" classification proposed by Chen et al. [15], incorporating specific temporal and dimensional attributes, while the latter expands upon the "shake-to-open" mechanism described by Long et al. [38] by introducing a complementary "hover-to-open" variant. The remaining six patterns constitute original contributions, representing previously undocumented deceptive advertising practices identified through our study.

4.3 Operation and Video Recording

Four researchers spent five months operating all selected 485 apps, documenting their interactions through iOS screen recordings. Specifically, these apps were divided into two groups of nearly equal size, with each group comprising approximately half Chinese and half English apps. For each group, two researchers were assigned to independently create App Store accounts, download, and operate the apps. This process yielded two approximately five-minute videos per app, ensuring minimal bias from limited interaction by a single researcher.

Prior to operating the iOS apps, researchers activated the built-in screen recording feature on iPhone XR devices to capture all app interfaces. Recording user interaction movements, such as clicks or swipes, is critical as these actions provide essential context for annotators to infer the researchers' intentions during app operation. However, a technical limitation arises: screen recording alone does not capture touch interactions, as finger touch points are not displayed on the screen. To address this, we enabled the iPhone's AssistiveTouch feature and connected a mouse to the device via an OTG adapter. This setup allowed the iPhone to be controlled by the mouse, with the cursor's position visible on the screen. Left mouse clicks emulated finger taps, while dragging the mouse simulated swipes. Full technical specifications appear in Appendix C. To mitigate the risk of exposing researchers' personal information during operation and recording, we adopted a privacy-preserving recording method recommended by [14, 56]. Researchers were permitted to pause and re-record videos at any point if they believed personal information was inadvertently revealed. Additionally, they could request the deletion or blurring of any video segments that might compromise their privacy.

Each app was operated for approximately five minutes, during which researchers performed a series of actions akin to a standard usability inspection [44] for ads: (1) Creating an account and logging into the app, with this process excluded from public videos to avoid privacy risks. (2) Using the app as a typical user, including interacting with various buttons and exploring different interfaces. (3) Temporarily navigating to the Home Screen or Control Center and then returning to the app. (4) Clicking buttons likely to display ads and viewing the ads. (5) During non-full-screen ad displays, attempting to interact with the Home Indicator and functional buttons within the app. (6) Simulating finger swipes on ads by dragging the mouse and shaking the iPhone.

(7) Attempting to close the ad after interaction. (8) Reopening the app and repeating steps (1)–(7) to ensure comprehensive exploration. Researchers were allowed to switch between actions flexibly and determine the sequence of actions independently.

4.4 Dark Pattern Annotation

Following the recording process, researchers annotated dark patterns in ads within each video according to the established codebook and manual. Specifically, each researcher independently reviewed and annotated all videos they recorded, as well as those by their counterpart in the same app group, resulting in two annotation drafts per video. Subsequently, researchers collaboratively reviewed each video alongside its two drafts, engaging in discussions to resolve discrepancies and achieve consensus. In cases where consensus could not be reached, an additional researcher acted as an arbitrator. This process yielded a single, unanimous annotation for each video.

Our annotation process revealed two significant challenges, for which we developed mitigation strategies.

- **Challenge 1**: The annotation process was tedious, repetitive, and time-consuming, leaving researchers fatigued and less attentive sometimes.

> **Strategy 1**: To address this challenge, we arranged for researchers to alternate between operating apps and recording videos (§4.3) and annotating videos (§4.4). Furthermore, we developed a Python-based annotation tool using the Tkinter library to help researchers format annotation files. This tool enabled researchers to concentrate on the annotation content, alleviating the mental strain associated with repetitive formatting.

- **Challenge 2**: Certain cases that were not encountered during the dark pattern identification phase (§4.2.1) emerged, requiring researchers to determine whether these cases qualified as dark patterns.

> **Strategy 2**: Regular meetings were conducted throughout the annotation process to discuss newly discovered cases and devise appropriate countermeasures. Adjustments were made to the codebook and manual as needed. Notably, during these discussions, we incorporated three additional dark patterns ("Button-Covering Ads," "Multiple Close Buttons," "Increased Ads with Use") while eliminating one ("Ad Reappearance"). A representative meeting record is made public in an online dataset repository [5]. It is important to note that both the codebook and manual were finalized before the joint review of annotation drafts to resolve discrepancies. This ensured that all researchers referred to the most up-to-date version of the codebook when filling in missing dark patterns or revising incorrect annotations.

4.4.1 Ensuring Annotation Reliability. Building upon the methodology of Di Geronimo *et al.* [20], we implemented a multi-stage verification protocol to ensure annotation accuracy. The process comprised three key components: (1) independent duplicate recordings of each app by separate researchers, (2) parallel annotation of video recordings by two distinct researchers, and (3) a reconciliation phase involving arbitration by a third researcher. This comprehensive approach ensured triple validation of all annotations.

To maintain annotation consistency, we conducted a two-stage reliability assessment measuring inter-rater agreement through Krippendorff's α coefficient. In each stage, four annotators independently assessed 100 time-coded dark pattern instances using a 5-point Likert scale (ranging from "strongly disagree" to "strongly agree"). Statistical analysis identified divergent cases, which were subsequently reviewed in calibration meetings. This systematic approach resulted in an 11% improvement in inter-annotator consistency between assessment stages, reflecting strengthened consensus in pattern recognition.

5 Validation

We performed a statistical analysis of AdsDP, revealing distinctive characteristics of ads dark patterns within the dataset. We further validated the proposed ads dark patterns across popular Android and iOS apps. Additionally, we evaluated the performance of a state-of-the-art dark pattern detection tool, UIGuard, on AdsDP, to determine whether the dataset presents significant challenges for such automated solutions. Furthermore, we extended our analysis to user reviews from a wider range of iOS applications to investigate the prevalence of the identified dark patterns in real-world app advertising practices.

5.1 Statistical Characterization of Ads Dark Patterns in AdsDP

Recall that AdsDP encompasses a total of 485 iOS apps, including 240 apps from the China App Store and 245 apps from the U.S. App Store. We analyze the prevalence of ads dark patterns at two levels, as outlined in Table 3: *app-level* (§5.1.2) and *ad-level* (§5.1.3), comparing their distribution across the two language-specific iOS app markets.

5.1.1 Distribution of iOS Apps in terms of App Category. Both Chinese and English iOS apps span diverse app categories, each falling within a total of 37 categories. The Top 10 app categories of those Chinese iOS apps include *Navigation* (7.5%), *Entertainment* (5.0%), *Word Puzzles* (5.0%), *Simulation* (4.6%), *Puzzle & Brain Teasers* (4.6%), *Books* (3.8%), *Lifestyle* (3.8%), *Productivity* (3.8%), *Photography & Video* (3.3%), and *Health & Fitness* (3.3%), totally accounting for 44.7%. The Top 10 app categories of English iOS apps include *Entertainment* (5.7%), *Social* (5.3%), *Photography & Video* (4.9%), *Puzzle & Brain Teasers* (4.9%), *Education* (4.5%), *Casual* (4.5%), *Simulation* (4.5%), *Shopping* (4.1%), *Health & Fitness* (4.1%), and *Action* (4.1%). Note that these two subsets of apps share 5 categories of the Top 10 categories, which are *Entertainment*, *Photography & Video*, *Puzzle & Brain Teasers*, *Simulation*, and *Health & Fitness*. The distribution results are further visualized in Figures 11 and 12 (Appendix D).

5.1.2 Prevalence of App-Level Ads Dark Patterns within Chinese and English iOS Apps. We present the statistics on app-level ads dark patterns within both Chinese and English iOS apps in Table 4. The data reveals that Chinese and English apps receive an average of 33.3 and 307.1 ad-related user comments, respectively, highlighting the genuine and significant user concerns regarding ads in these apps.

Table 4. **Statistics of App-Level Ads Dark Patterns in Apps (Chinese vs. English)**

App dataset	#Apps	#Avg. ad-related comments	#Apps w/ app resumpt. ads	#Apps w/ unexpected full-screen ads	#Apps w/ paid ad removal service	#Apps w/ barter for ad-free privilege	#Apps w/ increased ads with use	#Apps w/ reward-based ads
App Dataset (Chinese)	240	33.3	35 (14.6%)	94 (39.2%)	65 (27.1%)	9 (3.8%)	5 (2.1%)	88 (36.7%)
App Dataset (English)	245	307.1	18 (7.5%)	79 (32.9%)	86 (35.8%)	0 (0%)	2 (0.8%)	80 (33.3%)

App Resumption Ads. Through our manual analysis, we identified that 35 (14.6%) Chinese apps and 18 (7.5%) English apps feature *App Resumption Ads*. Among the two observed manifestations (*Return from Home Screen* and *Return from Control Center*), 32 Chinese and 15 English apps displayed full-screen ads upon resuming from the Home Screen, while 13 Chinese and 8 English apps showed a large ad when returning from the Control Center drop-down menu – both disrupting normal app interaction.

Unexpected Full-Screen Ads. Additionally, 94 (39.2%) Chinese apps and 79 (32.9%) English apps incorporate *Unexpected Full-Screen Ads*. The majority (66.2% Chinese, 60.9% English) appeared following a button click (*Button-Triggered Unexpected Ads*), while the remaining appeared during normal app usage without any interaction (*Unprompted Intrusive Ads*).

Paid Ad Removal. Furthermore, 65 (27.1%) Chinese apps and 86 (35.8%) English apps provide *Paid Ad Removal* services. The pricing for ad removal services is consistently higher in English apps than in Chinese apps. Specifically, the weekly cost is $1.21 (Chinese) versus $3.01 (English), and annual plans at $15.02 (Chinese) compared to $46.88 (English). The complete pricing structure across subscription periods is illustrated in Figure 13 (Appendix D).

Barter for Ad-Free Privilege. 9 (3.8%) Chinese apps offer *Barter for Ad-Free Privilege* options, such as watching videos, viewing ads, or leaving a five-star review, whereas such options were not observed in English apps.

Increased Ads with Use. Moreover, we identified that in 5 (2.1%) Chinese apps and 2 (0.8%) English apps, no startup ads were displayed upon the first entry, and there were almost no other ads. However, upon the second entry, a startup ad appeared, and the number of other ads significantly increased.

Reward-Based Ads. 88 (36.7%) of Chinese apps and 80 (33.3%) of English apps incorporate *Reward-Based Ads*. These predominantly appeared in two forms: 93.9% of Chinese and 85.8% of English implementations offered *earning game items* option, while the remainder provided *unlocking additional features*.

5.1.3 Prevalence of Ad-Level Dark Patterns within iOS in-app Ads. Table 5 presents a comparative statistical analysis of ad-level dark patterns across both datasets. Through direct app operation, we documented ads exhibiting at least one dark pattern from §4.2.2. Specifically, AdsDP contains 1,125 pattern-exhibiting ads in Chinese apps (4.7 per app) versus 850 in English apps (3.5 per app).

Table 5. **Ad-Level Dark Pattern Analysis (Chinese vs. English).** The columns from "Video ads" to "Video full-screen ads" describe ad formats rather than specific patterns. "Shake-to-open" and "Hover-to-open" represent *Gesture-Induced Ad Redirection* variants, while "Home indicator-covering Ads" exemplifies *Button-Covering Ads*.

App dataset	# Total ads	# Video ads	# Full-screen ads	# Video full-screen ads	# Shake-to-open ads	# Hover-to-open ads	# Home indicator-covering ads	# Ads w/ multiple close buttons	# Ads w/ closure failure	# Ads w/o exit option	# Auto-redirect ads	# Bias-driven UI ads	# Disguised ads
App dataset (Chinese)	1125	627 (55.7%)	897 (79.7%)	526 (46.8%)	354 (31.5%)	70 (6.2%)	9 (0.8%)	142 (12.6%)	258 (22.9%)	371 (33.0%)	24 (2.1%)	53 (4.7%)	7 (0.6%)
App dataset (English)	850	471 (55.4%)	476 (56.0%)	434 (51.1%)	6 (0.7%)	4 (0.5%)	2 (0.2%)	65 (7.6%)	377 (44.4%)	587 (69.1%)	90 (10.6%)	8 (0.9%)	8 (0.9%)

Video Ads and Full-Screen Ads. In the Chinese dataset, 627 (55.7%) ads are video ads, compared to 471 (55.4%) ads in the English dataset, while 897 (79.7%) ads in the Chinese dataset are full-screen ads, significantly higher than the 476 (56.0%) ads in the English dataset. Additionally, the Chinese dataset contains 526 (46.8%) video full-screen ads, slightly fewer than the 434 (51.1%) ads in the English dataset. For video ads with a claimed ad length, we examined the length of in-app ads in both Chinese and English iOS apps. Figure 14 (Appendix D) compares the distribution of ad durations between Chinese iOS apps (represented by the solid black line) and English iOS apps (represented by the dashed orange line). Our analysis reveals that ads in the Chinese dataset are generally shorter than those in the English dataset. Specifically, 49.3% of ads in the Chinese dataset are less than 5 seconds, and 53.3% last for less than 10 seconds. In contrast, only 22.9% of ads in the English dataset are shorter than 5 seconds, and 26.9% are under 10 seconds. For longer ads exceeding 30 seconds, both datasets exhibit a similar proportion, with 13.7% in the Chinese dataset and 13.4% in the English dataset.

Dark Patterns within iOS in-app Ads. The analysis reveals significant regional variations in ad-related dark patterns. *Shake-to-open* ads appear in 31.5% of Chinese apps compared to just 0.7% of English apps, while *hover-to-open* implementations are rare in both markets (6.2% Chinese vs. 0.5% English). The Chinese dataset contains 9 (0.8%) ads obscuring the home indicator (*Button-Covering Ads*), a pattern absent in English apps. Notable differences emerge in closure mechanisms: English apps show lower rates of *multiple close buttons* (7.6% vs 12.6%)

and more frequently lack exit options (69.1% vs 33.0%). *Auto-redirect ads* are more common in English apps (10.6% vs 2.1%), while *bias-driven UI ads* predominate in Chinese apps (4.7% vs 0.9%). Both markets show minimal *disguised ads*, with 0.6% in the Chinese dataset and 0.9% in the English dataset.

Ad closure failures prove particularly problematic, affecting 44.4% of English ads versus 22.9% of Chinese ads. Figure 2 illustrates common failure modes: *Multi-Step Ad Closure*, *Closure Redirect Ads*, *Unclosable Ads*, *Consecutive Ad Replacement*, and *Forced Ad-Free Purchase Prompts*. The *Multi-Step Ad Closure* pattern requires users to complete multiple dismissal actions, as the initial close button click merely redirects to another advertisement page. *Closure Redirect Ads* immediately direct users to promotional landing pages, typically app store destinations, upon closure attempts. *Unclosable Ads* remain persistently visible despite user attempts to dismiss them through interface controls. The *Consecutive Ad Replacement* pattern automatically displays new advertisements in the same location following the closure of previous ones. *Forced Ad-Free Purchase Prompts* present subscription offers immediately after ad dismissal, effectively transforming the closure action into a monetization opportunity. These patterns collectively degrade user experience through increased interaction costs and unintended outcomes, demonstrating widespread implementation of user-hostile ad closure practices.

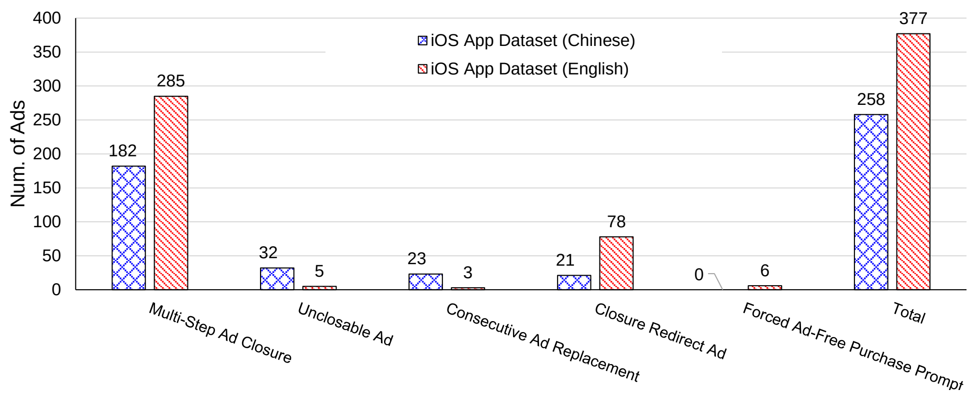

Fig. 2. Various Manifestations of Failed Ad Closure

In summary, our analysis reveals that certain dark patterns, such as, *shake-to-open ads*, and *bias-driven UI ads*, are more prevalent in the Chinese dataset. Conversely, *ads closure failure*, *ads without exit options*, and *auto-redirect ads* are more common in the English dataset. Additionally, *hover-to-open* and *disguised ads* are less prevalent in both datasets. These differences suggest potential variations in user experience and ad design strategies across the Chinese and English app markets.

5.2 Statistical Characterization of Ads Dark Patterns in Popular iOS and Android Apps

To validate the prevalence of identified dark patterns beyond our initial study of 485 iOS apps with significant ad-related comments, we extended our investigation to include popular apps across both Android and iOS platforms. Our selection methodology involved: (1) sourcing top downloaded U.S. apps from an app insights service[1]; and (2) sampling 20 apps each from three download rank tiers: top 100, next 900 (101-1,000), and subsequent 9,000 (1,001-10,000), totaling 60 apps per platform.

Following training using our established codebook and manual, two additional annotators conducted systematic evaluations using the procedure outlined in §4.3 and §4.4. Android app testing employed an emulator [8] running Android 16.0, consistent with prior research methodologies [37, 57].

[1] https://appwisp.com

5.2.1 Prevalence of App-Level Dark Patterns in Popular Apps (iOS vs. Android). Table 6 reveals that "Unexpected Full-Screen Ads" (iOS: 38.3%, Android: 48.3%) and "Paid Ad Removal" (iOS: 46.7%, Android: 38.3%) emerge as the most prevalent patterns. Compared to AdsDP's English apps, we observed a 5.4% (iOS) and 15.4% (Android) increase in "Unexpected Full-Screen Ads", while "Reward-based Ads" decreased by 3.3% (iOS) and 8.3% (Android). Notably, "Paid Ad Removal" appeared in 46.7% of iOS apps (10.9% higher than AdsDP), and Android showed a 19.2% higher incidence of "App Resumption Ads" (26.7% prevalence).

Table 6. **Comparative Statistics of App-Level Ads Dark Patterns in AdsDP versus Popular Android and iOS Apps.** The symbols ↑/↓ denote percentage deviations exceeding ±3% from the baseline, while ↑↑/↓↓ indicate more substantial deviations surpassing ±10%.

App dataset	#Apps	#Apps w/ app resumpt. ads	#Apps w/ unexpected full-screen ads	#Apps w/ paid ad removal service	#Apps w/ barter for ad-free privilege	#Apps w/ increased ads with use	#Apps w/ reward-based ads
Baseline: AdsDP Apps in English (data from Table 4)	245	18 (7.5%)	79 (32.9%)	86 (35.8%)	0 (0%)	2 (0.8%)	80 (33.3%)
Popular Apps of **iOS**	60	5 (8.3%)	23 (38.3%) ↑	28 (46.7%) ↑↑	0 (0%)	0 (0%)	18 (30%) ↓
Popular Apps of **Android**	60	16 (26.7%) ↑↑	29 (48.3%) ↑↑	23 (38.3%)	0 (0%)	0 (0%)	15 (25%) ↓

5.2.2 Prevalence of Ad-Level Dark Patterns in Popular Apps (iOS vs. Android). Table 7 reveals that popular apps display higher ad densities (iOS: 5.2, Android: 4.8 ads/app) compared to AdsDP (3.5 ads/app). While patterns like "Shake-to-Open" and "Disguised Ads" maintain consistent prevalence across datasets, notable differences emerge: "Ads Closure Failure" increases significantly in popular apps (iOS: 50.6%, Android: 52.0% vs AdsDP's 44.4%), and "Auto-Redirect Ads" show platform-specific surges (iOS +4.5%, Android +14.6%). Platform variations appear in "Hover-to-Open" (4.5% Android prevalence) and reduced "Ads Without Exiting Option" in iOS (-8.2%).

Table 7. **Comparative Statistics of Ad-Level Ads Dark Patterns in AdsDP versus Popular Android and iOS Apps.**

App dataset	# Total ads	# Shake -to- open ads	# Hover -to- open ads	# Home indicator -covering ads	# Ads w/ multiple close buttons	# Ads w/ closure failure	# Ads w/o exit option	# Auto-redirect ads	# Bias-driven UI ads	# Disguised ads
Baseline: AdsDP Apps in English (data from Table 5)	850	6 (0.7%)	4 (0.5%)	2 (0.2%)	65 (7.6%)	377 (44.4%)	587 (69.1%)	90 (10.6%)	8 (0.9%)	8 (0.9%)
Popular Apps of **iOS**	312	1 (0.3%)	2 (0.6%)	1 (0.3%)	21 (6.7%)	158 (50.6%) ↑	190 (60.9%) ↓	47 (15.1%) ↑	4 (1.3%)	10 (3.2%)
Popular Apps of **Android**	286	0 (0%)	13 (4.5%) ↑	1 (0.3%)	15 (5.2%)	149 (52.0%) ↑	205 (71.7%)	72 (25.2%) ↑↑	7 (2.4%)	2 (0.7%)

The analysis confirms the persistent prevalence of proposed dark patterns in popular apps, though with shifted emphases - decreased "reward-based ads" but increased deployment of disruptive patterns like "unexpected full-screen ads." While Android-focused studies [20] might yield comparable taxonomies, AdsDP's iOS-specific documentation addresses critical research gaps given the platform's distinct ecosystem and data collection constraints.

5.3 Ineffectiveness of Existing Dark Pattern Detection Solutions on AdsDP

Recent research has introduced two advanced dark pattern detection tools, UIGuard [15] and AidUI [39], which we assessed using our AdsDP dataset. Both tools identify dark patterns by analyzing individual screenshots, employing a combination of computer vision (CV) and natural language processing (NLP) techniques. These tools examine spatial distribution, color schemes, and textual content within UI elements to uncover evidence of dark patterns. For example, UIGuard comprises two primary components: the *Property Extraction* module and the *Knowledge-Driven Dark Pattern Checker*. In the *Property Extraction* module, UIGuard utilizes PaddleOCR [46] to extract text and calculates the bounding box of text using the Faster R-CNN model [48]. It also detects ad-related icons using the ResNet-18 model [33], extracts UI element colors, and conducts contrast analysis with adjacent elements. These features are subsequently used by the *Knowledge-Driven Dark Pattern Checker* to identify dark patterns through a series of predefined steps. For instance, the checker detects the dark pattern "False Hierarchy" (similar to our "Bias-Driven UI Ads") by identifying two clustered UI elements, one of which contains text matching patterns such as "`No/no thanks/close/next time/later/...`", which users are more likely to select.

Following the respective guidelines of UIGuard and AidUI, we implemented AidUI within a Docker container while utilizing the Python source code and pre-trained models from UIGuard. Evaluation on their original datasets yielded detection metrics consistent with published results: AidUI achieved precision=0.66, recall=0.67, and F1=0.65; UIGuard demonstrated precision=0.82, recall=0.77, and F1=0.79. These outcomes verify our successful replication of both detection systems.

5.3.1 Mapping of AdsDP Dark Patterns to the Dark Patterns Detectable by the Two Detection Tools. To assess whether AidUI and UIGuard can effectively detect the dark patterns identified in our study, we established a mapping between the dark patterns described in our work and those detectable by AidUI and UIGuard. This involved comparing the descriptions of each dark pattern in their publications with those we identified. If the descriptions were similar, we mapped the patterns accordingly. If not, we further analyzed their detection processes to determine their applicability to our identified patterns, creating mappings where appropriate.

For AidUI, we successfully mapped 3 of our dark patterns ("App Resumption Ads", "Unexpected Full-Screen Ads", and "Disguised Ads") to 2 dark patterns detectable by AidUI (namely, "Nagging" and "Disguised Ad"). In contrast, we mapped 7 of our dark patterns to 5 dark patterns detectable by UIGuard, as detailed in Table 8 (Appendix E). Given AidUI's limited focus on advertisement-related dark patterns, we derived an image dataset based on AdsDP and evaluated the performance of UIGuard on AdsDP.

5.3.2 Preparation of the AdsDP Dataset for Evaluating UIGuard. Since UIGuard detects dark patterns based on individual images, we derived an image dataset from our AdsDP video dataset to evaluate its performance on AdsDP. From videos annotated with the 7 dark patterns compatible with UIGuard's detection capabilities, we compiled 103 representative screenshots. The last column in Table 8 details the number of screenshots extracted for each dark pattern. We gathered 10 to 20 screenshots for 6 types of dark patterns, with "App Resumption Ads" further divided into two specific manifestations: "Return from Home Screen" (15) and "Return from Control Center" (10). Other types include "Unexpected Full-Screen Ads (Unprompted Intrusive Ads)" (15), "Reward-Based Ads" (15), "Paid Ad Removal" (20), "Bias-Driven UI Ads" (14), and "Disguised Ads" (10). Additionally, due to the limited instances of "Barter for Ad-Free Privilege (Viewing Ads)", we collected only 4 screenshots.

5.3.3 UIGuard's Performance on AdsDP. We evaluated UIGuard's performance on AdsDP by processing all 103 screenshots. Successful detection required UIGuard's output to match the annotated dark pattern type in AdsDP. As shown in Figure 3, gray bars indicate the 97 undetected cases out of 103 total screenshots. For these failures, we examined UIGuard's intermediate outputs containing extracted UI features (e.g., recognized ad-related icons and text) to identify missed characteristics, such as complex textual content or sophisticated UI designs, that contributed to detection failures.

We identified four potential causes of detection failures on those 97 screenshots: insufficient implementation, nonstandard icon, unexpected text pattern, and non-English text.

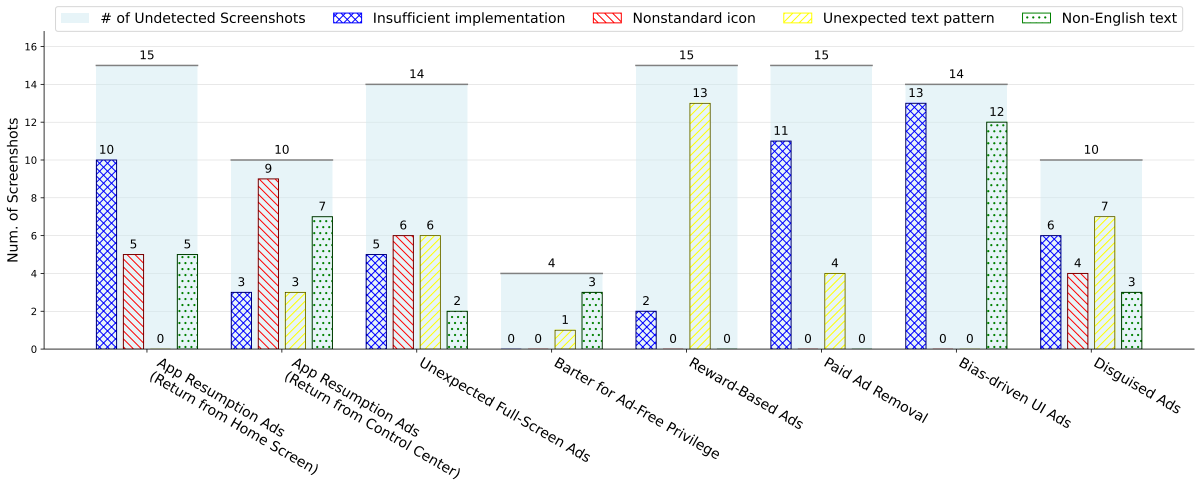

Fig. 3. **Analysis of Undetected Cases for 7 Ad Dark Patterns (8 Manifestations) in UIGuard Evaluation.** Gray bars indicate total undetected screenshots, while striped bars represent detection failures attributed to specific causes: **Insufficient Implementation** (tool limitations), **Nonstandard Icon** (atypical ad indicators), **Unexpected Text Pattern** (unconventional wording), and **Non-English Text** (language barriers). Multiple factors may contribute to individual detection failures.

- **Insufficient implementation**: Specifically, 50 screenshots were not detected by UIGuard due to this reason, with the most frequent errors arising in the detection of the "App Resumption Ads (Return from Home Screen)" (10), "Paid Ad Removal" (11), and "Bias-driven UI Ads" (13). When detecting "Return from Home Screen," the primary errors stemmed from the icon semantic module's inability to recognize ad-related icons, such as ▷ and ×. For "Paid Ad Removal," errors arose due to ads employing small font sizes or intricate typography, causing the text extraction module to misinterpret text or improperly segment phrases (*e.g.*, "Remove ads" being identified as "Bemove ads" or "Removeads"). For "Bias-driven UI Ads," the element grouping module failed to disclose the asymmetry relation of options, resulting in detection failures. Figure 4(a) illustrates an example of an "insufficient implementation" error encountered by UIGuard while detecting the "Bias-driven UI Ads" dark pattern. The text extraction module misinterpreted the text "close|5" in the left option as "closes," which caused failure in matching the "Bias-driven UI Ads" text pattern "`No/no thanks/close/next time`". Furthermore, the element grouping module did not compare the "Close" and "Learn More" buttons. These limitations contributed to the failure in this instance.
  Finding: The detection process proposed in UIGuard theoretically could detect these screenshots; however, in practice, errors accumulated across various modules (*e.g.*, the OCR failing to correctly extract text in screenshots), resulting in UIGuard's inability to detect these dark pattern instances.
- **Nonstandard icon:** We found that in 24 screenshots, UIGuard failed to detect the ads in the captured images due to this reason. For instance, Figure 4(b) shows an undetected "App Resumption Ads (Return from Control Center)" screenshot, where the ad does not feature an ad icon but prominently displays the text "AD" in the lower-left corner. For such images, although humans can easily identify them as advertisement screenshots, it is difficult for UIGuard to detect them.
  Finding: UIGuard is designed to detect ads that comply with the adChoices self-regulatory program [3], which includes a specific icon style. However, many ads in the AdsDP dataset use custom-designed icons, making it challenging for UIGuard to detect them.

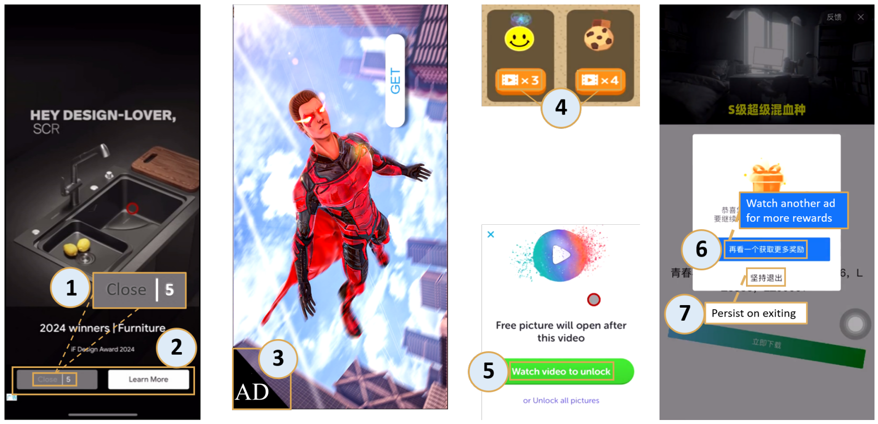

(a) Insufficient implementation (b) Nonstandard icon (c) Unexpected text pattern (d) Non-English text

Fig. 4. **Detection Failures in UIGuard.** (a) Misinterpretation of "close|5" (1) as "closes," while failing to detect the asymmetric design between "close|5" and "Learn More" buttons (2) that should be symmetrical. (b) Failure to recognize the ad due to absence of standard icon ▷, despite visible "AD" indicator (3). (c) Inability to interpret the reward icon (4), "Watch an ad to get 3 or 4 items," or identify "Watch video to unlock" (5) as "Tempting to Watch" pattern, as the text pattern "`...watch...ad/ads...`" used by UIGuard did not match this phrase. (d) Failure to analyze non-English button text (6-7).

- **Unexpected text pattern:** This detection limitation affected 34 screenshots, predominantly in "Reward-Based Ads" cases (13 instances). UIGuard's reliance on the exact text pattern "`...watch...ad/ads...`" proves insufficient for two common scenarios: (1) icon-based representations of reward mechanics, as illustrated in Figure 4(c), where an icon implies "Watch ads to earn three or four game items," and (2) textual variations like "`watch video to unlock`" (Figure 4(c)) that deviate from the predefined pattern while maintaining equivalent semantic meaning.
  Finding: UIGuard organizes common textual expressions of specific dark patterns into predefined text patterns, which are then matched against the extracted text from screenshots. However, many ads convey similar meanings using unexpected textual descriptions or even graphical elements.
- **Non-English text:** We found that in 32 screenshots, the text within the ads was not written in English, making it difficult for UIGuard to detect them. For example, Figure 4(d) shows an undetected "Bias-driven UI Ads" dark pattern screenshot. In this screenshot, the "Watch another ad" option is highlighted with a blue background (while the "Exit" option is not obvious and is with a normal background color). However, because the text in this screenshot is not written in English, UIGuard failed to detect this dark pattern.
  Finding: UIGuard focuses on detecting ads in English; however, some information in the screenshots (such as the content of options or text identifying the ads) is not written in English.

In summary, among the 15 dark patterns proposed in this paper, 12 are not covered by AidUI, and eight are not covered by UIGuard, highlighting that the majority of dark pattern types in AdsDP have not been addressed or incorporated into existing detection tools. Six of these unaddressed dark patterns occur dynamically, which cannot be detected by detectors based on static screenshots. Furthermore, due to the four aforementioned reasons, the performance of UIGuard in detecting the seven dark patterns it claimed to cover significantly decreased, and it failed to detect 94% of the sampled screenshots. Dark pattern instances in AdsDP with these characteristics present significant challenges for existing detectors which are all based on individual images (thus lacking the contextual information), demonstrating the urgent need of our video-based AdsDP dataset.

5.4 Usefulness of Our Codebook in Helping Identify Ads Dark Patterns in More Apps

To assess the prevalence of identified dark patterns in real-world advertising practices, we analyzed 10,000 randomly selected ad-related user reviews from 8,187 iOS apps. Using GPT-4, we classified each review according to our codebook's dark pattern definitions through an iterative refinement process. The initial classification prompt was progressively improved by incorporating misclassification examples, culminating in a final prompt with six counterexamples. We excluded reviews discussing "Long Ads/Many Ads" (already quantified at 36.9% in §3). The classification results are presented in Figure 5(a), with the prompt detailed in Figure 15 (Appendix F).

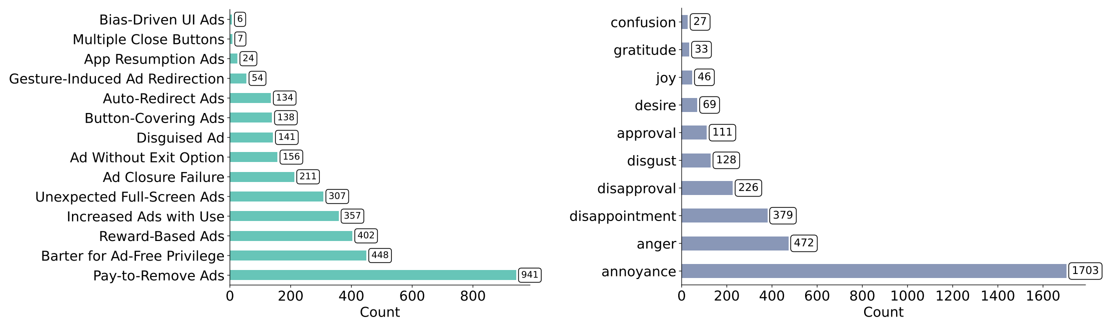

(a) # of Reviews Related to each Dark Pattern Type (b) Top 10 Sentiment Types in Reviews Describing Dark Patterns

Fig. 5. Count and Sentiment of Reviews Describing Ads Dark Patterns

The most prevalent patterns were: "Pay-to-Remove Ads" (941), "Barter for Ad-Free Privilege" (448), "Reward-Based Ads" (402), "Increased Ads with Use" (357), and "Unexpected Full-Screen Ads" (307). Analysis revealed that 75% of "Pay-to-Remove Ads" reviews criticized either the pricing rationale (e.g., *"Is it even fair to pay $4.99 for the app and then an additional 99 cents to remove ads? Seriously?"*) or reported non-functional purchases (e.g., *"I paid to remove ads, but they keep appearing."*). Notably, some users considered this pattern legitimate: *"Lots of ads, which I understand need to be monetized. I'd love the option of a paid ad-free service."*

Other frequently mentioned patterns included "Ad Without Exit Option" (156), "Disguised Ad" (141), "Button-Covering Ads" (138), and "Auto-Redirect Ads" (134). User reviews often described compound patterns, such as "Ad Closure Failure" combined with "Unexpected Full-Screen Ads" (*"Ads appear during video playback and redirect when closed"*). Less common patterns like "Gesture-Induced Ad Redirection" (54) and "App Resumption Ads" (24) appeared less frequently, likely due to their complex implementation. Some reviews provided detailed interaction descriptions, such as *"If I switch to another app and return to Yealico, an ad pops up even though I've paid to remove ads. Please fix."*, for "App Resumption Ads."

Sentiment analysis conducted using GPT-4 (the prompt is detailed in Figure 16, Appendix F) revealed strong user reactions to dark patterns, as visualized in Figure 5(b). Predominant negative sentiments included "annoyance" (1,703 instances) and "anger" (472), with users frequently warning others against app usage (*"Excessive time-wasting ads - avoid this app"*). Additional negative responses encompassed "disappointment" (379), "disapproval" (226), and "disgust" (128). Conversely, a minority expressed positive reactions: "approval" (111), "desire" (69), "joy" (46), and "gratitude" (33).

Our examination of 10,000 ad-related reviews revealed 3,326 instances aligning with our codebook's dark pattern definitions, demonstrating their prevalence beyond the AdsDP dataset across diverse applications. Applying the codebook enabled more granular analysis than §3, exposing previously obscured user experiences like specific ad interaction issues (2.1% reporting "Ad Closure Failure") and advertising volume trends (3.57% noting "Increased Ads with Use"), which were previously overshadowed by generic complaints about "Excessive Ads" (36.9%).

This approach demonstrates how codebooks can be effectively integrated with LLMs to analyze large-scale user feedback with reliable accuracy, providing deeper insights into user perceptions of advertising designs.

6 Discussion

This section delineates the necessity for establishing a refined dark pattern taxonomy, outlines practical applications of AdsDP, suggests mitigation approaches, discusses the study's limitations, and elucidates potential ethical considerations.

6.1 The Need for a Fine-grained Taxonomy

Building on prior research [15, 27, 28, 38], our taxonomy introduces two novel discriminative features—timing and required user actions—to systematically distinguish ads dark patterns. First, timing critically influences both user experience and pattern classification. While app-launch ads, recommended by Google [4], are generally acceptable (§4.2, Exemplar 1), disruptive patterns during app usage like "App Resumption Ads" (triggered upon returning from home screen) and "Unexpected Full-Screen Ads" (during active use) qualify as dark patterns due to their intrusive nature. Second, user-action requirements reveal design intent: "Gesture-Induced Ad Redirection" demands explicit engagement, whereas "Auto-Redirect Ads" occur passively. Regional adoption data strongly validates this distinction—Chinese apps favored gesture-based redirection (424 vs 10 instances), while English apps preferred automatic redirects (90 vs 24).

This refined taxonomy offers three key advances: (1) resolving definitional overlaps like the ambiguous "Bait-and-Switch" noted by Shi *et al.* [52]; (2) integrating contextual cues to address taxonomy fragmentation (Gray *et al.* [29]); and (3) delivering actionable guidance for developers to optimize ad implementation through precise timing and sizing specifications.

6.2 Applications of AdsDP

The AdsDP dataset presents valuable opportunities for multiple stakeholders. Machine learning researchers can utilize this resource to enhance ad detection algorithms [36, 53, 62] and develop video-based models for dynamic dark pattern identification [15, 39]. Advertising scholars may employ automated tools to analyze pattern distribution across apps and providers, revealing monetization strategies through involuntary user engagement. For HCI researchers, our methodology establishes a useful approach for future dark pattern investigations. End users can also benefit from studying representative dark pattern cases to improve their recognition and avoidance of deceptive advertising practices.

6.3 Implications for Mitigating the Widespread Dark Patterns

Our research reveals widespread implementation of dark patterns in mobile advertising, driven by revenue models prioritizing impressions and clicks. Developers routinely misuse advertising SDKs to serve disruptive ads, while certain SDK providers enable these practices through technical documentation [9, 61] despite regulatory concerns [25, 60]. The identified patterns enable practical countermeasures, such as new ads dark pattern detection tools. Educational initiatives [12] can incorporate these findings to demonstrate manipulative techniques, while ad-blocking communities [21] may develop targeted solutions against offending apps and SDKs.

Addressing these challenges requires coordinated efforts across the advertising ecosystem. Advertising SDK providers should implement and enforce standardized guidelines for ad implementations, rigorously test apps to prevent deceptive UI integrations, and monitor ad request frequency to mitigate excessive exposure. App developers must prioritize transparent advertising practices and actively incorporate user feedback, strategically optimizing ad timing and sizing based on our findings to enhance user perception. App store platforms should establish comprehensive policies featuring mandatory disclosures, automated pattern detection systems, streamlined reporting channels, and graduated enforcement measures. This multifaceted approach balances consumer protection with sustainable monetization while addressing current regulatory gaps.

6.4 Limitations

We acknowledge two primary limitations in our study. The first limitation pertains to potential biases in identifying and discussing dark patterns in ads, including observer bias, self-report bias, and recall bias. Observer bias may occur during app operation, as researchers, despite being well-trained and vigilant, may sometimes fail to recognize the presence of dark patterns, a phenomenon referred to as "DP-Blindness" in [20]. Self-report bias could arise during meetings, where, despite instructions for thorough reporting, fatigue or distractions may lead to incomplete discussion of certain observations. Recall bias may also affect participants who report dark patterns from previous sessions, leading to potential inaccuracies due to memory lapses. Crowdsourcing may be an effective way to mitigate some biases introduced by a small group of researchers. Following Chen *et al.*'s findings [15] about improved recognition through user education, future work could employ crowdsourced annotations with sufficiently trained participants to mitigate these biases. The second limitation is that our work focuses mainly on in-app advertisements within the iOS platform. Furthermore, our investigation was limited to apps from the China and U.S. App Stores, where distinct characteristics in the use of dark patterns in ads were observed. This highlights the importance of extending future research to include apps from more app stores and countries in order to gain a more comprehensive understanding of dark patterns in ads.

6.5 Ethical Considerations

The experimental procedures adhered to the guidelines established by the ethics committee of our university and complied with regulations concerning the handling of sensitive and private data, as well as anonymization protocols. Specifically, during the collection of anonymous user reviews, no personally identifiable information about the reviewers was obtained. For the video collection and annotation process, all participating researchers provided individual consent for data-sharing through the AdsDP dataset. They also agreed that the information captured in the videos could be publicly shared and utilized in future research. Regarding VPN usage for accessing the U.S. App Store, we followed those established research protocols from prior work [51, 55, 58, 59], and had all associated accounts and services permanently deactivated after study completion. Since the experiments conducted in this study did not involve additional human subjects, no further formal ethics approval was required under our university's policy.

7 Conclusion

In this paper, we introduce AdsDP, an annotated video dataset capturing dark patterns within iOS in-app advertisements, developed through a multi-collaborator workflow. To construct this database, we first developed a codebook and a manual for ads dark patterns by systematically investigating these patterns across 100 iOS applications in two language-specific app markets. This analysis identified 15 distinct types of ads dark patterns, 11 of which have not been previously documented. Using the codebook and manual as a guide, four researchers thoroughly operated each of the 485 applications on an iPhone, recording the entire app usage process on video and subsequently annotating the instances of dark patterns. In total, AdsDP contains 718 videos, with a cumulative duration of 60 hours, documenting 5,782 occurrences of the 15 identified ads dark patterns.

Additionally, we performed a statistical analysis of the dataset across the two types of language-specific apps, revealing the prevalence of ad-related dark patterns and the variations in their types and frequencies between the app categories. We also evaluated the performance of state-of-the-art dark pattern detection solutions using AdsDP, observing a significant decline in their effectiveness on our novel dataset, underscoring the urgent need for the development of new detection methods for ads dark patterns. Lastly, we compared user reviews from a broader range of iOS apps with our definition of ads dark patterns, discovering that such patterns may be widespread in in-app advertising practices. Our codebook can assist in their detection. We anticipate that AdsDP will aid researchers in systematically analyzing ads dark patterns, ultimately contributing to the enhancement of detection techniques and fostering more ethical advertising practices.

8 Dataset and Code Availability

The datasets generated during this study are publicly available at [5].

Acknowledgments

This work was supported in part by the National Key R&D Program of China (No.2023YFB3106800) and the National Natural Science Foundation of China (No.62272410).

References

[1] Jacob Aagaard, Miria Emma Clausen Knudsen, Per Bækgaard, and Kevin Doherty. 2022. A Game of Dark Patterns: Designing Healthy, Highly-Engaging Mobile Games. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*. 1–8.

[2] Ad units, ad formats and ad types 2024. *Ad units, ad formats and ad types*. Retrieved Feb, 2025 from https://support.google.com/admob/answer/6128738?hl=en-GB&sjid=2202240538547538821-NC

[3] adChoice 2024. *YourAdChoices*. Retrieved Feb, 2025 from https://youradchoices.com/

[4] AdMob Guidance 2024. *App open ad guidance and best practices*. Retrieved Feb, 2025 from https://support.google.com/admob/answer/9341964?hl=en&sjid=5295512244428965357-NC

[5] AdsDP Dataset 2025. *AdsDP Dataset*. Retrieved May, 2025 from https://zenodo.org/records/15316373

[6] Mohammed Aldeen, Jeffrey Young, Song Liao, Tsu-Yao Chang, Long Cheng, Haipeng Cai, Xiapu Luo, and Hongxin Hu. 2024. End-Users Know Best: Identifying Undesired Behavior of Alexa Skills Through User Review Analysis. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 8, 3 (2024), 1–28.

[7] Aashir Amaar, Wajdi Aljedaani, Furqan Rustam, Saleem Ullah, Vaibhav Rupapara, and Stephanie Ludi. 2022. Detection of fake job postings by utilizing machine learning and natural language processing approaches. *Neural Processing Letters* (2022), 1–29.

[8] Android Studio 2025. *Android Studio*. Retrieved Apr, 2025 from https://developer.android.com/studio

[9] AppLovin Ad Guidance 2024. *Banner & MREC Ads*. Retrieved Feb, 2025 from https://developers.applovin.com/en/max/android/ad-formats/banner-and-mrec-ads/

[10] Giorgio Brajnik and Silvia Gabrielli. 2010. A review of online advertising effects on the user experience. *International Journal of Human-Computer Interaction* 26, 10 (2010), 971–997.

[11] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.

[12] H Brignull, M Leiser, C Santos, and K Doshi. 2023. *Deceptive patterns – user interfaces designed to trick you*. https://www.deceptive.design/

[13] Harry Brignull, Marc Miquel, Jeremy Rosenberg, and James Offer. 2015. Dark patterns-user interfaces designed to trick people. In *Proceedings of the Poster Presentation, Australian Psychological Society Congress, Sydney, NSW, Australia*. 21–23.

[14] Barry Brown, Moira McGregor, and Donald McMillan. 2014. 100 days of iPhone use: understanding the details of mobile device use. In *Proceedings of the 16th international conference on Human-computer interaction with mobile devices & services*. 223–232.

[15] Jieshan Chen, Jiamou Sun, Sidong Feng, Zhenchang Xing, Qinghua Lu, Xiwei Xu, and Chunyang Chen. 2023. 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*. 1–20.

[16] Consumer Protection Report 2022. *Bringing Dark Patterns to Light*. Retrieved Sep, 2022 from https://www.ftc.gov/reports/bringing-dark-patterns-light

[17] Create a rewarded ad unit 2024. *Create a rewarded ad unit*. Retrieved Feb, 2025 from https://support.google.com/admob/answer/7311747?sjid=2202240538547538821-NC

[18] John W Creswell and Cheryl N Poth. 2016. *Qualitative inquiry and research design: Choosing among five approaches*. Sage publications.

[19] Peter R Darke and Robin JB Ritchie. 2007. The defensive consumer: Advertising deception, defensive processing, and distrust. *Journal of Marketing research* 44, 1 (2007), 114–127.

[20] Linda Di Geronimo, Larissa Braz, Enrico Fregnan, Fabio Palomba, and Alberto Bacchelli. 2020. 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*. 1–14.

[21] EasyList Community 2024. *EasyList*. Retrieved Feb, 2025 from https://easylist.to/

[22] Jill J Francis, Marie Johnston, Clare Robertson, Liz Glidewell, Vikki Entwistle, Martin P Eccles, and Jeremy M Grimshaw. 2010. What is an adequate sample size? Operationalising data saturation for theory-based interview studies. *Psychology and health* 25, 10 (2010), 1229–1245.

[23] Cuiyun Gao, Yichuan Man, Hui Xu, Jieming Zhu, Yangfan Zhou, and Michael R Lyu. 2017. IntelliAd: Assisting mobile app developers in measuring Ad costs automatically. In *2017 IEEE/ACM 39th International Conference on Software Engineering Companion (ICSE-C)*. IEEE, 253–255.

[24] Cuiyun Gao, Jichuan Zeng, David Lo, Xin Xia, Irwin King, and Michael R Lyu. 2022. Understanding in-app advertising issues based on large scale app review analysis. *Information and Software Technology* 142 (2022), 106741.

[25] General Data Protection Regulation 2018. *Art. 4 GDPR Definitions*. Retrieved Feb, 2025 from https://gdpr-info.eu/art-4-gdpr/

[26] Daniel G Goldstein, R Preston McAfee, and Siddharth Suri. 2013. The cost of annoying ads. In *Proceedings of the 22nd international conference on World Wide Web*. 459–470.

[27] Colin M Gray, Shruthi Sai Chivukula, and Ahreum Lee. 2020. What Kind of Work Do" Asshole Designers" Create? Describing Properties of Ethical Concern on Reddit. In *Proceedings of the 2020 acm designing interactive systems conference*. 61–73.

[28] Colin M Gray, Yubo Kou, Bryan Battles, Joseph Hoggatt, and Austin L Toombs. 2018. The dark (patterns) side of UX design. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–14.

[29] Colin M Gray, Cristiana Teixeira Santos, Nataliia Bielova, and Thomas Mildner. 2024. 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*. 1–22.

[30] Jiaping Gui, Stuart Mcilroy, Meiyappan Nagappan, and William GJ Halfond. 2015. Truth in advertising: The hidden cost of mobile ads for software developers. In *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, Vol. 1. IEEE, 100–110.

[31] Johanna Gunawan, Amogh Pradeep, David Choffnes, Woodrow Hartzog, and Christo Wilson. 2021. A comparative study of dark patterns across web and mobile modalities. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–29.

[32] Emil Häglund and Johanna Björklund. 2024. AI-driven contextual advertising: Toward relevant messaging without personal data. *Journal of Current Issues & Research in Advertising* 45, 3 (2024), 301–319.

[33] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.

[34] Barbara Johnstone and Jennifer Andrus. 2024. *Discourse analysis*. John Wiley & Sons.

[35] Hao-Ping Hank Lee, Jacob Logas, Stephanie Yang, Zhouyu Li, Nata Barbosa, Yang Wang, and Sauvik Das. 2023. When and why do people want ad targeting explanations? evidence from a four-week, mixed-methods field study. In *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2903–2920.

[36] Kiho Lee, Chaejin Lim, Beomjin Jin, Taeyoung Kim, and Hyoungshick Kim. 2024. AdFlush: A Real-World Deployable Machine Learning Solution for Effective Advertisement and Web Tracker Prevention. In *Proceedings of the ACM on Web Conference 2024*. 1902–1913.

[37] Zhe Liu, Chunyang Chen, Junjie Wang, Yuekai Huang, Jun Hu, and Qing Wang. 2022. Guided bug crush: Assist manual gui testing of android apps via hint moves. In *Proceedings of the 2022 CHI conference on human factors in computing systems*. 1–14.

[38] Mengyi Long, Yue Xu, Jiangrong Wu, Qihua Ou, and Yuhong Nan. 2023. Understanding Dark UI Patterns in the Mobile Ecosystem: A Case Study of Apps in China. In *Proceedings of the 2023 ACM Workshop on Secure and Trustworthy Superapps*. 33–40.

[39] SM Hasan Mansur, Sabiha Salma, Damilola Awofisayo, and Kevin Moran. 2023. Aidui: Toward automated recognition of dark patterns in user interfaces. In *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, 1958–1970.

[40] Arunesh Mathur, Gunes Acar, Michael J Friedman, Eli Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. 2019. Dark patterns at scale: Findings from a crawl of 11K shopping websites. *Proceedings of the ACM on human-computer interaction* 3, CSCW (2019), 1–32.

[41] Arunesh Mathur, Mihir Kshirsagar, and Jonathan Mayer. 2021. 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*. 1–18.

[42] Ngoc Thi Nguyen, Agustin Zuniga, Hyowon Lee, Pan Hui, Huber Flores, and Petteri Nurmi. 2020. (m) ad to see me? intelligent advertisement placement: Balancing user annoyance and advertising effectiveness. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 2 (2020), 1–26.

[43] Liming Nie, Yangyang Zhao, Chenglin Li, Xuqiong Luo, and Yang Liu. 2024. Shadows in the interface: a comprehensive study on dark patterns. *Proceedings of the ACM on Software Engineering* 1, FSE (2024), 204–225.

[44] Jakob Nielsen. 1994. Usability inspection methods. In *Conference companion on Human factors in computing systems*. 413–414.

[45] Lorelli S Nowell, Jill M Norris, Deborah E White, and Nancy J Moules. 2017. Thematic analysis: Striving to meet the trustworthiness criteria. *International journal of qualitative methods* 16, 1 (2017), 1609406917733847.

[46] PaddleOCR 2024. *PaddleOCR*. Retrieved Feb, 2025 from https://github.com/PaddlePaddle/PaddleOCR

[47] Personalised and non-personalised ads 2024. *Personalised and non-personalised ads*. Retrieved Feb, 2025 from https://support.google.com/admob/answer/7676680?hl=en-GB&sjid=2202240538547538821-NC

[48] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2016. Faster R-CNN: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence* 39, 6 (2016), 1137–1149.
[49] Precedence Research. 2024. *Mobile Application Market Size, Share, and Trends 2024 to 2034*. Retrieved Fab, 2024 from https://www.precedenceresearch.com/mobile-application-market
[50] Zhanna Sarsenbayeva, Niels van Berkel, Danula Hettiachchi, Weiwei Jiang, Tilman Dingler, Eduardo Velloso, Vassilis Kostakos, and Jorge Goncalves. 2019. Measuring the effects of stress on mobile interaction. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3, 1 (2019), 1–18.
[51] Ashley Sheil, Gunes Acar, Hanna Schraffenberger, Raphaël Gellert, and David Malone. 2024. 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*. 1–24.
[52] Zewei Shi, Ruoxi Sun, Jieshan Chen, Jiamou Sun, Minhui Xue, Yansong Gao, Feng Liu, and Xingliang Yuan. 2025. 50 Shades of Deceptive Patterns: A Unified Taxonomy, Multimodal Detection, and Security Implications. *arXiv preprint arXiv:2501.13351* (2025).
[53] Sandra Siby, Umar Iqbal, Steven Englehardt, Zubair Shafiq, and Carmela Troncoso. 2022. {WebGraph}: Capturing advertising and tracking information flows for robust blocking. In *31st USENIX Security Symposium (USENIX Security 22)*. 2875–2892.
[54] Natalia Silberstein, Oren Somekh, Yair Koren, Michal Aharon, Dror Porat, Avi Shahar, and Tingyi Wu. 2020. Ad close mitigation for improved user experience in native advertisements. In *Proceedings of the 13th International Conference on Web Search and Data Mining*. 546–554.
[55] Alexander Sjösten, Peter Snyder, Antonio Pastor, Panagiotis Papadopoulos, and Benjamin Livshits. 2020. Filter list generation for underserved regions. In *Proceedings of The Web Conference 2020*. 1682–1692.
[56] John C Tang, Sophia B Liu, Michael Muller, James Lin, and Clemens Drews. 2006. Unobtrusive but invasive: using screen recording to collect field data on computer-mediated interaction. In *Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work*. 479–482.
[57] Faiza Tazi, Suleiman Saka, Griffin Opp, Shradha Neupane, Sanchari Das, Lorenzo De Carli, and Indrakshi Ray. 2023. Accessibility evaluation of IoT Android mobile companion APPs. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–7.
[58] Van Hong Tran, Aarushi Mehrotra, Marshini Chetty, Nick Feamster, Jens Frankenreiter, and Lior Strahilevitz. 2024. Measuring compliance with the california consumer privacy act over space and time. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–19.
[59] Kieron Ivy Turk and Alice Hutchings. 2023. Click Here to Exit: An Evaluation of Quick Exit Buttons. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–15.
[60] Unfair Competition Law 1993. *Unfair Competition Law, CAL. BUS. PROF CODE 17200*. Retrieved Feb, 2025 from https://leginfo.legislature.ca.gov/faces/codes_displaySection.xhtml?sectionNum=17200.&lawCode=BPC
[61] Unity Ad Guidance 2024. *Implementing rewarded ads in iOS*. Retrieved Feb, 2025 from https://docs.unity.com/ads/en-us/manual/ImplementingRewardedAdsiOS
[62] Zhiju Yang, Weiping Pei, Monchu Chen, and Chuan Yue. 2022. Wtagraph: Web tracking and advertising detection using graph neural networks. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1540–1557.
[63] Eric Zeng, Tadayoshi Kohno, and Franziska Roesner. 2021. What makes a "bad" ad? user perceptions of problematic online advertising. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–24.

Appendix

A Prompts for Topic Modeling and Comments Classification (corresponding to §3)

The prompt shown in Figure 6 guides GPT to read user comments and summarize prevalent ad-related topics. We independently applied this prompt four times, asking GPT to generate four topic lists, which were then manually refined and merged into a final topic list. The prompt in Figure 7 instructs GPT to read a single comment at a time and categorize it under an ad-related topic based on its content.

B Instances of Ads Dark Patterns from AdsDP (corresponding to §4.2.2)

An example is provided for each type of dark pattern below to enhance the reader's understanding.

- **App Resumption Ads**: As depicted in Figure 8(a), the user exits a novel reading app by tapping the Home Indicator, returns to the app by clicking the app icon, and is presented with a large ad that obstructs their reading.

Prompt: Summarizing ad-related comments as topic list

Task: Please perform a thematic analysis of the provided input text (where each line represents a comment for an app), summarizing the content to extract several popular topics related to advertisements.
Notice: For each topic, kindly provide the following information: the *<topic name>*, the *<number of related comments>*, and the *<meaning of each topic>*.

Fig. 6. Prompt for Summarizing Ad-Related Comments as Topic List

Prompt: Classifying each comment into a topic in the proposed topic list

Here is the topics' list:
{The Whole Topic List}, supplemented with:
Miscellaneous Issues: Issues that cannot be classified into the aforementioned topics.
Every row of the list follows this format: *<Name of the topic>*: *<Definition of the topic>*.

Task: The user will now provide you with a comment on an app. Please read the comment carefully and determine which topic from the list it belongs to.
Notice:
1. Please only output the *<Name of the topic>*, such as "Excessive Ads."
2. Don't output the details of each topic.

Fig. 7. Prompt for Classifying Each Comment into a Topic

- **Unexpected Full-Screen Ads**: In the example shown in Figure 8(b), the user clicks the "Easy" button to start a game but is immediately shown a full-screen ad.
- **Auto-Redirect Ads**: Figure 8(c) illustrates a typical instance where the user watches an ad without interaction, and it automatically redirects to the iOS App Store's landing page.
- **Long Ad/Many Ads**: As shown in Figure 8(d), a user is subjected to 16 ads within five minutes of normal app usage, rendering the experience intolerable.
- **Barter for Ad-Free Privilege**: Figure 8(e) illustrates a scenario where the app developer allows users to skip all ads for the next three days by watching a single ad.
- **Paid Ad Removal**: As shown in Figure 8(f), a prominent interface encourages users to pay for the removal of ads.
- **Reward-Based Ads**: Figure 8(g) shows an app asking the user to watch an ad to unlock a feature.
- **Ad Without Exit Option**: For instance, the ad in Figure 8(h) lacks a close button, preventing the user from exiting and returning to the app.
- **Ad Closure Failure**: Figure 8(i) shows an instance where the user clicks the close button in the top-left corner, but the ad does not close, instead displaying another interface.
- **Increased Ads with Use**: Figure 9(a) illustrates this pattern, where no ads are shown upon the user's first entry, but ads appear and increase significantly on subsequent visits.
- **Gesture-Induced Ad Redirection**: Figure 9(c) shows an example where the user's finger hovers over the ad without clicking, yet the ad redirects the user to the landing page.

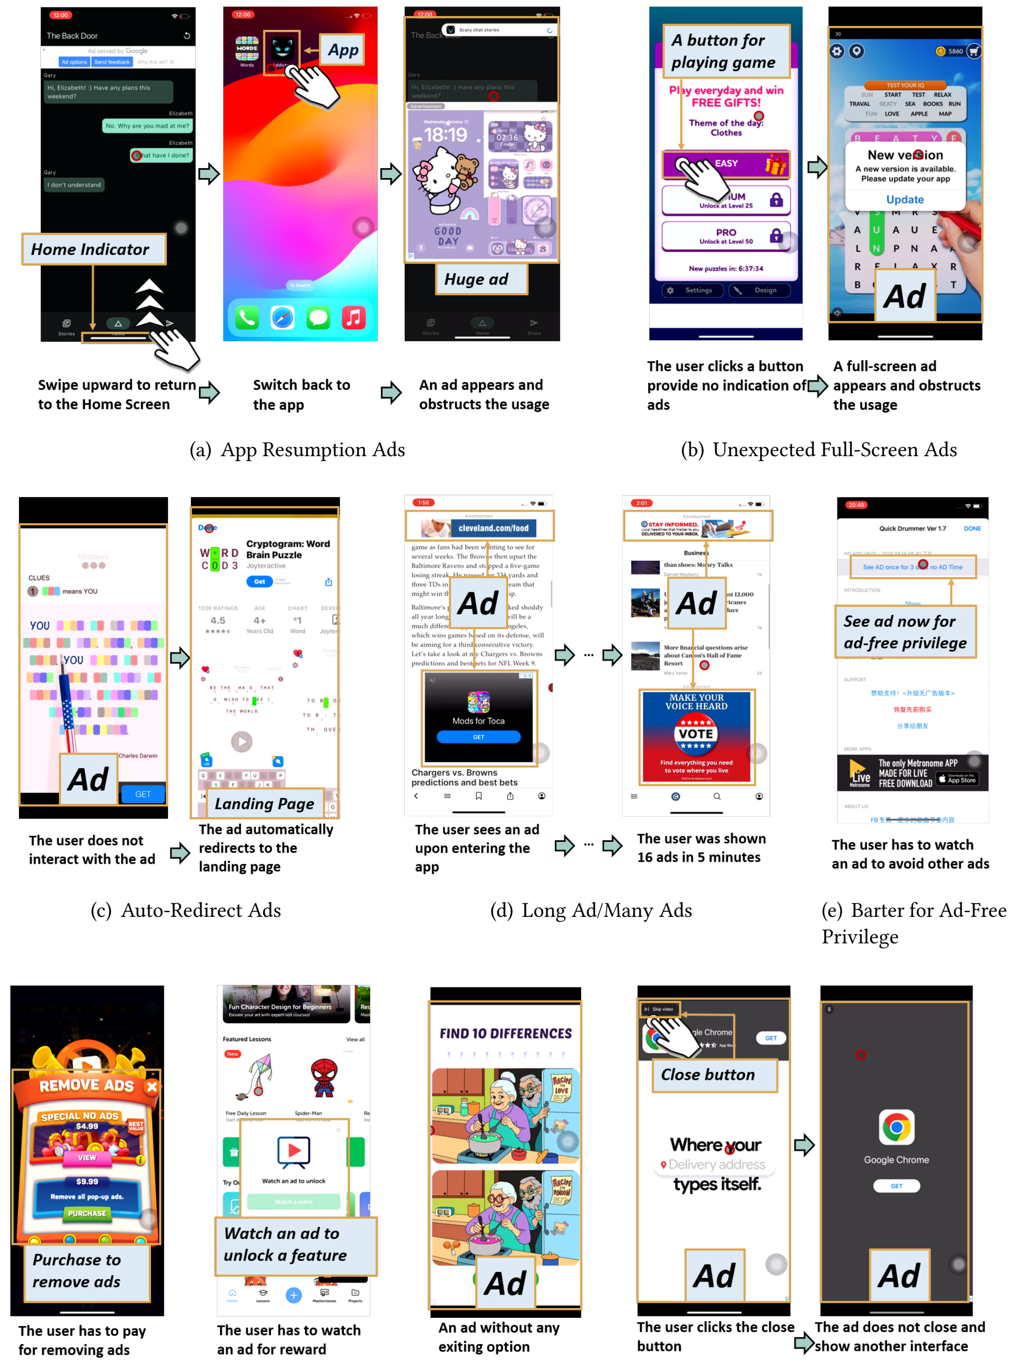

(a) App Resumption Ads

(b) Unexpected Full-Screen Ads

(c) Auto-Redirect Ads

(d) Long Ad/Many Ads

(e) Barter for Ad-Free Privilege

(f) Paid Ad Removal

(g) Reward-Based Ads

(h) Ad Without Exit Option

(i) Ad Closure Failure

Fig. 8. Cases of Dark Patterns in Theme "Nagging" and "Forced Action".

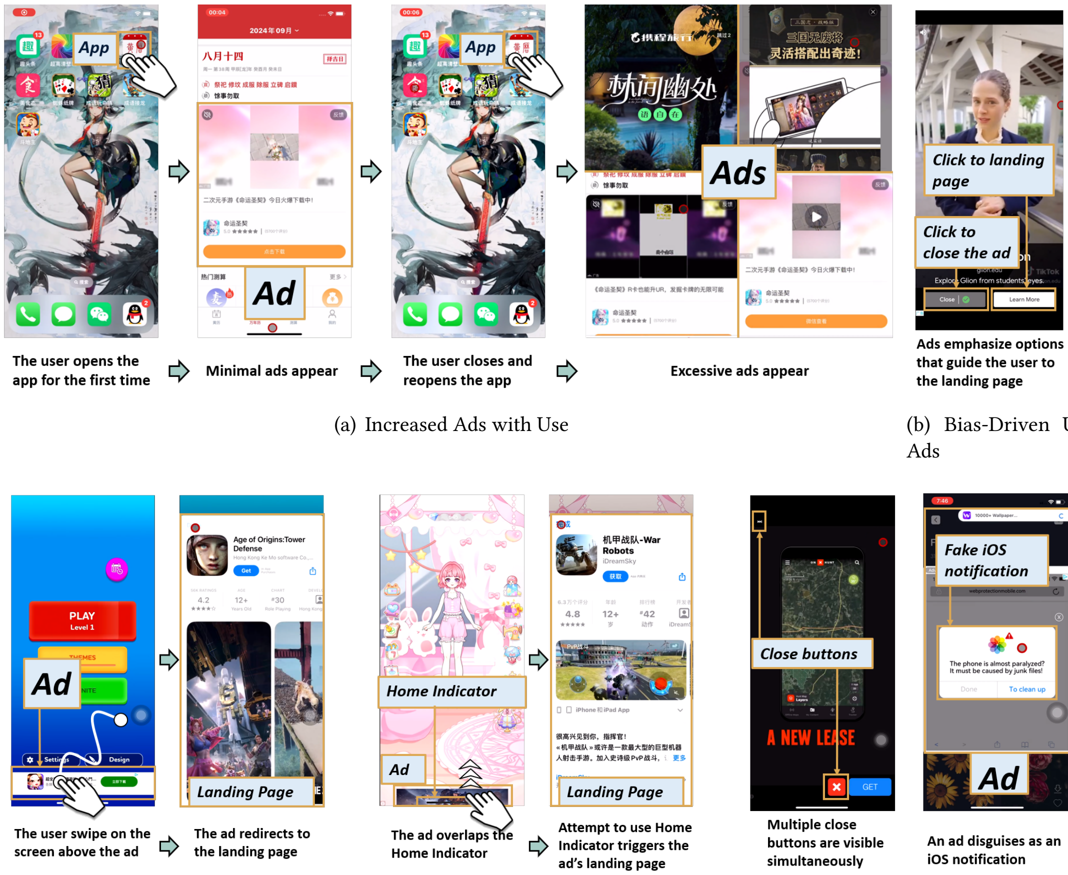

(a) Increased Ads with Use

(b) Bias-Driven UI Ads

(c) Gesture-Induced Ad Redirection

(d) Button-Covering Ads

(e) Multiple Close Buttons

(f) Disguised Ads

Fig. 9. Cases of Dark Patterns in Theme “Deception of Acquaintance”, “Sneaking” and “Interface Interference”. Each case is depicted using one or more keyframes extracted from a video clip in the AdsDP dataset, arranged chronologically from left to right. For static dark patterns, a single frame with the descriptive text above highlights the dark pattern's occurrence. For dynamic dark patterns, multiple frames capture key interactions and their effects, with text detailing them.

- **Button-Covering Ads**: Figure 9(d) illustrates how an ad overlaps the Home Indicator, causing an ad redirection when the user attempts to access it.
- **Multiple Close Buttons**: In Figure 9(e), the ad displays two buttons that both suggest they will “close the ad,” but only the hidden button in the top-left corner functions as expected.
- **Bias-Driven UI Ads**: For example, in Figure 9(b), the ad presents two options: one to close the ad (left) and one to go to the landing page (right). The right button is visually emphasized with bright colors, guiding users toward the landing page.
- **Disguised Ads**: Figure 9(f) shows an example where ads appear as iOS notifications, misleading users into thinking they are system-generated alerts.

(a) A full set of equipment utilized by one of the researchers to operate the apps

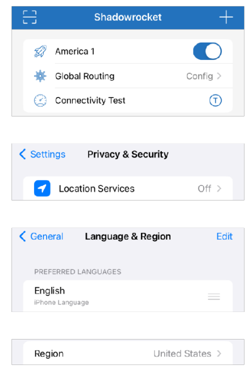

(b) Preparation by researchers before operating apps from the U.S. App Store

Fig. 10. The software and hardware configurations during operating and recording the apps

C Configurations for App Operation and Recording (corresponding to §4.3)

Figure 10 depicts the software and hardware configurations used for app operation and recording. As illustrated in Figure 10(a), each researcher used an iPhone XR running iOS 17.6 with AssistiveTouch enabled, displaying an AssistiveTouch menu on the screen. A mouse connected via an OTG adapter generated a visible cursor, customized to a high-contrast red color for enhanced visibility against UI elements. Screen recording was activated through the Control Center before app operation. Since the researchers were not located in the United States during the experiments, they utilized the ShadowRocket VPN and newly registered U.S. Apple accounts to download and operate apps from the U.S. App Store, as shown in Figure 10(b). Device settings, including language, country, and system region, were adjusted to simulate the behavior of a typical U.S. iPhone user, with location services disabled.

D Detailed Statistical Results on Characteristics of AdsDP (corresponding to §5.1)

The distribution of top app categories for Chinese and English markets is presented in Figure 11 and 12 respectively. Figure 13 analyzes the pricing models for ad removal services across various subscription durations, derived from annotation data in AdsDP. Figure 14 displays the cumulative distribution of advertisement durations (in seconds) for both Chinese and English applications.

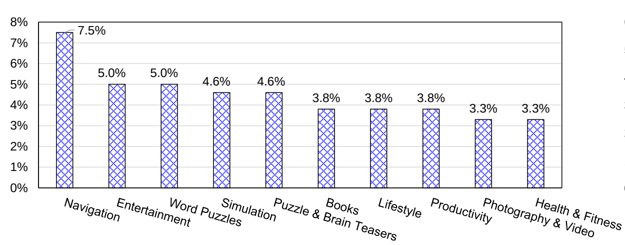

Fig. 11. Top 10 categories of Apps from the Chinese market

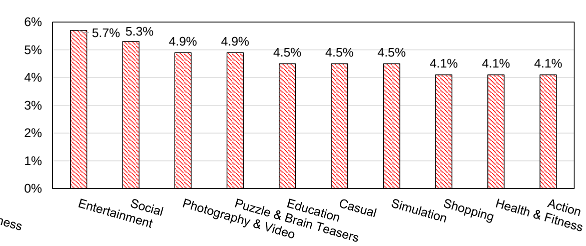

Fig. 12. Top 10 categories of Apps from the English market

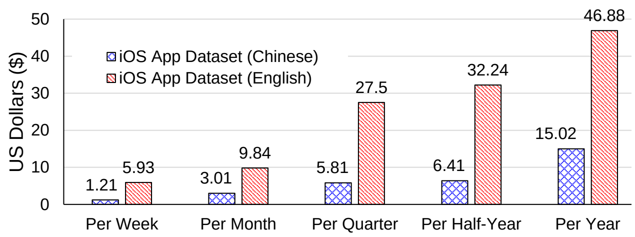

Fig. 13. Pricing of Ad Removal (Chinese vs. English)

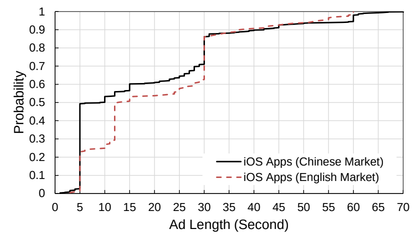

Fig. 14. CDF of Ad Length (seconds)

E Mapping of AdsDP Dark Patterns to Dark Patterns Detectable by UIGuard (corresponding to §5.3)

We compared the descriptions of each dark pattern in UIGuard's original paper with those we identified. If the descriptions were similar, we mapped the patterns accordingly. If not, we further analyzed their detection processes to determine their applicability to our identified patterns, creating mappings where appropriate. Through this process, we have established a mapping from AdsDP to UIGuard, as shown in Table 8.

Table 8. **Mapping of Dark Patterns from AdsDP to UIGuard**

Dark Pattern in AdsDP			**Corresponding Dark Pattern in UIGuard**		**# of Screenshots**
Code	Manifestations	Description	Code	Description	
App Resumption Ads	Return from Home Screen	Users return to the home screen temporarily by Home Indicator. Upon switching back to the app, they are forced to watch an ad	Pop-up Ads	A pop-up window unexpectedly and repeatedly appears and interrupts user tasks	15
	Return from Control Center	Users swipe upward to access the Control Center. Upon switching back to the app, they are forced to watch an ad			10
Unexpected Full-Screen Ads	Unprompted Intrusive Ads	A full-screen ad appears unprompted without any user interaction			15
Barter for Ad-Free Privilege	Watching ads	App offers ad-free options through viewing ads	General Types	Watch ads to unlock features or get rewards	4
Reward-Based Ads		App requires user to watch ads for benefits			15
Paid Ad Removal		App offers paid options to remove ads	General Types	Pay to avoid ads	20
Bias-driven UI Ads		Ad presents options to users and prominently highlights the option benefiting the advertiser	False Hierarchy	One option is more salient than other equal option	14
Disguised Ads		Ad visually resemble normal content in apps or system UI of iOS	Disguised Ad	Ads pretends to be normal content	10

F Prompts for Comments Classification and Sentiment Analysis Combined with Codebook (corresponding to §5.4)

The prompt shown in Figure 15 incorporates the codebook, instructing GPT to compare the comment content with the dark patterns described in the codebook to determine whether the comment reflects a dark pattern. The prompt in Figure 16 incorporates a list of human emotions, guiding GPT to analyze which user emotions are conveyed in the comment content.

Prompt: Determining whether a comment is related to a dark pattern

Here are some malicious or deceptive designs we have identified in advertisements, known as dark patterns:
{Codebook}
Every row of this list follows this format: <*Name of the dark pattern*>: <*Definition of the dark pattern*>.

Task: I will now provide you with a user comment discussing advertising practices.
Please read the comment carefully and determine if the comment describes a phenomenon that fully matches or closely resembles the definition of any of the dark patterns listed.
Notice:
1. You should not forcefully associate a comment with a specific dark pattern or overly interpret it. Here are the classification errors you previously made, which you should avoid in this task:
{Counterexamples List}
2. If you believe the comment does not relate to any of the dark patterns, simply output "N/A" with no additional content.
3. If you believe the comment relates to a specific dark pattern, only output the most relevant <*Name of the dark pattern*> without any explanation or additional content such as <*Definition of the dark pattern*>.

Fig. 15. Prompt for Determining whether a Comment is Related to a Dark Pattern

Prompt: Determining the sentiment expressed by the comment

Here is the sentiment' list:
{Sentiments List}
Every row of this list represents a human <*sentiment*>.

Task: The user will now provide you with a comment talking about the app's ad practice. Please read the comment carefully and determine which human sentiment is primarily reflected in the comment.
Notice:
1. Please only output the sentiment, such as "anger", don't output any other content.
2. You only need to output the <*sentiment*> that best represents this review, not multiple sentiments.

Fig. 16. Prompt for Determining the Sentiment Expressed by the Comment