International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0301
e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 2071 – 2080

Detection and Analysis of Dark Patterns in Modern Web Applications

Aadithya P[1], Kushal Kumar M[2], M V Puneeth[3], Nithesh Gowda G[4], Devadatta H S[5], Dr. Gurunath R[6]
[1,2,3,4,5]PG Scholar (MCA), Department of Master of Computer Applications, Dayananda Sagar College of Arts, Science and Commerce, Bengaluru, Karnataka, India.
[6]Associate Professor, Department of Master of Computer Applications, Dayananda Sagar College of Arts, Science and Commerce, Bengaluru, Karnataka, India.
Email ID: *aadipnabha@gmail.com[1], kushal7124@gmail.com[2], mvpuneeth2004@gmail.com[3], gowdanithesh28@gmail.com[4], Fatewritings@gmail.com[5]*

Abstract
Context: Dark patterns are user-interface design strategies deliberately engineered to coerce, mislead, or manipulate users into performing actions that serve business interests at the expense of user autonomy, privacy, and informed consent. As digital commerce and online services have proliferated, so too has the prevalence of such deceptive design choices, raising urgent ethical, legal, and technical questions. **Objective:** This paper presents a systematic review of dark patterns across modern web applications, proposes a multi-layered taxonomy, an automated detection framework (Dark Scan), and a severity-scoring model (DPSS). **Methods:** We employ a mixed-method approach combining automated DOM and behavioural analysis, expert interview synthesis, and controlled user-study evidence from published literature. **Results:** Dark patterns are identified in 78.4% of e-commerce platforms, 65.2% of SaaS services, and 54.7% of social-media applications. Exposure increases unintended purchase rates by 34.7% and reduces platform trust by 48.3%. **Conclusion:** The scale and harm of dark patterns demand coordinated responses. Our proposed Dark Scan pipeline achieves 91.4% precision and 88.7% recall, offering a practical compliance-auditing tool.
Keywords: dark patterns, deceptive design, user interface manipulation, web usability, privacy by design, HCI, UX ethics, automated detection, GDPR, consumer protection.

1. Introduction
The World Wide Web has fundamentally transformed how individuals access information, communicate, conduct commerce, and manage personal affairs. With this transformation has come an explosion in the sophistication of web application design both in service of genuine usability improvements and, increasingly, in service of manipulative strategies that exploit cognitive vulnerabilities for commercial gain. These strategies, collectively known as dark patterns, represent one of the most pressing challenges at the intersection of human-computer interaction (HCI), ethics, and digital regulation. The term 'dark pattern' was coined by UX designer Harry Brignull in 2010 to describe user interface elements deliberately crafted to trick users into doing things they did not intend or want to do. Since Brignull's initial formulation, the concept has been elaborated, critiqued, and empirically investigated by researchers across computer science, law, psychology, and consumer-protection studies. Legislators in the European Union, the United States, and numerous other jurisdictions have begun to respond with targeted regulatory frameworks, yet enforcement remains inconsistent and technically challenging. The fundamental problem is asymmetric: designers and developers possess deep technical and psychological knowledge that most users lack, and this knowledge differential is increasingly being weaponised. A progress bar that never quite completes, a cancellation flow that requires ten discrete steps, a pre-ticked consent checkbox in a seven-layer cookie dialogue — each of these represents not a design accident but an intentional choice made within a system whose incentive structures reward conversion rates over user wellbeing. This paper addresses three interconnected research questions:

- **RQ1:** What taxonomic framework best captures the full diversity of dark patterns observed in contemporary web applications?

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0301
e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 2071 – 2080

- **RQ2:** To what extent can dark patterns be detected automatically, and with what accuracy, using static DOM analysis and dynamic behavioural tracing?
- **RQ3:** What is the measurable impact of dark-pattern exposure on user decision quality, platform trust, and psychological wellbeing?

Our contributions are as follows. First, we present an empirically grounded, five-tier taxonomy of dark patterns. Second, we describe and evaluate Dark Scan, an automated detection pipeline. Third, we report findings on the behavioural and psychological impact of dark-pattern exposure. Fourth, we derive policy and design recommendations informed by existing regulatory frameworks.

2. Related Work

a) Foundations of Dark Pattern Research

Brignull's original catalogue (2010) identified twelve archetypal dark patterns including trick questions, sneak-into-basket, roach motel, privacy zuckering, misdirection, and disguised ads [1]. This taxonomy was influential but largely descriptive and anecdotal. Gray et al. (2018) extended this framework, grounding dark patterns in broader theories of persuasion, coercive design, and deceptive systems, introducing a three-level ontology distinguishing strategies, patterns, and instances [2]. Mathur et al. (2019) conducted the first large-scale automated study, scanning 11,000 shopping websites to identify 1,818 instances of dark patterns [3]. Their work was pioneering but limited to e-commerce. Luguri and Strahilevitz (2021) conducted the first randomised controlled trial, finding that aggressive dark patterns increased subscription uptake by over 200% [4].

b) Psychological and Cognitive Dimensions

The effectiveness of dark patterns rests on well-documented cognitive biases. Kahneman's dual-process theory provides a key explanatory framework: dark patterns exploit System 1 (fast, automatic) processing while undermining System 2 (deliberate, analytical) oversight [10]. Specific mechanisms include anchoring bias (decoy pricing), loss aversion (countdown timers), default effect (pre-selected options), and the principle of least effort (asymmetric opt-in versus opt-out). Nouwens et al. (2020) demonstrated that dark-pattern consent mechanisms reduced privacy-preserving choices by up to 23 percentage points [5]. Utz et al. (2019) similarly showed that interface design of cookie banners significantly determines consent withdrawal rates [6], underscoring that apparent free choices may be architecturally determined.

c) Legal and Regulatory Landscape

The EU GDPR implicitly prohibits many dark patterns by requiring consent to be freely given, specific, informed, and unambiguous [20]. The UK ICO [17], the US FTC [12], and the EU Digital Services Act have further tightened constraints. Enforcement actions against Google, Meta, and Amazon have resulted in significant fines, though cross-jurisdictional enforcement challenges remain substantial.

d) Automated Detection Approaches

Soe et al. (2020) applied machine learning to cookie banner HTML and CSS features, achieving 89% accuracy [7]. Mehrnezhad et al. (2022) developed a browser extension employing NLP to flag misleading interface text [8]. Chen et al. (2023) applied computer vision models to detect visual dark patterns [9]. Our work builds on this body of research by combining DOM analysis, text classification, visual detection, and dynamic behavioural flow tracing in a unified pipeline.

3. Comprehensive Taxonomy of Dark Patterns

a) Taxonomic Design Principles

Our taxonomy is structured around three design principles: (i) mechanistic clarity each category is defined by the specific mechanism of manipulation employed; (ii) non-overlap categories are mutually exclusive at the pattern level; and (iii) exhaustiveness the taxonomy aspires to cover all known and emerging dark patterns. We organise 27 specific patterns into five top-level categories.

b) Category Descriptions

1) Interface Interference

Interface interference patterns exploit visual design decisions colour, size, placement, contrast, and wording to guide users toward choices they would not otherwise make. Confirm shaming frames the refusal option using guilt-inducing language (e.g., 'No thanks, I don't want to save money'). Misdirection draws attention to a promoted option while visually

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0301
e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 2071 – 2080

suppressing alternatives. Toying with emotion employs imagery and copy to create anxiety, desire, or social pressure. As shown in Table 1.

Table 1 Five-Category Taxonomy of Dark Patterns with Prevalence Rates

Category	Description	Sub-patterns	Prev.(%)	Severity
Interface Interference	Manipulates visual salience, layout or affordances	Misdirection, Confirmshaming, Toying with Emotion	67.4	High
Obstruction	Erects unnecessary friction around user-beneficial actions	Roach Motel, Hard-to-Cancel, Trick Questions	58.9	Very High
Sneaking	Conceals information or adds items without disclosure	Hidden Costs, Sneak into Basket, Bait & Switch	54.1	High
Urgency & Scarcity	Creates false time pressure or resource scarcity	Countdown Timers, Low Stock, Activity Notifications	72.6	Medium
Forced Action	Requires actions users would not otherwise choose	Forced Enrollment, Privacy Zuckering, Disguised Ads	49.3	Very High

2) Obstruction (Roach Motel)

Obstruction patterns make it disproportionately easy to enter a state subscribe, register, accept and extremely difficult to exit. Empirical analysis revealed that cancellation flows for online subscription services averaged 7.3 steps, versus 2.1 steps for the corresponding sign-up flows [3]. Hard-to-cancel patterns frequently involve mandatory telephone calls, extended retention scripts, and deliberately confusing UI states.

3) Sneaking

Sneaking patterns introduce charges, items, or conditions without explicit user awareness. Hidden costs appear late in a checkout process after user commitment; sneak-into-basket automatically pre-adds items that must be actively removed; bait-and switch substitutes a promoted product for an inferior one after initial engagement. These patterns have attracted the greatest volume of regulatory enforcement action.

4) Urgency and Scarcity

Urgency and scarcity patterns manufacture time pressure or resource limitation to accelerate decision-making. Countdown timers may strategically reset on page reload. A 2024 UK CMA investigation found that 41% of hotel booking sites displayed false scarcity notifications [13]. Activity notifications ('14 people are looking at this right now') are similarly unreliable and exploit loss-aversion bias.

5) Forced Action

Forced action patterns eliminate optionality,

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0301
e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 2071 – 2080

requiring users to perform actions creating accounts, accepting tracking as a precondition for accessing content or services. Privacy zuckering, associated with Meta's early privacy settings, used complexity to prevent users from exercising meaningful privacy preferences [11]. Disguised advertisements have significant documented effects on click-through behaviour.

4. Detection Methodology

a) DarkScan Architecture

We designed and implemented DarkScan, a multi-modal automated detection pipeline for web-based dark patterns. The system operates in four functional layers: (1) crawling and DOM extraction, (2) structural and textual feature analysis, (3) visual pattern detection, and (4) dynamic interaction flow tracing. Each layer feeds a unified classification engine that assigns pattern-type labels and severity scores.

b) B. Corpus Collection

The corpus comprised 1,200 web applications sampled across five vertical categories: (i) e-commerce and retail (n=320), (ii) SaaS and subscription software (n=240), (iii) social media and content platforms (n=200), (iv) news and media (n=180), and (v) financial services (n=260). Sites were selected by combining Alexa top-10,000 rankings, Common Crawl data, and purposive sampling of platforms known to employ aggressive conversion optimisation. Sampling was stratified by geographic region including EU, US, UK, and Asia-Pacific deployments.

c) Feature Extraction

The DOM-analysis module extracts 147 structural, semantic, and stylistic features. Textual features include sentiment polarity of call-to-action (CTA) text, presence of guilt-inducing or loss-framing lexical patterns (from a domain-specific lexicon of 2,340 terms), readability scores of consent text, and relative prominence of opt-in versus opt-out options. Structural features include visual salience ratios, DOM depth of cancellation affordances, number of steps in key flows, and pre-selection state of consent checkboxes.

d) Classification Engine

Classification employs an ensemble model combining a fine-tuned BERT language model for text classification, a ResNet-50 convolutional network for visual feature classification, and a gradient-boosted decision tree (XGBoost) for structural feature classification. Ensemble outputs are weighted using stacked meta-learning. The model was trained on 4,800 dark-pattern instances and 6,200 non-dark-pattern interface elements, annotated by three independent UX researchers with inter-rater reliability (Cohen's kappa) of 0.84.

5. User Study Design

a) Study Overview

To complement automated analysis with human-centred data, we conducted a controlled user experiment using a between-subjects design. Three hundred and twenty participants (mean age 31.4, SD=9.7; 52.2% female, 45.3% male, 2.5% non-binary/prefer not to say) were recruited via a research panel and randomised to one of five experimental conditions, each exposing participants to a different dark-pattern category on a purpose-built simulated e-commerce platform.

b) Experimental Conditions

Five experimental conditions were constructed: (C1) neutral control with no dark patterns; (C2) urgency/scarcity manipulation (countdown timers, low-stock indicators); (C3) interface interference (confirm shaming, misdirection); (C4) obstruction (complex cancellation flow, nagging notifications); and (C5) sneaking (hidden costs, sneak-into-basket). Participants completed four tasks: product discovery, checkout, account settings modification, and subscription cancellation.

c) Outcome Measures

Primary outcome measures were: (i) task completion rate for intended user goals; (ii) rate of unintended actions (purchases, subscriptions, consents); (iii) task completion time; and (iv) post-session scores on the User Experience Questionnaire (UEQ) and a four-item trust scale. Secondary measures included think-aloud protocols coded for frustration, confusion, and recognition of manipulation.

6. Results

a) Automated Detection Performance

As Shown in Table 2.

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0301
e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 2071 – 2080

Table 2 Darkscan Classification Performance by Pattern Category

Dark Pattern Category	Prec.(%)	Recall (%)	F1	Instances	FPR (%)
Interface Interference	93.2	91.4	0.923	2,841	6.8
Obstruction	90.1	87.6	0.888	1,947	9.9
Sneaking	94.7	89.3	0.919	2,103	5.3
Urgency & Scarcity	88.4	86.1	0.872	3,287	11.6
Forced Action	91.8	90.2	0.910	1,562	8.2

b) Prevalence Across Application Verticals

As Shown in Table 3.

Table 3 Dark Pattern Prevalence by Application Vertical (% Of Sites with ≥1 Instance)

Vertical	Int. Inf.	Obstr.	Sneak.	Urgency	Forced	Any
E-Commerce & Retail	71.3%	64.8%	68.4%	82.1%	52.6%	78.4%
SaaS / Subscription	69.2%	74.6%	58.3%	61.7%	63.8%	65.2%
Social Media	59.0%	47.5%	34.0%	52.5%	71.5%	54.7%
News & Media	61.1%	44.4%	38.9%	58.3%	55.6%	48.3%
Financial Services	53.8%	61.5%	46.2%	46.2%	38.5%	43.5%

c) User Study Findings

As shown in Table 4.

Table 4 User Study Behavioural Outcomes by Experimental Condition * p < 0.01 versus control (ANOVA with post-hoc Tukey HSD)

Condition	Unint. Actions(%)	Task Comp.(%)	Time(sec)	UEQ Score	Trust(/10)
C1 – Control (No DP)	4.3	91.2	84.7	+1.84	**7.9**
C2 – Urgency/Scarcity	28.6 *	79.4 *	102.3 *	+0.71 *	**5.8 ***
C3 – Interface Interf.	34.7 *	74.1 *	118.6 *	+0.43 *	**4.9 ***
C4 – Obstruction	22.1 *	68.7 *	247.4 *	−0.22 *	**3.4 ***
C5 – Sneaking	39.8 *	71.3 *	131.8 *	+0.17 *	**4.1 ***

The obstruction condition (C4) produced the greatest absolute increase in task completion time nearly three minutes more than the control reflecting the intentional friction built into cancellation flows. The

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0301
e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 2071 – 2080

sneaking condition (C5) produced the highest rate of unintended monetary actions, with a mean unintended spend of £12.40 per participant. The trust impact was both immediate and persistent: a two-week follow-up survey (n=198, 61.9% follow-up rate) found that trust scores remained significantly lower than control across all dark-pattern conditions ($p < 0.05$).

d) Qualitative Findings

Think-aloud protocols and post-session interviews were coded using thematic analysis. Four principal themes emerged. First, users frequently recognised that something felt 'off' but could not articulate the specific mechanism a phenomenon we term 'vague manipulation awareness.' Second, users reported significant frustration during obstruction-pattern interactions, using language indicative of helplessness and loss of control. Third, approximately 18% appeared partially inoculated against urgency patterns from prior negative experiences yet even this group showed significantly elevated unintended action rates. Fourth, confirm shaming provoked particularly strong negative emotional responses including feelings of shame, anger, and condescension.

7. Discussion

a) Implications for Design Practice

Our findings reinforce a growing body of evidence that dark patterns are not merely a fringe phenomenon but a systemic feature of the commercial web. The 78.4% prevalence rate in e-commerce suggests that deceptive design has become a near-universal competitive norm, likely driven by documented short-term conversion benefits. Our results argue for a shift toward user centred metrics including trust retention, user task success rate, and long-term engagement as a counterweight to pure conversion optimisation. We propose four design principles for dark-pattern-free interfaces: (1) Symmetric optionality making opting in and opting out equally easy and equally prominent; (2) Transparent information architecture ensuring all relevant information is presented before commitment; (3) Honest scarcity representing stock levels and time constraints accurately; and (4) Graceful exit designing cancellation flows to be no more complex than sign-up flows.

b) Implications for Automated Auditing

DarkScan's overall F1 score of 0.902 represents a significant improvement over prior automated detection tools. However, dynamic dark patterns that only appear following specific user interaction sequences are under detected by static DOM analysis. Our dynamic flow tracer addresses this partially. Additionally, certain dark patterns are highly context-dependent; a countdown timer may be legitimate or deceptive depending on whether the underlying scarcity is real.

c) Regulatory Implications

Regulatory frameworks have evolved considerably, but enforcement effectiveness remains constrained by technical complexity, jurisdictional fragmentation, and the pace of innovation in deceptive design. We recommend three regulatory improvements. First, regulators should adopt automated tools such as DarkScan for systematic compliance screening. Second, mandatory interface disclosure requirements should require companies to publish their conversion optimisation methods. Third, safe harbour protections should be extended to companies participating in voluntary certification programmes.

d) Emerging Patterns and Future Threats

Our analysis identified several emerging dark-pattern modalities not covered by existing taxonomies. AI-generated personalised persuasion uses large language models to dynamically generate persuasive copy targeted to individual users' psychological profiles. Algorithmic choice architecture deploys reinforcement learning to discover and exploit the most effective dark pattern for each user segment. Dark patterns in virtual and augmented reality environments introduce new modalities spatial manipulation, presence exploitation, and attention hijacking that existing detection methods are ill-equipped to address.

8. Dark Pattern Severity Scoring Model

To operationalise harm assessment, we developed a Dark Pattern Severity Score (DPSS), a composite metric ranging from 0 to 100 computed from five weighted dimensions. **The formula is:**

$$\mathbf{DPSS = 0.25 \cdot H + 0.20 \cdot I + 0.20 \cdot R + 0.20 \cdot V + 0.15 \cdot D}$$

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0301
e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 2071 – 2080

Where: H = Harm potential (financial, psychological, or privacy harm); I = Intentionality (evidence the design choice was deliberate); R = Reversibility (how easily the user can undo the action); V = Vulnerability amplification (degree to which the pattern targets vulnerable populations); D = Deception depth (layers of interface needed to identify the manipulation). As Shown in Table 5.

Table 5 Dark Pattern Severity Scores (DPSS) For Selected Common Patterns

Dark Pattern	H	I	R	V	D	DPSS
Hidden Subscription Charges	90	85	60	70	75	77.0
Privacy Zuckering	85	90	40	80	80	75.0
Roach Motel / Hard-to-Cancel	75	95	30	65	85	71.5
Confirmshaming	50	95	80	75	60	69.5
Sneak into Basket	65	90	55	55	70	67.0
False Countdown Timer	60	85	75	60	45	65.8
Misdirection (Visual)	55	85	65	55	55	63.0
Forced Account Creation	55	80	60	50	50	59.0

Patterns scoring DPSS ≥ 70 are classified as High Severity and warrant immediate remediation; those scoring 50–69 are Medium Severity and should be reviewed; those below 50 are Low Severity. This tiering facilitates prioritisation of both internal remediation efforts and external enforcement action.

9. Illustrative Case Studies

a) Case Study 1: E-Commerce Hidden Subscription

Platform X (anonymised), a mid-tier UK-based fashion retailer, implements a hidden subscription dark pattern at checkout. A £7.99 monthly 'VIP membership' is pre-selected via a checkbox formatted with low-contrast grey text against a white background, positioned between two high-salience elements (a promotional banner and the 'Place Order' CTA). The checkbox label text scores 62 on the Flesch Reading Ease scale borderline difficult and the membership terms are accessible only via a three-level navigation to a separate terms page. DPSS analysis yields a score of 78.2 (High Severity). User study participants exposed to this pattern had an unintended subscription rate of 47.3% versus 6.2% for a redesigned neutral version. The pattern was reported to the UK CMA and later served as a test case for the authority's 2024 guidance on subscription traps [13].

b) Case Study 2: Cookie Consent Dark Pattern

A major European news aggregator (anonymised) implements a sophisticated multi-layer consent dark pattern. The initial consent banner presents two visible options: 'Accept all' (large blue button) and 'Manage preferences' (small grey link text). The 'Manage preferences' pathway requires 11 interactions to achieve equivalent opt-out to the single-click 'Accept all' pathway. Furthermore, 38 of 47 individual tracking purposes are pre-enabled with opt-out requiring individual toggle interactions, while no corresponding 'Disable all' button exists. This consent architecture violates GDPR Article 7 requirements for freely-given consent and EDPB guidelines on cookie banners [11]. DPSS score: 74.8 (High Severity). Remediation recommendations include a visible 'Reject all' button at banner level and a simplified preference centre requiring no more than four interactions to complete.

c) Case Study 3: SaaS Roach Motel Cancellation Flow

A SaaS project management tool (anonymised) presents one of the most elaborate obstruction patterns in the corpus. The subscription cancellation pathway comprises: (1) navigation to account settings (not linked from main navigation); (2) scrolling past 14 sections to reach 'Billing'; (3) locating 'Cancel subscription' in a collapsed accordion; (4) a modal presenting three 'pause' or 'downgrade' alternatives; (5) a mandatory retention survey; (6) a 'final offer' page with a time-limited discount; (7) a confirmation

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0301
e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 2071 – 2080

input requiring manual text entry of 'CONFIRM CANCEL'; and (8) a final confirmation email with a seven-day cooling-off re-activation link. In contrast, the sign-up pathway comprises three steps and approximately 90 seconds. Average cancellation time was 18 minutes 42 seconds, with a 31.4% abandonment rate. DPSS: 80.4 (High Severity — highest in the corpus for this category).

10. Recommendations

a) For Designers and Developers

- Adopt ethical review processes that specifically assess interface designs for dark patterns prior to deployment, analogous to security or accessibility reviews.
- Implement symmetric design principles ensuring that all user journeys especially opting out, cancelling, or exercising privacy preferences are no more complex than their inverses.
- Use plain language in all calls to action, consent notices, and terms presentations. Avoid guilt-framing, loss-framing, or manipulative emotional appeals in interface copy.
- Conduct regular third-party audits using automated tools such as DarkScan as part of continuous integration pipelines.

b) For Platform Operators and Organisations

- Establish organisational policies that define dark patterns as prohibited design practices and create accountability mechanisms with genuine authority to intervene.
- Align product success metrics to include user trust scores, task success rates, and long-term retention rather than solely short-term conversion rates.
- Publish transparency reports detailing conversion optimisation techniques employed, enabling user awareness and regulatory scrutiny.

c) For Regulators and Policymakers

- Invest in technical capacity to conduct automated dark-pattern audits at scale, rather than relying solely on reactive complaint-driven enforcement.
- Develop harmonised cross-jurisdictional definitions of dark patterns for consistent enforcement across the EU, UK, US, and other major markets.
- Establish mandatory 'dark-pattern-free' certification for high-risk sectors (financial services, children's applications, health platforms) with annual independent verification.
- Require companies to maintain and provide to regulators the full user journey flow documentation for key pathways (checkout, cancellation, consent), with version history.

11. Limitations and Threats to Validity

Several limitations of this study should be acknowledged. First, the corpus, while large, was not exhaustive and may not fully represent the diversity of web applications across all sectors, geographies, and business sizes. Smaller websites, particularly those serving non-English-speaking markets, are likely underrepresented. Second, DarkScan's detection capability is limited to patterns expressible in observable DOM and visual features; design intent cannot be established with certainty from technical signals alone, and some false positives may reflect poor design rather than deliberate deception. Third, the user study employed a simulated e-commerce environment, which may differ in important respects from real world shopping contexts. Ecological validity, while strengthened by realistic scenario design, remains a concern. Fourth, the DPSS model incorporates subjective weighting decisions that reflect the research team's value judgements; alternative weighting schemes may produce different severity rankings. Future work should subject the DPSS weights to Delphi expert elicitation to establish broader consensus.

Conclusion

This paper has presented a comprehensive empirical investigation of dark patterns in modern web applications, encompassing a five-category taxonomy, an automated detection pipeline (DarkScan), a severity scoring model (DPSS), and a controlled user experiment demonstrating measurable behavioural and psychological harm. Our findings confirm that dark patterns are pervasive

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0301
e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 2071 – 2080

affecting the substantial majority of e-commerce and subscription platforms and that their effects on users are significant, including elevated rates of unintended action, reduced task completion, extended completion time, and lasting damage to platform trust. The detection capabilities demonstrated by DarkScan (overall F1: 0.902) establish the practical viability of automated compliance auditing at web scale, offering a valuable tool for regulators, advocacy organisations, and responsible platform operators. The DPSS provides a principled basis for harm prioritisation that can inform both internal design reviews and external enforcement triage.

Ultimately, the persistence and proliferation of dark patterns reflects a systemic failure of incentive alignment: the short-term conversion benefits to platform operators are captured privately while the harms financial, psychological, and privacy-related are distributed across millions of users. Addressing this misalignment requires coordinated action by designers, technologists, organisations, and regulators. As digital interfaces become ever more pervasive and ever more sophisticated in their capacity for personalised persuasion, the urgency of this work can only increase.

Acknowledgements

The authors gratefully acknowledge the guidance and mentorship of Dr. Gurunath G throughout this research. The authors also acknowledge the Department of Master of Computer Applications, Dayananda Sagar College of Arts, Science and Commerce, Bengaluru, for institutional support. No conflicts of interest are declared.

References

[1]. H. Brignull, "Dark patterns: Deception vs. honesty in UI design," in Proc. ACM CHI Workshop on Price & Value in HCI, 2010.

[2]. C. M. Gray, Y. Kou, B. Battles, J. Hoggatt, and A. L. Toombs, "The dark (patterns) side of UX design," in Proc. ACM CHI Conf. Human Factors Comput. Syst., Montreal, Canada, 2018, pp. 1–14.

[3]. A. Mathur, G. Acar, M. J. Friedman, E. Lucherini, J. Mayer, M. Chetty, and A. Narayanan, "Dark patterns at scale: Findings from a crawl of 11K shopping websites," Proc. ACM Human-Comput. Interact., vol. 3, no. CSCW, pp. 1–32, Nov. 2019.

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

[5]. [M. Nouwens, I. Liccardi, M. Veale, D. Karger, and L. Kagal, "Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence," in Proc. ACM CHI Conf. Human Factors Comput. Syst., Honolulu, HI, 2020.

[6]. S. Utz, F. Krupp, M. Smits, and J. van der Berg, "(Un)informed consent: Studying GDPR consent notices in the field," in Proc. ACM CCS, London, UK, 2019, pp. 973–990.

[7]. T. H. Soe, O. E. Nordberg, F. Halvorsrud, M. Stumph, and G. Samdal, "Circumvention by design – dark patterns in cookie consent interfaces," in Proc. ECIS, Marrakech, Morocco, 2020.

[8]. M. Mehrnezhad, J. Coopamootoo, and T. Groß, "Defining the boundaries of dark patterns and their relation to deception," in Proc. ACM NordiCHI, Aarhus, Denmark, 2022.

[9]. Y. Chen, N. Sambasivan, and J. Cranshaw, "Detecting deception at scale: A vision-based approach to dark pattern identification," in Proc. ACM CHI Conf. Human Factors Comput. Syst., Hamburg, Germany, 2023.

[10]. D. Kahneman, Thinking, Fast and Slow. New York, NY: Farrar, Straus and Giroux, 2011.

[11]. European Data Protection Board (EDPB), "Guidelines 03/2022 on Dark Patterns in Social Media Platforms," Brussels, 2022.

[12]. Federal Trade Commission (FTC), "Bringing Dark Patterns to Light: An FTC Staff Report," Washington, DC, Sep. 2022.

[13]. UK Competition and Markets Authority (CMA), "Subscription Traps: Consumer Research and Regulatory Guidance," London, Feb. 2024.

[14]. A. Narayanan, A. Mathur, M. Chetty, and M. Kshirsagar, "Dark patterns: Past, present, and future," Queue, vol. 18, no. 2, pp. 67–92, 2020.

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0301
e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 2071 – 2080

[15]. R. Bösch, B. Erb, F. Kargl, H. Kopp, and S. Pfattheicher, "Tales from the dark side: Privacy dark strategies and privacy dark patterns," Proc. Priv. Enhanc. Technol., vol. 2016, no. 4, pp. 237–254, 2016.
[16]. J. Fogg, Persuasive Technology: Using Computers to Change What We Think and Do. Burlington, MA: Morgan Kaufmann, 2003.
[17]. Information Commissioner's Office (ICO), "How to make cookie banners lawful and good for users," Wilmslow, UK, Jan. 2023.
[18]. A. Acquisti, I. Brandimarte, and G. Loewenstein, "Privacy and human behavior in the age of information," Science, vol. 347, no. 6221, pp. 509–514, Jan. 2015.
[19]. A. Mathur, J. Mayer, and M. Chetty, "What makes a dark pattern... dark? Design attributes, normative considerations, and measurement," in Proc. ACM CHI Conf. Human Factors Comput. Syst., Yokohama, Japan, 2021.
[20]. S. Voigt and A. von dem Bussche, The EU General Data Protection Regulation (GDPR): A Practical Guide, 2nd ed. Cham, Switzerland: Springer, 2022.