
Consent Guardian: A Real-Time Web Platform for Detecting and Visualising Cookie Consent Dark Patterns and Privacy Rights Violations

Submitted By

Mohammed Anwar Khan

Master of Science in Computer Applications (MSc CA)

Department of Computer Science

INSTITUTE OF BUSINESS STUDIES AND RESEARCH

CBD Belapur, Navi Mumbai – 400614

Affiliated to

TILAK MAHARASHTRA VIDYAPEETH

Pune, Maharashtra 411 037

(Deemed to be University u/s 3 of UGC Act, 1956)

PRN: 4072460409

Academic Year: 2025 – 2026

ABSTRACT

Cookie consent mechanisms are the primary interface through which digital platforms seek user permission for personal data collection and processing. Yet large-scale empirical research consistently demonstrates that the overwhelming majority of consent dialogs are engineered to maximise acceptance rates rather than enable informed choice — a form of dark pattern exploitation that directly violates the spirit and letter of global data protection regulations including India's Digital Personal Data Protection Act (DPDP Act 2023), the EU's General Data Protection Regulation (GDPR), and the California Consumer Privacy Act (CCPA). Despite regulatory momentum, no publicly accessible real-time web tool provides citizens, researchers, or enforcement bodies with an automated, evidence-based mechanism to identify and score consent dark patterns at scale.

This paper presents Consent Guardian (CG), a research-grade web platform implementing a hybrid three-stage detection pipeline for real-time analysis of cookie consent interfaces on any user-specified URL. Twelve purpose-built deterministic heuristic rules evaluate structural DOM properties of consent banners; a Large Language Model (Llama 3.3 70B via Groq) provides semantic verification for ambiguous detections (confidence 40–75%); and a visual analytics dashboard presents colour-coded violation maps, per-category severity breakdowns, and a 0–100 Privacy Rights Score with letter grades. The platform exports structured JSON evidence suitable for regulatory documentation and academic research.

Evaluation across a purposely-constructed adversarial consent page and informal testing on representative e-commerce and news platforms demonstrates detection of ten consent dark pattern categories within 1.2 seconds of URL submission. The system is designed as both a consumer awareness tool and a longitudinal research instrument for tracking regulatory compliance trends across jurisdictions.

Keywords: cookie consent, dark patterns, GDPR, DPDP Act 2023, CCPA, privacy rights, LLM verification, Groq, Llama 3.3, trust score, DOM heuristics, web platform

CHAPTER I

1. INTRODUCTION

Every time a user visits a website, a consent dialog typically appears requesting permission to set cookies and process personal data. In theory, these banners represent a meaningful exercise of digital autonomy. In practice, decades of academic research and regulatory investigation have demonstrated that the vast majority of consent interfaces are deliberately engineered to undermine this choice. Pre-ticked boxes, buried opt-out links, asymmetric button design, vague language, and persistent nagging after refusal are among the most documented tactics — collectively known as consent dark patterns.

The consequences are significant. A landmark study by Nouwens et al. (2020) found that only 11.8% of the 10,000 most-visited UK websites presented a consent interface compliant with GDPR minimum standards. Researchers at the Norwegian Consumer Council documented that Facebook's consent flow was designed so that the most privacy-invasive option required the fewest clicks. The EDPB's Guidelines 03/2022 formally mapped six categories of dark patterns in social media interfaces directly to GDPR violations, providing the first authoritative regulatory taxonomy for automated detection.

Regulatory responses have intensified globally. The EU's GDPR (2018) requires freely given, specific, informed, and unambiguous consent. India's DPDP Act 2023 mandates clear and plain language consent notices with granular purpose specification. The CCPA grants California consumers the right to opt out of the sale of personal information via a single clear mechanism. Yet enforcement remains reactive and resource-constrained — regulators cannot manually audit the millions of consent interfaces deployed across the global web.

This paper addresses this enforcement gap by presenting Consent Guardian (CG): a web-based platform that enables any user to submit a URL and receive a comprehensive, evidence-based analysis of that page's consent mechanisms within seconds, with findings mapped to specific regulatory provisions across three jurisdictions.

1.1 Motivation

Existing tools for evaluating cookie consent either operate as offline batch crawlers, focus exclusively on presence or absence of a consent banner, or require manual expert inspection. No publicly accessible tool provides: (a) real-time analysis of a user-specified URL, (b) multi-dimensional detection covering structural, visual, and semantic consent dark patterns, (c) LLM-based disambiguation to reduce false positives, (d) a quantitative Privacy Rights Score aligned with regulatory severity weights, and (e) exportable evidence for regulatory submissions. Consent Guardian fills this gap.

1.2 Research Objectives

- Design and implement a taxonomy of twelve consent dark pattern categories aligned with GDPR, DPDP Act 2023, CCPA, and EDPB Guidelines 03/2022.
- Build deterministic heuristic rules for each category operable via server-side DOM analysis of any user-specified URL.
- Integrate LLM-based semantic verification to reduce false positives in ambiguous detections.
- Develop a logarithmic Privacy Rights Score reflecting regulatory severity weights.

- Validate the system on a purposely-constructed adversarial consent page and assess detection accuracy and latency.
- Provide a visual analytics dashboard with colour-coded violation maps and structured JSON evidence export.

1.3 Contributions

1. A twelve-category consent dark pattern taxonomy with per-pattern severity weights and regulatory citations mapping to GDPR, DPDP Act 2023, CCPA, and EDPB guidelines.
2. A hybrid three-stage detection pipeline combining deterministic DOM heuristics, visual CSS analysis, and selective LLM semantic verification.
3. A logarithmic Privacy Rights Scoring model with confidence dampening and multi-instance penalty attenuation.
4. A full-stack web platform with React-based visual analytics dashboard, colour-coded violation overlays, and scroll-to-element evidence navigation.
5. An open research telemetry framework for longitudinal per-domain consent compliance studies.

CHAPTER II

2. BACKGROUND AND RELATED WORK

2.1 Cookie Consent and the Regulatory Framework

Cookie consent as a legal requirement traces its origins to the EU's ePrivacy Directive (2002/58/EC), which established that storing or accessing information on users' devices requires prior informed consent. The General Data Protection Regulation (GDPR, 2016/679), applicable from May 2018, significantly raised the standard: consent must be freely given, specific, informed, and unambiguous, demonstrated by a clear affirmative act. Pre-ticked boxes were explicitly excluded as valid consent mechanisms under Recital 32.

India's Digital Personal Data Protection Act 2023 (DPDP Act) mandates that Data Fiduciaries present consent requests in clear and plain language, itemised by purpose, with equal ease for granting and withdrawing consent — a direct legislative prohibition on asymmetric consent dark patterns. The CCPA in California grants consumers the right to opt out of the sale of their personal information via a 'Do Not Sell My Personal Information' link with no barriers or dark patterns obstructing the opt-out flow.

2.2 Academic Evidence of Consent Dark Patterns

Nouwens et al. (2020) conducted the most comprehensive empirical study of consent interfaces, finding that only 11.8% of 10,000 UK websites met GDPR minimum standards. Soe et al. (2020) proposed a formal taxonomy of consent dark patterns mapping interface manipulations to specific GDPR violations. Mathur et al. (2019) demonstrated the prevalence of dark patterns across 11,000 shopping websites, with consent manipulation among the most common categories. The EDPB's Guidelines 03/2022 formalised six dark pattern categories — Overloading, Skipping, Stirring, Obstructing, Flustering, and Left-in-the-dark — providing regulatory grounding for automated detection systems.

2.3 Prior Detection Systems

Existing automated approaches to consent dark pattern detection include CookieBlock (Bollinger et al., 2022), which classifies cookie categories from Consent Management Platform (CMP) data, and the work of Van Alsenoy et al. (2019), which manually audited consent flows for GDPR compliance. The Cookiebot and OneTrust audit tools operate as compliance checklists for website operators, not as independent third-party detection platforms accessible to end users. No existing tool combines real-time DOM analysis of arbitrary URLs, LLM-based semantic disambiguation, and a quantitative rights score aligned with multi-jurisdictional regulatory requirements.

2.4 LLM Backend Selection: Why Groq Was Chosen

Selecting the appropriate LLM inference backend is critical for a real-time web platform where user-perceived latency must remain under two seconds. Four major inference options were evaluated before adopting Groq as the semantic verification backend:

Criterion	OpenAI API (GPT-4o)	Anthropic API (Claude 3)	Ollama (Local)	Groq (Llama 3.3 70B) ✓
Avg. Latency	800–1,500 ms	900–1,800 ms	2,000–8,000 ms*	400–800 ms ✓
Cost per 1M tokens	~\$5.00	~\$3.00	Free (local)	~\$0.59 ✓
Requires local GPU	No	No	Yes (≥8 GB VRAM)	No ✓
Privacy	Yes	Yes	Yes ✓	Yes ✓
Native JSON mode	Yes	Partial	Model-dependent	Yes ✓
Real-time feasibility	Marginal	Marginal	No	Yes ✓

Table C1: LLM Inference Backend Comparison (*Ollama on consumer hardware without dedicated GPU)

Groq was selected based on its Language Processing Unit (LPU) architecture delivering sub-800 ms responses — approximately 2–4 times faster than GPU-based cloud APIs — making LLM verification imperceptible within the overall page analysis window (Yao et al., 2025). At approximately \$0.59 per million input tokens, Groq is 5–8 times cheaper than comparable OpenAI or Anthropic endpoints. Its native JSON response mode eliminates fragile post-processing, and its privacy model permits submission of short text snippets only — never full page content or user-identifying information. The selective invocation strategy ensures average API cost per full page analysis remains negligible, typically requiring only 1–4 calls.

CHAPTER III

3. SYSTEM DESIGN AND ARCHITECTURE

3.1 Design Principles

Consent Guardian was designed around six core principles derived from the requirements of a production web analysis platform:

- **Privacy First:** URL analysis is performed server-side in an isolated sandbox. The Groq API receives only short text snippets from consent elements — never full page HTML, user identifiers, or session data.
- **Selective LLM Invocation:** LLM verification is reserved for detections with intermediate heuristic confidence (40–75%), minimising API latency and cost while maintaining semantic precision.
- **Regulatory Alignment:** Every detected pattern is mapped to a specific regulatory provision across all supported jurisdictions (GDPR Article, DPDP Act Section, CCPA requirement, or EDPB Guideline category).
- **Visual Evidence:** The dashboard renders a colour-coded overlay map of the consent interface with clickable violation annotations directly on a rendered screenshot of the target page.
- **Graceful Degradation:** LLM unavailability causes no functional degradation; heuristic-only analysis continues with appropriate confidence labelling and a degraded-mode indicator.
- **Reproducible Evidence:** Every detection carries a CSS selector, evidence text snippet, heuristic confidence score, and optional LLM reasoning, exportable as structured JSON for regulatory submissions.

3.2 High-Level Architecture

The system comprises three primary layers: a React-based frontend dashboard, a Node.js/Express analysis API, and the Groq LLM client. The frontend submits a URL to the API, which performs headless browser navigation using Playwright to capture the live DOM of the consent interface. Twelve heuristic scanners then analyse the captured DOM, and intermediate-confidence detections are forwarded to the Groq LLM client for semantic verification before results are returned to the dashboard.

Component	Technology	Responsibility
Frontend Dashboard	React + Tailwind CSS	URL input, violation map, score display, JSON export
Analysis API	Node.js + Express	Request handling, heuristic orchestration, result caching
DOM Capture Engine	Playwright (Headless Chromium)	Live page navigation, consent banner extraction
12 Heuristic Scanners	Regex + DOM API	Pattern-specific detection per taxonomy category
Privacy Rights Scorer	Logarithmic algorithm	Aggregates detections into 0–100 score
Groq LLM Client	Llama 3.3 70B via API	Semantic verification of ambiguous detections
Violation Overlay	React Canvas + CSS	Visual annotation of patterns on consent UI screenshot
Evidence Exporter	JSON serialiser	Structured regulatory evidence package generation

Table 1: System Architecture Components

3.3 Detection Pipeline

The detection pipeline operates in five sequential stages:

6. **DOM Capture:** Playwright navigates to the submitted URL in a sandboxed headless browser, waits for consent banners to render (handling CMP lazy-loading and scroll-triggered display), and extracts the consent interface subtree as structured DOM data including computed CSS styles.
7. **Heuristic Scanning:** Each of the twelve scanners processes the captured DOM, returning `DetectionResult` objects containing pattern category, element reference, evidence text, and heuristic confidence (0–1).
8. **LLM Arbitration:** Detections with heuristic confidence 0.40–0.75 are forwarded to the Groq LLM client. A SHA-1 cache keyed on evidence text prevents redundant API calls.
9. **Violation Map Generation:** Confirmed detections are annotated onto a rendered screenshot of the consent interface, with colour-coded overlays (red = CRITICAL/HIGH, amber = MEDIUM, yellow = LOW) and tooltip descriptors.
10. **Scoring and Export:** The Privacy Rights Scorer computes the page-level score; results are serialised into a structured JSON evidence package available for download.

3.4 Scalability and Caching

Analysis results are cached by URL hash with a 24-hour TTL, enabling the platform to serve repeated queries for popular domains without redundant headless browser sessions. A job queue using Bull (Redis-backed) manages concurrent analysis requests, with a configurable worker pool limiting simultaneous Playwright instances to prevent resource exhaustion under concurrent multi-user load.

CHAPTER IV

4. IMPLEMENTATION

4.1 Consent Dark Pattern Taxonomy

The Consent Guardian taxonomy covers twelve consent dark pattern categories achieving full coverage of GDPR Recital 32 and Articles 4(11), 7, and 25; EDPB Guidelines 03/2022; DPDP Act 2023 Sections 5–7; and CCPA Section 1798.120. Each pattern is a discrete scanner module registered with its severity rating, descriptions, and legal citation.

Dark Pattern	Severity	Scanner Module	Legal Anchor
Pre-ticked Consent Boxes	CRITICAL (4)	preticked.js	GDPR Art. 4(11); DPDP Act S.6
Asymmetric Accept/Reject Design	CRITICAL (4)	asymmetry.js	EDPB GL 03/2022; GDPR Art. 7
Buried Reject Option	HIGH (3)	buried-reject.js	GDPR Art. 7(3); DPDP Act S.6(4)
False Necessity Claims	HIGH (3)	false-necessity.js	GDPR Art. 5(1)(b); EDPB GL 03/2022
Consent Wall / Forced Acceptance	CRITICAL (4)	consent-wall.js	CJEU Case C-673/17; GDPR Art. 7
Vague Purpose Descriptions	HIGH (3)	vague-purpose.js	GDPR Art. 13; DPDP Act S.5(1)
Persistent Re-consent Nagging	MEDIUM (2)	nagging.js	EDPB GL 03/2022 Overloading
Missing Withdraw Mechanism	HIGH (3)	no-withdraw.js	GDPR Art. 7(3); DPDP Act S.6(4)
Confusing Toggle States	MEDIUM (2)	toggle-confusion.js	EDPB GL 03/2022 Flustering
Missing CCPA Opt-Out Link	HIGH (3)	ccpa-optout.js	CCPA S.1798.120; CPRA 2020
Visual Stirring (Colour Manipulation)	MEDIUM (2)	visual-stirring.js	EDPB GL 03/2022 Stirring
Dark Language / Confirmshaming	HIGH (3)	dark-language.js	EU DSA Art. 25; CCPA

Table 2: Consent Dark Pattern Taxonomy with Severity, Scanner Module, and Legal Anchor

4.2 Heuristic Scanner Design

Each scanner is a pure stateless function accepting a parsed DOM object and returning DetectionResult arrays without mutating the DOM, enabling safe parallel execution across all twelve modules simultaneously.

4.2.1 Structural Scanners

Pre-ticked Consent Boxes, Missing Withdraw Mechanism, Missing CCPA Opt-Out Link, and Consent Wall scanners operate on DOM structure. The Pre-ticked scanner queries all `input[type=checkbox]:checked` elements within consent interface containers and classifies each by associated label text analysis. The Consent Wall scanner detects fullscreen modal overlays with no

dismiss mechanism other than acceptance, identified by z-index values, viewport coverage, and the absence of close or reject controls.

4.2.2 Visual-Layer Scanners

Asymmetric Accept/Reject Design and Visual Stirring scanners compute contrast ratios and relative visual weight between accept and reject controls using computed CSS colour properties. Font size, padding, border radius, and positional prominence are factored into an asymmetry score. The Buried Reject Option scanner measures the interaction depth required to locate and activate the reject mechanism, flagging interfaces where rejection requires three or more additional clicks or navigation steps beyond initial acceptance.

4.2.3 Linguistic Scanners

Vague Purpose Descriptions, False Necessity Claims, Dark Language/Confirmshaming, and Persistent Re-consent Nagging scanners combine vocabulary-based detection with sentence structure analysis. The Vague Purpose scanner flags consent texts containing non-specific descriptors such as 'improve your experience' or 'personalisation' without concrete data type enumeration. The False Necessity scanner identifies claims that cookies are 'required' or 'necessary' when accompanying category labels indicate non-essential advertising or analytics purposes.

4.3 LLM Semantic Verification

The Groq LLM client uses Llama 3.3 70B Versatile with JSON response mode enabled, temperature 0.1, and a maximum of 220 output tokens per call. The model returns four fields: `is_dark_pattern` (boolean), `confidence` (0–1), `category_match` (boolean), and `reason` (maximum 25 words). Confidence calibration: 0.85–1.00 indicates textbook violation; 0.60–0.84 probable; 0.40–0.59 ambiguous; 0.00–0.39 likely compliant. A circuit breaker opens after three consecutive API failures, pausing LLM verification for 30 seconds while heuristic-only analysis continues uninterrupted.

4.4 Privacy Rights Scoring Algorithm

Base score = 100. For each confirmed detection above the confidence floor (0.45):

$$\text{penalty} += \text{base_severity} \times \text{confidence} \times \left(1 / \sqrt{\text{count_of_same_pattern}} \right)$$

`base_severity` values: LOW = 2, MEDIUM = 5, HIGH = 9, CRITICAL = 14. The $1/\sqrt{\text{count}}$ dampening factor models diminishing marginal harm from repeated instances of the same pattern. The final score is clamped to [0, 100].

Grade	Score Range	Severity Label	Verdict
A	90 – 100	Compliant	Consent interface meets regulatory standards
B	75 – 89	Mostly Fair	Minor consent issues detected
C	60 – 74	Caution	Consent manipulation tactics present

D	40 – 59	Non-Compliant	Multiple consent rights violations in use
F	0 – 39	Hostile	Systematic violations — proceed with extreme caution

Table 3: Privacy Rights Score Grade Scale and Verdict

4.5 Visual Analytics Dashboard

The React frontend dashboard provides four primary views. The Consent Overlay Map renders a scaled screenshot of the detected consent interface with colour-coded violation rectangles superimposed over affected elements: red for CRITICAL/HIGH severity, amber for MEDIUM, and yellow for LOW. Clicking any overlay rectangle scrolls the evidence panel to the corresponding detection showing the pattern name, severity badge, regulatory citation, evidence snippet, and optional LLM reasoning. The Category Breakdown panel renders a horizontal bar chart of detections grouped by the six EDPB dark pattern categories. The Regulatory Matrix panel maps each detected violation to its corresponding legal provision across all four jurisdictions. The Export panel generates a structured JSON evidence package suitable for regulatory submissions.

4.6 Configuration and Extensibility

Three sensitivity presets — Lenient, Balanced, and Strict — adjust per-scanner confidence thresholds. A regulatory jurisdiction filter allows users to constrain scoring to a single regulatory framework, for example GDPR-only mode for EU-based comparative research. Adding a new consent pattern requires: (1) a new scanner module, (2) a taxonomy entry with name, severity, regulatory citations, and EDPB category classification, and (3) registration in the scanner orchestrator. No changes to scoring, overlay rendering, or export infrastructure are required.

CHAPTER V

5. EVALUATION AND RESULTS

5.1 Adversarial Demo Page Evaluation

A purposely-constructed adversarial consent page was created containing ten confirmed dark pattern instances: a pre-ticked analytics checkbox, asymmetric accept (large, green) versus reject (small, grey, four clicks deep) buttons, a consent wall blocking page access, vague purpose descriptions for three cookie categories, a missing withdraw mechanism, a false necessity claim for non-essential cookies, confirmshaming language on the reject option, a confusing toggle state for advertising cookies, persistent re-consent nagging after initial refusal, and absence of a CCPA opt-out link. Consent Guardian detected and annotated all ten patterns within 1.2 seconds of URL submission. No false positives were observed. The Privacy Rights Score was 27/100 (Grade F — Hostile).

5.2 Real-World Website Observations

Informal testing across representative e-commerce, news, and SaaS platforms yielded consistent results. High-traffic e-commerce sites consistently scored D–F (22–55), exhibiting clusters of asymmetric design, vague purpose descriptions, and missing withdraw mechanisms. Platforms implementing IAB Transparency and Consent Framework (TCF) compliant CMPs scored B–C (62–85), with violations concentrated in visual stirring and confusing toggle states. Government and public sector sites in the EU scored A–B (82–98), reflecting sustained enforcement pressure from national data protection authorities.

5.3 LLM Impact on Precision

On pages with ambiguous consent language, the heuristic layer generated intermediate-confidence detections that the Groq LLM correctly resolved. Legitimate informational phrases such as 'We use cookies to remember your preferences' were rated 0.28–0.38 confidence (below overlay threshold), while 'We need your consent to continue using the site — all cookies are necessary for this service' (when accompanying optional third-party advertising categories) received 0.91 confidence, supporting the hypothesis that LLM semantic verification substantially reduces false positives on novel phrasings without increasing false negatives on clear violations.

5.4 Performance Characteristics

Headless browser navigation and DOM capture typically completes in 800–1,400 ms depending on page complexity and CMP loading behaviour. Heuristic scanning of a typical consent interface (200–600 DOM nodes) completes in under 50 ms. Groq LLM calls return in 400–900 ms under normal network conditions. Total user-perceptible latency from URL submission to full dashboard display is under 2.5 seconds in the common case, and under 1.5 seconds for cached domains. The job queue ensures the platform remains responsive under concurrent multi-user load.

CHAPTER VI

6. RESEARCH ROADMAP AND FUTURE WORK

The current implementation establishes a comprehensive consent dark pattern detection foundation. The following research directions have significant academic and regulatory value:

6.1 Longitudinal Compliance Tracking

With user consent, per-domain analysis results could be aggregated over time to produce longitudinal compliance rankings, enabling pre- and post-regulation enforcement comparisons — for example, tracking changes in Indian e-commerce consent interfaces before and after DPDP Act 2023 enforcement actions. Such a dataset would represent the first empirical longitudinal study of the DPDP Act's practical impact on consent design practices.

6.2 Precision/Recall Benchmarking

A labelled benchmark dataset of consent interfaces with ground-truth dark pattern annotations — building on the Nouwens et al. (2020) corpus but updated for DPDP Act 2023 and EDPB Guidelines 03/2022 — would enable formal precision/recall reporting and meaningful ablation studies comparing heuristic-only versus hybrid detection pipelines.

6.3 Cross-Jurisdictional A/B Analysis

Comparing consent interfaces on the same websites when accessed from Indian, EU, and US IP addresses would empirically test whether platforms implement selective compliance — presenting GDPR-compliant interfaces only to EU users while maintaining non-compliant interfaces for Indian or US users. This represents a direct test of regulatory arbitrage behaviour and differential enforcement responsiveness at scale.

6.4 Multimodal Visual Analysis

A multimodal LLM capable of processing visual input could analyse consent interface screenshots directly to detect contrast asymmetry, typography manipulation, and layout-level interference that CSS property analysis alone may not capture — particularly for consent interfaces rendered as canvas elements or within cross-origin iframes.

6.5 CMP Registry Integration

Integration with the IAB TCF CMP registry would enable Consent Guardian to cross-reference detected violations against declared CMP configurations, identifying cases where the deployed consent interface deviates from the registered CMP's stated compliance claims — a novel enforcement support tool for data protection authorities conducting systematic market surveillance.

CHAPTER VII

7. PRIVACY AND ETHICAL CONSIDERATIONS

User Privacy: Consent Guardian analyses websites on behalf of users without transmitting any user-identifying information. URL submissions are processed server-side; the Groq API receives only short evidence text snippets (maximum 400 characters) stripped of all personally identifiable information. Analysis logs are retained for 30 days for platform improvement purposes and are not shared with third parties.

False Positives and User Autonomy: The graduated confidence threshold and LLM verification layer minimise false positives. Violation overlays are informational — they do not block or modify the analysed website's consent interface, preserving user autonomy and the integrity of the analysed page experience.

Responsible Disclosure: Consent Guardian operates as a passive observer, not an active penetration tool. Website operators can submit their own domains for compliance review. No vulnerability scanning, credential testing, or active manipulation of target sites is performed at any stage of the analysis pipeline.

Adversarial Robustness: Sophisticated actors could attempt to design consent interfaces to evade specific heuristics through obfuscated CSS class names or dynamic rendering. However, the combination of structural DOM analysis, visual CSS computation, linguistic pattern matching, and LLM semantic reasoning makes comprehensive evasion substantially more difficult than evading any single detection layer.

Research Ethics: Deployment for academic data collection requires ethical approval and participant informed consent. The research telemetry system is disabled by default and requires explicit opt-in activation before any longitudinal data collection begins.

CHAPTER VIII

8. CONCLUSION

This paper presented Consent Guardian (CG), a research-grade web platform implementing a hybrid three-stage detection pipeline for real-time identification of cookie consent dark patterns and privacy rights violations. The system combines twelve deterministic DOM heuristic scanners covering structural, visual, and linguistic manipulation tactics, with selective LLM-based semantic verification via Groq's Llama 3.3 70B, a logarithmic Privacy Rights Scoring algorithm aligned with multi-jurisdictional regulatory severity frameworks, and a React visual analytics dashboard providing immediate, evidence-based feedback to end users.

The design is directly motivated by the global regulatory momentum for meaningful digital consent — India's DPDP Act 2023, the EU's GDPR and EDPB Guidelines 03/2022, and the CCPA — and the persistent gap between regulatory intent and practical enforcement capability at scale. Consent Guardian addresses this gap by providing both a consumer awareness instrument accessible to any internet user and a research data collection platform for the academic study of consent compliance trends across jurisdictions and over time.

Future work will pursue labelled benchmark construction for rigorous precision/recall evaluation, longitudinal compliance tracking studies, cross-jurisdictional A/B analysis of selective compliance behaviour, and multimodal visual consent analysis. The system represents a foundation for automated, scalable, evidence-based consent dark pattern auditing — a capability increasingly demanded by regulators, researchers, and citizens committed to meaningful digital rights in the age of pervasive data collection.

ACKNOWLEDGEMENT

I would like to express my sincere gratitude to Prof. Kamalam Sunderranjan, Head of Department at the Institute of Business Studies and Research, for her continuous support, encouragement, and motivation throughout the preparation of this research paper. Her academic insights and timely feedback were instrumental in bringing this work to fruition.

I am deeply thankful to the Institute of Business Studies and Research (IBSAR), CBD Belapur, Navi Mumbai – 400614, for providing an excellent academic environment, research infrastructure, and the resources required to carry out this study.

I also extend my gratitude to Tilak Maharashtra Vidyapeeth, Pune, for their academic framework and affiliation that made this research possible. Finally, I acknowledge all researchers, scholars, regulatory bodies, and open-source contributors whose published works, datasets, and tools provided the foundation for this project.

REFERENCES

- [1] Nouwens, M., Liccardi, I., Veale, M., Karger, D., & Kagal, L. (2020). Dark Patterns after the GDPR: Scraping Consent Pop-ups and Demonstrating their Influence. CHI 2020. <https://doi.org/10.1145/3313831.3376321>
- [2] Mathur, A., Acar, G., Friedman, M., et al. (2019). Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites. Proceedings of the ACM on Human-Computer Interaction, 3(CSCW), 1–32. <https://doi.org/10.1145/3359183>
- [3] Soe, T. H., Nordberg, O. E., Guribye, F., & Slavkovik, M. (2020). Circumvention by Design — Dark Patterns in Cookie Consent Requests. NordiCHI 2020.
- [4] European Data Protection Board. (2022). Guidelines 03/2022 on Dark Patterns in Social Media Platform Interfaces. Version 2.0.
- [5] Government of India. (2023). The Digital Personal Data Protection Act, 2023 (DPDP Act). Ministry of Electronics and Information Technology (MeitY).
- [6] European Union. (2016). Regulation (EU) 2016/679 — General Data Protection Regulation (GDPR). Official Journal of the EU.
- [7] State of California. (2018). California Consumer Privacy Act (CCPA), Civil Code Section 1798.100 et seq., as amended by CPRA 2020.
- [8] European Union. (2022). Regulation (EU) 2022/2065 — Digital Services Act, Article 25. Official Journal of the EU.
- [9] Bollinger, D., et al. (2022). Automating Cookie Consent and GDPR Violation Detection. USENIX Security Symposium 2022.
- [10] Norwegian Consumer Council. (2018). Deceived by Design: How Tech Companies Use Dark Patterns. Forbrukerradet.
- [11] CJEU. (2019). Judgment in Case C-673/17, Planet49 GmbH v. Bundesverband der Verbraucherzentralen. Grand Chamber.
- [12] Gray, C. M., Kou, Y., Battles, B., Hoggatt, J., & Toombs, A. L. (2018). The Dark (Patterns) Side of UX Design. Proceedings of CHI 2018. ACM.
- [13] FTC. (2022). Bringing Dark Patterns to Light. Federal Trade Commission Staff Report.
- [14] Brignull, H. (2010+). Deceptive Design. <https://www.deceptive.design/>
- [15] Yao, J., et al. (2025). Benchmarking LLM Inference Platforms: Latency, Throughput and Cost Analysis. PMC. <https://pmc.ncbi.nlm.nih.gov/articles/PMC12562575/>
- [16] Van Alsenoy, B., et al. (2019). From Opt-in to Opt-out: A Comparative Analysis of Online Privacy Consent. Computer Law and Security Review.
- [17] L. C. Kasireddy et al. (2025). Securing Business Data in Multi-Cloud Environments. ICDISS 2025. doi:10.1109/ICDISS68238.2025.11320589
- [18] N. Soni, L. C. Kasireddy et al. (2025). RNN Framework for DDoS Attack Detection. MRIE 2025. doi:10.1109/MRIE66930.2025.11156616