

Towards a Preliminary Ontology of Dark Patterns Knowledge

Colin M. Gray
gray42@purdue.edu
Purdue University
West Lafayette, Indiana, USA

Cristiana Santos
c.teixeirasantos@uu.nl
Utrecht University
The Netherlands

Nataliia Bielova
nataliia.bielova@inria.fr
Inria Centre at Université Côte d'Azur
France

ABSTRACT

Deceptive design practices are increasingly used by companies to extract profit, harvest data, and limit consumer choice. *Dark patterns* represent the most common contemporary amalgamation of these problematic practices, connecting designers, technologists, scholars, regulators, and legal professionals in transdisciplinary dialogue. However, a lack of universally accepted definitions across the academic, legislative and regulatory space has likely limited the impact that scholarship on dark patterns might have in supporting sanctions and evolved design practices. In this late breaking work, we seek to harmonize regulatory and academic taxonomies of dark patterns, proposing a preliminary three-level ontology to create a shared language that supports translational research and regulatory action. We identify potential directions for scholarship and social impact building upon this ontology.

CCS CONCEPTS

• **Human-centered computing** → **Human computer interaction (HCI); Empirical studies in HCI.**

KEYWORDS

dark patterns, deceptive design, regulation, ontology

ACM Reference Format:

Colin M. Gray, Cristiana Santos, and Nataliia Bielova. 2023. Towards a Preliminary Ontology of Dark Patterns Knowledge. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23)*, April 23–28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3544549.3585676>

1 INTRODUCTION

Deceptive design practices are increasingly common in digital environments, impacting digital experiences on social media [26, 31], e-commerce [24], mobile devices [20], cookie consent banners [18], and gaming [37], among others. An increasingly dominant framing of these deceptive practices is known as “dark patterns”—describing instances where design choices subvert, impair, or distort the ability of a user to make autonomous and informed choices in relation to digital systems regardless of the designer’s intent [3, 4, 9].

While the origins of dark patterns as a concept to describe manipulative design practices goes back over a decade to when the term was coined by practitioner and scholar Harry Brignull [11],

in the past five years there has been growing momentum in the use of the term to unite scholars, regulators, and designers in transdisciplinary dialogue to identify problematic practices and find ways to prevent or discourage the use of these patterns. According to a recent study of the historical evolution of #darkpatterns on Twitter by Obi and colleagues [29], since 2019, conversations have included stakeholders not only from design and technology but also social scientists, lawyers, journalists, lawmakers, and members of regulatory bodies and consumer protection organizations.

Within the regulatory space, in 2022 alone, the term “dark patterns” was codified into EU law in the Digital Services Act [9], the Digital Markets Act [5], and the Data Act proposal [8], and into US law in the California CPRA [4]. Regulatory bodies such as the US Federal Trade Commission (FTC), the UK Competition and Market Authority (CMA), the EU Commission, the European Data Protection Board (EDPB) and the Organisation for Economic Co-operation and Development (OECD) have released guidance on specific types of dark patterns with various levels of overlap with definitions from academic scholarship [2, 3, 6, 12, 30]. In addition, the concept of dark patterns has been leveraged in sanctions against companies that have relied upon manipulative practices. Recent actions include a \$245 million USD judgment against Fortnite, a product from Epic Games, for their use of manipulative practices to encourage the purchase of content [36] and multiple settlements by various US states against Google for their use of dark patterns to obtain location data [7, 35]. In the EU, both Data Protection Authorities (DPAs) and court decisions have forbidden certain practices related to dark patterns, including: pre-selection of choices [1]; refusing consent if it is more difficult than giving it [15, 16]; and misinforming users on the purposes of processing data and how to reject them [16, 23].

As part of this convergent discourse, HCI scholars have addressed the threat of dark patterns in a wide range of publications, proposing definitions and types of dark patterns [10, 17, 22, 24, 25]. However, the specific forms that dark patterns can take, the role of context, the ubiquity of the practices, the technolog(ies) used or application area, the comparative harms of different patterns, remedies, and the role of user education are still a topic of ongoing research and the consequence of this dynamic topic is of an ever-expanding list of categories and variants whose scale continues to grow. Two large challenges face an ongoing transdisciplinary engagement with the concept of dark patterns. First, the literature has grown quickly and is siloed, often lacking accurate citation provenance trails of given typologies and definitions, making it difficult to trace where (new) specific types of patterns emerged and under which conditions. The space that dark patterns scholars have sought to cover is also vast, with important research occurring in specific domains (e.g., games, e-commerce, privacy and data protection) and across

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CHI EA '23, April 23–28, 2023, Hamburg, Germany

© 2023 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9422-2/23/04.

<https://doi.org/10.1145/3544549.3585676>

different technologies and modalities (e.g., mobile, desktop, conversational user interfaces (CUIs), AR/VR). This diversity of research has led some scholars to propose fragmentary, domain-specific typologies without necessarily finding commonalities across domains. Second, regulators—the ultimate decision makers that could provide legal certainty to this landscape— and policy makers have been interested in the scholarly conversation regarding dark patterns, but have in some cases created wholly new domain-related terminology to describe types already known in the academic literature. In other cases, regulators and policymakers have inconsistently cited academic sources making connections across the regulatory, legal, and academic spaces fraught.

We seek to support these challenges and ongoing conversations by building the foundation for a common ontology of dark patterns. By taking the first steps towards building an *ontology*, we seek to create a shareable, extendable, and reusable knowledge representation of dark patterns conceptualization. This groundwork for an ontology is both domain and application agnostic though it has potential utility in domain or context-specific instances. As noted by Fonseca [14], these types of ontologies can be useful in supporting social science research by “creating better conceptual schemas and applications.” To create this preliminary ontology, we build upon ten contemporary taxonomies of dark patterns from both the academic and regulatory literature, and thereafter we identify three levels of hierarchy for pattern types. Hence we harmonize concepts across these taxonomies to provide a consistent and consolidated, shared, and reusable dark patterns ontology for future research, regulatory action, and sanctions.

We make three contributions in this late-breaking work.

- (1) First, we introduce the hierarchical concepts of low-level, meso-level, and high-level dark patterns to the literature, disambiguating UI-level patterns that may lead to opportunities for detection (low-level) and strategies that may be targeted by policy and legislation (meso- and high-level).
- (2) Second, we identify when patterns first emerged and how naming has evolved over time and across sources.
- (3) Third, we describe a common language and hierarchy of dark patterns that aligns disparate terminology from scholars and regulators. This enabled us to propose the foundation for an ontology of dark patterns knowledge that can be utilized by practitioners, scholars, regulators, and legal professionals.

2 METHOD

We used a content analysis approach [27] to identify and characterize elements of existing dark patterns taxonomies. As a research team, we leveraged our collective experiences in human-computer interaction, design, computer science, law, and regulation. Specifically, our team included established dark patterns scholars, including one with a focus on human-computer interaction and design (Colin), one with a focus on computer science and web measurement and experience in regulation (Nataliia), and one with a background in computer science and data protection law (Cristiana). Across these perspectives, in accordance with previous scholarship, we seek to characterize dark patterns in a transdisciplinary space where multiple disciplinary perspectives are needed [18].

2.1 Data Collection

We collected dark patterns taxonomies from a total of 10 sources, including: 1) a stable set of patterns shared on <https://darkpatterns.org> since 2018 by Harry Brignull¹; 2) scholarly academic sources that present a distinct and comprehensive taxonomy and have been regularly cited [10, 17, 22, 24] in scholarly and regulatory literature; and 3) public reports from stakeholders and regulators in the EU and USA that include a dark patterns taxonomy [2, 3, 6, 12, 30]. Our selection of these 10 sources (summarized in supplemental material) encompass the most authoritative and commonly cited taxonomies in the research and regulatory literature, and all taxonomies that were previously cited in regulatory reports were also present in our set of scholarly academic sources.

2.2 Data Analysis

We began our analysis by identifying the constitutive components of each taxonomy without considering overlaps across sources through a bottom-up approach. Across the ten taxonomies from academic and regulatory sources, we identified 186 low-level and 59 high-level patterns (a total of 245 patterns) (see Tables 1 and 2 in the appendix). The number of discrete elements is perhaps unsurprising, since each typology author has used a different point of focus and categorization based on the sector they sought to describe or support (e.g., Mathur et al. [25] and the CMA [6] focus on e-commerce; the EDPB’s focus on data protection practices within social media platform interfaces [12]²). All taxonomy elements are included in supplemental material for other scholars to build upon. We used the following procedure to carefully identify individual taxonomy components, their source, and similarities of components across taxonomies:

- **Aggregating existing patterns.** We first listed all high- and low-level patterns verbatim in the structure originally indicated in the textual source. *High-level patterns* include any instances where the pattern is denoted as a category, strategy, goal, intention, or other parent in a parent-child relationship and *low-level patterns* indicate specific patterns that are included as a child in a parent-child relationship.
- **Identifying provenance through direct citations and inference.** Based on citations provided in the source-document, we indicated any instances where patterns were *directly cited* or otherwise duplicated from previous sources. Because many patterns were uncited—particularly in regulatory reports—we also relied upon citations elsewhere in the document or explicit use of existing pattern vocabulary and definitions from previously published sources, which we indicate as *inferential*. We used these direct and inferential citation patterns to identify where patterns were first introduced, even if they appeared alongside other patterns that had been published previously.

¹This collection of dark patterns was moved to <https://www.deceptive.design> in 2022, but the patterns have been stable since 2018 when the final pattern, “confirmshaming,” was added.

²The EDPB guidelines were updated to a new version in February 2023 after our initial analysis was complete [13]. Future versions of the ontology will need to address any differences.

- **Clustering similar patterns.** We grouped patterns that appeared either to be identical or similar (in a *is-a* or *equivalent-to* relationship), using definitions to identify affinities among patterns that did not have identical names. This portion of the analysis was the most extensive, including in depth conversations between an HCI (Colin) and legal (Cristiana) scholar and a careful reading of the definitions as they might be understood by designers and lawyers.
- **Creating meso-level patterns.** From the findings of this content analysis procedure, we recognized that there were not only low- and high-level patterns present, but also a “*meso*” level of pattern knowledge. By recognizing similarities among low-level patterns, we introduced *meso-level patterns* into our analysis, identifying these patterns by using the names or elements of existing taxonomies where possible, or coining new pattern names to characterize the low-level patterns we grouped together.
- **Finalizing the draft ontology.** Across these three levels of hierarchy, we grouped 233 of the 245 taxonomy elements³. This structure resulted in a draft ontology which includes 6 high-level patterns, 24 meso-level patterns, and 34 low-level patterns.

3 MAPPING THE EVOLUTION OF DARK PATTERNS

Pattern types have largely stabilized in the past five years, including high-level pattern types or “strategies” (e.g., nagging, obstruction, sneaking, interface interference, forced action) and low-level patterns (including Brignull’s [11] and those introduced by Gray et al. [17] and Mathur et al. [24]).

High-level patterns were most likely to co-occur across multiple sources. For instance, Gray et al.’s [17] original five high-level “dark pattern strategies” were found across multiple other sources, even if they were not consistently cited: nagging [2, 22], obstruction [2, 3, 22, 24], sneaking [2, 3, 22, 24], interface interference [2, 3, 22], and forced action [2, 3, 22, 24] (FTC uses “coerced action” instead). After their introduction in Mathur et al. [24], newly introduced categories relating to social psychology or behavioral economics also became common: urgency [2, 3, 22], scarcity [3, 22], and social proof [2, 3, 22] (the FTC bundles “Endorsements” with “social proof”). We have grouped these types together as part of a sixth high-level pattern of “social engineering.”

Domain or context-specific patterns. The most volatility has occurred in relation to *domain-* or *context-specific patterns*. These include expansions of Mathur et al.’s [24] high-level patterns of “social proof” and “scarcity,” which have since been reiterated by the EU Commission [2] and OECD [30] and extended by the CMA [6] and FTC [3] taxonomies. In addition, the EDPB guidance on dark patterns in social media [12] included a wholly new set of 6 high-level and 15 low-level patterns, although the majority of these could be inferred as similar to already existing patterns proposed in the

³Four ungrouped elements were from the CMA report [6] and described generic elements of digital systems which were not explicitly framed as deceptive or manipulative: Choice Structure, Choice Information, Feedback, and Messengers. All eight high-level patterns from Bösch [10] were also excluded since they were not reiterated in any downstream literature.

academic literature. Importantly, though, the EDPB taxonomy included multiple patterns which we found to be new low-level or meso-level additions, including “privacy maze,” “dead end,” “conflicting information,” “information without context” (which we renamed from the EDPB pattern “decontextualizing”), and “visual prominence” (which we renamed from the EDPB pattern “look over there”). Similarly, the CMA taxonomy focused on choice architecture as a guiding structure with three categories focused on choice “structure,” “information,” and “pressure.” This taxonomy structure also yielded new patterns, including “bundling,” “complex language,” and “personalization.”

Our analysis demonstrates the value in classifying or generating context-specific patterns that illuminate gaps in current taxonomies, and also the benefit of mapping these patterns within larger ontologies (such as the one we propose in this LBW) to identify abstractions of patterns that may apply across many domains, contexts, and legal fields.

4 TOWARDS A SHARED ONTOLOGY OF DARK PATTERNS

In our draft ontology, summarized in Figures 1 and 2, we seek to synthesize and harmonize all existing taxonomies while adding useful and consistent structure to allow for other scholars and regulators to build on a shared description of dark patterns knowledge. Our ontology includes three different levels of hierarchy:

- **High-level patterns** are the most abstracted form of knowledge, including general *strategies* that characterize the inclusion of manipulative, coercive, or deceptive elements that might limit user autonomy and decision making. These patterns are context-agnostic and can be employed through a range of modalities and technologies (e.g., desktop, mobile, VUIs, VR/AR) and application types (e.g., e-commerce, gaming, social media).
- **Meso-level patterns** bridge high- and low-level forms of knowledge and describe an *angle of attack* or specific approach to limiting, impairing, or undermining the ability of the user to make autonomous and informed decisions or choices. These patterns are content-agnostic and may be interpreted in a contextually-appropriate way based on the specific context of use or application type.
- **Low-level patterns** are the most situated and contextually dependent form of knowledge, including specific *means of execution* that limits or undermines user autonomy and decision making, is described in visual and/or temporal form(s), and is likely to be detectable through algorithmic, manual, or other technical means.

4.1 Problematizing and Evolving the Ontology

Not all of our mappings were clear-cut and some may be productively extended or disputed in future versions of this ontology. Through dialogue, we sought to locate existing patterns within our ontology based on our best understanding of the pattern as described by its name and definition in the source taxonomy. We describe some of our challenges in mapping and present these as opportunities to refine the ontology in future work:

High-Level Pattern	Meso-Level Pattern	Low-Level Pattern
Nagging D: <u>Gr</u> <u>Lu</u> <u>EUCOM</u> <u>FTC</u> <u>OECD</u> I: <u>EDPB</u> <u>CMA</u>	–	–
Obstruction D: <u>Gr</u> <u>Lu</u> <u>Ma</u> <u>EUCOM</u> <u>FTC</u> <u>OECD</u> I: <u>EDPB</u> <u>CMA</u>	Roach Motel (D: <u>Br</u> <u>Gr</u> <u>Lu</u> <u>EUCOM</u> I: <u>Bö</u> <u>Ma</u> <u>FTC</u> <u>OECD</u>)	Immortal Accounts (D: <u>Bö</u> <u>FTC</u>) Dead End (D: <u>EDPB</u>)
	Creating Barriers	Price Comparison Prevention (D: <u>Br</u> <u>Gr</u> <u>Lu</u> <u>FTC</u> <u>EUCOM</u> <u>OECD</u>) Intermediate Currency (D: <u>Gr</u> <u>Lu</u> <u>FTC</u> <u>EUCOM</u> <u>OECD</u> ; I: <u>CMA</u>) Privacy Maze (D: <u>EDPB</u>)
Sneaking D: <u>Gr</u> <u>Lu</u> <u>Ma</u> <u>EUCOM</u> <u>OECD</u> I: <u>EDPB</u> <u>CMA</u> <u>FTC</u>	Bait and Switch (D: <u>Br</u> <u>Gr</u> <u>Lu</u> <u>FTC</u> <u>EUCOM</u> I: <u>OECD</u>)	Disguised Ad (D: <u>Br</u> <u>Gr</u> <u>Lu</u> <u>FTC</u> <u>EUCOM</u> <u>OECD</u>) Sneak into Basket (D: <u>Br</u> <u>Gr</u> <u>Ma</u> <u>Lu</u> <u>FTC</u> <u>EUCOM</u> <u>OECD</u>) Drip Pricing, Hidden Costs, or Partitioned Pricing (D: <u>Br</u> <u>Gr</u> <u>Ma</u> <u>Lu</u> <u>CMA</u> <u>FTC</u> <u>EUCOM</u> <u>OECD</u>) Reference Pricing (D: <u>CMA</u> <u>OECD</u>) Conflicting Information (D: <u>EDPB</u>) Information without Context (I: <u>EDPB</u>)
	Hiding Information	
Interface Interference D: <u>Gr</u> <u>Lu</u> <u>EUCOM</u> <u>FTC</u> <u>OECD</u> I: <u>Br</u> <u>Ma</u> <u>EDPB</u> <u>FTC</u>	(De)contextualizing Cues	False Hierarchy (D: <u>Gr</u> <u>Lu</u> <u>FTC</u> <u>OECD</u> I: <u>EDPB</u>) Visual Prominence (I: <u>EDPB</u>) Bundling (D: <u>CMA</u>) Pressured Selling (D: <u>Lu</u> <u>FTC</u>)
	Manipulating Visual Choice Architecture (I: <u>CMA</u>)	
	Bad Defaults (D: <u>Bö</u> <u>Gr</u> <u>Lu</u> <u>FTC</u> <u>OECD</u> ; I: <u>CMA</u> <u>EUCOM</u> <u>EDPB</u>)	–
	Emotional or Sensory Manipulation (I: <u>Gr</u> <u>Lu</u> <u>CMA</u> <u>EUCOM</u> <u>OECD</u>)	Cuteness (D: <u>Lu</u>) Positive or Negative Framing (I: <u>Gr</u> <u>EDPB</u>)
	Trick Questions (D: <u>Br</u> <u>Gr</u> <u>Ma</u> <u>Lu</u> <u>FTC</u> <u>EUCOM</u> <u>OECD</u>)	–
	Choice Overload (I: <u>EDPB</u> <u>CMA</u>)	–
	Hidden Information (D: <u>Gr</u> <u>Lu</u> <u>FTC</u> <u>EUCOM</u> <u>OECD</u> ; I: <u>EDPB</u>)	–
	Language Inaccessibility	Wrong Language (I: <u>EDPB</u>) Complex Language (D: <u>CMA</u>)
	Feedforward Ambiguity (I: <u>EDPB</u>)	–

Figure 1: Draft ontology of dark patterns organized by level of pattern. “D” indicates a *direct* use of the pattern language in the original source(s) and “I” indicates an *inferred* similarity between different terminology used across two or more pattern types. Sources are indicated by abbreviation and are colored cyan if they are regulatory reports or magenta if they are academic articles. *Italicized pattern names* indicate new pattern types introduced in this paper while all other text relies upon the sources indicated. Underlined sources indicate the earliest mention of that pattern or patterns in the sources we analyzed. A full description of the inferred pattern names is included in supplemental material to support future work.

High-Level Pattern	Meso-Level Pattern	Low-Level Pattern	
Forced Action D: Gr Lu Ma EUCOM OECD I: CMA FTC	Forced Continuity (D: Br Gr Lu FTC EUCOM OECD ; I: Ma)	–	
	Forced Registration (D: Bö Lu FTC EUCOM OECD ; I: Ma CMA FTC)	–	
	<i>Forced Communication or Disclosure</i>		Privacy Zuckering (D: Br Gr Lu OECD ; I: FTC)
			Friend Spam (D: Br Lu FTC OECD)
			Address Book Leeching (D: Bö Lu FTC OECD)
			Social Pyramid (D: Gr Lu FTC OECD)
		Pay-to-Play (D: FTC)	
		Grinding (D: FTC)	
	<i>Attention Capture</i>		Auto-Play (D: FTC)
	Social Engineering	Scarcity and Popularity Claims (D: Ma Lu FTC CMA)	High Demand (D: Ma Lu FTC EUCOM OECD)
Social Proof (D: Ma Lu EUCOM OECD)		Low Stock (D: Ma Lu FTC EUCOM OECD)	
		Endorsements and Testimonials (D: Ma Lu FTC EUCOM OECD ; I: FTC)	
		Parasocial Pressure (I: FTC)	
Urgency (D: Ma Lu FTC EUCOM OECD)		Activity Messages (D: Ma Lu FTC EUCOM OECD)	
		Countdown Timer (D: Ma Lu FTC EUCOM OECD)	
<i>Shaming</i>		Limited Time Message (D: Ma Lu FTC EUCOM OECD)	
Personalization (D: CMA)		Confirmshaming (D: Br Ma Lu FTC EUCOM OECD)	

Figure 2: Draft ontology of dark patterns organized by level of pattern, continued.

Combinations of patterns evolve over time. Some combinations of patterns have evolved over time. For instance, Mathur et al.’s [24] high-level pattern “social proof” originated with two sub-patterns, “activity messages” and “testimonials.” The FTC created new low-level patterns, introducing “endorsements” (we bundled with testimonials as one low-level pattern) and more specific types of endorsement or testimonials (e.g., “deceptive celebrity endorsements,” “false activity messages”). Future work could identify the most useful level of abstraction for such patterns.

Novel and context-specific names for dark patterns have been introduced. The use of novel names for patterns (particularly by the EDPB and CMA) or the use of patterns in specific contexts (e.g., e-commerce, social media) caused us to consider both the presence of granular low-level patterns and the relation of these low-level patterns to inferred meso-level patterns. In particular, the use of novel names for patterns types and definitions was a challenge from an analytic perspective, resulting in: i) instances where a wholly new pattern was introduced (e.g., CMA’s “information overload” which we leveraged to create a new meso-level

pattern of “choice overload”); ii) instances where a new high-level strategy was highly similar to an existing high-level strategy (e.g., EDPB’s “skipping” which we subsumed within “sneaking”); and iii) instances where existing patterns included both a generalizable pattern and domain-specific information which may need to be captured in specific low-level patterns in future work (e.g., EDPB’s “left in the dark” is a form of “hidden information” but implies specific low-level patterns that are specific to data protection).

These analytic challenges strengthen some of the key definitional components of dark patterns as they have previously been described in the literature [25] but inform areas where definitions could be further problematized. For instance, many pattern definitions mixed *characteristics* of dark patterns (e.g., information hiding, manipulation of the choice architecture), *impact on user behavior* (e.g., confusion, obstruction), *manifestation in the user interface* (e.g., visibility of elements, interpretability by users), and *context-specific issues* (e.g., impacts on data protection, presentation in social media interfaces). Future definitions of patterns at all levels could benefit from separating out these elements at appropriate

levels of abstraction to increase the potential for utility across many contexts and domains—particularly at the high- and meso-levels of pattern description.

5 USING THE ONTOLOGY TO SUPPORT TRANSDISCIPLINARY DIALOGUE

Building on our preliminary ontology, we identify multiple benefits for all actors that use a shared nomenclature of dark patterns. This can help not only to anchor patterns in existing literature, but also to enable actors to extend or further characterize them. Further alignment of dark patterns definitions can be achieved through multiple types of practices that could further help establish dialogue between actors from different disciplinary and professional roles, including:

1) Identifying multiple levels of granularity. This ontology can be used to identify abstractions of dark patterns across many domains, contexts, or application types (i.e., high- and meso-level patterns) while creating more robust characterizations of low-level patterns that may be highly domain or context specific. By indicating gaps in our draft ontology, we show that many additional low-level patterns need to be identified, which would aid computer scientists and engineers in building precise detection tools and provide guidance for designers. For instance, “nagging” only exists as a high-level pattern and it is unclear which kinds of nagging should be considered problematic, and under which conditions. Also, the ontology can be extended to include potential harms in relation to specific dark patterns types [19]. Other meso-level patterns that lack any defined low-level patterns (e.g., “trick questions,” “choice overload,” “hidden information”) raise the question of which specific cases—as inscribed through low-level patterns—need to be articulated as particularly problematic or indicative. This work will require a shared language among regulators and scholars to inform disciplined inquiry of specific domains and use contexts.

2) Engaging in more robust citation practices to trace the evolution of pattern types and their presence over time. This ontology—and its extension over time—can be supported through robust citation patterns across regulatory and academic literature, including indications where definitions of specific patterns have been altered, extended, or synthesized. While definitions of some low and high-level patterns have remained relatively consistent over time (in particular, Brignull’s original patterns), the rapid growth of new pattern types and the introduction of hierarchy injects added complexity that can inform more disciplined and transparent knowledge generation practices.

3) Connecting regulatory sanctions to multiple levels of patterns and combinations of pattern types. By mapping case law to the ontology, regulators could gain additional knowledge and support in identifying where dark patterns practices at multiple levels and in multiple combinations are at play, and are deemed to be illegal per jurisdiction [21]. Current descriptions of the use of dark patterns often report infringements only at the lowest level (often re-using Brignull’s original pattern types) and in doing so, could miss broader abstractions that may translate across domains. Early enforcement actions in relation to subscription services and online tracking (a focus in both EU and FTC enforcement) has started to reveal the interplay of multiple pattern types at multiple

levels in causing user confusion, manipulation, or coercion, and other broad genres of manipulation that might benefit from similar characterizations.

4) Using levels of pattern granularity to inform the detection of dark patterns through automated or manual approaches. The detection of dark patterns could also be more robustly supported by our ontology, with our assertion that low-level patterns show the most promise in being detectable. Existing detection efforts (e.g., [24, 28, 32–34]) have shown that higher-level patterns are difficult or impossible to detect at scale due to their abstract nature that requires interpretation, while low-level UI elements with discrete and known qualities (e.g., cookie consent banners, elements of the checkout process) are more accessible with software tools for automated detection. Our ontology of low-level patterns and gaps creates a foundation for future detection efforts, allowing computer science scholars to focus on pattern types which are most likely to be detectable.

6 CONCLUSION

In this late-breaking work, we present our analysis of ten existing taxonomies of dark patterns and propose a foundation for a shared ontology of dark patterns knowledge. Building on our analysis, future scholars, regulators, and legal professionals can benefit from our process of aggregation and consolidation, using our hierarchical organization of dark patterns types to indicate links to existing and similar concepts. This description encourages the establishment of provenance in future work, allowing scholars and regulators to identify pattern types and their origins and provide an audit trail to connect specific contextually-bound instances with broader categorizations. Finally, this draft ontology creates a foundation for a shared and reusable knowledge source, allowing many stakeholders to work together in building a shared, explicit and precise conceptualization of what is already known in the literature and which can be further refined and extended.

ACKNOWLEDGMENTS

This work is funded in part by the National Science Foundation under Grant No. 1909714 and the ANR 22-PECY-0002 IPOP (Interdisciplinary Project on Privacy) project of the Cybersecurity PEPR.

REFERENCES

- [1] 2019. Judgment in Case C-673/17 Bundesverband der Verbraucherzentralen und Verbraucherverbände – Verbraucherzentrale Bundesverband eV v Planet49 GmbH. <http://curia.europa.eu/juris/documents.jsf?num=C-673/17>.
- [2] 2022. *Behavioural study on unfair commercial practices in the digital environment : dark patterns and manipulative personalisation : final report*. Publications Office of the European Union. <https://op.europa.eu/en/publication-detail/-/publication/606365bc-d58b-11ec-a95f-01aa75ed71a1/language-en/format-PDF/source-257599418>
- [3] 2022. *Bringing Dark Patterns to Light Staff Report*. Technical Report. Federal Trade Commission. https://www.ftc.gov/system/files/ftc_gov/pdf/P214800%20Dark%20Patterns%20Report%209.14.2022%20-%20FINAL.pdf
- [4] 2022. California Privacy Rights Act. https://coppa.ca.gov/meetings/materials/20220608_item3.pdf
- [5] 2022. Digital Markets Act - Regulation (EU) 2022/1925 of the European Parliament and of the Council of 14 September 2022 on contestable and fair markets in the digital sector and amending Directives (EU) 2019/1937 and (EU) 2020/1828 (Digital Markets Act) (Text with EEA relevance). <http://data.europa.eu/eli/reg/2022/1925/oj>

- [6] 2022. *Evidence review of Online Choice Architecture and consumer and competition harm*. Technical Report. <https://www.gov.uk/government/publications/online-choice-architecture-how-digital-design-can-harm-competition-and-consumers/evidence-review-of-online-choice-architecture-and-consumer-and-competition-harm> Accessed: 2022-4-13.
- [7] 2022. Press Release: AG Racine announces Google must pay \$9.5 million for using “dark patterns” and deceptive location tracking practices that invade users’ privacy. <https://thedcline.org/2022/12/30/press-release-ag-racine-announces-google-must-pay-9-5-million-for-using-dark-patterns-and-deceptive-location-tracking-practices-that-invade-users-privacy/> Accessed: 2022-12-31.
- [8] 2022. Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on harmonised rules on fair access to and use of data (Data Act). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=COM%3A2022%3A68%3AFIN>
- [9] 2022. Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market For Digital Services and amending Directive 2000/31/EC (Digital Services Act). *Official Journal of the European Union* 65, L277 (Oct. 2022).
- [10] Christoph Bösch, Benjamin Erb, Frank Kargl, Henning Kopp, and Stefan Pfattheicher. 2016. Tales from the Dark Side: Privacy Dark Strategies and Privacy Dark Patterns. *Proceedings on Privacy Enhancing Technologies* 2016, 4 (2016). <https://doi.org/10.1515/popets-2016-0038>
- [11] Harry Brignull. 2018. Dark Patterns: User Interfaces Designed to Trick People. <http://darkpatterns.org/>
- [12] European Data Protection Board. 2022. *Guidelines 3/2022 on Dark patterns in social media platform interfaces: How to recognise and avoid them*. Technical Report Version 1.0. https://edpb.europa.eu/system/files/2022-03/edpb_03-2022_guidelines_on_dark_patterns_in_social_media_platform_interfaces_en.pdf
- [13] European Data Protection Board. 2023. *Guidelines 3/2022 on Dark patterns in social media platform interfaces: How to recognise and avoid them*. Technical Report Version 2.0. https://edpb.europa.eu/system/files/2023-02/edpb_03-2022_guidelines_on_deceptive_design_patterns_in_social_media_platform_interfaces_v2_en_0.pdf
- [14] Frederico Fonseca. 2007. The double role of ontologies in information science research. *Journal of the American Society for Information Science and Technology* 58, 6 (April 2007), 786–793. <https://doi.org/10.1002/asi.20565>
- [15] French DPA (CNIL). 2021. Deliberation of the restricted committee No. SAN-2021-023 of 31 December 2021 concerning GOOGLE LLC and GOOGLE IRELAND LIMITED. https://www.cnil.fr/sites/default/files/atoms/files/deliberation_of_the_restricted_committee_no_san-2021-023_of_31_december_2021_concerning_google_llc_and_google_ireland_limited.pdf.
- [16] French DPA (CNIL). 2022. Deliberation of the restricted committee No. SAN-2021-024 of 31 December 2021 concerning FACEBOOK IRELAND LIMITED. https://www.cnil.fr/sites/default/files/atoms/files/deliberation_of_the_restricted_committee_no_san-2021-024_of_31_december_2021_concerning_facebook_ireland_limited.pdf.
- [17] 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* (Montreal QC, Canada) (CHI '18). dl.acm.org, New York, NY, USA, 534:1–534:14. <https://doi.org/10.1145/3173574.3174108>
- [18] Colin M Gray, Cristiana Santos, Nataliia Bielova, Michael Toth, and Damian Clifford. 2021. Dark Patterns and the Legal Requirements of Consent Banners: An Interaction Criticism Perspective. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (CHI'21). ACM Press. <https://doi.org/10.1145/3411764.3445779>
- [19] Gunawan, Santos, and Kamara. 2022. Redress for Dark Patterns Privacy Harms? A Case Study on Consent Interactions. *2022 ACM Symposium on (2022)*. <https://johannagunawan.com/assets/pdf/gunawan-22-cslaw.pdf>
- [20] Johanna Gunawan, Amogh Pradeep, David Choffnes, Woodrow Hartzog, and Christo Wilson. 2021. A Comparative Study of Dark Patterns Across Web and Mobile Modalities. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2 (Oct. 2021), 1–29. <https://doi.org/10.1145/3479521>
- [21] Mark Leiser. 2020. ‘Dark Patterns’: the case for regulatory pluralism. LawArXiv ea5n2. Center for Open Science. <https://doi.org/10.31219/osf.io/ea5n2>
- [22] Jamie Luguri and Lior Jacob Strahilevitz. 2021. Shining a Light on Dark Patterns. *Journal of Legal Analysis* 13, 1 (March 2021), 43–109. <https://doi.org/10.1093/jla/laaa006>
- [23] Luxembourg DPA. 2021. Decision regarding Amazon Europe Core S.À RL. <https://cnpd.public.lu/fr/actualites/international/2021/08/decision-amazon-2.html>
- [24] Arunesh Mathur, Gunes Acar, Michael J Friedman, Elena 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 (Nov. 2019), Article No. 81. <https://doi.org/10.1145/3359183>
- [25] 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*. Association for Computing Machinery, New York, NY, USA, 1–18. <https://doi.org/10.1145/3411764.3445610>
- [26] Thomas Mildner and Gian-Luca Savino. 2021. Ethical User Interfaces: Exploring the Effects of Dark Patterns on Facebook. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI EA '21, Article 464). Association for Computing Machinery, New York, NY, USA, 1–7. <https://doi.org/10.1145/3411763.3451659>
- [27] Kimberly A Neundorff. 2001. *The Content Analysis Guidebook*. Sage.
- [28] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. 2020. Dark Patterns after the GDPR: Scraping Consent Pop-ups and Demonstrating their Influence. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376321>
- [29] Ikechukwu Obi, Colin M Gray, Shruthi Sai Chivukula, Ja-Nae Duane, Janna Johns, Matthew Will, Ziqing Li, and Thomas Carlock. 2022. Let’s Talk About Socio-Technical Angst: Tracing the History and Evolution of Dark Patterns on Twitter from 2010-2021. (July 2022). arXiv:2207.10563 [cs.SI] <http://arxiv.org/abs/2207.10563>
- [30] OECD. 2022. *Dark commercial patterns*. Technical Report. <https://doi.org/10.1787/44f5e846-en>
- [31] Brennan Schaffner, Neha A Lingareddy, and Marshini Chetty. 2022. Understanding Account Deletion and Relevant Dark Patterns on Social Media. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2 (Nov. 2022), 1–43. <https://doi.org/10.1145/3555142>
- [32] Than Htut Soe, Cristiana Teixeira Santos, and Marija Slavkovic. 2022. Automated detection of dark patterns in cookie banners: how to do it poorly and why it is hard to do it any other way. (April 2022). arXiv:2204.11836 [cs.LG] <http://arxiv.org/abs/2204.11836>
- [33] Ioannis Stavrakakis, Andrea Curley, Dympha O’Sullivan, Damian Gordon, and Brendan Tierney. 2021. A Framework of Web-Based Dark Patterns that can be Detected Manually or Automatically. (2021). <https://doi.org/10.21427/20g8-d176>
- [34] Marieke Van Hofslot, Almila Akdag Salah, Albert Gatt, and Cristiana Santos. 2022. Automatic Classification of Legal Violations in Cookie Banner Texts. In *Proceedings of the Natural Legal Language Processing Workshop 2022*. Association for Computational Linguistics, Abu Dhabi, United Arab Emirates (Hybrid), 287–295. <https://aclanthology.org/2022.nllp-1.27>
- [35] Jess Weatherbed. 2022. Google is paying a \$85m settlement to Arizona to end user-tracking suit. <https://www.theverge.com/2022/10/5/23389331/google-settlement-arizona-user-tracking-privacy-suit>. <https://www.theverge.com/2022/10/5/23389331/google-settlement-arizona-user-tracking-privacy-suit> Accessed: 2023-1-4.
- [36] Shoshana Wodinsky. 2022. The ‘dark patterns’ in Fortnite that led to the largest FTC penalties ever. <https://www.marketwatch.com/story/the-dark-patterns-in-fortnite-that-led-to-the-largest-ftc-penalties-ever-11671488228>. <https://www.marketwatch.com/story/the-dark-patterns-in-fortnite-that-led-to-the-largest-ftc-penalties-ever-11671488228> Accessed: 2022-12-20.
- [37] José P Zagal, Staffan Björk, and Chris Lewis. 2013. Dark Patterns in the Design of Games. In *Foundations of Digital Games*. <http://www.diva-portal.org/smash/record.jsf?pid=diva2%3A1043332&dsid=1018>

Table 1: Academic taxonomies of dark patterns.

	High-Level Pattern	Low-Level Pattern
Brignull [11]	—	Sneak into Basket, Bait and Switch, Roach Motel, Price Comparison Prevention, Disguised Ads, Privacy Zuckering, Trick Questions, Hidden Costs, Confirmshaming, Friend Spam, Forced Continuity, Misdirection
Bösch et al. [10]	Obscure Maximize Deny Preserve Centralize Publish, Violate, Fake	Privacy Zuckering, Immortal Accounts, Hidden Legalese Stipulations, Bad Defaults Shadow User Profiles, Address Book Leeching, Forced Registration Immortal Accounts Shadow User Profiles, Address Book Leeching Shadow User Profiles —
Gray et al. [17]	Nagging Sneaking Obstruction Interface Interference Forced Action	— Intermediate-Level Currency, Roach Motel, Price Comparison Prevention Bait and Switch, Sneak into Basket, Hidden Costs, Forced Continuity Toying with Emotion, Aesthetic Manipulation, Trick Questions, Preselection, Disguised Ad, Hidden Information, False Hierarchy Gamification, Privacy Zuckering, Social Pyramid
Mathur et al. [24]	Sneaking Urgency Misdirection Social Proof Scarcity Obstruction Forced Action	Sneak into Basket, Hidden Costs, Hidden Subscription Limited-time Message, Countdown Timer Confirmshaming, Visual Interference, Trick Questions, Pressured Selling Activity Message, Testimonials Low-stock Message, High-demand Message Hard to Cancel Forced Enrollment
Luguri et al. [22]	Nagging Social Proof Obstruction Sneaking Interface Interference Forced Action Scarcity Urgency	— Testimonials, Activity Messages Immortal Accounts, Intermediate-Level Currency, Roach Motel, Price Comparison Prevention Bait and Switch, Sneak into Basket, Hidden Costs, Hidden Subscription / Forced Continuity Cuteness, False Hierarchy / Pressured Selling, Toying with Emotion, Trick Questions, Preselection, Disguised Ad, Hidden Information / Aesthetic Manipulation, Confirmshaming Friend spam/social pyramid/address book leeching, Privacy Zuckering, Gamification, Forced Registration High Demand Message, Low Stock Message Countdown Timer, Limited Time Message

Table 2: Regulatory taxonomies of dark patterns.

	High-Level Pattern	Low-Level Pattern
EDPB [12]	Overloading Skipping Stirring Hindering Fickle Left in the Dark	Continuous Prompting, Privacy Maze, Too Many Options Deceptive Snuggness, Look Over There Emotional Steering, Hidden in Plain Sight Dead End, Longer than Necessary, Misleading Information Lacking Hierarchy, Decontextualizing Language Discontinuity, Conflicting Information, Ambiguous Wording or Information
EU Commission [2]	Nagging Social Proof Obstruction Sneaking Interface Interference Forced Action Urgency	— Testimonials, Activity Messages Intermediate-Level Currency, Roach Motel / Difficult Cancellations, Price Comparison Prevention Bait and Switch, Sneak into Basket, Hidden Costs, Hidden Subscription / Forced Continuity Toying with Emotion, Trick Questions, Preselection (default), Disguised Ad, Hidden Information / False Hierarchy, Confirmshaming Forced Registration Countdown Timer / Limited Time Message, Low Stock / High Demand Message
OECD [30]	Forced Action Interface Interference Nagging Obstruction Sneaking Social Proof Urgency	Forced Registration, Forced Disclosure / Privacy Zuckering, Friend Spam / Social Pyramid / Address Book Leeching, Gamification Hidden Information, False Hierarchy, Preselection, Misleading Reference Pricing, Trick Questions, Disguised Ads, Confirmshaming / Toying with Emotion Nagging Hard to Cancel or Opt Out / Roach Motel / Click Fatigue / Ease, (Price) Comparison Prevention, Immortal Accounts, Intermediate Currency Sneak into Basket, Hidden Costs / Drip Pricing, Hidden Subscription / Forced Continuity, Bait and Switch (including Bait Pricing) Activity Messages, Testimonials Low Stock / High Demand Message, Countdown Timer / Limited Time Message
UK CMA [6]	Choice Structure Choice Information Choice Pressure	Defaults, Ranking, Partitioned Pricing, Sludge, Bundling, Dark nudge, Choice overload and decoys, Virtual currencies in gaming, Sensory manipulation, Forced outcomes Drip pricing, Reference pricing, Framing, Complex language, Information overload Scarcity and popularity claims, Prompts and reminders, Messengers, Commitment, Feedback, Personalisation
US FTC [3]	Endorsements (Social Proof) Scarcity Urgency Obstruction Sneaking or Information Hiding Interface Interference Coerced Action Asymmetric Choice	False Activity Messages, Deceptive Consumer Testimonials, Deceptive Celebrity Endorsements, Parasocial Relationship Pressure False Low Stock Message, False High Demand Message False Discount Claims, False Limited Time Message, Baseless Countdown Timer Immortal Accounts Roadblocks to Cancellation, Price Comparison Prevention Intermediate Currency, Hidden Subscription or Forced Continuity, Drip Pricing, Hidden Costs, Hidden Information, Sneak-into-Basket Bait and Switch, Disguised Ads, False Hierarchy or Pressured Upselling, Misdirection Friend Spam, Social Pyramid Schemes, and Address Book Leeching, Pay-to-Play or Grinding, Forced Registration or Enrollment, Nagging, Auto-Play, Unauthorized Transactions Subverting Privacy Preferences, Preselection, Confirm Shaming, Trick Questions