

ACDPs Transparency Labels: Supporting User Agency Against Attention-Capture Damaging Patterns

*Original*

ACDPs Transparency Labels: Supporting User Agency Against Attention-Capture Damaging Patterns / Monge Roffarello, A., Graziosi, R., Passarella, L.. - STAMPA. - (2026), pp. 380-384. (The 29th ACM Conference on Computer-Supported Cooperative Work and Social Computing Salt Lake City, Utah (USA) October 10-14, 2026) [10.1145/3785651.3831549].

*Availability:*

This version is available at: 11583/3012848 since: 2026-10-12T14:58:10Z

*Publisher:*

ACM

*Published*

DOI:10.1145/3785651.3831549

*Terms of use:*

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

*Publisher copyright*

(Article begins on next page)

# ACDPs Transparency Labels: Supporting User Agency Against Attention-Capture Damaging Patterns

Alberto Monge Roffarello  
Dipartimento di Automatica e  
Informatica  
Politecnico di Torino  
Torino, Italy  
alberto.monge@polito.it

Riccardo Graziosi  
Dipartimento di Automatica e  
Informatica  
Politecnico di Torino  
Torino, Italy  
riccardo.graziosi@polito.it

Luca Passarella  
Politecnico di Torino  
Torino, Italy  
s333229@studenti.polito.it

## Abstract

Social platforms increasingly rely on Attention-Capture Damaging Patterns (ACDPs) — interface mechanisms such as infinite scroll, autoplay, and recommendation feeds — to extend user engagement beyond users’ original intentions. While emerging regulations require platforms to be more transparent about how they shape users’ behaviors and attitudes, how to concretely map this requirement to ACDPs remains unclear. We propose ACDPs Transparency Labels (ATLs): contextual, in-app interface notices that make ACDPs visible, explain how they may affect users’ digital well-being, and offer a clear call to action to regain control. We developed the ATL concept through two co-design workshops with 13 users of TikTok and YouTube. This poster reports the most generative findings from the workshops, distills them into five preliminary design guidelines, and position ATLs as a socio-technical proposal that connects user agency, platform design responsibility, and regulatory enforcement.

## CCS Concepts

• **Human-centered computing** → *Participatory design; Interaction design; Human computer interaction (HCI)*; • **Computing methodologies** → *Philosophical/theoretical foundations of artificial intelligence*.

## Keywords

Attention-Capture Damaging Patterns, Dark Patterns, Regulation

### ACM Reference Format:

Alberto Monge Roffarello, Riccardo Graziosi, and Luca Passarella. 2026. ACDPs Transparency Labels: Supporting User Agency Against Attention-Capture Damaging Patterns. In *Companion of the Computer-Supported Cooperative Work and Social Computing (CSCW Companion '26)*, October 10–14, 2026, Salt Lake City, UT, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3785651.3831549>

## 1 Introduction

Over the past decade, scholars have expressed increasing concern about unethical design practices adopted by digital service providers, framing them as deceptive or “dark” patterns [12] and studying their negative effects on users. Among these, Attention-Capture Damaging Patterns (ACDPs) [22] — an integrated design

apparatus that combines interaction-design components (infinite scroll, autoplay, pull-to-refresh) with an AI-driven algorithmic substrate (engagement-optimised recommendation, variable reward scheduling) — are widely used on social media platforms to prolong engagement beyond users’ original intentions, negatively affecting users’ attention, agency, and overall digital wellbeing [15, 18, 22]. As social media platforms deploy ACDPs at scale in social environments for advertising revenue, a structural tension becomes evident: maximizing engagement to sustain advertising income may come at the expense of how collective time and attention are allocated, and of users’ individual control over their own [22]. Existing remedies have so far proven inadequate. Tools for digital self-control are largely ineffective in the long term [19], mainly because they do not target ACDPs specifically and leave users to fight engagement-optimised systems with blunt timers and blockers [16]. In parallel, regulatory initiatives such as the European Digital Services Act (DSA) [6] and AI Act [10] have targeted overtly deceptive design practices, e.g., sneaking items into a shopping basket or making the unsubscribing process unnecessarily complex, leaving ACDPs in a regulatory gray zone [26]. ACDPs are indeed harder to frame: they persuade rather than deceive, exploiting users’ cognitive biases and psychological vulnerabilities to create attention traps that erode users’ sense of agency and lead them to lose sight of their goals [22].

To ease these tensions, we propose the concept of ACDPs Transparency Labels (ATLs). ATLs are contextual, in-app interface notices that inform users about the ACDPs being employed and their potential impact on digital wellbeing, while also providing practical calls to action that enable users to regain control. ATLs are intended to be integrated directly within the target interface. In this respect, they are inspired by transparency mechanisms such as cookie notices, but they extend the concept beyond data collection to interface-level persuasion. For example, the ATL in Figure 2 raises awareness about YouTube’s autoplay by showing how many videos have played in a row and offering a one-tap toggle to disable the feature. We developed the ATL concept through two co-design workshops with end users of TikTok and YouTube, chosen as familiar exemplars of social and video-streaming ecosystems in which ACDPs are well-documented. In each workshop, participants first reflected on their experiences with the target service, then collaboratively explored how ATLs could expose specific ACDPs and offer users meaningful ways to respond. This paper reports work-in-progress on three contributions: (i) we introduce ATLs as a transparency concept that targets ACDPs at the interface level on social platforms; (ii) we present the most generative findings from



This work is licensed under a Creative Commons Attribution 4.0 International License. *CSCW Companion '26*, Salt Lake City, UT, USA  
© 2026 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2378-0/26/10  
<https://doi.org/10.1145/3785651.3831549>

the co-design workshops; and (iii) we distill these into five preliminary design considerations, and discuss how ATLs may support user agency, help platforms build trust capital, and operationalize specific transparency obligations under the DSA and AI Act.

## 2 Background and Related Works

In the contemporary attention economy, social media platforms are engineered for promoting endless engagement via Attention-Capture Damaging Patterns (ACDPs), often at the expense of users' digital wellbeing. ACDPs are a subclass of deceptive design patterns that exploit cognitive biases to keep users engaged beyond their original intentions [22]. They include, among others, "seductive" patterns such as endless autoplay of videos, infinite ranked feeds, and guilty-pleasure recommendations [22], and have been described as strategies that "keep users occupied and entertained" (engaging) and "influence users' decision-making ability" (governing) [18]. Recent HCI work has documented the prevalence of ACDPs on social media [18, 22] and video-streaming platforms [4], as well as their consequences for users' digital wellbeing, including reduced attention, a diminished sense of agency, and a recurring feeling of regret after use [4, 15, 22].

Digital wellbeing — the impact of digital technologies on people's ability to live good and fulfilling lives [2] — has traditionally been framed as an individual responsibility, often reduced to engagement metrics and to the expectation that users will exercise self-control in the face of highly effective ACDPs. Recent work, however, challenges this view and shifts responsibility from users to platforms and designers, recognizing that individual-level interventions such as usage timers and app blockers cannot address an ecosystem permeated by ACDPs in the long term [9]. This shift has produced design frameworks, heuristics, and tools that help designers identify and avoid wellbeing-harming choices in the first place [20, 21]. Despite these advances, social media platforms often struggle to implement these principles in practice, as engagement-driven business models may conflict with the goal of promoting user wellbeing.

In parallel, regulators and policymakers have begun to address deceptive design practices by introducing new regulatory frameworks centered on transparency, fairness, and user protection [27]. European-level examples include the EU DSA [6], the EU AI Act [10], and the upcoming European Digital Fairness Act [11], while similar initiatives are ongoing in the US [3, 5] and globally [7]. When addressing challenges posed by ACDPs, however, current regulations primarily articulate high-level principles and obligations rather than prescribing concrete solutions. Indeed, ACDPs, from infinite scroll to algorithmically curated feeds, operate in a "grey zone:" they may subtly undermine autonomy and mental wellbeing without amounting to overt deception or clearly unlawful coercion, making them difficult to classify, detect, and pre-empt through legal criteria alone [26]. As a result, there is a growing consensus that legal requirements must be complemented by proactive design interventions that operationalize regulatory principles in practice [28, 29].

ATLs are intended to be one such intervention, sitting at the intersection of design and regulatory responses to ACDPs. Their primary aim is not to restrict use or impose hard limits, but rather

to foster awareness and gradual learning. At the same time, they do not require platforms to radically redesign their services or abandon engagement-based features. Instead, ATLs preserve existing interface structures while making ACDPs more visible, understandable, and accountable to users — and, in doing so, give regulators a concrete interface-level mechanism through which the transparency and user-protection obligations of emerging regulations can be enacted.

## 3 Method

To develop the concept of ATLs, we conducted two co-design workshops, each focused on a different social media platform, i.e., TikTok (n=6) and YouTube (n=7). TikTok was chosen given its widespread adoption among adolescents and young adults and its well-documented use of ACDPs [18, 19]. Recent EU investigations under the DSA and AI Act have further scrutinized TikTok's interface for potentially manipulative design choices [25]. YouTube was selected as a complementary platform reaching a substantially broader age range. YouTube is designated as a Very Large Online Platform under the DSA, making it directly subject to the transparency and risk-mitigation obligations, and YouTube's design patterns have been shown to undermine users' sense of agency and autonomy [4, 15].

A total of 13 participants (8 women, 5 men) were recruited through university mailing lists and social media in Italy and screened to ensure regular use of the targeted platform; ages ranged from 18 to 55+, with most participants in the 18–24 and 25–34 brackets. Each workshop lasted approximately two and a half hours and combined a focus group eliciting participants' perceptions of the target platform and its ACDPs, followed by a co-design session in which participants drafted and discussed ATLs. In the co-design phase, a researcher introduced a set of ACDPs selected for the target platform using visual examples, along with the general idea of ATLs. We deliberately provided high-level, generic explanations to avoid biasing participants, and used familiar examples from other domains, such as cookie notices and nutrition labels, to clarify the concept. Participants then collaboratively designed ATLs on a Figma template seeded with platform screenshots and ACDP examples, exploring tone of language, timing of appearance, and the degree of user control the labels should afford. Participants worked on researcher-provided laptops running a shared Figma template; for those without prior Figma experience, a researcher operated the tool on their behalf, translating verbal input and sketches into mockups. Workshops were conducted in Italian and audio-recorded. Participants produced 23 ATL mockups in total (11 for TikTok and 12 for YouTube). Two researchers analyzed transcripts, mockups, and facilitator notes through reflexive thematic analysis [1], combining inductive coding with a deductive scaffold drawn from prior taxonomies of ACDPs [18, 22]. Quotes and materials were translated into English by the authors for inclusion in this paper. Materials are openly available at [https://osf.io/m5jw6/overview?view\\_only=962d3134f74745b997ecdbcd3ca88483](https://osf.io/m5jw6/overview?view_only=962d3134f74745b997ecdbcd3ca88483).

The study was compliant with Politecnico di Torino's ethical guidelines for human research; all participants gave written informed consent prior to participation and could withdraw at any time. Recordings and transcripts were anonymized, and participants are referred to only by anonymized identifiers throughout.

## 4 Findings and Design Considerations

Across the two workshops, participants converged on a shared understanding of how ACDPs shape their behaviour, often describing an “autopilot” mode of engagement on TikTok and YouTube that they realized only in retrospect — “I’ll start scrolling, and then I look up and realize a whole hour has gone by without me noticing” (P3). Through discussion, they recognized features such as infinite scroll, autoplay, and personalized recommendations as engineered to prolong engagement, and not as accidental. In this section, we distill the most generative themes from the workshops into five preliminary design considerations for ATLs on social media platforms (Table 1).

According to our participants, ATLs should be coherent with the look and feel of the platform on which they appear (DC1). This was seen as essential for user receptivity: a label that visually “belongs” to the interface is more likely to be noticed, trusted, and engaged with. This perspective positions ATLs as an alternative to external digital self-control tools and highlights the need for regulatory approaches that require technology companies themselves to integrate such mechanisms, much like the standardized inclusion of cookie notices. Within this framing, ATLs can be understood as *internal design interventions*: following Lukoff et al.’s categorization [15], mechanisms that reshape the experience from within a platform show particular promise for strengthening user agency. The co-designed concepts reflected this preference for visual coherence. For example, participants suggested using TikTok’s color palette to ensure visual continuity with the app’s branding (Figure 1), while maintaining sufficient contrast for the label to remain noticeable.




**Figure 1:** Example of a “do” and “don’t” for DC1 from two mockups drafted in the TikTok workshop. The left mockup adopts TikTok’s typography and color palette; the right uses alarm-style colors that clash with the platform’s look and feel and risk feeling judgmental.

The tone of the labels’ messaging emerged as a second major point of emphasis (DC2). Participants consistently advocated for language that raises awareness without invoking guilt or judgment, favoring phrasing that is neutral and empathetic. As one participant noted, if a ATL “made me feel guilty or at fault, absolutely no, I would not find it useful,” whereas something that simply provides information “without limiting [the experience]” would be perceived as genuinely helpful (P9). Participants also preferred questions over directives — prompts such as “Why am I seeing these suggestions?” were proposed to invite reflection rather than alarm. Furthermore, despite being largely criticized in the digital wellbeing literature [8, 9, 23], screen time emerged across both workshops as a potential informative hook for the labels. Participants noted that they often realized only retrospectively how much




**Figure 2:** An ATL for YouTube developed by one of the participants that warns that autoplay is currently activated (A); tapping it opens an explanatory panel showing how many videos in a row the user has watched, along with a call to action to disable autoplay (B).

time had passed — “five minutes or half an hour, it does not matter, I am not able to realize it” (P5) — and acknowledged that even “just one or two hours a day” felt like time that could be better spent (P7). Time spent was therefore widely perceived as the most tangible indicator of diminished control and a powerful trigger for self-reflection. Importantly, however, participants differentiated time-spent considerations from the concept of time modeled by traditional digital self-control tools. They described time not as a constraint to be enforced, but as a self-reflection opportunity that should be surfaced in a non-intrusive manner and contextualized through explanations of the ACDPs driving it.

Participants underscored the importance of when and how often ATLs should appear to maximize their impact (DC3): a well-timed ATL that surfaces precisely when an ACDP is at play can “shake” users out of an engagement trance, whereas a poorly timed one is likely to be ignored. Timing was closely linked to frequency. Participants cautioned that if ATLs appear too often or at fixed intervals, they risk becoming predictable and annoying. The consensus was that transparency cues should be shown at a moderate and context-sensitive frequency: frequent enough to be impactful, yet not so repetitive that they fade into background noise. Some participants even proposed incorporating a degree of randomness or “controlled irregularity” (P12) into the schedule to prevent users from anticipating the label. Participants also stressed that ATLs should not remain static, but instead be dynamic and adaptive to avoid fading into the background through repeated exposure (DC4). They recognized the risk of habituation, noting that a repetitive, decontextualized warning would eventually be dismissed automatically. To counter this, participants proposed varying the label’s visual form, textual content, or tone to maintain a “constant impact on attention” (P1). This resonates with a broader desire for labels that “learn” from user behavior — for instance, by escalating their messaging or visual salience when gentler cues do not prompt reflection. This stands in contrast to existing label-based approaches such as cookie notices, which typically appear only once, outside the core user task, and rapidly lose relevance. It is also aligned with evidence showing that varying the form of digital self-control interventions over time can enhance their effectiveness [13]. Beyond simply informing users, participants emphasized that an effective ATL should empower them to take action over ACDPs (DC5). Direct call-to-actions — be it a button to disable autoplay (Figure 2, B) or a setting that modifies how recommendations are generated — were seen as crucial.

**Table 1: Overview of the five design considerations (DC) for ATLs emerging from our work.**

| Design Consideration             | Description                                                                                                                                                                                                                                                                         |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>DC1 - Visual identity</b>     | ATLs should respect the platform’s visual identity by adhering to its brand style, including typography, color palette, and interaction patterns.                                                                                                                                   |
| <b>DC2 - Language</b>            | ATLs should adopt a concise, jargon-free tone while supporting meaningful reflection via informative hooks. They should avoid blaming or shaming users and emphasize agency by stating that users can modify the feature.                                                           |
| <b>DC3 - Timing and Position</b> | ATLs should appear at meaningful interaction points, usually when an ACDP is activated or begins to shape user behaviour. Spatially, ATLs should appear close to the relevant interface element they refer to, supporting users in connecting the notice to the underlying feature. |
| <b>DC4 - Adaptability</b>        | ATLs are dynamic notices that should adapt their tone and visual design based on contextual and session-based information, in order to counter habituation.                                                                                                                         |
| <b>DC5 - Interactivity</b>       | ATLs should offer users a clear and immediate way to act on the relevant ACDP directly (e.g., disabling autoplay, pausing recommendations, switching to a non-profiled feed).                                                                                                       |

## 5 Discussion

ATLs could foster a win-win dynamic for both businesses and users. As individuals become more aware of the strategies employed to capture and monetize their attention, they may grow more sensitive to how different platforms shape their behaviour, and may factor a platform’s respect for their wellbeing into their usage decisions. Transparency about ACDPs can therefore function as a signal of accountability and ethical commitment, and integrating ATLs may be understood not merely as a compliance effort but as a strategic investment in trust capital that strengthens user satisfaction and long-term retention. The framing shares analogies with nutrition labels in the food domain: rather than eliminating unhealthy products, such labels reframed competition around informed consumer choice. This dynamic may become particularly relevant as alternative and decentralized social networking protocols, such as Bluesky and Mastodon, gain mainstream adoption [14, 17]: as switching costs fall and competition opens up, differentiation may no longer revolve solely around features or network effects, but also around perceived alignment with users’ values and well-being.

Beyond that, transparency is a common thread across recent regulatory frameworks for online platforms [6, 10]. Because the ATL concept is grounded in transparency, it may be particularly well suited to supporting compliance with several specific provisions. For instance, DSA Article 27 requires platforms to disclose the main parameters of recommender systems, and Article 38 requires Very Large Online Platforms to offer at least one non-profiled option, both of which fit naturally with an in-feed ATL that points at the personalization mechanism and exposes the non-profiled alternative. Similarly, DSA Articles 34–35 treat negative effects on physical and mental well-being arising “from online interface design that may stimulate behavioural addictions” as systemic risks requiring mitigation. In this respect, ATLs like the ones in Figure 2 and Figure 1 could represent a concrete mitigation strategy that platforms can adopt to fulfill this obligation at the interface level. ATLs can be similarly mapped onto AI Act provisions on AI disclosure and manipulation (Articles 5 and 50) and DSA Article 28 on the safety of minors. For instance, social media platforms such as TikTok allow users to apply AI-based beauty filters to photos and videos. Because such filters can promote unrealistic beauty standards — an effect

to which younger users may be especially vulnerable — platforms could deploy ATLs that disclose when a filter has been applied and clarify how the image has been altered, helping young users distinguish between reality and digitally manipulated representations.

## 6 Risks, Caveats, and Future Works

This paper introduced the concept of ATLs: contextual, interface-level notices designed to make ACDPs visible, understandable, and steerable. A number of open questions and limitations remain. First, the concept itself requires further investigation and elaboration, and whether ATLs can meaningfully influence user behavior — for example, by reducing time spent in infinite scrolls — remains an open empirical question. Second, cookie notices, our closest analogy, are a well-documented transparency failure: habituation, dark-pattern consent flows, and minimal compliance have eroded their usefulness [24], and an ATL mandate without robust specification could repeat this trajectory. Third, platforms whose revenue depends on engagement may have an incentive to comply with such a mandate in the most minimal way that survives audit; how ATLs should be specified — by regulators, third-party auditors, or independent design bodies — to be robust to adversarial implementation is an open question we hope to pursue in future work.

## References

- [1] Virginia Braun and Victoria Clarke. 2023. Thematic analysis. In *APA Handbook of Research Methods in Psychology: Research Designs: Quantitative, Qualitative, Neuropsychological, and Biological* (2 ed.), Harris Cooper, Marc N. Coutanche, Linda M. McMullen, Abigail T. Panter, David Rindskopf, and Kenneth J. Sher (Eds.). American Psychological Association, 65–81. doi:10.1037/0000319-004
- [2] Christopher Burr, Mariarosaria Taddeo, and Luciano Floridi. 2020. The Ethics of Digital Well-Being: A Thematic Review. *Science and Engineering Ethics* (2020), 2313–2343. doi:10.1007/s11948-020-00175-8
- [3] California Department of Justice, Office of the Attorney General. 2024. California Consumer Privacy Act (CCPA). <https://oag.ca.gov/privacy/ccpa>. Accessed 2025-01-24.
- [4] Akash Chaudhary, Jaivrat Saroha, Kyzyl Monteiro, Angus G. Forbes, and Aman Parnami. 2022. “Are You Still Watching?”: Exploring Unintended User Behaviors and Dark Patterns on Video Streaming Platforms. In *Proceedings of the 2022 ACM Designing Interactive Systems Conference (Virtual Event, Australia) (DIS ’22)*. Association for Computing Machinery, New York, NY, USA, 776–791. doi:10.1145/3532106.3533562
- [5] Colorado Department of Law, Colorado Attorney General’s Office. 2025. Colorado Privacy Act Resource. <https://coag.gov/resources/colorado-privacy-act/>. Accessed: 2025-08-26.

- [6] European Commission. 2022. The Digital Services Act package. Accessed: 2024-07-11. <https://digital-strategy.ec.europa.eu/en/policies/digital-services-act-package>
- [7] Department of Consumer Affairs, Ministry of Consumer Affairs, Food & Public Distribution, Government of India. 2025. Consumer Protection Act(s). <https://consumeraffairs.gov.in/pages/consumer-protection-acts>. Accessed 2025-01-24.
- [8] Niall Docherty. 2021. Digital self-control and the neoliberalization of social media well-being. *International Journal of Communication* 15 (2021), 3827–3846. Open Access. <https://ijoc.org/index.php/ijoc/article/view/17721>
- [9] Niall Docherty and Asia J. Biega. 2022. (Re)Politicizing Digital Well-Being: Beyond User Engagements. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 573, 13 pages. doi:10.1145/3491102.3501857
- [10] European Commission. 2024. Artificial Intelligence Act. Accessed: 2024-09-11. <https://artificialintelligenceact.eu/>
- [11] European Commission. 2025. Digital Fairness Act: Call for Evidence and Public Consultation. Official Website. Accessed: 2026-01-09. Reference Ares(2025)5829481. Part of the Fitness Check on EU consumer law.. [https://ec.europa.eu/info/law/better-regulation/have-your-say/initiatives/14622-Digital-Fairness-Act\\_en](https://ec.europa.eu/info/law/better-regulation/have-your-say/initiatives/14622-Digital-Fairness-Act_en)
- [12] 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). Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3173574.3174108
- [13] Geza Kovacs, Zhengxuan Wu, and Michael S. Bernstein. 2018. Rotating Online Behavior Change Interventions Increases Effectiveness But Also Increases Attrition. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 95 (Nov. 2018), 25 pages. doi:10.1145/3274364
- [14] Yuhan Liu, Emmy Song, Owen Xingjian Zhang, Jewel Merriman, Lei Zhang, and Andrés Monroy-Hernández. 2025. Understanding Decentralized Social Feed Curation on Mastodon. *Proc. ACM Hum.-Comput. Interact.* 9, 7, Article CSCW507 (Oct. 2025), 25 pages. doi:10.1145/3757688
- [15] Kai Lukoff, Ulrik Lyngs, Himanshu Zade, J. Vera Liao, James Choi, Kaiyue Fan, Sean A. Munson, and Alexis Hiniker. 2021. How the Design of YouTube Influences User Sense of Agency. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 368, 17 pages. doi:10.1145/3411764.3445467
- [16] Ulrik Lyngs, Kai Lukoff, Petr Slovak, Reuben Binns, Adam Slack, Michael Inzlicht, Max Van Kleek, and Nigel Shadbolt. 2019. Self-Control in Cyberspace: Applying Dual Systems Theory to a Review of Digital Self-Control Tools. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3290605.3300361
- [17] Erika Melder, Ada Lerner, and Michael Ann DeVito. 2025. “A Blocklist is a Boundary”: Tensions between Community Protection and Mutual Aid on Federated Social Networks. *Proc. ACM Hum.-Comput. Interact.* 9, 2, Article CSCW021 (May 2025), 30 pages. doi:10.1145/3710919
- [18] Thomas Mildner, Gian-Luca Savino, Philip R. Doyle, Benjamin R. Cowan, and Rainer Malaka. 2023. About Engaging and Governing Strategies: A Thematic Analysis of Dark Patterns in Social Networking Services. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 192, 15 pages. doi:10.1145/3544548.3580695
- [19] Alberto Monge Roffarello and Luigi De Russis. 2023. Achieving Digital Wellbeing Through Digital Self-control Tools: A Systematic Review and Meta-analysis. *ACM Trans. Comput.-Hum. Interact.* 30, 4, Article 53 (Sept. 2023), 66 pages. doi:10.1145/3571810
- [20] Alberto Monge Roffarello, Luigi De Russis, and Kai Lukoff. 2025. The Digital Attention Heuristics: Supporting the User's Attention by Design. *ACM Trans. Comput.-Hum. Interact.* 32, 4, Article 38 (Aug. 2025), 41 pages. doi:10.1145/3725215
- [21] Alberto Monge Roffarello, Luigi De Russis, and Massimiliano Pellegrino. 2024. Digital Wellbeing Lens: Design Interfaces That Respect User Attention. In *Proceedings of the 2024 International Conference on Advanced Visual Interfaces* (Arenzano, Genoa, Italy) (AVI '24). Association for Computing Machinery, New York, NY, USA, Article 51, 5 pages. doi:10.1145/3656650.3656674
- [22] Alberto Monge Roffarello, Kai Lukoff, and Luigi De Russis. 2023. Defining and Identifying Attention Capture Deceptive Designs in Digital Interfaces. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 194, 19 pages. doi:10.1145/3544548.3580729
- [23] Alberto Monge Roffarello, Monica Molino, and Luigi De Russis. 2026. What is Digital Wellbeing? A Leverage Points Framework to Guide Research and Action. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems* (CHI '26). Association for Computing Machinery, New York, NY, USA, Article 481, 21 pages. doi:10.1145/3772318.3793192
- [24] Midas Nouwens, Ilaria Liscardi, 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. doi:10.1145/3313831.3376321
- [25] Official Journal of the European Union. 2024. Case DSA.100121 – TikTok Lite Rewards programme. [https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ:C\\_202405942](https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ:C_202405942). Accessed: 2025-07-22.
- [26] Lorena Sánchez Chamorro, Kerstin Bongard-Blanchy, and Vincent Koenig. 2023. Ethical Tensions in UX Design Practice: Exploring the Fine Line Between Persuasion and Manipulation in Online Interfaces. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference* (Pittsburgh, PA, USA) (DIS '23). Association for Computing Machinery, New York, NY, USA, 2408–2422. doi:10.1145/3563657.3596013
- [27] C. Santos, V. Morozovaite, and S. De Conca. 2025. No harm no foul: how harms caused by dark patterns are conceptualised and tackled under EU data protection, consumer and competition laws. *Information & Communications Technology Law* 34, 3 (2025), 329–375. doi:10.1080/13600834.2025.2461958
- [28] Caroline Sindere and Claire Pershan. 2022. Why Europe's Digital Services Act Regulators Need Design Expertise. *Tech Policy Press* (2022). Accessed: 2025-02-10. <https://www.techpolicy.press/why-europes-digital-services-act-regulators-need-design-expertise/>
- [29] Weiwei Yi and Zihao Li. 2025. Mapping the scholarship of the regulation of dark patterns: A systematic review of concepts, regulatory paradigms, and solutions from law and HCI perspectives. *Computer Law & Security Review* 59 (2025), 106225. doi:10.1016/j.clsr.2025.106225