

Consent Cannot Be Delegated

Why AI shouldn't decide what we share with AI – and what to do instead

Author

Rafael Vrecar · CTSi · [ORCID 0000-0001-6572-5254](#)

Sources and co-authors

This white paper condenses two of the author's own publications for a broader readership:

(1) Position paper (main basis). Vrecar, R. (2026). *Why AI Shouldn't Decide What We Share with AI and How We Can Do Better* [position paper and presentation], workshop at CHI '26, Barcelona. Sole author. Open access under CC BY-NC-ND 4.0: hdl.handle.net/20.500.12708/228428 [1]

(2) Empirical basis (co-authored). Vrecar, R., Geld, B., Dobrosovestnova, A., Čolakovová, M. G., Weiss, A. (2026). *Do Cookies Taste Different? The Impact of Options and Website Type When Setting Cookie Preferences*. In: CHIRA 2025, CCIS 2834, pp. 56–77, Springer Nature Switzerland, [doi:10.1007/978-3-032-16448-3_4](https://doi.org/10.1007/978-3-032-16448-3_4) [2]. All empirical results in this white paper come from this joint work with Barbara Geld, Anna Dobrosovestnova, Michaela Gaea Čolakovová and Astrid Weiss.

ABSTRACT

Cookie banners and privacy labels are meant to let us make informed decisions about our data. So far, both fall short of that. A study with 136 valid responses shows that how people decide depends above all on which options a banner offers. This was the only statistically significant factor; the type of website, prior knowledge and age made no measurable difference. It seems obvious to hand such decisions over to AI systems. This white paper shows why that does not work legally: the GDPR understands consent as a freely given, informed and unambiguous statement by the data subjects themselves. If an AI system decides on their behalf, that condition is not met. On top of that, large language models are known to produce errors. As an alternative, this white paper presents a single consent window in the browser, which performed better than ordinary banners in a first, small and non-representative evaluation. AI should support the decision, not make it.

1 Introduction and background

Cookie banners have been part of daily browsing since the GDPR took effect in May 2018. They are meant to let us make an informed decision about our data. In practice, that rarely works. Many banners make rejecting more laborious than accepting [3, 4]. Some websites do not technically implement the choice that was made [5]. And dark patterns, design tricks that push people towards a particular answer, are widespread in banners [6, 7].

The second major push for transparency is privacy labels: short overviews of the data an app collects. Apple introduced them for the App Store in 2020 [8]. Here too the evidence is mixed.

Many users do not fully understand the labels because they are complex, use unfamiliar terms and arrange the information confusingly [9, 10]. They are not without effect, though: they influence whether someone installs an app [11]. Both instruments show the same pattern: they do have an effect, and they are still not properly understood.

That leads to the question at the centre of this white paper: if people are overwhelmed by such decisions, should AI take them over?

2 Facts and findings

2.1 What actually drives the decision

The basis is an online study using mock-up websites (a vignette study): 150 participants recruited via the Prolific platform, 136 of whose responses were valid [2]. Each of them worked through a news site, an online shop, a weather portal and an information site about cookies, deciding on the cookie banner every time, 544 decisions in total. Which options the banner offered changed randomly between three variants:

- *Variant A:* Accept all | Reject non-necessary | Customize settings
- *Variant B:* Accept all | Customize settings
- *Variant C:* Accept all | Close page

Of all the factors examined, only the variant of options had a statistically significant effect ($\chi^2 = 44.1$, $df = 2$, $p < 0.001$): with Variant A, the only one offering a one-click rejection, noticeably fewer participants accepted all cookies than would be expected under statistical independence; with Variants B and C, noticeably more did. Neither website type nor familiarity with the term “cookie” nor age had any measurable influence. Across all 544 decisions, 47.8 % went to “Accept all”. A typo had also been placed in the banner text deliberately; only 11.8 % of participants passed the corresponding attention check, and guessing cannot be ruled out.

2.2 What follows from this

- Where cookies can be rejected in one click, people use that. Without such a button, hardly anyone takes the detour through the settings. Most accept everything instead.
- Education alone does not change this. Those who could explain what cookies are decided just like everyone else.
- What matters is the design of the banner, not what users know. Other studies reach the same conclusion: a clearly visible “reject” noticeably changes the decision [12], personalised banners help people decide more knowledgeably [13].

2.3 A single consent window in the browser

This is where a prototype comes in that Andreas Schlößl developed and evaluated in his bachelor’s thesis, primarily supervised by Astrid Weiss and co-supervised by the author. The idea: instead of every website building its own banner, the browser provides one window that looks the same everywhere. Once people have understood how it works, they find their way on every site.

The results he reports are positive: participants found the window more credible and trustworthy

than a classical banner. On the System Usability Scale, a standard usability questionnaire, it reached 72.25 out of 100 points [14], and 80 % preferred it over the classical solution.

These numbers should be read with care, though. The evaluation was carried out within a bachelor's thesis and is not representative: ten participants, all male, aged between 18 and 30 ($M = 23.3$, $SD = 3.13$), half of them with a computer science background. A larger follow-up study is planned as part of the author's doctoral research [1].

3 Implications

3.1 Why AI cannot replace consent

The obvious answer to overwhelmed users would be to leave the decision to an AI system. The law does not allow this. Article 4(11) GDPR describes consent as a freely given, specific, informed and unambiguous indication of the data subject's wishes: they themselves signify agreement to the processing, through a statement or a clear affirmative action [15]. If an AI system decides on their behalf, none of these conditions is met, because the indication no longer comes from the person.

There is also a technical risk: large language models are known to produce inaccurate or fabricated answers [16]. Where legal and security consequences are involved, that weighs heavily. One could restrict such a system to a verified set of responses; whether that is enough for the variety of real cases is an open question.

3.2 What AI can help with instead

AI is useful here as an assistant that guides people through the decision rather than making it: explaining privacy policies in plain terms, clarifying terms from privacy labels, that is, addressing precisely the comprehension problems that are documented [9, 10]. The decision itself stays with the person.

3.3 What to do

A single standard. The first sensible step is a mandatory consent window that looks the same on every website. It makes the situation recognisable and is the precondition for setting preferences once and having them applied automatically.

Consistent terms and education. If privacy policies use the same terms and people are taught what those terms mean, that solves the comprehension problem instead of working around it.

Security matters too. Consent is not only about data protection. People who consent to tracking load 45 % more third-party scripts and 63 % more security-sensitive data flows on average [17]. A banner that encourages acceptance therefore also enlarges the attack surface.

Limits. The study used static mock-up websites and a non-representative sample; ten people took part in the prototype evaluation. Both results point in a direction, but do not yet prove it.

4 Conclusion

- **What matters is the options on offer.** Of all the factors examined, only the choice in the banner was statistically significant ($p < 0.001$), not what users know.
- **Consent cannot be handed over.** If an AI decides on our behalf, that is not consent under Article 4(11) GDPR, however good the system may be.
- **The way forward is a single standard.** A consent window that is the same everywhere, comprehensible language and AI in a supporting role enable informed decisions without taking them out of people's hands.

References

- [1] Rafael Vrecar. Why AI Shouldn't Decide What We Share with AI and How We Can Do Better. Position paper and presentation, workshop at CHI '26, Barcelona, Spain, May 2026. Open access, CC BY-NC-ND 4.0, <https://hdl.handle.net/20.500.12708/228428>.
- [2] Rafael Vrecar, Barbara Geld, Anna Dobrosovstnova, Michaela Gaea Čolaková, and Astrid Weiss. Do Cookies Taste Different? The Impact of Options and Website Type When Setting Cookie Preferences. In *Computer-Human Interaction Research and Applications (CHIRA 2025)*, volume 2834 of *Communications in Computer and Information Science*, pages 56–77, Cham, 2026. Springer Nature Switzerland.
- [3] Georgios Kampanos and Siamak F. Shahandashti. Accept All: The Landscape of Cookie Banners in Greece and the UK. In *ICT Systems Security and Privacy Protection*, pages 213–227, Cham, 2021. Springer International Publishing.
- [4] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. (Un)informed Consent: Studying GDPR Consent Notices in the Field. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (CCS '19)*, pages 973–990. ACM, 2019.
- [5] Célestin Matte, Nataliia Bielova, and Cristiana Santos. Do Cookie Banners Respect My Choice? Measuring Legal Compliance of Banners from IAB Europe's Transparency and Consent Framework. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 791–809, 2020.
- [6] Philip Hausner and Michael Gertz. Dark Patterns in the Interaction with Cookie Banners, 2021. Workshop "What Can CHI Do About Dark Patterns?" at CHI 2021.
- [7] Oksana Kulyk, Willard Rafnsson, Ida Marie Borberg, and Rene Hougard Pedersen. "So I Sold My Soul": Effects of Dark Patterns in Cookie Notices on End-User Behavior and Perceptions. In *Proceedings of the 2022 Symposium on Usable Security and Privacy (USEC)*. Internet Society, 2022.
- [8] Apple Inc. App Privacy Labels Now Live on the App Store. Apple Developer News, 2020. <https://developer.apple.com/news/?id=3wann9gh>.
- [9] Ishika Keswani, Kerick Walker, Adrian Clement, Eusila Kitur, Nannapas Wonghirundacha, Ryan Aubrey, Vivien Song, and Eleanor Birrell. User Understandings of Technical Terms in App Privacy Labels. In *Twenty-First Symposium on Usable Privacy and Security (SOUPS 2025)*, pages 279–298, 2025.
- [10] Shikun Zhang, Lily Klucinec, Kyerra Norton, Norman Sadeh, and Lorrie Faith Cranor. Exploring Expandable-Grid Designs to Make iOS App Privacy Labels More Usable. In *Twentieth Symposium on Usable Privacy and Security (SOUPS 2024)*, pages 139–157. USENIX Association, 2024.
- [11] David G. Balash, Mir Masood Ali, Chris Kanich, and Adam J. Aviv. "I Would Not Install an App with This Label": Privacy Label Impact on Risk Perception and Willingness to Install iOS Apps. In *Twentieth Symposium on Usable Privacy and Security (SOUPS 2024)*, pages 413–432. USENIX Association, 2024.
- [12] Nataliia Bielova, Laura Litvine, Anyisia Nguyen, Mariam Chammat, Vincent Toubiana, and Estelle Hary. The Effect of Design Patterns on (Present and Future) Cookie Consent Decisions. In *33rd USENIX Security Symposium (USENIX Security 24)*, pages 2813–2830. USENIX Association, 2024.
- [13] Tom Biselli, Laura Utz, and Christian Reuter. Supporting Informed Choices about Browser Cookies: The Impact of Personalised Cookie Banners. *Proceedings on Privacy Enhancing Technologies*, 2024(1):171–191, 2024.
- [14] James R. Lewis and Jeff Sauro. The Factor Structure of the System Usability Scale. In Masaaki Kurosu, editor, *Human Centered Design*, pages 94–103. Springer Berlin Heidelberg, 2009.
- [15] European Parliament and Council of the European Union. Regulation (EU) 2016/679 (General Data Protection Regulation). Official Journal of the European Union, 2016. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32016R0679>.

- [16] Gabrijela Perković, Antun Drobnjak, and Ivica Botički. Hallucinations in LLMs: Understanding and Addressing Challenges. In *2024 47th MIPRO ICT and Electronics Convention (MIPRO)*, pages 2084–2088, 2024.
- [17] David Klein, Marius Musch, Thomas Barber, Moritz Kopmann, and Martin Johns. Accept All Exploits: Exploring the Security Impact of Cookie Banners. In *Proceedings of the 38th Annual Computer Security Applications Conference (ACSAC '22)*, pages 911–922. ACM, 2022.

Acknowledgements

The prototype described in Section 2.3, and its evaluation, were built by Andreas Schlögl as part of his bachelor’s thesis at TU Wien, primarily supervised by Astrid Weiss and co-supervised by the author of this white paper. The figures given there (System Usability Scale, preference, sample) are from his thesis as well. The rights to the screenshots are his (© Andreas Schlögl); they are therefore not reproduced here, but the original position paper shows them.

The study reported in Section 2.1 was conducted together with Barbara Geld, Anna Dobrosovestnova, Michaela Gaea Čolakovová and Astrid Weiss and was published by Springer. Only individual figures and test results from that chapter are reported here: no passages of text, no artwork or tables, and no charts derived from them. The full study, including its methodology, vignettes and contingency tables, is available under the DOI given above.

Rafael Vrecar is a recipient of a DOC Fellowship of the Austrian Academy of Sciences (DOC/27292).

Use of AI tools

An AI assistant (Claude, Anthropic) was used to condense, translate, structure and simplify the wording of the author’s two underlying publications for the format of this series. The ideas, arguments and conclusions come from the referenced own work; all content, figures and references were checked by the author against the original publications. In the underlying publications, DeepL, ChatGPT, Apple Intelligence Writing Tools and Overleaf Writefull were used solely for linguistic refinement.

Suggested citation

Vrecar, R. (2026). *Consent Cannot Be Delegated*. CTSi White Paper No. 01. TU Wien.

Contact: Rafael Vrecar · rafael.vrecar@tuwien.ac.at · CTSi · cts.wien