AI-ASSISTED DESIGN FOR BETTER PRIVACY COMMUNICATION IN HUMAN-ROBOT CO-EXISTENCE

Hilja Autto / Helena Haapio / Yueh-Hsuan Weng

Hilja Autto, Doctoral Researcher, University of Lapland, Faculty of Law | Project Researcher, University of Vaasa, School of Accounting and Finance, Business Law, Yliopistonkatu 8, 96300 Rovaniemi, FI; hilja.autto@ulapland.fi; https://research.ulapland.fi/en/persons/hilja-autto

Helena Haapio, Professor of Practice, University of Lapland, Faculty of Law | Associate Professor, University of Vaasa, School of Accounting and Finance, Business Law | Contract Strategist, Lexpert Ltd, Ritarikatu 7 A 2, 00170 Helsinki, FI; helena.haapio@lexpert.com; http://www.uva.fi/en/profile/?view=104151 | http://www.lexpert.com

Yueh-Hsuan Weng, Associate Professor, Inamori Frontier Program, Institute of Advanced Study, Kyushu University, 744 Motooka, Nishi-ku, Fukuoka 819-0395, JP; weng.yueh.hsuan.647@m.kyushu-u.ac.jp; http://www.robolaw.asia

Keywords: ***Artificial Intelligence, Generative AI, Human-robot Interaction, Information Design, Privacy Communication, Social Robots***

Abstract: *Information design focuses on making communication as effective and understandable as possible. Using elements such as plain language and design patterns it has proven successful in various contexts, including legal communication. Generative AI tools such as ChatGPT can help make the design process more efficient. This paper explores how these tools can be used to improve privacy-related communication, with a particular focus on human-robot interactions. To demonstrate the possibilities, we used OpenAI's ChatGPT and DeepL Write to create the conversational text of a wireframe sketch for a human-friendly privacy policy for RoBot, a (fictional) friendly and playful social robot.*

1. Privacy Communication Challenges in Human-robot Interaction

In many applications, human-robot interaction raises privacy concerns that may inhibit the wider use of robots. Traditional legal language may not be the best way to address these concerns – at least not if we want to communicate in a user-friendly and effective way. Such language may work well for communication between lawyers, but it is problematic for communication to others.

Robotics will change our daily lives, and we may face new challenges in this context: people may not understand that there are robots among them collecting privacy-sensitive data. Legal language is not the easiest way to communicate this and obtain informed consent. In the EU, the GDPR (General Data Protection Regulation) aims to give individuals access to and control over how their personal data is used. The GDPR includes the principle of transparency: it requires data controllers to disclose full and accurate information about their practices, their purposes and the rights of individuals (called data subjects), so that they can understand, consent to and, if necessary, challenge the processing of their data. Currently, even when privacy communications contain these elements, they do not necessarily meet the ideal of transparency due to the lack of user-friendliness of the language and the design (or lack thereof) of conventional privacy policies.

Privacy communication in general seems to suffer from some common problems of legal language, such as complex wording, abstract terminology, striving for precision while containing ambiguity, and irrelevant information, which often leads to rejection by the addressee [Mattila, 2017, 60, Passera 2017, 50]. Another major challenge is that people tend to ignore much of this information, especially in the online environment [Salo/Haapio 2017, 2]. Language is not the only cause, of course. Research and practice tell us that there are many reasons why traditional ways of providing information to individuals do not work [Haapio et al. 2018, Rossi et

al. 2019]. Privacy notices tend to be extremely long and take a long time to read. Too much information can lead to information overload and discourage users from reading. Privacy notices can appear as a "wall of text" that is impenetrable to the human eye [Passera 2017]. Information hierarchy, meaningful headings and visual organization are often lacking. And the list goes on. The good news is that a number of solutions to these communication problems already exist. Advances in robotics and AI bring new challenges, but also new opportunities.

2. Next Generation Robots

The term "robot" was coined by the Czech novelist Karel Čapek in his work Rossum's Universal Robots, meaning "forced labor" or "dreary work" [Čapek 1920]. This definition can still be used to describe the robots used in the 20th Century, especially in the manufacturing area. Then, AI became a transition point from the 20th Century robots that work in closed, restricted areas to the next-generation robots that interact with humans to help perform daily life tasks, for example in public places like hospitals, schools and shopping malls, and even in private places like our homes. These next-generation robots are highly adaptable to sophisticated human living environments and aim to serve the Human-Robot Co-Existence Society. Functionality, mobility and sociability are key features of the next generation of robots. They may be equipped with sensors, cameras and microphones to gather information about their surroundings and about people in order to function optimally in human communities. These robots are powered by artificial intelligence, which enhances their ability to function socially [Weng/Chen/Sun 2009].

Currently, there are many research projects related to the co-evolution of AI and robots. One example is the Moonshot Goal 3 Research and Development Program funded by the Japan Science and Technology Agency (JST). Under the JST Moonshot Goal 3, the team led by Tohoku University proposes to create a next-generation healthcare robot that uses AI for high-level interaction to improve patient self-efficacy during robotic care. While there are many benefits, the AI features of robots' social interactions also raise a number of legal and ethical concerns [Weng et al. 2020]. Humans may not always recognize and understand robotic data collection capabilities and privacy and data protection issues; these are not as obvious as on a computer, where you click "I agree" on a privacy notice. If they do not understand, they may not be able to assess the impact of the data collection or give informed consent. In addition, as the Moonshot research has shown, in the presence of multiple people, a robot may misinterpret consent as having been given by people who did not actually give it. This raises the question of how to ensure data security and effective privacy communication in human-robot interactions, and how to ensure informed consent. Research conducted by one of the authors of this paper (Weng) shows that the physical presence of future robots in the home poses privacy risks by facilitating the collection of personal data, while existing privacy regulations have not kept pace with these advances [Weng et al. 2020]. In addition to data privacy and security, there are other ethical and legal issues related to robots. Sometimes enhancing human self-efficacy may mean that the robot has to hide the truth of patients' real health situations during healthcare [Weng/Hirata 2022]. In addition, the emotional capabilities of robots, such as their ability to recognize human emotions and adapt their actions accordingly, can raise ethical concerns. These concerns encompass the potential for AI to manipulate human emotions as well as the gathering and use of such data, especially within the context of marketing [McStay 2020].

3. Responding to the Challenges through Information Design and Plain Language

Over the past decade, there has been a growing interest in a different, modern way of communicating legal information, one that focuses on the needs of the users, especially those without legal training. Designers and proactive legal thinkers have developed legal information design, or, at a more general level, legal design, which applies design to broader issues in the legal domain to prevent and solve legal problems [Passera 2017]. These approaches give priority to the perspective of citizens, consumers, businesses, etc., the "users", and not just to the perspective of lawyers and judges. Information design is all about communicating in the most

effective and understandable way possible. Hayhoe defines information design as "the process of identifying, selecting, organizing, synthesizing, and presenting information to an audience so that the audience can use it effectively and efficiently to achieve a particular purpose" [Hayhoe 2012].

Plain language and language simplification are an important part of information design. According to the International Plain Language Federation, "a communication is in plain language if its wording, structure, and design are so clear that the intended readers can easily find what they need, understand what they find, and use that information".[1] Plain language means paying attention not only to wording and structure, but also to document design (font, line length, white space, etc.). It includes using short sections or breaking up longer sections, using meaningful headings, linking group ideas, arranging sections in a logical order, adding a summary at the beginning of documents, and using examples, tables and charts. Several case studies have demonstrated that readers clearly prefer plain language to legal and official language: they understand it better and faster and are more likely to read and follow it [Butt 2001]. It also almost always improves the content, and readers are more likely to enjoy plain language documents [Kimble 1996–1997]. When information is given to the reader in appropriate doses, the reader can select the information that is relevant. This helps avoid information overload and self-defeating over-clarity, trying to cover every possible detail. This way, the reader is not intimidated by obscurity and size of the information. The principles of plain language can also be applied to legal and technical communication [Kimble 1996–1997].

For architects, interaction designers and software engineers, design patterns are a common way to share transferable solutions to common problems. Information designers also use design patterns, to solve usability and comprehension problems that often occur in complex communication. Design patterns provide a systematic way of identifying, collecting and sharing good practice. They can be said to provide a *lingua franca*, common language, for sharing knowledge and innovation on similar challenges between and across disciplines and domains [Waller et al. 2016]. The concept of design patterns is relatively new in law and legal communication, while templates, model contracts, standard terms and boilerplate are well known. A design pattern is not a template intended to be copied and pasted. Instead, it illustrates a model approach and provides examples that help identify a problem and develop an effective solution. Haapio and Hagan argue that lawyers should develop and share design patterns, "to promote high-quality, efficient and consistent work and language use between collaborators from different disciplines" [Haapio/Hagan 2016]. Design patterns can aid communication, allowing people working on similar challenges to share best practices and to collaborate more smoothly, also in multi-disciplinary settings. Examples and prototypes of early legal design pattern libraries can be found at Stanford Legal Design Lab's Legal Communication Design web pages,[2] and contract design patterns can be found at WorldCC Contract Design Pattern Library.[3] The latter currently contains ten different pattern families: emphasis, explainers, layering, layout, navigation, organizing, reviewing, summarizing, tone of voice and visuals.

4. Information Design for Better Privacy Communication

The shortcomings of traditional privacy communication are well known and documented [Haapio et al. 2018]. In recent years, design patterns have found their way into privacy design [Haapio et al. 2018, Rossi et al. 2019]. Several initiatives have explored icons to summarize data practices. For example, the DaPIS project has developed a set of data protection icons to help communicate privacy concepts [Cavazzuti 2021]. Visuals and other design patterns can help individuals to be more aware of their rights, while also helping data controllers to be transparent, as required by the GDPR.

A well-known early adopter of a layered approach is the Creative Commons (CC) license system.[4] The icons indicate that the material is covered by CC licences. It provides a link to a "human-readable summary" of the

1 International Plain Language Federation n.d.
2 Stanford Legal Design Lab n.d.
3 WorldCC Foundation/Passera/Haapio n.d.
4 WorldCC Foundation/Passera/Haapio n.d.

licence in plain language, which explains what can and cannot be done with the licensed material. By clicking on the link, the user can read the full text of the licence. There is also a machine-readable action layer: a summary of the key freedoms granted and obligations imposed, presented in a format that applications, search engines and other kinds of technology can act upon. Different layers can also be used to address different audiences: the first layer can be for those who want an overview, while the second layer can be for those who, for various reasons, prefer the full text (e.g. lawyers, judges or regulators) [Rossi et al. 2019].

Another widely-cited example of a layered approach is the Juro Privacy Policy. On its home page, https://juro.com, Juro provides a brief overview called "Your privacy at a glance". Further information is made available in manageable bits using a design pattern known as accordion: key information is presented at the top, which, when clicked, displays further details inside expandable panels. The Juro Privacy Policy also uses further design patterns and pattern families, such as tone of voice (conversational style), navigation and organizing (FAQ-style headings) and visuals (companion icons and a timeline showing the data collection process). It illustrates the core idea of layering: overview first, details only later. The layered approach can also help to pursue the GDPR's goals: transparency and user-friendliness.[5]

5. AI-Assisted Privacy Communication Design

Generative AI tools, such as OpenAI's ChatGPT, can be and have been used to accelerate and enhance the adoption of information design.[6] For example, they can help with content, language, structure and presentation. In the case of highly text-intensive documents, the capabilities of these tools are particularly useful: they can translate legalese into conversational and understandable language, and suggest summaries and headings. Even before ChatGPT was launched, there were AI-powered tools, such as DeepL Translator and DeepL Write:[7] these can help information designers translate text from one language to another, correct grammar and punctuation errors, adjust a tone of voice, and express nuance.

To explore the possibilities of these tools, we tested them in the context of privacy communication for RoBot, a (fictional) friendly and playful social robot. Our goal was to see how, using earlier work as inspiration, AI tools could help us apply design tools and techniques that have proven useful in the context of privacy communication. Building on Creative Commons licenses and the Juro privacy policy, we wanted to create a privacy policy that was both user-friendly and transparent and used a conversational and cheerful tone. It was designed to work on small screens, so people can easily browse the privacy policy on their phone. We were aware of the fact that on small screens, it is useful to rely on interactive elements such as accordions, links and panels where users can open and collapse information layers as needed, depending on their tasks, needs and preferences. In the layered approach, information is revealed step by step to the reader.[8] When opened, the content of a pop-up note or an accordion tab coexists on the same plane as the main text; when closed, their content is fully filtered out.

Based on these ideas, we applied the layered approach and created a rough idea of the look and feel of the privacy policy. Then, using OpenAI's ChatGPT, we created the conversational text of a wireframe sketch where the robot introduces itself, what data it collects, why it collects it, and how to opt out of data collection. In the beginning of our experiment, we gave ChatGPT a brief context in our prompt: "The RoBot Privacy Policy describes how RoBot collects, uses, and discloses personal information and the choices available to users. The privacy policy is written in plain language. Can you suggest a privacy policy like this?"[9] As shown in the conversation, ChatGPT's initial suggestions were somewhat vague, so we prompted it to further develop the text. To refine it, we provided more context about the capabilities of next-generation robots: "Functionality,

5 Juro n.d.; Passera n.d.; Juro/Passera n.d.

6 See, e.g., Corrales Compagnucci, M./Fenwick, M./Haapio, H. (2022) and Autto/Haapio/Nuottila forthcoming.

7 DeepL Translator, https://www.deepl.com/translator; DeepL Write, https://www.deepl.com/write.

8 Juro n.d.; Passera n.d.; for their free privacy notice template, see *Juro-Privacy*, n.d.

9 The full conversation, including the prompts used, can be viewed at "User-friendly RoBot Privacy Policy", https://chat.openai.com/share/ec505889-03e9-4479-8704-8fbc98c741e1.

mobility and sociability are key features of the next generation of robots. They may be equipped with sensors, cameras and microphones to gather information about their surroundings and about people in order to function optimally in human communities." We asked ChatGPT to use this information to complete the text. ChatGPT's revised response was more informative. We then asked it to focus specifically on the robot's data collection capabilities and to present this information in a conversational way, as if it were coming directly from the robot. This was done to make the text easier to approach and more understandable. We slightly shortened, corrected and supplemented the output of ChatGPT and merged it with our earlier draft. Finally, we ran the text through DeepL Write to improve its grammar. After a few rounds of prompts shown in the conversation, we ended up with a sketch for a human-friendly privacy policy for RoBot, an excerpt of which is shown in Figure 1. The Figure shows a combination of our own draft, combined with the output of two AI-powered tools, namely ChatGPT and DeepL Write.

The Privacy Policy is designed using the accordion pattern which allows the user to open and close layers of data. In the Figure, the "screen" has the "Why we collect your data" section open, and the other layers of data are hidden. It is assumed that the robot's data collection feature can be turned off, should the user so wish. For the purposes of our experiment, we also assumed that there is a company privacy policy in the background. For example, we assumed that when someone buys or rents a robot, he or she will first encounter a user-friendly Juro-style privacy policy, where the reader is not overwhelmed with data and a wall of text full of legalese. Instead, the reader can focus on one thing at a time, and a busy reader can quickly scan the most important points. Our sketch is not intended to be the final product, it is a wireframe for further development. We envision the functions of the display and the details of our wireframe sketch to go through a careful localization and compliance assessment before becoming a feature of the robot.

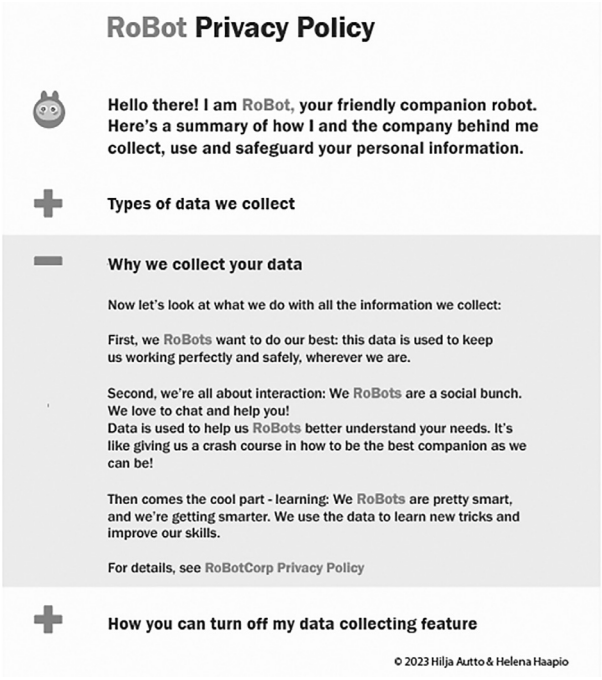

Figure 1: Excerpt from the wireframe sketch for a privacy policy for RoBot[10]

[10] RoBot is a (fictional) friendly and playful robot. Conversational text created using OpenAI's GPT-4 (ChatGPT September 25, 2023 version) and information design elements (layered layout, accordion, FAQ subheadings). Licensed under CC BY-NC

There are many reasons why we propose a layered approach to robot privacy policies. First, it may reduce the transparency problems associated with privacy communications. Second, layering makes the policy more understandable and user-friendly. It can even engage users who are not normally interested in reading about privacy. Perhaps the tone of the text and the adorable robot character can get them to pay attention and become interested and engaged. Either way, there is no question that the policy will be more appealing to most users than one presented in text-only and legalese. And when the design is user-centred, information transparency, fairness, privacy, and other values are more likely to be achieved. After all, privacy communications can and should be engaging and designed for the end users. Article 12 of the GDPR requires that privacy information be provided "in a concise, transparent, intelligible and easily accessible form, using clear and plain language" – all the more reason to adopt plain language and solutions that researchers and practitioners have developed to privacy communication problems, such as layered notices, colour coding and privacy icons [Haapio et al. 2018; Rossi et al. 2019].

6. What Does the Future Hold?

We are moving to a world that is no longer just documents and screens. Screens are getting smaller, and the use of virtual reality is becoming part of everyday life [Haapio et al. 2018; Rossi et al. 2019]. In this changing world, innovative and appropriate designs are needed to communicate privacy effectively and successfully. The tools we have presented can be applied to a wide range of situations beyond human-robot interaction, between professionals in different fields, as well as between professionals and laypeople. In this paper, our primary perspective is that of communicating privacy choices from robots to humans through written or digital privacy policies. We also see many other opportunities for the application of information design in the context of privacy communication. When humans use laptops and cell phones, their interactions are often based on information displayed on digital screens. For robots, the situation is different: the embodiment of robots allows them to convey infomation through a variety of visual, audio, and physical means. The embodiment of robots may raise new privacy issues [Weng et al. 2020], but the diversification of means of communication may also open up new possibilities for information design. Privacy can now be communicated beyond the documentary form: perhaps the robot can communicate about privacy through speech and answer users' questions [Waller/Passera/Haapio 2022].

Artificial intelligence (AI) has entered many industries, and the world of information design is no exception. Information designers and writers are often faced with tedious tasks such as creating the same text or graphics in multiple languages, translating legalese into plain language, and creating summaries of long texts. AI-empowered tools such as Google Translate, DeepL, GPT-4, and ChatGPT can help with such tasks. These tools can also help humans with creativity issues such as fear of blank paper and writer's block. The same technologies could also be used by readers to cut through the impenetrable walls of text and the fog of legalese: an "intelligent reader" app could provide plain-language translations to help readers better understand their privacy choices and their implications.

With design patterns, combined with auto-designers and text generation tools that have been around for a while, the possibilities seem endless [Corrales Compagnucci/Fenwick/Haapio 2022]. Perhaps soon we will be able to just give a short description of what we want, such as a privacy policy with a friendly robot image, skimmable headings, and an accordion-based three-layer layout, and say this in a prompt to OpenAI's Codex or ChatGPT, which will then generate the code to produce what we want. We could also add a fourth, machine-readable layer, similar to that used in Creative Commons licenses. For now, these developments are in the early stages. But it may not be long before we have access to tools that make it easy to translate text or voice prompts into user-friendly privacy communications that meet regulatory requirements and include plain language, layered layouts, meaningful headings, and explanatory diagrams: privacy policies that people can and will read (or listen to), understand, and act upon.

7. Conclusion

If robots themselves are friendly, cute, and playful, their privacy policies could and should be, too. This paper argues that privacy policies don't have to be just text written in legalese – they can be user-friendly and engaging. Even the GDPR allows for this – in fact, it requires that privacy information be provided in an easily accessible form, using clear and plain language. We have only scratched the surface of what plain language and design patterns can do. The early evidence is encouraging: information design tools offer a way to balance the needs and interests of robot users and builders – and privacy communication readers and providers – so that privacy information can truly serve both. In the future, as robots become more commonplace, more attention will need to be paid to privacy communication in human-robot interaction. As robots become mainstream, they will be accessible to everyone. This will require privacy information that everyone can understand, and information design will become increasingly important. New emerging technologies may soon provide new tools for users and builders, readers and writers, for the benefit of all.

8. Acknowledgement

The contribution of Yueh-Hsuan Weng to this work was mainly supported by JSPS KAKENHI Grant Number 22K18513.

9. References

Autto, Hilja/Haapio, Helena/Nuottila, Jouko, Contracts Rethought and Redesigned: A New Era with AI, Journal of Strategic Contracting and Negotiation (forthcoming).

Butt, Peter, Legalese versus Plain Language, Amicus Curiae, No. 35, June/July 2001, 28–32. https://doi.org/10.14296/ac.v2001i35.1332 (2001).

Čapek, Karel, R.U.R. – Rossum's Universal Robots (1920).

Cavazzuti, Chiara, DaPIS: The Data Potection Icon Set, May 2021, http://gdprbydesign.cirsfid.unibo.it/dapis-the-data-protection-icon-set/ last accessed 30.3.2023 (2021).

Corrales Compagnucci, Marcelo/Fenwick, Mark/Haapio, Helena, Digital Technology, Future Lawyers and the Computable Contract Designer of Tomorrow. In: Corrales Compagnucci, Marcelo/Haapio, Helena/Fenwick, Mark (Eds.), Research Handbook on Contract Design. Edward Elgar Publishing, Cheltenham, 421–444 (2022).

Haapio, Helena/Hagan, Margaret, Design Patterns for Contracts. In: Schweighofer, Erich/Kummer, Franz/Hötzendorfer, Walter/Borges, Georg (Eds.), Networks. Proceedings of the 19th International Legal Informatics Symposium IRIS 2016. Österreichische Computer Gesellschaft OCG /books@ocg.at, Wien, 381–388 (2016).

Haapio, Helena/Hagan, Margaret/Palmirani, Monica/Rossi, Arianna, Legal Design Patterns for Privacy. In: Schweighofer, Erich/Kummer, Franz/Saarenpää, Ahti/Schafer, Burkhard (Eds.), Data Protection/LegalTech. Proceedings of the 21th International Legal Informatics Symposium IRIS 2018. Editions Weblaw, Bern, 445–450 (2018).

Hayhoe, George F., Telling the Future of Information Design, Communication Design Quarterly Review, Vol. 1, No. 1, 23–26 (2012).

International Plain Language Federation, Plain Language Definitions, https://www.iplfederation.org/plain-language/ last accessed 31.10.2023.

Juro, The Juro Privacy Policy, https://juro.com/privacy, last accessed 30.3.2023.

Juro-Privacy (Juro), free-privacy-notice / Design-first privacy notice template. In: Github. https://github.com/juro-privacy/free-privacy-notice, last accessed 31.3.2023.

Kimble, Jonathan, Writing for Dollars, Writing to Please, Scribes Journal of Legal Writing, Vol. 6, 1–38 (1996–1997).

Mattila, Heikki E. S., Vertaileva oikeuslingvistiikka: juridinen kielenkäyttö, lakimieslatina, kansainväliset oikeuskielet, Alma Talent, Helsinki (2017).

McStay, Andrew, Emotional AI, Soft Biometrics and the Surveillance of Emotional Life: An Unusual Consensus on Privacy. In: Big Data & Society, Vol. 7, No. 1, https://doi.org/10.1177/2053951720904386 (2020).

Passera, Stefania, Beyond the Wall of Contract Text: Visualizing Contracts to Foster Understanding and Collaboration within and across Organizations. Doctoral Dissertation, Aalto University (2017).

Passera, Stefania, Juro Privacy Policy, https://stefaniapassera.com/portfolio/juro/, last accessed 30.3.2023.

Rossi, Arianna/Ducato, Rossana/Haapio, Helena/Passera, Stefania, When Design Met Law: Design Patterns for Information Transparency. In: Droit de la Consommation – Consumenterecht, No. 122–123, 79–121, Mars–Juin/Maart–Juni (2019).

Salo, Marika/Haapio, Helena, Robo-Advisors and Investors: Enhancing Human-Robot Interaction through Information Design. In: Schweighofer, Erich/Kummer, Franz/Hötzendorfer, Walter/Sorge, Christoph (Eds.), Trends and Communities of Legal Informatics. Proceedings of the 20th International Legal Informatics Symposium IRIS 2017, Österreichische Computer Gesellschaft / books@ocg.at, Wien, 441–448 (2017).

Stanford Legal Design Lab, Legal Design Pattern Libraries, https://legaltechdesign.com/communication-design/legal-design-pattern-libraries/, last accessed 31.10.2023.

Waller, Robert/Passera, Stefania/Haapio, Helena, Layered Contracts: Both Legally Functional and Human-Friendly. In: Jacob, Kai/Schindler, Dierk/Strathausen, Roger/Waltl, Bernhard (Eds.), Liquid Legal – Humanization and Law, Springer, Cham, 269–287 (2022).

Waller, Robert/Waller, Jenny/Haapio, Helena/Grag, Gary/Morrisseau, Sandi, Cooperation through Clarity: Designing Simplified Contracts, In: Journal of Strategic Contracting and Negotiation, Vol. 2, No. 1–2, , 48–68 (2016).

Weng, Yueh-Hsuan/Chen, Chien-Hsun/Sun, Chuen-Tsai, Toward the Human-Robot Co-Existence Society: On Safety Intelligence for Next Generation Robots, In: International Journal of Social Robotics, Vol. 1, No. 4, 267–282 (2009).

Weng, Yueh-Hsuan/Hirata, Yasuhisa, Design-Centered HRI Governance for Healthcare Robots. In: Journal of Healthcare Technology, 7 January 2022, Article ID 3935316. https://doi.org/10.1155/2022/3935316 (2022).

Weng, Yueh-Hsuan/Gulyaeva, Svetlana/Winter, Jana/Slavescu Andrei/Hirata, Yasuhisa, HRI for Legal Validation: On Embodiment and Data Protection. In: 2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN). IEEE, 1387–1394 (2020).

WorldCC Foundation/Passera, Stefania/Haapio, Helena, Pattern Families. In: Contract Design Pattern Library. https://contract-design.worldcc.com/families-overview, last accessed 29.3.2023.