Towards Assessing Features of Dark Patterns in Cookie Consent Processes

Emre Kocyigit(✉), Arianna Rossi, and Gabriele Lenzini

SnT, University of Luxembourg, 4364 Esch-sur-Alzette, Luxembourg
{emre.kocyigit,arianna.rossi,gabriele.lenzini}@uni.lu

Abstract. There has been a burst of discussions about how to characterize and recognize online dark patterns — i.e., web design strategies that aim to steer user choices towards what favours service providers or third parties like advertisers rather than what is in the best interest of users. Dark patterns are common in cookie banners where they are used to influence users to accept being tracked for more purposes than a data protection by default principle would dictate. Despite all the discussions, an objective, transparent, and verifiable assessment of dark patterns' qualities is still missing. We contribute to bridging this gap by studying several cookie processes, in particular their multi-layered information flow —that we represent as message sequence charts—, and by identifying a list of observable and measurable features that we believe can help describing the presence of dark patterns in digital consent flows. We propose thirty one of such properties that can be operationalised into metrics and therefore into objective procedures for the detection of dark patterns.

Keywords: Dark patterns · Deceptive designs · Cookie consent · Features · Dark pattern detection

1 Introduction

Dark patterns are manipulative design techniques that aim to favour certain purposes of digital services at the price of user's autonomy, such as collecting as many personal data as possible, and are very often illegal. Discussions about what exactly dark patterns are have arisen from different perspectives, for instance user experience, data protection law, interaction design. Each view has added useful insights on their determining characteristics that can help researchers, legislators, regulators, product developers and designers alike to recognize them and avoid their use.

Dark patterns in cookie banners have been particularly studied [6,15–17] due to their ubiquitous presence, which has increased dramatically since the GDPR strengthened the consent and transparency requirements for personal data processing, therefore new ways of sidestepping the rules and extorting consent have been experimented. Still, even within this very specific domain, there is a lack of criteria capable of capturing the essence of dark patterns and that can be operationalised into measurable variables. Now, even if not all the characteristics of

F. Bieker et al. (Eds.): Privacy and Identity 2022, IFIP AICT 671, pp. 165–183, 2023.
https://doi.org/10.1007/978-3-031-31971-6_13

dark patterns can be measured, like the intentionality of the deception, this approach would offer a framework that can be leveraged to reliably discuss, argue, verify claims on the supposed presence of dark patterns on a digital service and, eventually, detect them.

Moreover, dark patterns are not only related with Human-Computer Interaction (HCI) layers, like the graphical user interface, but also with the back-end of applications. Even though inspecting the user interface level can help to determine some dark pattern characteristics such as unbalanced weight of options or hard-to-notice buttons, there are certain manipulative strategies that are hidden in the back-end and are therefore invisible without inspecting the elements of a web page such as the cookie content, the cookie size, etc. For example, some websites do not respect the user's consent refusal decision and continue to collect their data nevertheless, like Matte et al. indicated in [13]. Therefore, we consider both HCI and Machine-to-Machine Interaction (MMI) layers, like the Web Server, in this study.

Our Contribution. For the purpose of providing objective descriptions on whether there are dark patterns in cookie consent flows, we study different websites' cookie consent processes and extract their activity diagrams through HCI and MMI layers, i.e. User Action, User Interface, Browser and Web Server. Consequently, we propose thirty-one features of cookie consent processes that could be useful to recognize dark patterns via the objective assessment of their characteristics. We define such features based on the information flow layers like User Interface, Web Server etc. and label them according to their values, which can be quantitative or binary. By leveraging these objective descriptions which can be operationalised into metrics, automated detection and analysing tools can be developed to help different stakeholders to detect, prevent and avoid dark patterns.

2 Related Work

The history of deceptive designs influencing people's decisions dates back to the days when digital services did not exist. For instance, fake advertisements of store closings existed even before the internet with the goal of attracting customers and boosting purchases [14]. Today's online deceptive designs can be more complicated and hard to respond because digital services affect thousands, even millions of people, unlike a traditional store. Moreover, these designs are omnipresent: for example, 97% of the most common websites and applications exploit at least one deceptive design in the European Union [10].

Especially after 2010, when the UX expert Harry Brignull coined the term "dark pattern" [3], the approval of regulations like the GDPR, an increased public scrutiny and the visible pervasiveness of these design strategies prompted researchers to analyze dark patterns. Mathur et al. [11] defined dark patterns' attributes in seven categories: *"Sneaking", "Urgency", "Misdirection", "Social Proof", "Scarcity", "Obstruction", "Forced Action"*, after scraping 11K

e-commerce websites and having examined their functionalities. In [5], M.C. Gray et al. categorized the dark patterns into five types as *"Nagging", "Obstruction", "Sneaking", "Interface Inference" and "Forced Action"* from an interface design and user experience perspective and covered all eleven types - e.g. *"Trick questions", "Sneak into basket", and "Roach motel"* - that were categorized in [1].

Even if many efforts have been made to create taxonomies where a strict categorization of different types of dark patterns is proposed, we believe it is more useful and objective to identify those attributes, i.e. functional characteristics, that can help detect a dark pattern, no matter the category to which it belongs. Mathur et al [11] defined dark pattern attributes as *"Asymmetric", "Restrictive", "Covert", "Deceptive", "Information Hiding", and "Disparate Treatment"*. Even though these definitions are helpful to describe some dark patterns, they are not yet measurable, actionable nor objectively portrayed. For example, the asymmetric attribute is defined as "unequal burden on options" provided to the user, but the burden is not explicitly measurable, also because it is not defined whether it concerns the amount of time, steps, cognitive effort, etc. that a user takes to perform a certain action.

Dark patterns can be faced in various parts of websites, applications, videogames, and other digital services. Researchers often focus on specific phases of the user experience, such as consent management flows, to analyze and understand their influence on users. For example, M. Nouwens et al. [15] found that dark patterns are ubiquitous in the designs of consent banners by scraping 10000 websites and their consent management platforms in the UK. In another work [13], Matte et al. detected legal violations related to dark patterns in cookie banners. For example, they proved that 141 websites behaved like they had user's positive consent though the user did not take any consent action. On the other hand, M.C. Gray et al. [6] highlighted that dark patterns in consent banners require an holistic approach due to their n-dimensional structure. C. Krisam et al. examined German websites and their cookies, and found that dark patterns that nudge users towards accepting cookies exist in most consent processes [8]. P. Hausner et al. also worked on the cookie banners of some German websites and inspected html and css elements of different examples as a first step of building automated detection tool for dark patterns [7].

In order to prevent and detect manipulative designs and legal violations in the cookie consent flow, developing automated detection tools is an emerging solution, like in [2]. However, there is still a long way to go in the automated detection studies. T.H. Soe et al. [17] used a data set [16] that contains 300 news websites, to train their machine learning model for automated detection of dark patterns in cookie banners such as "nagging, obstruction, sneaking" etc. They also pointed out the representation challenges and detection challenges of automated detection. For example, machine learning models can be quite successful for specific data types like image or text. Recent studies on image-based tasks or text-based tasks displayed promising results [4]. However, the image is usually not enough to understand if there is a dark pattern or not. Two banners that are visually similar but contain different textual contents can

be classified under the same label by the model, even if one of them contains manipulative text, i.e., a dark pattern. Another challenge in automated detection is that data set bias can affect the model's decision-making mechanism [2], therefore data collection process and sources should be designed regarding real-life data and possible bias.

As a result, the automated detection of dark patterns is still an unsolved issue, and an objective definition of features, at date missing, can support automated detection processes.

3 Background

3.1 Cookie Consent Flow

In order to obtain measurable criteria of the cookie consent process that provide an objective description of the dark patterns presence in this process, we firstly studied the cookie consent flow and defined its phases as shown in Fig. 1. While the flows have different interface design elements such as buttons, banner size, text etc. or provide various user actions such as "Decline all", "Refuse unnecessary cookies" etc., they show a regular/common structure, probably also as a result of regulations like the GDPR.

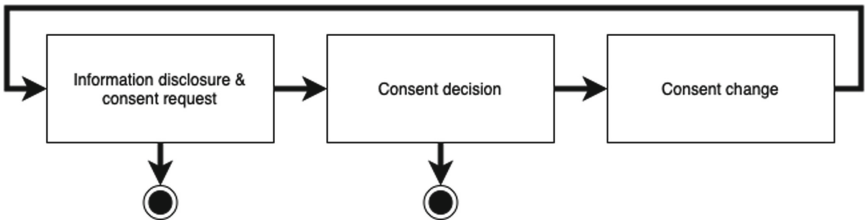

Fig. 1. Cookie consent flow phases

When users visit a website for the first time, they get informed about the presence of cookies and are asked to consent, following Article 7.3 of the GDPR and Article 6.3 and 6.4 of the ePrivacy Directive. This is why, we defined the first phase of the cookie consent flow as "Information disclosure and consent request". At this phase, cookie banners usually present links or buttons that redirect users towards a complete privacy policy, if they initiate that further action. Often there is also the possibility to customize consent options through user configuration.

After users are informed about the cookie policy, they usually have two options: leave the website without any further action or decide whether to grant or refuse consent. For this reason, the cookie consent process may end after the first phase or continue with the second phase, namely "Consent decision". At this stage, users can grant their consent for all data processing purposes, for some

(with customization of consent option), or refuse consent. After the consent decision, user can surf the website without any further consent action. According to the user's decision, cookie-related requests and responses run between browser and web server. Thus, cookies which will be loaded and saved to the user's device are known when user revisit the website after the consent decision.

Finally, users may change their previous consent decision (i.e., right to consent withdrawal, Article 7 GDPR). Hence, we defined the third phase as "Consent decision change". At this stage, users may withdraw their consent, provide consent for previously refused processing purposes, or give consent for all or some of the previously refused purposes. This change reshapes cookie content which will be handled between the browser and web server for the user. Also, the third phase triggers the first phase as shown in Fig. 1.

We created a matrix to extract all possible cookie consent scenarios as shown in Appendix A. Because there is a loop from the third phase to the first phase, infinite scenarios are possible, but we ignored the last phase because it only determines if the loop is over or not. Eventually, we listed 14 possible scenarios in the consent management flow in terms of user actions and extracted the features by considering these scenarios.

3.2 Human-Computer and Machine-to-Machine Interaction Layers of the Cookie Consent Process

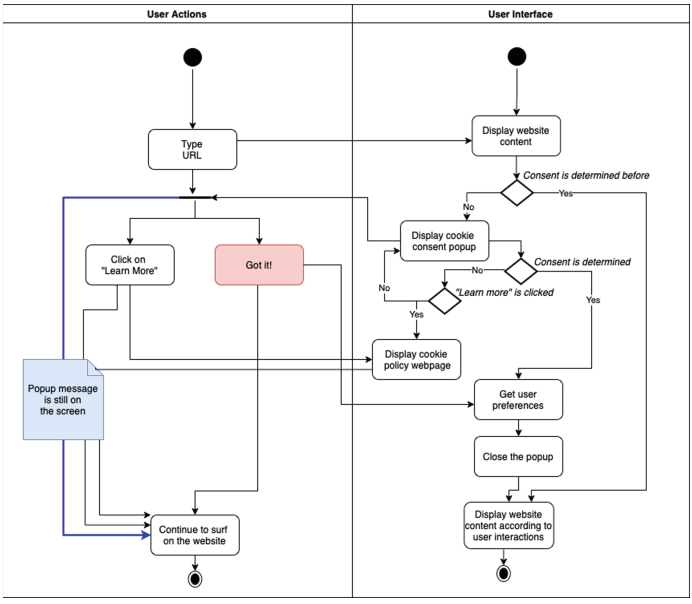

Fig. 2. User Action and User Interface layers of Cookie Consent Flow of Website 2

A web flow is performed via a harmony of different components by synchronous and asynchronous requests and responses online. We drew a Unified Modeling Language (UML) activity diagram which has different layers such as "User Action", "Browser - User Interface", "Browser - Engine" and "Server" as shown in Appendix B to inspect human-computer interaction and machine-to-machine interaction of consent management flows. While vertical flows display alternative cases level by level for each layer, horizontal flows show interactions between different layers. For example, request and response between web engine and web server can be seen in the horizontal flows.

Various consent management flows are available on websites, thus we selected 10 websites (see Appendix C) that provide different consent management flows to examine and drew a diagram for each to make a comparative study. This comparison contributed to define measurable features in the cookie consent flow which may be descriptive for Dark Patterns. For example, while the "User Action" layer of a website's cookie consent flow has 4 actions as shown in Fig. 2, another website's cookie consent flow is more complex because it has more than 20 actions as shown in Fig. 3. We focused on this type of differences to obtain the features. The feature extraction steps and the methodology are detailed in the next section.

4 Research Questions and Methodology

4.1 Research Questions

In our research, we aim to define the functional characteristics of dark patterns. As a first step towards this goal, we choose the cookie consent process as a use case to elaborate the objective assessment of dark patterns' characteristics. Thereby, our first research question is: *"What are the features in cookie consent processes to objectively describe the presence of dark patterns?"*. The answer to this question will be a starting point to formulate the metrics that can be helpful to define the functional characteristics of dark patterns. Secondly, most studies have only looked at the most visible side of the online world, like the user interface as we mentioned in Sect. 1, thus we have an imperfect understanding of dark patterns. Another motivation of this paper is to find out if there are features that are visible/measurable only by examining the HCI and MMI layers of the cookie consent flow. A holistic approach is essential to define a comprehensive feature list. Accordingly, another question is: *"Are there any dark pattern-related features of cookie consent processes that can be identified only by inspecting multiple layers of HCI and MMI flows?"*. We believe that answers to these questions will aid to formulate the characteristics of dark patterns and to extract suitable features for the automated detection models in future work.

4.2 Methodology

As a first step, we selected cookie consent flows of 10 different websites which are listed in Table 2. Some of them contain dark patterns like "Forced Action"

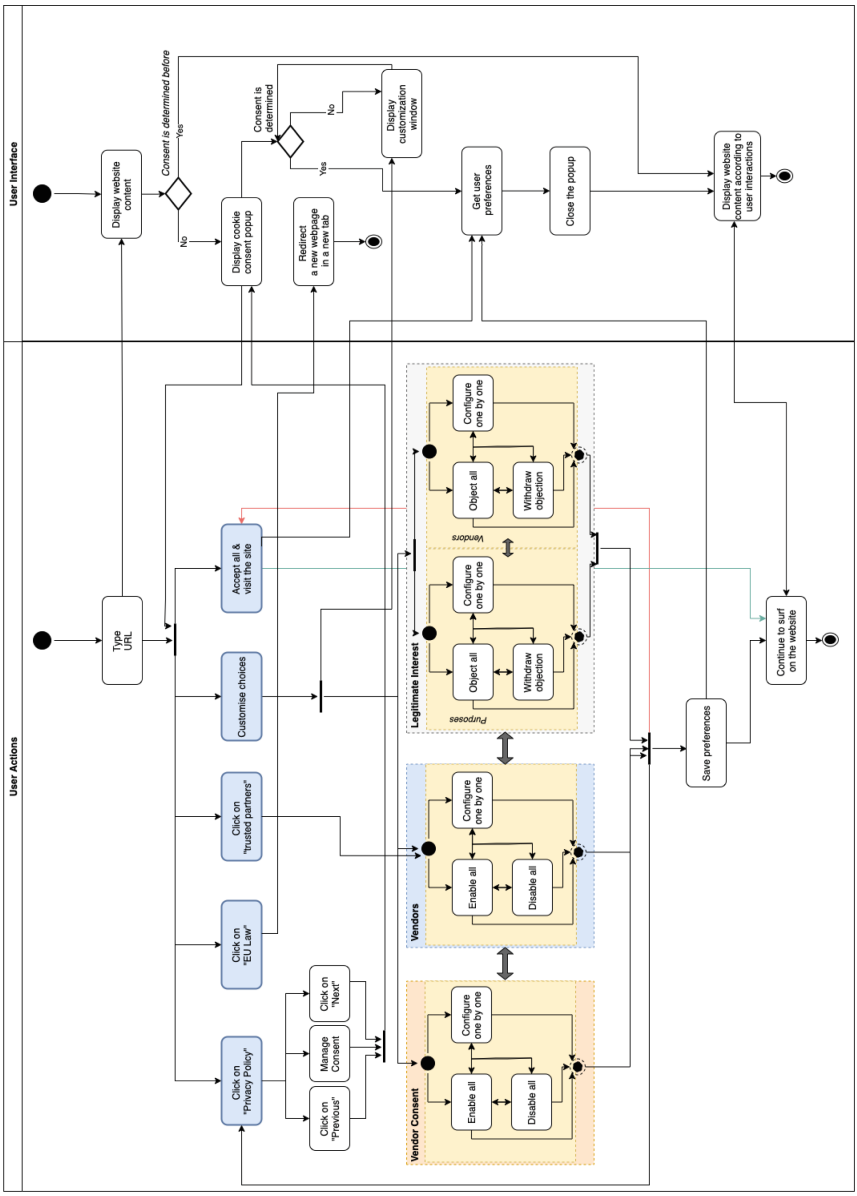

Fig. 3. User Action and User Interface layers of Cookie Consent Flow of Website 3

[5]. The second step is creating UML activity diagrams for each cookie consent flow to represent information flows in multiple layers as: User Action - User Interface - Browser - Server. By these diagrams, we inspected horizontal and vertical flows between and within layers to understand which patterns can be associated with dark patterns. We organized the analysis of information flow into two subsections. First, we inspected the number of elements, the possible routes from the starting user action to the final one, the patterns possibly related to dark pattern types, and cookie related information such as cookie size and cookie number of each cookie consent flow separately. Then, we compared the same layers of different cookie consent flows with each other. For example, we extracted routes for same action from different cookie consent processes and compared them to find their differences, such as number of actions required to the user. After the analysis steps, we defined features according to the HCI and the MMO layers by using the analysis outputs. Finally, we discussed the relationship between existing dark pattern attributes and our proposed features to answer the research questions presented in Sect. 4.

As a first step, we focused on "User Action" and "User Interface" layers to extract measurable features at the moment of information disclosure and consent request. For example, if we look at Fig. 2, there is a bold, blue line that highlights the path from the "Type URL" action to the "Continue to surf on the website" action. However, in Fig. 3, a direct line is not available between these two actions. This difference indicates that the user does not need to make a consent decision on the first website, while the second one forces the user to decide. This difference can be an indicator of forced action [5], which is a restrictive property [12] that is common in dark pattern examples. Additionally, while there are two user options available in Fig. 2 after the "Type URL" action, five options are present in Fig. 3. This also affects the difference in complexity between the two examples, which may be an important metric to define the asymmetric property, i.e., when the complexity of two flows does not have equal load [12]. Consequently, we defined the features "Forced decision" and "User options at first connection", respectively, in Sect. 5.

5 Features in Cookies Consent Processes Relevant for Dark Patterns Detection

After extracting the multiple layers of the cookie consent flows of ten different websites that are listed in B and carrying out the extraction process that was detailed in Sect. 3, we listed thirty one features according to different layers of cookie consent information flow – i.e., User Action, User Interface, Browser Engine, Web Server – and types – i.e., Binary and Quantitative – as given in the Appendix B. As a result, we categorized the features into: Human-Computer(User Action and User Interface) Interaction-based in Sect. 5.1 and Machine-to-Machine(Browser Engine and Web Server) Interaction-based in Sect. 5.2.

5.1 Human-Computer Interaction-Based Features

1. **Forced decision**: This feature describes if the cookie consent banner prevents website usage without a consent decision. In the UML activity diagram, if there is no direct link from the "Type URL" user action to the "Continue to surf on the website" user action, it means that the website forces the user to make a consent decision. The value of the feature would be "1" then. Some websites allow to navigate the website while keeping the consent banner on the screen without requiring a user action. In this case, the value would be "0".
2. **User options at first visit**: Consent banners vary in terms of the first possible actions available. While some of them provide three options such as "Agree", "Disagree" and "Customise", others can provide only one option like "Got it", or two options "Accept all" and "Learn more". This feature describes the number of user actions at the first connection.

Full consent

3. **Availability**: This feature describes if there is a full consent option that allows the user to accept all processing purposes at once, like an "Accept all" button. The value 1 is assigned when this option is available, while 0 is assigned when it is not. Generally, consent banners provide this option.
4. **Total number of routes**: This feature gives the total number of possible routes from the "Type URL" user action to the "Accept all/Agree all etc." user action. Some websites provide more than one way to consent to all cookies, while providing one or no way to reach the option of refusal.
5. **Minimum length of routes**: This feature describes the shortest route from the "Type URL" user action to the "Accept all" action, and can be helpful to define if a flow provides asymmetric options.
6. **Maximum length of routes**: This feature describes the longest route from the "Type URL" user action to the "Accept all" action and can be helpful to define if a flow provides asymmetric options. If loops exist, then only the first loop will be considered for this feature.

Full consent refusal

7. **Availability**: This feature describes if there is a full consent refusal option like the "Decline all" button, where 1 means that that option is available. While websites generally have a full cookie acceptance option, it is not always the case that they have a full consent refusal option on the first layer. Since websites may use various phrases to express consent refusal (for example, while one shows "Refuse unnecessary cookies", another one can show "Refuse all"), we accept all of these as full consent refusal indications.
8. **Total number of routes**: This feature gives the total number of all possible routes from the "Type URL" user action to "Decline all/Disagree all" user action. Therefore, this feature can be helpful to understand if both options are provided to user in a symmetric way.
9. **Minimum length of routes**: Users perform actions to complete a consent decision flow and this process contains a different number of steps. While

some websites provide shorter or easier ways to get to consenting, they may make the consent refusal path harder. This feature describes the shortest route from the "Type URL" user action to the "Decline all" action and can be helpful to define asymmetry within the website or comparatively with other websites.

10. **Maximum length of routes**: This feature describes the longest route from "Type URL" user action to "Decline all" action, and can be helpful to define if a flow provides asymmetric options. If loops exist, then only the first loop will be considered for this feature.

Customised consent Some websites provide the customization of consent preferences to users according to the data processing purposes. This process usually takes more time and requires more user actions than granting full consent.

11. **Availability**: This feature describes if the cookie consent flow contains a configuration or customization option concerning processing purposes and/or vendors, where 1 means that the customization is available on the cookie banner, while 0 signifies that it is not available.
12. **Total number of routes**: This feature gives the total number of all possible routes from "Type URL" user action to "Customize/preferences/options etc." user action. Therefore, this feature can be helpful to understand the complexity of the flow.
13. **Minimum length of routes**: This feature describes the shortest route from the "Type URL" user action to the "Partly accept/save and exit" action after configuring the consent settings. This feature can also be helpful to define if a flow provides asymmetric options.
14. **Maximum length of routes**: This feature describes the longest route from the "Type URL" user action to "Partly accept/save and exit" action. If loops exist, then only the first loop will be considered for this feature.

15. **Total user action options**: This feature describes the total number of actions on the User Actions layer of cookie consent flow. It can be helpful to express the complexity of consent management flow.
16. **Total consent flow routes**: This feature describes all possible routes from the first user action to the end on the User Actions layer. Website designs can provide different consent scenarios, where some of them offer "Accept all" scenarios more often than "Decline all" scenarios. This feature can be helpful to understand asymmetry through all possible scenarios.
17. **Total hyperlinks on user interface**: This feature concerns the Browser - User Interface layer, and describes the total number of hyperlinks (e.g. buttons) on it. It can be helpful to understand the complexity of interactions on the user interface.
18. **Consent decisions management availability**: This feature describes if there is a consent change button, link, dashboard etc. on the website to edit a previously given consent decision, e.g. withdrawing consent. If the value

of this feature is 1, that means that the consent decision can be updated by the user.

19. **Dead end**: This feature describes if there is a link or button which redirects user from the cookie consent flow to a new web page outside of the flow without providing the possibility to come back into the cookie consent flow. For example, a cookie banner can open a new window when a user clicks on the "privacy policy" link and the new page doesn't provide any "back" button.

5.2 Machine-to-Machine Interaction-Based Quality Features

In the network section of browser inspection, requests between browser and web server can be observed. The number of requests, the size and number of cookies, and the size of transferred files vary according to whether the visit is the first one or not and whether consent decisions are taken or not.

20. **Number of cookies at first visit**: The cookies that are installed on a device and their details can be unveiled when the web browsers are inspected. This feature describes the total number of cookies installed when a user visits a website for the first time.
21. **Number of cookies in full consent**: This feature describes the total number of cookies installed when a user visits the website after full consent grant.
22. **Number of cookies in full consent refusal**: This feature describes the total number of cookies when a user visits the website after full consent refusal.
23. **Cookie size at first visit**: This feature describes the total cookie size when a user visits the website for the first time. Each cookie has its own size value, but we obtain the total size to simplify the feature and generate a quantitative one.
24. **Cookie size in full consent**: The cookie size value should be associated with consent grant decision. This feature describes the total cookie size when consent is fully granted.
25. **Cookie size in full consent refusal**: This feature is the opposite of cookie size in consent grant feature, and describes the total cookie size when consent is refused. When user does not allow all cookies, this feature should be lower from the above one in most cases.
26. **Total requests at first visit**: This feature describes the total number of requests between browser and web server at the first stage, before the consent decision. Some of the requests are not related with consent but with functionality of the website, e.g., source files of the web page, while others are directly related with the consent management process or the consent decision. Therefore, this feature can be helpful when it is used with the below two to check if the consent decision is correctly implemented.
27. **Total requests in full consent**: This feature describes the total number of requests between browser and web server when consent is granted.

28. **Total requests in full consent refusal**: This feature describes the total number of requests between browser and web server when consent is refused.
29. **Transferred file size at first visit**: Transferred files between browser and web server can vary based on consent decision. This feature describes the total transferred file size in the first interaction between the user and the website.
30. **Transferred file size in full consent**: This feature describes the total transferred file size when the user types a URL after consent is granted.
31. **Transferred file size in full consent refusal**: This feature describes the total transferred file size when the user types a URL after consent is refused.

6 Discussion

We have listed 31 features which can be operationalizable into metrics to assess dark patterns after analysing the cookie consent processes in the previous section. One of the purposes of this work is to promote objective grounds to discuss qualities about online patterns and eventually provide evidence to discuss, confirm or refute claims on whether a particular design pattern is "Dark". To the best of our intention, the features that we have identified are all operationalizable. It is possible, in other words, to define measures for the features that can be evaluated on the layered message sequence charts that we have provided. Some of them can be already used to objectively capture certain qualities and characteristics that have been only informally described in the literature. For instance, in [12], Mathur et al. defined six attributes and two of them, i.e. asymmetric (unequal weights on the provided options to the user) and restrictive (discarding specific options that should be available to the user) can be described with the features we proposed. For example, asymmetric can be evaluated by comparing these features: "Total number of routes in full consent grant (#4)", "Total number of routes in full consent refusal (#8)", "Total number of routes in customized consent grant (#12)". If their values are not equal, we can conclude that the design is asymmetric in an objective way. The restrictive attribute can be described by the following ones: "Forced decision (#1)", "Full consent availability (#3)", "Full consent refusal availability (#7)", and "Customized consent availability (#11)". For instance, if the value of "Full consent refusal availability (#7)" is zero, then we can argue that users do not have the option to reject cookies in one action, i.e., the cookie consent flow contains a restrictive pattern.

The features can provide descriptions of dark patterns when they are used in different cookie consent flows, not only within one cookie consent flow. We believe that our proposals are not only providing direct description of dark patterns' attributes such as "asymmetric" or "restrictive", but also indirect metrics like "complexity". For example, the complexity of a large amount of different flows can be measured with the help of "Total consent flow routes (#16)", "Total user action options (#15)", and "Total buttons on User Interface (#17)". In this case, reasonable inferences can be made for individual samples after calculating the mean values. For example, "choice complexity", which is defined as required effort for an action or selection in [10], can be measured with these features.

Although dark patterns can be hidden in the back-end of the cookie consent flows, the MMI-based features can evaluate their presence since the features check the interaction within/between Browser and Web Server layers of the cookie consent flow. For example, when "Cookie size in full consent (#24)" and "Cookie size in full consent refusal (#25)" are assessed together, if they are equal we can deduce that the consent refusal option is not working or that the cookie consent flow is deceptive, as it shows to users the option of refusal but user's choice is not correctly registered. We measured these features of different websites when users firstly arrived to the website before consent decision and after full consent grant or full consent refusal as shown in the Fig. 4. As it can be observed, in most instances the values of the features are unsurprisingly lower, rather than equal, when a user refuses to consent comparing to the full consent grant.

	Cookie number			Cookie size			Request number			Transferred file size		
Website	in first visit	full consent	consent refusal	in first visit	full consent	consent refusal	in first visit	full consent	consent refusal	in first visit	full consent	consent refusal
w3schools.com	3	5	5	66	519	401	46	39	37	1.32 MB	144.01 KB	47.13 KB
researchgate.net	4	10	6	387	946	663	41	35	31	1.18 MB	15.31 KB	14.40 KB
stack overflow.com	5	6	6	360	379	403	44	36	33	609.51 KB	126.69 KB	104.17 KB
dropbox.com	10	65	46	677	2796	1757	108	405	90	1.80 MB	1.59 MB	375.26 KB

Fig. 4. Extracted values of machine-to-machine interaction-based cookie consent features

As we mentioned in Sect. 2, previous categorization and definition studies offer a variety of perspectives to examine dark patterns, but the literature on the topic still lacks objective and measurable criteria to assess the presence of dark patterns on online services. We believe that measurable metrics are not only needed for an objective evaluation, but also to develop automated detection tools, which require a concrete and measurable feature set and can be helpful for developers and designers during testing processes, and to regulators and researchers during evaluation and inspection processes. Automated detection tools like machine learning-based systems require structured data in the training process. These metrics can be helpful to constitute feature sets and data sets thanks to their measurable nature.

7 Limitations and Future Work

We examined the cookie consent flow as a use case of dark patterns and proposed thirty one features that can be operationalised into metrics based on our multi-layers interaction flow analyses. We provided clear descriptions for them, but we have not yet defined procedures, quality measures, nor guidelines to interpret

them for any of the features that we have described. We have informally given a few exemplifying arguments to show that this exercise is possible, but its full development is left for future work.

We selected ten different websites and extracted the interaction flows of their cookie consent processes. We believe this sample is sufficiently representative for this initial analysis, but of course additional different examples may help to explore other features. We excluded the "legitimate interest" paths in the cookie consent flow and the flows that repeat over time, e.g., nagging users to consent in their later visits. Moreover, the choices we made in selecting certain elements of the process in the message sequence chart affect what features we can define. For instance the edges in our charts are unweighted. One could, for instance, consider the cost of a certain choice which can provide new features.

Furthermore, while studying multiple layers of interaction flows can contribute to obtain metrics describing dark pattern characteristics, we excluded the language. Therefore these features are not capable of describing examples like "confirmshaming" [9], which exploits human emotions like guilt to manipulate users. We also did not include visual elements such as button size, colour, contrast ratio etc. into this study but we intend to consider them in our future work.

The back-end of websites would need further investigations to cover other aspects of Dark Patterns and extend their features. For instance, we focused on the total cookie size in different consent conditions, but we did not extract the content of the cookies, and we observed the total number of requests between browser and web server, but we did not detect the ones that are directly related with cookies. Hence, the features can be elaborated into their breakdowns such as "Total cookie-based requests after full consent" or "Advertisement-based requests after full consent refusal". Since all cookies do not represent same mission, e.g., some can be designed the functionality of cookie consent mechanism.

8 Conclusion

We extracted the information flows of ten cookie consent processes and studied the multiple layers as "User Action", "Browser-User Interface", "Browser-Engine" and "Server" to obtain objective and measurable criteria to assess the presence of dark patterns. In this direction, we proposed thirty one features of dark patterns in the cookie consent flow. They can be operationalised into metrics not only to evaluate the presence of dark patterns on online services, but also to be build automated detection models.

Acknowledgement. This paper is published as a part of the project Decepticon (grant no. IS/14717072) supported by the Luxembourg National Research Fund (FNR). We would like to thank Nataliia Bielova for her helpful comments.

A Appendix A

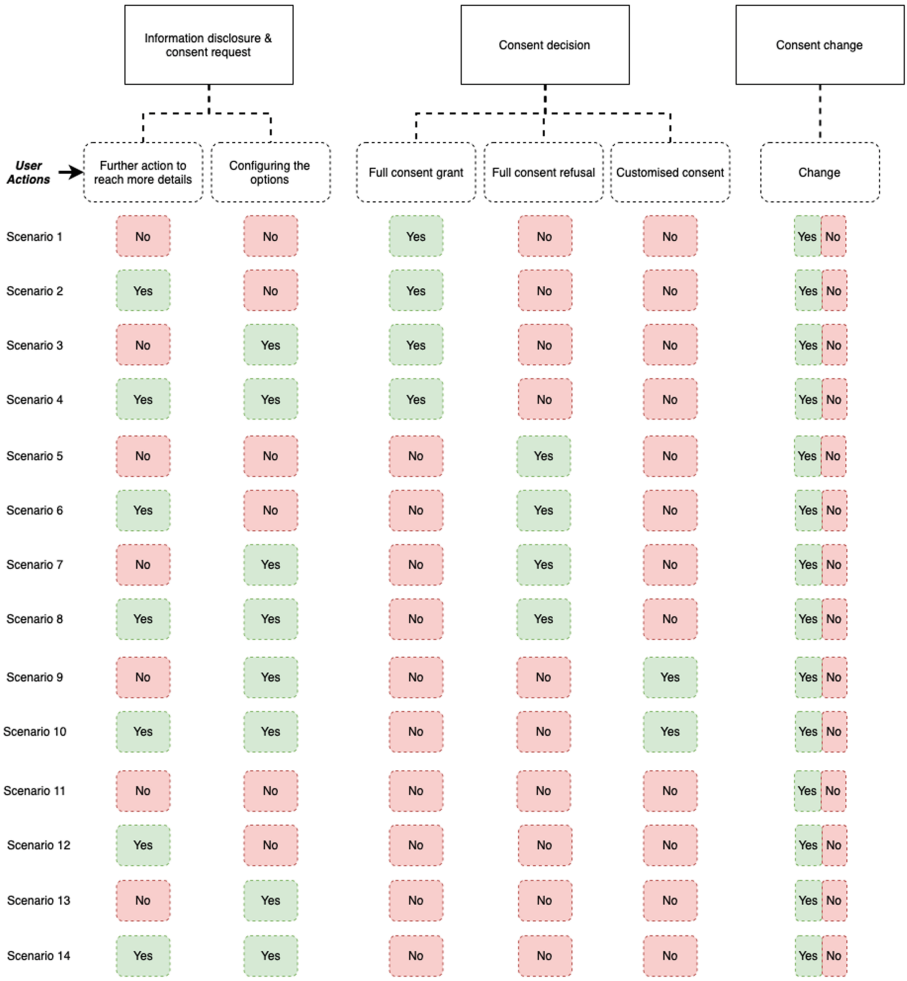

Fig. 5. Cookie consent scenarios based on user actions in the first two phases

B Appendix B

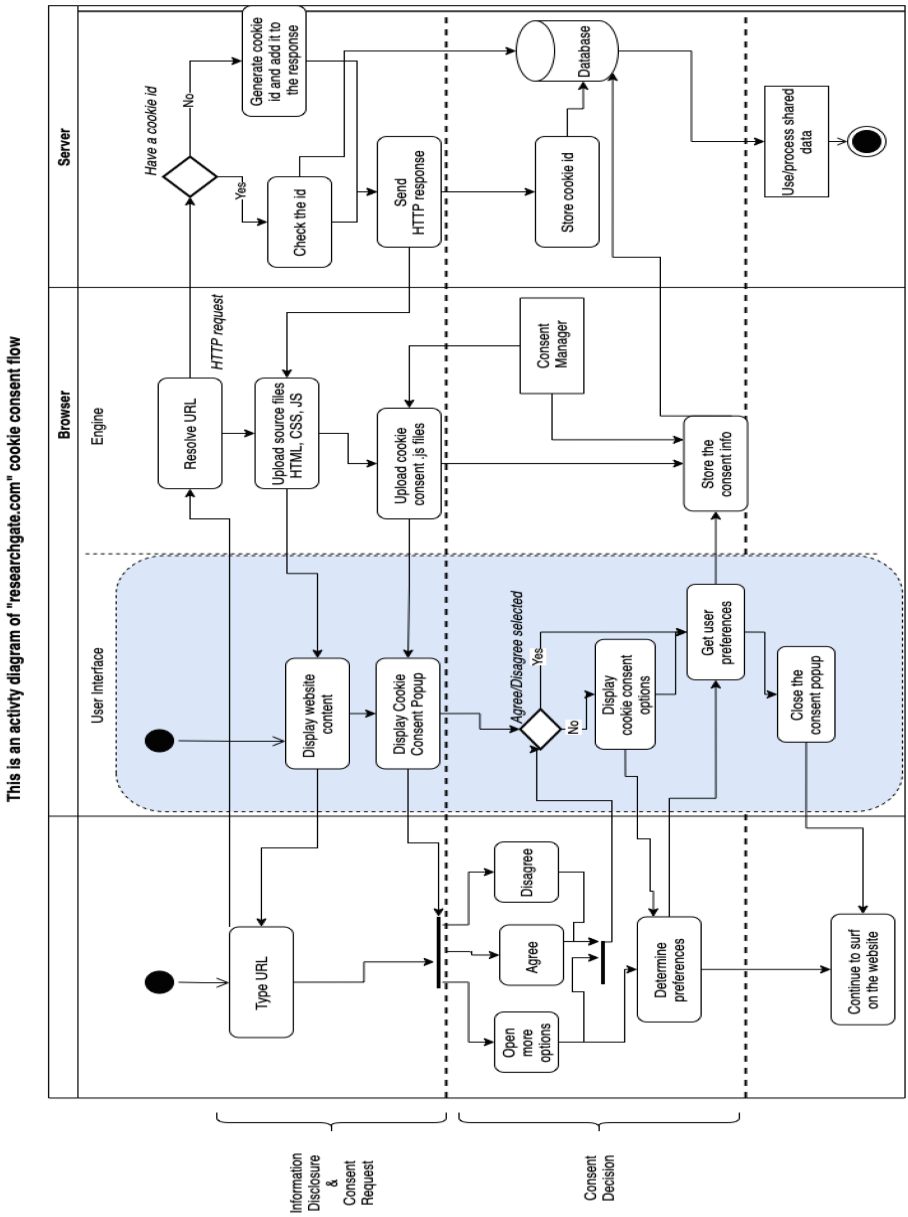

Fig. 6. Multi-layer presentation of human-computer and machine-to-machine interaction flow of cookie consent flow.

C Appendix C

Table 1. Features of Cookie Consent Process Relevant for Dark Patterns Detection

No	Feature	Layer	Value
1	Forced decision	User Action	Binary
2	User options at first visit	User Action	Quantitative
3	Full consent availability	User Action	Binary
4	Total number of routes in full consent grant	User Action	Quantitative
5	Minimum length of routes in full consent grant	User Action	Quantitative
6	Maximum length of routes in full consent grant	User Action	Quantitative
7	Full consent refusal availability	User Action	Binary
8	Total number of routes in full consent refusal	User Action	Quantitative
9	Minimum length of routes in full consent refusal	User Action	Quantitative
10	Maximum length of routes in full consent refusal	User Action	Quantitative
11	Customised consent availability	User Action	Quantitative
12	Total number of routes in customised consent	User Action	Quantitative
13	Minimum length of routes in customised consent	User Action	Quantitative
14	Maximum length of routes in customised consent	User Action	Quantitative
15	Total user action options	User Action	Quantitative
16	Total consent flow routes	User Action	Quantitative
17	Total hyperlinks on user interface	User Interface	Quantitative
18	Consent decisions management availability	User Action	Binary
19	Dead end	User Action	Binary
20	Number of cookies at first visit	Web Engine	Quantitative
21	Number of cookies after full consent	Web Engine	Quantitative
22	Number of cookies after full consent refusal	Web Engine	Quantitative
23	Cookie size at first visit	Web Engine	Quantitative
24	Cookie size after full consent	Web Engine	Quantitative
25	Cookie size after full consent refusal	Web Engine	Quantitative
26	Total requests at first visit	Web Engine-Web Server	Quantitative
27	Total requests after full consent	Web Engine-Web Server	Quantitative
28	Total requests after full consent refusal	Web Engine-Web Server	Quantitative
29	Transferred file size at first visit	Web Engine-Web Server	Quantitative
30	Transferred file size after full consent	Web Engine-Web Server	Quantitative
31	Transferred file size after full consent refusal	Web Engine-Web Server	Quantitative

Table 2. Websites That Are Selected To Inspect Cookies Consent Processes

No	URL	No	URL
1	https://www.researchgate.net	6	https://twitter.com
2	https://www.kaggle.com	7	https://www.dropbox.com
3	https://www.w3schools.com	8	https://wwwfr.uni.lu
4	https://stackoverflow.com	9	https://github.com
5	https://www.linkedin.com	10	https://medium.com

References

1. Deceptive design - user interfaces crafted to trick you. https://www.deceptive.design/
2. Bollinger, D., Kubicek, K., Cotrini, C., Basin, D.: Automating cookie consent and gdpr violation detection. In: 31st USENIX Security Symposium (USENIX Security 22). USENIX Association (2022)
3. Brignull, H.: Dark patterns: dirty tricks designers use to make people do stuff. Retrieved Sept 29 2019 (2010)
4. Géron, A.: Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow. "O'Reilly Media, Inc." (2022)
5. Gray, C.M., Kou, Y., Battles, B., Hoggatt, J., Toombs, A.L.: The dark (patterns) side of ux design. In: Proceedings of the 2018 CHI conference on human factors in computing systems. pp. 1–14 (2018)
6. Gray, C.M., Santos, C., Bielova, N., Toth, M., Clifford, D.: 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. pp. 1–18 (2021)
7. Hausner, P., Gertz, M.: Dark patterns in the interaction with cookie banners. arXiv preprint arXiv:2103.14956 (2021)
8. Krisam, C., Dietmann, H., Volkamer, M., Kulyk, O.: Dark patterns in the wild: Review of cookie disclaimer designs on top 500 german websites. In: European Symposium on Usable Security 2021, pp. 1–8 (2021)
9. Luguri, J., Strahilevitz, L.J.: Shining a light on dark patterns. Journal of Legal Analysis **13**(1), 43–109 (2021)
10. Lupiáñez-Villanueva, F., Boluda, A., Bogliacino, F., Liva, G., Lechardoy, L., de las Heras Ballell, T.R.: Behavioural study on unfair commercial practices in the digital environment: dark patterns and manipulative personalisation. European Commission, Directorate-General for Justice and Consumers, final report. May (2022)
11. Mathur, A., et al.: Dark patterns at scale: Findings from a crawl of 11k shopping websites. In: Proceedings of the ACM on Human-Computer Interaction 3(CSCW), pp. 1–32 (2019)
12. Mathur, A., Kshirsagar, M., Mayer, J.: What makes a dark pattern... dark? design attributes, normative considerations, and measurement methods. In: Proceedings of the 2021 CHI Conference on Human Factors in Systems, pp. 1–18 (2021)

13. Matte, C., Bielova, N., Santos, C.: 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), pp. 791–809. IEEE (2020)
14. Narayanan, A., Mathur, A., Chetty, M., Kshirsagar, M.: Dark patterns: past, present, and future: the evolution of tricky user interfaces. Queue **18**(2), 67–92 (2020)
15. Nouwens, M., Liccardi, I., Veale, M., Karger, D., Kagal, L.: 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, pp. 1–13 (2020)
16. Soe, T.H., Nordberg, O.E., Guribye, F., Slavkovik, M.: Circumvention by design-dark patterns in cookie consent for online news outlets. In: Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society, pp. 1–12 (2020)
17. Soe, T.H., Santos, C.T., Slavkovik, M.: Automated detection of dark patterns in cookie banners: how to do it poorly and why it is hard to do it any other way. arXiv preprint arXiv:2204.11836 (2022)