Goldsmiths Research Online

*Goldsmiths Research Online (GRO)
is the institutional research repository for
Goldsmiths, University of London*

Citation

Jamalifard, Marzieh; Russell-Rose, Tony and Brignull, Harry. 2024. 'Exploring Deceptive Patterns: Insights from Eye Tracking, EMG and Sentiment Analysis'. In: 37th International BCS Human-Computer Interaction Conference. University of Central Lancashire, Preston, United Kingdom 15 - 17 July 2024. [Conference or Workshop Item]

Persistent URL

https://research.gold.ac.uk/id/eprint/37707/

Versions

The version presented here may differ from the published, performed or presented work. Please go to the persistent GRO record above for more information.

If you believe that any material held in the repository infringes copyright law, please contact the Repository Team at Goldsmiths, University of London via the following email address: gro@gold.ac.uk.

The item will be removed from the repository while any claim is being investigated. For more information, please contact the GRO team: gro@gold.ac.uk

Goldsmiths
UNIVERSITY OF LONDON

Exploring Deceptive Patterns: Insights from Eye Tracking, EMG and Sentiment Analysis

Marzieh Jamalifard
Goldsmiths, University of London
Computing Department
London SE14 6NW UK
mjama001@gold.ac.uk

Dr Tony Russell-Rose
Goldsmiths, University of London
Computing Department
London SE14 6NW UK
t.russell-rose@gold.ac.uk

Dr Harry Brignull
Testimonium Ltd
Unit 219, Foundry 78 The Beacon
Eastbourne, BN21 3NW UK
harry@testimonium.co

A Deceptive Pattern (DP) is a user interface design that intentionally misleads users by manipulating their decision-making process. In this study, we investigate the effects of two DPs, namely 'Hard to cancel' and 'Hidden subscription'. We adopt a mixed-method approach that combined eye tracking, electromyography (EMG), and sentiment analysis to explore user behaviour, cognitive processing, and emotional responses. Our findings show that explicit consent mechanisms result in higher user engagement and quicker noticeability of terms and conditions compared to implicit consent, and obfuscated cancellation processes elicit stronger negative emotional responses. Sentiment analysis further supports these findings, showing a higher proportion of negative sentiment towards deceptive practices. These results underscore the value of transparent and ethical design practices and the importance of regulatory frameworks to mitigate the adverse effects of deceptive UX patterns.

Deceptive Pattern, User experience, Eye tracking, EMG analysis.

1. INTRODUCTION

In recent years, the field of user experience (UX) design has seen remarkable growth, with designers striving to create interfaces that prioritize clarity, efficiency, and user trust. However, alongside this progress, there has been a concerning rise in the use of DPs aimed at manipulating user behaviour and perceptions (Brignull 2023). These deceptive practices, ranging from subtle nudges to outright misinformation, pose significant ethical challenges for designers and raise questions about the integrity and transparency of digital experiences. This paper investigates the effects of specific DPs in UX design. We adopt a mixed-method approach, combining eye tracking, electromyography (EMG), and sentiment analysis to study the effects of DPs on user behaviour, cognitive processing, and emotional responses. EMG facilitates the quantitative measurement of emotional responses, while eye tracking provides insights into attention patterns and interaction behaviours. These quantitative analyses are augmented by qualitative methods in the form of retrospective think aloud sessions to provide insight into the subjective experience. The transcripts from these sessions are further analysed for sentiment to provide additional quantitative insight into effective states. We focus specifically on DPs known as 'Hard to cancel' and 'Hidden subscription'. We investigate these patterns in a lab-based study, using live instances of contemporary online products and services.

2. BACKGROUND

The term 'DPs' is used to refer to the practice of using digital interfaces to 'steer, deceive, coerce, or manipulate consumers into making choices that often are not in their best interests' (Lupiáñez-Villanueva et al. 2022). These practices have been increasingly documented in recent years. (Di Geronimo et al. 2020) analyse DPs in 240 popular mobile apps and conduct an online experiment with 589 users to determine how they perceive them. They found that 95% of apps have DPs and that popular apps have at least seven deceiving interfaces. The online experiment shows that most users do not recognise DPs, but informed users can spot malicious designs. (Di Geronimo et al. 2020) continue the academic discourse on DPs by examining two new dimensions: (1) the prevalence of DPs in popular mobile apps and (2) users' awareness and ability to recognise these patterns. Their focus on mobile apps is particularly important due to the widespread use and significant impact of these apps on social life.

© The Authors. Published by BCS Learning and Development Ltd. Proceedings of . . .

(Gray et al. 2018) delved into dark patterns through the lens of UX practitioners' experiences and discussions. They compiled a collection of examples that practitioners identified as representative of DPs. (Luguri and Strahilevitz 2021) discuss user interface DPs that manipulate or confuse users. Experiments show that these patterns boost service subscriptions, especially among less educated users.

A study by (M. Bhoot et al. 2020) identified five critical elements that significantly aid users in recognising DPs, even when they are not fully aware of the unethical intentions behind these designs. The elements that aided in the identification were the frequency of occurrence, trustworthiness, level of frustration, misleading behavior, and physical appearance. Furthermore, the study developed a taxonomy of factors that influence users' susceptibility to DPs. The research revealed strong correlations between these five elements and users' abilities to identify DPs. Additionally, it was noted that the nature of these correlations varies depending on the specific type of DPs under consideration.

A comparison between a website with dark patterns (DPs) and one without them found no significant difference in users' willingness to revisit the site. Most participants were moderately aware of the concept of DPs, despite not being familiar with the specific term (Zac et al. 2023) examined the impact of DPs on various consumer groups and concluded that all individuals are susceptible to manipulative digital design techniques.

3. METHOD

This study investigates the effects of two specific DPs: (a) 'Hidden subscription', in which the user is unknowingly enrolled in a recurring subscription or payment plan without clear disclosure or their explicit consent, and (b) 'Hard to cancel', in which the user finds it easy to sign up or subscribe but difficult to cancel. A study by (Luguri and Strahilevitz 2021) identified hidden subscription as one of the most effective dark pattern strategies. Likewise, studies by (Ibarra 2017) have shown that while platforms need to prevent accidental account deletions, it's unclear if some of this friction is intentionally manipulative to retain users and continue monetising their information. Therefore 'Hidden subscription' and 'Hard to Cancel' DPs were selected as the focus of this study.

We adopt a mixed-method approach, combining eye tracking, electromyography (EMG), and sentiment analysis to study the effects of DPs on user behaviour, cognitive processing, and emotional responses. EMG facilitates the quantitative measurement of emotional responses, while eye tracking provides insights into attention patterns and interaction behaviours. These quantitative analyses are augmented by retrospective think-aloud sessions to provide insight into the subjective experience. The transcripts from these sessions are further analysed for sentiment to provide additional quantitative insight into affective states. Electromyography (EMG) assesses muscle activity by detecting the surface voltages generated during muscle contraction. Positive emotions are identified by recording the activation of the zygomaticus major muscle, which engages when smiling.

Conversely, negative emotions are gauged through the activation of the corrugator supercilii muscle, which is involved in frowning. EMG has been employed in numerous studies to evaluate the valence of emotions. (Cacioppo et al. 1986); (Hassenzahl and Sandweg 2004) Hassenzahl and Sandweg (2004); (Lv et al. 2008); (McCarthy and Wright 2004).

It can be challenging to fully comprehend user behaviour on the web without observing their behaviour in a natural environment (Kellar et al. 2008). For this study we therefore chose to investigate the patterns by using live instances of real platforms rather than mock-ups of fictitious websites. This approach allowed us to capture authentic user interactions and responses, reflecting realistic scenarios that users encounter in their daily online experiences.

For the first pattern (Hidden subscription), we compare the sign-up processes of Facebook and Salesforce. Facebook's sign-up button assumes implicit consent, while Salesforce's checkbox requires explicit consent for the "Hidden Subscription" pattern. This contrast between direct and indirect methods of consent shows the online environment's spectrum from transparent to assumed agreements. These methods reveal the variety of online user consent strategies and their effects on user awareness and legal compliance. For the second pattern ('Hard to cancel), we compare the cancellation processes of Facebook and Pinterest. Again, these sites employ contrasting approaches, with the former being much more prolonged and obfuscated.

3.1. Study Protocol

Prior to the study participants completed a pretest questionnaire to elicit demographic data and prior experiences. Hidden subscription patterns were studied, with participants signing up for various services. Salesforce represents the control condition and uses an explicit consent checkbox, while Facebook represents the DP condition, assuming implicit consent via the sign-up button. Eye-tracking data was collected to capture participants' visual attention during the tasks.

For the 'Hard to cancel' pattern, in which participants tried to delete their accounts, Pinterest represents the control condition and uses a simpler cancellation process, while Facebook represents the DP condition with a complex and obfuscated cancellation process (Mathur et al. 2019). Facial EMG data were collected to capture the muscle activity. The order of presentation of conditions was permuted in both cases to mitigate learning effects and associated bias.

After the tasks, participants were engaged in a retrospective think-aloud session and a semi-structured interview discussing their experiences and thought processes related to the tasks, providing qualitative insights for further analysis. For the ethical considerations of this study, it was ensured that all interactions with the websites were conducted using fictitious account details, ensuring that no actual user data was used at any point during the research. This study utilised a within-subjects experimental design. The independent variable was the UI design (DPs or no DPs). The dependent variables were visual attention, measured through Fixation Duration and Time to First Fixation, emotional responses measured through peak amplitudes from facial EMG; and sentiment, measured using retrospective think-aloud sessions and interview transcripts.

3.2. Participants and Data Collection

The study involved ten individuals, six female and four male master's students from Goldsmiths, University of London. Their ages ranged from 18 to 40. Tobii Pro Lab screen-based eye-tracking software was used to collect eye movement data. The sampling frequency was 60 Hz with a 10 ms interval. A presentation screen with a resolution of 1920x1080 was used to display the web pages. EMG responses were measured via 24 mm Facial EMG electrodes using the ANT Neuro eego™sports amplifier, with a sampling frequency rate of 1000 Hz. The signal was recorded from two muscles, ZM (zygomaticus muscle) and CS (Corrugator Supercilii muscle), to assess emotional mimicry from positive and negative emotional expressions, respectively (Rutkowska et al. 2024).

3.3. EMG Analysis

An essential initial step in analysing EMG signals involved preprocessing the data to eliminate noise and highlight the signal of interest. Noise, including artifacts from electrode movement relative to the skin or interference from electrical equipment, was identified and addressed (Kale and Dudul 2009). The EMG signal was typically filtered using a 20-500 Hz bandpass filter to capture the optimal bandwidth for facial EMG (Van Boxtel 2001). Additionally, a 50 Hz notch filter was used to remove power-line interference. Relevant data segments, corresponding to experimental trials were selected for further processing, and segments with motion artifacts were identified and removed. The data were then full wave rectified, converting negative values to positive.

To smooth the high-frequency rectified EMG signal, it was subsequently passed through a low-pass filter. The EMG signals were analysed by the EMG NeuroKit2 python library by computing the peak amplitudes from the signal as a measure of muscle activation intensity (Schumann et al. 2010). Event related analysis was conducted to highlight the differences between the peak amplitude statistics of the muscle activity of the users while they were completing the 'Hard to Cancel' task.

3.4. Eye tracking analysis

The data was cleaned by removing invalid data points and imputing the missing values using a linear interpolation technique. Noise was filtered to correct artifacts, including blinks, by applying a low-pass Butterworth filter of the 5th order with a 0.5 Hz cut-off, all using Python (Ouzts and Duchowski 2012). The eye-tracking data were analysed by calculating two metrics: Fixation Duration and Time to First Fixation (within the area of interest (AOI)). The dataset was filtered to isolate gaze points within the specified AOI, and the Time to First Fixation values were calculated by measuring the mean time elapsed from when the user entered the webpage to the first fixation on the specified AOI. Fixation Duration was calculated by computing the mean of the total duration of all fixations.

The eye-tracking data were further analysed using visualisations generated by the Tobii Pro Lab software, which recorded participant gaze data and produced heat maps to signify the frequency and duration of gaze fixations. The areas of interest (AOI) included the text above the sign-up button, which indicates the agreement to the Terms and Conditions (via a checkbox in the control condition).

3.5. Sentiment analysis

To enable efficient analysis of the interviews, preprocessing was applied to the raw interview text data using the Python NLTK library (Bird et al. 2009). This phase involved removal of spurious characters including numbers, punctuation and stop words. This was followed by normalisation, folding to lowercase. Finally, tokenisation was applied to identify the individual terms and phrases. Sentiment scores were calculated from the interview verbatims to classify user feedback into positive, negative, and neutral, to represent the emotional tone of the comments and views of the users.

4. RESULTS

4.1. Hard To Cancel

Analysis of the EMG data was based on measuring the highest spike of muscle activity as a proxy for emotional response. Figure 1 shows the average peak amplitude of facial EMG signals for each participant when using the two websites. It can be seen that Facebook elicited higher peak amplitudes, especially when participants tried to delete their accounts, suggesting stronger emotional responses compared to the milder reactions associated with Pinterest. Table 1 shows the distribution of sentiment associated with the two websites, Pinterest and Facebook, categorised into Negative, Positive, and Neutral sentiments. It can be seen that Facebook elicits a greater proportion of negative sentiment and a lower proportion of positive sentiment than Pinterest.

4.2. Hidden subscription

Table 5 shows the distribution of sentiment associated with the two websites, across three categories: Negative, Positive, and Neutral. It can be seen that Facebook elicits a greater proportion of negative sentiment and a lower proportion of positive sentiment than Salesforce. Figure 2 shows fixations on the Facebook sign-up page. These heat maps show that users are most engaged with the central section of the sign-up form, while peripheral information received considerably less visual attention. Figure 3 shows fixations on the Salesforce sign-up page. Cooler colors indicate less interaction, while warmer colors highlight areas of high user engagement, showing a greater focus on the email input field.

5. DISCUSSION

In our study, we used eye-tracking, facial EMG and sentiment analysis to investigate how DPs change the experience of users. The eye-tracking analysis showed how individuals focused more on input fields than on essential sections notably terms and conditions. This represents an opportunity for exploitation via DPs notably hidden subscription. We found that a checkbox makes the terms and conditions more noticeable.

Tables 2 and 3 indicate that the Fixation Duration within Salesforce's sign-up page AOIs was significantly longer, suggesting that users gave more attention to the Terms and Conditions. This is also reflected in the shorter Time to First Fixation for Salesforce AOIs. The Mann-Whitney U test was used to compare the Time to First Fixation and Fixation Duration between Facebook and Salesforce. Table 4 indicates that there is a statistically significant difference in both metrics between Facebook and Salesforce. Participants fixate on Salesforce AOIs significantly faster than on Facebook AOIs. In addition, the result for Fixation Duration indicates that participants spend significantly more time fixating on Salesforce AOIs than on Facebook AOIs.

The result for Fixation Duration indicates that participants spend significantly more time fixating on Salesforce elements than on Facebook elements. The heat map in Figure 2 shows that users are most engaged with the central section of the sign-up form, while the expected AOIs received considerably less visual attention. By contrast, Figure 3 shows the converse pattern. The quantitative data are supported by sentiment analysis which summarises what participants are thinking and feeling. The results show that for the specific tasks investigated, Salesforce and Pinterest offer more positive user experiences than Facebook.

This study highlights the importance of educating individuals about DPs, enabling them to more effectively recognise and avoid being misled by these patterns. However, it is clear that we also need regulatory frameworks that can deal with DPs in a way that takes their wider effects into account. The sentiment analysis and facial EMG results show that the latter can be used as a proxy for affective state, in that differences in muscle activity can align with differences in attitudes.

The results show that, through triangulation with sentiment analysis and eye tracking, EMG can be used as a robust method for investigating DPs. The study examined 'Hidden subscription' and 'Hard to cancel' DPs, but many other types were not explored. Future research should investigate a broader range of DPs to fully understand their impacts. Additionally, the study's small sample size of ten participants may limit the generalizability of the findings. Including participants with no prior experience with live instances of real platforms in future studies would provide additional insights.

6. CONCLUSION

This study examined the influence of deceptive design patterns on user experience in digital platforms, utilising eye-tracking, facial electromyography

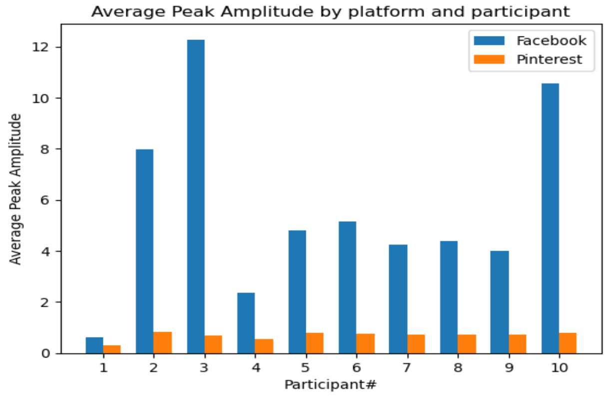

Figure 1: *EMG Response Comparison for Facebook and Pinterest (in Seconds)*

Table 1: *Pinterest & Facebook Sign-up Sentiment Distribution*

	Negative	Positive	Neutral
Pinterest	7.2%	23.5%	69.3%
Facebook	12.7%	8.7%	78.6%

(EMG), and sentiment analysis to uncover how these patterns affect user attention, emotional responses, and perceptions. Facebook elicited higher negative sentiment and lower positive sentiment compared to Salesforce. During account deletion tasks, Facebook induced stronger emotional responses, with a higher peak amplitude of facial EMG signals. Sentiment analysis underscored these effects, showing a clear preference for platforms that adopt more transparent practices. The results advocate for the necessity of ethical design principles, enhanced regulatory frameworks, and increased user education to mitigate the adverse effects of deceptive designs and promote digital environments that prioritise user welfare and autonomy.

Table 2: *Salesforce & Facebook Time to First Fixation & Fixation Duration Standard Deviation (in seconds)*

	Time to First Fixation	Fixation Duration
Salesforce	*0.60*	*0.23*
Facebook	*1.14*	*0.12*

Table 3: *Salesforce & Facebook Time to First Fixation Fixation Duration Mean (in seconds)*

	Time to First Fixation	Fixation Duration
SalesForce	*1.15*	*0.77*
Facebook	*5.43*	*0.36*

Table 4: *Mann-Whitney U and p-value comparison between Facebook and Salesforce (in Seconds)*

	U statistic	P value
Time to First Fixation	*25.0*	*0.011*
Fixation Duration	*20.0*	*0.013*

Table 5: *Salesforce & Facebook Sign-up Sentiment Distribution*

	Negative	Positive	Neutral
Salesforce	0.9%	22.6%	76.5%
Facebook	13.4%	11.4%	75.2%

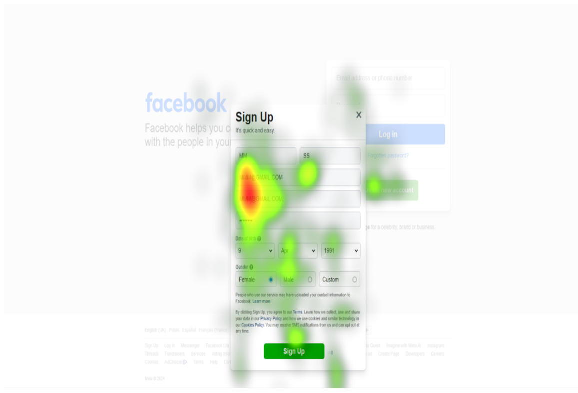

Figure 2: *Heat map of Facebook's sign-up page*

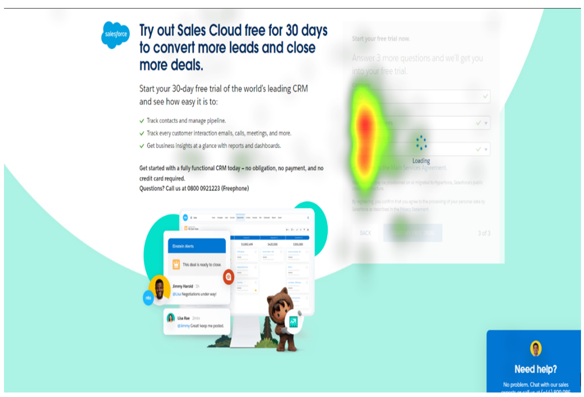

Figure 3: *Heat map of Salesforce sign-up page*

REFERENCES

Bird, S., E. Klein, and E. Loper (2009). *Natural language processing with Python: analyzing text with the natural language toolkit.* " O'Reilly Media, Inc.".

Brignull, H. (2023). Deceptive patterns: Exposing the tricks tech companies use to control you. *(No Title)*.

Cacioppo, J. T., R. E. Petty, M. E. Losch, and H. S. Kim (1986). Electromyographic activity over facial muscle regions can differentiate the valence and intensity of affective reactions. *Journal of personality and social psychology 50*(2), 260.

Di Geronimo, L., L. Braz, E. Fregnan, F. Palomba, and A. Bacchelli (2020). Ui dark patterns and where to find them: a study on mobile applications and user perception. In *Proceedings of the 2020 CHI conference on human factors in computing systems*, pp. 1–14.

Gray, C. M., Y. Kou, B. Battles, J. Hoggatt, and A. L. Toombs (2018). The dark (patterns) side of ux design. In *Proceedings of the 2018 CHI conference on human factors in computing systems*, pp. 1–14.

Hassenzahl, M. and N. Sandweg (2004). From mental effort to perceived usability: transforming experiences into summary assessments. In *CHI'04 extended abstracts on Human factors in computing systems*, pp. 1283–1286.

Ibarra, I. A. (2017). *Should we treat data as labor? Moving beyond'free'*. SSRN.

Kale, S. N. and S. V. Dudul (2009). Intelligent noise removal from emg signal using focused time-lagged recurrent neural network. *Applied Computational Intelligence and Soft Computing 2009*(1), 129761.

Kellar, M., K. Hawkey, K. M. Inkpen, and C. Watters (2008). Challenges of capturing natural web-based user behaviors. *Intl. Journal of Human–Computer Interaction 24*(4), 385–409.

Luguri, J. and L. J. Strahilevitz (2021). Shining a light on dark patterns. *Journal of Legal Analysis 13*(1), 43–109.

Lupiáñez-Villanueva, F., A. Boluda, F. Bogliacino, G. Liva, L. Lechardoy, and T. R. de las Heras Ballell (2022). *Behavioural study on unfair commercial practices in the digital environment: dark patterns and manipulative personalisation*. Publications Office of the European Union.

Lv, H.-R., Z.-L. Lin, W.-J. Yin, and J. Dong (2008). Emotion recognition based on pressure sensor keyboards. In *2008 IEEE international conference on multimedia and expo*, pp. 1089–1092. IEEE.

M. Bhoot, A., M. A. Shinde, and W. P. Mishra (2020). Towards the identification of dark patterns: An analysis based on end-user reactions. In *Proceedings of the 11th Indian Conference on Human-Computer Interaction*, pp. 24–33.

Mathur, A., G. Acar, M. J. Friedman, E. Lucherini, J. Mayer, M. Chetty, and A. Narayanan (2019). Dark patterns at scale: Findings from a crawl of 11k shopping websites. *Proceedings of the ACM on human-computer interaction 3*(CSCW), 1–32.

McCarthy, J. and P. Wright (2004). Technology as experience. *interactions 11*(5), 42–43.

Ouzts, A. D. and A. T. Duchowski (2012). Comparison of eye movement metrics recorded at different sampling rates. In *Proceedings of the symposium on eye tracking research and applications*, pp. 321–324.

Rutkowska, J. M., T. Ghilardi, S. V. Vacaru, J. E. van Schaik, M. Meyer, S. Hunnius, and R. Oostenveld (2024). Optimal processing of surface facial emg to identify emotional expressions: A data-driven approach. *Behavior Research Methods*, 1–14.

Schumann, N. P., K. Bongers, O. Guntinas-Lichius, and H. C. Scholle (2010). Facial muscle activation patterns in healthy male humans: A multi-channel surface emg study. *Journal of neuroscience methods 187*(1), 120–128.

Van Boxtel, A. (2001). Optimal signal bandwidth for the recording of surface emg activity of facial, jaw, oral, and neck muscles. *Psychophysiology 38*(1), 22–34.

Zac, A., Y.-C. Huang, A. von Moltke, C. Decker, and A. Ezrachi (2023). Dark patterns and online consumer vulnerability. *Available at SSRN 4547964*.