"Doctors with Borders"

- Exploring the Perception of Cursor Nudging in Clinical Domains -

Project Report

cs-23-hci-9-08

Aalborg University
Department of Computer Science

SUMMARY

This thesis explores the concept of cursor nudging, a subtle form of digital nudging that influences user decisions through cursor movements and visual cues. The study aims to evaluate the impact of cursor nudging on user behaviour and decision-making in general and clinical settings.

The primary objectives of this research are:

To assess the impact of different types of cursor nudges on user behaviour and decision-making in general contexts.

To evaluate the applicability and effectiveness of cursor nudging in clinical decision support systems (CDSS).

The research is divided into two studies:

Study 1: This study involved general users interacting with various cursor nudges (Border, Crawl, Point, Shape) in a controlled web application environment. The study measured the nudges' impact on user decision-making and overall experience.
Study 2: This study focused on medical professionals using cursor nudges in a clinical decision-making context. It assessed the efficacy of the Border nudge in supporting diagnostic accuracy and gathered qualitative feedback on user experiences and ethical concerns.
Study 1 was focused on gathering quantitative data while Study 2 was focused on qualitative data

Study 1: Results indicated a generally positive reception to cursor nudging. The Shape nudge was particularly effective, receiving high ratings for both impact and user satisfaction. Participants reported feeling guided by the nudges without feeling in control, highlighting the importance of balance between guidance and autonomy.
Study 2: In the clinical setting, the Border nudge was found to significantly support clinicians by managing extensive medical knowledge and causing participants to be more deliberate with their choices. However, concerns about potential over-reliance on AI and the need for transparency were prominent. Qualitative feedback emphasised the benefits of time-saving and additional support.

The discussion integrates findings from both studies, emphasising the potential of cursor nudging to enhance decision-making. In general contexts, cursor nudging effectively influenced user behaviour subtly. In clinical settings, it provided valuable support for medical professionals.

Limitations include small sample sizes and a focus on immediate impacts without assessing long-term effects. Future research should involve larger, more diverse participant groups and explore the long-term impacts of cursor nudging.

This research demonstrates the potential of cursor nudging as a valuable tool for enhancing decision-making in both general and clinical contexts. It highlights the importance of ethical considerations and user acceptance in the responsible implementation of AI-driven cursor nudges. The findings contribute to the fields of HCI and healthcare technology, providing a foundation for future research and development of intuitive, supportive, and ethical AI systems.

"Doctors With Borders": Exploring the Perception of Cursor Nudging in Clinical Domains

Saif Ahmed Khan
sak22@student.aau.dk
Aalborg University
Aalborg, Denmark

ABSTRACT

This thesis explores the concept of cursor nudging, a subtle form of digital nudging that influences user decisions through cursor movements and visual cues. Two studies were conducted to address key questions regarding the impact of cursor nudges on user behaviour and decision-making, and the potential applicability of such nudges in clinical decision support systems (CDSS). Study 1 focused on general users, assessing their attitudes towards different types of cursor nudges and their influence on decision-making processes. The results indicated a generally positive reception, with visual nudges receiving the highest ratings for effectiveness and user satisfaction. Study 2 extended the investigation to the clinical domain, involving medical professionals to evaluate the Border nudge's effectiveness in aiding diagnostic decisions. This study revealed that cursor nudging could support clinicians by managing extensive medical knowledge and reducing cognitive load, although concerns about over-reliance and the need for transparency were noted. Our results provides insights into the efficacy of cursor nudging, and extending its application to clinical settings.

CCS CONCEPTS

• **Human-centered computing → Empirical studies in HCI**.

KEYWORDS

Nudging, Cursor Nudge, AI, Persuasive Technology, Clinical Decision Support Systems (CDSS)

1 INTRODUCTION

In the evolving field of human-computer interaction (HCI), persuasive technologies have emerged as powerful tools for influencing user behaviour and decision-making processes. These technologies leverage psychological principles to guide users towards desired actions without overtly restricting their choices, a concept known as "nudging" as introduced by Thaler and Sunstein [18]. Nudging, particularly through digital interfaces, has gained significant attention for its potential applications across various domains, including healthcare, where decision support systems (DSS, or clinical decision support systems CDSS) play a crucial role in enhancing clinical outcomes.

This thesis explores the concept of cursor nudging, a subtle form of digital nudging that influences user decisions through cursor movements and visual cues. Previous research has demonstrated the design of cursor nudging in web applications, but its potential, and application in clinical settings remains unexplored. This study aims to fill this gap by investigating the effectiveness and perception of cursor nudging among both general users and clinicians.

2 RELATED WORK

2.1 Previous Research

In our previous work, "Intelligent Cursors: An Implementation of AI Nudging Through Cursors" [12], we explored the application of cursor-based nudges to influence user behaviour in a web application context. This study laid the groundwork for the current research by detailing the design and implementation of various nudge types and their effects on user interaction.

2.1.1 Nudge Types. The study categorised nudges into two main types: physical and visual. Each type aimed to influence user behaviour through different mechanisms, which were further categorised according to the groups created by Caraban et al. [3].

Border Nudge (Friction): This physical nudge temporarily prevented users from moving the cursor onto a button that the AI did not select. By creating a temporary barrier, the nudge introduced friction, making it more difficult for users to choose less favourable options. This nudge aligns with the concept of "confront" nudges, which aim to slow down decision-making and prompt more thoughtful consideration.

Crawl Nudge (Facilitate): This physical nudge guided the cursor towards a button that the AI deemed desirable. The cursor moved towards the button when the user was near it, reducing the effort required to make the preferred choice. This type of nudge fits within the "facilitate" category, as it makes it easier for users to select the preferred option.

Pointer Nudge (Facilitate): This visual nudge used the cursor's natural arrow shape to point towards the AI's preferred button. The cursor continuously oriented itself towards the desired option, providing a subtle visual cue to guide user attention. This nudge also falls under the "facilitate" category, helping users make the preferred choice with minimal effort.

Shape Nudge (Friction): Another visual nudge, Shape Nudge, changed the cursor's shape to indicate approval or disapproval of the hovered option. When the cursor was over an AI-preferred button, it displayed a positive symbol, and a negative symbol otherwise. This nudge introduced a moment of friction by making users reconsider their choice upon seeing the visual feedback.

2.1.2 Implementation. The implementation utilised modern web technologies such as React and Next.js to create an interactive and responsive web application. The nudges were designed to be transparent, ensuring that users were aware of the AI's influence and could make informed decisions.

2.1.3 Task Design. The task design in the previous study involved a quiz format to evaluate the effectiveness of different nudge types. The quiz consisted of multiple sections, each presenting different types of questions ranging from basic mathematical problems to more complex tasks involving Raven's Progressive Matrices [10]. The structure allowed users to engage with the nudges in a controlled environment, providing valuable data on their effectiveness.

This previous work allowed for the design of this current research. It also delves into the intentional design of the cursor nudges in more detail, justifying the choices by utilising Caraban et al. [3], Hansen and Jespersen [7], and Thaler and Sunstein [18], and their work within Persuasive Technology and Nudging, in the paper.

2.1.4 Recap on Nudging. Thaler and Sunstein's 2008 book [18] introduces the concept of "libertarian paternalism" or "nudge," which alters behaviour predictably without restricting choices or significantly changing economic incentives. They define a nudge as part of choice architecture influencing decision-making. Caraban et al.'s [3] systematic review identified 23 nudging mechanisms across six categories, finding no direct link between mechanisms and effectiveness; instead, effectiveness depends on context. Failures may stem from educational deficiencies, sustainability issues, timing, strength, unexpected effects, and strong preferences. Hansen and Jespersen [7] discuss dual process theory [5, 11], dividing nudges into automatic (System 1) and reflective (System 2) minds, with Hansen and Jespersen's [7] discussion of transparency to avoid manipulation.

2.2 Clinical Domain

2.2.1 Nudging in Health Care. Nudging has become a significant strategy to influence the behaviour of healthcare professionals in clinical settings, aiming to improve decision-making and adherence to best practices without restricting choices [14–16]. The review by Sant'Anna et al. [16] offers a comprehensive overview of various nudging strategies and their impacts. This review identified eleven different strategies including accountable justification, goal setting, suggested alternatives, feedback, information transparency, peer comparison, active choice, alerts and reminders, environmental cueing/priming, defaults/pre-orders, and education. Defaults/Pre-Orders, pre-selects the most appropriate medication or treatment, making it the default choice. The authors also state "*We have also included in this category all user-interface design strategies for electronic systems that aim to make a given behaviour more likely*" [16].

Numerous studies have demonstrated the effectiveness of nudging in improving clinical decision-making. For instance, a systematic review found that nudges like changing default options and providing feedback are highly effective in improving clinical practice [13]. Another study highlighted that nudging is more likely to be successful when they fit into the workflow of decision makers [8].

2.2.2 AI and CDSS. There has been an increase of AI related tools in Healthcare, AI algorithms can analyse vast amounts of medical data, identifying patterns that may not be immediately evident to human clinicians, thereby supporting more accurate and timely diagnoses [4, 6]. Esteva et al. [4] developed a deep neural network for skin cancer classification that achieved dermatologist-level accuracy, highlighting the potential of AI to augment clinical expertise in diagnostic tasks. Beyond diagnostics, AI-driven CDSS have shown promise in predicting patient outcomes through predictive analytics [2, 17]. By leveraging historical data and current patient information, these systems can forecast disease progression, potential complications, and treatment responses. This proactive approach enables healthcare providers to implement preventive measures and personalised treatment plans, ultimately improving patient outcomes. Gulshan et al. demonstrated the effectiveness of a deep learning algorithm in detecting diabetic retinopathy, achieving high sensitivity and specificity. [6]. However, the implementation of AI in CDSS raises important considerations related to cognitive biases, ethical implications, and the balance between AI recommendations and clinical autonomy [1]. Ensuring transparency in AI systems is crucial to maintaining clinician trust and ensuring that AI serves as a supportive tool. AI models must be interpretable and explainable, allowing clinicians to understand the rationale behind AI-driven suggestions. Jiang et al. [9] emphasised the need for robust regulatory frameworks to assess the safety and efficacy of AI systems in healthcare. They argued that current regulations are inadequate to fully evaluate AI applications in clinical settings. In conclusion, the integration of AI into CDSS holds significant potential to enhance clinical decision-making by providing accurate diagnoses, predicting patient outcomes, and offering personalised treatment recommendations. However, the successful implementation of these systems depends on addressing ethical considerations, ensuring transparency, and achieving seamless integration into clinical workflows.

3 STUDY 1 - NUDGES ON THE PUBLIC

3.1 Method

For Study 1 we focus on testing the system and suite of nudges mentioned before from the previous research [12]. This study aims to explore several key questions regarding cursor nudging and its impact on user behaviour:

- What are the participants' overall attitudes toward cursor nudging after experiencing it?
- How do participants generally react to different types of cursor nudging?
- Does cursor nudging influence their behaviour and decision-making process during the study?
- What are the motivations for using a system with cursor nudges?

The study was deployed online in the form of a website, The technical details of which can be found in the previous research paper [12]. This study primarily focuses on the gathering of quantitative data, and thus a larger outreach was necessary. The study was online and ran for 2 weeks, to allow enough time for the recruited participants to access it. A total of 35 participants were recruited for this study through online platforms, references, and reaching out to universities and companies. The inclusion criteria for a participant was simple, just the requirement of access to a computer or laptop and basic proficiency in using the internet. The ages of the individuals were between 16 to 44 years. All participants provided informed consent before participating in the study. Participants accessed the study via a provided link, which they opened on a computer or laptop. Upon opening the link, they were greeted with a homepage that provided a brief description of the study,

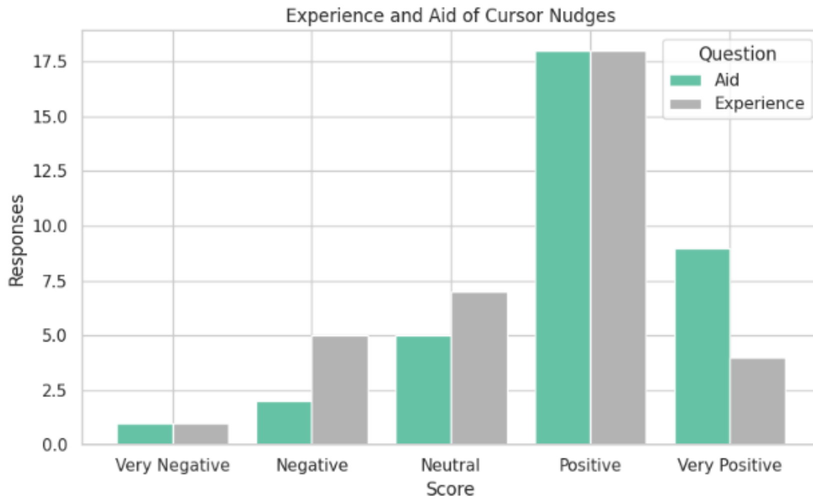

Figure 1: Feedback on Experience and Aid of cursor nudges

including how the tasks will be presented, the use of AI for cursor nudging, and assurances of data confidentiality. Participants then proceeded to the quiz, which consisted of five sections in total. The first section featured no nudges, to allow for the user to become familiar with the format of the tasks. Then each section that followed, featured a different cursor nudging mechanism as outlined in the previous research [12]. Before each section, participants encountered a screen that described the specific cursor nudge they would experience and explained how the AI would interact with them, via the cursor, during that section. Each section contained five questions, making a total of 25 questions across all sections. Participants answered the questions while experiencing the respective cursor nudging mechanism for that section. After completing the quiz, participants were directed to a survey, a Google form embedded on the final page of the website, consisting of 18 questions. This survey primarily utilised the Likert scale to gather quantitative data on participants' experiences and perceptions. Additionally, three open-ended questions were included to collect a small amount of qualitative feedback. Throughout the process, various data points were collected. Quiz performance data, including the accuracy of answers and the time taken to complete each section, were recorded to assess the impact of different cursor nudges on task performance. The data in reference to the performance of the user was sent directly from the website to a PostgreSQL database running on Supabase, an Open-Source backend cloud computing service. Quantitative data from the Likert scale questions, in the survey, provided insights into participants' perceptions of the nudging mechanisms. This data included their effectiveness, intrusiveness, and perceived influence on behaviour. As most questions were using the Likert scale, there may be thoughts the participants would wish to express, so a few open ended questions at the end of the survey provided qualitative data regarding their overall thoughts, opinions, and any ethical worries or implications they perceived. The collected data was anonymised to ensure participant confidentiality and analysed using statistical methods to identify significant patterns and trends in participants' reactions and perceptions.

Following data collection, the results were processed and analysed using Google Colab with Python notebooks. Matplotlib was utilised to create visual representations of the data, making the results digestible. For qualitative data, a Ground up coding approach was employed. Initial codes were created based on common themes and patterns observed in the data. The qualitative data was then reviewed and assigned appropriate codes, allowing for a structured analysis of participants' feedback and opinions. This systematic approach allowed for a comprehensive understanding of the data, highlighting key insights and trends related to the research questions.

3.2 Results

Next we present the most relevant results of this study. These provide insights into participants' overall attitudes towards cursor nudging, their reactions to different types of nudges, the influence of nudging on behaviour and decision-making, and motivations for using a system with cursor nudges. These results are derived from the 35 participants. Every participant solved the same set of questions.

3.2.1 Attitudes and Influence. Participants' overall attitudes toward cursor nudging were assessed through survey questions focused on their general experience and the perceived aid provided by the nudges. The responses indicate a generally positive reception to the concept of cursor nudging. Figure 1 illustrates participants' ratings on their overall experience with cursor nudges and the extent to which they felt the nudges aided their performance. The survey responses were categorised into a five point Likert scale: Very Negative, Negative, Neutral, Positive, and Very Positive. The figure indicates a high approval rating for the experience of the nudges, and similarly for the perceived aid received by the nudges. The majority of participants reported a positive overall experience, and how helpful cursor nudges were, with over 17 responses in this category.

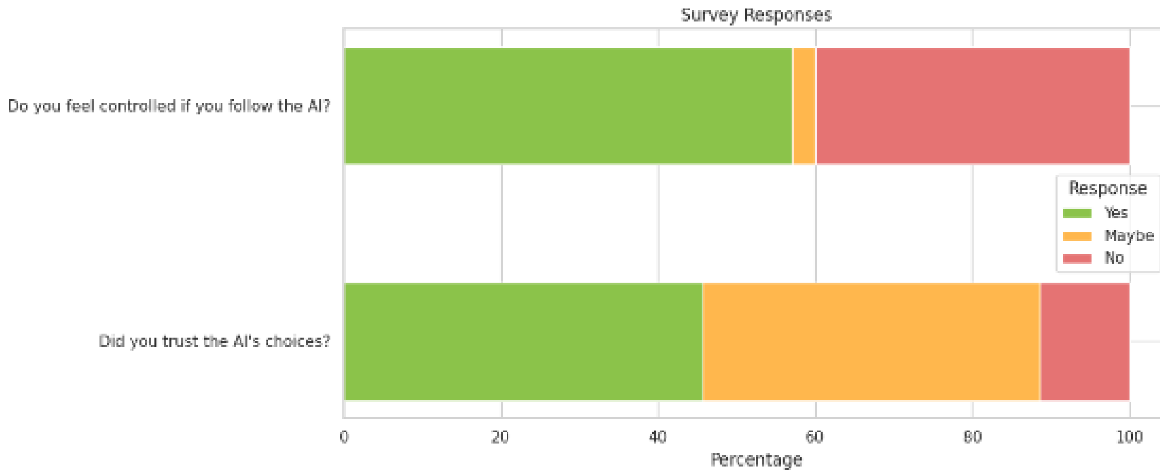

Figure 2: Control and Trust Evaluation

Figure 2 presents the results of two key questions, from the survey, related to participants' feelings of control and trust in the AI's choices. The responses for these questions were taken and are now represented by a stacked bar graph. This is to allow an intuitive understanding of the spread of opinion and allows for better comparison between the two as opposed to pie charts. The majority of participants felt controlled by the AI and the cursor nudges, with a rough 60:40 split between those who feel controlled and those who do not. In regards to trusting the AI's choices and believing they were correct, there is roughly 45 % of participants who did, 10 % who did not completely trust the AI, and roughly 45 % who were unsure and selected maybe.

The influence of cursor nudging on participants' behaviour and decision-making processes during the study was assessed by asking whether they noticed any changes in their choices after interacting with a nudge. This is reflected in figure 3. A majority of participants admit that they noticed a change in their choice with almost 80 % admitting to being either swayed, led, or convinced. 20 % claim they did not perceive any changes in their decision making process. These results suggest that most participants were aware of the influence of the cursor nudging on their choices and behaviour.

Participants' reactions to different types of cursor nudges were assessed using a 5-point Likert scale, with questions focused on the impact the nudge type had on the participant's behaviour, with a higher score indicating a higher impact, and the how the participant perceives the nudge type, with a higher score indicating a more positive experience. Figure 5 presents the average scores for these responses, along with the standard deviation represented by error bars. For the Border nudge, the average impact score was around 3.5 (SD = 1.1), while the perception score was slightly lower, close to 3 (SD = 1.3). The Crawl nudge received an average impact score of approximately 3.5 (SD = 1.06) and a perception score just below 3 (SD = 1.2). The Point nudge was rated higher, with both impact (SD = 0.98) and perception (SD = 1) scores just below 4. The Shape nudge received the highest ratings, with an average impact score above 4 (SD = 0.77) and a perception score close to 4 (SD = 1.2). The standard deviation values, suggests a moderate level of agreement amongst participants. The Shape nudge had the highest average scores and relatively smaller error bars, indicating a more consistent positive reception compared to the other types. These results highlight that while different nudge types were generally perceived positively, there was slight variation in their perceived impact and effectiveness, with the Shape nudge being the most favourable.

Participants were asked about the potential usefulness of cursor movement or changes in their mouse for aiding tasks in their daily life. This question aimed to understand if there is motivations behind using systems with cursor nudges and their perceived practicality. Figure 4 displays the distribution of responses to the question "Would cursor movement or changes in your mouse aid you in any tasks in your daily life?" The responses were categorised as Yes, No, and Maybe. There is an almost even split between all three choices, this indicates a diverse range of opinions on this topic.

3.2.2 Quiz Performance Data. To assess the influence of different nudge types on quiz performance, participants' answers were recorded and analysed for correctness across various nudge conditions. Figure 6 illustrates the distribution of correct and incorrect answers for each nudge type, including No Nudge, Border, Crawl, Point, and Shape. The green bars represent the percentage of correct answers, while the red bars indicate the percentage of incorrect answers. No Nudge: Participants had a high percentage of correct answers (approximately 95%), indicating strong performance without any nudging intervention. Border Nudge: The correct answers decreased slightly, with around 80% correctness, suggesting some influence on participants' performance. Crawl Nudge: This nudge type resulted in a notable drop in correct answers, with approximately 60% correctness, indicating a potential negative impact on performance. Point Nudge: The performance improved significantly with this nudge, showing around 90% correct answers, suggesting this nudge type was effective in guiding participants to correct answers. Shape Nudge: The Shape nudge resulted in the highest

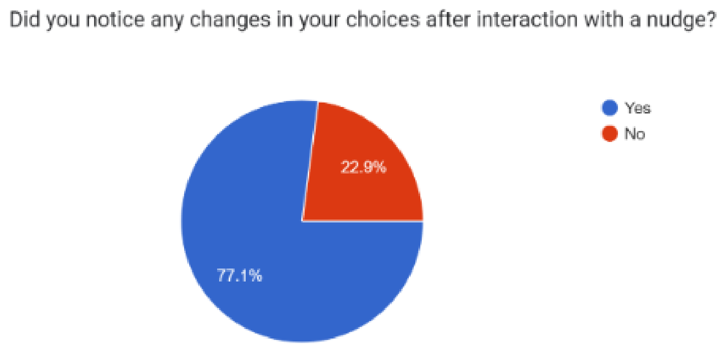

Figure 3: Self Perceived Behavioural Change

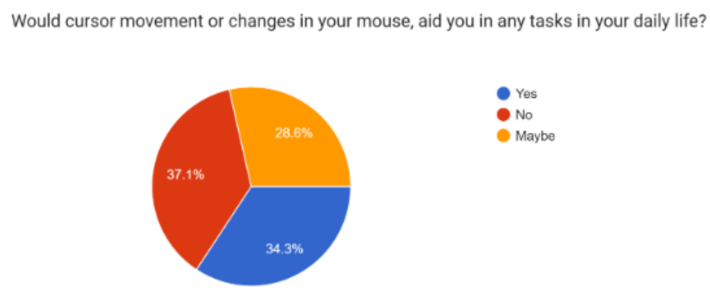

Figure 4: Motivation for Use

performance among the nudge types, with nearly 95% correctness, closely matching the No Nudge condition.

The high correctness rate in the No Nudge condition serves as a baseline for comparison. It indicates that participants performed well without any nudging assistance. Border and Crawl nudge types, showed a decrease in the percentage of correct answers, with the Crawl nudge having the most pronounced negative impact. This suggests that these types of nudges might have introduced distractions or confusion, negatively affecting performance. The next two nudges improved on performance levels. The Point nudge showed significant effectiveness, guiding participants towards the correct answers. The Shape nudge demonstrated the highest performance, indicating its strong positive impact on guiding participants effectively. An ANOVA, using python library scipy, was conducted to determine if there were statistically significant differences in quiz performance across the different nudge types. The results of the analysis indicated a significant effect of nudge type on quiz performance: ($F = 8.63$, $p > 0.001$). The F-statistic of approximately 8.63 suggests a significant variation between the mean percentages of correct answers for the different nudge types. The p-value is extremely low >0.001. This indicates that the differences in quiz performance across the nudge types are statistically significant.

Completion time for participants, aids in assessing the influence of different nudge types. Figure 7 illustrates the quiz completion time per section, in seconds. All 5 sections are visualised as box plots to show the distribution of completion times. The figure shows the differences between each nudge type with the highest median completion time belonging to the Crawl nudge and the lowest to the Shape nudge. The No Nudge section can help to serve as a base case as there was no influence on the participants. The median completion time was around 75 seconds ($SD = 50.64$), this shows moderate variability in the completion times. Border completion time was slightly lower than No Nudge, ($SD = 55.56$), indicating similar variability, the higher standard deviation and the outliers can represent the participants who may have had a slightly more difficult experience in navigating the nudge. Crawl as mentioned before had the highest median completion time ($SD = 53.84$), however it is similar to No Nudge and Border. Pointer sees the first significant reduction in the median completion time ($SD = 45.28$) with a lower standard deviation, indicating lower variability. The Shape nudge, with its lowest median completion time ($SD = 30.82$) and smallest standard deviation, indicates the most consistent performance in the study. ANOVA was used to analyse the results of the times between the groups, ($F = 4.375$, $p = 0.002$). The F-statistic indicates there is a significant amount of variance between the means of the differing nudge types, as opposed to the variance within each nudge type. The p-value shows there is a statistically significant difference in completion times between the nudge types. These results show that the nudge types can have an impact on quiz performance. Tukey's HSD, using python library statsmodels, was used to perform post-hoc analysis to help identify the pairs of nudge types that have significant differences in completion time. This analysis shows a significant difference with Shape in comparison to Border ($MD = 32.6$, $p = 0.03$) or No Nudge ($MD = 41.3$, $p = 0.003$),

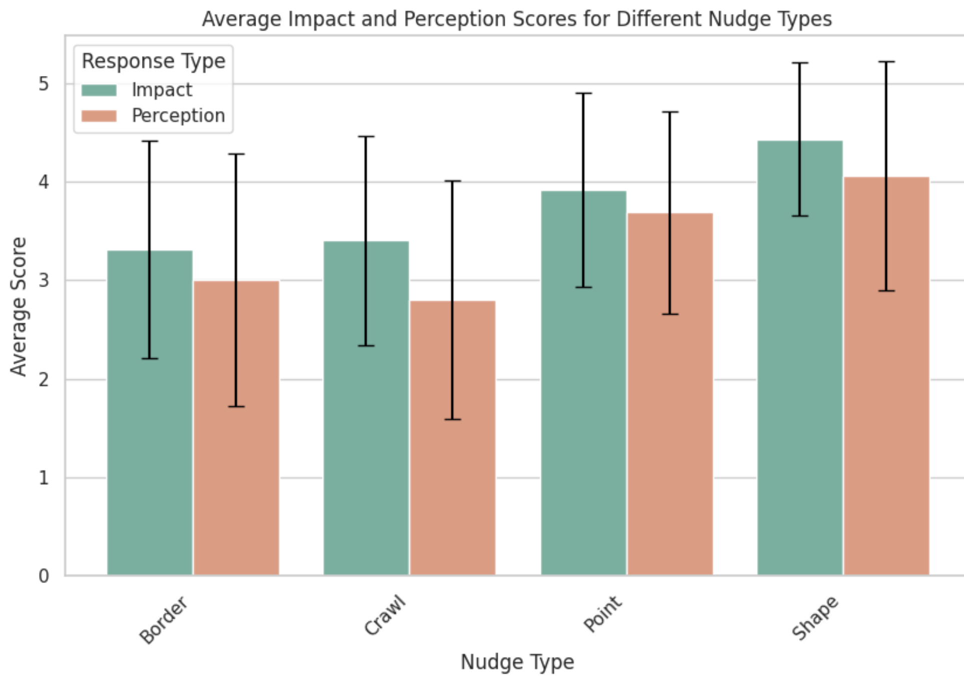

Figure 5: Impact and Perception Graph of each Nudge Type

the the mean difference in favour of Shape. These mean differences indicate Shape's significantly lower completion times compared to the case with no influence, and with the friction creating Border nudge.

3.2.3 Qualitative Responses. The qualitative responses at the end of the survey allowed for participants to provide their perceptions of the cursor nudges and any ethical concerns they could have. The participants were asked if they have any suggestions or what their thoughts were, so they have a chance to share any extra thoughts or considerations that were not addressed. These would be further categorised into, suggestions for nudges visibility, critique, and positive feedback. When looking at the question "Do you have any ethical concerns about AI cursor nudges?", there was roughly 50% of participants who expressed they had no concerns at this current moment. The other responses were categorised into three main themes: Control/Responsibility, Bias/Transparency, and Trust/Difficulties.

Roughly seven participants expressed no changes were needed to the system of nudges, with a few more praising the visual types of nudges over the physically controlling ones. There were also suggestions for improving the visibility of the nudges, as some participants had not noticed it that much, or their own preference to button highlighting to indicate choice instead of cursors. There were also suggestions for allowing the nudges to be activated via the will of the user, as in when the user asks for them. The overall critiques of the nudges were centred around the first two types of nudges, Border and Crawl, where they felt intrusive or seemed confusing.

Participant 3: "maybe the AI could have been more active as I didn't notice it as much "

Participant 7: "The "pointing" and the simple "X" and "correct" marks were the most visually interesting and intuitive for me."

Participant 9: "I prefer button highlighting than cursor nudges to show the user the recommended (AI or not) option."

Participant 17: "the crawl nudge seemed like it was all over the place . at least for a person attempting it for the first time."

Participant 22: "The first movement type nudge felt very intrusive and would be a nightmare to use."

Participant 23: "If you mean the nudges you've implemented, I think they do what they are designed to do well."

Participant 32: "It works very good, but can be a bit too intrusive in day-to-day tasks"

Participant 35: "If nudges are sort of meant as a hint type of thing, perhaps there should instead be an active will from the user to get it (e.g. hold mouse 1 for slight gravitation)"

Control and Responsibility. Participants expressed concerns regarding the potential loss of control and increased reliance on AI, which could lead to laziness and a lack of independent thinking. Some participants felt that the nudges made them dependant, reducing their autonomy and encouraging a lazy approach to the tasks, essentially blindly following the nudges suggestion without critical thinking. Participants were also concerned about the possibility of nudges steering people towards specific actions, manipulating the user's behaviour or tricking the user into selecting other options. One participant raised concerns on responsibility and that users should be responsible for their own inputs.

Participant 1: "Yeah, lets not do that. Removes ones control fully. Good to have as a method to check, but easy to get lazy and just depend on what option it shows you."

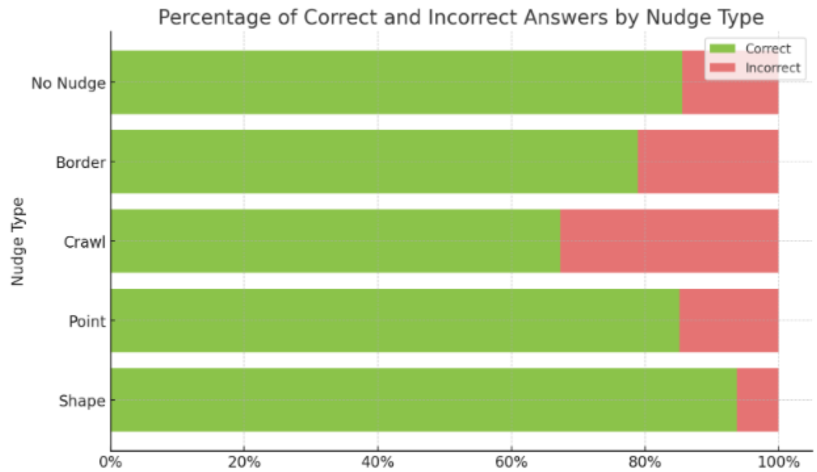

Figure 6: Correct Answer Performance by Nudge Type

Participant 9: "Users should be responsible for their inputs, and should be liable for them. If AI nudges do exist they should be very obvious and opt-in."

Participant 11: "On important forms it would be like sleight of hand in getting a user to select options other than by their own hand."

Bias and Transparency. Participants highlighted issues related to the biases inherent in AI systems and the need for transparency. The lack of transparency can create concerns, and participants highlighted how transparency is important specifically where an AI is trained. This also relates to other concerns on bias. Many participants raised concerns on bias, such as bias of the developer being pushed through, or if there was biased data it was trained on. There is also the fear of spreading misinformation, if not checked for accuracy and fairness.

Participant 7: "...I think transparency on how an AI is taught is extremely important, to avoid problems such as those with certain AI Art..."

Participant 16: "It does have biases that can reflect the developer's agenda"

Participant 20: "That depends on the task. There are the typical concerns regarding AI that has been trained on biased data. Its suggestions might proliferate that bias."

Trust and Difficulties. Participants highlighted concerns regarding trust and difficulties in navigating the nudges. Participants mentioned the importance of trust in AI systems, especially for tasks that require intuition. There was also the matter of if there is a lack of knowledge by the user, it is difficult to evaluate the AI's suggestions critically. This leads to the difficulties in the use of the nudges as well. The nudges disrupted some participant's thought process, leading difficulties in productivity. Some even found it difficult to go against the nudges.

Participant 10: "It disrupted my own thinking, did not feel I could trust what I saw of nudging"

Participant 13: "I feel like I trusted the AI too much, as I found myself not putting as much effort to answer the questions myself. it felt mentally impossible to go against the nudges..."

Participant 29: "I think it really depends on the context of which it is used. A worry might be that you know the answer it nudges you to, is correct (example - capital of England), and then if you are prompted something you don't know, you kinda just trust it to be correct."

4 STUDY 2 - NUDGE ON DOMAIN (CLINICAL)

Building on the findings from Study 1, Study 2 investigates the potential applicability of cursor nudging within the clinical domain, specifically in Clinical Decision Support Systems (CDSS). Despite the generally positive reception of cursor nudging and its demonstrated impact on performance and decision-making in a general population, there was a split in opinions on its usefulness in specific professional domains. Given the critical nature of decision-making in the medical field, we chose to explore the efficacy and perceptions of cursor nudging in supporting clinical tasks. The focus of this study is on the Border nudge type, selected for its characteristic friction effect that encourages deliberate, system 2 thinking without directly suggesting specific answers.

4.1 Method

For this study, we aim to explore several questions to be able to determine if the impact cursor nudges had on the previous study are consistent when moving to a specific domain, and to answer the weaknesses of study 1:

- What is the perception and impact of the border nudge when applied in the clinical domain?
- How are medical professionals' behaviour and decision-making influenced by a cursor nudge?
- What are the motivations for using a system with cursor nudges in the clinical field?

The study followed the within-subjects design where all participants received the same set of questions, answer options, and

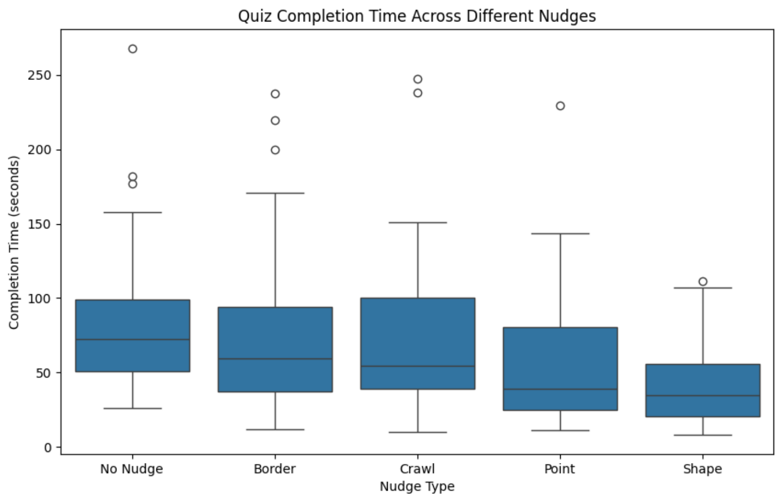

Figure 7: Time to Complete, Grouped by Nudge Types

predetermined incorrect AI suggestions. The only varying factor was the participants' individual experience levels. This study was conducted on a smaller scale, with a total of six participants, with varying levels in the medical/clinical field. Two participants were early medical school students, two were late medical students with experience in the field through clinical rotations, the last two were experienced practitioners with over 11+ years in the field. The inclusion criteria outside the experience in the medical/clinical field, was similar to the previous study, requirement of a computer or laptop, and understanding of basic internet use, with the added requirement of not having taken part in the first study.

4.1.1 Task Changes. The study utilised a modified version of the software used in the previous study. This is due to adjusting it to a clinical context. Figure 8 reflects the new task. The alterations are as follows:

Selection Option: Instead of four buttons, a radio list of options was presented for each question, with a submit button on the right hand side. This was to ensure a participant would not accidentally experience a border nudge before thinking about the questions as they are more complex, and require more time.

Question Content: Questions were sourced from the Step 1 USMLE sample exam [19] and focused on diagnosing patients based on symptom descriptions. This choice reflects the complexity and critical nature of clinical decision-making.

Number of Questions: The number of questions were reduced to 10 per section with only the one section (compared to five per section totalling 25 in Study 1) to account for the increased difficulty and cognitive load.

Supporting Images: Some questions included relevant images alongside symptom descriptions to aid in diagnosis, these were placed on the left hand side.

Nudge Type: Only the border nudge was used, due to its ability to induce thoughtful reconsideration without leading users directly to an answer. This approach aligns with the ethical standpoint that AI should not dictate clinical decisions but rather prompt careful evaluation. Any other cursor nudge could not fill this role as they all have high degrees of "Facilitation", where as Border leads itself to creating "Friction" instead.

Supporting Images: Some questions included relevant images alongside symptom descriptions to aid in diagnosis, these were placed on the left hand side.

4.1.2 Procedure. The participants accessed the study through a provided link, where they were greeted with a homepage that provided a brief description of the study, including how the tasks will be presented, the use of AI for cursor nudging, and assurances of data confidentiality. Participants then proceeded to the quiz, which consisted of a single section this time. This time the only section featured the Border nudging mechanism. The study consisted of 10 total questions centred around the clinical field, specifically the questions listed where listing symptoms of a patient, with relevant images if needed on the left, for the participant to diagnose. The cursor nudges, would be provide incorrect suggestions 30% of the time. This is to test the participants' ability to identify and resist incorrect nudges. As we are testing this in the clinical domain, we do not want to give the assumption to the participants, that the AI in infallible. After completing the quiz, participants were directed to a survey, consisting roughly 14 questions. A few of the questions had sections for the participant to explain their thoughts or answer after each one. There was also significantly more open ended questions requiring a qualitative analysis, in this survey compared to study 1, this is due to the lower sample size, and the interest in discovering if the motivations for cursor nudging exists in the clinical domain.

Data collection was similar to the previous study, in which the answers that were chosen by the user were recorded, as well as the time to complete the tasks. However, there is a greater focus on the qualitative data gathered from the survey responses. These were collected and analysed, with graphs being created and grouping and coding being done on the relevant answers for the key questions.

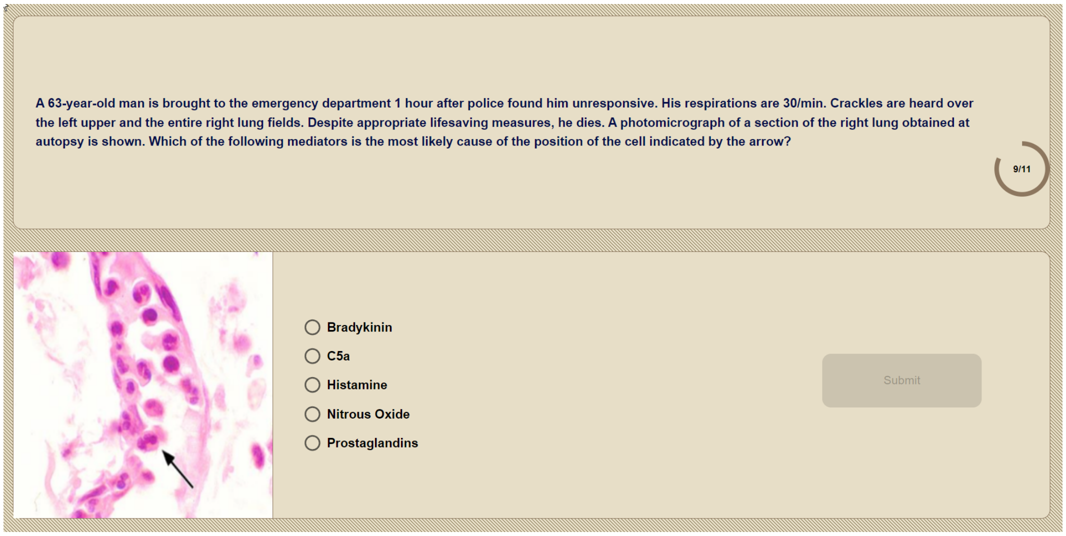

Figure 8: New Task Setup for Study 2

4.2 Results

This section presents the findings from Study 2, focusing on the impact and perception of the border nudge in a clinical context, the influence of cursor nudging on medical professionals' behaviour and decision-making, and the motivations for using cursor nudging systems in the clinical field. The data was gathered from the six participants, varying in their professional experience, including medical school students, residents, and experienced practitioners. There were questions, using a five-point Likert scale, to gauge opinions from participants in a digestible format however, we will focus on the qualitative responses that came with each question. This is to gain a better understanding of the thought of the clinician participants. Any responses that are shown as quotes may have slight alterations such as, grammar changes and removal of punctuation or additional spaces, this is just for readability and the content has not been changed.

4.2.1 Overall experience. Participants overall experience towards the cursor nudging was assessed through a few survey questions. When asked to explain their overall experience with the cursor nudges, majority of participants, generally found the nudges to be a helpful guide, providing clarity and aiding in their decision-making process. This also includes comments due to participants' lack of confidence in their own answers, as well as description of AI contributing to a positive overall experience. There were also a few comments in the form of feedback, referring to incorrect answers.

Participant 1: "I wasn't really confident with the answers so i trusted the AI"

Participant 2: "The experience was really smooth, I was confused in one question and I felt the resistance caused by AI, which made my answer choice clear."

Participant 3: "I think it showed right answers for some questions and wrong for others"

Participants were asked about navigating the study, and how easy it was to understand the nudges. No grouping was done for this as all participants responded with positive feedback and the ease in navigating through it. All responses were either saying "it was easy" or "pretty straightforward".

The responses to the question "How did you feel about the AI blocking the submit button when it disagreed with your choice?" was grouped into two main themes based on the participants' feedback, Helpful and Annoying. Participants who found the AI blocking the submit button helpful mentioned that it made them reconsider their choices and guided them to the right answers. This theme also includes mixed feelings where participants acknowledged the helpfulness but also noted potential over-reliance on AI. Participants who found the nudge blocking the submit button annoying mentioned that it led to self-uncertainty. There was a split of four users referring it to as helpful, and two who mentioned it is annoying.

Participant 2: "I have mixed opinions. It can be helpful as it will guide you to the right answer, but that also makes the possibility of too much dependence on it in the future."

Participant 5: "Made me rethink my choice and think about the answer again"

Participant 3: "It definitely is a bit annoying but that leads to self uncertainty"

Participant 6: "Extremely annoying"

4.2.2 Behaviour and Decision-Making. Participants were asked several questions to gain an understanding of how their behaviour or decision-making may have been affected. There were a few questions that we analysed, to help better understand their process. The questions were: "How did the AI's nudges affect your decision-making process during the quiz?" "To what extent did you rely on your own knowledge versus the AI's suggestions, and why?"

Decision-Making. All participants stated that there was a perceived impact or change in their decision-making. These range from making it easier for them, to making them uncertain about their own choices. This suggests that although all participants were impacted, there was a difference in how this was received.

Participant 1: "it helped me. Guided me to the right answers"

Participant 2: "It made decision-making easier for me"

Participant 3: "It definitely made me uncertain about my choices that lead me to go through the question again just to be sure"

Participant 4: "Affected a lot as i changed my answers sometimes when it blocked my answer"

Participant 5: "Helped me a lot in navigating what the right answer is"

Participant 6: "Sometimes impacted."

Knowledge Base. Participants were asked on if they relied on their own knowledge, or if they would be more reliant on the AI. There was an observable difference between all responses, three main groups formed with differing opinions. One group (AI users) were the participants who relied on the nudging and AI, one participant's response was "Some of the questions i did based on my knowledge, but mostly relied on AI as it has more data and information". The second group (Both) were roughly in between, partially relying on both, one participant's response for this was "partially relied on both. More on AI in some questions". The final group (Self-knowledge) were those who relied heavily on their own knowledge, one participant's response was "I relied on my knowledge to the extent until I was confused about two options and I couldn't recall the right one. I think mostly relying on self-knowledge makes learning better". In each group there were two participants based on their responses, showing a clear split between their use. If we introduce the experience of the participants into this question, to get a better understanding if that had a correlation to the reliance on AI, we would see that the groups formed are not based on experience. Table 1 outlines the formation of the three groups. Participant experience refers to their experience in the medical field, and the second column refers to which knowledge base they relied on more.

Table 1: Participant Experience Level and AI Reliance

Participant Experience	Reliant On
Low	AI
Low	Self
Medium	Self
Medium	Both
High	Both
High	AI

4.2.3 Motivations. Participants were asked whether they saw a potential use for cursor nudging in the clinical field. Their responses have been analysed and categorised into several themes, that reflect their perceptions and justifications.

Managing Medical Knowledge. Participants highlighted the extensive nature of medical knowledge, indicating that cursor nudging could help manage this vast information.

Participant 1: "It is very helpful in guiding you what the right answer is and saves time as well. Since medicine is a very vast subject and all the knowledge can not be remembered by a human mind..."

Participant 4: "...You don't remember everything in the top of your head so instead of researching everything you cant remember, this nudging is a good way."

Time and Efficiency. The potential for cursor nudging to improve efficiency was frequently noted. Participants suggested that nudges could streamline the diagnostic process, saving valuable time for healthcare professionals.

Participant 1: "It is very helpful in guiding you what the right answer is and saves time as well..."

Participant 4: "Its time saving..."

Participant 5: "As it guides you so its efficient and helpful"

Assistance and Support System. The idea of cursor nudging acting as a supportive tool was another common theme. Participants felt that nudges could assist in making more informed decisions by providing an additional layer of support, being there to also reduce bias from the doctor or due to the problems arising from working long hours.

Participant 2: "Cursor nudging can be useful in clinical field. As doctors are checking many patients on a daily basis and there can be a confirmation bias. This feature of cursor nudging can assist the doctor in the right diagnosis through judging the sign and symptoms in the clinical notes."

Participant 3: "Maybe it could be used in diagnostic lab test. Based on patient's symptoms it could help which tests to perform and which not, just so nothing is missed. But of course with much more precision."

Participant 6: "Perhaps doctors working long hours can use a sound double check system to improve efficacy of diagnosis"

The majority of participants agreed that there is a use for cursor nudging in the clinical field. When asked "Is there a use of cursor nudging in a clinical field?", participants responded with 66% stating there is a use, with 33% selecting "maybe", none of the participants selected "no". The participants who selected maybe also provided insight in where this could operate or its use. This data suggest generally a positive outlook on the integration of cursor nudging into clinical decision support systems, with a focus on its potential to aid in managing extensive medical knowledge, saving time, providing assistance, provided however, the precision is increased.

5 DISCUSSION

Two studies were carried out to explore the effectiveness and perception of cursor nudging in both general and clinical contexts. Through these two distinct studies, we aimed to answer key questions regarding the impact of cursor nudges on user behaviour and decision-making, and the potential applicability of such nudges in clinical decision support systems (CDSS).

5.1 User Behaviour and Decision-Making

The first study focused on understanding the participants' overall attitudes towards cursor nudging, their reactions to different types of nudges, and whether these nudges influenced their behaviour and decision-making processes. The results, as shown in Figure 1, indicated a generally positive reception to cursor nudging, with most participants reporting that the nudges aided their performance. Roughly, 60% of participants rated their experience with cursor nudges as positive or very positive.

Figure 2 highlights participants' feelings of control and trust in the AI's choices. While approximately 60% of participants felt controlled by the AI, about 45% trusted the AI's choices. This split highlights the delicate balance between guidance and autonomy that cursor nudging must maintain to be effective and ethically acceptable. This is consistent with Hansen and Jespersen's [7] emphasis on transparency to avoid manipulation.

The effectiveness of different nudge types is detailed in Figure 5. The Shape nudge, which provided visual feedback through cursor shape changes, received the highest ratings for both impact and perception, with an average impact score above 4.0. This supports the idea that visual cues can be potent in influencing user behaviour subtly yet effectively. In contrast, the Border and Crawl nudges received lower scores, indicating some participants found them intrusive or confusing.

5.2 Performance

Figures 6 and 7 provide insights into the impact of different nudge types on quiz performance and completion times. Figure 6 illustrates the distribution of correct and incorrect answers for each nudge type. The Shape nudge resulted in the highest performance, with nearly 95% correctness, closely matching the No Nudge condition, which serves as a baseline. This indicates that the Shape nudge effectively guided participants towards correct answers without introducing significant distractions. However, there may be the case that the Shape nudge was the easiest to confirm which option the AI had chosen, as such, participants may have just quickly chosen without thinking, as Shape was also situated at the end of the quiz, so users may have also gotten restless. The completion time in figure 7 may also corroborate this, as it is significantly lower than the rest. In contrast, the Crawl nudge showed a notable drop in correct answers, with approximately 60% correctness, suggesting that this type of nudge might have introduced confusion or distraction, negatively affecting performance. The movement of the mouse towards an option, may have slowed down the process of selecting their choice, where they may have finished thinking about the answers, only to have been dragged out and confused by the interaction. The Border nudge also resulted in a decrease in correctness, though to a lesser extent than the Crawl nudge, emphasising the need for careful design and implementation of such interventions to avoid reducing decision accuracy. These results can also suggest that the harder types of nudges to navigate, and that do not present the answer immediately, lead to the user being more deliberate, as the others may promote laziness.

5.3 Motivations

Study 2 was designed to address the mixed feedback from Study 1 regarding the applicability of cursor nudging in professional domains. The split in opinions, combined with the observed improvement in performance and decision-making, warranted a focused investigation in a domain. The clinical domain was selected due to the high stakes and the potential for cursor nudging to enhance careful consideration in diagnostic tasks. By focusing on the border nudge, we aimed to explore its utility in promoting reflective thinking [11], which is crucial in clinical settings. The feedback from medical professionals in Study 2 indicated a clear potential for cursor nudging to enhance clinical decision-making. Participants appreciated the ability of cursor nudges to help manage the extensive knowledge required in medicine, reduce time spent on decision-making, and act as a support system to mitigate cognitive biases and decision fatigue. For instance, participants noted that cursor nudges could assist in quickly guiding them to the most relevant diagnostic options, thereby saving time and improving efficiency. This is particularly important in fast-paced clinical environments where timely decisions can significantly impact patient outcomes. While some participants expressed trust in the AI's recommendations, others were concerned about potential over-reliance. For example, a participant mentioned, "I wasn't really confident with the answers, so I trusted the AI." This highlights the need for balance between AI assistance and clinician autonomy, ensuring that AI serves as a supportive tool rather than a crutch.

The positive reception of cursor nudging in the clinical domain, despite the initial concerns about over-reliance and transparency, highlights its potential as a valuable tool in CDSS. By enhancing the decision-making process and providing an additional layer of support, cursor nudging can help healthcare professionals make more informed and accurate decisions.

5.4 Future work and Limitations

Despite the promising findings, this research has several limitations that should be addressed in future studies. First, the sample sizes for Study 2, which included only six participants. This limited sample size leads to a low statistical power. It also limits the ability to compare the findings of opinions and performance in the test across studies. Future research should involve larger, more diverse participant groups to validate the findings and ensure broader applicability. The studies focused primarily on the immediate impact of cursor nudging on decision-making and did not assess long-term effects or user behaviour changes over time. Longer studies would need to be carried out to see the effect a cursor nudge could have on a user once they use it more frequently. Another limitation is that the study did not explore the potential variations in effectiveness among different types of cursor nudges across various clinical tasks. Future research should investigate which specific nudge types are most effective for different decision-making scenarios within clinical settings. This could help in tailoring nudging strategies to specific tasks, enhancing their overall effectiveness. Additionally, the ethical concerns raised by participants, such as the potential for over-reliance on AI and the need for transparency, highlight the importance of addressing these issues comprehensively.

6 CONCLUSION

This research explored the effectiveness and perception of cursor nudging in influencing user behaviour and decision-making, both in general contexts and within clinical contexts. By conducting two studies, we examined how different types of cursor nudges impact user interactions and assessed the potential applicability of these nudges in clinical settings.

This research also makes contributions to the field of human-computer interaction. One is the demonstration of cursor nudging efficacy. There is empirical evidence on the effectiveness of cursor nudging in influencing user behaviour and decision-making. The findings also show the opinions of users on what type of cursor mechanisms they would prefer, and how the physical or visual nudges impact user choices. This work also extends the application of cursor nudging to clinical settings. Nudging is already a concept in the clinical domain, however not through cursor nudging. Border type nudge shows that there is value in aiding the clinician in making informed decisions, by promoting deliberate thinking.

The findings from Study 1 demonstrated that cursor nudges, particularly the Shape nudge, positively influenced user behaviour by subtly guiding their choices without overwhelming them. The study highlighted a delicate balance between providing helpful guidance and maintaining user autonomy, emphasising the importance of transparency and user control in the design of persuasive technologies.

Study 2 extended this investigation to the clinical domain, focusing on the Border nudge's effectiveness in aiding medical professionals' decision-making processes. The results revealed that cursor nudging could significantly support clinicians by managing extensive medical knowledge and reducing cognitive load. Participants appreciated the time-saving benefits and additional support provided by the AI, although they also similarly raised concerns about potential over-reliance and the need for transparency.

Despite the promising results, this research has several limitations as mentioned in the future work section. The small sample size in Study 2 may affect the generalisability of the findings. Future research should involve larger, more diverse participant groups and explore long-term effects of cursor nudging. Additionally, variations in the effectiveness of different nudge types across various clinical tasks warrant further investigation.

Cursor nudging presents a valuable tool for enhancing decision-making in both general and clinical contexts. The positive reception and significant impact on user behaviour, highlights its potential for integration into decision support systems. However, ethical considerations and user acceptance are critical factors that must be addressed to ensure the responsible and effective use of cursor nudges. Continued research and innovation in this field will be essential to fully realise the benefits of persuasive technology in improving decision-making and outcomes in healthcare.

REFERENCES

[1] Anne Kathrine Petersen Bach, Trine Munch Nørgaard, Jens Christian Brok, and Niels Van Berkel. 2023. "If I Had All the Time in the World": Ophthalmologists' Perceptions of Anchoring Bias Mitigation in Clinical AI Support. *Conference on Human Factors in Computing Systems - Proceedings* (4 2023). https://doi.org/10.1145/3544548.3581513

[2] Wouter Bulten, Hans Pinckaers, Hester van Boven, Robert Vink, Thomas de Bel, Bram van Ginneken, Jeroen van der Laak, Christina Hulsbergen-van de Kaa, and Geert Litjens. 2020. Automated deep-learning system for Gleason grading of prostate cancer using biopsies: a diagnostic study. *The Lancet. Oncology* 21, 2 (2 2020), 233–241. https://doi.org/10.1016/S1470-2045(19)30739-9

[3] Ana Caraban, Evangelos Karapanos, Daniel Gonçalves, and Pedro Campos. 2019. 23 Ways to Nudge: A review of technology-mediated nudging in human-computer interaction. *Conference on Human Factors in Computing Systems - Proceedings* (5 2019). https://doi.org/10.1145/3290605.3300733

[4] Andre Esteva, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun. 2017. Dermatologist–level classification of skin cancer with deep neural networks. *Nature* 542, 7639 (2 2017), 115. https://doi.org/10.1038/NATURE21056

[5] Bertram Gawronski and Laura A. Creighton. 2013. (PDF) Dual Process Theories. https://www.researchgate.net/publication/281178902_Dual_Process_Theories

[6] Varun Gulshan, Lily Peng, Marc Coram, Martin C. Stumpe, Derek Wu, Arunachalam Narayanaswamy, Subhashini Venugopalan, Kasumi Widner, Tom Madams, Jorge Cuadros, Ramasamy Kim, Rajiv Raman, Philip C. Nelson, Jessica L. Mega, and Dale R. Webster. 2016. Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs. *JAMA* 316, 22 (12 2016), 2402–2410. https://doi.org/10.1001/JAMA.2016.17216

[7] Pelle Guldborg Hansen and Andreas Maaløe Jespersen. 2013. Nudge and the Manipulation of Choice: A Framework for the Responsible Use of the Nudge Approach to Behaviour Change in Public Policy. *European Journal of Risk Regulation* 4, 1 (3 2013), 3–28. https://doi.org/10.1017/S1867299X00002762

[8] Joseph D Harrison and Mitesh S Patel. [n. d.]. AMA Journal of Ethics ® September 2020, Volume 22, Number 9: E796-801 MEDICINE AND SOCIETY Designing Nudges for Success in Health Care. ([n. d.]). www.journalofethics.org796

[9] Fei Jiang, Yong Jiang, Hui Zhi, Yi Dong, Hao Li, Sufeng Ma, Yilong Wang, Qiang Dong, Haipeng Shen, and Yongjun Wang. 2017. Artificial intelligence in healthcare: past, present and future. *Stroke and Vascular Neurology* 2, 4 (12 2017), 230–243. https://doi.org/10.1136/SVN-2017-000101

[10] John and Jean Raven. 2003. Raven Progressive Matrices. *Handbook of Nonverbal Assessment* (2003), 223–237. https://doi.org/10.1007/978-1-4615-0153-4{_}11

[11] Daniel Kahneman. 2017. *Thinking Fast and Slow.*

[12] Saif Ahmed Khan. 2024. *Intelligent Cursors: An Implementation of AI Nudging Through Cursors.* Technical Report. Aalborg University. https://drive.google.com/file/d/1sQbfjMHDFJ7tAMj8k3IkvUxcpXivk3NV/view?usp=sharing

[13] Briana S Last, Alison M Buttenheim, Carter E Timon, Nandita Mitra, and Rinad S Beidas. 2021. Systematic review of clinician-directed nudges in healthcare contexts. *BMJ Open* 11 (2021), 48801. https://doi.org/10.1136/bmjopen-2021-048801

[14] Rosanna Nagtegaal, Lars Tummers, Mirko Noordegraaf, and Victor Bekkers. 2019. Nudging healthcare professionals towards evidence-based medicine: A systematic scoping review. *Journal of Behavioral Public Administration* 2, 2 (10 2019). https://doi.org/10.30636/JBPA.22.71

[15] Luisa M. Pettigrew and Nicholas Mays. 2021. What role is there for 'nudging' clinicians? *British Journal of General Practice* 71, 703 (2 2021), 82–85. https://doi.org/10.3399/BJGP21X714857

[16] Anita Sant'Anna, Andreas Vilhelmsson, and Axel Wolf. 2021. Nudging healthcare professionals in clinical settings: a scoping review of the literature. *BMC Health Services Research* 21, 1 (12 2021), 1–14. https://doi.org/10.1186/S12913-021-06496-Z/FIGURES/5

[17] Peter Ström, Kimmo Kartasalo, Henrik Olsson, Leslie Solorzano, Brett Delahunt, Daniel M Berney, David G Bostwick, Andrew J Evans, David J Grignon, Peter A Humphrey, Kenneth A Iczkowski, James G Kench, Glen Kristiansen, Theodorus H van der Kwast, Katia R M Leite, Jesse K McKenney, Jon Oxley, Chin-Chen Pan, Hemamali Samaratunga, John R Srigley, Hiroyuki Takahashi, Toyonori Tsuzuki, Murali Varma, Ming Zhou, Johan Lindberg, Cecilia Lindskog, Pekka Ruusuvuori, Carolina Wählby, Henrik Grönberg, Mattias Rantalainen, Lars Egevad, and Martin Eklund. 2020. Artificial intelligence for diagnosis and grading of prostate cancer in biopsies: a population-based, diagnostic study. *The Lancet. Oncology* 21, 2 (2 2020), 222–232. https://doi.org/10.1016/S1470-2045(19)30738-7

[18] C.R. Thaler and R. Sunstein. 2008. *Nudge: Improving Decisions about Health, Wealth and Happiness.*

[19] USMLE. [n. d.]. *Sample Test Questions Step 1.* https://www.usmle.org/sites/default/files/2021-10/Step_1_Sample_Items.pdf

Received 21 June 2024