Responsible Human-Centered Artificial Intelligence for the Cognitive Enhancement of Knowledge Workers*

Troy R. Weekes[0000−0002−0371−1997] and Thomas C. Eskridge[0000−0002−2117−529]

L3 Harris Institute for Assured Information,
Department of Human-Centered Design,
Florida Institute of Technology, Melbourne FL 32901, USA
tweekes@fit.edu, teskridge@fit.edu

Abstract. Over the past decade, the demand for high-performing knowledge workers (KWs) has grown at an unprecedented rate and shows no signs of slowing. Researchers, designers, engineers, and executives are examples of KWs that perform non-routine, creative work. The work outcomes of KWs as individuals, teams, and organizations play a vital role in the global economy and quality of life. One of the most significant challenges KWs face is balancing stressors on their cognitive and emotional well-being while seeking high productivity. Human cognitive enhancement proposes improving human abilities to acquire and generate knowledge and understand the world. Our cognitive enhancement application for KWs, called the Flow Choice Architecture (FCA), senses their cognitive and affective states, adds context, and recommends appropriate nudges to maximize their healthy flow time. This study provides insights into how FCA implements Human-Centered Design and Responsible Artificial Intelligence (RAI) principles as an interactive AI-powered application that promotes healthy flow performance during knowledge work. FCA applied the RAI tools from Microsoft's Human-AI eXperience Toolkit to evaluate FCA-specific scenarios. By defining FCA as a hybrid recommendation system and conversational AI agent, we found the following categories of human-AI failure scenarios in FCA: input errors, trigger errors, delimiter errors, and response generation errors. We recommend simulating these errors and undesirable behaviors to improve the design of explainable nudges, meaningful metrics, and well-tuned triggers. The outcome of this RAI evaluation was a robust FCA system design that meets the needs of KWs and enhances their capability to thrive and flourish at work.

Keywords: Knowledge Worker · Cognitive Enhancement · Responsible Artificial Intelligence.

* Supported by L3 Harris Institute for Assured Information.

1 Introduction

The transition from a manufacturing economy to a service economy has caused the demand for high-performing knowledge workers (KWs) to grow at an unprecedented rate. The digitalization of productivity tools to support KWs has delivered cloud-based artificial intelligence (AI) services [4], remote presence [7], and workplace analytics [12]. Although these productivity multipliers have been impressive in achieving performance gains, there is a gap in harmonizing these technologies with personal effectiveness and well-being.

As the number of job opportunities for KWs continues to grow, the requirement to quantify and qualify knowledge work will increasingly become the new normal. Coupled with the confluence of AI and bio-sensing technology, we anticipate an avalanche of opportunities to provide KWs with bio-signal analytics. Cognitive enhancement proposes to improve human abilities to acquire knowledge and understand the world and improve their individual performance.

One of the most significant challenges KWs face while seeking high productivity is balancing stressors on their cognitive and emotional well-being. This research has the goal of developing a personalized service that enhances the cognitive abilities and emotional well-being of individual KWs. How might we design and evaluate this service according to human-centered AI (HCAI) interaction best practices and the Responsible AI (RAI) principles of reliability, fairness, explainability, and privacy?

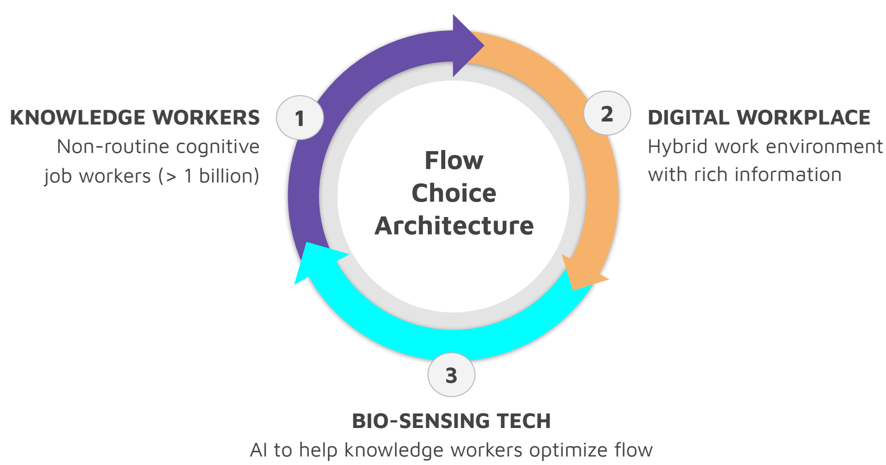

Fig. 1. FCA enhances KWs in Digital Workplaces

This paper describes the human-aware and context-aware neurotechnology AI system known as the **Flow Choice Architecture (FCA)** shown in Figure

1. FCA seeks to increase the healthy time KWs spend in the flow state. FCA's central hypothesis states that if the cognitive and affective states associated with high-performance knowledge work can be measured and contextualized, then timely "nudges" [29, 28] may modify the KW's flow experience to facilitate entry into flow and to extend its duration.

In Section 2, we describe KWs by their characteristics, features, and capabilities. We outline their knowledge work environments, work artifacts, and work resources. Importantly, we provide an understanding of what knowledge work is and what distinguishes deep work from shallow work. We identify the challenges of distractions, interruptions, and fatigue confronting KWs.

In Section 3, we delve into the concept of flow to understand why the KW desired to be in this state. We articulate the high-level framework of FCA as a personalized HCAI, explaining how FCA enhances KWs in the digital workplace.

In Section 4, we explore the concept of a choice architecture, which is the taxonomy of nudges and their affordances. We explain the significant difference between FCA and traditional neurofeedback. We define specific research questions around the potential failures of FCA as a recommendation system and conversational AI.

In Section 5, we apply the Human-AI Experience (HAX) Playbook to the design of FCA. We examine FCA as a recommendation system and conversational AI using the playbook. There was significant benefit in identifying human-AI failure scenarios before coding FCA.

In Section 6, we review the results from the HAX Playbook. There are potential failures for FCA's recommendation system and conversational AI components, and recommendations for each failure source. The recommendations provide simulations to understand better the errors that cause the failures.

In Section 7, we discuss how responsible AI tools provide a cost-effective method to mitigate risks early in the design phase. We identify how a responsible FCA helps the KW experience more healthy flow by avoiding distractions, interruptions, and fatigue. We explain how this research on responsible AI fosters reliability, fairness, explainability, and privacy in the design of FCA.

In the final section, we summarize the findings from the study, which highlight the role of designing error-free tools that safeguard and genuinely enhance the KW's cognitive and emotional well-being. We conclude with insights about future work to use simulations and synthetic data to overcome prototyping challenges and refine the pool of nudges.

2 Who are the Knowledge Workers?

Knowledge Workers (KWs) who are potential FCA operators include researchers, engineers, architects, accountants, writers, and artists. They perform complex tasks requiring considerable concentration and creativity. KWs are highly mobile individuals that may work in one or more enterprises, which can be government, commercial, or non-profit in nature [19]. KWs must create, distribute, and apply knowledge in different contexts under conditions with varying workloads and interruptions. The KW can be a novice or an expert who may be succeeding on tasks, making errors, overloaded, or distracted.

2.1 The Knowledge Work Environment

While the work environments of KWs vary widely across different industries, we focused on the digital office workspace where the KW operates a computer system on a desk to complete a range of work activities. The computer and the desk are primary artifacts of the knowledge work environment (KWE). Secondary artifacts in the KWE such as lamps, toys, books, posters, and windows may be used for switching the focus, taking a break, or sparking creativity. The KWE may include supervisors and teammates interacting with the KW to perform work tasks. Situations in the KWE may be normal, abnormal, or emergency scenarios that determine the priority and relevance of interactions during operation time. Interruptions, context-switching, and high workload conditions cause situation complexity in the KWE. One predominant challenge with cluttered, information-rich, and dynamic KWEs is the likelihood of the KW becoming distracted and interrupted to the detriment of work completion.

2.2 What is Knowledge Work?

Knowledge work involves interactions with a variety of tasks with different requirements and demands. Knowledge work tasks vary from writing documents and computing calculations to discovering novel patterns and testing unknown concepts. Characteristic features of knowledge work include the challenge, goals, feedback, progress, interest, demand, success, failure, time spent, and bodily needs during the tasks.

Deep knowledge work is the practice of KWs focusing primarily on a complex task over an uninterrupted period [21]. Mastering these complex tasks requires intense focus indicative of the flow state. FCA trains KWs to do more deep knowledge work by focusing on human performance and well-being improvement. Shallow knowledge work is relatively low importance and low priority tasks completed in small pieces without demanding full attention [21]. It is because of the easy and low demand preference of the KW that shallow work consumes much of most KWs' working hours [34]. FCA was designed to draw attention to and reverse this KW behavior.

3 Flow Experience

Flow is a subjective sense of high control, concentration, and absorption in a task [32]. In the workplace, personal flow occurs when individuals, acting solo or in teams, operate with optimal focus and skill without apparent effort or self-consciousness, which yields a heightened sense of satisfaction, intrinsic motivation, and peak performance [8, 9, 20].

KWs are prone to distractions [25, 16, 23] and interruptions [17, 1] during knowledge work. KWs are also prone to cognitive fatigue, which leads to exhaustion, increased disinterest in a job, and decreased feelings of personal efficacy related to work [3, 5, 6]. Task demands are likely to significantly influence the imposition of fatigue, assuming that the KW has limited cognitive resources. This research focuses on the experience of personal flow and not team flow [14] or collective flow among workgroups [26]. The operator may customize FCA's nudges to facilitate meaningful and powerful cues that drive personal flow.

3.1 Neurofeedback for Flow Experience

Figure 2 illustrates how FCA continuously analyzes the human-task interaction within the given environmental conditions. FCA monitors cognitive workload and affective state transitions as electroencephalography (EEG) signals, task state, and task duration vectors. FCA computes a dashboard with indicators about how the individual performs at work and summarizes them in a work-day visualization of metrics with a proactive assessment of well-being, engagement, and flourishing.

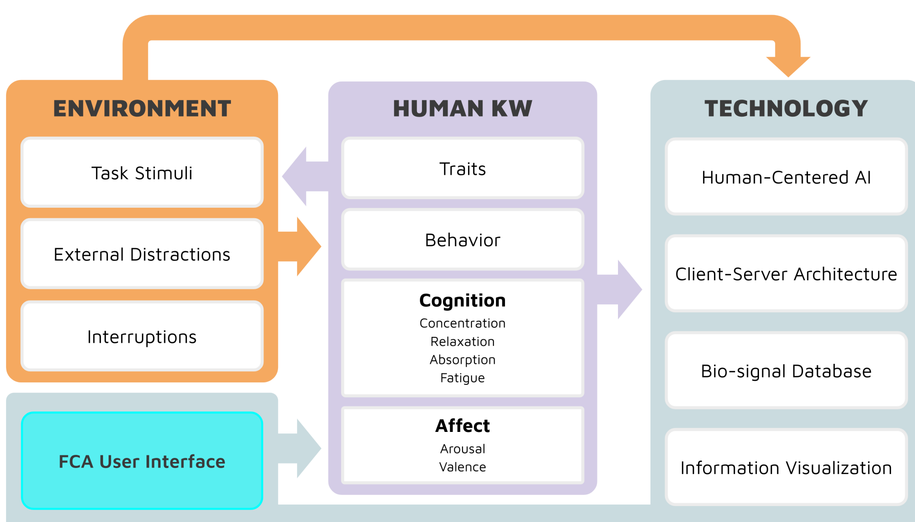

Fig. 2. High-level Framework of FCA

4 Choice Architecture of Nudges

A significant contribution of Thaler and Sunstein's Nudge Theory [29] is the generalization that "nudges" are a viable approach to promote behavior change. Schneider et al. [27] demonstrated that nudging could be performed by employing user interface (UI) design elements to guide people's behavior in digital choice environments. FCA applies several types of nudges, e.g., decision assistance nudges using defaults and decision structure nudges using convenience.

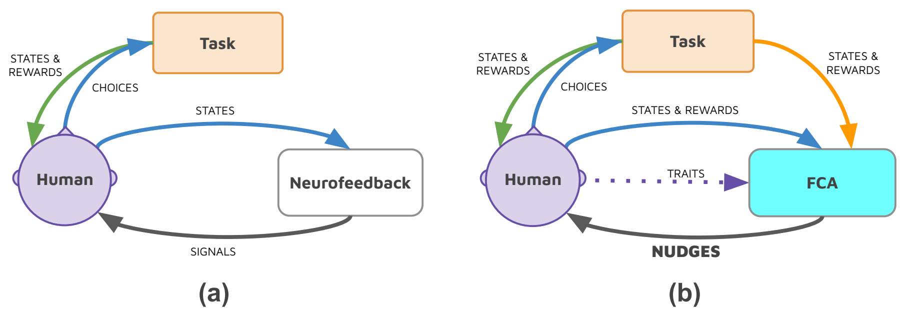

Fig. 3. Traditional Neurofeedback vs FCA [10]

Figure 3 (a) shows how standard neurofeedback is used by an operator while performing a task. The operator makes choices based on task demands and then update their mental state based on the new task state and the rewards received. The operator must then integrate the neurofeedback signals with task feedback to improve task performance, which could be a divided attention task [10].

Figure 3 (b) shows how FCA is used by an operator while performing a task. FCA uses states and rewards from the operator and the task to determine if a nudge is warranted. This information, including if there were previous nudges, is contextualized to determine which nudges are most likely to influence the operator to enter flow. Nudges with little effect on the operator's state will be selected less often than those with a rapid and positive effect [10].

FCA employs a gamified, multi-modal interface to present a hierarchical choice architecture of flow and cognitive well-being nudges that KWs may personalize. Nudges are unique distractions presented as contextual recommendations to guide the operator toward healthy flow. Nudges are external stimuli that consume some level of attention and cognitive resources. FCA learns which nudges, if any, are effective for individual KWs given a specific context.

Some FCA nudges will have uncertain effects. Nudges will also be effective less than 100% of the time, meaning that even if FCA recommends the "correct" nudge, the user may not give the desired response [10]. By retrieving cases similar to the current bio-signals, task, and context, the nudge suggestions made in FCA can examine alternatives and reason effectively through an exploration-exploitation trade-off [33]. FCA nudges differ based on their UI modality or presentation method, e.g., speech, text, music, ambient sounds, and rituals. If the KW appropriately receives the nudges, they may help them eliminate distractions and maintain focus so that KWs can spend more time in deep work such as learning new skills and mastering complex tasks [21].

4.1 FCA as a Recommendation System

A recommendation system is a software tool that recommends suitable items to a user or group of users [15]. Many modern digital platforms embody recommendation engines as a way of personalizing their services for users [24]. Traditionally, recommendation systems recommend to individual users the most relevant items based on the representation of item features. These traditional content-based and user-centric recommendation systems do not adapt to contextual information, such as time, place, and the presence of other people [2]. FCA recommends nudges to guide the operator to the flow state based on contextualized bio-signal data and trait information. The following research questions arise: what are the potential failures of FCA as a recommendation system, and how can these failures be simulated to detect and mitigate them?

4.2 FCA as a Conversational AI

FCA uses conversational AI to interact with KWs to perform command-type tasks and speech-to-text operations. The communication modalities of FCA nudges are multifarious and multi-modal. FCA uses metaphors such as the state metaphor that encodes the operator's attention and flow states and the background color metaphor that encodes the operator's flow performance records, such as the shortest flow onset time and the longest flow dwell time. FCA uses a responsive conversational agent for meaningful operator state storytelling, work session briefings, debriefings, summaries, and control interactions. The operator talks to FCA, and it responds in words and performs corresponding actions. The following research questions arise: what are the potential failures of FCA as a conversational AI, and how can these failures be simulated to detect and mitigate them?

5 Methodology

After an initial focus on user scenarios that are traceable to critical user problems, we focused on planning for failures and considering errors. We applied the responsible AI (RAI) tools from Microsoft's Human-AI eXperience (HAX) Toolkit [18] to evaluate a set of FCA's failure scenarios. The HAX Toolkit provides interaction guidelines, design patterns, workbooks, and a playbook for generating and testing human-centered AI experiences [13]. We selected the HAX Playbook to systematically explore common human-AI interaction failures, define user scenarios that cause potential shortcomings, and recommend simulations of system behaviors for early user testing. This scenario-based design (SBD) approach facilitated the HCAI design of FCA in ways that mitigate failures and provide affordances that help KWs avoid and recover from errors.

The HAX Playbook began with classifying FCA by its primary functions. The categories of AI provided in the HAX Playbook were search AI, recommendation system, conversational AI, text prediction and assistance, and classification. FCA was evaluated as a hybrid interaction system with a direct human-facing recommendation system and conversational AI functions.

As defined in the HAX Playbook, a recommendation system makes preferred-content suggestions, sometimes by predicting the user's rating of the content. The recommendation function predicted which nudge would effectively move the operator from the estimated state to the desired state. The effects of the nudges were monitored and used for training the recommendation model.

FCA featured conversational interactions with the operator through exchanges of natural-language dialogue, similar to conversing with a person. Application of the HAX Playbook focused on interactions that simulated FCA interacting with an operator as a workplace coach. FCA spoke particular nudges and anticipated a natural language response from the operator.

We sought to design a responsible HCAI that effectively understands the bio-signals of KWs and interacts with them to enhance their performance. The desired outcome of implementing the HAX Playbook was to identify and mitigate failures in the human-facing FCA. Table 1 and Table 2 highlight the taxonomies of human-AI failure scenarios relevant to FCA as a recommendation system and conversational AI, respectively.

Table 1: Taxonomy of FCA Recommendation System Failure Scenarios

Failure Source	**Failure Scenario**
Trigger Errors	
Missed trigger	[FCA] fails to detect a valid triggering event and misses the opportunity to nudge.
Spurious trigger	[FCA] triggers in the absence of a valid triggering event (it triggers when not intended).
Delayed trigger	[FCA] detects a valid triggering event but nudges too late to be useful.
Input Errors	
Spurious events	KWs may trigger accidental actions which they may try to undo leading to spurious events that can confuse [FCA].
Delimiter Errors	
Truncation	[FCA] begins capturing input too late, or stops capturing input too early, and thus acts only on partial input.
Overcapture	[FCA] begins capturing input too early, or stops capturing input too late, and thus acts on spurious data
Response Generation Errors	
Ambiguities	[FCA] chooses an ambiguous response or wrong interpretation for a given scenario.
Wrong item	[FCA] may return the wrong nudge from the choice architecture.
Poor precision	[FCA] returns a result list that includes many non-relevant nudges.
Poor recall	[FCA] returns a result list that excludes relevant nudges.
Poor ranking	[FCA] returns an order of nudges in the results list that does not match an intended natural order.
Low result diversity	[FCA] returns all the nudges in the list that are similar to one another.

Table 2: Taxonomy of FCA Conversational AI Failure Scenarios

Failure Source	**Failure Scenario**
Trigger Errors	
Missed trigger	[FCA] fails to detect a valid triggering event and misses the opportunity to nudge.
Spurious trigger	[FCA] triggers in the absence of a valid triggering event (it triggers when not intended).
Delayed trigger	[FCA] detects a valid triggering event but nudges too late to be useful.
Input Errors	
Transcription	Transcription errors are common in systems that rely on speech recognition.
Noisy channel	KW input is corrupted by background noise, including by capturing other sounds in the background.
Delimiter Errors	
Truncation	[FCA] begins capturing input too late, or stops capturing input too early, and thus nudges on partial input.
Overcapture	[FCA] begins capturing input too early, or stops capturing input too late, and thus nudges on spurious data
Response Generation Errors	
Ambiguities	[FCA] chooses an ambiguous nudge or wrong interpretation for a given scenario.
No understanding	[FCA] fails to map the user's input to any known nudge and thus takes no action.
Misunderstanding	[FCA] maps the user's input to the wrong nudge.
Partial understanding	Although [FCA] has the correct interpretation of intent, it could fail to effect the correct nudge.

6 Results

Table 3 outlines recommendations for FCA as a recommendation system. There were twelve failures generated.

Table 3: Recommendations for Recommendation System Failures

Failure Source	**Recommendation**
Trigger Errors	
1. Missed trigger	Simulate this error by intentionally ignoring a triggering event and continuing to process input as if no trigger had occurred. Consider simulating this error at different rates to understand how the triggering false-negative rate impacts the interaction.
2. Spurious trigger	Simulate this error by triggering the system unexpectedly. Consider simulating this error at different rates to understand how the triggering false-positive rate impacts the interaction.
3. Delayed trigger	Simulate this error by artificially inserting a short delay between the user's input and the system's output. Experiment with the different delay lengths to understand the abovementioned trade-off.
Input Errors	
4. Spurious events	Simulate accidental clicks on nearby, or neighboring buttons or links, or buttons that have similar iconography to the intended buttons.
Delimiter Errors	
5. Truncation	Simulate this error by intentionally leaving out the first or last words of user input.
6. Overcapture	Simulate this error by intentionally including extra words at the start or end of input.
Response Generation Errors	
7. Ambiguities	Simulate such errors by intentionally leaving inputs ambiguous and forcing the system to choose the wrong interpretation for a given scenario.
8. Wrong item	Simulate this error by randomly selecting an item from the choice architecture and returning it instead of the intended item.
9. Poor precision	Simulate this error by randomly selecting items from the choice architecture and adding them to the results list.
10. Poor recall	Simulate this low recall by intentionally leaving out key results, perhaps going so far as to prevent the user from completing their task.
11. Poor ranking	Simulate this error by shuffling or reversing the order of the ranked list.
12. Low result diversity	Simulate this situation by adding near-duplicate items to the set of items being ranked, and then including a scenario where these items all appear in the ranked results.

Table 4 outlines recommendations for FCA as a Conversational AI. There were eleven failures generated.

Table 4: Recommendations for Conversational AI Failures

Failure Source	Recommendation
Trigger Errors	
1. Missed trigger	Simulate this error by intentionally ignoring a triggering event and continuing to process input as if no trigger had occurred. Consider simulating this error at different rates to understand how the triggering false-negative rate impacts the interaction.
2. Spurious trigger	Simulate this error by triggering the system unexpectedly. Consider simulating this error at different rates to understand how the triggering false-positive rate impacts the interaction.
3. Delayed trigger	Simulate this error by artificially inserting a short delay between the user's input and the system's output. Experiment with the different delay lengths to understand the abovementioned trade-off.
Input Errors	
4. Transcription	Simulate transcription errors by using an automated speech-to-text transcriber to convert the user's utterance to text or using any of the four techniques (i.e., truncation, substitution, insertion, or extension) to manipulate the user's utterance.
5. Noisy channel	Simulate this error by including unrelated text in the transcription, or by removing portions of correctly transcribed text.
Delimiter Errors	
6. Truncation	Simulate this error by intentionally leaving out the first or last words of user input.
7. Overcapture	Simulate this error by intentionally including extra words at the start or end of input.
Response Generation Errors	
8. Ambiguities	Simulate such errors by intentionally leaving inputs ambiguous and forcing the system to choose the wrong interpretation for a given scenario.
9. No understanding	Simulate this error by intentionally returning a non-answer response to a valid, well-formed input.
10. Misunderstanding	Simulate this error by intentionally processing a user's input with the wrong intent or action category.
11. Partial understanding	Simulate this error by replacing a default attribute originally assigned to the specified component.

7 Discussion

Microsoft's HAX Playbook offered a cost-effective method to identify and mitigate risks in human-AI interaction. The playbook aided in the design of FCA to minimize failures and improve the human-AI cognitive enhancement experience. The outcome of this RAI evaluation is a set of targeted failure scenario simulations to provide insights into the identified errors and undesirable behaviors. Applying RAI tools in the early design stage of FCA helped address potential AI flaws by using the criteria of reliability, fairness, explainability, and privacy.

7.1 Reliability

The desired outcome of implementing the HCD and RAI is to maximize the frequency of use and retention of FCA. To achieve this outcome, FCA has to perform with reliability and safety. For example, the choice architecture of nudges presented in FCA should be appropriate for KWs and designed to fit into the knowledge work environment safely. We recommend tuning the algorithms to minimize trigger errors and response generation errors. Offline model evaluations should be conducted periodically to monitor their performance on trending behavioral patterns and determine when online models need to be updated to meet the changes in individual KW choices and activities.

7.2 Fairness

FCA should treat all KWs fairly. For example, the personal profile customized by FCA operators should suggest the personalized nudges rather than propagating undesired biases about other KWs, groups, domains, and types of KWs. One key area of fairness is to substantiate the use of certain features for the profile forms and algorithms. For example, the misuse of demographics from the samples that algorithms use to make their decisions may skew towards specific trigger errors in recommendations and conversations. The notion of fairness is to consider and mitigate these algorithmic biases to the greatest extent possible.

7.3 Explainability

FCA should minimize value capture [22] in its nudges and well-being metrics by providing KWs with adequate information about why specific nudges were recommended and what benefits the metrics provide. For example, when FCA nudges a KW to relax, the UI should explain to the KW that the nudge was triggered because the KW appeared overwhelmed by the task, and how performing the nudge may resolve the issue. FCA should provide transparency by auditing its personalization, classification, contextualization, metrics, and nudge features. Periodic reviews with the operator should analyze how FCA computes states, contexts, and nudges. The design process should include evaluations to remove ambiguities and determine which bio-signals are effective and eliminate unnecessary features from computations.

7.4 Privacy

There are rules and guidelines for appropriately handling personal data linked or linkable to any individual, even if the individual is unknown. We prioritize taking steps to design suitable mitigations and controls to reduce privacy risks. For example, FCA addresses potential operator sensitivities by limiting the amount of personally identifiable information that it collects and allows operators to control their data expiration rates. FCA places a high value on privacy and autonomy by treating bio-signal data as protected health information and alerting operators to obtain their consent before streaming data into storage and algorithms. The bio-signal data and FCA processing are meant for the use of the individual operator only and should not be used in a supervisory fashion to enforce external work efficiency goals.

8 Conclusion

There are significant implications for KWs who work in demanding environments with unclear goals, unstructured tasks, and high workloads. They tend to suffer from work-related anxiety, apathy, and boredom. FCA demonstrates that these unfavorable states and conditions cost KWs performance and well-being. To reverse this behavior, FCA helps KWs measure their work and avoid situations that induce undesirable states.

Given that nudges will be effective less than 100% of the time, our findings in this study highlighted the vital need for explicitly testing the failure scenarios of FCA to deliver an error-free human-AI experience. It was challenging to prototype and test the personalized and contextualized experiences of FCA. The complexity of simulating the personalized and contextualized experiences arise from the need for generalizability. To overcome this challenge, we used synthetic and augmented data to prototype dynamic, personalized experiences early. This remedy sufficed for initial iterations until we refined the datasets with more representative priors from experiments with human KWs.

We propose to explore additional nudging strategies that may reduce the unwanted effects of anxiety, boredom, and apathy, such as providing KWs with well-timed breaks [30], granting autonomy over how to use the breaks [31], and reducing emotional demands on KWs [11]. Other nudging strategies may include high-quality humor, functional music, and adapting the task so that the task demand matches the skill of the KW. Future work will consider the mechanisms underlying different nudging techniques, such as reducing the undesirable state, improving cognitive well-being, and modifying other mediating factors.

References

1. Adler, R.F., Benbunan-Fich, R.: Self-interruptions in discretionary multitasking. Computers in Human Behavior **29**(4), 1441–1449 (2013)
2. Adomavicius, G., Tuzhilin, A.: Context-aware recommender systems. In: Recommender systems handbook, pp. 217–253. Springer (2011)
3. Alarcon, G.M.: A meta-analysis of burnout with job demands, resources, and attitudes. Journal of vocational behavior **79**(2), 549–562 (2011)
4. Attaran, M., Attaran, S., Kirkland, D.: The need for digital workplace: increasing workforce productivity in the information age. International Journal of Enterprise Information Systems (IJEIS) **15**(1), 1–23 (2019)
5. Bakker, A.B., Demerouti, E., Sanz-Vergel, A.I.: Burnout and work engagement: The jd–r approach. Annu. Rev. Organ. Psychol. Organ. Behav. **1**(1), 389–411 (2014)
6. Bakker, A.B., Emmerik, H.v., Euwema, M.C.: Crossover of burnout and engagement in work teams. Work and occupations **33**(4), 464–489 (2006)
7. Choudhury, P., Foroughi, C., Larson, B.: Work-from-anywhere: The productivity effects of geographic flexibility. Strategic Management Journal **42**(4), 655–683 (2021)
8. Csikszentmihalyi, M.: Flow: The psychology of optimal performance.(1990) (1990)
9. Csikszentmihalyi, M.: Play and intrinsic rewards. In: Flow and the foundations of positive psychology, pp. 135–153. Springer (2014)
10. Eskridge, T.C., Weekes, T.R.: Opportunities for case-based reasoning in personal flow and productivity management. In: International Conference on Case-Based Reasoning. pp. 349–354. Springer (2020)
11. Goldberg, L.S., Grandey, A.A.: Display rules versus display autonomy: emotion regulation, emotional exhaustion, and task performance in a call center simulation. Journal of occupational health psychology **12**(3), 301 (2007)
12. Hayes, J.: Workplace analytics: Surveillance or saviour? Engineering & Technology **12**(9), 46–48 (2017)
13. Hong, M.K., Fourney, A., DeBellis, D., Amershi, S.: Planning for natural language failures with the ai playbook. In: Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. pp. 1–11 (2021)
14. van den Hout, J.J., Davis, O.C., Weggeman, M.C.: The conceptualization of team flow. The Journal of psychology **152**(6), 388–423 (2018)
15. Lu, J., Wu, D., Mao, M., Wang, W., Zhang, G.: Recommender system application developments: a survey. Decision Support Systems **74**, 12–32 (2015)
16. Manly, T., Robertson, I.H., Galloway, M., Hawkins, K.: The absent mind:: further investigations of sustained attention to response. Neuropsychologia **37**(6), 661–670 (1999)
17. Mark, G., Gudith, D., Klocke, U.: The cost of interrupted work: more speed and stress. In: Proceedings of the SIGCHI conference on Human Factors in Computing Systems. pp. 107–110 (2008)
18. Microsoft: What is the hax toolkit? (Oct 2021), https://www.microsoft.com/en-us/haxtoolkit/, https://www.microsoft.com/en-us/haxtoolkit/, accessed on 2022-03-08
19. Murray, A.J., Greenes, K.A.: From the knowledge worker to the knowledge economy: six billion minds co-creating the future. Vine (2007)
20. Nakamura, J., Csikszentmihalyi, M.: The concept of flow. In: Flow and the foundations of positive psychology, pp. 239–263. Springer (2014)

21. Newport, C.: Deep work: Rules for focused success in a distracted world. Hachette UK (2016)
22. Nguyen, C.T.: How twitter gamifies communication. Applied epistemology pp. 410–436 (2021)
23. Pattyn, N., Neyt, X., Henderickx, D., Soetens, E.: Psychophysiological investigation of vigilance decrement: boredom or cognitive fatigue? Physiology & behavior **93**(1-2), 369–378 (2008)
24. Resnick, P., Varian, H.R.: Recommender systems. Communications of the ACM **40**(3), 56–58 (1997)
25. Robertson, I.H., Manly, T., Andrade, J., Baddeley, B.T., Yiend, J.: Oops!': performance correlates of everyday attentional failures in traumatic brain injured and normal subjects. Neuropsychologia **35**(6), 747–758 (1997)
26. Salanova, M., Rodríguez-Sánchez, A.M., Schaufeli, W.B., Cifre, E.: Flowing together: A longitudinal study of collective efficacy and collective flow among workgroups. The Journal of psychology **148**(4), 435–455 (2014)
27. Schneider, C., Weinmann, M., Vom Brocke, J.: Digital nudging: guiding online user choices through interface design. Communications of the ACM **61**(7), 67–73 (2018)
28. Sunstein, C.R.: Nudging: a very short guide. In: The Handbook of Privacy Studies, pp. 173–180. Amsterdam University Press (2018)
29. Thaler, R.H., Sunstein, C.R.: Nudge: Improving decisions about health, wealth, and happiness. Penguin (2009)
30. Trougakos, J.P., Hideg, I.: Momentary work recovery: The role of within-day work breaks. In: Current perspectives on job-stress recovery. Emerald Group Publishing Limited (2009)
31. Trougakos, J.P., Hideg, I., Cheng, B.H., Beal, D.J.: Lunch breaks unpacked: The role of autonomy as a moderator of recovery during lunch. Academy of Management Journal **57**(2), 405–421 (2014)
32. Ullén, F., de Manzano, Ö., Almeida, R., Magnusson, P.K., Pedersen, N.L., Nakamura, J., Csíkszentmihályi, M., Madison, G.: Proneness for psychological flow in everyday life: Associations with personality and intelligence. Personality and Individual Differences **52**(2), 167–172 (2012)
33. Weekes, T., Eskridge, T.C.: A neurofeedback-driven humanoid to support deep work. In: Proceedings of the 33rd Florida Conference on Recent Advances in Robotics. pp. 14–16 (2020)
34. Zaman, S., Wesley, A., Silva, D.R.D.C., Buddharaju, P., Akbar, F., Gao, G., Mark, G., Gutierrez-Osuna, R., Pavlidis, I.: Stress and productivity patterns of interrupted, synergistic, and antagonistic office activities. Scientific data **6**(1), 1–18 (2019)