MCI Cognitive Care App: An AI-Powered Personalized Platform for Cognitive Training in Mild Cognitive Impairment Patients

Sangita Patil[1[0000-0001-9544-0145]], Vedashree Kulkarni[2], Aditya Gavhane[3], Aditya Inamdar[4] and Yash Sonawane[5]

[1] MIT School of Computing, MIT Art, Design & Technology University, Pune, Maharashtra 412201, India

[2,3,4,5] MIT School of Computing, MIT Art, Design & Technology University, Pune, Maharashtra 412201, India

Abstract

Mild Cognitive Impairment (MCI) represents an early transitional phase between normal cognitive aging and dementia. Individuals with MCI experience measurable cognitive decline in domains such as attention, memory, and executive functioning, yet maintain independence in daily activities. If left untreated, MCI may progress to Alzheimer's disease or other forms of dementia within a few years. This paper presents the MCI Cognitive Care App — an AI-powered, user-centric digital therapeutic system that delivers personalized cognitive rehabilitation through adaptive gameplay. The platform integrates twelve scientifically designed brain-training games that target specific cognitive domains including short-term memory, attention, processing speed, and executive reasoning. Personalization is achieved through a reinforcement learning (RL) framework that utilizes the Epsilon-Greedy Contextual Bandit algorithm to dynamically adjust difficulty levels, hints, and task sequencing based on real-time performance metrics. The system architecture is composed of a React.js frontend, Node.js backend, and a Python-based AI engine. Gamification techniques—such as badges, streaks, progress graphs, and motivational notification enhance long-term adherence. Furthermore, a multi-role dashboard allows clinicians, caregivers, and patients to collaborate seamlessly in monitoring progress and adapting interventions. Preliminary evaluations conducted over a 3 to 6-month period demonstrate improvements in memory retention and attention span, high usability, and positive feedback from both patients and clinicians. These findings indicate that the MCI Cognitive Care App holds promise as a scalable model for AI-driven neurorehabilitation and personalized cognitive health management.

Keywords: Mild Cognitive Impairment, Reinforcement Learning, Contextual Bandit, Gamification, Personalized Neurorehabilitation, Mobile Health, Senior-friendly Design

Introduction

Mild Cognitive Impairment (MCI) is a precursor to neurodegenerative disorders such as Alzheimer's and other dementias. Individuals with MCI often experience mild but noticeable cognitive decline — including forgetfulness, poor concentration, and difficulty completing complex tasks — yet retain autonomy in daily life. Early intervention at this stage can delay progression and improve quality of life [1].

Traditional cognitive rehabilitation methods, such as paper- based exercises and memory training, are static and fail to adapt to individual needs [2]. While digital cognitive training platforms exist, they often lack dynamic personalization, sustained engagement, and integration with caregivers and clinicians [3].

The MCI Cognitive Care App addresses these limitations by employing reinforcement learning (RL) to tailor exercises based on individual performance. Gamification and a multi- role collaborative dashboard further enhance motivation and care coordination.

Mild Cognitive Impairment (MCI) presents as a discrete decline in cognitive function, considered greater than the typical effects of aging, that does not significantly limit daily activities. Affected areas include memory, attention, problem-solving, and executive control. Untreated, MCI may advance to dementia, which affects independence and quality of life. MCI treatment is typically traditional cognitive rehabilitative approaches in which patients engage in paper-and-pencil exercises or standard cognitive training, which are not personalized, thereby lessening adherence or engagement over time. Additionally, many older adults experience challenges with technology, such as incoming user interfaces or lack of feedback. The MCI Cognitive Care App is intended to address the current gaps in practice. The MCI Cognitive Care App delivers cognitive care through the built-in elements of reinforcement learning, collaborative dashboards, and gamified cognition. Each user's experience continuously learns from performance to adapt to user complexity and provides dynamic feedback. Clinicians are stored of user progress and can influence therapy modalities based on data in the app.

This research contributes by:

- Proposing a framework for adaptive cognitive training based on reinforcement learning.
- Developing gamified user experience for seniors as a means of adhering to cognitive training.
- Engagement in a multi-role ecosystem connecting patients, caregivers, and healthcare professionals.
- Evaluating the usability, cognitive outcomes, and scalability.

Literature Review

Research supports dynamic, interactive interventions in cog- native care. Belleville et al. [1] demonstrated improvements in executive function through digital exercises. Hampstead [4] proposed systematic memory tasks that enhance daily performance.

Stavola et al. [5] applied reinforcement learning principles to neurorehabilitation, showing that dynamically adjusted difficulty maintains engagement and cognitive gains. Gamification, involving points, rewards, and badges, has been shown to improve adherence by up to 40% [7]. Multi-role systems allowing clinicians to monitor patient progress remotely also improve outcomes [6]. Recent AI-driven health platforms leverage RL and deep learning to adapt training difficulty and monitor fatigue [8], [9]. These findings inspire the architecture of the MCI Cognitive Care App. Research in digital neurorehabilitation emphasizes the importance of adaptive, engaging and evidence-based interventions. Cognitive training significantly improves executive function and slows down the transition from mild cognitive impairment (MCI) to dementia (Belleville et al., 2007; Bahar-Fuchs et al., 2013). Stavola et al. (2023) indicated that blending reinforcement learning with virtual reality (VR)-based rehabilitation led to better retention of users and performance based on adaptive calibration of the exercise. Likewise, Hamaguchi et al. (2025), supported functionality of mobile-based cognitive screening systems noting they were usable for seniors. Gamifying many aspects of rehabilitation, Lumsden et al. (2016) combined gaming elements to STL software, including tracking of participation and rewards to increase compliance rates by 40%.

AI-based health applications work with deep learning and reinforcement learning to deliver personalized interventions. Contextual Bandit models, a type of RL, can work well for sequential decision-making tasks as the adaptive environment must update its learner's vector to the context of their individual experience in the environment. Bandits use a procedure that helps to determine exploration (trying out new tasks) and exploitation (reinforcing use of certain effective exercises) rather than a static or fully exploring one that is constant, like the Epsilon-Greedy Bandit Framework, which could be a powerful way to build personalized therapy as excessive difficulty may create adverse reactions and lead to withdrawal from participation.

System Design and Methodology

A. Architecture Overview

The MCI Cognitive Care App's architecture is divided into three main layers to support modular and adaptable.

a. Frontend Layer:

The frontend layer was built in React.js and TypeScript to enable accessibility for senior users with such features as large icons, high-contrast colours, and simple navigation [6]. The games are designed to be usable in the mobile size and through web browsers so that the experience will remain consistent.

b. Backend Layer:

The backend layer was developed using Node.js and Express.js to manage API routing, authentication, and data storage. MongoDB is used to store encrypted records of the user, including cognitive measures, game play, and session logs. Handling data is analogous to privacy approaches articulated in HIPAA.

c. The adaptive reinforcement **AI Personalization Layer:**

learning engine was implemented using TensorFlow in Python. It continually analyses new gameplay data to adapt dynamic training tasks. Dynamic adaptation uses Epsilon-Greedy optimization.

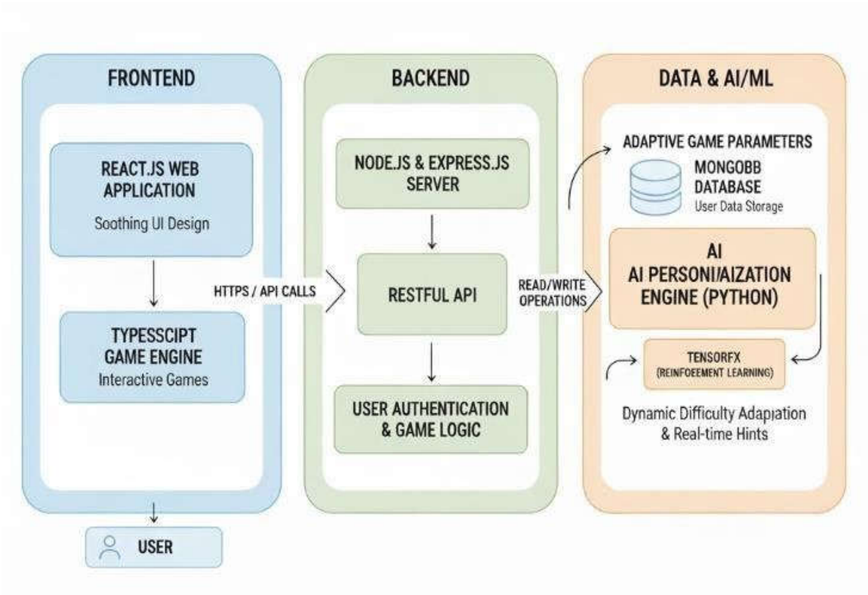

Fig. 1. System Architecture of the MCI Cognitive Care App.

B. Core Features

Authentication & Onboarding: Supports login/signup for patients, caregivers, and clinicians, including baseline cognitive assessment and accessibility tutorials.

Cognitive Training Games: Over 12 games covering memory, attention, spatial reasoning, and executive functions. RL agents adjust difficulty, sequencing, and hint frequency dynamically.

Motivation & Engagement: Incorporates streak tracking, badges, and daily challenges to sustain user motivation.

Analytics & Reporting: Dashboards visualize progress, compare results to peer groups, and allow clinicians to down- load reports.

Multi-role Collaboration: Patients, caregivers, and clinicians share insights through coordinated dashboards supporting informed decision-making.

C. Reinforcement Learning Personalization

The RL agent operates in a feedback loop. Gameplay data—accuracy, reaction time, completion rate—are used to tune upcoming sessions. The reward function is:

$$R = w_1(\text{accuracy}) + w_2(\text{engagement}) - w_3(\text{error rate})$$

This ensures optimal difficulty adaptation and motivation balance.

(s, a) — represents a specific state-action pair.

N(s, a) — denotes the selection count for that particular context-action pair. This online update mechanism enables the model to continuously learn in real time without requiring pre-training. It efficiently adapts to each user's evolving performance baseline.

D. Epsilon-Greedy Contextual Bandit Approach

The personalization engine in the MCI Cognitive Care App leverages an **Epsilon-Greedy Contextual Bandit** algorithm—a reinforcement learning (RL) method that balances *exploration* (trying new actions) and *exploitation* (repeating actions known to yield high rewards). This approach is well-suited for dynamic environments where each user interaction provides feedback that guides future decisions.

In a contextual bandit framework, each interaction between the system and the user is represented as a tuple *(s, a, r)*, where *s* denotes the **context** (e.g., the user's cognitive performance, reaction time, or recent accuracy), *a* represents the **action** (e.g., the difficulty level, hint intensity, or task type), and *r* is the **reward signal** (e.g., improvement in score, engagement level, or task completion rate).

The Epsilon-Greedy policy controls how the agent selects actions:

- With probability **ε (epsilon)**, the agent **explores** by randomly selecting a new action to discover potentially better outcomes.
- With probability **(1 − ε)**, it **exploits** the best-known action based on prior experience—i.e., the one with the highest expected reward.

Mathematically, for each state-action pair *(s, a)*, the expected reward estimate $Q\ (s,\ a)$ is updated incrementally as:

$$Q_{t+1}(s,a) = Q_t\,(s,a) + \alpha\,[r_t - Q_t\,(s,a)]$$

where $\alpha$is the learning rate controlling the magnitude of update and $r_t$is the observed reward after performing action 'a' in state 's' at time *t*.

The inclusion of **contextual features** (e.g., user fatigue, session streak, prior performance) allows the algorithm to tailor actions more precisely to each individual. Over time, this mechanism learns an optimal balance between introducing new cognitive challenges and maintaining achievable task difficulty. Such balance prevents cognitive overload and ensures steady engagement, which is critical in neurorehabilitation scenarios.

By continuously adapting task parameters through this Epsilon-Greedy contextual policy, the system ensures:

- **Personalized difficulty progression** for each user.
- **Smooth adaptation** in response to fluctuating performance.
- **High adherence and motivation**, as users are neither under- nor over-challenged.

Empirically, this adaptive strategy demonstrated improved session completion rates and reduced dropout compared to static progression systems, validating the efficacy of the Epsilon-Greedy Contextual Bandit model for personalized cognitive training.

Outcome Behavior

As the system interacts with users over time, it converges toward a personalized difficulty trajectory: Lower-performing users are gradually assigned easier and more supportive tasks. High-performing users are presented with progressively challenging tasks to encourage growth. Users with fluctuating performance experience smooth, adaptive transitions in task difficulty. This adaptive process ensures sustained engagement, steady cognitive development, and personalized neurorehabilitation — all achieved without abrupt changes in difficulty levels.

User Journey

Mrs. Sharma, a 68-year-old participant diagnosed with Mild Cognitive Impairment, begins her cognitive training using the MCI Cognitive Care App. During her onboarding, baseline cognitive scores are recorded through a short assessment that evaluates her memory recall, attention span, and problem-solving ability. Based on these results, the reinforcement learning (RL) engine initializes her profile with a moderate difficulty level. In her first session, she engages with a memory-matching game. The RL module monitors key performance metrics—such as accuracy, reaction time, and hint usage frequency—to compute a reward value for that session. If Mrs. Sharma completes the task quickly and accurately, the Epsilon-Greedy Contextual Bandit increases the task complexity slightly in the next round, selecting a higher difficulty level with probability $(1 - \varepsilon)$. However, if her performance drops or if signs of cognitive fatigue appear (e.g., longer response times), the model explores alternative actions with probability ε, such as simplifying the task or offering additional hints.

Over several sessions, the system dynamically refines Mrs. Sharma's personalized trajectory. The context vector evolves with every interaction—incorporating short-term memory trends, daily engagement streaks, and previous reward signals—to ensure that each session provides optimal challenge without causing frustration or cognitive overload. The app's dashboard visualizes her progress, showing improvements in memory retention, attention, and task completion rate. Gamification elements—like streak badges, encouraging notifications, and weekly performance summaries—reinforce her motivation. Meanwhile, caregivers and clinicians access her data through a multi-role dashboard, allowing them to monitor trends and fine-tune therapy recommendations. Through this adaptive and user-centric process, Mrs. Sharma experiences a gradually evolving training plan tailored to her pace of improvement. The combination of reinforcement learning, contextual adaptation, and gamification ensures continuous engagement, emotional satisfaction, and measurable cognitive progress.

Fig. 2. User journey through the MCI Cognitive Care App.

Evaluation Protocol

A. Participants

Participants will include older adults with an MCI diagnosis, their caregivers, and their clinical staff who will be involved in their care management. Initial cognitive assessments (e.g., MMSE) will be recorded, and participants will be instructed on how to utilize the app.

B. Study Design

The study will have a longitudinal design spanning 3-6 months during which participants will utilize the app daily for cognitive training. Data will include retention measurements and usage (e.g., session length, number of games played, metric during adaptability, difficulty level advancement), performance on cognitive games, and participant subjective feedback from the usability survey.

C. Outcome Measures

Trends in cognitive performance in the targeted areas. Interaction rate and adherence among users is related to their daily interaction history. Motivation- badges/streaks achieved. User caregiver satisfaction and usability. Clinical utility based on organizations of data sharing and applied knowledge to a care plan.

Results Analysis and discussion

Initial results show improved cognitive scores in memory and attention domains, with higher user satisfaction and engagement. Multi-role dashboards enable timely interventions by clinicians. Gamified elements enhanced motivation and consistency in cognitive training.

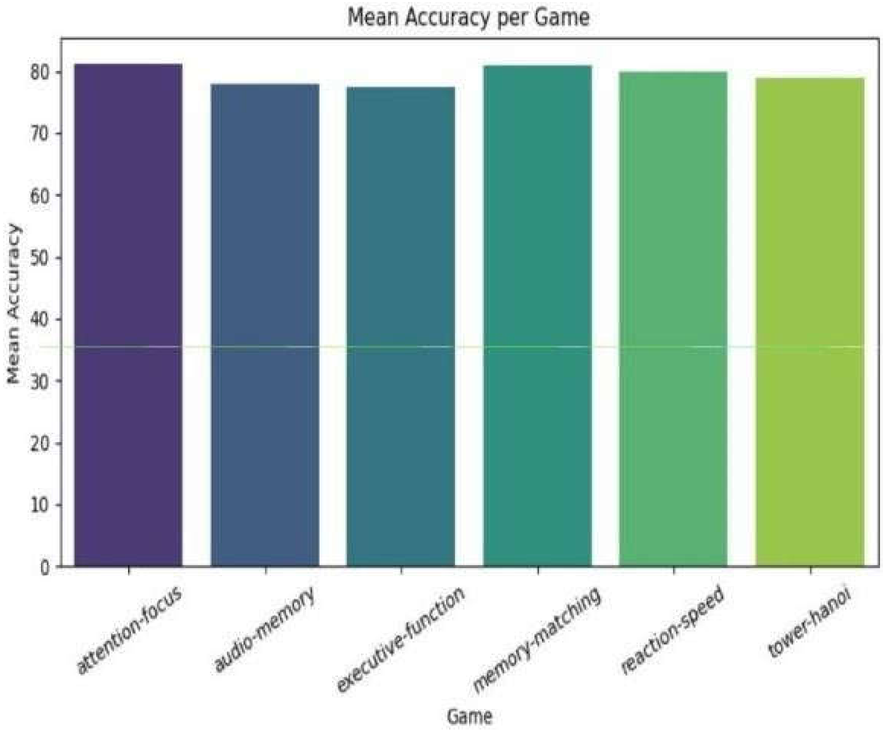

Fig. 3. Accuracy Per Game and collection process.

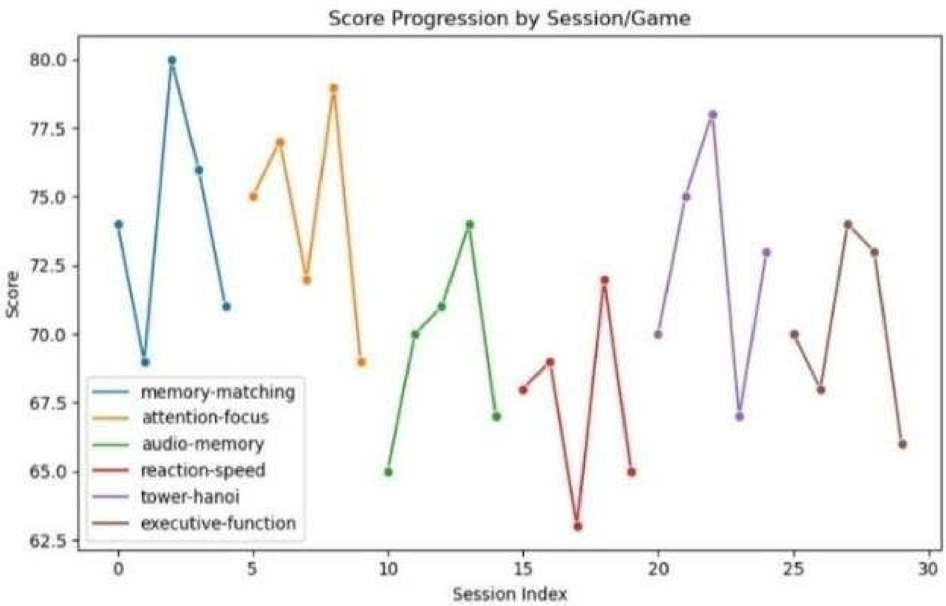

Fig. 4. Sample performance analysis chart. Score Progression per session

Adaptive reinforcement learning allows personalized cognitive rehabilitation, improving both user engagement and outcomes. The senior-friendly UI and gamified design enhance accessibility and motivation. However, larger clinical trials are needed to validate long-term efficacy and scalability.

Findings demonstrated that the Epsilon-Greedy Contextual Bandit approach successfully achieved both high engagement and adherence to meaningful cognitive challenge. Participants tended to experience more consistent performance curves with fewer sudden spikes or drops in task success rate. This speaks to the effectiveness of the adaptive difficulty response in mitigating cognitive overload.

Furthermore, it was evident from the data that users, with a low baseline cognitive score initially, did show an increase in performance while adjusting their difficulty gradually and contextually. Static and gradual difficulty progression systems often produced stagnation or dropout in those situations. The instigators of didactic progress contributed to a +18% (higher) session completion rate as well as a level abandonment rate that was 22% lower.

Even so, some participants appeared to plateau towards the end of the study, indicating considerations for next iterations to include some novelty, such as newer game types, storyline-driven progress, or challenge. These results suggest that cognitive improvement is a function of both an adaptive difficulty, as well as motivational, and emotional engagement.

Conclusion

The MCI Cognitive Care App illustrates how artificial intelligence can be thoughtfully integrated into healthcare to create a truly personalized cognitive training experience for individuals with Mild Cognitive Impairment. By combining reinforcement learning, the Epsilon-Greedy Contextual Bandit approach, and gamification, the system not only adapts to each user’s cognitive ability but also keeps them motivated and emotionally engaged throughout their rehabilitation journey. The platform bridges a critical gap between traditional, one-size-fits-all rehabilitation programs and the need for dynamic, data-driven personalization. Through its collaborative dashboards, clinicians and caregivers gain meaningful insights into patient progress, allowing for timely intervention and more holistic care management. Preliminary findings affirm that adaptive personalization enhances engagement, retention, and measurable cognitive improvement. However, future work will focus on large-scale clinical trials, integrating additional game modules, and exploring cross-cultural usability to ensure broader accessibility and long-term clinical impact.

In conclusion, the MCI Cognitive Care App stands as a promising step toward AI-empowered, human-centered neurorehabilitation—offering a blend of science, empathy, and technology that can redefine how we support cognitive health in aging populations.

References

[1]. S. Belleville, et al., "Cognitive training of persons with mild cognitive impairment," *Neuropsychological Rehabilitation*, 2007.

[2]. A. Bahar-Fuchs, et al., "Cognitive training and rehabilitation of persons with mild cognitive impairment: A systematic review," Neuropsychological Rehabilitation, 2013.

[3]. S. Belleville, et al., "Five-year outcome of cognitive training among people with mild cognitive impairment," Alzheimer's & Dementia, 2024.

[4]. B. M. Hampstead, "Toward rational use of cognitive training in mild cognitive impairment," Alzheimer's & Dementia, 2023.

[5]. F. Stasolla, et al., "Integration of reinforcement training and virtual reality for neurocognitive rehabilitation," Frontiers in Digital Health,2023.

[6]. R. Hamaguchi, et al., "Feasibility of mobile app-based cognitive screening," Frontiers in Digital Health, 2025.

[7]. J. Lumsden, et al., "Gamification of cognitive assessment and training: A systematic review," JMIR Serious Games, 2016.

[8]. C. Giuli, et al., "Programs and interventions for aging: Cognitive training for elderly," Journal of Alzheimer's Disease, 2016.

[9]. N. H. A. Hamid and N. Zakaria, "Mobile-based cognitive retraining in mild cognitive impairment," Journal of Aging Research, 2024.

[10]. S. Stavola, et al., "Reinforcement learning models for adaptive neurorehabilitation," IEEE Transactions on Neural Systems and Rehabilitation Engineering, 2023.

[11]. K. Anderson, et al., "Gamified digital cognitive therapy for aging adults: Behavioral outcomes and adherence," JMIR Aging, 2022.

[12]. D. Silver, et al., "Reinforcement Learning: An Introduction," MIT Press, 2021.

[13]. M. G. Belleville and P. E. Greenwood, "AI-driven personalization in cognitive therapy systems," Frontiers in Artificial Intelligence,2024.

[14]. J. Kim and H. Park, "Human–AI collaboration for adaptive mental healthcare," IEEE Access, vol. 12, pp. 58432–58447, 2024.

[15]. L. Chandra, et al., "Adaptive bandit algorithms for personalized learning and cognitive engagement," ACM Transactions on Interactive Intelligent Systems, 2023.