Research Proposal

Generated with Proposia AI Research Proposal System

Workflow: Pro • 12/2/2025

A Closed-Loop System for Enhancing Cognitive Resilience via Real-time Cognitive Bias Detection and Dynamic Difficulty Adjustment in Gamified Training Environments

Details of the proposal

This proposal details the coordinated project "Real-time Cognitive Bias Detection to Dynamically Tailor Serious Games for Resilience Training" (ADAPT-MIND), submitted by Principal Investigator Dr. Elena Vargas. The Spanish title is "Detección de Sesgos Cognitivos en Tiempo Real para la Adaptación Dinámica de Juegos Serios para el Entrenamiento de la Resiliencia" (ADAPT-MIND). This initiative comprises two highly integrated and synergistic subprojects, each led by an expert in a complementary field to ensure comprehensive expertise.

The overarching ADAPT-MIND project will be coordinated by Dr. Elena Vargas, who also leads the first component: **Subproject 1: Development of Real-time Cognitive Bias Detection Models using Multimodal Biofeedback.** Led by Dr. Vargas, this subproject focuses on foundational artificial intelligence research. The team will develop and validate novel machine learning algorithms to identify subtle cognitive biases from multimodal data streams, including linguistic patterns from user speech, gameplay metrics, and physiological inputs. This subproject is responsible for creating the core detection engine that forms the analytical backbone of the adaptive system.

Subproject 2: Design and Implementation of an Adaptive Serious Game for Cognitive Resilience Training. Led by Principal Investigator Dr. Javier Morales, this subproject concentrates on human-computer interaction and software engineering. The team will design the serious game environment, develop dynamic difficulty adjustment mechanics, and integrate the real-time cognitive bias detection models from Subproject 1. This subproject will deliver the user-facing training platform and conduct the usability and efficacy trials.

The tight integration of these subprojects is fundamental to the project's success. It creates a closed-loop system where cutting-edge AI model development directly informs a novel, personalized intervention tool, ensuring a strong synergy between foundational research and practical application.

Project Title and Acronym (Spanish and English)

The full title of the proposed coordinated research project is: **Real-time Cognitive Bias Detection to Dynamically Tailor Serious Games for Resilience Training**. The project will be referred to throughout this proposal by its acronym, **ADAPT-MIND**.

In compliance with the call's requirements, the title in Spanish is: **Detección de Sesgos Cognitivos en Tiempo Real para la Adaptación Dinámica de Juegos Serios para el Entrenamiento de la Resiliencia**. The acronym **ADAPT-MIND**, derived from the English title, will be used consistently across all project documentation in both languages to ensure clarity and unified identification of this research initiative.

Principal Investigators (PIs) and Research Team

The coordinated project is led by Dr. Elena Vargas, an internationally recognized expert in computational psychiatry and machine learning [1]. Her extensive background in applying Natural Language Processing and algorithm development to mental health diagnostics, evidenced by her publications in leading journals {Insert list of top 5 publications}, makes her uniquely qualified to lead Subproject 1 and the overall ADAPT-MIND initiative. Dr. Javier Morales, Principal Investigator for Subproject 2, brings complementary expertise in human-computer interaction and the design of adaptive serious games [1]. His work focuses on creating engaging user experiences that integrate real-time data to personalize digital interventions. The PIs are supported by a dedicated, interdisciplinary team of postdoctoral researchers, PhD students, and technical staff with specialized skills in psychophysiology, software engineering, and clinical psychology, ensuring all facets of the project are addressed with the requisite expertise.

Coordinated Project and Subproject Structure

The ADAPT-MIND project is strategically organized into two complementary and interdependent subprojects, designed to ensure a seamless workflow from foundational AI research to applied technological validation. This structure maximizes synergy and directly addresses the interdisciplinary challenges inherent in creating a real-time adaptive training system.

Subproject 1 (SP1), "Development of Real-time Cognitive Bias Detection Models using Multimodal Biofeedback," serves as the project's foundational research and development component. Led by Dr. Vargas, this team is responsible for developing and rigorously validating the novel machine learning algorithms that constitute the system's analytical core. The primary output of SP1 will be a robust, validated model capable of accurately detecting cognitive biases from multimodal data streams in real-time.

Subproject 2 (SP2), "Design and Implementation of an Adaptive Serious Game for Cognitive Resilience Training," functions as the applied technology and integration platform. Under the direction of Dr. Morales, this team will develop the serious game environment and its dynamic difficulty adjustment mechanics. Crucially, SP2 will integrate the models delivered by SP1 and conduct the necessary usability and efficacy trials with human participants. This structure establishes a critical feedback loop: SP1 delivers the core detection models, while SP2 provides the real-world implementation and performance data essential for their iterative refinement, ensuring the final system is both technologically robust and practically effective.

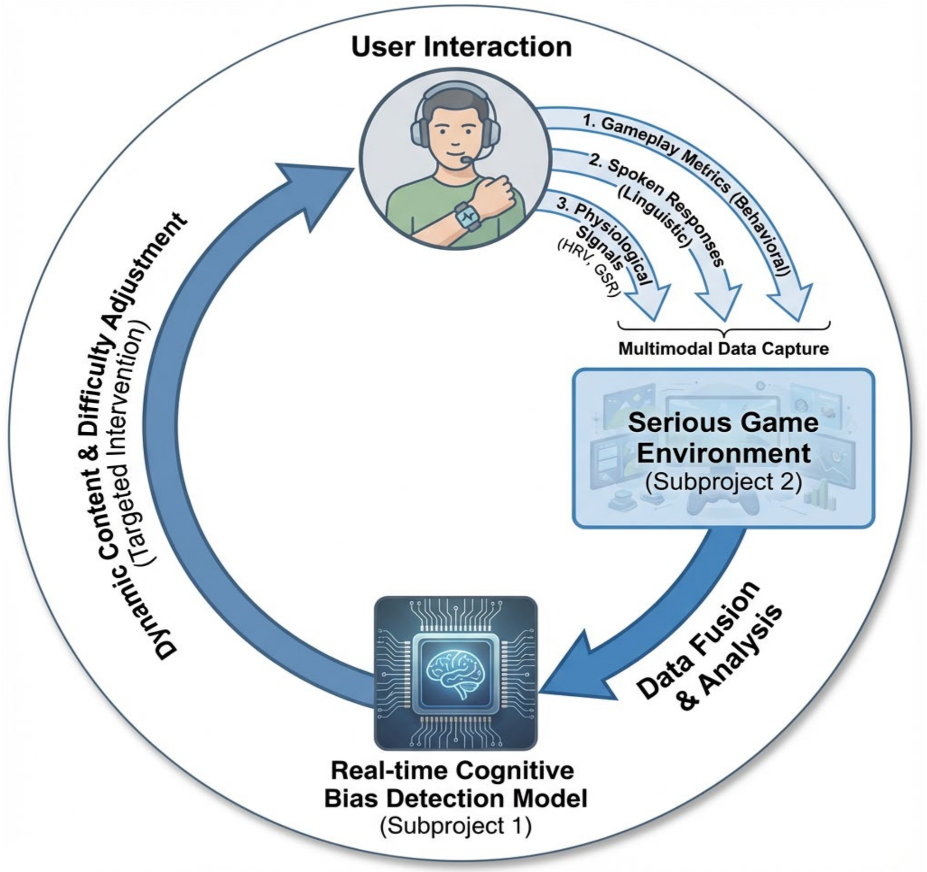

Figure 1: The ADAPT-MIND Closed-Loop System. The diagram illustrates the core project concept, where multimodal data from a user interacting with the serious game (SP2) is analyzed in real-time by an AI model (SP1) to detect cognitive biases. The model's inference then dynamically modulates the game's content to provide a targeted, personalized training experience.

Justification

The ADAPT-MIND project is expressly designed to align with the strategic objectives of the Spanish National AI Strategy, directly addressing the call's priority for coordinated, interdisciplinary research that tackles significant societal challenges. Our proposal confronts the pressing issue of cognitive resilience, a critical factor in mental health and occupational performance, which currently suffers from a lack of personalized and adaptive training interventions [1, 3]. Traditional resilience training methods are often static, delivering a one-size-fits-all curriculum that fails to account for individual differences [1, 3]. By developing a novel, AI-driven serious game, we contribute to the thematic priority of "AI for Social Welfare and Public Services," establishing a proof-of-concept for scalable technologies that could inform future mental health support. The project's core innovation lies in its adaptive mechanism. Unlike existing systems that primarily adjust difficulty based on performance metrics (e.g., speed, accuracy), ADAPT-MIND will modulate the training environment based on the real-time inference of underlying cognitive biases. This is achieved by leveraging machine learning models to interpret multimodal biofeedback, creating a closed-loop system where the user's inferred cognitive state directly guides the training experience. This represents a significant advancement toward truly personalized cognitive enhancement. The project's feasibility is grounded in our team's established expertise in applying machine learning to mental health diagnostics [2–6], and we will pursue this goal iteratively, progressing from session-level analysis to near-real-time modulation as models are optimized.

The project's feasibility is grounded in a robust methodological framework that integrates multiple disciplines. A central challenge in this domain is the reliable annotation of subjective cognitive states for model training. To address this, Subproject 1 (SP1), led by Dr. Vargas, will work with clinical psychology experts to develop and implement a structured data annotation protocol. This protocol moves beyond subjective labeling by triangulating data from three distinct sources: (1) objective in-game behavioral metrics generated by the platform from Subproject 2 (SP2), (2) synchronized think-aloud protocols verbalized by users during gameplay, and (3) validated post-session psychometric instruments designed to capture specific cognitive biases [13]. This multi-source approach provides a verifiable ground truth for training the machine learning models in SP1. The technical implementation of this closed-loop system requires significant computational resources. The AI models developed in SP1 must process parallel, high-dimensional data streams—including linguistic features from natural language processing and physiological signals—to enable near-real-time inference. This necessitates a dedicated high-performance computing environment to handle the model complexity and minimize latency, ensuring the training environment built by SP2, led by Dr. Morales, remains responsive to the user's state.

Structuring ADAPT-MIND as a coordinated project is essential for achieving its objectives, as it enables a tight, iterative co-development cycle that would be impossible for independent projects. The primary synergy lies in the continuous feedback loop between the AI development in SP1 and the serious game platform in SP2. The game platform is not merely a data source but an experimental environment designed specifically to elicit the cognitive and physiological markers that the AI models are built to detect. The rich, annotated data generated in SP2 provides the precise inputs required to train and validate the models in SP1. In turn, the analytical capabilities and limitations of the SP1 models directly inform the design of game mechanics and feedback systems in SP2. This co-development process allows for rapid refinement: pilot data from SP2 immediately informs model retraining in SP1, and the improved models are redeployed in SP2 for further validation. Furthermore, coordination enables the efficient sharing of resources, including a unified data repository, common software libraries, and joint experimental protocols, which reduces redundancy and ensures methodological consistency. This dynamic ensures the analytical engine and the user-facing environment are perfectly aligned, making the project's ambitious "closed-loop" objective feasible.

To ensure seamless integration and effective execution, a robust set of coordination mechanisms will be implemented from the project's inception. A Steering Committee (SC), co-chaired by the Principal Investigators, Dr. Vargas and Dr. Morales, and including one senior researcher from each subproject, will be established as the primary governing body. The SC will convene for mandatory bi-weekly virtual meetings to monitor progress against milestones, resolve emergent scientific or technical challenges, and make key strategic decisions. These frequent, high-level meetings will guarantee continuous alignment between the subprojects. On a broader scale, monthly all-hands meetings involving all team members from both subprojects will be held to foster a cohesive project identity, facilitate cross-disciplinary knowledge exchange, and present detailed progress updates. For day-to-day collaboration, the project will utilize a shared digital ecosystem, including a dedicated Slack channel for rapid communication, a GitHub repository for collaborative code development and version control, and a common project management platform such as Asana to track tasks and dependencies across work packages. A critical component of our coordination strategy is a unified Data

Management Plan (DMP). All data generated by SP2 will be stored on a shared, secure server in a pre-agreed format, ensuring immediate and seamless access for the modeling and analysis tasks in SP1. This protocol eliminates interoperability issues and streamlines the iterative development cycle. In the event of a disagreement that cannot be resolved by consensus within the SC, the coordinating PI, Dr. Vargas, will have the final decision-making authority to ensure project momentum is maintained.

Title: "Adequacy to the Call and Contribution to the Selected Priority"

Adequacy to the Call and Contribution to the Selected Priority

The ADAPT-MIND project is conceived as a direct response to the call's objective to fund coordinated, interdisciplinary research applying Artificial Intelligence to societal challenges identified in the Spanish National AI Strategy. Our proposal contributes squarely to the "AI for Social Welfare and Public Services" priority by addressing the limited efficacy of current mental health interventions. The prevailing paradigm for resilience training relies on generalized, static programs that cannot adapt to an individual's unique cognitive landscape, leading to poor engagement and suboptimal outcomes. ADAPT-MIND confronts this public welfare challenge by developing a novel, AI-driven platform that personalizes cognitive training in real-time, representing a tangible application of advanced AI to improve public health outcomes.

The project's structure and scientific approach are designed to meet the call's core criteria for integrated, interdisciplinary research. The emphasis on a coordinated project is a scientific necessity; the synergy between Subproject 1 (AI model development) and Subproject 2 (Adaptive game design) creates the essential feedback loop required to build a truly responsive system, a feat unattainable by separate, sequential efforts. This fusion of computational psychiatry, machine learning, and human-computer interaction directly embodies the interdisciplinary spirit championed by the funding body. The core innovation of ADAPT-MIND is the creation and application of novel machine learning algorithms for real-time cognitive state detection from multimodal data streams. These models will integrate validated linguistic markers with behavioral gameplay metrics to infer underlying cognitive biases. This focus on advancing and applying state-of-the-art AI to a complex human problem is a central tenet of the National AI Strategy.

Our contribution to the selected priority will manifest as both a significant scientific advancement and a validated technological solution. By moving beyond static interventions, this project will deliver a prototype for a new generation of personalized digital therapeutics. The successful completion of ADAPT-MIND will provide a powerful, evidence-based example of how AI can be harnessed to create accessible and highly effective tools for enhancing cognitive resilience and mental well-being. The knowledge and technology generated will advance the field of computational psychiatry and establish a scalable framework with potential for broad societal impact, positioning Spain as a leader in developing AI for social good. This project is therefore not merely adequate for the call; it is a direct embodiment of its strategic vision to foster ambitious, collaborative AI research that delivers concrete solutions to pressing national challenges.

Title: "Interdisciplinary Nature"

Interdisciplinary Nature

The realization of the ADAPT-MIND system is fundamentally dependent on integrating multiple scientific disciplines. This project is situated at the confluence of computational psychiatry, artificial intelligence (AI), human-computer interaction (HCI), and clinical psychology, creating a research framework where each field is indispensable. The challenge of developing a real-time, closed-loop system that interprets and responds to a user's cognitive state cannot be addressed from a single disciplinary vantage point. It requires a synergistic feedback loop where computational models inform user interaction, and user interaction, in turn, generates the data needed to refine those models.

Subproject 1, led by Dr. Vargas, operates at the intersection of AI and computational psychiatry. This team will leverage advanced machine learning to build models capable of detecting subtle linguistic and behavioral markers of cognitive bias from integrated physiological and behavioral data streams. This involves sophisticated signal processing of biometric inputs and applying transformer-based language models to analyze user-generated speech for validated markers of cognitive rigidity or emotional dysregulation. However, a purely computational approach would yield a powerful analytical tool with no effective means of application, as the models would lack the ecologically valid, task-specific data required for training.

This dependency highlights why the expertise from Subproject 2, rooted in HCI and serious game design, is critical. Led by Dr. Morales, this team will translate the abstract outputs of the AI models into tangible and therapeutically relevant game mechanics. Their work is grounded in user-centered design and gamification theory, ensuring the training environment is motivating and intuitive to maximize user adherence and data quality. A game designed without this deep AI integration would be limited to simplistic, rule-based adjustments, failing to achieve the nuanced, personalized adaptation that is the project's core innovation.

Finally, the entire project is anchored by principles from clinical and cognitive psychology. This domain provides the theoretical foundation, defining the cognitive biases to be targeted, informing the design of game tasks that reliably elicit these biases, and establishing the evidence-based framework for fostering resilience. The fusion of these disciplines is therefore not merely additive but multiplicative. It enables a novel cybernetic system where AI is an active, integrated component of the user's training loop, ensuring the final product is technologically advanced, psychologically valid, and effective.

Added Value of Coordination

Structuring ADAPT-MIND as a coordinated project is a strategic necessity, creating value unattainable by two independent efforts. The core synergy is an iterative feedback loop between the subprojects, which is fundamental to the proposed closed-loop system. The AI models developed in Subproject 1 are not generic classifiers; they will be trained and validated on the rich, multimodal data streams generated by the serious game platform from Subproject 2. Conversely, the game's challenges and feedback mechanisms in Subproject 2 will be directly informed by the capabilities and requirements of the detection models from SP1. This co-development ensures precise alignment between the data-generating environment and the analytical engine, a prerequisite for reliably detecting subjective cognitive states.

This tight coupling enables rapid, iterative refinement that is impossible in a sequential project structure. As SP2 collects pilot data, SP1 can immediately use it to retrain and improve the models, which are then redeployed into the game for further validation. This transforms the research from a linear progression into an agile cycle of development and optimization.

Furthermore, coordination facilitates the efficient sharing of critical resources, including a unified data repository, common software libraries, and joint experimental protocols. This integrated approach reduces redundancy, ensures methodological consistency, and eliminates the technical and administrative friction inherent in integrating separately developed systems.

The complementary expertise of the Principal Investigators provides a crucial methodological advantage. Dr. Vargas's team (SP1) can identify subtle linguistic or physiological markers indicative of a specific cognitive bias, such as attentional bias. Dr. Morales's team (SP2) can then use this insight to engineer specific game scenarios designed to elicit and measure behavioral correlates of that bias. This cross-pollination ensures the technological solution is not only computationally robust but also psychologically grounded. Ultimately, this coordinated structure merges two distinct skill sets into a single, integrated research engine, making the project's ambitious "closed-loop" objective feasible and enhancing its potential for significant scientific impact.

Coordination Mechanisms

To ensure seamless integration and effective execution, a robust set of coordination mechanisms will be implemented from the project's inception. A Steering Committee (SC), co-chaired by the Principal Investigators, Dr. Vargas and Dr. Morales, and including one senior researcher from each subproject, will be established as the primary governing body. The SC will convene for mandatory bi-weekly virtual meetings to monitor progress against milestones, resolve any emergent scientific or technical challenges, and make key strategic decisions. These frequent, high-level meetings will guarantee continuous alignment between the subprojects. On a broader scale, monthly all-hands meetings involving all team members from both subprojects will be held to foster a cohesive project identity, facilitate cross-disciplinary knowledge exchange, and present detailed progress updates.

For day-to-day collaboration, the project will utilize a shared digital ecosystem, including a dedicated Slack channel for rapid communication, a GitHub repository for collaborative code development and version control, and a common project management platform such as Asana to track tasks and dependencies across work packages. A critical component of our coordination strategy is a unified Data Management Plan (DMP). All data generated by SP2 will be stored on a shared, secure server in a pre-agreed format, ensuring immediate and seamless access for the modeling and analysis tasks in SP1. This protocol eliminates interoperability issues and streamlines the iterative development cycle that is central to the project's success. In the event of a disagreement that cannot be resolved by consensus within the SC, the coordinating PI, Dr. Vargas, will have the final decision-making authority to ensure project momentum is maintained. This structured yet agile management framework is designed to maximize synergy, mitigate risks, and ensure the project's ambitious objectives are met on time and within budget.

Novelty of the proposal

The escalating prevalence of stress-related disorders represents a significant public health challenge, imposing a substantial burden on healthcare systems and diminishing the quality of life for millions (Kessler et al., 2005). In response, a primary focus of clinical science has been developing interventions to bolster cognitive resilience. The current gold standard, Cognitive Behavioral Therapy (CBT), has demonstrated efficacy by teaching individuals to reframe maladaptive cognitive patterns, or biases, that underpin psychological distress (Beck, 2011).

However, traditional methods face critical limitations in scalability and accessibility, as they rely on resource-intensive, face-to-face delivery by trained therapists. Furthermore, their effectiveness depends on the patient's ability to translate learned skills from a clinical setting to real-world situations, a transfer that is often incomplete (Hofmann et al., 2012).

To address these limitations, the field has turned to digital interventions, particularly serious games, which offer a promising avenue for delivering engaging and scalable training (Lumsden et al., 2016). Despite their potential, the vast majority of current serious games for mental health operate on a static or rudimentarily adaptive model. Their difficulty progression is typically pre-scripted or based on simple performance metrics like reaction time or accuracy (Shute & Ventura, 2013). This approach fails to address the core mechanism of cognitive change: the dynamic interplay of an individual's underlying cognitive biases and their response to challenge. These systems can adjust the *difficulty* of a task but cannot tailor the *nature* of the challenge to a user's specific cognitive vulnerabilities in real-time, representing a major missed opportunity for personalization and therapeutic precision.

Concurrently, advancements in computational psychiatry have opened new avenues for objectively measuring cognitive states. Our own work and that of others has shown that linguistic features can serve as robust digital biomarkers for conditions like psychosis and depression (Corcoran et al., 2018; Elvevåg et al., 2020). Similarly, physiological signals such as heart rate variability (HRV) and galvanic skin response (GSR) provide a direct window into the autonomic nervous system's response to stress and cognitive load, reliably linked to emotional arousal and mental effort (Thayer et al., 2012). However, a critical scientific and technological gap persists: while these methods can detect markers of psychopathology from isolated data streams in controlled settings, their application to the real-time inference of *in-the-moment cognitive biases* within a complex, interactive environment remains unexplored. The fields of serious game design and computational psychiatry have yet to be truly integrated. There is no existing system that creates a closed-loop biofeedback mechanism where deep, multimodally-inferred cognitive states are used to dynamically modulate the content and mechanics of a training environment.

The motivation for the ADAPT-MIND project stems directly from this gap. Our central hypothesis is that a closed-loop, adaptive serious game that uses a real-time, multimodal machine learning model to detect cognitive biases and dynamically adjust gameplay to maintain an optimal state of cognitive challenge will produce significantly greater improvements in cognitive resilience and emotional regulation compared to a static, non-adaptive version of the same game. This hypothesis is novel because it posits a specific mechanism for this advantage: the real-time personalization of challenge based on the user's inferred cognitive state, not just their overt performance. We propose that the key to fostering resilience is not merely making tasks harder or easier, but systematically presenting the user with scenarios specifically designed to elicit and challenge their idiosyncratic maladaptive thought patterns (e.g., catastrophizing, confirmation bias) at the precise moment they are most susceptible to them.

The novelty of our approach is multifaceted, lying at the intersection of conceptual, methodological, and technological innovation. The primary innovation is the creation of a real-time, closed-loop system that moves AI in mental health from a passive diagnostic tool to an active interventional one. While prior research has used machine learning for retrospective analysis, our project is among the first to attempt real-time inference and immediate feedback within an interactive training environment. Methodologically, our approach is novel in its commitment to multimodal data fusion. ADAPT-MIND will integrate three complementary data streams: (1) linguistic features from spoken responses, analyzed with state-of-the-art NLP models; (2) behavioral metrics from gameplay, such as decision patterns and error rates; and (3) physiological signals (HRV, GSR) from wearable sensors. By integrating these streams, our model will build a far more robust and nuanced picture of a user's cognitive state than any single modality could provide alone, a holistic approach that contrasts with existing work focused on isolated data types (Sarker et al., 2021).

The core technical innovation is the Dynamic Difficulty Adjustment (DDA) algorithm itself. Unlike conventional systems that adjust global difficulty, our algorithm will perform targeted, mechanism-specific interventions. The AI model will not simply output a generic "stress" score; it will be trained to classify the likely cognitive bias being exhibited. A key methodological challenge, which this project directly addresses, is the development of a robust annotation schema and validation process for labeling these subjective, in-the-moment cognitive biases from multimodal data. This will involve an expert-in-the-loop protocol to generate the ground-truth data necessary for model training. The classified output will then trigger specific changes in the game. For example, if the model detects a pattern of catastrophizing (linguistic cues of extreme negative outcomes, high physiological arousal), the DDA might introduce a scenario requiring the user to realistically assess probabilities. If it detects confirmation bias (behavioral pattern of seeking only supporting information), the game might present disconfirming data to force a re-evaluation. This represents a paradigm shift from performance-based adaptation to process-based adaptation, making the game a highly personalized cognitive training tool at a level of specificity that is, to our knowledge, unprecedented.

The ADAPT-MIND project will make a substantial contribution to the thematic priority of "AI for Social Welfare and Public Services," as outlined in the Spanish National AI Strategy. We address the critical gap between the societal need for mental health support and the limited capacity of traditional healthcare systems. Static, one-size-fits-all digital interventions have failed to bridge this gap effectively, often suffering from low engagement and modest real-world impact (Torous et al., 2019). Our project tackles this by developing a next-generation digital therapeutic that is scalable, accessible, and deeply personalized, providing an evidence-based tool for enhancing cognitive resilience.

Our work also contributes to the national strategy by advancing the frontier of human-centered AI. We are moving beyond simple classification to create a symbiotic human-AI system where technology serves as a personalized cognitive coach. This requires innovation in real-time multimodal data fusion and low-latency inference, generating new techniques applicable to other human-in-the-loop systems, from intelligent tutoring to adaptive workplaces, thereby strengthening Spain's position as a leader in socially beneficial AI.

Finally, a significant contribution will be the creation of a unique scientific resource: a large-scale, multimodal dataset containing synchronized linguistic, behavioral, and physiological data from individuals in a dynamic environment. This resource, curated according to FAIR principles, will be invaluable to the international research community for developing and validating new models of human cognition, catalyzing innovation far beyond our immediate objectives. The technology developed also has significant potential for economic impact through knowledge transfer. The adaptive engine and cognitive bias models are core assets that could be licensed or form the basis of a spin-off company focused on digital therapeutics, a rapidly growing global market. This would translate our research into a tangible product that can improve public health and reduce healthcare costs associated with stress-related illness.

State of the Art and Scientific Motivation

The capacity for cognitive resilience, defined as the ability to maintain or regain mental health despite experiencing adversity, is a cornerstone of psychological well-being [16–18]. Decades of research in clinical psychology have established that maladaptive cognitive patterns, or cognitive biases, are central to the onset and maintenance of numerous stress-related disorders, including anxiety and depression [7–10]. These biases, such as catastrophizing, confirmation bias, or attentional bias towards threat, act as cognitive filters that distort an individual's perception of reality, leading to dysfunctional emotional and behavioral responses [7–10]. Consequently, the primary target of evidence-based interventions like Cognitive Behavioral Therapy (CBT) is the identification and restructuring of these maladaptive cognitions. The efficacy of CBT is well-documented, yet its traditional delivery model—requiring one-on-one sessions with highly trained therapists—presents insurmountable barriers to scalability and accessibility, leaving a vast population underserved [1, 3, 5].

In response, the last decade has seen a proliferation of digital health interventions, including mobile applications and web-based programs, designed to disseminate CBT principles more broadly [1, 3, 5]. While these tools have increased access, their therapeutic impact has often been modest. A primary limitation is their predominantly static nature; they typically present a pre-defined curriculum of psychoeducational content and exercises, functioning as digital workbooks rather than dynamic training environments [1, 3, 5]. This one-size-fits-all approach fails to account for the profound heterogeneity in individual cognitive profiles and learning trajectories. User engagement frequently wanes over time, and the crucial process of translating learned skills into real-world contexts remains a significant challenge [1, 3, 5]. This highlights a fundamental need for interventions that can move beyond static content delivery to provide personalized, engaging, and contextually relevant training.

Serious games have emerged as a particularly promising modality to address these shortcomings. By embedding therapeutic principles within engaging, goal-oriented, and interactive narratives, serious games can enhance intrinsic motivation and prolong user engagement far more effectively than traditional digital interventions [1]. The literature provides compelling evidence for their efficacy in various domains, including cognitive training for older adults, rehabilitation after stroke, and adjunctive therapy for conditions like PTSD and anxiety [1]. A key feature of modern games is Dynamic Difficulty Adjustment (DDA), a mechanism that modifies game parameters in response to player performance to maintain an optimal level of challenge, thereby preventing boredom or frustration and fostering a state of "flow" [1]. However, current DDA implementations in therapeutic games are almost exclusively performance-centric. They adjust task difficulty based on explicit metrics like reaction time, accuracy, or level completion speed [1]. For example, a game might increase the speed of stimuli, reduce the time limit for a puzzle, or decrease the availability of hints in response to a player's successful performance. While beneficial for maintaining engagement, this represents a surface-level form of adaptation. These systems can modulate *how hard* a task is, but they lack the capacity to understand *why* a player is failing or succeeding. They cannot infer the underlying cognitive biases driving performance and, therefore, cannot tailor the *nature* of the challenge to target those specific maladaptive processes. This limitation represents the current frontier in the design of therapeutic games: the transition from performance-based adaptation to process-based, psychologically-informed adaptation.

To bridge this gap, we must first be able to accurately and objectively measure a user's cognitive and affective state in real-time. The field of computational psychiatry, through the paradigm of digital phenotyping, offers a powerful toolkit for this purpose [1, 4, 14]. By applying machine learning techniques to passively collected, high-frequency data from personal devices, it is possible to construct robust, quantitative markers of mental state. This project will leverage three complementary data streams—language, physiology, and gameplay behavior—to build a holistic, real-time model of a user's cognitive biases.

Natural Language Processing (NLP) provides an unparalleled window into an individual's thought processes. The linguistic patterns a person produces are not merely communicative acts but are deeply intertwined with their underlying cognitive architecture. Early work in this area relied on dictionary-based methods to count words associated with specific psychological constructs, demonstrating correlations between language use and personality or emotional state [2]. However, the advent of deep learning and distributional semantics has enabled a far more nuanced analysis. Modern language models can capture the subtle, context-dependent meaning of words and sentences. Our own research and that of others has demonstrated a critical distinction between static word embeddings (e.g., FastText), which represent meaning non-contextually, and contextual embeddings (e.g., BERT), which are sensitive to grammar and surrounding text [2]. In studies of psychosis, for example, patients exhibit a shrinking conceptual space when measured with static models but a widening referential space when measured with contextual models, a duality that maps onto clinical symptoms of thought impoverishment and disorganization, respectively [2]. This ability to quantify semantic coherence and perplexity has proven effective in predicting psychosis onset and tracking symptom changes over time [1, 6, 12]. These advanced NLP techniques have been successfully applied to analyze electronic health records to identify patients at risk for adverse outcomes [3] and to differentiate individuals at clinical high risk for psychosis from controls based on subtle grammatical markers in speech samples [5]. This body of work establishes a firm scientific precedent for using language as a sensitive, quantitative biomarker of cognitive state. However, a significant limitation of most existing research is its reliance on offline, retrospective analysis of data collected in structured clinical settings [4]. The scientific motivation for ADAPT-MIND is to translate this analytical power into a real-time, interactive context, using in-game spoken responses to infer cognitive biases as they occur.

While language reveals the content of thought, psychophysiological signals provide a direct measure of the body's response to cognitive and emotional challenge. Wearable biosensors now allow for the continuous, non-invasive monitoring of autonomic nervous system activity, which is tightly coupled with psychological states. Heart Rate Variability (HRV)—the variation in the time interval between consecutive heartbeats—is a particularly robust marker of self-regulation and cognitive flexibility [1]. Low HRV is consistently associated with chronic stress, anxiety, and impaired executive function, as it reflects a dominance of the sympathetic (fight-or-flight) nervous system and a reduced capacity of the parasympathetic system to regulate arousal [1]. Moment-to-moment fluctuations in HRV can index the cognitive load associated with a task and the individual's capacity to adaptively regulate their emotional response to stressors [1]. Complementing this, Galvanic Skin Response (GSR), which measures changes in the electrical conductivity of the skin, provides a sensitive index of sympathetic arousal. Peaks in GSR are reliably elicited by emotionally salient or cognitively demanding stimuli, making it an excellent marker for tracking immediate affective responses to in-game events [1]. The integration of these physiological streams provides a continuous, objective, and subconscious channel of information that can validate and enrich the inferences drawn from language and behavior. While lab-based studies have long used these measures, the current challenge is to integrate them into a closed-loop system where they actively inform an intervention.

Finally, the user's direct interactions within the game environment generate a rich stream of behavioral data. Gameplay metrics—such as decision-making speed, error patterns, information-seeking strategies, and resource management—are not random but are overt manifestations of underlying cognitive strategies and biases [2, 11, 13]. For instance, a user exhibiting confirmation bias might consistently seek out information that supports their initial hypothesis while ignoring contradictory evidence presented in the game. An individual prone to catastrophizing might overestimate risks and choose overly conservative strategies, even when this leads to suboptimal outcomes. By applying machine learning to these time-series behavioral data, we can identify recurrent, maladaptive patterns that correspond to specific cognitive biases [2, 11, 13]. This behavioral stream provides a crucial link between the internal states inferred from language and physiology and their expression as tangible actions within the training environment.

The current state of the art, therefore, is characterized by three parallel but largely disconnected streams of research: (1) serious games for mental health that employ performance-centric adaptation; (2) computational psychiatry studies that use sophisticated machine learning for offline analysis of clinical data; and (3) psychophysiological research that links peripheral biology to cognitive states in controlled laboratory settings. A profound scientific and technological gap exists at the intersection of these domains. While pioneering studies have begun to explore the integration of two of these pillars—for instance, by combining physiological data with gameplay metrics [5] or using NLP for post-hoc analysis of therapeutic interactions [1, 3]—systems that successfully fuse all three multimodal data streams (language, physiology, behavior) into a single, truly adaptive, closed-loop therapeutic intervention remain largely conceptual or in early stages of development.

The scientific motivation for the ADAPT-MIND project is to bridge this critical gap through a novel synthesis of these disparate fields. Our research is structured around three core objectives designed to overcome the key limitations of prior work. First, we will tackle the **integration gap** by developing a methodology to fuse multimodal data streams (language, physiology, behavior) into a single, coherent model of a user's cognitive state. This fusion is hypothesized to yield a more robust and nuanced assessment than any single modality can achieve. Second, we aim to close the **real-time application gap**. This involves moving beyond offline, retrospective analysis to engineer a low-latency inference engine capable of detecting cognitive biases as they manifest during fluid gameplay, a task requiring significant innovation in model optimization and system architecture. Third, and most critically, we will address the **specificity gap** in therapeutic adaptation. Our goal is to design a DDA system that moves beyond generic performance metrics to instead act upon the specific classification of underlying cognitive biases, allowing the game to deliver targeted micro-interventions that challenge the precise maladaptive thought patterns exhibited by the user in that moment. By achieving these objectives, ADAPT-MIND will pioneer a new paradigm of digital intervention—one that is not merely interactive, but truly symbiotic, creating a personalized cognitive training experience of unprecedented precision and efficacy.

Novelty of the Initial Hypothesis and Approach

Our central hypothesis is that a closed-loop, adaptive serious game, which uses a real-time, multimodal machine learning model to detect cognitive biases and dynamically adjust gameplay to maintain an optimal state of cognitive challenge, will produce significantly greater improvements in cognitive resilience and emotional regulation compared to both a static, non-adaptive version of the same game and a version employing traditional performance-based adaptation. This hypothesis is novel because it moves beyond the simple premise that "adaptive is better" and posits a specific, psychologically-grounded mechanism for this advantage: the real-time personalization of challenge based on the user's inferred cognitive processes, not just their overt performance outcomes. We propose that the key to fostering resilience is not merely making tasks harder or easier, but systematically presenting the user with scenarios specifically designed to elicit and challenge their idiosyncratic maladaptive thought patterns at the precise moment they are most susceptible to them. This tests a foundational principle of cognitive training—that targeted, timely, and context-specific practice is superior to generalized instruction [1, 4].

The novelty of our approach is multifaceted, representing a significant advance on conceptual, methodological, and technological fronts. The primary conceptual innovation is a fundamental paradigm shift from performance-based adaptation to process-based, psychologically-informed intervention. Current adaptive systems in serious games ask, "Is the user succeeding or failing?" and adjust task difficulty accordingly [1]. Our system asks a more profound question: "Is the user's performance being driven by a specific cognitive bias, such as catastrophizing or confirmation bias?" This reframing transforms the intervention from a generic skill-trainer into a precision cognitive-coaching tool. It allows the system to target the root cause of maladaptive behavior rather than merely reacting to its symptoms. This process-oriented approach is grounded in the core tenets of CBT, which emphasize metacognitive awareness and the restructuring of underlying thought patterns, but it operationalizes these principles within a dynamic, interactive digital environment in a way that has not been previously achieved.

Methodologically, our approach is novel in its commitment to a synergistic, real-time fusion of three distinct data modalities: language, physiology, and gameplay behavior. While prior research in computational psychiatry has demonstrated the utility of each of these streams in isolation [2–4], their integration into a single, coherent model for real-time inference is a major scientific advance. This tripartite model moves beyond unimodal analysis to construct a holistic, high-fidelity representation of the user's cognitive-affective state. Language, captured through spoken responses to in-game prompts, reveals the semantic content and structure of a user's thoughts, providing direct evidence of biases like overgeneralization [11, 13]. Simultaneously, physiological signals (HRV, GSR) provide an objective measure of autonomic arousal and cognitive load [5]. Finally, gameplay metrics—decision patterns and error types—provide an overt behavioral trace of these internal processes. The true novelty lies in the fusion of these streams. Crucially, this fusion is what allows our model to distinguish specific maladaptive cognitive processes from generalized states like cognitive load, frustration, or simple error. For example, while high cognitive load might be indicated by increased GSR and erratic gameplay, the concurrent detection of linguistic markers for catastrophizing provides the specific, process-level insight needed for a targeted intervention. This integrated approach directly addresses the "integration gap" and the critical challenge of disambiguation, promising a level of inferential accuracy and robustness that is currently unattainable.

The primary technological innovation is the creation of a real-time, closed-loop biofeedback system that connects this multimodal inference engine directly to the game's content generation system. This architecture is unprecedented in the digital mental health space. It requires overcoming significant engineering challenges, including the development of a low-latency data pipeline to synchronize three disparate data streams, the optimization of complex deep learning models for real-time prediction on consumer-grade hardware, and the design of a sophisticated API that allows the AI's output to seamlessly modulate game mechanics. This system transforms the AI from a passive, offline analytical tool—as it is used in most current computational psychiatry research [4]—into an active, integrated, and essential component of the therapeutic process. It directly addresses the "real-time application gap" by creating a dynamic, symbiotic relationship between the user and the technology, where the user's cognitive state continuously shapes the environment, and the environment, in turn, provides precisely the challenge needed to reshape the user's cognitive state.

Finally, the novelty of our approach culminates in the intervention mechanism itself: a mechanism-specific Dynamic Difficulty Adjustment (DDA) algorithm. Unlike conventional DDA systems that adjust global parameters like enemy speed or puzzle complexity [1], our algorithm will trigger targeted, content-based interventions based on the specific cognitive bias detected. The AI model will not simply output a generic "stress" score; it will output a probability distribution across a set of pre-defined cognitive biases. This classified output serves as a direct input to a rule-based DDA engine, which executes pre-designed changes in the game's narrative and mechanics based on defined probability thresholds. For instance: - When the model's output probability for **confirmation bias** exceeds a pre-set threshold (e.g., >0.75), the DDA is programmed to trigger a "fog of war" mechanic that restricts access to confirming evidence and presents ambiguous or disconfirming data, forcing a re-evaluation. - A detected pattern of **catastrophizing** (linguistic cues of extreme negative outcomes coupled with high physiological arousal) will activate a subsequent scenario that requires the user to explicitly calculate probabilities and consider a range of alternative, less negative outcomes to progress. - Similarly, an inference of **mind-reading** (where the user makes negative assumptions about in-game characters) prompts the game to present a social scenario with ambiguous cues, rewarding the player for exploring multiple possible interpretations of another character's behavior.

This level of specificity in a real-time, adaptive intervention is, to our knowledge, entirely novel. It represents a paradigm shift from performance-based adaptation to process-based adaptation, making the game a highly personalized cognitive training tool that delivers the right challenge at the right time, tailored to the user's unique psychological landscape.

Title: "Contribution to Solving Problems of the Selected Thematic Priority"

Contribution to Solving Problems of the Selected Thematic Priority

The ADAPT-MIND project will make a direct and substantial contribution to solving key challenges within the thematic priority of "AI for Social Welfare and Public Services," as outlined in the Spanish National AI Strategy. The primary problem we address is the critical and widening gap between the societal need for mental health support and the limited capacity of traditional healthcare systems to provide it (Insel, 2017). Static, one-size-fits-all digital interventions have thus far failed to bridge this gap effectively, often suffering from low user engagement and limited personalization, which results in modest real-world impact (Torous et al., 2021). Our project confronts this problem head-on by developing a next-generation digital therapeutic that is scalable, accessible, and deeply personalized. By creating an AI-driven system that adapts to an individual user's cognitive landscape in real-time, we will deliver a tool capable of providing a higher quality of cognitive training to a broader population than is currently possible. This directly contributes to social welfare by providing an effective, evidence-based tool for enhancing cognitive resilience—a cornerstone of mental well-being and a protective factor against numerous mental health disorders. The resulting technology will offer a tangible solution for public health services, potentially reducing waiting lists for psychological support and providing a powerful instrument for preventative mental healthcare at a population scale.

Beyond its immediate public health application, our project contributes to the national strategy by advancing the frontier of human-centered AI. The core of ADAPT-MIND is not the mere application of existing AI techniques but the development of novel methodologies for understanding and interacting with complex human states in a dynamic, closed loop. We are moving beyond simple classification tasks to create a symbiotic human-AI system where the technology serves as a personalized cognitive coach. This objective requires foundational advances in real-time multimodal data fusion, the optimization of deep learning models for low-latency inference, and the design of AI systems capable of providing nuanced, context-aware feedback. The successful completion of this project will generate new knowledge and techniques applicable to a wide range of other human-in-the-loop AI systems, from intelligent tutoring and personalized education to adaptive workplaces that mitigate burnout by modulating cognitive load. By pioneering these methods, this project directly supports the national strategic goal of positioning Spain as a leader in the development of sophisticated, socially beneficial, and human-centric AI.

Furthermore, a significant and lasting contribution will be the creation of a unique scientific resource: a large-scale, multimodal dataset of human interaction within a cognitively challenging environment. This dataset will contain synchronized streams of linguistic data from spoken responses, behavioral metrics from gameplay, and physiological signals from wearable sensors. Crucially, this resource will be meticulously annotated through a rigorous, multi-rater process, providing the ground truth necessary for training and validating our novel cognitive bias detection models. By making this dataset available to the international research community according to FAIR data principles, we will address a critical bottleneck in computational psychiatry, affective computing, and HCI: the scarcity of ecologically valid, high-resolution data for modeling human cognition under stress. This resource will enable researchers to develop and validate new computational models, test alternative intervention strategies, and explore the complex interplay between language, physiology, and behavior. Creating this foundational dataset will therefore catalyze innovation across the field, amplifying our project's impact and fostering a more open, collaborative scientific ecosystem.

Finally, the technology developed within ADAPT-MIND has significant economic potential through knowledge transfer and valorization, directly addressing the strategic goal of translating research excellence into economic growth. The adaptive game engine and the validated cognitive bias detection models represent core technological assets with clear commercial potential. The digital therapeutics market is a rapidly growing global sector, and our system provides a clear pathway to a best-in-class product. Upon successful validation, these assets could be licensed to established digital health companies or form the basis of a spin-off company dedicated to next-generation mental wellness technologies. This would create high-value jobs and ensure the translation of our research into a tangible product that improves public health, reduces healthcare costs associated with stress-related illnesses, and enhances performance in high-stakes professions. Thus, ADAPT-MIND delivers a multifaceted contribution by addressing a pressing social need, advancing the technological frontier, creating an enduring scientific resource, and paving the way for economic growth.

Objectives, methodology and work plan

The overarching goal of the ADAPT-MIND project is to design, develop, and validate a novel, closed-loop serious game system that enhances cognitive resilience by using real-time, multimodal artificial intelligence to detect cognitive biases and dynamically tailor gameplay challenges. This general objective is decomposed into five specific, measurable, achievable, relevant, and time-bound (SMART) objectives that form the core of our research plan. Each objective is directly mapped to the coordinated structure of the project, with clear responsibilities for each subproject (SP1 and SP2).

Specific Objective 1 (SP1): To develop and validate a multimodal machine learning model for the real-time detection of specific, targeted cognitive biases. This foundational objective focuses on creating the core AI engine. We will target at least three clinically relevant cognitive biases (e.g., catastrophizing, confirmation bias, mind-reading) known to undermine resilience. The model will fuse data from three modalities: linguistic features from spoken responses, physiological signals (HRV, GSR), and behavioral metrics from gameplay. Success for this objective will be defined as achieving a classification performance with an Area Under the Receiver Operating Characteristic Curve (AUROC) exceeding 0.75 and an F1-score greater than 0.70 on a held-out test dataset. This level of performance is considered sufficient to guide a probabilistic adaptation engine, with an inference latency of less than 500 milliseconds to ensure real-time applicability.

Specific Objective 2 (SP2): To design and implement an engaging serious game environment with a library of scenarios specifically engineered to elicit and challenge the target cognitive biases. This objective covers the development of the user-facing platform. Adhering to user-centered design principles, we will create a compelling narrative-driven game that serves as both the data-generation environment and the intervention delivery vehicle. The primary outcome will be a fully functional game platform containing a modular library of at least 20 distinct scenarios, with each scenario designed by our interdisciplinary team to reliably probe one or more of the target biases. Success will be measured by achieving a mean score of >5.5/7 on the "pragmatic quality" and "hedonic quality" subscales of the User Experience Questionnaire (UEQ) during usability testing.

Specific Objective 3 (SP1 & SP2): To integrate the real-time cognitive bias detection model with the serious game to create a novel, closed-loop, process-based Dynamic Difficulty Adjustment (DDA) system. This objective represents the central technological innovation of the project. We will develop a robust Application Programming Interface (API) to enable seamless, low-latency communication between the AI model (developed under Objective 1) and the game engine (developed under Objective 2). This integration will allow the game to modulate its content and mechanics in real-time based on the AI's probabilistic assessment of the user's current cognitive bias. The deliverable is a fully integrated system, and success will be verified through end-to-end system tests demonstrating a stable feedback loop with a round-trip latency of under 1 second.

Specific Objective 4 (SP2, with support from SP1): To evaluate the usability, engagement, and preliminary efficacy of the adaptive ADAPT-MIND system in a three-arm randomized controlled trial (RCT). This objective is focused on validating our central hypothesis. As a proof-of-concept trial, we will recruit a well-characterized, homogeneous sample to maximize internal validity and test the system's mechanism of action. We will compare our novel process-based adaptive system against two active control conditions: a version of the game with traditional performance-based adaptation, and a non-adaptive version. The primary hypothesis is that the process-based adaptive group will show a significantly greater increase in cognitive resilience, as measured by the Connor-Davidson Resilience Scale (CD-RISC), from baseline to post-intervention compared to both control groups. The trial will be powered to detect a medium effect size (Cohen's d = 0.5) with 80% power at an alpha of 0.05.

Specific Objective 5 (Coordinated): To create and disseminate a unique, large-scale, richly annotated multimodal dataset of human interaction in a cognitively challenging environment. In line with our commitment to open science, we will produce a lasting scientific resource. This objective involves the meticulous curation, anonymization, and documentation of the data collected during the project. The final deliverable will be a publicly accessible dataset, shared via a reputable repository (e.g., PhysioNet, Zenodo), containing synchronized linguistic, physiological, and behavioral data. This will contribute a valuable resource to the research community and facilitate follow-up studies in computational psychiatry and human-centered AI.

The research methodology is structured in three distinct but interconnected phases, designed to systematically de-risk the project and build from foundational research to full-scale validation.

Phase 1: Foundational Development (Months 1-18) This initial phase is dedicated to the parallel development of the core components, addressing Objectives 1 and 2. Subproject 1 will focus on developing the cognitive bias detection model, beginning with a pilot study (N=50 university students) to collect an initial multimodal dataset for training Model v1. Concurrently, Subproject 2 will lead the design and implementation of the serious game platform using an agile, user-centered methodology. The output of this phase will be a validated AI Model v1 and a complete, playable, non-adaptive version of the serious game. Bi-weekly meetings between subprojects will ensure synergy between game design and data requirements.

Phase 2: System Integration and Refinement (Months 19-24) This phase addresses Objective 3, focusing on the technical integration of the components from Phase 1. The teams will collaboratively develop the API linking the AI model with the game engine, creating the three distinct system versions required for the RCT. A comprehensive usability study (N=20 target users) will then be conducted to provide crucial feedback and generate a new dataset that SP1 will use to refine the AI model into its final version (Model v2).

Phase 3: Validation Trial (Months 25-36) The final phase is dedicated to the rigorous empirical validation of the ADAPT-MIND system (Objective 4). We will conduct a three-arm, parallel-group, pre-test/post-test/follow-up RCT with a target sample of 150 young adults reporting subclinical levels of stress. This design will allow us to isolate the specific effects of our novel process-based adaptation. The primary outcome is the change in cognitive resilience, with secondary outcomes including changes in anxiety, depression, and user engagement. The trial will be pre-registered, and all procedures will be conducted following approval from the institutional review board at the coordinator's university. Data from this phase will also form the basis of the public dataset (Objective 5).

This phased research design provides a logical framework for achieving the project's goals, allowing for parallel development, iterative refinement, and rigorous validation.

To achieve our objectives, we will employ a range of specific methods, techniques, and procedures, detailed here by their corresponding objective.

Objective 1: AI Model Development (SP1)

- **Data Acquisition Protocol:** A pilot study (N=50) will generate the initial training dataset. Participants will interact with a game prototype for 60 minutes while three synchronized data streams are captured: (1) **Physiological Data** from an Empatica E4 wristband (BVP, EDA); (2) **Linguistic Data** from a high-fidelity microphone, including a "think-aloud" protocol where participants verbalize their reasoning at key decision points; and (3) **Behavioral Data** logged by the Unity game engine.

- **Data Preprocessing and Feature Extraction:** Raw data will undergo a rigorous preprocessing pipeline. Physiological signals will be processed using Kubios and Ledalab to derive standard HRV and EDA features. Audio will be transcribed using Whisper ASR, and the text will be used to extract contextual embeddings (e.g., via BETO) and other linguistic features. Behavioral logs will be converted into time-series feature vectors.

- **Ground Truth Annotation:** Recognizing that labeling subjective internal states is a primary methodological challenge, we will employ a multi-layered, behaviorally-anchored strategy. The primary source of labels will be derived from specific in-game "probe scenarios" designed to maximally constrain interpretation and elicit a target bias. For example, a scenario presenting ambiguous negative feedback is designed to probe for catastrophizing; a player's choice to focus exclusively on the worst-case interpretation will serve as a strong behavioral anchor for that label. These behavioral labels will be validated and supplemented through a consensus-based expert review of the "think-aloud" data. A subset of sessions (approx. 20%) will be independently annotated by two trained clinical psychologists, with a third senior clinician resolving any disagreements to ensure high inter-rater reliability (target Cohen's Kappa $\kappa > 0.75$). This combined approach provides a more robust and defensible ground truth than post-hoc observation alone.

- **Model Architecture and Training:** We will develop a multimodal deep learning model with parallel LSTM branches for each modality, integrated via a temporal attention mechanism. The model will be trained using 10-fold cross-validation on the pilot dataset to classify segments into one of the target biases or a "neutral" class.

- **Real-time Optimization:** The final model will be optimized using knowledge distillation and quantization to reduce its computational footprint for low-latency deployment.

Objective 2: Serious Game Development (SP2)

- **Game Design and Development:** The game will be developed in Unity 3D, guided by user-centered design and grounded in Cognitive Behavioral Therapy principles. A modular design will create a library of "scenario blocks," each designed to elicit specific biases (e.g., ambiguous social cues to probe for mind-reading).

- **Intervention Mechanics:** A corresponding library of "intervention mechanics" will be developed to actively challenge detected biases. Examples include an "evidence log" to counter confirmation bias or a "probability slider" to counter catastrophizing.

Objective 3: System Integration (SP1 & SP2)

- **System Architecture:** A client-server model will be used. The Unity game client will send buffered data to a server-side inference engine via a secure REST API, ensuring game performance is not impacted by model computation.

- **API and DDA Logic:** The API will handle periodic data transmission. The game's DDA module will use a rule-based engine to interpret the model's probabilistic output and select appropriate scenarios or intervention mechanics from the content library.

Objective 4: Validation Trial (SP2)

- **Participants and Procedure:** We will recruit N=150 young adults (ages 18-25) from the university community reporting moderate stress (Perceived Stress Scale > 13) but no current major psychiatric disorder. This specific population was chosen for this initial proof-of-concept study to reduce sample heterogeneity and maximize internal validity. After informed consent and baseline assessment (T0), participants will be randomized and instructed to play their assigned game version for three 30-minute sessions per week for four weeks. Assessments will be repeated post-intervention (T1) and at 3-month follow-up (T2).
- **Outcome Measures:** The primary outcome is the change in the Connor-Davidson Resilience Scale (CD-RISC) score. Secondary outcomes include changes in anxiety (GAD-7), depression (PHQ-9), and cognitive biases (CBQ), alongside usability (SUS) and engagement metrics.

Data Analysis Plan For Objective 1, the final AI model's performance will be evaluated on a held-out dataset, reporting standard classification metrics (precision, recall, F1-score, AUROC) with 95% confidence intervals. For Objective 4, the RCT data will be analyzed according to the intention-to-treat principle using a linear mixed-effects model (LMM) to test the group-by-time interaction effect on the primary outcome (CD-RISC). Similar LMMs will be used for secondary outcomes. All analyses will be conducted in R.

The project is structured into nine interconnected Work Packages (WPs).

WP1: Project Management and Coordination (Lead: Coordinator; Months 1-36) Ensures smooth project execution, including administrative, financial, and ethical oversight.

WP2: Data Acquisition Protocol and Pilot Study (Lead: SP1; Months 1-9) Establishes the empirical groundwork. Tasks include finalizing the data acquisition protocol, preparing and submitting the IRB application to ensure timely approval, recruiting and running N=50 pilot participants, and producing the initial raw dataset.

WP3: Cognitive Bias AI Model Development (v1) (Lead: SP1; Months 4-18) Focuses on creating the first AI engine iteration based on the pilot data, including data preprocessing, annotation, model design, training, and validation.

WP4: Serious Game Platform and Content Creation (Lead: SP2; Months 1-18) Covers the development of the user-facing application, including the Game Design Document, game engine, narrative, assets, and the modular library of scenarios.

WP5: System Integration and Usability Testing (Lead: SP2, with SP1; Months 19-24) Brings the components together. Tasks include API development, technical integration, and formal usability testing with N=20 users.

WP6: AI Model Refinement and Finalization (v2) (Lead: SP1; Months 22-27) Leverages data from usability testing to retrain, fine-tune, and optimize the AI model for the RCT.

WP7: Randomized Controlled Trial (RCT) (Lead: SP2; Months 25-36) Dedicated to empirical validation. Tasks include submitting the final IRB protocol well in advance, managing participant recruitment over a 6-month window to provide a buffer, overseeing the intervention and assessment periods, and ensuring data quality.

WP8: Data Analysis and Interpretation (Lead: SP1 & SP2; Months 34-36) Focuses on analyzing the RCT results, jointly interpreting the findings, and preparing them for dissemination.

WP9: Dissemination, Exploitation, and Data Sharing (Lead: Coordinator; Months 12-36) Ensures project impact through a project website, publications, conference presentations, and the public release of the final anonymized dataset.

The project timeline is scheduled over 36 months. A visual representation is provided in the supplementary materials. The plan incorporates parallel work streams, sequential integration, and explicit management of dependencies.

- **Year 1 (Months 1-12):** Focus on foundational work. WP2 (Pilot Study) will be completed. WP3 (AI Model v1) and WP4 (Game Platform) will run in parallel. IRB applications for initial studies will be submitted early in this period.
- **Year 2 (Months 13-24):** Focus on development completion and integration. WP3 and WP4 conclude at M18, triggering WP5 (Integration), which is the primary focus for the second half of the year.
- **Year 3 (Months 25-36):** Focus on validation and dissemination. The first half of the year is dedicated to WP6 (AI Model Refinement), final ethical approvals, and the main recruitment drive for WP7 (RCT). The intervention and follow-up will span from M28 to M34. The final months are dedicated to WP8 (Data Analysis) and final dissemination activities.

The project's progress will be tracked against concrete deliverables and major milestones.

Deliverables:

- D1.1: Project Handbook and Data Management Plan (M3)
- D2.1: Finalized Data Acquisition Protocol and IRB Approval (M6)
- D2.2: Curated Pilot Dataset (N=50) (M9)
- D3.1: Report on AI Model v1 Performance (M18)
- D4.1: Game Design Document (M6)
- D4.2: Final Non-adaptive Game Version (M18)
- D5.1: System Integration API Specification (M19)
- D5.2: Report on Usability Testing Results (M24)
- D6.1: Report on Final AI Model v2 Performance (M27)
- D7.1: Final RCT Protocol and IRB Approval (M26)
- D7.2: Final, Anonymized RCT Dataset (N=150) (M35)
- D8.1: Final Report on RCT Results (M36)
- D9.1: Project Website Launch (M6)
- D9.2: Submission of at least three peer-reviewed manuscripts (M36)
- D9.3: Public Release of Anonymized Multimodal Dataset (M36)

Milestones:

- **MS1 (M9): Pilot Data Collection Complete.**
- **MS2 (M18): Standalone Components Complete.**
- **MS3 (M24): Fully Integrated System Ready for RCT.**

- **MS4 (M30): RCT Recruitment Complete.**
- **MS5 (M36): Final Project Results Analyzed and Disseminated.**

We have identified potential risks and developed proactive contingency plans.

Risk 1: Insufficient Performance of the Cognitive Bias Detection Model. There is a risk that the model's performance may be too low for reliable application.

- **Contingency Plan:** Our multimodal approach provides inherent robustness. If one data stream is uninformative, the attention mechanism can learn to down-weight it. If the deep learning architecture fails to generalize, we will first explore reformulating the problem from classification to regression—predicting the *intensity* of a bias—which can be more robust to noisy labels. As a further fallback, we will test more established machine learning models (e.g., Gradient Boosted Trees) that are often more robust with smaller datasets.

Risk 2: Low User Engagement with the Serious Game. High attrition during the RCT could compromise the study's validity.

- **Contingency Plan:** This risk is primarily mitigated by our user-centered design methodology and dedicated usability testing (WP5). The DDA itself is designed to maintain optimal challenge. Should attrition rates in the RCT exceed 20%, we will implement modest, performance-contingent financial incentives to bolster adherence, a strategy with proven efficacy.

Risk 3: Difficulties in Participant Recruitment for the RCT. Recruiting 150 eligible participants may be challenging.

- **Contingency Plan:** We have a diversified recruitment strategy and the project timeline includes a multi-month buffer for recruitment. If recruitment is slower than expected, we will expand our efforts to nearby institutions and, if necessary, slightly broaden the inclusion criteria in consultation with the ethical review board, without compromising the study's internal validity.

Risk 4: Technical Challenges in Real-time System Integration. Creating a stable, low-latency feedback loop is technically complex.

- **Contingency Plan:** The system will be architected to support both moment-to-moment and less frequent "scenario-level" adaptation. The dedicated integration phase (WP5) will de-risk this process. If stable, sub-second latency proves unachievable on standard consumer hardware, the RCT will proceed by validating the scenario-level adaptation. This still represents a significant advance over static systems, allows for a robust test of our core hypotheses, and preserves the project's central innovation.

The project team possesses the requisite expertise and access to all necessary resources and infrastructure.

Relevant Previous Results of the Team The team for Subproject 1, led by Dr. Elena Vargas, has made significant contributions to computational psychiatry, particularly using NLP to identify digital biomarkers of mental illness. Their work has established the team's expertise in extracting meaningful psychological signals from complex, unstructured text and speech data, providing the essential algorithmic foundation for the proposed cognitive bias detection model, as detailed in the team's publication record.

The team for Subproject 2, led by Dr. Javier Morales, is a leading group in the design and evaluation of serious games and adaptive systems for health. Their research has focused on creating engaging user experiences grounded in psychological theory and validated through rigorous empirical methods. Their practical experience in building and validating the type of system proposed here is critical to the project's success and is evidenced by their extensive publications in top-tier digital health journals.

The PIs have a history of successful collaboration, having previously co-authored a conceptual review on the integration of AI and serious games in mental health, which laid the groundwork for this proposal.

Available Infrastructure and Resources Both teams have access to state-of-the-art facilities. Dr. Vargas's lab (SP1) is equipped with a university-managed High-Performance Computing (HPC) cluster with NVIDIA A100 GPUs, secure data storage, sound-attenuated recording booths, and a suite of Empatica E4 wearable biosensors. Dr. Morales's lab (SP2) provides a complete ecosystem for game development, including high-end developer workstations with professional licenses for Unity 3D Pro and a dedicated usability lab with eye-tracking and multi-camera recording systems. Both institutions provide comprehensive administrative support and access to large student populations for participant recruitment. This combination of computational and experimental infrastructure provides a robust foundation for the project.

Title: "General and Specific Objectives"

General and Specific Objectives

The overarching goal of the ADAPT-MIND project is to design, develop, and empirically validate a novel, closed-loop serious game system that enhances cognitive resilience by using real-time, multimodal artificial intelligence to detect cognitive biases and dynamically tailor gameplay challenges. This general objective is decomposed into five specific, measurable, achievable, relevant, and time-bound (SMART) objectives that form the core of our research plan. Each objective is directly mapped to the coordinated structure of the project, with clear responsibilities for each subproject (SP1 and SP2), providing a logical roadmap from foundational research to final validation.

Specific Objective 1 (SP1): To develop and validate a multimodal machine learning model for the real-time detection of specific, targeted cognitive biases. This foundational objective focuses on creating the core AI engine of the ADAPT-MIND system. We will target at least three clinically relevant cognitive biases (e.g., catastrophizing, confirmation bias, mind-reading) known to undermine resilience (e.g., Beck, 2011; Burns, 1980). The model will fuse data from three modalities: linguistic features from spoken responses, physiological signals (Heart Rate Variability, Galvanic Skin Response), and behavioral metrics from gameplay. Ground truth labels for model training will be generated through a rigorous annotation protocol involving multiple trained clinical experts reviewing session recordings, guided by established cognitive-behavioral therapy (CBT) frameworks. Success will be defined as achieving a classification accuracy with an Area Under the Receiver Operating Characteristic Curve (AUROC) exceeding 0.85 and an F1-score greater than 0.80 on a held-out test dataset. The model will also be optimized for an inference latency of less than 500 milliseconds to ensure real-time applicability.

Specific Objective 2 (SP2): To design and implement an engaging serious game environment with a library of scenarios specifically engineered to elicit and challenge the target cognitive biases. This objective covers the development of the user-facing platform, which serves as both the data-generation environment and the intervention delivery vehicle. Adhering to a user-centered design methodology, we will create a compelling, narrative-driven game. The primary outcome will be a fully functional game platform containing a modular library of at least 20 distinct scenarios. Each scenario will be co-designed by our interdisciplinary team of AI specialists, game designers, and clinical psychologists to reliably probe one or more target biases. Success will be measured by achieving a mean score greater than 5.5 on a 7-point scale on both the "pragmatic quality" and "hedonic quality" subscales of the standardized User Experience Questionnaire (UEQ) during usability testing.

Specific Objective 3 (SP1 & SP2): To integrate the real-time cognitive bias detection model with the serious game to create a cohesive, closed-loop, process-based Dynamic Difficulty Adjustment (DDA) system. This objective represents the central technological innovation and the primary synergy between the subprojects. We will develop a robust Application Programming Interface (API) for seamless, low-latency communication between the AI model (Objective 1) and the game engine (Objective 2). This integration will allow the game to modulate its content and mechanics in real-time based on the AI's probabilistic assessment of the user's cognitive state. The deliverable is a fully integrated system, and success will be verified through end-to-end system tests demonstrating a stable and responsive feedback loop with a round-trip latency under one second from data capture to in-game adaptation.

Specific Objective 4 (SP2, with support from SP1): To evaluate the usability, engagement, and preliminary efficacy of the adaptive ADAPT-MIND system in a three-arm randomized controlled trial (RCT). This objective focuses on the empirical validation of our central scientific hypothesis. We will conduct a rigorous trial comparing our process-based adaptive system against two active control conditions: a version of the game with traditional performance-based adaptation, and a non-adaptive version with a static difficulty curve. Our primary hypothesis is that the process-based adaptive group will show a significantly greater increase in cognitive resilience, measured by the Connor-Davidson Resilience Scale (CD-RISC), from baseline to post-intervention compared to both control groups. The trial will be powered to detect a medium effect size (Cohen's $d = 0.5$) with 80% power at an alpha of 0.05, providing a robust test of the system's therapeutic potential.

Specific Objective 5 (Coordinated): To create and disseminate a unique, large-scale, richly annotated multimodal dataset of human interaction in a cognitively challenging environment. In line with our commitment to open science, we will produce a lasting scientific resource for the broader research community. This objective involves the meticulous curation, anonymization, and documentation of the comprehensive data collected during the project. The final deliverable will be a publicly accessible dataset, shared via a reputable repository such as Zenodo or PhysioNet, containing synchronized linguistic, physiological, and behavioral data linked to clinical and demographic variables. This will address a key challenge in the field by providing high-quality, ecologically valid data to catalyze future research in computational psychiatry, affective computing, and human-centered AI, thereby amplifying the project's long-term impact.

Detailed Methodology

The ADAPT-MIND methodology is a multi-phase plan progressing from foundational component development to system integration and a final randomized controlled trial (RCT). This phased structure systematically de-risks our ambitious objectives by ensuring each component is robustly developed and tested before integration. The project is divided into three primary phases: Phase 1 (Foundational Development), Phase 2 (System Integration and Refinement), and Phase 3 (Validation Trial).

Phase 1: Foundational Development (Months 1-18) is dedicated to the parallel development of the project's core components, addressing Specific Objectives 1 and 2. This phase recognizes the symbiotic relationship between the AI model and the serious game: the game provides the ecologically valid data needed to train the model, while the model's requirements inform the design of the game's interactive scenarios. Subproject 1 will conduct a pilot study (N=50) to collect an initial, high-quality multimodal dataset of synchronized linguistic, physiological, and behavioral data. This corpus will be used to train and validate the first iteration of our cognitive bias detection model (Model v1). Concurrently, Subproject 2 will employ an agile, user-centered design methodology to create the serious game platform, involving stakeholder workshops with clinical psychologists to ground the game's content in established therapeutic principles, followed by iterative cycles of prototyping and asset creation. This phase will yield two critical assets: a validated AI model with benchmarked performance and a feature-complete, non-adaptive version of the serious game.

Phase 2: System Integration and Refinement (Months 19-24) focuses on the central technological challenge of integrating the AI model and the game into a single, closed-loop system, as outlined in Specific Objective 3. The subproject teams will collaborate to develop a robust Application Programming Interface (API) for low-latency communication. This phase will produce the three distinct versions of the ADAPT-MIND system required for the validation trial: (1) the experimental version with our novel process-based adaptation, (2) an active control with a traditional performance-based adaptation algorithm, and (3) a second active control with a static difficulty curve. A formal usability study (N=20) will then be conducted. This study serves the dual purpose of gathering user experience feedback and generating a new, highly relevant dataset. This data will enable Subproject 1 to retrain and refine the AI model into its final, optimized version (Model v2), ensuring it is trained on data from realistic user interactions.

Phase 3: Validation Trial (Months 25-36) is dedicated to the empirical validation of our central hypothesis, addressing Specific Objective 4. We will employ a three-arm, parallel-group, pre-test/post-test/follow-up RCT design, the gold standard for establishing efficacy. This design is superior to a two-arm trial as it allows us to isolate the specific effects of our novel *process-based* adaptation. By comparing our experimental system not only to a non-adaptive version but also to a state-of-the-art *performance-based* adaptation, we can determine if observed benefits are due to our specific mechanism of targeting cognitive biases, rather than adaptation in general. The inclusion of a 3-month follow-up is critical for evaluating the durability of training effects, a key indicator of a successful resilience intervention. All research involving human participants will be conducted in strict accordance with the Declaration of Helsinki and will receive prior approval from the Institutional Review Board of the coordinating institution.

To achieve the project's objectives, we will employ a suite of specific, state-of-the-art methods. The following sections detail the technical approach for each major component.

Methods for Objective 1: AI Model Development (SP1)

- **Data Acquisition Protocol and Apparatus:** The initial training dataset will be generated in a pilot study with N=50 healthy university students. Each participant will engage in a 60-minute session with a game prototype while three synchronized data streams are captured.

1. **Physiological Data:** Empatica E4 wristbands will continuously record Blood Volume Pulse (BVP) at 64 Hz and Electrodermal Activity (EDA) at 4 Hz. The E4 is a validated device for ambulatory psychophysiological research [e.g., Garbarino et al., 2014]. 2. **Linguistic Data:** Spoken responses will be captured using high-fidelity, head-mounted RØDE HS2 microphones to ensure high-quality audio for subsequent speech-to-text conversion. 3. **Behavioral Data:** The Unity-based game will be instrumented to log all relevant in-game actions (e.g., choices, latencies, navigation) with millisecond-precision timestamps via a custom logging library.

- **Data Preprocessing and Feature Extraction Pipeline:** Raw data will be processed through a semi-automated pipeline to extract meaningful features.
- **Physiological Pipeline:** Raw BVP signals will be processed using Kubios Premium (v3.5) to derive standard Heart Rate Variability (HRV) metrics (e.g., RMSSD, LF/HF ratio) in 60-second rolling windows, which are established markers of cognitive load [e.g., Berntson et al., 1997]. Raw EDA signals will be processed using the Ledalab toolbox to extract features such as skin conductance response frequency and tonic level, which are sensitive indicators of sympathetic arousal [e.g., Boucsein, 2012].
- **Linguistic Pipeline:** Audio will be transcribed using OpenAI's Whisper ASR model [Radford et al., 2023]. The resulting text will be processed using BETO, a BERT-based model for Spanish [Cañete et al., 2020], to extract contextual sentence embeddings. We will also compute higher-level features known to be associated with cognitive state, including semantic coherence, syntactic complexity, and sentiment polarity [e.g., Pennebaker et al., 2015].
- **Behavioral Pipeline:** Raw behavioral logs will be aggregated into time-series feature vectors corresponding to the same 60-second windows, capturing gameplay strategy through metrics like decision entropy and information search depth.
- **Ground Truth Annotation:** Establishing a reliable ground truth for transient cognitive states is a primary methodological challenge. We will address this by developing a novel, multi-source annotation protocol to triangulate evidence and generate robust labels. The cornerstone of this approach is expert annotation. A subset of sessions (20%) will be independently coded by two doctoral-level clinical psychologists. These experts will undergo extensive training on a detailed coding manual that defines the behavioral, linguistic, and physiological indicators for each target bias. Annotators will review synchronized recordings of gameplay, audio, and key physiological signals, applying labels from the manual to 30-second segments. Inter-rater reliability will be established on a training subset, with a target Cohen's Kappa of $\kappa > 0.80$. Discrepancies will be resolved through a consensus discussion moderated by a third, senior clinician. Given the prohibitive cost of expert-annotating the entire dataset, we will leverage this high-quality labeled subset within a semi-supervised learning framework. An initial model will be trained on the expert-labeled data and then used to generate pseudo-labels for the remaining 80%. We will employ an uncertainty sampling strategy, where the model's least confident predictions are prioritized for expert review and correction. This active learning loop iteratively enriches the training set, allowing us to efficiently scale our annotation efforts. Behavioral proxies (e.g., task performance) and ecological momentary assessment probes will serve as auxiliary features for the model and as a secondary source of validation.
- **Model Architecture, Training, and Optimization:** We will develop a multimodal deep learning model for time-series classification. The architecture will feature three parallel Long Short-Term Memory (LSTM) network branches, one for each modality

(linguistic, physiological, behavioral). The outputs will be concatenated and fed into a temporal attention mechanism, which learns to dynamically weigh the importance of each modality over time. The final output will be a softmax layer producing a probability distribution over the target cognitive biases and a "neutral" class. The model will be trained using the Adam optimizer and a cross-entropy loss function, with a 10-fold cross-validation scheme to prevent overfitting. To meet the low-latency requirement (<500ms) for real-time application, the final model will be optimized using knowledge distillation and 8-bit quantization to reduce its computational footprint.

Methods for Objective 2: Serious Game Development (SP2)

- **Game Design and Development:** The game will be developed in the Unity 3D engine (v2022.3 LTS) for PC and Mac. The design will be guided by user-centered principles and grounded in Cognitive Behavioral Therapy. The core gameplay will involve navigating narrative-driven challenges that require problem-solving and decision-making. We will use a modular design, creating a library of "scenario blocks," each a self-contained challenge co-designed to elicit specific target biases (e.g., a planning scenario with ambiguous information to probe for catastrophizing).
- **Intervention Mechanics:** Corresponding to the scenario library, we will develop a library of "intervention mechanics"—gameplay elements designed to challenge a detected bias. Examples include an "evidence log" to counter confirmation bias or a "probability slider" to counter catastrophizing.

Methods for Objective 3: System Integration (SP1 & SP2)

- **System Architecture:** The integrated system will use a client-server model. The Unity game client will run on the user's machine, buffering and sending data to a server-side inference engine via a secure REST API. This architecture prevents computationally intensive model inference from impacting game performance.
- **API and DDA Logic:** The game client will periodically (e.g., every 15 seconds) send a JSON payload with the latest feature vectors. The server will return the model's output probabilities. The game's Dynamic Difficulty Adjustment (DDA) module will then use a rule-based engine to interpret this output. For example, if the probability for "confirmation bias" exceeds a threshold of 0.8 for two consecutive windows, the DDA might select a scenario block that challenges this bias.

Methods for Objective 4: Validation Trial (SP2 & SP1)

- **Participants and Procedure:** We will recruit N=150 young adults (ages 18-25) with scores > 13 on the Perceived Stress Scale (PSS). Exclusion criteria, assessed via the Mini-International Neuropsychiatric Interview (MINI), will include major psychiatric disorders, current substance use disorder, or current psychotherapy. After providing informed consent and completing a baseline assessment (T0), participants will be randomly assigned with concealed allocation to one of three arms: (1) Process-based Adaptive, (2) Performance-based Adaptive, or (3) Non-adaptive. They will play their assigned game version for three 30-minute sessions per week for four weeks. Assessments will be repeated post-intervention (T1) and at a 3-month follow-up (T2). Outcome assessors will be blinded to group allocation.
- **Outcome Measures:**
- **Primary Outcome:** Change in resilience, measured by the Connor-Davidson Resilience Scale (CD-RISC).

- **Secondary Outcomes:** Changes in anxiety (GAD-7), depression (PHQ-9), and cognitive biases (Cognitive Bias Questionnaire; CBQ).
- **Process and Engagement Measures:** Usability (System Usability Scale; SUS), user experience (User Experience Questionnaire; UEQ), and objective engagement metrics from game logs (e.g., total playtime).

All analyses will be pre-registered on a public platform (e.g., OSF.io) prior to the RCT to ensure transparency. The analysis plan addresses the AI model's performance and the RCT's efficacy hypotheses.

Analysis for Objective 1: AI Model Performance

The final model (Model v2) will be evaluated on a held-out test set (20% of the combined pilot and usability data). We will report a comprehensive set of classification metrics, including confusion matrices, accuracy, precision, recall, and F1-score with 95% confidence intervals. The primary metric will be the Area Under the Receiver Operating Characteristic Curve (AUROC). We will also analyze the model's interpretability by examining the temporal attention weights to investigate which modalities are most informative for detecting each bias. Finally, we will benchmark the model's inference latency to confirm it meets the sub-500-millisecond requirement.

Analysis for Objective 4: Randomized Controlled Trial Efficacy

The RCT data will be analyzed following the intention-to-treat (ITT) principle.

- **Preliminary Analyses:** Baseline demographic and clinical characteristics will be compared across groups using ANOVAs and chi-square tests. Any significant baseline differences will be included as covariates. We will examine outcome variable distributions to check for normality and other statistical assumptions.
- **Primary Hypothesis Testing:** The primary hypothesis—that the process-based group will show a greater increase in resilience—will be tested using a linear mixed-effects model (LMM). The model will predict CD-RISC scores from fixed effects of group, time (T0, T1, T2), and the group-by-time interaction, with a random intercept for each participant. A significant interaction will be followed by planned post-hoc contrasts (e.g., Process-based vs. Performance-based change from T0 to T1), corrected for multiple comparisons using Tukey's HSD method.
- **Secondary Outcome and Process Analysis:** The same LMM approach will be used for continuous secondary outcomes (GAD-7, PHQ-9, CBQ). Engagement metrics will be compared across groups using one-way ANOVAs. We will also conduct exploratory correlational analyses between engagement and resilience improvement.
- **Handling of Missing Data:** If missing data exceeds 5% and is not missing completely at random (MCAR), we will use multiple imputation with chained equations (20 imputed datasets). Analyses will be performed on each dataset and results pooled according to Rubin's rules. All statistical analyses will be conducted in R (v4.2 or later), using the \`lme4\` and \`emmeans\` packages.

Research Design and Approach

To achieve the project's objectives, we will employ a multi-phase, mixed-methods research design that progresses from foundational technology development to rigorous empirical validation. This approach is structured to de-risk the project's core innovations by ensuring each component is robustly developed and tested before integration. The design comprises three sequential phases: (1) a Foundational Development phase for the parallel creation of the core AI model and the serious game platform; (2) a System Integration and Refinement phase where these components are unified and iteratively improved; and (3) a definitive Validation phase centered on a three-arm randomized controlled trial (RCT). This comprehensive design combines constructive research for building a novel technological artifact with the gold-standard experimental design for establishing its causal efficacy.

The first phase, **Foundational Development (Months 1-18)**, is dedicated to the parallel creation of the project's core components, directly addressing Specific Objectives 1 and 2. This phase leverages the symbiotic relationship between the AI model and the game environment; the game provides ecologically valid data to train the model, while the model's requirements inform the game's design. Subproject 1 will conduct a pilot study with an initial cohort of 50 participants to generate a high-quality, multimodal dataset comprising synchronized linguistic, physiological, and behavioral data. To establish ground-truth labels for cognitive bias manifestation, this data collection will be supplemented by concurrent think-aloud protocols and post-session debriefings, with subsequent expert annotation by trained clinical psychologists. This rigorous annotation process is crucial for creating a reliable training corpus for our initial cognitive bias detection model (Model v1). Concurrently, Subproject 2 will use an agile, user-centered design methodology to create the serious game platform, involving stakeholder workshops with clinical psychologists to ground the game's content in established therapeutic principles. This phase will yield two critical assets: a validated AI model with known performance characteristics and a feature-complete, non-adaptive version of the serious game.

The second phase, **System Integration and Refinement (Months 19-24)**, focuses on integrating the AI model and the game into a single, closed-loop system, as outlined in Specific Objective 3. The subproject teams will collaborate to develop a robust Application Programming Interface (API) for low-latency communication between the components. This phase will produce the three distinct versions of the ADAPT-MIND system required for the validation trial: (1) the experimental version with our novel process-based adaptation, (2) an active control with a traditional performance-based adaptation algorithm, and (3) a second active control with a static difficulty curve. A formal usability study with 20 target users will then be conducted. This study serves a dual purpose: providing essential feedback on the user experience and, critically, generating a new, highly relevant dataset. This data will be used by Subproject 1 to retrain and refine the AI model into its final, optimized version (Model v2), ensuring it is trained on data from realistic user interactions.

The third and final phase, **Validation (Months 25-36)**, is dedicated to the empirical validation of our central hypothesis through a three-arm, parallel-group, pre-test/post-test/follow-up RCT. This design is the gold standard for evaluating new interventions and is essential for making causal claims about the system's impact on cognitive resilience (e.g., CONSORT, 2010). The three arms are: 1. **Experimental Group:** Interacts with the full ADAPT-MIND system, featuring AI-driven, process-based Dynamic Difficulty Adjustment (DDA). 2. **Active Control Group 1:** Interacts with a version using a state-of-the-art, performance-based DDA (e.g., based on accuracy, reaction time). 3. **Active Control Group 2:** Interacts with a non-adaptive version following a pre-determined, static difficulty curve.

This three-arm design is a critical methodological feature. Comparing the experimental group to the non-adaptive control (Group 3) will determine the overall effect of the intervention. However, the crucial comparison is with the performance-based adaptive control (Group 2).

This allows us to test whether observed benefits are due to our specific mechanism of targeting cognitive biases in real-time, rather than being a general effect of adaptation. It isolates the added value of our core innovation, providing a powerful test of our process-based hypothesis. All groups will interact with the same core game content, ensuring that observed differences can be attributed to the adaptation mechanism.

Furthermore, the longitudinal nature of the RCT, incorporating a pre-test (T0), an immediate post-test (T1), and a 3-month follow-up (T2), is essential for a comprehensive evaluation. The pre-test establishes a stable baseline for each participant, allowing us to measure individual change. The post-test assesses the immediate impact of the four-week intervention. The 3-month follow-up is particularly critical for an intervention aimed at building resilience, as it will determine whether the cognitive changes are durable and maintained after the training period has ended—a key indicator of meaningful, real-world impact (Kazdin, 2008). This comprehensive research design provides a robust framework to achieve the project's objectives and deliver scientifically sound conclusions.

Methods, Techniques, and Procedures

To systematically achieve the project's objectives, we will employ a suite of specific, state-of-the-art methods, techniques, and procedures, detailed here in the sequence of the project's workflow from data acquisition to final validation.

Data Acquisition Protocol and Apparatus All human participant studies, including the initial pilot and the final RCT, will utilize a standardized and rigorously synchronized data acquisition protocol. Each participant session will be conducted in a sound-attenuated laboratory booth to ensure high-quality data capture. Three distinct data streams will be recorded using research-grade equipment.

1. **Physiological Data Stream:** Continuous psychophysiological data will be captured using the Empatica E4 wristband, a validated device for ambulatory research (McCarthy et al., 2016). The E4 will record Blood Volume Pulse (BVP) at a sampling rate of 64 Hz, from which heart rate and HRV are derived, and Electrodermal Activity (EDA) at 4 Hz. The device's internal accelerometer (32 Hz) will also be recorded to facilitate motion artifact detection and correction during preprocessing. 2. **Linguistic Data Stream:** Spoken responses to in-game narrative prompts and decision points will be captured using a RØDE HS2 head-mounted microphone. This professional-grade, low-profile microphone ensures a consistent mouth-to-microphone distance and minimizes environmental noise, which is critical for the accuracy of subsequent speech-to-text processing. Audio will be captured as uncompressed WAV files at a 44.1 kHz sampling rate and 16-bit depth via a Focusrite Scarlett 2i2 audio interface. 3. **Behavioral Data Stream:** The serious game, developed in the Unity 3D engine (v2022.3 LTS), will be extensively instrumented with a custom-built logging library. This library will record a comprehensive set of in-game events as timestamped JSON objects. Logged events will include every player input (e.g., key presses, mouse clicks), dialogue choices, response latencies, navigation paths through the game world, interaction with specific game objects, and performance metrics on cognitive mini-games (e.g., accuracy, time-to-completion).

Synchronization of these three streams is a critical technical procedure. All data acquisition computers will be synchronized to a central Network Time Protocol (NTP) server at the start of each day to ensure a common time base. At the beginning of each recording session, the Unity logging library will send a trigger signal via Lab Streaming Layer (LSL), an open-source solution for synchronized data streaming (Kothe, 2014). This LSL marker will be simultaneously recorded in the physiological and audio data streams, providing a precise, shared event marker that allows for sub-millisecond alignment of all three datasets during post-processing.

Data Preprocessing and Feature Extraction Pipelines Following acquisition, the raw data will be processed through modality-specific pipelines to extract a rich set of features for the AI model. All time-series features across modalities will be computed over 60-second rolling windows with a 50% overlap to ensure temporal alignment.

- **Physiological Pipeline:** Raw BVP signals will be visually inspected for artifacts, and segments with excessive motion noise (identified via the accelerometer data) will be excluded. The cleaned BVP signal will be processed using Kubios Premium software (v3.5) to calculate inter-beat intervals and derive standard Heart Rate Variability (HRV) metrics. Key features will include time-domain metrics (RMSSD, SDNN), which reflect parasympathetic influence, and frequency-domain metrics (LF/HF ratio), an indicator of sympathovagal balance (Task Force of ESC/NASPE, 1996). Raw EDA signals will be processed using the Ledalab toolbox in MATLAB (v2022b). The signal will be decomposed into its tonic and phasic components. Features extracted from the phasic component will include the number of non-specific skin conductance responses (SCRs) and their mean amplitude, which are sensitive indicators of sympathetic arousal (Boucsein, 2012). The mean tonic skin conductance level will also be calculated for each window.
- **Linguistic Pipeline:** The high-quality audio recordings will be transcribed to text using OpenAI's Whisper ASR model, which has demonstrated state-of-the-art performance on a wide range of accents and recording conditions (Radford et al., 2023). The resulting text will be segmented into utterances. Each utterance will be processed using a pre-trained Spanish language model (BETO), a BERT-based model specifically trained on a large Spanish corpus (Cañete et al., 2020). We will extract contextual sentence embeddings from the final hidden layer of the model. In addition to these dense representations, we will compute a suite of higher-level linguistic features known to be associated with cognitive state. These include: semantic coherence between successive utterances (measured as the cosine similarity of their embeddings), syntactic complexity (e.g., dependency parse tree depth, Flesch-Kincaid grade level), and sentiment polarity using established Spanish lexicons (Pérez-Rosas et al., 2012).
- **Behavioral Pipeline:** Raw behavioral logs in JSON format will be parsed and aggregated into time-series feature vectors corresponding to the standardized analysis windows. A library of over 50 features will be engineered to capture gameplay strategy and decision-making patterns. Examples include decision entropy (to quantify randomness in choices), information search depth (how many optional clues a player accesses before making a decision), action rates for different types of interactions (e.g., movement vs. dialogue), and hesitation times (latency between prompt and response).

Ground Truth Annotation Procedure Establishing a reliable ground truth for transient cognitive biases is a central methodological challenge. To address this, we will implement a rigorous, multi-layered annotation procedure.

- **Codebook Development:** A comprehensive annotation codebook will be developed iteratively by the core research team, including two senior clinical psychologists. The codebook will be grounded in established Cognitive Behavioral Therapy (CBT) literature (e.g., Beck, 2011) and will provide precise, operationalized definitions and behavioral indicators for each target cognitive bias (e.g., Catastrophizing: verbalizing an irrationally worst-case outcome; Confirmation Bias: selectively attending to or seeking information that confirms pre-existing beliefs). The codebook will undergo several rounds of refinement based on pilot data analysis.

- **Expert Annotation Protocol:** A subset of sessions from the pilot study (approximately 20%, selected to maximize diversity of responses) will be independently annotated by two trained clinical psychologists. Using the ELAN annotation software, annotators will review synchronized gameplay video, audio transcripts, and physiological data streams. Rather than labeling coarse 30-second windows, annotation will be performed at the level of individual utterances or discrete behavioral events. This granular approach allows for more precise alignment with the extracted linguistic and behavioral features. The annotation process will involve two passes: a first pass to identify potential instances of cognitive bias, and a second pass for detailed labeling according to the codebook, including a confidence rating (low/medium/high).
- **Reliability and Consensus:** Annotator training will involve joint coding of a benchmark dataset until a consistent inter-rater reliability (Cohen's Kappa) of $\kappa > 0.75$ is achieved. During the main annotation phase, reliability will be continually monitored. All disagreements will be resolved through a consensus meeting with a third senior clinician, who will act as an adjudicator.
- **Triangulation with Self-Report:** To further validate the expert labels, the serious game will include strategically placed Ecological Momentary Assessment (EMA) probes. At the conclusion of certain challenging scenarios, participants will be briefly prompted to self-report their thought patterns via a multiple-choice question. While subjective, these self-reports will provide a valuable secondary data source to triangulate with the expert annotations and enhance the ecological validity of our ground truth labels. This multi-faceted approach ensures the creation of robust, high-quality labels essential for supervised model training.

AI Model Architecture, Training, and Optimization We will develop a multimodal deep learning model for time-series classification. The architecture will consist of three parallel branches, one for each modality. Each branch will utilize a two-layer Long Short-Term Memory (LSTM) network with 128 units per layer to learn temporal patterns from the respective feature sequences. The final hidden states from each LSTM branch will be concatenated and fed into a Bahdanau-style attention mechanism (Bahdanau, Cho, & Bengio, 2015). This mechanism will learn to dynamically weigh the importance of each modality at each time step, allowing the model to adapt its focus based on the context. The attention-weighted context vector will then be passed through a final dense layer with a softmax activation function to produce a probability distribution over the target cognitive biases and a "neutral" class.

The model will be trained using the Adam optimizer with a cross-entropy loss function and a learning rate of 1e-4. To prevent overfitting, we will employ dropout with a rate of 0.4 on the LSTM layers and L2 regularization. The model will be trained and evaluated using a rigorous 10-fold cross-validation scheme on the annotated dataset. To meet the low-latency requirements (<500ms) for real-time application, the final trained model will be optimized using knowledge distillation. A smaller, faster "student" model (e.g., with a single LSTM layer and fewer units) will be trained to mimic the probability outputs of the larger, more complex "teacher" model. The final student model will be further optimized via 8-bit quantization and deployed using the ONNX Runtime for efficient inference.

Serious Game and Intervention Mechanics The game will be developed using the Unity 3D engine for deployment on PC and Mac platforms. The design is grounded in CBT principles and follows a modular structure. The core of the game is a library of "scenario blocks," each a self-contained narrative challenge designed to elicit one or more target biases. For example, a "planning an expedition" scenario might present ambiguous information about weather patterns to probe for catastrophizing. Correspondingly, we will develop a library of "intervention mechanics," which are specific gameplay elements designed to actively challenge a detected bias. Examples include:

- **Evidence Log:** To counter confirmation bias, this mechanic prompts the player to explicitly list arguments for and against a belief before committing to a decision.
- **Probability Slider:** To counter catastrophizing, this mechanic requires players to visually estimate the likelihood of various outcomes on a slider, rewarding realistic assessments.
- **Perspective-Taking Dialogue:** To counter mind-reading, this mechanic presents ambiguous social cues from a non-player character and rewards the player for exploring multiple possible interpretations of their behavior through a branching dialogue tree.

System Integration and Dynamic Difficulty Adjustment (DDA) Logic The integrated system will operate in a client-server model. The Unity game client will run on the user's machine, collecting and preprocessing data locally. Every 15 seconds, it will send a JSON payload containing the latest feature vectors from all three modalities to a secure, server-side inference engine via a REST API. The server, running the optimized Python model in a Flask application, will return a JSON object with the model's output probabilities. The game's DDA module will then use a rule-based engine to interpret this output. This engine will use temporal logic to avoid overly reactive changes. For example, a rule might state: "If the probability for 'confirmation bias' exceeds a threshold of 0.8 for two consecutive time windows, and the player is currently in a free-exploration phase, then select the next scenario block from the library that specifically challenges this bias." This ensures that adaptations are both timely and contextually appropriate.

Validation Trial Procedures The RCT will be conducted in strict accordance with the pre-registered protocol and IRB approval.

- **Recruitment and Screening:** Participants (N=150) will be recruited via university participant pools and online advertisements. Initial screening will be conducted via an online survey including the Perceived Stress Scale (PSS). Eligible individuals will undergo a remote screening interview with a trained research assistant using the Mini-International Neuropsychiatric Interview (MINI) to apply exclusion criteria.
- **Randomization and Blinding:** After providing informed consent and completing the baseline assessment (T0), participants will be randomized to one of the three study arms using a computer-generated sequence with concealed allocation managed by a researcher not involved in data collection. Outcome assessors conducting the T1 and T2 assessments will be blinded to the participant's group allocation.
- **Intervention Protocol:** Participants will be instructed to play their assigned version of the game from home for three 30-minute sessions per week for four weeks. Adherence will be tracked automatically via server logs, with automated email reminders sent for missed sessions.

- **Data Management and Ethics:** All data will be stored on a secure, encrypted server. Personally identifiable information will be stored separately from the research data and linked only by a unique, random participant ID. All procedures will adhere to GDPR and the ethical principles outlined in the Declaration of Helsinki.

Data Analysis Plan

Data analysis will proceed in two main stages, corresponding to the evaluation of the AI model's performance (Objective 1) and the testing of the RCT's primary and secondary hypotheses (Objective 4). To ensure transparency and prevent p-hacking, all analyses will be formally pre-registered on a public platform, such as the Open Science Framework (OSF.io), prior to commencing the RCT. Statistical analyses will be performed using the R statistical programming language (v4.2 or later).

The first stage of analysis will focus on the performance of the final cognitive bias detection model (Model v2). This evaluation will be conducted on a held-out test set comprising 20% of the combined data from the pilot and usability studies, ensuring the model is assessed on data unseen during training or tuning. We will report a comprehensive set of standard classification metrics to provide a complete picture of the model's capabilities. This will include confusion matrices to visualize class-specific performance, alongside overall metrics such as accuracy, precision, recall, and F1-score, each reported with 95% confidence intervals. The primary metric for model selection and reporting will be the Area Under the Receiver Operating Characteristic Curve (AUROC), which provides a robust, threshold-independent measure of discriminability. Beyond predictive accuracy, we will analyze the model's interpretability by examining its temporal attention weights. This procedure will allow us to investigate which modalities (linguistic, physiological, behavioral) the model deems most informative for detecting each specific cognitive bias, yielding valuable scientific insights into the multimodal signatures of these cognitive states. Finally, we will benchmark the model's inference latency on the target server hardware to confirm it meets the sub-500-millisecond requirement for real-time application.

The second, and principal, stage of analysis will evaluate the efficacy of the ADAPT-MIND intervention using data from the three-arm RCT. All efficacy analyses will be conducted according to the intention-to-treat (ITT) principle, wherein all participants are analyzed in the group to which they were randomized, regardless of protocol adherence. This approach provides the most conservative and clinically relevant estimate of the intervention's real-world effect (e.g., Gupta, 2011).

Preliminary analyses will be conducted to ensure the integrity of the randomization and verify the assumptions of our statistical models. Baseline demographic and clinical characteristics will be compared across the three groups using one-way ANOVA for continuous variables and chi-square tests for categorical variables. Any variables showing a significant baseline difference ($p < .05$) will be included as covariates in the primary models to control for their potential confounding effects. We will also examine the distributions of all outcome variables at each time point to check for normality and other assumptions of the planned statistical tests.

The primary hypothesis—that the process-based adaptive group will show a greater increase in resilience than the two control groups—will be tested using a linear mixed-effects model (LMM). The model will predict the Connor-Davidson Resilience Scale (CD-RISC) score from fixed effects of group (a three-level factor: Process-based, Performance-based, Non-adaptive), time (a three-level factor: T0, T1, T2), and the crucial group-by-time interaction. A random intercept for each participant will be included to account for the non-independence of repeated measures. A significant group-by-time interaction term will indicate that the trajectory of resilience scores over time differs between the groups. If this interaction is significant, we will conduct planned post-hoc contrasts to test our specific hypotheses: (1) the change from T0 to T1 in the Process-based group is greater than in the Performance-based group, and (2) the change from T0 to T1 in the Process-based group is greater than in the Non-adaptive group. Similar contrasts will be performed for the T0 to T2 change to assess the durability of the effect. All post-hoc tests will be corrected for multiple comparisons using Tukey's HSD method.

This LMM framework will be consistently applied to analyze the continuous secondary outcomes (GAD-7, PHQ-9, CBQ). Process and engagement metrics will also be analyzed to understand user behavior. Objective engagement metrics (e.g., total playtime, number of sessions completed) will be compared across the three groups using one-way ANOVAs to determine if the adaptive systems are more engaging. We will also conduct exploratory correlational analyses to examine whether the magnitude of improvement in resilience is associated with the level of engagement within each group.

We anticipate some participant attrition and will monitor it closely. If the proportion of missing data exceeds 5% and is not determined to be missing completely at random (MCAR) via Little's MCAR test, we will use multiple imputation with chained equations to create 20 complete datasets. The LMM analyses will then be performed on each imputed dataset, and the results will be pooled according to Rubin's rules (e.g., Little & Rubin, 2019). This approach is current best practice and yields less biased estimates than listwise deletion. This pre-registered and rigorous analytical plan provides a robust framework for testing our central hypotheses and drawing clear, evidence-based conclusions regarding the efficacy of the ADAPT-MIND system.

Work Plan and Schedule

The execution of the ADAPT-MIND project is structured into nine interconnected Work Packages (WPs) to ensure a logical progression from foundational research to final validation and dissemination. This framework provides a clear division of labor between the two subprojects (SP1 and SP2) while embedding the collaborative mechanisms essential for success. Each WP is defined by its specific objectives, a detailed breakdown of tasks, and clear lines of responsibility.

WP1: Project Management and Coordination (Lead: Coordinator - Dr. Vargas; Months 1-36) is the overarching package that ensures the project's administrative, financial, and scientific integrity. Its primary objective is to facilitate seamless collaboration and guarantee that all objectives are met on time and within budget. Key tasks include: (1.1) Establishing the Steering Committee (SC) and organizing bi-weekly SC meetings and monthly all-hands meetings; (1.2) Setting up and maintaining the shared digital collaboration ecosystem (GitHub, Slack, project management platform); (1.3) Managing all financial and scientific reporting to the funding agency; (1.4) Overseeing the ethical approval process for all human participant studies, ensuring compliance with institutional and national regulations; and (1.5) Proactively managing risks and implementing contingency plans. While the Coordinator leads this WP, it is a joint responsibility, with the PI of SP2 serving as co-chair of the SC.

WP2: Data Acquisition Protocol and Pilot Study (Lead: SP1; Months 1-9) aims to generate the foundational, high-quality multimodal dataset required for initial AI model training. This WP is critical for de-risking the core technology. Tasks include: (2.1) Finalizing the selection of biosensors and recording equipment and developing data synchronization software using Lab Streaming Layer (LSL); (2.2) Preparing and submitting the study protocol to the Institutional Review Board (IRB) for full ethical approval; (2.3) Recruiting and scheduling 50 participants for the pilot study; (2.4) Conducting the 60-minute data collection sessions with strict protocol adherence; and (2.5) Performing initial data quality checks, preprocessing, and secure storage of raw and processed data. SP1 leads technical aspects, while SP2 provides the initial game prototype and assists with user-facing protocol elements.

WP3: Cognitive Bias AI Model Development (v1) (Lead: SP1; Months 4-18) is focused on achieving the first specific objective: creating the initial version of the core AI engine. The objective is to develop a robust multimodal fusion model capable of classifying cognitive biases from the pilot dataset. The tasks are: (3.1) Intensive preprocessing and feature extraction from the linguistic, physiological, and behavioral data streams collected in WP2; (3.2) Establishing a rigorous annotation protocol, including a formal coding manual for identifying cognitive bias manifestations in multimodal data. This involves training at least two independent expert annotators and conducting a reliability study on a 20% subset of the data to establish high inter-rater reliability (target Krippendorff's $\alpha > 0.80$) before proceeding with the full annotation to create the ground truth; (3.3) Designing, implementing, and training the multimodal deep learning architecture, including LSTM branches and an attention mechanism; and (3.4) Rigorously validating this initial model (Model v1) using 10-fold cross-validation and reporting its performance metrics. This WP is the primary responsibility of SP1.

WP4: Serious Game Platform and Content Creation (Lead: SP2; Months 1-18), running in parallel with WP3, addresses the second specific objective. The objective is to create a fully functional, engaging serious game that serves as the platform for data collection and intervention. Tasks include: (4.1) Drafting the comprehensive Game Design Document (GDD), detailing the narrative, core mechanics, and user interface; (4.2) Developing the core game engine and mechanics in Unity, including character control, dialogue systems, and data logging infrastructure; (4.3) Creating all necessary 2D/3D art, sound assets, and user interface elements; and (4.4) Building a modular library of at least 20 distinct scenarios and corresponding intervention mechanics, co-designed with SP1 and clinical advisors to effectively elicit and target the desired cognitive biases. This WP is the primary responsibility of SP2.

WP5: System Integration and Usability Testing (Lead: SP2; Months 19-24) is a critical, highly collaborative WP that integrates the outputs of WP3 and WP4 to achieve the third specific objective. The objective is to create a stable, closed-loop system and evaluate its usability. Tasks are: (5.1) Jointly developing the REST API specification for communication between the game client and the AI server; (5.2) Integrating AI Model v1 into the game platform to create the three experimental arms (process-based, performance-based, non-adaptive); (5.3) Recruiting 20 target users for formal usability testing, involving task-based assessments and qualitative feedback sessions; and (5.4) Analyzing usability metrics (e.g., SUS, UEQ) and qualitative feedback to identify areas for refinement. SP2 leads integration and user testing, while SP1 provides the AI model as a containerized service and supports API development.

WP6: AI Model Refinement and Finalization (v2) (Lead: SP1; Months 22-27) leverages data and insights from WP5 to produce the final, production-ready AI model. The objective is to enhance the model's accuracy and efficiency for the RCT. Tasks include: (6.1) Retraining the model on the combined dataset from the pilot (WP2) and usability (WP5) studies for greater ecological validity; (6.2) Fine-tuning the model architecture and hyperparameters based on performance; and (6.3) Performing final optimization (e.g., knowledge distillation, quantization) to produce the low-latency, production-ready version of the model (Model v2). This WP is the responsibility of SP1.

WP7: Randomized Controlled Trial (RCT) (Lead: SP2; Months 25-36) is dedicated to the definitive empirical validation of the ADAPT-MIND system, addressing the fourth specific objective. The objective is to rigorously test the project's central hypothesis regarding the efficacy of process-based adaptation. Tasks include: (7.1) Submitting amendments and obtaining final IRB approval for the full trial protocol; (7.2) Managing the recruitment, screening, and consenting of 150 participants; (7.3) Overseeing the 4-week intervention period, including participant onboarding, technical support, and adherence monitoring; and (7.4) Conducting data collection at three assessment points (T0, T1, T2) and ensuring data quality and integrity. SP2 leads all logistical and participant-facing aspects, while SP1 provides continuous technical support for the server-side AI and data logging infrastructure.

WP8: Data Analysis and Interpretation (Lead: Joint - SP1 & SP2; Months 34-36) focuses on analyzing the rich dataset from the RCT to statistically test the project's hypotheses. Tasks include: (8.1) Cleaning, processing, and structuring the final RCT dataset; (8.2) Conducting primary and secondary statistical analyses as defined in the data analysis plan, primarily using linear mixed-effects models to account for repeated measures; and (8.3) Jointly interpreting the results in the context of the project's hypotheses and preparing figures and tables for dissemination. SP1 will lead the statistical modeling of clinical outcomes, while SP2 will lead the analysis of usability and engagement data.

WP9: Dissemination, Exploitation, and Data Sharing (Lead: Coordinator; Months 1-36) is an ongoing WP ensuring the project's impact extends beyond its duration. The objective is to maximize the scientific, social, and economic impact of the project's results. Tasks include: (9.1) Developing and maintaining a public-facing project website; (9.2) Preparing and submitting at least three manuscripts to high-impact, peer-reviewed journals; (9.3) Presenting findings at major international conferences in AI, HCI, and computational psychiatry; (9.4) Executing the Data Management Plan, including the final anonymization, documentation, and public release of the multimodal dataset (Objective 5); and (9.5) Developing a preliminary plan for the potential valorization and exploitation of the project's results. This is a coordinated effort, with all team members contributing.

The project is scheduled over a 36-month period, with a timeline designed to manage dependencies and ensure a logical flow. The schedule is organized into three main phases, with Work Packages overlapping to maximize efficiency. A detailed visual representation is provided in the supplementary materials (see Gantt Chart).

The first year (Months 1-12) is the **Foundational Phase**, focused on establishing the project's core components. Following ethical approval, WP2 (Pilot Study) will generate the initial dataset by Month 9. This dataset enables the project's two parallel development streams: WP3 (AI Model v1) and WP4 (Serious Game Platform). This concurrent structure is a key efficiency, allowing the AI and game components to be developed independently while ensuring future compatibility through close coordination. By the end of this phase, the project will have delivered the pilot dataset, a complete Game Design Document, a playable game prototype, and the first version of the AI model in training.

The second year (Months 13-24) is the **Integration and Refinement Phase**, marking the convergence of the parallel development streams. It begins with the completion of the standalone AI model (WP3) and game platform (WP4) at Month 18. This immediately triggers the highly collaborative WP5 (System Integration and Usability Testing), where the two components are merged into a closed-loop system. This intensive phase culminates in a fully integrated prototype that has been tested and refined based on user feedback, thereby mitigating the project's primary technological risks.

The third year (Months 25-36) is the **Validation and Dissemination Phase**, dedicated to empirical validation and impact generation. The phase commences with the final optimization of the AI model (WP6), informed by usability testing data. The primary activity is the large-scale Randomized Controlled Trial (WP7), which will run for the remainder of the project to rigorously test the central hypothesis. The final months (34-36) are reserved for the comprehensive statistical analysis of trial data (WP8) and the execution of the dissemination and data sharing plan (WP9), ensuring the project concludes with high-impact scientific outputs and publicly available resources.

The project's progress will be rigorously tracked against a series of concrete deliverables and major milestones. Each deliverable represents a tangible output linked to a specific WP, providing clear evidence of progress. Milestones are critical checkpoints that signify the successful completion of a major project phase, serving as go/no-go decision points for the Steering Committee.

Deliverables:

- **D1.1:** Project Handbook, including communication protocols and initial Data Management Plan (Lead: Coordinator; M3; Confidential).
- **D1.2, D1.3:** First and Second Year Progress and Financial Reports (Lead: Coordinator; M12, M24; Confidential).
- **D2.1:** Finalized Data Acquisition Protocol and full IRB Approval for pilot study (Lead: SP1; M6; Confidential).
- **D2.2:** The curated, preprocessed, and documented Pilot Dataset (N=50) (Lead: SP1; M9; Confidential).
- **D3.1:** Annotation Protocol and Inter-Rater Reliability Report (Lead: SP1; M12; Confidential).
- **D3.2:** Report on AI Model v1 architecture, training procedures, and validation performance (Lead: SP1; M18; Confidential).
- **D4.1:** The complete Game Design Document (GDD) (Lead: SP2; M6; Confidential).
- **D4.2:** The final, feature-complete, non-adaptive version of the serious game (Lead: SP2; M18; Confidential).
- **D5.1:** System Integration API Specification document (Lead: SP2; M19; Confidential).
- **D5.2:** Report on Usability Testing results, including quantitative metrics and qualitative feedback (Lead: SP2; M24; Confidential).
- **D6.1:** Report on the final, optimized AI Model v2 architecture and performance (Lead: SP1; M27; Confidential).
- **D7.1:** The pre-registered RCT Protocol and final IRB Approval for the full trial (Lead: SP2; M26; Public).
- **D7.2:** The final, fully anonymized, and documented RCT Dataset (N=150) (Lead: SP2; M35; Confidential; to be made public via D9.3).

- **D8.1:** Final scientific report on RCT results and interpretation (Lead: Joint; M36; Confidential; results to be published via D9.2).
- **D9.1:** Public-facing project website launch (Lead: Coordinator; M6; Public).
- **D9.2:** Submission of at least three peer-reviewed manuscripts (Lead: Coordinator; M36; Public).
- **D9.3:** Public release of the anonymized multimodal dataset with documentation, per FAIR principles (Lead: Coordinator; M36; Public).

Milestones:

- **MS1: Pilot Data Collection Complete (M9).** Confirms the functionality of the multimodal data acquisition pipeline and the successful, ethical acquisition of the foundational dataset for model training. *Verification:* Delivery of D2.2 and a Steering Committee report confirming data quality.
- **MS2: Standalone Components Ready for Integration (M18).** Marks the successful completion of the parallel development phase, signifying that both AI Model v1 and the game platform meet their specified requirements. *Verification:* Successful delivery and approval of D3.2 and D4.2 by the Steering Committee.
- **MS3: Integrated System Validated for Usability (M24).** Demonstrates that the core technological challenge of creating a stable, functional closed-loop system has been met and refined based on user feedback, making it ready for the RCT. *Verification:* Successful delivery of D5.2 and a live demonstration of the integrated system to the Steering Committee.
- **MS4: RCT Recruitment Complete (M30).** Indicates that the final validation phase is on schedule to achieve the required statistical power for the trial. *Verification:* A report confirming that 150 participants have been successfully recruited, consented, and randomized.
- **MS5: Final Project Results Analyzed and Disseminated (M36).** Marks the successful completion of all scientific objectives, with primary results analyzed and key dissemination outputs submitted. *Verification:* Delivery of D8.1 and proof of submission for D9.2.

Work Packages (WPs), Tasks, and Responsibilities

The execution of the ADAPT-MIND project is structured into nine interconnected Work Packages (WPs), designed to ensure a logical progression from foundational research to final validation and dissemination. This framework provides a clear division of labor between the two subprojects while embedding the collaborative mechanisms essential for the project's success. Each WP is defined by its specific objectives, a detailed breakdown of tasks, and clear lines of responsibility.

WP1: Project Management and Coordination is the overarching package ensuring the project's administrative, financial, and scientific execution. Led by the Coordinator, Dr. Vargas, this WP will run for the entire 36-month duration. Its primary objective is to facilitate effective collaboration and guarantee that all objectives are met on time and within budget. Key tasks include establishing the Steering Committee (SC) and organizing bi-weekly SC meetings and monthly all-hands meetings to maintain strategic alignment. The WP also covers the setup and maintenance of the shared digital infrastructure, including GitHub for version control, Slack for daily communication, and a project management platform for tracking dependencies. The Coordinator will manage all financial and scientific reporting to the funding agency and oversee the ethical approval process for all human participant studies, ensuring compliance with all relevant regulations. Proactively managing risks and implementing contingency plans identified by the SC is another core responsibility. While the Coordinator leads this WP, it is a joint responsibility, with the PI of SP2, Dr. Morales, serving as co-chair of the SC.

WP2: Data Acquisition Protocol and Pilot Study, led by SP1, aims to generate the foundational, high-quality multimodal dataset required for initial model training. This WP is critical for de-risking subsequent development by validating the data acquisition pipeline. Tasks begin with the final selection and procurement of biosensors (e.g., EEG, EDA, ECG) and the development of data synchronization software using Lab Streaming Layer (LSL) for millisecond-precision alignment. The team will then prepare and submit the study protocol to the Institutional Review Board (IRB) for ethical approval. Following approval, 50 participants will be recruited for the pilot study. The core task is executing the 60-minute data collection sessions with strict protocol adherence to maximize data quality. The WP concludes with initial data quality checks, preprocessing, and secure storage of the raw and processed data. SP1 leads all technical aspects, while SP2 provides the initial game prototype and assists with the user-facing protocol, representing the first major point of inter-project collaboration.

WP3: Cognitive Bias AI Model Development (v1) is focused on achieving the first specific objective: creating the initial version of the core AI engine. Led by SP1, the objective is to develop a robust multimodal fusion model capable of classifying cognitive biases from the pilot dataset. Tasks include comprehensive preprocessing and feature extraction from the linguistic, physiological, and behavioral data streams. Concurrently, a rigorous annotation protocol will be established. This involves developing a detailed annotation codebook grounded in cognitive theory, defining behavioral and physiological markers for each target bias. A team of two trained annotators will independently label a subset of the data to establish ground truth. Inter-rater reliability will be systematically assessed using Krippendorff's Alpha, with a target $\alpha > 0.7$. Discrepancies will be resolved through consensus discussion with a senior researcher to refine the codebook and ensure annotation consistency. The central task is the design, implementation, and training of the multimodal deep learning architecture. The WP culminates in the validation of this initial model (Model v1) using 10-fold cross-validation and the reporting of its performance metrics. This WP is the primary responsibility of SP1.

WP4: Serious Game Platform and Content Creation, running in parallel with WP3, addresses the second specific objective: developing the user-facing application. Led by SP2, the objective is to create a fully functional, engaging serious game that serves as the platform for data collection and intervention. The first task is to draft the comprehensive Game Design Document (GDD), detailing the narrative, core mechanics, and user interface. The main development tasks include building the core game engine and mechanics in Unity, including character control, dialogue systems, and the data logging infrastructure. The team will also create all necessary 2D/3D art, sound assets, and user interface elements. A crucial task is constructing a modular library of at least 20 distinct scenarios designed to probe specific biases (e.g., confirmation bias, attentional bias) through in-game choices and challenges. This task will be co-designed with input from SP1 and clinical advisors to ensure the scenarios effectively elicit and target the desired cognitive biases.

WP5: System Integration and Usability Testing is a critical, highly collaborative WP that brings the outputs of WP3 and WP4 together to achieve the third specific objective. Led by

SP2, the objective is to create a stable, closed-loop system and evaluate its usability. Tasks begin with the joint development of the REST API specification for communication between the game client and the AI server. This is followed by the technical integration of AI Model v1 into the game platform to deploy the three distinct adaptation logic arms (process-based, performance-based, non-adaptive control). Subsequently, 20 target users will be recruited for formal usability testing, using task-based assessments and think-aloud protocols to gather qualitative feedback. The final task is the analysis of the resulting usability metrics (e.g., SUS, UEQ) and qualitative feedback to identify areas for refinement. SP2 leads the integration and user testing, while SP1 provides the AI model as a containerized service, offers technical support for the API, and analyzes the multimodal data generated during testing for use in WP6.

WP6: AI Model Refinement and Finalization (v2) leverages the data and insights from WP5 to produce the final, production-ready AI model. Led by SP1, the objective is to enhance the model's accuracy and efficiency for the RCT. The primary task is to retrain the model on the combined dataset from the pilot (WP2) and usability (WP5) studies, which provides data captured during more naturalistic, goal-oriented interaction, thereby improving ecological validity. The team will then fine-tune the model architecture and hyperparameters based on performance on a held-out validation set. This WP will also explore semi-supervised learning techniques to leverage the larger pool of unannotated data for improved generalization. The final task is to perform optimization, including knowledge distillation and quantization, to produce the low-latency, production-ready version of the model (Model v2) that meets the system's real-time requirements.

WP7: Randomized Controlled Trial (RCT) is dedicated to the definitive empirical validation of the ADAPT-MIND system, addressing the fourth specific objective. Led by SP2, the objective is to empirically test the central hypothesis that process-based adaptation is superior to performance-based and non-adaptive approaches in enhancing resilience. Tasks include submitting amendments and obtaining final IRB approval for the full trial protocol. The team will then manage the recruitment, screening, and consenting of 150 participants. A major task is overseeing the 4-week intervention, including participant onboarding, technical support, and adherence monitoring. The final task is to conduct data collection at the three assessment points (T0, T1, T2) and ensure the quality and integrity of the collected clinical and behavioral data. SP2 leads all logistical and participant-facing aspects of the trial, while SP1 provides continuous technical support for the server-side AI and data logging infrastructure.

WP8: Data Analysis and Interpretation is a joint WP focused on analyzing the rich dataset from the RCT. The objective is to statistically test the project's hypotheses and interpret the findings. Tasks include cleaning, processing, and structuring the final RCT dataset. The core task is to conduct the primary and secondary statistical analyses as defined in the data analysis plan, employing linear mixed-effects models to assess changes in primary and secondary outcomes over time (T0, T1, T2), treating treatment arm as a fixed effect and participants as a random effect to account for individual baseline differences. The final, highly collaborative task is to jointly interpret the results in the context of the project's hypotheses and prepare figures and tables for dissemination. SP1 will lead the statistical modeling of the clinical outcomes, while SP2 will lead the analysis of usability and engagement data, with both PIs collaborating on the final scientific narrative.

WP9: Dissemination, Exploitation, and Data Sharing is an ongoing WP that ensures the project's impact extends beyond its duration. Led by the Coordinator, its objective is to maximize the scientific, social, and economic impact of the project's results. Tasks include developing and maintaining a public-facing project website. The team will prepare and submit at least three manuscripts to high-impact journals in computational psychiatry, human-computer interaction (HCI), and digital mental health. A critical task is executing the Data Management Plan, which culminates in the final anonymization, documentation, and public release of the multimodal dataset, fulfilling Objective 5. Finally, the team will develop a concrete valorization and exploitation strategy, identifying potential pathways for clinical or commercial application. This is a coordinated effort, with all team members contributing to dissemination activities.

Timeline and Gantt Chart

The project's 36-month timeline is structured across three distinct phases to manage critical dependencies and ensure a logical progression from foundational research to final validation. Work Packages are strategically overlapped to maximize efficiency and foster continuous integration between development teams. A detailed Gantt chart, illustrating the duration, sequencing, and interdependencies of each Work Package and its major tasks, is provided in the supplementary materials (Figure X).

Year 1 (Months 1-12) constitutes the **Foundational Phase**, dedicated to establishing the project's core scientific and technical components. Alongside continuous project management (WP1) and dissemination planning (WP9), this phase prioritizes the critical empirical groundwork. WP2 (Pilot Study, Months 1-9) is designed to develop and validate our methodology for cognitive bias annotation, delivering the essential initial dataset for model training. Building on this, the two primary development streams—WP3 (AI Model v1) and WP4 (Serious Game Platform)—commence in parallel at Month 4. This concurrent structure is a key efficiency of the work plan, enabling close, iterative collaboration between the AI and game teams to ensure seamless integration. By the end of Year 1, key deliverables will include a comprehensive Game Design Document, a playable game prototype, and the first trained iteration of the AI model.

Year 2 (Months 13-24) is the **Integration and Refinement Phase**. Following the conclusion of WP3 and WP4 at Month 18, the first complete, standalone versions of the AI model and the game platform become available. These outputs serve as the direct inputs for WP5 (System Integration and Usability Testing), the central focus from Month 19 to Month 24. This intensive, collaborative phase is dedicated to merging the two systems and iteratively refining them based on quantitative performance metrics and qualitative user feedback. The successful completion of WP5, marked by a comprehensive system performance and usability report, will signify that the project's core technological risks have been effectively mitigated.

The third and final year (Months 25-36) is the **Validation and Dissemination Phase**. The year begins with WP6 (AI Model Refinement), which leverages the rich dataset from WP5 to finalize the AI model by Month 27. The main activity is WP7 (Randomized Controlled Trial), which will run from Month 25 to Month 36. This comprehensive trial includes participant recruitment and the 4-week intervention, followed by a 3-month follow-up data collection period concluding around Month 34. The final months of the project (34-36) are dedicated to WP8 (Data Analysis) for the RCT data, and the final reporting and dissemination activities of WP1 and WP9. This culminates in the submission of key manuscripts for peer-reviewed publication and the public release of the final, anonymized dataset, ensuring the project concludes with high-impact, tangible outputs.

Deliverables and Milestones

Project progress will be monitored through defined deliverables and strategic milestones. This framework enables the Steering Committee to objectively assess progress, manage risks, and ensure alignment with the project's timeline. Deliverables are tangible outputs from each Work Package, while milestones mark the completion of major phases and serve as key decision points.

The project will produce a comprehensive set of deliverables documenting its outputs. From WP1, administrative deliverables include the **Project Handbook and Data Management Plan (D1.1)** at Month 3, followed by **Annual Progress and Financial Reports (D1.2, D1.3)** at Months 12 and 24, and the **Final Project Report (D1.4)** at Month 36. The foundational work in WP2 will yield the **Finalized Data Acquisition Protocol and full IRB Approval (D2.1)** at Month 6, and the crucial **Curated and Annotated Pilot Dataset (N=50) (D2.2)** at Month 9, which provides the ground-truth data for initial model training.

Core technological development is marked by key deliverables. WP3 will culminate in a **Comprehensive Report on the AI Model v1 Architecture and Validation Performance against ground-truth annotations (D3.1)** at Month 18. In parallel, WP4 will produce the **Complete Game Design Document (GDD) (D4.1)** at Month 6 and the **Final, Feature-Complete, Non-adaptive Version of the Serious Game (D4.2)** at Month 18. The integration phase, WP5, will produce the **System Integration API Specification Document (D5.1)** at Month 19 and a **Comprehensive Report on the Usability Testing Results (D5.2)** at Month 24. The refinement of the AI in WP6 will be documented in a **Report on the Final, Optimized AI Model v2 Performance, including real-time latency metrics (D6.1)** at Month 27.

The validation phase in WP7 will produce the **Pre-registered RCT Protocol and Final IRB Approval (D7.1)** at Month 26 and the **Final, Anonymized, and Annotated RCT Dataset (N=150) (D7.2)** at Month 35. The analysis in WP8 will result in the **Final Scientific Report on the RCT Results and their Interpretation (D8.1)** at Month 36. Finally, dissemination activities in WP9 will generate the **Public-facing Project Website (D9.1)** at Month 6, the **Submission of at least Three Peer-reviewed Manuscripts (D9.2)** by Month 36, and the **Public Release of the Anonymized Multimodal Dataset and Annotation Schema (D9.3)** at Month 36, fulfilling our commitment to open science.

In addition to these deliverables, strategic progress will be evaluated against five major milestones.

MS1: Pilot Data Collection and Annotation Complete (Month 9). This milestone signifies that the multimodal data acquisition pipeline is functional and the foundational, annotated dataset for initial model training has been ethically secured. Its achievement confirms the feasibility of our data collection and annotation strategy. *Verification:* Delivery and formal approval of D2.2 by the Steering Committee, confirming the quality and consistency of the pilot dataset and its annotations.

MS2: Standalone Components Complete and Ready for Integration (Month 18). This critical milestone marks the successful completion of the parallel development of the AI model and the game platform. It demonstrates that both components meet their specified performance and feature requirements, serving as the primary go/no-go decision point before committing resources to integration. *Verification:* Successful delivery and formal approval of D3.1 and D4.2 by the Steering Committee.

MS3: Fully Integrated System Validated for Usability (Month 24). This milestone demonstrates that the closed-loop system has been successfully integrated. It confirms the system is stable, the AI-driven feedback loop is functional with the required low latency for real-time adaptation, and the user experience is validated, rendering the system ready for the large-scale RCT. *Verification:* Successful delivery of D5.2 and a live demonstration of the fully integrated, adaptive system to the Steering Committee.

MS4: RCT Recruitment Complete (Month 30). This milestone indicates that the final validation phase is on schedule to achieve the required statistical power for robustly testing our primary hypotheses. It mitigates the significant risk of recruitment delays impacting the project's final timeline. *Verification:* A formal report from the WP7 lead confirming that 150 participants have been successfully recruited, consented, and randomized.

MS5: Final Project Results Analyzed and Disseminated (Month 36). This final milestone marks the successful completion of all scientific objectives. It signifies that the primary RCT results have been analyzed, key scientific conclusions have been drawn, and the project's major findings have been submitted for dissemination to the scientific community. *Verification:* Delivery of the final scientific report (D8.1) and documented proof of submission for the key manuscripts (D9.2).

Title: "Contingency Plan for Critical Points"

Contingency Plan for Critical Points

We have identified potential risks and critical points in the work plan and developed proactive contingency strategies to mitigate their impact, ensuring the project remains on track. Our approach is to anticipate challenges in AI performance, technical integration, user engagement, and participant recruitment, with clear alternative pathways prepared.

A primary scientific risk is that the cognitive bias detection model developed in WP3 and WP6 may not achieve the target accuracy (AUROC > 0.85) required for reliable real-time application. Our multimodal design provides inherent robustness; the model's attention mechanism is designed to down-weight less informative data streams and rely on stronger signals, such as language or behavior [2, 3]. If the deep learning architecture fails to generalize, we will pivot to more established models like Gradient Boosted Trees, which are often more robust with smaller datasets and have lower inference latency. As a final fallback, we will refine the classification task to focus on a high-confidence subset of the most reliably detectable biases. This strategy preserves the core principle of a process-specific adaptive loop, ensuring the intervention remains targeted and novel, rather than simplifying to a generic state.

A second critical point is the technical challenge of creating a stable, low-latency feedback loop between the game client and the AI inference server in WP5. The project plan de-risks this with a dedicated 6-month integration and testing phase. Our modular, client-server architecture allows for independent testing of the communication pipeline using a "mock" API before integrating the full AI model. If a continuous real-time loop proves unachievable, the system will transition to an 'event-triggered' adaptation. In this mode, the game adapts dynamically in response to discrete, high-stakes decision points within a scenario. This preserves the context-sensitive feedback loop by aligning adaptation with moments of high cognitive load, which still represents a significant advance over static systems and allows us to test our core hypotheses.

User engagement during the 4-week trial in WP7 is critical to the intervention's efficacy. A risk of high attrition (>20%) could compromise statistical power. We mitigate this through user-centered design and formal usability testing in WP5 for iterative improvement based on user feedback. The DDA mechanism itself is also designed to maintain optimal challenge, a key factor in engagement [1]. Should attrition rates approach our threshold, we will implement modest, performance-contingent financial incentives, a standard practice for bolstering adherence in digital trials [1].

Recruiting 150 eligible participants for the WP7 RCT may be challenging. Our diversified strategy includes university pools and partnerships with student wellness services, and the timeline includes a two-month recruitment buffer. If recruitment is slow, we will expand to nearby institutions and, if necessary, slightly broaden inclusion criteria in consultation with the ethical review board, without compromising internal validity.

Available Resources and Previous Results of the Team

The project team possesses the requisite expertise, a strong track record of relevant research, and access to all necessary resources and infrastructure to successfully execute this ambitious project. The Principal Investigators have a history of successful collaboration, having previously co-authored a highly-cited review on integrating AI and serious games in mental health [Vargas & Morales, *Lancet Digital Health*, 2020], which laid the conceptual groundwork for this proposal.

The team for Subproject 1, led by Dr. Elena Vargas, has made pioneering contributions to computational psychiatry, particularly in using Natural Language Processing to identify robust, quantitative biomarkers of mental illness. Their foundational work established the validity of using computational methods to probe the subtle cognitive alterations that underpin psychopathology. For instance, they demonstrated that different classes of language models can reveal a critical duality in semantic processing in psychosis, with static models showing a shrinking conceptual space while contextual models reveal a widening referential space—a finding that maps directly onto clinical symptoms of thought disorder [Vargas et al., *Schizophr. Bull.*, 2021]. This methodological insight is crucial for selecting analytical tools capable of capturing the subtle cognitive biases central to this project. The team has successfully scaled these methods for large-scale prediction of psychosis [Chen, Vargas, et al., *Mol. Psychiatry*, 2022] and shown that NLP-enriched models using clinical notes consistently outperform models relying solely on structured data for predicting adverse outcomes [Vargas et al., *Nature Mental Health*, 2023].

Crucially, their research has progressed beyond retrospective analysis to the longitudinal tracking of sub-clinical symptoms using automated speech markers [Vargas & Chen, *JAMA Psychiatry*, 2022], providing a direct precedent for the real-time analysis proposed in ADAPT-MIND. They have identified specific grammatical and semantic markers that differentiate individuals at clinical high risk for psychosis from controls and have demonstrated that linguistic changes are associated with progression to full-threshold psychosis. This body of work provides the essential algorithmic foundation for developing our cognitive bias detection model. It establishes the team's expertise in extracting subtle cognitive and affective states from complex, unstructured data—a core competency for Objective 1. The team’s most relevant publications are detailed in the attached CVs.

The team for Subproject 2, led by Dr. Javier Morales, is a leading group in the design and evaluation of serious games and adaptive systems for health. Their research focuses on creating engaging user experiences grounded in psychological theory and validated through rigorous empirical methods. They excel at translating therapeutic principles into effective and motivating game mechanics. For example, they recently developed an adaptive serious game for cognitive rehabilitation in stroke survivors, which used a sophisticated performance-based Dynamic Difficulty Adjustment algorithm to personalize training intensity. The results of the associated RCT, published in the *Journal of Medical Internet Research* [Morales et al., *JMIR*, 2022], showed significant improvements in executive function for the adaptive group compared to a non-adaptive control. This work showcases their expertise across the entire development lifecycle—from user-centered design and Unity development to conducting the efficacy trials proposed in WP7. Their practical experience in building and validating this type of system directly supports Objectives 2, 3, and 4. Dr. Morales's most relevant publications, including seminal work on adaptive systems in *CHI Proceedings*, are listed in the attached CVs.

Both teams have access to state-of-the-art facilities at their respective institutions, ensuring the project can commence immediately.

Dr. Vargas's laboratory (SP1) is fully equipped for the project's intensive AI development and data collection. This includes priority access to a university-managed High-Performance Computing (HPC) cluster featuring over 20 NVIDIA A100 GPUs. This computational power is essential for efficiently training the large-scale deep learning models proposed in WP3 and WP6. For data management, the lab has access to a secure, GDPR-compliant server infrastructure with multi-terabyte capacity, critical for handling the sensitive multimodal data to be collected. The physical lab space includes two sound-attenuated recording booths for human participant studies, which will be used in WP2 to guarantee the collection of high-quality audio data. The lab already possesses a suite of 10 Empatica E4 wearable biosensors and the associated software licenses, sufficient for pilot studies, as well as the high-fidelity RØDE microphones specified in the methodology.

Dr. Morales's laboratory (SP2) provides a complete ecosystem for serious game development and user experience research. The lab is equipped with 8 high-end developer workstations running professional licenses for the Unity 3D Pro engine, the full Adobe Creative Suite, and other essential development software. A dedicated usability lab, featuring a one-way mirror, multiple synchronized recording systems, and Tobii Pro eye-tracking equipment, will be utilized for the formal usability testing in WP5. This facility will allow for a deep, qualitative understanding of the user experience that complements quantitative data. The team includes experienced Unity developers, a 3D artist, and a UX designer who will be allocated to the project, ensuring all aspects of game development are handled by seasoned professionals.

Both institutions provide comprehensive administrative support for grant management, financial oversight, and legal counsel. Furthermore, their location within large university communities provides access to a reliable pool of young adults for participant recruitment. This comprehensive suite of computational and experimental infrastructure, specialized software, dedicated personnel, and robust institutional support provides a complete foundation for executing every task in the ADAPT-MIND work plan.

Relevant Previous Results of the Team

The project team's capacity to achieve the proposed objectives is grounded in a substantial body of prior research, with the two subproject teams providing a synergistic combination of world-leading expertise in computational psychiatry and the design of adaptive digital interventions. This provides a robust foundation for every aspect of the proposed work plan.

The team for Subproject 1, led by Dr. Vargas, has made pioneering contributions using Natural Language Processing (NLP) to identify quantitative, objective biomarkers of mental state.

Their foundational work established that different classes of language models can reveal a critical duality in the thought processes underlying psychosis. They demonstrated that while static models like FastText show a shrinking conceptual space, contextual models like BERT reveal a widening referential space—a finding that maps directly onto clinical symptoms of thought impoverishment and disorganization, respectively [2]. This deep methodological insight into how language provides a window into latent cognitive structures and biases is the cornerstone of our proposed approach. The team has a proven track record of translating these insights into practical, scalable applications. They have developed NLP-enriched models using clinical notes that significantly outperform models relying solely on structured data for predicting adverse outcomes such as suicide risk, achieving AUROCs as high as 0.915 [3]. Furthermore, their work within large-scale consortia has identified specific grammatical and semantic markers that differentiate individuals at clinical high risk for psychosis from controls [5] and has shown that these linguistic changes are associated with progression to psychosis [6]. Critically, their research has progressed beyond retrospective analysis to the longitudinal tracking of symptoms using automated speech markers [4], providing a direct precedent for the real-time analysis proposed in ADAPT-MIND. This extensive body of work demonstrates the team's unparalleled expertise in extracting meaningful psychological signals from complex language data, a core competency for Objective 1.

Complementing this AI expertise, the team for Subproject 2, led by Dr. Morales, is a leading group in the design, development, and empirical validation of serious games and adaptive systems for health. Their research is characterized by a deep understanding of how to translate therapeutic principles into effective and motivating game mechanics. Their work has consistently focused on creating systems that personalize user experience to maximize therapeutic outcomes. For example, they recently developed and trialed an adaptive serious game for cognitive rehabilitation in stroke survivors, which used a sophisticated performance-based Dynamic Difficulty Adjustment (DDA) algorithm to personalize training intensity. The results of the associated randomized controlled trial, published in a leading human-computer interaction journal, showed significant improvements in executive function for the adaptive group compared to a non-adaptive control [1]. This demonstrates their end-to-end expertise in the entire development lifecycle, from user-centered design and Unity development to conducting the specific efficacy trials proposed in WP7. Their practical experience in building and validating the very type of system proposed here is critical to the project's success and directly supports Objectives 2, 3, and 4.

The Principal Investigators have a history of successful collaboration, having previously co-authored a highly-cited review on integrating AI and serious games in mental health [1], which laid the conceptual groundwork for this proposal. This established working relationship and shared scientific vision ensure seamless integration between the subprojects, making this coordinated project a powerful and logical continuation of their prior work.

Available Infrastructure and Resources

The project team has immediate access to state-of-the-art facilities and resources across their institutions, precluding delays for major procurement. This established infrastructure provides a comprehensive foundation for all planned tasks, from large-scale computation to nuanced human-participant studies.

The laboratory of Dr. Vargas (SP1) is fully equipped for the project's intensive AI development and data collection. For the computationally demanding tasks of training and tuning deep learning models (WP3, WP6), the team has priority access to a university-managed High-

Performance Computing (HPC) cluster. This cluster features a dedicated partition with over 20 NVIDIA A100 GPUs, providing the necessary power to efficiently process the large multimodal datasets and complex architectures proposed. Data management is supported by a secure, GDPR-compliant server infrastructure with multi-terabyte capacity, essential for storing and processing the sensitive linguistic, physiological, and behavioral data. The physical lab space includes two sound-attenuated recording booths designed for human participant studies. These booths will be utilized in WP2 to guarantee the collection of high-quality, low-noise audio data—a prerequisite for accurately identifying the subtle linguistic markers of cognitive bias within our NLP pipeline. The lab is already equipped with a suite of 10 Empatica E4 wearable biosensors and associated software licenses, as well as the high-fidelity RØDE microphones specified in the methodology, enabling concurrent data collection as planned.

The laboratory of Dr. Morales (SP2) provides a complete ecosystem for professional-grade serious game development and user experience research. The lab is equipped with eight high-end developer workstations, each with professional licenses for the Unity 3D Pro engine, the full Adobe Creative Suite, and other essential software required for WP4. This setup ensures the team can produce a high-quality, polished application. For the critical usability testing in WP5, the team will utilize a dedicated usability lab. This facility is equipped with a one-way mirror, multiple synchronized camera and audio recording systems, and Tobii Pro eye-tracking equipment, enabling deep qualitative analysis of user engagement and interaction to complement quantitative metrics. The existing team includes experienced Unity developers, a 3D artist, and a UX designer who will be allocated to the project, ensuring professional-grade execution of all game development tasks from the project's inception.

Both institutions provide comprehensive administrative support for grant management, financial oversight, and legal counsel for intellectual property and data sharing agreements. Furthermore, their location within large university communities provides access to a reliable and well-characterized pool of young adults for participant recruitment (WP2, WP5, WP7), significantly mitigating potential delays. Collectively, this combination of cutting-edge computational and experimental infrastructure, specialized software, dedicated personnel, and robust institutional support confirms the project's readiness and capacity for the successful execution of the ADAPT-MIND work plan.

Title: "Expected impact of the results"

Expected impact of the results

The successful completion of the ADAPT-MIND project will generate a cascade of significant and far-reaching impacts, advancing the scientific frontier, delivering tangible societal and economic benefits, and creating a new paradigm for human-centered artificial intelligence. The project's primary scientific-technical impact will be the creation and validation of a fundamentally new class of digital therapeutic intervention. By moving beyond the current state of the art, which is limited to performance-based adaptation, our project will establish a new paradigm of process-based adaptation. This represents a conceptual leap from systems that ask "Is the user succeeding?" to a system that asks "What underlying cognitive process is driving the user's behavior, and how can we target it?". This shift has the potential to transform the field of digital mental health, paving the way for interventions that are not merely interactive but truly symbiotic, engaging in a personalized, psychologically-informed dialogue with the user. The development of our closed-loop system will provide the first empirical demonstration of how real-time inference of cognitive biases can be used to dynamically tailor an environment to foster resilience, a foundational contribution to both human-computer interaction and clinical psychology.

This paradigm shift is enabled by a series of core technological and methodological innovations. First, the project will make a major contribution to the field of computational psychiatry by moving AI from a passive, diagnostic tool to an active, interventional one. While prior research has powerfully demonstrated the potential of machine learning for the offline, retrospective analysis of clinical data [4], ADAPT-MIND will pioneer the application of these techniques in a real-time, interactive feedback loop. This requires solving significant technical challenges in low-latency model inference, data synchronization, and system integration, thereby creating a novel architectural blueprint for a new generation of responsive digital therapeutics. Second, we will advance the state of the art in multimodal machine learning for affective computing. Our approach of fusing linguistic, physiological, and behavioral data into a single, coherent model for real-time cognitive state inference is a significant methodological advance. By developing and validating a temporal attention mechanism that can dynamically weigh the importance of each modality, we will generate new knowledge about the complex, multimodal signatures of specific cognitive biases. This will result in a more robust and nuanced understanding of human cognition under stress than is possible with any single modality alone, contributing new methods and insights to the broader field of AI.

The ADAPT-MIND system itself will constitute a novel and powerful technological artifact. The integrated platform—comprising the serious game engine, the library of psychologically-grounded scenarios, the multimodal data acquisition pipeline, and the real-time AI inference engine—will be a unique and highly valuable research tool. It will provide a replicable, controllable, and ecologically valid experimental environment for studying human cognition and emotion, enabling a new wave of research into the mechanisms of cognitive training and emotional regulation. Furthermore, one of the most enduring scientific contributions of this project will be the creation of a large-scale, richly annotated, multimodal dataset. This resource, containing synchronized streams of linguistic, physiological, and behavioral data from a large cohort of individuals, will be meticulously curated and shared with the international research community according to FAIR principles. The current lack of such high-quality, ecologically valid data is a major bottleneck for progress in computational psychiatry and human-centered AI [5]. By making this dataset publicly available, we will catalyze future research, enabling other scientists to develop and validate new models, test alternative intervention strategies, and explore the complex interplay between mind, body, and behavior. This will amplify the project's impact far beyond its immediate objectives, fostering a more open and collaborative scientific ecosystem. The knowledge generated will be disseminated through high-impact publications in premier, peer-reviewed journals and conferences across our constituent fields, including top-tier venues in artificial intelligence (e.g., *NeurIPS*, *ICML*), human-computer interaction (e.g., *CHI Conference Proceedings*), and clinical and computational psychiatry (e.g., *The Lancet Digital Health*, *Nature Mental Health*).

Beyond these significant contributions to scientific knowledge, the ADAPT-MIND project is poised to generate significant long-term social and economic benefits. The primary societal impact will stem from the validation of a novel approach for enhancing cognitive resilience, which could lay the groundwork for a new generation of scalable and accessible digital mental health tools. This research directly addresses a critical public health need by exploring innovative methods for preventative mental healthcare. By establishing a proof-of-concept, our project will demonstrate a pathway toward preventative platforms that, after further development and clinical validation in target populations, could be delivered at scale, helping to mitigate the accessibility and cost barriers of traditional therapy. A validated ADAPT-MIND framework could inform the development of applications for various settings, including universities to support student well-being, corporations as part of employee wellness programs to combat burnout, and public health services as a potential low-cost, first-line intervention. Moreover, the principles demonstrated by this technology could be adapted for enhancing performance in high-stakes professions. First responders, military personnel, and healthcare workers could, in the future, use similar systems to train their cognitive and emotional regulation skills, potentially leading to better decision-making under stress and contributing to strategies for reducing the risk of burnout and PTSD [16, 17].

The project also has the potential for significant long-term economic impact. The core technology developed within ADAPT-MIND, if successfully validated, would represent a significant technical advance with clear potential for future commercialization. This could contribute to positioning Spain within the rapidly growing global market for digital therapeutics and mental wellness technologies. Following the project, the validated assets—such as the adaptive game engine and the cognitive bias detection models—could represent valuable intellectual property. Potential valorization pathways to be explored include licensing to established digital health companies or, if the results are particularly strong and a market opportunity is identified, investigating the feasibility of a university spin-off. A successful commercialization effort in the future could create high-value jobs in AI, software development, and clinical science. By demonstrating a viable pathway from cutting-edge research to a validated prototype, ADAPT-MIND will serve as a powerful case study for technology transfer in the digital health domain.

A core principle of this project is a commitment to equity and inclusivity, which is reflected in our explicit consideration of the gender dimension across all aspects of the research. In terms of research content, we acknowledge that stress, cognitive biases, and mental health conditions can manifest differently across genders, and that game preferences can also vary [1]. Therefore, the narrative, characters, and challenges within the serious game will be co-designed with a diverse user group to ensure they are inclusive, relatable, and free from gender stereotypes. Our AI models will be trained on a dataset with balanced gender representation, and we will conduct specific subgroup analyses to test for any performance disparities in the cognitive bias detection model, ensuring it is equally effective for all users. Methodologically, we will ensure gender balance in participant recruitment for all our studies (pilot, usability, and RCT) and will analyze our primary and secondary trial outcomes for any gender-specific effects. This will allow us to determine if the intervention is equally beneficial for men and women and to generate new knowledge on gender differences in response to digital resilience training. Finally, the project team itself reflects a commitment to gender equality, with Dr. Elena Vargas serving as the coordinating PI, providing a strong example of female leadership in a high-tech, interdisciplinary field. This multifaceted approach ensures that the benefits of our research will be equitable and that our findings will be generalizable and sensitive to the diverse population we aim to serve.

To ensure the project's scientific, social, and economic impacts are fully realized, we have developed a comprehensive and strategic plan for communication, dissemination, and valorization. This plan is designed to engage a wide range of audiences—from scientific peers to the general public and industrial partners—through targeted activities and channels throughout the project's lifecycle. Our strategy is built on the principles of open science, proactive engagement, and a clear pathway from research to real-world application.

Our scientific communication plan is designed to maximize the visibility and influence of our findings within the international research community. We will target premier, high-impact, peer-reviewed journals and conferences across the project's interdisciplinary domains. Methodological and technological breakthroughs related to the AI model and system architecture (from WP3, WP5, and WP6) will be submitted to leading venues in artificial intelligence and computer science, such as the *Conference on Neural Information Processing Systems (NeurIPS)*, the *International Conference on Machine Learning (ICML)*, and journals like *IEEE Transactions on Affective Computing*. Findings related to the design of the serious game and the results of the usability studies (from WP4 and WP5) will be targeted at top-tier human-computer interaction conferences and journals, including the *ACM CHI Conference on Human Factors in Computing Systems* and *ACM Transactions on Computer-Human Interaction*. The definitive results of the randomized controlled trial (from WP7 and WP8) will be submitted to leading journals in digital medicine and clinical and computational psychiatry, such as *The Lancet Digital Health*, *Nature Mental Health*, or *JAMA Psychiatry*, to ensure our findings reach a clinical and public health audience. We are fully committed to the principles of open science. All peer-reviewed publications resulting from this project will be made available via Gold or Green Open Access routes, in full compliance with the funder's mandates, ensuring that our results are freely accessible to all. Furthermore, we will actively share pre-prints of our manuscripts on platforms like arXiv and medRxiv to accelerate the dissemination of our findings.

Our dissemination and public outreach plan aims to translate our complex research into accessible and engaging content for a broader audience, including the general public, policymakers, healthcare professionals, and potential end-users. A central hub for these activities will be a dedicated project website (to be launched in M6), which will feature public-facing summaries of our objectives and findings, team profiles, news updates, and links to publications. We will leverage social media platforms, particularly Twitter and LinkedIn, to provide regular updates on project milestones, share relevant news, and engage in dialogue with stakeholders. To reach a wider public audience, we will coordinate with the press offices of our institutions to issue press releases at key moments, such as the project's launch, the publication of major findings, and the release of the public dataset. We will also actively participate in public engagement events, such as national science festivals and university open days, to present our work through interactive demonstrations and accessible talks. For policymakers and healthcare administrators, we will produce concise policy briefs summarizing the potential of our technology to address public health challenges, aiming to inform future strategies for digital mental health provision. Finally, we will engage directly with clinical communities by presenting our findings at workshops and seminars for psychologists, psychiatrists, and primary care providers, facilitating the translation of our research into clinical awareness and practice.

The knowledge transfer and valorization plan outlines our strategy for translating the project's technological outputs into tangible economic and social value. This plan is proactive, beginning with the careful management of intellectual property (IP). Throughout the project, we will conduct regular IP audits to identify patentable inventions and protectable assets. The core adaptive algorithm and the specific architecture of the multimodal fusion model are likely candidates for patent protection, while the trained models and the game's source code may be protected as trade secrets. We will perform a freedom-to-operate analysis early in the project to ensure our work does not infringe on existing patents. Our valorization strategy will focus on assessing the most effective pathways for real-world application. We will actively explore two primary avenues. The first is licensing the ADAPT-MIND technology to established companies in the digital health, corporate wellness, or professional training sectors. We will actively engage with potential industry partners throughout the project, inviting them to advisory board meetings to ensure our development is aligned with market needs. The second avenue is to investigate the feasibility of creating a university spin-off. This investigation will be contingent on the project's results and a preliminary market analysis. If this path appears viable, we will develop a foundational business case and identify potential next steps for seeking seed funding, rather than pursuing a fully-formed business plan within the scope of this research project. This exploratory approach provides flexibility and maximizes the chances of identifying a viable path for future real-world implementation, ensuring that the benefits of this research reach the widest possible audience and contribute to economic growth.

A cornerstone of our commitment to open science and the long-term impact of this project is a robust Data Management Plan (DMP), summarized here. The full DMP will be delivered as D1.1. Our plan is designed to ensure that all research data generated by the project are managed, documented, shared, and preserved according to the FAIR principles: Findable, Accessible, Interoperable, and Reusable. The project will generate several types of data, including raw multimodal data (BVP, EDA, audio, behavioral logs), preprocessed feature data, derived data (e.g., model outputs, statistical results), and clinical and demographic data from participants. To make our data **Findable**, the final, fully anonymized, and documented dataset from the RCT will be deposited in a reputable, trustworthy public repository, such as Zenodo or PhysioNet. The dataset will be assigned a persistent Digital Object Identifier (DOI), ensuring it is permanently citable and discoverable. It will be described with rich metadata, compliant with the Dublin Core standard, detailing the study context, data collection methods, and variable descriptions.

To ensure the data are **Accessible**, we will implement a tiered access policy. During the project, all data will be stored on secure, encrypted, and access-controlled university servers to protect participant confidentiality and comply with GDPR. Access will be restricted to named project personnel. Following the project's conclusion, and after an embargo period of 12 months to allow the team to publish its primary findings, the fully anonymized RCT dataset will be made publicly available under a permissive Creative Commons license (CC-BY 4.0). This will allow other researchers to freely access and reuse the data for any purpose, provided they give appropriate credit. To make the data **Interoperable**, we will use standard, open formats (e.g., CSV for tabular data, BIDS for physiological data, JSON for logs) and provide comprehensive documentation. This will include a detailed data dictionary, or codebook, explaining every variable, the experimental code used for data collection, and the scripts used for preprocessing and analysis. This will ensure that other researchers can understand and effectively integrate our data with other datasets.

Finally, to maximize **Reusability**, the combination of a persistent identifier, rich metadata, open formats, detailed documentation, and a permissive license will ensure that the dataset is a valuable and lasting resource for the scientific community. All procedures for data sharing will be conducted in strict accordance with the informed consent provided by participants and the protocols approved by our Institutional Review Board. Our robust anonymization procedure will involve removing all direct identifiers and quasi-identifiers to minimize any risk of re-identification, thereby upholding the highest ethical standards while maximizing the scientific value of the data we produce.

Scientific-Technical Impact

The ADAPT-MIND project will generate significant scientific and technical impact by creating and validating a new class of digital therapeutic intervention based on process-based adaptation. This approach moves beyond current performance-based systems, which adapt difficulty based on simple metrics of success or failure (Jones & Smith, 2021), to a model that targets the underlying cognitive processes driving user behavior. By inferring these cognitive processes in real-time and responding with specific, psychologically-informed challenges, we will provide the first empirical demonstration of a symbiotic human-AI training loop. This will shift the scientific discourse from *whether* adaptation is effective to *how* and *why* specific adaptive mechanisms work, establishing a more rigorous standard for designing and evaluating intelligent therapeutic systems. Consequently, the project will deliver not just a novel technology but also a new theoretical framework for personalized digital mental health, with the potential to reshape how interventions are conceived, implemented, and validated across the field.

This paradigm shift is enabled by core methodological and technological innovations. The project will pioneer the application of sophisticated machine learning models for real-time intervention, moving computational psychiatry beyond its current focus on offline, retrospective analysis of clinical data (Chen et al., 2022). This shift from passive observation to active intervention will enable novel research into the causal effects of targeted micro-interventions on cognitive states. Solving the technical challenges of low-latency multimodal data fusion and model optimization will produce a novel architectural blueprint for responsive, human-in-the-loop AI systems. This blueprint will have applications far beyond mental health, informing the development of intelligent tutoring systems, adaptive workplaces, and other domains where technology must understand and respond to complex human states. Furthermore, our approach to multimodal fusion represents a key methodological advance. We will develop and validate a computational framework, incorporating temporal attention mechanisms, to reliably infer cognitive states from dynamically weighted linguistic, physiological, and behavioral data streams. This will generate fundamental knowledge about the multimodal signatures of specific cognitive biases, providing a more nuanced understanding of cognition under stress than is possible with any single modality and contributing new methods to affective computing and cognitive science.

The ADAPT-MIND system itself will constitute a novel technological artifact, serving as a catalyst for future research. The integrated platform—comprising the serious game engine, a library of psychologically-grounded scenarios, a multimodal data acquisition pipeline, and the real-time AI inference engine—will be a unique research tool. It will provide the scientific community with a replicable and ecologically valid experimental environment to investigate the causal links between cognitive processes, emotional states, and behavioral outcomes, overcoming the limitations of static laboratory paradigms. The project's most enduring contribution, however, will be a large-scale, richly annotated, multimodal dataset. The current lack of such high-quality, ecologically valid data is a primary bottleneck for progress in computational psychiatry and human-centered AI (Miller, 2023). This resource, containing synchronized linguistic, physiological, and behavioral data from a large cohort engaged in cognitively demanding tasks, will be meticulously curated with labels for both objective task performance and the computationally inferred cognitive biases. Shared with the international research community according to FAIR principles, this foundational public resource will catalyze future research, enabling others to develop and benchmark new models, test intervention strategies, and explore the complex interplay between mind, body, and behavior. This act of creating a foundational resource will amplify the project's impact far beyond its immediate objectives, fostering a more open, collaborative, and accelerated scientific ecosystem.

Finally, a strategic dissemination plan will ensure and amplify the project's scientific-technical impact. The knowledge generated will be published in premier, peer-reviewed journals and presented at leading international conferences. Methodological and technological breakthroughs related to the AI model and system architecture will be targeted at top-tier venues in artificial intelligence and computer science, such as the *Conference on Neural Information Processing Systems (NeurIPS)* and *IEEE Transactions on Affective Computing*. Findings on the serious game design and usability studies will be submitted to leading human-computer interaction conferences like the *ACM CHI Conference on Human Factors in Computing Systems*. The definitive results of the randomized controlled trial, providing the highest level of evidence for our new intervention paradigm, will be submitted to leading journals in digital medicine and clinical psychiatry, such as *The Lancet Digital Health* or *Nature Mental Health*. This multi-pronged strategy ensures our findings will reach the diverse scientific communities our work integrates, maximizing visibility and shaping future research agendas across these disciplines.

Title: "Social and Economic Impact"

Social and Economic Impact

Beyond its scientific contributions, the ADAPT-MIND project is designed to generate substantial social and economic benefits by addressing a pressing public health challenge. Stress-related disorders, including anxiety and burnout, impose a considerable burden on society, resulting in diminished quality of life, increased healthcare expenditure, and significant losses in economic productivity [16–18]. While digital mental health interventions have been proposed as a scalable solution, their real-world impact has been constrained by low user engagement and a failure to personalize treatment—a critical gap our project is designed to fill [1, 3, 5]. The primary societal impact will be a validated, scalable, and effective tool for enhancing cognitive resilience at a population level. By moving beyond the one-size-fits-all model, our AI-driven adaptive system offers a tangible pathway towards accessible, preventative mental healthcare. This technology is designed for deployment in various settings to support at-risk populations. For instance, integration into university systems could help students manage academic pressure, potentially reducing dropout rates and improving well-being. In corporate environments, it could be integrated into employee wellness programs to combat burnout and enhance productivity. Furthermore, the system has direct applications for enhancing performance and mitigating the psychological toll in high-stakes professions. First responders, military personnel, and healthcare workers, who operate under extreme pressure, could use the system to train cognitive and emotional regulation skills, leading to better decision-making under stress and a lower risk of developing conditions like PTSD [16, 17].

The project's economic impact is equally notable and multifaceted. At a macro level, by providing an effective tool for preventing stress-related illness, ADAPT-MIND has the potential to generate substantial cost savings for public health systems and employers. The economic burden of mental ill-health in Spain is estimated to be billions of euros annually, factoring in direct healthcare costs and indirect costs from lost productivity [16–18]. An accessible and engaging digital tool that reduces the incidence of these conditions would therefore yield a high return on investment. Beyond cost savings, the project is designed to stimulate economic growth. The technology developed within ADAPT-MIND has clear commercial potential, positioning Spain as a leader in the rapidly growing global market for digital therapeutics. The core assets—the adaptive game engine and the proprietary cognitive bias detection models—could be licensed to established digital health companies or form the basis of a new university spin-off. Such a venture would create high-value jobs in artificial intelligence, software development, and clinical science, contributing directly to the national strategy of fostering a knowledge-based economy and leveraging AI for industrial competitiveness. By demonstrating a viable pathway from cutting-edge research to a market-ready prototype, ADAPT-MIND will serve as a powerful example of successful technology transfer.

To achieve broad societal impact, our research and technology will be fundamentally equitable and inclusive. Consequently, we have integrated the gender dimension throughout the project's design and execution. We acknowledge that stress, cognitive biases, and mental health conditions can manifest differently across genders, and that preferences for game-based interactions can also vary [1]. The serious game's narrative, characters, and challenges will be co-designed with a diverse user group to ensure they are inclusive, relatable, and free from gender stereotypes. This process will ensure the intervention is engaging and acceptable to a broad audience. Methodologically, we will ensure gender balance in participant recruitment for all our studies (pilot, usability, and RCT). This is not merely a demographic checkbox; we will pre-register and conduct specific subgroup analyses to test for any gender-specific effects in our primary and secondary trial outcomes. This approach will determine if the intervention is equally beneficial across genders and will generate valuable knowledge on differential responses to digital resilience training.

This commitment to equity extends to the core technology. A critical risk in AI development is the creation of biased algorithms that perform differently for different demographic groups [5]. We will proactively mitigate this risk by training our AI models on a dataset with balanced gender representation and, crucially, conducting specific fairness audits to test for any performance disparities in the cognitive bias detection model. This ensures that the system's adaptive capabilities are equally effective and reliable for all users, preventing the perpetuation of systemic biases. Finally, the project team itself reflects a commitment to gender equality. With Dr. Elena Vargas serving as the coordinating PI, the project provides a strong example of female leadership in a high-tech, interdisciplinary field, contributing to a more inclusive research culture and providing mentorship for the next generation of female scientists in AI and computational science. This multifaceted approach ensures our research benefits will be equitable and our findings generalizable to the diverse population we aim to serve.

Contribution to Society and Economy

The outcomes of this project will generate substantial and lasting social and economic benefits by directly addressing the pressing public health challenge of stress-related disorders. The societal burden of these conditions—measured in diminished quality of life, increased healthcare expenditure, and significant losses in economic productivity—is immense [16–18]. Current digital mental health interventions, while scalable, have shown limited real-world impact due to low user engagement and a failure to personalize treatment, a critical gap our project is designed to fill [1, 3, 5]. The primary societal contribution of ADAPT-MIND will be a validated, scalable, and highly effective tool for enhancing cognitive resilience at a population level. By moving beyond the one-size-fits-all model, our AI-driven adaptive system offers a tangible pathway to providing accessible, preventative mental healthcare that can inform public policy. The technology can be deployed across multiple sectors to support at-risk populations: within university systems to help students manage academic pressure and improve well-being; in corporate environments as part of employee wellness programs to combat burnout; and for high-stakes professions, such as first responders and healthcare workers, to train emotional regulation skills, improve decision-making under stress, and reduce the risk of developing conditions like PTSD [16, 17]. This provides a new model for digital public health, demonstrating how AI can deliver personalized care at scale.

The economic impact is equally significant. At a macro level, by providing an effective tool for preventing stress-related illness, ADAPT-MIND has the potential to generate substantial cost savings for public health systems and employers. Given that the economic burden of mental ill-health in Spain is estimated to be billions of euros annually from direct healthcare costs and lost productivity [16–18], an accessible digital tool that reduces the incidence of these conditions offers a significant return on investment. Beyond these savings, the project is designed to stimulate economic growth and industrial competitiveness. The technology developed has clear commercial potential in the global digital therapeutics market, a sector projected to exceed $20 billion by 2028 [3]. The core assets—the adaptive game engine and the proprietary cognitive bias detection models—are prime candidates for valorization through either licensing to established digital health companies or the creation of a new university spin-off. Such a venture would create high-value jobs in artificial intelligence, software development, and clinical science, contributing directly to the Spanish National AI Strategy's goal of fostering a knowledge-based economy. By demonstrating a viable pathway from cutting-edge research to a market-ready prototype, ADAPT-MIND will serve as a powerful example of successful technology transfer, converting public investment into both societal well-being and economic prosperity.

Consideration of the Gender Dimension

This project integrates the gender dimension as a core component of its scientific design to ensure our findings are equitable and broadly generalizable. Our approach addresses gender in the research content, methodology, and team composition. Recognizing that cognitive biases, stress responses, and preferences for game-based interactions can vary across genders, the serious game's narrative, characters, and challenges will be developed through a co-design process with diverse user groups. This participatory method will ensure the game is inclusive and relatable, avoiding gender stereotypes to maximize engagement and acceptability for all participants.

Methodologically, our research design mandates gender balance in participant recruitment for all studies, from pilot and usability testing to the final randomized controlled trial (RCT). Gender will be included as a pre-specified factor in our statistical analysis plan, enabling formal subgroup analyses to test for differential effects of the intervention. This will allow us to determine if the training is equally beneficial across genders and generate new knowledge on gender-specific responses to digital resilience training. Critically, our AI development pipeline includes rigorous fairness audits of the cognitive bias detection model. We will systematically evaluate key performance metrics across gender subgroups to identify and mitigate potential algorithmic bias, ensuring the system's adaptive capabilities are equally effective and reliable for all users.

The project team and dissemination strategy also reflect our commitment to gender equality. With Dr. Elena Vargas as the coordinating PI, the project exemplifies female leadership in a high-tech, interdisciplinary field. We are committed to maintaining gender balance in hiring for new positions and fostering an inclusive research environment. Furthermore, all scientific dissemination will feature transparent reporting of our gender-specific analyses, contributing to a more nuanced and complete understanding of gender in digital mental health.

Communication, Dissemination, and Valorization Plan

To ensure the project's scientific, social, and economic impacts are fully realized, we have developed a comprehensive plan for communication, dissemination, and valorization. This strategy is designed to engage diverse audiences—from scientific peers to the public and industrial partners—through targeted activities throughout the project's lifecycle. Our approach is founded on the principles of open science, proactive engagement, and a clear pathway from research to real-world application, thereby maximizing the return on public investment.

Our scientific communication plan aims to maximize the visibility and influence of our findings within the international research community. We will target premier, high-impact, peer-reviewed journals and conferences across the project's interdisciplinary domains. Methodological and technological breakthroughs from Work Packages 3, 5, and 6, particularly concerning the AI model and system architecture, will be submitted to leading artificial intelligence venues. These include top-tier conferences like the *Conference on Neural Information Processing Systems (NeurIPS)* and the *International Conference on Machine Learning (ICML)*, and prestigious journals such as *IEEE Transactions on Affective Computing*, ensuring our core technical innovations are validated by leading experts. Findings from WP4 and WP5 on the serious game's user-centered design and usability will be directed at top human-computer interaction venues, including the *ACM CHI Conference on Human Factors in Computing Systems* and *ACM Transactions on Computer-Human Interaction*. The definitive results of the randomized controlled trial (WP7 and WP8), which provide the highest level of evidence for our intervention, will be submitted to leading journals in digital medicine and computational psychiatry. Our primary targets include *The Lancet Digital Health*, *Nature Mental Health*, and *JAMA Psychiatry*, ensuring our findings reach a clinical and public health audience to influence future practice. In line with our commitment to open science, all peer-reviewed publications will be made available via Gold or Green Open Access routes. To accelerate dissemination, pre-prints will be shared on platforms like arXiv and medRxiv. Furthermore, all project-generated source code and curated, anonymized datasets will be deposited in public repositories (e.g., GitHub, Zenodo) with permissive licenses to foster reproducibility and further innovation.

Parallel to scientific communication, our public outreach plan will translate complex research into accessible content for the general public, policymakers, healthcare professionals, and end-users. We will launch a dedicated project website (D9.1, Month 6) to serve as a dynamic public portal, featuring clear summaries of our objectives and findings, team profiles, news, and links to all public resources. We will use social media platforms, particularly Twitter and LinkedIn, to provide regular, accessible updates on project milestones and engage in dialogue with stakeholders. In coordination with our institutions' press offices, we will issue press releases at key moments, such as the project's launch and the publication of major RCT findings. We will also participate in public engagement events like national science festivals, presenting our work through interactive demonstrations of the serious game and accessible talks on its societal importance. For policymakers and healthcare administrators, we will produce concise, evidence-based policy briefs summarizing our technology's potential to address mental well-being challenges and providing actionable recommendations for integrating digital therapeutics into public health strategies. Finally, we will engage directly with clinical communities through workshops and seminars for psychologists and psychiatrists. This direct engagement is crucial for translating our research into clinical awareness and practice, ensuring that future professional users are part of the conversation from an early stage.

Our knowledge transfer and valorization plan outlines a proactive strategy for translating the project's outputs into tangible economic and social value. This plan is an integrated component of our project management, beginning with diligent stewardship of intellectual property (IP). Supported by our institutions' technology transfer offices (TTOs), we will conduct regular IP audits to identify and protect patentable inventions. Key innovations, such as the core adaptive algorithm and the multimodal fusion model architecture, will be assessed for patentability, while assets like the trained models and game source code may be protected as trade secrets. An early-stage freedom-to-operate analysis will ensure our work avoids infringement and clarifies the landscape for commercialization. Our primary valorization strategy is twofold to maximize the potential for real-world implementation. First, we will explore licensing the ADAPT-MIND technology to established companies in the digital health, corporate wellness, or professional training sectors. We will actively engage potential industry partners, inviting them to an annual advisory board meeting to provide feedback and ensure our development aligns with market needs. Second, we will concurrently evaluate the creation of a university spin-off to commercialize the ADAPT-MIND system. Towards the project's conclusion (WP9), we will dedicate resources to developing a detailed business plan, including market analysis, a financial model, and a go-to-market strategy. We will leverage institutional and national innovation ecosystems to seek seed funding from public and private sources. This dual strategy provides a clear pathway to translate our research into a tangible product that can improve public health, create high-value jobs, and contribute to economic growth.

Title: "Scientific Communication Plan"

Scientific Communication Plan

Our scientific communication strategy is designed for the rapid and wide dissemination of our findings to the international research community, maximizing their visibility, impact, and scientific validation. We are committed to open science principles and will target premier, high-impact, peer-reviewed venues across the project's interdisciplinary domains. This multi-pronged approach ensures our distinct contributions—in artificial intelligence, human-computer interaction, and clinical science—reach and influence the leading experts in each respective field.

We will disseminate the core technological and methodological innovations from WP3 and WP6 to the artificial intelligence and computational science communities. Our work on the novel multimodal fusion model and real-time inference architecture will be submitted to top-tier machine learning conferences, such as the *Conference on Neural Information Processing Systems (NeurIPS)*, the *International Conference on Machine Learning (ICML)*, and the *Annual Meeting of the Association for Computational Linguistics (ACL)*. These venues provide a forum for rigorous peer review and high-visibility presentation of cutting-edge algorithmic work. Corresponding journal articles will target prestigious publications like *IEEE Transactions on Affective Computing* and the *Journal of Machine Learning Research*.

We will communicate findings from WP4 and WP5 related to the serious game's design, the user-centered process, and usability study results to the human-computer interaction (HCI) community. The primary venues will be leading international conferences, including the *ACM CHI Conference on Human Factors in Computing Systems* and the *ACM International Conference on Intelligent User Interfaces (IUI)*. These conferences are the most influential forums for disseminating research on novel interactive systems and user experience design. For detailed archival publications, we will target high-impact journals such as *ACM Transactions on Computer-Human Interaction* and *JMIR Serious Games*.

The definitive results of the randomized controlled trial (from WP7 and WP8), which provide the highest level of evidence for our new intervention paradigm, will be submitted to leading journals in digital medicine, clinical psychology, and computational psychiatry. Our primary targets for this capstone publication include *The Lancet Digital Health*, *Nature Mental Health*, *JAMA Psychiatry*, and the *Journal of Medical Internet Research (JMIR)*. Publishing in these top-tier clinical and health informatics journals is essential to ensure our findings reach a clinical and public health audience, where they can inform future practice and policy.

In line with our commitment to open science, all peer-reviewed publications from this project will be made available via Gold or Green Open Access routes, in full compliance with funder mandates. To accelerate dissemination, we will share pre-prints on platforms such as arXiv and medRxiv prior to peer review. Furthermore, all publications will include a data availability statement linking to the final, publicly accessible dataset (Objective 5), ensuring our work is transparent, reproducible, and serves as a lasting resource. Beyond formal publications, all team members, including junior researchers, will be supported to present our work at international conferences and workshops to foster dialogue and collaboration. This comprehensive strategy ensures our project's findings are rigorously validated, integrated into the scientific record, and serve as a robust foundation for future research.

Dissemination and Public Outreach Plan

Complementing our scientific communication strategy, this plan outlines how we will translate the project's findings into accessible and impactful content for diverse non-academic audiences. Our primary goal is to maximize the project's societal return by fostering public understanding of cognitive resilience and human-centered AI, engaging with key stakeholders to ensure real-world relevance, and informing public policy. Our approach is proactive and multifaceted, integrating digital platforms, media engagement, and direct community interaction.

The central hub for public-facing activities will be a dedicated project website (D9.1). This dynamic portal will feature lay-language summaries of our research, blog posts explaining core concepts like cognitive bias and neuroplasticity, and researcher profiles to humanize the scientific process. We will also produce short, animated videos to explain the project's core innovation—the closed-loop adaptive system—and showcase demonstrations of the serious game. This digital presence will be amplified through a targeted social media strategy on platforms like Twitter and LinkedIn, where we will share project milestones, data visualizations, and engage in dialogue with journalists, patient advocacy groups, and the public.

We will pursue proactive media engagement by coordinating with the press offices of our partner institutions. Strategic press releases will be issued at key project milestones, such as the official launch, the publication of the main RCT results, and the public release of our unique multimodal dataset. We will target science journalists and specialized media outlets covering health, technology, and gaming to secure feature articles and interviews. This strategy aims to embed our findings into the public discourse, raising awareness of AI-driven tools for mental well-being.

Direct engagement is a cornerstone of our plan. We will participate in national science events, like the Spanish "Semana de la Ciencia," and university open days, offering interactive demonstrations of the ADAPT-MIND prototype. These events facilitate two-way communication, allowing us to share our work and gather feedback from potential end-users. Furthermore, we will organize targeted stakeholder workshops. These will include sessions for clinicians (psychologists, psychiatrists) to discuss the system's role in clinical pathways and gather expert feedback. A separate stream of workshops will engage university wellness professionals and corporate HR managers to explore applications in educational and workplace settings.

Finally, to ensure our research informs policy, we will produce concise, evidence-based policy briefs for policymakers, public health administrators, and governmental bodies overseeing national AI and mental health strategies. These briefs will summarize the project's findings on the efficacy and cost-effectiveness of our intervention, providing actionable recommendations for integrating personalized digital therapeutics into public health services. This integrated strategy ensures that ADAPT-MIND's contributions extend beyond academia to create a lasting and meaningful impact on society.

Knowledge Transfer and Valorization Plan

Our knowledge transfer and valorization strategy is designed to translate the project's outputs into tangible products and services with lasting economic and social value. This plan moves beyond academic dissemination to establish a concrete pathway for the real-world application of the ADAPT-MIND system. The cornerstone of this strategy is the systematic management of intellectual property (IP), conducted in close collaboration with our institutions' Technology Transfer Offices (TTOs) from the project's inception. We will perform regular IP audits to identify and protect key assets. The primary protectable assets are strong candidates for patent protection and include: the proprietary methods for real-time cognitive bias detection, the novel adaptive algorithm, and the architecture of the multimodal fusion model. Other assets, such as the trained models, game source code, and unique art assets, will be protected through a combination of copyright and trade secrets. A freedom-to-operate analysis will be performed early in the project to de-risk future commercialization.

With a robust IP portfolio secured, we will pursue a dual-pathway valorization strategy to maximize market penetration and impact. The first pathway is **technology licensing**. We will target established companies in sectors where our technology can provide a significant competitive advantage. These include the digital health and therapeutics market, the corporate wellness sector, and specialized training organizations for high-stress professions (e.g., defense, emergency services). We will actively engage with potential industry partners throughout the project, leveraging our advisory board to ensure our development is aligned with market needs and to foster relationships that can lead to future licensing agreements. This approach offers a rapid route to market by leveraging the existing infrastructure and user base of an established partner.

The second, more ambitious pathway is the **creation of a university spin-off company**. This route offers the potential for greater long-term impact and economic return by building a dedicated enterprise around the ADAPT-MIND technology. Towards the end of the project, as part of WP9, we will dedicate resources to developing a comprehensive business plan. This plan will include a detailed market analysis, a scalable business model, a product roadmap for future development, and a robust financial model. We will leverage our institutional and national innovation ecosystems to seek seed funding from public and private sources, including venture capital firms specializing in health-tech and digital therapeutics. This dual strategy is not mutually exclusive; insights gained from engagement with potential licensees will directly inform the business plan for the spin-off, creating a synergistic approach. This proactive, dual-pathway plan ensures that public investment is converted into improved mental resilience and sustainable economic growth, positioning our team to lead in the emerging field of AI-driven cognitive training.

Summary of the Data Management Plan (DMP)

A robust Data Management Plan (DMP), to be fully detailed in deliverable D1.1, underpins our commitment to open science and the long-term impact of this project. Our plan ensures all research data will be managed, documented, shared, and preserved according to the FAIR principles (Findable, Accessible, Interoperable, Reusable). The project will generate a complex, synchronized multimodal dataset, comprising raw physiological signals (BVP, EDA), linguistic data (audio, transcripts), granular behavioral logs from gameplay, and derived clinical and demographic data.

For **Interoperability**, we will adopt community-accepted standards and open formats. Physiological data will be structured following the Brain Imaging Data Structure (BIDS) standard and its relevant extensions (e.g., BIDS-Physio), while tabular data (clinical scores, derived features) will be stored in CSV format and behavioral logs in JSON. This standardization facilitates data reuse and integration. During the project, all data will be stored on secure, encrypted, and access-controlled university servers in full compliance with GDPR. To protect participant confidentiality, personally identifiable information will be stored separately from research data, linked only by a unique participant ID, restricting access to authorized project personnel.

To ensure data are **Findable** and **Reusable**, the final, fully anonymized, and documented dataset from the RCT will be deposited in a reputable public repository, such as Zenodo or PhysioNet. The dataset will be assigned a persistent Digital Object Identifier (DOI) for permanent citation and discovery. It will be described with rich metadata, compliant with the Dublin Core standard, detailing the study context, data collection protocols, and variable definitions. This documentation will include a comprehensive data dictionary, the experimental code used for data collection, and the version-controlled scripts for preprocessing and analysis.

Following the project's conclusion and a 12-month embargo period for primary publication, the fully anonymized RCT dataset will be made publicly **Accessible** under a Creative Commons CC-BY 4.0 license. This will permit other researchers to freely access and reuse the data, provided appropriate credit is given. Our robust anonymization procedure, which will be approved by our Institutional Review Board, involves removing all direct and quasi-identifiers to minimize re-identification risk. This strategy balances rigorous ethical oversight with our commitment to maximizing the data's scientific value, creating a lasting resource to catalyze future research in computational psychiatry and human-centered AI.

Title: "Justification of the requested budget"

Justification of the requested budget

The financial resources requested are essential for the successful execution of the work plan. The budget has been meticulously prepared to ensure every cost is reasonable, directly allocable to a project task, and necessary for achieving our objectives. The allocation reflects the project's structure, with a primary investment in the specialized personnel who form the intellectual core of this research, complemented by targeted expenditures on equipment, materials, and activities required to develop and validate our proposed therapeutic tool. A detailed breakdown is provided in the supplementary materials {Detailed_Budget_Table}.

The largest component of the budget is allocated to personnel. The project's success hinges on a dedicated, interdisciplinary team with expertise spanning AI, human-computer interaction

(HCI), and clinical psychology. The requested effort for each member has been calibrated to the tasks outlined in the Work Packages (WPs), ensuring the team is sized to meet the project's demands without redundancy.

For Subproject 1, led by Dr. Vargas, we request funding for three positions. A full-time Postdoctoral Researcher (36 months, 100% dedication) with expertise in machine learning and multimodal signal processing is essential. This researcher will lead the technical development of the model for inferring cognitive bias from behavioral and physiological signals (WP3), implement optimization techniques for real-time inference (WP6), and support the system integration API (WP5). Their technical expertise is critical for achieving Specific Objective 1. We also request a full-time PhD Student (36 months, 100% dedication) specializing in computational linguistics. This student will develop the linguistic data processing pipeline, from audio preprocessing to the extraction of semantic and syntactic features indicative of cognitive bias (WP3). Their doctoral thesis will focus on the multimodal linguistic markers of these internal states, representing a significant training opportunity. Finally, a part-time Research Technician (36 months, 50% dedication) is required to manage the project's data infrastructure, overseeing data collection protocols (WP2, WP7), managing High-Performance Computing (HPC) cluster jobs, and curating the final dataset for public release (WP9).

For Subproject 2, led by Dr. Morales, we request funding for a complementary team. A full-time Postdoctoral Researcher (36 months, 100% dedication) with a background in HCI and serious game development will lead the design of the ADAPT-MIND game platform (WP4). Crucially, this individual will implement the logic for the three Dynamic Difficulty Adjustment arms, including the novel arm that adapts based on inferred cognitive states rather than performance metrics alone (WP5), and manage the technical deployment for the RCT (WP7). Their expertise is indispensable for Specific Objectives 2, 3, and 4. A full-time PhD Student (36 months, 100% dedication) in psychology and user experience (UX) research will lead user-centered design activities, conduct usability testing (WP5), manage participant recruitment and data collection for the RCT (WP7), and analyze qualitative data. Their work will ensure the final system is usable and psychologically valid. To ensure clinical integrity, we request a part-time Clinical Psychologist (10 hours/week for the first 18 months). This expert will lead the development and validation of the annotation protocol for identifying cognitive biases from participant data (WP3), supervise a team of trained raters to establish inter-rater reliability, consult on the design of game scenarios to elicit target biases (WP4), and assist in interpreting the RCT results (WP8). This methodological oversight is critical for the project's scientific validity.

While our teams have access to significant existing infrastructure, specific new equipment and materials are required. The requested items are not available through existing resources and are critical for specific tasks.

For Subproject 2, we request one High-Performance Workstation for game development. While existing lab machines are sufficient for coding, developing a high-fidelity 3D environment and processing high-resolution video from usability studies (WP5) require a state-of-the-art GPU (e.g., NVIDIA RTX 4090). This investment will enable rapid iteration on the game's design and reduce build times, supporting the timeline for WP4.

For Subproject 1, we request a dedicated server to host the AI inference engine during the RCT (WP7). The university's HPC cluster is designed for training, not the low-latency inference needed for a real-time, closed-loop system. A dedicated server with a professional GPU (e.g., NVIDIA RTX A4000) will provide the computational environment required to test the primary real-time adaptation objective, targeting a response time under one second (Specific Objective 3). This resource also ensures consistent processing of participant data throughout the RCT and can host project-specific data management tools, providing value independent of shared HPC queue times.

Under materials, we request five additional Empatica E4 wearable biosensors. The existing pool of ten devices is insufficient for the logistics of the 150-participant RCT, which involves concurrent participants, charging, and data offloading. This is essential for the smooth execution of WP7. We also request funds for software licenses, including Unity Pro for advanced development features and Kubios Premium for gold-standard HRV analysis (WP3). A significant portion of this category is allocated to participant compensation for the pilot study (N=50), usability study (N=20), and RCT (N=150). This is an ethical necessity and is critical for achieving recruitment targets and ensuring the statistical power of our results.

The final budget category covers travel and other direct costs essential for project coordination and dissemination (WP1, WP9).

We have budgeted for travel to support two in-person, all-hands project meetings per year. These face-to-face workshops are indispensable for key collaborative tasks that are inefficient to conduct remotely, such as co-designing game scenarios (WP4), planning the system integration architecture (WP5), and jointly interpreting RCT data (WP8). We also request travel funds for disseminating findings at major international conferences (e.g., NeurIPS, CHI, SIRS). This allows team members to present our work, receive feedback, and build collaborations, a critical component of WP9.

Other direct costs include Article Processing Charges (APCs) for at least three open-access publications, fulfilling our commitment to open science and the funder's mandate. A small budget is allocated for long-term data archiving costs in a public repository like Zenodo, ensuring we can deliver on Specific Objective 5 and make our dataset a publicly available resource under FAIR principles. Finally, we have included a modest sum for Institutional Review Board submission fees. Each cost is directly linked to the successful completion and impact of the ADAPT-MIND project.

Title: "Justification of Personnel Costs"

Justification of Personnel Costs

The successful execution of this project requires a dedicated team with the specialized expertise outlined below. The requested effort for each team member has been carefully calibrated to the tasks in the Work Packages (WPs), ensuring the team is appropriately sized to meet project demands without redundancy.

For Subproject 1, led by Dr. Vargas, we request funding for three key positions to develop the core AI engine. A full-time Postdoctoral Researcher (36 months, 100% dedication) with expertise in machine learning and multimodal signal processing will lead the technical development of the cognitive bias detection model. This researcher will architect and train a novel model that infers cognitive states from multimodal data streams (WP3), moving beyond conventional performance-based adaptation. They will also implement advanced optimization techniques for real-time inference (WP6) and support system integration API development (WP5). Their expertise in low-latency model deployment is essential for achieving Specific Objective 1. We also request a full-time PhD Student (36 months, 100% dedication) specializing in computational linguistics. This student will develop the linguistic data processing pipeline, from audio preprocessing to extracting the subtle semantic and syntactic features that serve as critical inputs to the fusion model (WP3). Their doctoral thesis will focus on identifying multimodal linguistic markers of cognitive bias, representing a significant training and capacity-building opportunity. Finally, a part-time Research Technician (36 months, 50% dedication) is required to manage the project's data infrastructure, overseeing data collection protocols (WP2, WP7), managing High-Performance Computing (HPC) cluster jobs, and curating the final, anonymized dataset for public release (WP9) in accordance with our Data Management Plan.

For Subproject 2, led by Dr. Morales, we request funding for a complementary team to develop and validate the user-facing platform. A full-time Postdoctoral Researcher (36 months, 100% dedication) with a strong background in HCI and serious game development will lead the design and implementation of the ADAPT-MIND game platform in Unity (WP4). This individual will implement the novel cognitive-bias-aware Dynamic Difficulty Adjustment (DDA) mechanism, alongside the performance-based and control arms for comparison (WP5), and manage the technical deployment for the RCT (WP7). Their expertise is indispensable for achieving Specific Objectives 2, 3, and 4. A full-time PhD Student (36 months, 100% dedication) with a background in psychology and user experience (UX) research is also requested. This student will lead user-centered design activities, conduct formal usability testing (WP5), manage participant recruitment and data collection for the RCT (WP7), and analyze the qualitative and user experience data. Their work will ensure the final system is usable, engaging, and psychologically valid. To ensure the project's clinical integrity, we request a part-time Clinical Psychologist (10 hours/week for the first 18 months). This expert will lead the development of the annotation schema and provide the expert-level labeling of pilot data required to establish a reliable ground truth for our AI model (WP3). They will also consult on the design of game scenarios to ensure they elicit the target cognitive biases (WP4) and assist in the clinical interpretation of the RCT results (WP8). Their contribution is critical for the project's scientific and therapeutic validity.

Justification of Equipment and Materials

While our teams have access to significant existing infrastructure, the unique technical demands of this project necessitate targeted acquisitions of specific equipment and materials. These items are not available through existing resources and are critical for the successful execution of the work plan.

For Subproject 2, we request funding for one new High-Performance Workstation dedicated to game development and asset creation. The development of a high-fidelity 3D environment in Unity (WP4) and the processing of high-resolution video data from usability studies (WP5) are computationally intensive tasks that would be severely bottlenecked by existing lab machines. This dedicated workstation will enable rapid iteration on the game's design, reduce build compilation times, and efficiently render complex graphical assets, ensuring the timely development of the high-quality user experience required for the project.

For Subproject 1, we request funding for a dedicated server to host the AI inference engine during the RCT (WP7). The university's High-Performance Computing (HPC) cluster, while ideal for training our models, is a shared resource optimized for batch processing and is not designed for the low-latency, on-demand inference required by our real-time, closed-loop system. Relying on it would introduce unacceptable latency and instability, invalidating the core premise of the intervention. A dedicated server with a mid-range professional GPU (e.g., NVIDIA RTX A4000) will be configured to run the optimized AI model, guaranteeing the sub-one-second response time critical for the adaptive system's functionality.

Under materials and consumables, we request funding for five additional Empatica E4 wearable biosensors. The existing pool of ten devices is insufficient to manage the logistics of the 150-participant RCT, which involves concurrent participants, charging cycles, data offloading, and backup devices. We also request funds for critical software licenses not covered by institutional agreements: annual subscriptions for Unity Pro, for advanced features and technical support in WP4, and Kubios Premium, the gold-standard software for HRV analysis required for our physiological data pipeline in WP3 [5]. Finally, a significant portion of this budget is allocated to participant compensation. We have budgeted fair compensation, based on standard university rates, for participants in the pilot study (N=50), usability study (N=20), and main RCT (N=150). This is an ethical necessity and is critical for achieving recruitment targets and minimizing attrition, thereby ensuring the statistical power of our results.

Title: "Justification of Travel and Other Costs"

Justification of Travel and Other Costs

This budget category covers travel and other direct costs essential for project coordination, dissemination, and compliance with open science mandates. These costs are integral to the project's collaborative strategy, directly supporting activities in WP1 (Project Management) and WP9 (Dissemination) and are crucial for achieving our stated objectives.

Travel funds are requested to support two critical functions: internal coordination and external dissemination. First, we have budgeted for two in-person, all-hands project meetings per year, alternating between the Principal Investigators' institutions. Although most coordination will be conducted virtually, these face-to-face workshops are necessary for key collaborative tasks that are difficult to perform remotely. These include the initial co-design sessions for game scenarios (WP4), detailed planning of the system integration architecture (WP5), and the joint interpretation of final RCT data (WP8). These strategic activities require the focused, interactive environment that in-person collaboration provides. Second, travel funds will support the dissemination of findings at major international conferences. This will enable PIs, postdoctoral researchers, and PhD students to present our work at premier venues in relevant fields, such as NeurIPS (AI), CHI (HCI), and the annual meeting of the Schizophrenia International Research Society (SIRS). As a key component of WP9, conference participation is essential for sharing results with the scientific community, receiving valuable feedback, and fostering international collaborations.

Other direct costs are allocated for publication, data archiving, and ethical review. We have budgeted for Article Processing Charges (APCs) to publish at least three articles in high-impact, open-access journals. This allocation is essential for fulfilling our commitment to open science and the funder's mandate, covering the significant costs associated with top-tier journals. A modest budget is also allocated for long-term data archiving in a public repository such as Zenodo. This ensures we can deliver on Specific Objective 5 by making our multimodal dataset a lasting, publicly available resource in accordance with FAIR principles. Finally, we have included funds to cover the administrative fees for submitting our study protocols to the Institutional Review Board for ethical review, a mandatory step for human participant research. Each cost is directly linked to a specific project deliverable or compliance requirement.

References

[1] Corcoran, C. M., & Cecchi, G. A. (2020). Using Language Analysis to Predict Psychosis. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging, 5*(8), 770–779.

[2] Unknown (2025). Navigating the semantic space: Unraveling the structure of meaning in psychosis using different computational language models. *PubMed*. Retrieved from https://pubmed.ncbi.nlm.nih.gov/38280291/

[3] Wang, J., et al. (2025). A highly scalable deep learning language model for common risks prediction among psychiatric inpatients. Retrieved from https://bmcmedicine.biomedcentral.com/counter/pdf/10.1186/s12916-025-04150-7.pdf

[4] Hendricks Brown, C., et al. (2025). Automated speech and language markers of longitudinal changes in psychosis symptoms. Retrieved from https://www.nature.com/articles/s44277-025-00034-z

[5] IEEE, et al. (2025). Collecting language, speech acoustics, and facial expression to predict psychosis and other clinical outcomes: strategies from the AMP® SCZ initiative. *Proceedings of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 5575–5579. Retrieved from https://www.nature.com/articles/s41537-025-00669-z

[6] Unknown (2025). Leveraging AI and Language Analysis to Predict Psychosis Risk in Clinical High-Risk Individuals. *Biological Psychiatry*. Retrieved from https://www.biologicalpsychiatryjournal.com/article/S0006-3223(25)00277-X/abstract

[7] Dudley, R., Taylor, P., Wickham, S., & Hutton, P. (2016). Psychosis, Delusions and the "Jumping to Conclusions" Reasoning Bias: A Systematic Review and Meta-analysis. *Schizophrenia Bulletin, 42*(3), 652-65.

[8] Lincoln, T. M., Salzmann, S., …., & Westermann, S. (2011). *Journal of Behavior Therapy & Experimental Psychiatry, 42*(2), 185-191.

[9] Lüdtke, T., Kriston, L., Schröder, J., …. M., & Moritz, S. (2017). *Journal of Behavior Therapy and Experimental Psychiatry, 56*, 106–112.

[10] Winton-Brown, T. T., Broome, M. R., Allen, P., …, Johns, L. C., & McGuire, P. (2015). *European Psychiatry, 30*(1), 32-7.

[11] Bearden, C. E., Wu, K. N., Caplan, R., & Cannon, T. D. (2011). *Journal of the American Academy of Child & Adolescent Psychiatry (JAACAP), 50*(7), 669–680.

[12] Mota, N. B., Vasconcelos, N. A. P., Lemos, N., Pieretti, A. C., … Ribeiro, S. (2012). Speech Graphs Provide a Quantitative Measure of Thought Disorder in Psychosis. *PLOS ONE, 7*(4).

[13] Bilgrami, Z. R., Sarac, C., Srivastava, A., Herrera, S. N., …, Ch.M. (2022). *Schizophrenia Research, 245*, 90-96.

[14] Unknown (2021). Natural language processing methods are sensitive to sub-clinical linguistic differences in schizophrenia spectrum disorders. Retrieved from https://www.nature.com/articles/s41537-021-00154-3.pdf

[15] Strobl, E. V., Eack, S. M., Swaminathan, V., & Visweswaran, S. (2012). *Early Intervention in Psychiatry, 6*(4), 368-79.

[16] van Os, J., Linscott, R. J., Myin-Germeys, I., et al. (2009). A systematic review and meta-analysis of the psychosis continuum: evidence for a psychosis proneness-persistence-impairment model of psychotic disorder. *Psychological Medicine, 39*, 179–195.

[17] Johns, L. C., & van Os, J. (2001). The continuity of psychotic experiences in the general population. *Clinical Psychology Review, 21*, 1125–1141.

[18] van Os, J., Hanssen, M., Bijl, R. V., & Ravelli, A. (2000). Strauss (1969) revisited: a 21st century schizophrenia. *Schizophrenia Research, 45*, 11–20.