An Object-Oriented Architecture for a Nostalgia-Driven Adaptive Persuasive System: Design and Field-Validated Implementation

Remi A. Ikechukwu, Nuka Nwiabu, Daniel Matthias, Emmanuel Bennett
Department of Computer Science, Rivers State University, Port Harcourt, Nigeria

ABSTRACT
Digital persuasive systems that rely on static, non-personalized reinforcement (reminders, rule-based tracking) frequently fail to sustain engagement once user motivation declines, particularly under acute stress. This paper presents the object-oriented architecture of a Nostalgic Persuasive System, an adaptive engine that detects a user's emotional and stress state from reflective journaling text and responds with a personalized, generationally-targeted nostalgic media recommendation. The architecture integrates six functional layers: data acquisition, model training/preprocessing, text-based emotion classification (a fine-tuned DistilRoBERTa model), text-based stress detection (a fine-tuned RoBERTa model), a dual content-recommendation layer (LightFM for movies, content-based vector similarity for songs), and a contextual-bandit adaptive-learning layer (LinUCB) that refines recommendations from user feedback. The system's functional/non-functional requirements, class structure, and behavioral design are detailed following Design Science Research guidelines. The architecture was implemented and deployed as a live web application, generating 268 real interaction logs over a 7-day field evaluation with 49 active participants, showing substantially higher habit adherence in the personalized condition (90.1%) than in a non-personalized control (41.3%). This paper contributes a reusable architectural blueprint for emotion-aware, adaptive persuasive systems.

General Terms
Design, Algorithms, Experimentation

Keywords
Object-oriented design, persuasive systems architecture, contextual bandit, recommendation systems, affective computing, software engineering

1. INTRODUCTION
Recent advances in affective computing, intelligent recommendation, and persuasive technologies have enabled increasingly personalized digital interventions. Still, many reported systems are presented either as isolated machine-learning models or as conceptual frameworks offering limited implementation guidance. Relatively few studies describe a modular software architecture that integrates emotion recognition, stress detection, recommendation, adaptive learning, and user interaction into one deployable application. This gap limits reproducibility, maintainability, and practical deployment. Object-oriented software engineering addresses this by promoting modularity, loose coupling, and component reuse which is essential where multiple intelligent components must operate together while remaining independently maintainable as models evolve.

This paper presents the object-oriented architecture of a nostalgia-driven adaptive persuasive system integrating transformer-based emotion classification, transformer-based stress detection, hybrid recommendation, and contextual-bandit adaptive learning into a unified framework. Three contributions evolves from this paper: first, a modular object-oriented architecture integrating these components within a unified persuasive system. Second, a specification of the functional requirements, non-functional requirements, structural architecture, and behavioral interactions necessary to implement such a system in practice. Third, a demonstration of practical feasibility through implementation and live deployment, evidencing that the design functions successfully in a real-world environment rather than remaining conceptual.

2. RELATED LITERATURE
The theoretical foundation for emotion-aware persuasive systems is rooted in the Persuasive Systems Design (PSD) framework, which organizes behavior-change strategies into primary task support, dialogue support, system credibility, and social support [1]. Although widely adopted, PSD's emphasis has traditionally been structural rather than attentive to users' dynamic emotional states during interaction. Alahäivälä et al. [2] demonstrated that a layered architecture for Behavior Change Support Systems is essential for translating persuasive design into deployable applications. In mobile health, McGowan et al. [3] showed personalization significantly improves engagement, while Zhang et al. [4] observed that social-support and credibility features remain underutilized in many PSD-based interventions.

User disengagement remains one of the most persistent challenges facing digital behavior-change technologies, irrespective of the underlying design framework. Kidman et al. [5] reported a median 70% mobile-health app abandonment within 100 days, with declining motivation and poor experience as primary causes. To Achieve an adaptive, sustained engagement will require an architecture that supports modularity and continuous evolution; Object-Oriented Analysis and Design via UML provides an established methodology for this [6], and forms the structural foundation for the class and sequence diagrams in this study.

Personalized recommendation is critical since content relevance strongly influences engagement. Hybrid approaches combining collaborative filtering with content-based features have been shown to alleviate the cold-start problem common to purely collaborative methods [7], a challenge widely recognized in recommender-system research [8]; content-based techniques relying on item similarity and are particularly suitable where descriptive metadata is available [9]. The proposed architecture combines both paradigms.

Adaptive content selection under uncertainty has been studied through the multi-armed bandit framework, which provides a principled balance between exploration of new alternatives and exploitation of previously successful recommendations [10]. The Linear Upper Confidence Bound (LinUCB) algorithm extends this by modeling expected reward as a function of context [11], making it well suited to persuasive systems where preferences evolve and require continuous personalization.

The affective-sensing component is founded on transformer-based language models: RoBERTa [12] provides a robust architecture for text classification, while DistilBERT [13] achieves comparable performance at lower computational cost, suiting real-time use. The importance of emotional awareness in intelligent systems originates from affective computing, where Picard [14] argued that emotionally-aware systems support natural interactions better. The system further employs nostalgia as an emotion-regulation mechanism, this draws on evidence that nostalgic reflection restores psychological equilibrium and strengthens self-continuity under stress [15]. Existing research has largely examined these areas separately. Few studies actually integrate them into a unified, field-validated object-oriented architecture, the gap this study addresses.

3. SYSTEM REQUIREMENTS SPECIFICATION

Following Design Science Research guidelines, the system's operational boundaries were defined prior to architectural design, serving as evaluation criteria for the resulting artifact.

Table 1. Functional Requirements

ID	Requirement
FR-1	Reflective journal ingestion: capture open-ended user journaling text via a secure text interface.
FR-2	Self-reported indicator capture: record user-reported stress/mood metrics via standardized scales.
FR-3	Real-time emotion classification: classify journal text into categorical affective states (Joy, Love, Sadness, Anger, Neutral, Fear, Surprise).
FR-4	Continuous stress quantification: compute a normalized stress score bound between 0.0 and 1.0 from text and self-reported inputs.
FR-5	Dynamic context mapping: evaluate whether the stress score breaches the intervention threshold (> 0.30).
FR-6	Nostalgic profile retrieval: match a user's birth year and preferences to culturally/generationally aligned media candidates.
FR-7	Content ranking and selection: score, filter, and select the highest-value nostalgic media item via the recommendation and bandit layers.
FR-8	Habit alignment interface: surface the target habit action alongside the nostalgic trigger.
FR-9	Interaction log generation: automatically generate structured telemetry for each session.
FR-10	Feedback capture: record the user's response to the post-content “does this bring back memories?” prompt as a reward signal.

Table 2. Non-Functional Requirements

ID	Requirement
NFR-1 (Performance)	Emotion/stress analysis completes within 2s; recommendations within 5s.
NFR-2 (Security)	All user data encrypted in transit (TLS).
NFR-3 (Scalability)	Multiple concurrent users; 99.5% uptime.
NFR-4 (Maintainability)	Modular; independent model upgrades without UI/schema changes.
NFR-5 (Usability)	Usable by technical and non-technical users.

4. SYSTEM ARCHITECTURE

The Nostalgic Persuasive System is structured into six layers, forming a closed adaptive loop that senses user emotional state and delivers, then refines, personalized nostalgic content: (i) data source/user-input; (ii) preprocessing/model-training; (iii) emotion classification (DistilRoBERTa [13]); (iv) stress detection (RoBERTa [12]); (v) dual recommendation (LightFM [7] for movies, content-based for songs); and (vi) a hierarchical contextual-bandit layer (LinUCB [11]) for final adaptive selection.

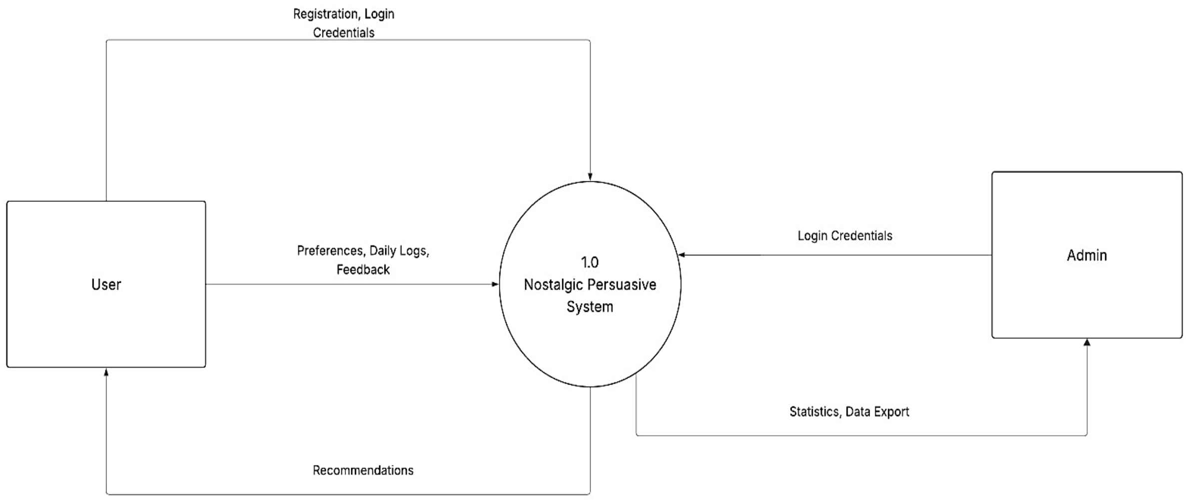

Fig 1: Context flow diagram, showing signal flow between text-analysis, recommendation, and bandit layers.

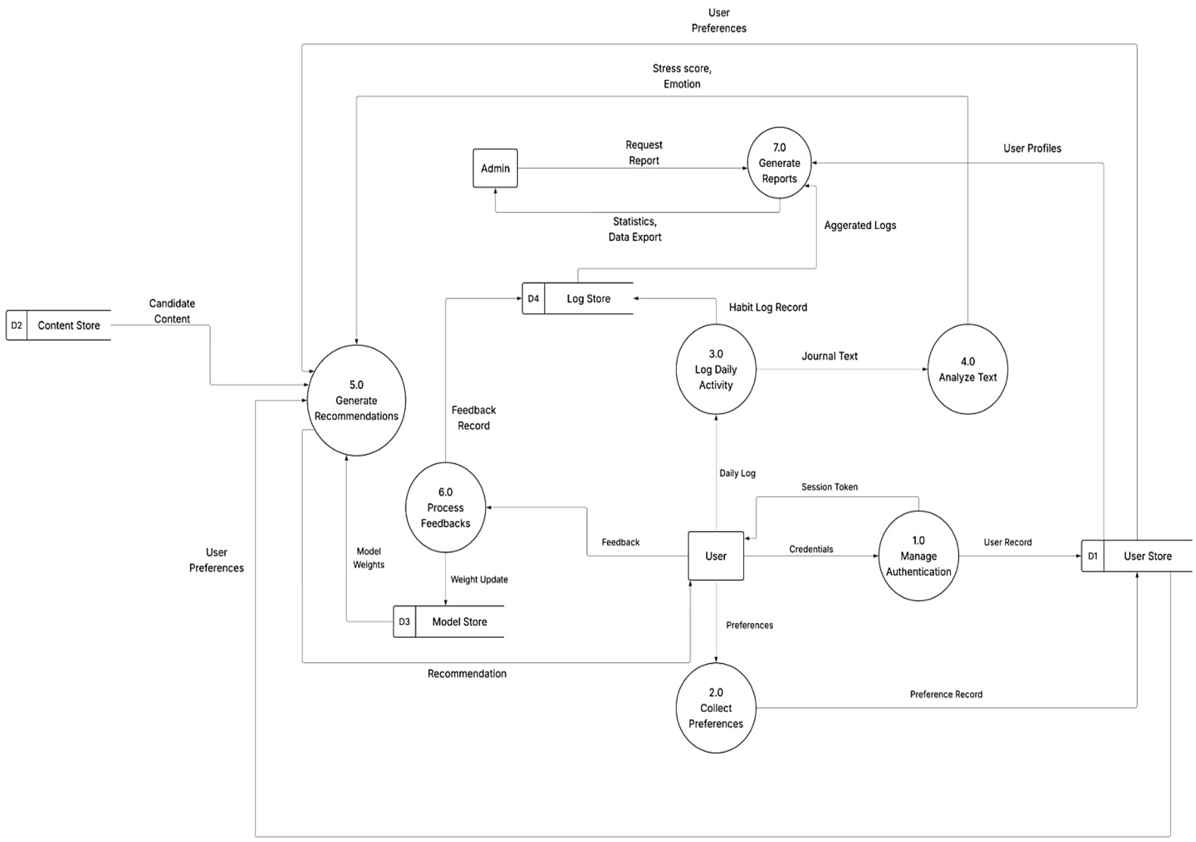

Fig 2: Data flow diagram of the Nostalgic Persuasive System

Both recommenders enforce a ten-year recency filter so all candidates are old enough to evoke nostalgia. Candidate content is scored using a nostalgia formula combining a personal-nostalgia term (weighted toward the user's formative years, ages 10–22 i.e the "reminiscence bump") and a cultural-nostalgia term (a popularity-based boost for pre-birth content): Final Score = Personal × (0.7 + 0.3 × Popularity) + Cultural. (1). Candidates below a minimum score are filtered before reaching the contextual bandit, which uses a context vector of stress level, emotion distribution, historical feedback rate, and normalized birth year.

5. OBJECT-ORIENTED DESIGN

The static structure comprises four component categories with high internal cohesion and minimal external coupling. The Recommendation Services i.e MovieRecommender and SongRecommender. Text Analysis Services: StressDetector and EmotionDetector, running on-device via PyTorch. Contextual Bandit System i.e LinUCBBandit (via MABWiser) and HierarchicalBandit, composing one global instance with many per-user instances persisted via joblib for warm restarts. Domain Entities includes User, UserPreferences, DailyHabitLog, ContentFeedback, and the Movie/Song catalogues. The Key relationships includes: HierarchicalBandit which is composed of multiple LinUCBBandit instances; each User has one UserPreferences record and generates multiple DailyHabitLog/ContentFeedback records over time; Then the recommenders, bandit, and text-analysis services form a dependency chain from candidate generation through context supply to selection and feedback.

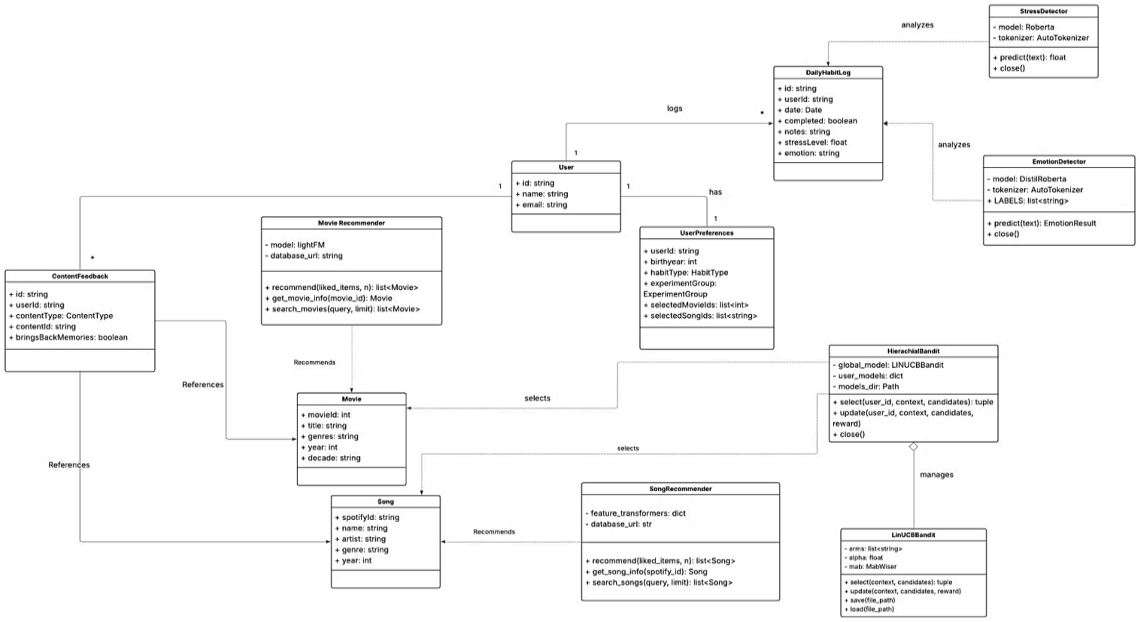

Fig 3: UML class diagram of the Nostalgic Persuasive System, showing the four component categories and their relationships

The runtime recommendation sequence proceeds are as follows: (1) user request, (2–4) context retrieval (e.g preferences, prior likes, most recent habit log), (5–11) real-time text analysis through the stress and emotion models, (12–17) candidate generation (up to 50 candidates per recommender, filtered to content ≥10 years old), (18–19) nostalgia scoring and filtering, (20–24) contextual-bandit selection through the upper-confidence-bound scoring, with random selection among the top five similarity-ranked items to preserve variety, (25–26) response delivery.

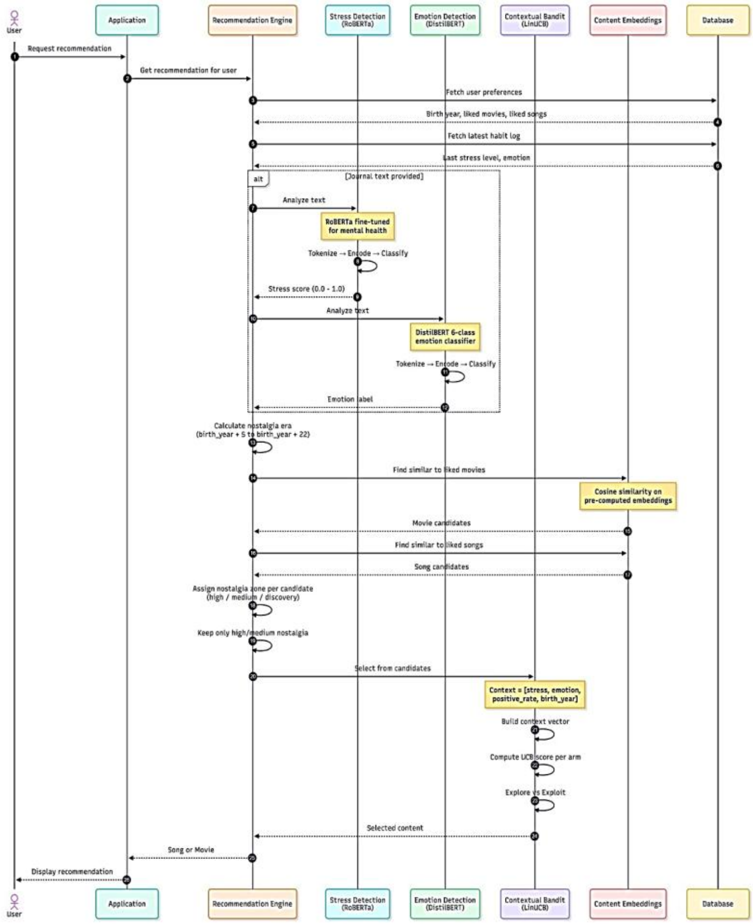

Fig 4: UML sequence diagram for the recommendation flow, fr om user request through contextual-bandit selection to response delivery

After viewing a recommendation, the participant answers "does this bring back memories?" A positive response yields a reward of 1, while the negative yields a reward of 0, this (user, content, reward, context) tuple is persisted and the LinUCB model performs an online parameter update, closing the reinforcement learning cycle and allowing future recommendations to better reflect the individual user's demonstrated preferences.

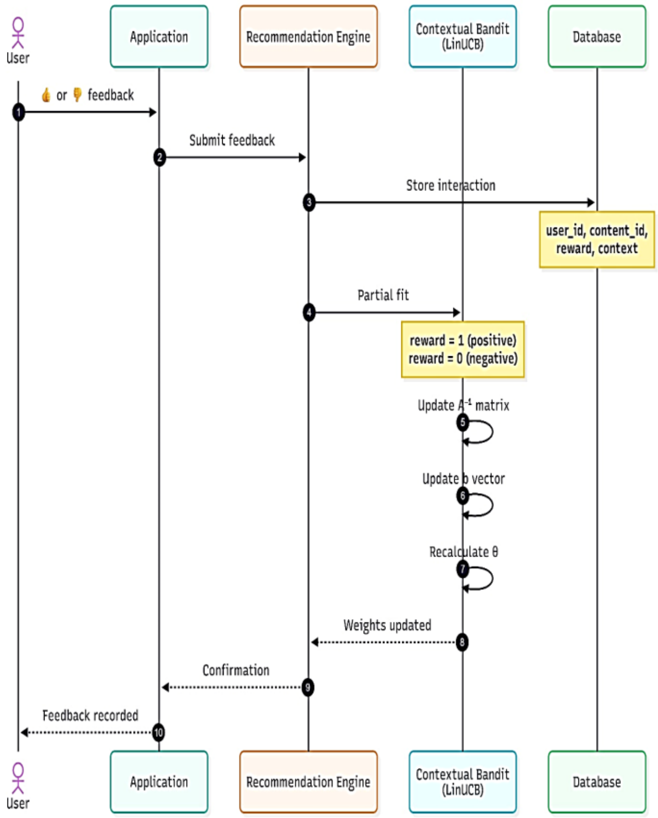

Fig 5: UML sequence diagram for the feedback loop, showing the online bandit parameter update following a user's nostalgia-relevance response.

6. IMPLEMENTATION AND FIELD DEPLOYMENT

The system was implemented as a web-integrated application with an administrative interface for monitoring engagement and content-effectiveness. Model training for the emotion and stress classifiers was conducted offline against public benchmark datasets prior to deployment, allowing low-latency runtime inference within the NFR-1 budget. To validate operational feasibility under realistic conditions, the complete system was deployed in a 7-day field study with 49 active participants across a personalized (treatment) and non-personalized (control) condition, generating 268 interaction logs.

The personalized condition achieved a habit adherence rate of 90.1%, compared to 41.3% in the control condition, a difference confirmed statistically significant by a chi-square test of independence ($\chi 2(1, N=268) = 70.10$, $p < .001$). Because the two conditions differed specifically in the presence of the architecture's adaptive personalization layer, this result indicates the architecture's components functioned together as intended under real usage, not merely in isolated unit testing.

To evaluate the architecture across more than a single aggregate scenario, adherence and stress outcomes were further examined across four naturally occurring participant sub-scenarios in the deployment: generational cohorts spanning Baby Boomers, Generation X, Millennials, and Generation Z (Table 3), each representing a distinct combination of nostalgic-content preferences and cold-start conditions for the recommendation and bandit layers.

Table 3. Adherence and stress outcomes by generational cohort

Cohort	n	Adherence	Mean Stress
Baby Boomers	5	89.3%	0.18
Generation X	6	90.3%	0.19
Millennials	4	89.3%	0.24
Generation Z	9	90.9%	0.24

A one-way ANOVA found no statistically significant difference in adherence across cohorts ($F(3,20) = 0.12$, $p = .945$), indicating the architecture's recommendation and bandit layers generalized consistently despite each cohort requiring different nostalgic content (different eras, media, and cultural references) and despite younger cohorts necessarily presenting a colder start to the per-user bandit model than older, more established users. Stress-reduction outcomes showed more variation by cohort ($F(3,20) = 1.67$, $p = .206$); given the modest per-cohort sample sizes ($n = 4$-9), this is reported as a suggestive rather than conclusive finding, and is identified as a direction for evaluation with larger cohorts in future work. Taken together, these sub-scenario results provide a more granular evaluation of the architecture than the aggregate result alone, showing that its cold-start handling (Section 4) performed comparably across participants with substantially different content-preference profiles and different amounts of prior interaction history.

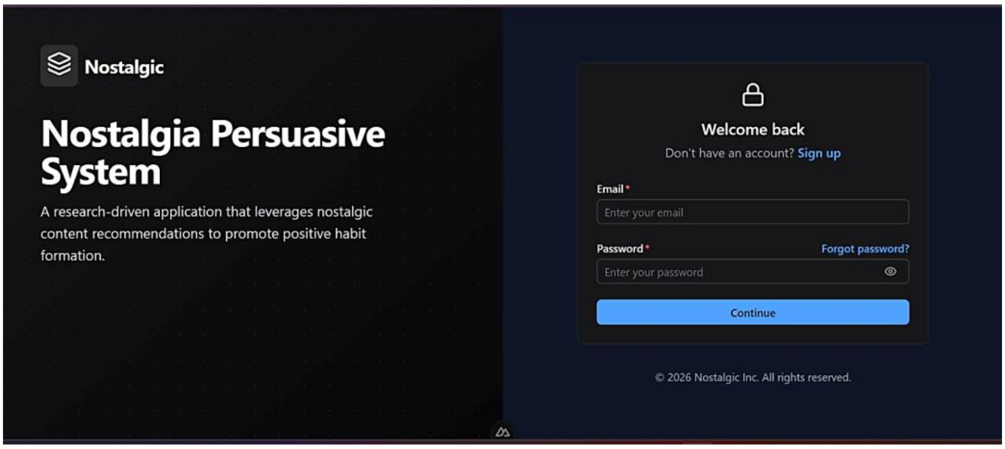

Fig 6: User-facing interface of the deployed Nostalgic Persuasive System.

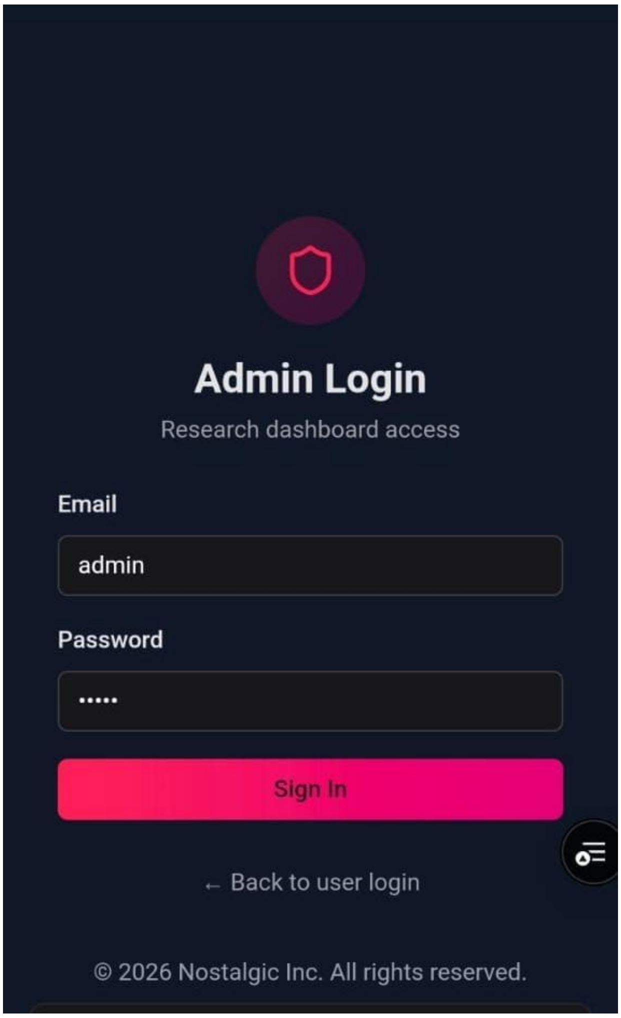

Fig 7: Administrative interface used for monitoring participant engagement and system performance during the field deployment

7. DISCUSSION

Decoupling Text Analysis Services from Recommendation Services and the Contextual Bandit layer (NFR-4) means any one component, for example the emotion classifier, can be replaced or upgraded without requiring changes elsewhere, this is a direct benefit of the object-oriented design in Section 5. Both the LightFM recommender (via user folding) and the hierarchical bandit (via a shared global model) explicitly address the cold-start problem, This is architecturally significant since every user begins as a cold-start case. The cohort-level analysis in Section 6 provides direct evidence for this design choice showing that adherence remained statistically indistinguishable across four cohorts with markedly different cold-start profiles, from Generation Z participants (necessarily the coldest starts for a nostalgia-content bandit) to Baby Boomer participants with the most established generational content history, suggesting the global-model fallback in the hierarchical bandit (Section 5) is doing meaningful architectural work rather than being a redundant safeguard.

Beyond nostalgia-driven persuasive technology, this architecture offers a reusable software-engineering framework for adaptive systems across domains including digital health, education, and productivity support, since each loosely-coupled, modular component supports future extension without disproportionate rework. The weaker, more variable stress-reduction result across cohorts (Section 6) is analytically informative in its own right suggesting that while the architecture's content-selection mechanism generalizes well for the behavioral outcome, it was primarily optimized against (habit adherence), its secondary psychological outcome (stress mitigation) may be more sensitive to cohort-specific, such as differing baseline stress-coping styles across generations. This distinction between a robust primary outcome and a more variable secondary outcome is a more precise characterization of the architecture's demonstrated capabilities than the aggregate result alone provides.

A limitation of this study, is that the architecture was validated over a 7-day window, this is well sufficient for a short-term behavioral impact though, but not for a long-term content fatigue effects, also, the fixed ten-year recency filter and nostalgia-scoring formula (Section 4) would benefit from longer-horizon evaluation to detect novelty decay. More work could be done in this area, when we consider this study's field evaluation, although it covered four distinct generational sub-scenarios and was conducted within a few application domains (habit formation i.e. exercise completion and smoking cessation), evaluating the architecture against more additional behavior-change domains and against alternative bandit algorithms (e.g., Thompson Sampling) would strengthen the claims of general applicability beyond what the present deployment supported. Architectural work also, could introduce explicit content-rotation logic within the bandit's reward shaping, to proactively mitigate novelty decay, and a dedicated ablation isolating the bandit layer's own marginal contribution, relative to the upstream recommenders would more precisely quantify its independent value.

8. CONCLUSION

This paper presented the functional requirements, object-oriented architecture, and behavioral design of a Nostalgic Persuasive System integrating transformer-based emotion and stress detection, a dual content-recommendation layer, and a contextual-bandit adaptive-learning mechanism. Unlike a purely theoretical proposal, this design was implemented as a live system and field-deployed over seven days, producing a measurable behavioral effect (90.1% vs. 41.3% habit adherence) consistent with its intended design goals. The resulting architecture provides a reusable blueprint for developing emotion-aware adaptive persuasive technologies while maintaining modularity, extensibility, and long-term maintainability.

9. REFERENCES

[1] H. Oinas-Kukkonen and M. Harjumaa, “Persuasive systems design: Key issues, process model, and system features,” Communications of the Association for Information Systems, vol. 24, no. 1, Article 28, 2009.

[2] T. Alahäivälä, H. Oinas-Kukkonen, and T. Jokelainen, “Software architecture design for health BCSS: Case Onnikka,” in Persuasive Technology, LNCS vol. 7822, Springer, pp. 3–14, 2013.

[3] A. McGowan, S. Sittig, D. Bourrie, R. Benton, and S. Iyengar, “The intersection of persuasive system design and personalization in mobile health: Statistical evaluation,” JMIR mHealth and uHealth, vol. 10, no. 9, e40576, 2022.

[4] K. Zhang, Y. Jiang, K. Xi, and N. Sun, “Application of persuasive system design in mobile health interventions for chronic disease management: A mini review,” Frontiers in Public Health, vol. 13, 1718371, 2025.

[5] P.G. Kidman, R.G. Curtis, A. Watson, and C.A. Maher, “When and why adults abandon lifestyle behavior and mental health mobile apps: Scoping review,” Journal of Medical Internet Research, vol. 26, e56897, 2024.

[6] A. Dennis, B.H. Wixom, and D. Tegarden, Systems Analysis and Design: An Object-Oriented Approach with UML, 6th ed., Wiley, 2020.

[7] M. Kula, “Metadata embeddings for user and item cold-start recommendations,” arXiv:1507.08439, 2015.

[8] A.I. Schein, A. Popescul, L.H. Ungar, and D.M. Pennock, “Methods and metrics for cold-start recommendations,” in Proc. 25th ACM SIGIR Conf., pp. 253–260, 2002.

[9] P. Lops, M. de Gemmis, and G. Semeraro, “Content-based recommender systems: State of the art and trends,” in Recommender Systems Handbook, Springer, pp. 73–105, 2011.

[10] P. Auer, N. Cesa-Bianchi, and P. Fischer, “Finite-time analysis of the multiarmed bandit problem,” Machine Learning, vol. 47, no. 2–3, pp. 235–256, 2002.

[11] L. Li, W. Chu, J. Langford, and R.E. Schapire, “A contextual-bandit approach to personalized news article recommendation,” in Proc. 19th Intl. Conf. World Wide Web (WWW '10), ACM, pp. 661–670, 2010.

[12] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, “RoBERTa: A robustly optimized BERT pretraining approach,” arXiv:1907.11692, 2019.

[13] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, “DistilBERT, a distilled version of BERT: Smaller, faster, cheaper and lighter,” arXiv:1910.01108, 2019.

[14] R.W. Picard, Affective Computing, MIT Press, 1997.

[15] C. Sedikides and T. Wildschut, “The sociality of personal and collective nostalgia,” European Review of Social Psychology, vol. 30, no. 1, pp. 123–173, 2019.