

University of Tartu
Faculty of Social Sciences
Narva College

Information Technology Systems Development

Denys Neliubov

**GYMBRO: TEHISINTELLEKTIL PÕHINEVA
TREENINGRAKENDUSE ARENDAMINE JA HINDAMINE
KONTEKSTITUNDLIKU AVATARI KAASLASEGA**

Bachelor's Thesis

Supervisor: Nicolai Morozov, Lecturer in Information Technology

Narva 2026

ABSTRACT

Neliubov, D. (2026). GymBRO: Development and Evaluation of an AI-Powered Fitness Application with a Context-Aware Avatar Companion. Bachelor's Thesis, University of Tartu, Narva College.

The increasing adoption and rapid abandonment of mobile fitness applications highlights a persistent gap in user motivation and long-term engagement. This thesis addresses this problem by proposing and implementing a living avatar concept — a virtual companion whose visual state is directly driven by the user's real fitness behaviour, powered by a context-aware AI assistant. The resulting application, GymBRO, integrates nutrition tracking, workout logging, and sleep monitoring into a single system where an AI chatbot, embodied as a responsive avatar character, provides personalized coaching based on the user's current data.

The theoretical framework draws on Self-Determination Theory (Deci & Ryan, 1985, 2000), the Tamagotchi effect (Frude & Jandric, 2015), and recent research on AI-powered conversational agents in health behaviour change (Aggarwal et al., 2023; McFadyen et al., 2024). These converging lines of research predict that combining contextual AI intelligence with a persistent visual companion should produce stronger motivation outcomes than either component alone.

The application was developed over six months using React, FastAPI, and the OpenAI GPT-4o-mini API, with an offline-first architecture. User evaluation was conducted through a structured survey instrument distributed via Google Forms to participants with prior experience using the application. At the time of the preliminary defence, data collection is ongoing; full results will be presented in the final submission.

Keywords: gamification, virtual companion, fitness application, AI assistant, motivation, user engagement, Self-Determination Theory, living avatar

AUTHORSHIP DECLARATION

I hereby declare that this thesis is my own work. All sources used in writing the thesis have been indicated and referenced correctly.

This thesis and the GymBRO application were developed with the assistance of AI tools, in accordance with the University of Tartu guidelines on AI use in academic work (Tartu Ülikooli tekstiroboti suunis, 2023). The specific AI tools and their roles are disclosed below:

Application development: ChatGPT (OpenAI) and Gemini (Google) were used as coding assistants for generating and debugging application code. Claude (Anthropic) was used for code review and refactoring.

Visual assets: Avatar companion images were initially generated using AI image generation models (Gemini, Nano, Banana) and subsequently refined by a collaborating designer who corrected visual artefacts and standardised style consistency.

Design structure: Grok (xAI) was used for brainstorming the initial application design structure and UI layout concepts.

Thesis writing: Claude (Anthropic) was used for structuring the thesis, drafting text, and organising the theoretical framework. All AI-generated text was reviewed, verified, and edited by the author.

All research design decisions, technical architecture choices, implementation logic, survey design, and academic conclusions are the author's own. The author takes full responsibility for the content and accuracy of this work.

Denys Neliubov

05.04.2026

ABBREVIATIONS AND DEFINITIONS

AI — Artificial Intelligence; computational systems capable of performing tasks that typically require human intelligence.

API — Application Programming Interface; a set of rules allowing software components to communicate.

GPT — Generative Pre-trained Transformer; the underlying architecture of the OpenAI language models used in GymBRO.

IndexedDB — a low-level browser API for client-side storage of structured data, used in GymBRO for offline-first local data persistence.

MVP — Minimum Viable Product; an early-stage product with sufficient features to be usable and testable.

SDT — Self-Determination Theory; a macro theory of human motivation developed by Deci and Ryan (1985, 2000).

UI — User Interface; the visual and interactive layer of a software application.

UX — User Experience; the overall experience of a person using a product, particularly in terms of ease and satisfaction.

XP — Experience Points; a numerical value representing accumulated user activity used to advance the avatar's level in GymBRO.

TABLE OF CONTENTS

ABSTRACT.....	2
AUTHORSHIP DECLARATION.....	4
ABBREVIATIONS AND DEFINITIONS.....	5
TABLE OF CONTENTS.....	6
INTRODUCTION.....	7
1. THEORETICAL BACKGROUND.....	9
1.1 Gamification in Health and Fitness Applications.....	9
1.2 Virtual Companions and the Tamagotchi Effect.....	10
1.3 Motivation Theories Applied to Fitness Behaviour.....	11
1.4 AI Conversational Agents in Mobile Health Applications.....	12
1.5 Review of Existing Fitness Applications.....	13
1.6 The Living Avatar Concept: Theoretical Synthesis.....	15
2. APPLICATION DEVELOPMENT.....	17
2.1 Development Methodology.....	17
2.2 Technology Stack and Justification.....	17
2.3 System Architecture.....	19
2.4 Avatar System and Living Avatar Logic.....	21
2.5 AI Integration and Context-Aware Chat.....	22
2.6 Supporting Features.....	23
2.7 Analytics System.....	25
2.8 Challenges and Solutions.....	25
2.9 Implementation Limitations.....	26
3. TESTING AND EVALUATION.....	27
3.1 Research Methodology.....	27
3.2 Survey Instrument.....	28
3.3 Results.....	29
3.4 Discussion.....	30
3.5 Limitations of the Study.....	30
CONCLUSIONS.....	31
RESÜMEE.....	33
REFERENCES.....	34
APPENDICES.....	37
LIHTLITSENTS.....	40

INTRODUCTION

I have been involved in sports since primary school — martial arts, gym training, and physical discipline have been a constant part of my life. In 2024, I reached what I consider my personal peak: strong physical shape, consistent training, and genuine love for the gym. Then I lost it. No training partner, plateauing numbers, and a central nervous system that stopped recovering the way it used to. The motivation I had built over years collapsed within weeks. I tried several fitness applications to compensate — calorie trackers, workout loggers, running apps — but they all existed in isolation from each other, and none of them felt like they understood my situation. They tracked data. They did not care.

That experience led to the research question addressed in this thesis: could a mobile application replace — or at least partially substitute — the social and motivational role of a training partner and a personal coach? Not just track data, but actually know the user, remember context, and respond in a way that feels personal. In September 2025, during a startup programme at Ida-Virumaa in Narva, I began building that idea. The result is GymBRO — a fitness application combining an AI assistant with a virtual avatar companion whose state is driven by the user's real behaviour. The project was presented at the Spinner24 competition and won first place. This thesis is the academic documentation of that work.

The problem GymBRO addresses is well-established in the literature. Mobile fitness applications suffer from a chronic retention problem — users adopt them enthusiastically and abandon them within weeks (Krebs & Duncan, 2015, as cited in Bitrián et al., 2021). Existing approaches to improving retention — gamification mechanics such as badges and streaks, or AI coaching assistants — are typically implemented in isolation. Gamification lacks personalisation. AI assistants lack emotional presence. No application I encountered combined both into a single coherent experience.

GymBRO proposes a specific solution to this gap: a "living avatar" — a virtual companion whose visual state directly reflects the user's fitness behaviour, powered by an AI assistant that maintains contextual awareness of the user's nutrition, training history, and goals. The avatar is not decorative. It is the interface through which the AI's knowledge becomes directly perceivable by the user.

The research question guiding this thesis is:

How does an AI assistant embodied as a context-aware virtual avatar companion affect user motivation and engagement in a fitness application?

The aim of this thesis is to design, develop, and evaluate a fitness application that integrates an AI-powered avatar companion, and to assess its effect on user motivation and engagement among fitness application users.

To achieve this aim, the following objectives were set:

- To review existing research on gamification, virtual companions, and AI assistants in fitness applications;
- To design and develop the GymBRO application implementing the living avatar concept;
- To evaluate the application through user testing and assess its impact on motivation and engagement.

The thesis is structured as follows. Chapter 1 presents the theoretical background, covering gamification in fitness applications, virtual companion research, motivation theories, and AI assistants in mobile health. Chapter 2 describes the development methodology, system architecture, and key features of the GymBRO application. Chapter 3 presents the user testing methodology, results, and discussion. The thesis concludes with a summary of findings and proposals for further development.

The author's contribution includes the full design and development of the application — React frontend, FastAPI backend, avatar logic, AI integration, and analytics system — developed over six months from October 2025 to April 2026. The authentication module was implemented as shared infrastructure and is not claimed as the author's original contribution.

1. THEORETICAL BACKGROUND

This chapter reviews existing research relevant to the GymBRO application across four interconnected areas: gamification in health and fitness applications, virtual companions and emotional attachment, motivation theories applied to physical activity, and AI-powered conversational agents in mobile health. The chapter concludes by synthesising these areas into the theoretical basis for the living avatar concept implemented in GymBRO.

1.1 Gamification in Health and Fitness Applications

Gamification is broadly defined as the use of game design elements in non-game contexts (Deterding et al., 2011). Common elements include points, badges, leaderboards, progress bars, challenges, streaks, and levels — mechanics originally designed to drive engagement in games, repurposed to motivate behaviour in everyday applications. In the health and fitness domain, gamification has become one of the most widely adopted strategies for increasing physical activity and sustaining user engagement with mobile applications (Lister et al., 2014).

Evidence for the effectiveness of gamified interventions on physical activity is growing but remains mixed. A systematic review and meta-analysis by Xu et al. (2024), covering studies up to May 2024, found statistically significant improvements in behaviour change outcomes compared to control conditions. However, several studies highlight important limitations. Huang et al. (2024) found that many users experience motivation fatigue — a gradual decline in engagement as the novelty of gamification elements wears off. Koivisto and Hamari (2019) note that the long-term effects of gamification remain insufficiently studied.

A key insight from the literature concerns the type of game elements used. Bitrián et al. (2021) demonstrated in a longitudinal field study that incommensurate elements — those not directly tied to performance metrics, such as social recognition and relationship-oriented feedback — were more effective at building intrinsic motivation and long-term engagement than commensurate elements such as points directly awarded for steps taken. This finding has important implications for GymBRO's design, in which the avatar's emotional state and visual appearance — rather than numerical scores — serve as the primary feedback mechanism.

The dominant theoretical framework applied to gamification in physical activity research is Self-Determination Theory (SDT; Deci & Ryan, 1985, 2000), reviewed in detail in Section 1.3. A systematic review by Xu et al. (2021) found that SDT was the most commonly applied theoretical framework in gamified physical activity interventions, used in 32% of studies reviewed.

1.2 Virtual Companions and the Tamagotchi Effect

A virtual companion, or digital pet, is a software-based artificial entity designed to simulate aspects of a companion relationship with the user. The concept entered mainstream consumer culture in 1996 with the introduction of the Tamagotchi by Bandai — a handheld device containing a simple digital creature that required regular feeding, cleaning, and attention, and would deteriorate or "die" if neglected. By 2010, tens of millions of Tamagotchi units had been sold worldwide (Giasiranis & Sofos, 2025), demonstrating an extraordinary capacity for digital artefacts to attract sustained human attention and care.

The psychological mechanism underlying this attachment has been termed the Tamagotchi effect — the development of emotional attachment to machines, robots, or software agents (Frude & Jandric, 2015). The phenomenon builds on the concept of animism, or the human tendency to attribute life and consciousness to inanimate objects. Research has shown that this attachment persists even when users explicitly recognise that the object is not alive — a phenomenon consistent with what Reeves and Nass (1996) termed the Media Equation: humans respond to computers and media as if they were real social actors.

Research on virtual companions in health contexts supports this dynamic. Studies have found that users form genuine attachment to virtual companions, associating care of the pet with positive emotional states. Critically for this thesis, virtual pets have been found to be particularly effective when the user's care actions have visible consequences for the companion's state — a design principle central to GymBRO's living avatar logic.

A narrative review of virtual pet research by Giasiranis and Sofos (2025) synthesised 45 academic articles and identified six major themes, confirming that virtual pets serve purposes beyond entertainment and can contribute to health habit formation and emotional well-being management. The review also noted that virtual pets remain under-researched relative to their cultural and commercial significance — a gap that this thesis contributes to addressing in the specific context of fitness motivation.

A relevant recent development is the application of virtual companions in commercial fitness contexts. The mobile game Pokémon GO demonstrated that avatar-like characters connected to physical world activity can produce large-scale increases in physical activity among players (Althoff et al., 2016). While Pokémon GO represents a full game rather than a gamified application, its success illustrates the motivational power of character

attachment combined with activity-contingent rewards — a combination that GymBRO implements in a personalised, AI-enhanced form.

1.3 Motivation Theories Applied to Fitness Behaviour

Understanding why people adopt and sustain physical activity habits requires a robust theoretical framework. Self-Determination Theory (SDT), developed by Deci and Ryan (1985, 2000), is the most extensively applied motivational framework in exercise psychology and the most commonly used theoretical basis in gamified fitness intervention research (Xu et al., 2021).

SDT proposes that all humans have three basic psychological needs whose satisfaction is essential for sustained motivation and psychological well-being: autonomy — the need to feel that one's behaviour is self-directed and volitional; competence — the need to feel effective and capable of achieving goals; and relatedness — the need to feel connected to others in a meaningful way (Ryan & Deci, 2000). Critically, SDT distinguishes between intrinsic motivation — engaging in an activity for its inherent interest and enjoyment — and extrinsic motivation — engaging to obtain external rewards or avoid negative outcomes. While both can drive behaviour, intrinsic motivation is associated with greater persistence, well-being, and quality of engagement.

SDT's relevance to gamification lies in how game design elements map onto these needs. Bitrián et al. (2021) found that competence and autonomy needs were satisfied by achievement-oriented affordances, while relatedness was satisfied by social and immersive features — and that satisfying these needs produced greater long-term health outcomes through intrinsic motivation. GymBRO's design maps directly onto this framework: the avatar's growth and level progression address competence; the companion's social-like responsiveness addresses relatedness; and the flexible offline-first tracking addresses autonomy.

A second relevant framework is habit formation theory, as articulated by Duhigg (2012). Habit formation proposes a three-component loop: cue (a trigger that prompts the behaviour), routine (the behaviour itself), and reward (the outcome that reinforces the behaviour). GymBRO functions as a habit-formation system: the avatar's visible needs serve as cues, fitness activities constitute the routines, and avatar improvement and AI positive feedback serve as rewards. This loop-based design aligns with research showing

that visible progress feedback is one of the most effective drivers of fitness app engagement (Lister et al., 2014).

It is important to note that SDT cautions against over-reliance on extrinsic rewards, as these can undermine intrinsic motivation through what Deci et al. (1999) termed the overjustification effect. GymBRO mitigates this risk by designing the avatar system as a relational companion rather than a points-and-prizes reward system — following Bitrián et al.'s (2021) recommendation to prioritise relationship-oriented design elements over direct performance-point mapping.

1.4 AI Conversational Agents in Mobile Health Applications

Artificial intelligence-powered conversational agents have attracted significant research attention as a vehicle for health behaviour change interventions. A systematic review by Aggarwal et al. (2023) examined 15 empirical studies on AI chatbots for health behaviour change and found that across diverse health domains, AI chatbots demonstrated high efficacy. The review found that free-flow, unconstrained conversations produced better user experience through personalisation, and that chatbots incorporating established behaviour change theories consistently outperformed those without theoretical grounding.

A complementary scoping review by Fu et al. (2026), covering studies through March 2024, confirmed that AI chatbots showed positive outcomes in 81.67% of comparisons across 43 studies, with the strongest evidence for physical activity, smoking cessation, and stress management. However, the review also identified that sleep, weight management, and dietary behaviour remain underexplored — domains that GymBRO specifically addresses.

A key limitation of existing AI health assistants identified across the literature is the absence of visual identity and emotional engagement. This pattern — where perceived personality sustains engagement while generic responses erode it — directly motivates GymBRO's living avatar approach: combining contextual AI intelligence with a persistent visual character so that each component addresses the known limitations of the other.

A randomised controlled trial by McFadyen et al. (2024) found that a generative AI-enabled application produced a threefold increase in user engagement compared to static content delivery, with 2.4 times higher usage frequency and 3.8 times longer usage

durations, providing strong evidence that generative AI-powered personalisation produces qualitatively different engagement dynamics.

1.5 Review of Existing Fitness Applications

To contextualise the problem that GymBRO addresses, this section surveys the three main categories of consumer fitness applications — workout tracking, nutrition logging, and sleep monitoring — with reference to representative applications used by the author prior to and during the development of this project. The review focuses on motivation and engagement mechanics rather than technical implementation, following the analytical framework of Lister et al. (2014), who evaluated fitness applications specifically through the lens of their capacity to sustain user behaviour change.

Workout tracking applications represent the most mature and commercially developed segment. Fitbod and Muscle Booster are representative examples of the dominant model: structured AI-assisted workout planners that generate personalised training programmes and track completed exercises. Both applications offer well-designed interfaces, evidence-based training logic, and strong feature sets. The primary value proposition is structure — the application tells the user what to train and tracks their progress against a plan. This model works well for users who are already committed to training and need an organised system to follow. The limitation, however, is not functional but motivational: the application is entirely passive between sessions. It does not initiate contact, does not know whether the user trained yesterday or last month, and cannot express anything resembling concern or encouragement. The application waits. Whether the user opens it is entirely up to the user — and for many, that decision is precisely where motivation collapses (Teixeira et al., 2012).

Nutrition tracking applications represent a similarly mature category. КБЖУ: счётчик калорий, a widely used Russian-language calorie counter, exemplifies the strengths of this segment: a comprehensive food database, accurate macronutrient tracking, a clear daily summary interface, and no significant usability friction. Used consistently, such applications provide a genuinely useful picture of nutritional intake. The limitation is structural rather than technical — the application records what happened but does not respond to it. Logging a nutritionally poor day produces the same interface experience as logging an excellent one. There is no feedback loop that connects the data to action, and no persistent entity whose state reflects the quality of the user's choices.

Sleep monitoring applications occupy a distinct position in the landscape. Widgetable, a mobile application featuring a virtual companion that sleeps and lives on the device's home screen widget, represents an interesting hybrid: it is primarily a companion and customisation application rather than a health analytics tool, but its sleep tracking feature — through which the companion "sleeps" alongside the user and its widget state reflects rest patterns — demonstrates that sleep data can be delivered through an emotionally engaging rather than purely analytical interface. The application's commercial success suggests that users respond positively to sleep monitoring when it is embedded in a companionship context rather than presented as a clinical metric.

Across all three categories, the applications reviewed are technically capable and, within their respective domains, genuinely useful. The issue is not quality — it is fragmentation and passivity. A user managing their fitness seriously must operate three or four separate applications simultaneously, none of which share data or context, and none of which initiate any form of engagement. Each application responds to user input but never generates it. The motivational burden falls entirely on the user: to remember to open each application, to interpret what the data means across multiple disconnected dashboards, and to translate that interpretation into action — alone, without guidance or social feedback (Middelweerd et al., 2014). The applications reviewed here are not an exhaustive survey of the market but are representative of the categories most relevant to the problem this thesis addresses..

This structural gap — the absence of a persistent, context-aware, proactively engaging layer across the fitness data the user already produces — is the specific problem GymBRO is designed to address. The hypothesis is not that existing applications are inadequate, but that the same data they collect could produce qualitatively different motivation outcomes if it were mediated through a relational interface: a companion that knows the user's nutrition, training history, and sleep, responds to that context intelligently, and whose own visible state reflects the user's real behaviour. The question this thesis investigates is whether that relational layer — an AI assistant embodied as a virtual companion — produces a measurable effect on user motivation and engagement.

1.6 The Living Avatar Concept: Theoretical Synthesis

Gamification research establishes that game elements can increase physical activity motivation, but their long-term effectiveness is undermined by motivation fatigue when

elements are primarily extrinsic (Huang et al., 2024; Bitrián et al., 2021). Relationship-oriented elements produce more durable engagement — pointing toward gamification approaches that centre companionship rather than competition.

Virtual companion research establishes that digital entities can create genuine emotional attachment through the Tamagotchi effect (Frude & Jandric, 2015), particularly when the companion's state visibly responds to the user's actions. However, traditional virtual pets lack contextual intelligence — they respond to predefined inputs but cannot engage in personalised dialogue or adapt to the user's specific situation.

AI chatbot research establishes that conversational agents can effectively support health behaviour change, especially through free-flow, personalised interactions. However, purely text-based chatbots lack visual presence and emotional anchoring, limiting the depth of the user's perceived relationship with the system.

SDT provides the motivational framework connecting these elements. A living avatar — an AI assistant embodied as a responsive visual character whose state reflects the user's actual behaviour — simultaneously addresses all three basic psychological needs: autonomy (the user freely decides which activities to track), competence (avatar growth reflects real achievements), and relatedness (the avatar functions as a persistent social-like companion).

The research gap that GymBRO addresses is therefore specific and well-grounded: no existing fitness application combines contextually aware AI assistance with a persistent visual avatar whose state is driven by real user behaviour data. This combination is theoretically predicted to produce stronger motivation and engagement than either component alone.

Based on this theoretical synthesis, the following hypothesis is proposed: *An AI assistant embodied as a context-aware visual avatar companion, whose state reflects the user's real fitness behaviour, will positively affect user motivation and engagement in a fitness application, as measured by self-reported motivation, avatar attachment, and perceived AI usefulness.*

This hypothesis is evaluated in Chapter 3 through a user survey administered to participants who used the GymBRO application.

2. APPLICATION DEVELOPMENT

This chapter describes the development of the GymBRO application. It covers the methodology used, justification of the technology stack, system architecture, and the implementation of core features including the avatar system, AI integration, and analytics.

2.1 Development Methodology

The application was developed using an Agile methodology with iterative prototyping over six months, from October 2025 to April 2026. The development began within the Spinner24 startup programme at Ida-Virumaa in Narva, where the project was pitched before a jury of five to six competing teams and won first place. This competitive context shaped the development priorities: the MVP had to be visually compelling, technically functional, and demonstrably unique within a short timeline.

A key early decision was the adoption of an offline-first architecture, storing user records locally and routing only AI and analytics requests through the backend. The rationale for this approach is discussed in detail in Section 2.2.

The visual assets for the avatar companions were generated using AI image generation models (Gemini, Nano, Banana) and subsequently refined by a collaborating designer who corrected visual artefacts, standardised style consistency across pose variants, and created additional position states. The author was the primary and sole developer of all application code throughout the project.

2.2 Technology Stack and Justification

The selection of the technology stack for GymBRO was governed by three primary constraints: the developer's existing technical expertise, a strict six-month academic timeline, and the requirement for a functional prototype deployable to physical Android devices for user testing. Each decision was made to maximize development velocity without compromising the application's core functionality.

React 19 was selected as the primary UI framework, leveraging the author's prior experience. While alternatives such as Flutter or Native Android (Kotlin) were considered, they would have necessitated a significant learning curve, which was incompatible with the project's timeline.

To bridge the gap between web and mobile, Capacitor 7 was chosen over React Native. This allowed for a standard React web application to be wrapped in a native Android container, maintaining a single JavaScript codebase. This approach ensured that the build pipeline remained unified and efficient, allowing for rapid deployment to physical devices. The development process was further accelerated by Vite 6, which provided near-instant Hot Module Replacement (HMR) compared to traditional build tools like Webpack, facilitating rapid UI iterations.

The application implements an offline-first architecture using IndexedDB (via the Dexie.js wrapper) for primary data storage. This decision was strategically driven by three factors. First, performance: all user interactions, such as logging meals or workouts, execute instantly on the client side without network latency. Second, infrastructure costs: by storing user records locally, the need for expensive cloud hosting was eliminated, making the project sustainable within an academic research budget. Third, research focus: since the study's primary objective is to evaluate the living avatar concept rather than cloud synchronisation, a local-first approach was deemed sufficient for the MVP. The primary tradeoff is that records do not synchronise automatically across devices — a known limitation discussed in Section 2.9.

While the core application logic resides on the client side, GymBRO integrates with a FastAPI server and a PostgreSQL database for specific administrative and analytical functions. It is important to clarify that the backend infrastructure is part of a shared development ecosystem. The server-side logic and the PostgreSQL schema were developed by a fellow student within the same study programme for a parallel research project. GymBRO utilises this shared infrastructure primarily for two purposes: user authentication — credentials are transmitted over HTTPS and stored as hashed values, never in plaintext — and centralised storage of the user profile data entered during onboarding, including name, weight, height, fitness goals, and activity level. This profile data forms the persistent personal context injected into every AI system prompt, enabling the avatar companion to maintain awareness of the user's physical characteristics across sessions. PostgreSQL was the preferred choice for this shared environment due to its ability to handle complex relational queries, such as demographic aggregations and onboarding funnel analysis, which are essential for the evaluation phase of this thesis.

The living avatar intelligence is powered by the OpenAI Chat Completions API using the gpt-4o-mini model. This model provides the optimal balance between high-quality, context-aware responses and low token costs, which is crucial for maintaining the avatar's personality within an academic budget. To ensure the avatar remains aware of the user's physical state, health data such as sleep sessions is retrieved through the Android Health Connect API via the @capgo/capacitor-health plugin. This integration allows the application to synchronise with various wearable devices, providing the necessary context for the avatar's feedback and state changes.

2.3 System Architecture

The system is organised into four functional layers. The frontend layer contains all screens and UI components. The local data layer manages IndexedDB tables via Dexie.js. The API client layer handles communication over HTTPS; notably, the implementation of the OpenAI API integration and chat logic was developed by the author as part of the client-side infrastructure. The backend layer provides global state management and analytics. User authentication is handled by the shared infrastructure module described in Section 2.2.

Figure 1 illustrates the separation between client and server responsibilities. The client stores all user records locally and contacts external services only for AI-specific requests (processed via the author's custom integration), pet data retrieval, and analytics events. This separation ensures a responsive UI even in low-connectivity environments and eliminates the need for cloud storage infrastructure at the MVP stage.




Figure 1 separation schema between client and server responsibilities (Source: Author)

The local database schema uses Dexie.js versioning across three schema versions, which allowed safe migrations as the data model evolved during development. The application maintains seven local tables: meals (nutrition log entries), trainingDays (completed workout days with unique date constraint), sleepRecords (daily sleep aggregates), chatMessages (AI conversation history managed by the author's AI module), customFoods (user-defined food items), foodOverrides (nutrient value overrides), and settings (application configuration). Figure 2 provides a schematic overview of the schema.




Figure 2 data tables schema (Source: Author)

2.4 Avatar System and Living Avatar Logic

The core innovation of GymBRO is the living avatar — a virtual companion whose visual state and behaviour are directly driven by the user's fitness data. The application currently offers two companion characters: a wolf and a rabbit. Both characters were initially generated using AI image generation models (Gemini, Nano, Banana) and subsequently refined by a collaborating designer, who corrected visual artefacts, standardised the visual style across character states, and created additional pose variants to represent different activity and emotional states.

Before accessing the main application, a new user completes an onboarding flow through the CreatePetScreen. The user assigns a name to their companion and provides basic personal information — weight, height, fitness goals — which forms the initial context injected into every subsequent AI system prompt. This onboarding step is critical to the living avatar concept: it establishes the companion's identity from the first interaction and ensures the AI assistant has meaningful personal context before any conversation begins.

The avatar's ongoing state is determined by a composite care score calculated from three dimensions of user activity: nutrition tracking, workout consistency, and sleep quality. Healthy food entries contribute XP to the avatar, advancing it through growth levels. The XP award system includes caps, bonuses, and penalties to prevent gaming: a single meal cannot award unlimited XP, bonus XP is granted for hitting daily nutritional targets, and skipping days reduces the streak multiplier. Training day streaks are tracked through the calendar interface, with the avatar's visual state reflecting current consistency. Sleep data imported from Android Health Connect contributes to the avatar's overall wellbeing score. The care scoring function maps these inputs to one of three states — strong, moderate, or weak — which directly determines the AI's emotional mode in the next conversation.

The avatar is rendered through the PetImage component, which selects the appropriate sprite based on the current character type and activity state. This visual feedback mechanism creates a direct, visually meaningful representation of the user's fitness behaviour — the defining design goal of the living avatar concept.

2.5 AI Integration and Context-Aware Chat

The AI chat functionality is implemented through the ChatScreen interface, which communicates with the backend `/ai/query` endpoint. Each request includes a dynamically constructed system prompt generated from the user's profile data and current activity context — including recent nutrition, training history, and avatar state.

The system prompt is built by the `generateSystemPrompt()` function. This function receives the pet's name and a context object containing all available user data, and assembles a structured prompt that defines the AI's identity, voice, behavioural rules, and the user's current state. This is the technical core of the living avatar concept: the AI does not receive a generic instruction set, but a personalised context rebuilt on every request.

The avatar level determines the coaching style: users at levels 1–3 receive gentle guidance suited to early habit formation, levels 4–7 receive structured coaching with warmth, and users at level 8 and above receive more direct, performance-oriented feedback. The pet care status — the composite signal that drives both avatar state and AI tone — is computed from five binary inputs: whether the user trained today, whether they have eaten, their 7-day average sleep duration, and whether a training schedule is active. A score of 4–5 maps to a 'strong' care state, triggering proud and affectionate AI responses; 2–3 maps to 'moderate', triggering encouraging coaching; 0–1 maps to 'weak', triggering comfort mode focused on one small achievable next step. This graduated system follows SDT's principle of supporting competence without undermining autonomy (Ryan & Deci, 2000; Deci et al., 1999).

The final assembled prompt defines the AI's identity, voice constraints, emotional modes, and a knowledge contract — an explicit instruction set preventing the AI from inventing data not present in the context. A selection of 14 dialogue examples is embedded in the prompt to guide tone consistency across interaction types. Conversation history is persisted locally in the `chatMessages` table, enabling continuity across sessions. AI requests are rate-limited to 15 requests per hour per user on the backend.

2.6 Supporting Features

GymBRO consists of ten screens, each mapped to a specific user task: `HomeScreen` (pet status, XP level, activity summary), `CalendarScreen` (training log and streak tracking), `MealsScreen` (nutrition journal with macronutrient view), `AddMealScreen` (food search, custom foods, calorie calculation), `SleepScreen` (sleep data from Health Connect), `ChatScreen` (AI companion conversation), `CreatePetScreen` (onboarding: pet name, type, user profile), `EnterScreen` (entry point navigation), `SettingsScreen` (application configuration), and `AuthScreen` (login and registration, shared infrastructure). Figure 3 and 4 provides a visual overview of the primary screens.




Figure 4 main screens on app, part 1 (Source: Author)




Figure 5 main screens on app, part 2 (Source: Author)

The nutrition module enables users to log meals from a pre-configured food database covering major food categories with calorie and macronutrient data. Users can also define custom food items. The XP award system is intentionally non-trivial: a single meal cannot award unlimited XP, bonus XP is granted for hitting daily nutritional targets, and repeated logging of the same low-quality food is subject to diminishing returns. This design prevents the system from being gamed and ensures that XP progression reflects genuine nutritional effort. The direct link between food quality and avatar visual state makes the companion's appearance an honest reflection of the user's habits rather than their logging frequency.

The training module provides a calendar interface for marking completed training days and computes workout streaks. The sleep module integrates with Android Health Connect to import sleep session data, aggregating incoming sessions into a one-record-per-day model using a transactional upsert that preserves existing records and inserts new dates atomically.

2.7 Analytics System

A telemetry system captures user behaviour data for both research and development purposes. The frontend sends structured events to the `/analytics/events` endpoint following key user actions. Twenty-five distinct event types are tracked across the user journey, covering onboarding completion, pet creation, AI chat interactions, meal logging, training entries, and sleep synchronisation.

2.8 Challenges and Solutions

The development of GymBRO presented several challenges that required iterative problem-solving. This section documents the most significant ones.

The first challenge was integrating the avatar companion into the application without it feeling decorative or disconnected from the functional screens. The design decision — making the avatar the visual background of the main activity screens (nutrition, sleep, calendar) rather than a separate dedicated screen — required careful layering of UI components. The avatar needed to be visually present and immediately readable without competing with interactive elements in the foreground. Achieving the right balance required multiple iterations of the component hierarchy in the React layout.

A related UI challenge was the FloatingFlameIcon — a small floating icon, similar in concept to the floating button in TikTok, that remains visible on top of all screens and provides quick navigation to the AI chat. Implementing a movable overlay element that stays within screen boundaries on mobile required handling touch position relative to screen dimensions and preventing the icon from being positioned off-screen. This is a small detail but an important one for mobile UX: an unreachable icon would defeat its purpose.

The avatar's visual state system — reflecting the user's fitness quality through the companion's appearance across states ranging from lean and active to visibly unhealthy — required designing a scoring function that felt motivating rather than punishing. An early version of the scoring was too sensitive: a single missed training day would visibly degrade the avatar. User feedback during development showed this felt discouraging. The final system uses a rolling 7-day average across three dimensions (nutrition, training, sleep), which smooths single bad days and rewards consistent improvement — a decision directly informed by SDT's principle of supporting competence without undermining motivation (Deci et al., 1999).

On the AI side, the main challenge was producing consistently warm, in-character responses across a wide range of user inputs. Early prompt versions produced responses that felt generic — corporate wellness language that contradicted the companion personality the application was trying to establish. The solution was adding explicit forbidden phrases to the prompt ("fitness journey", "overall well-being", "How can I assist you today?") and embedding 14 dialogue examples covering key interaction types to anchor tone consistency.

2.9 Implementation Limitations

Several limitations of the current implementation should be noted. First, the offline-first architecture means that records do not synchronise automatically across devices. Second, Android Health Connect integration limits sleep tracking to Android devices only. Third, the current avatar system supports one pet per user. These limitations are discussed further in the Conclusions.

3. TESTING AND EVALUATION

This chapter presents the methodology used to evaluate the GymBRO application, the results obtained from user testing, and a discussion of the findings in relation to the theoretical framework established in Chapter 1.

3.1 Research Methodology

This study employs a descriptive evaluative research design, combining quantitative and qualitative data collection through a structured self-report survey. The primary research instrument was an online questionnaire distributed via Google Forms to participants who had prior experience using the GymBRO application. This approach is consistent with established methods for evaluating prototype fitness applications, where post-use surveys are widely used to assess perceived motivation, usability, and feature-specific engagement (Lister et al., 2014; Bitrián et al., 2021).

The survey was administered following a period of application use. Participants were asked to use the application for at least one day prior to completing the survey, with a recommended usage period of one week to allow sufficient experience with the avatar progression system and AI chat across multiple sessions. The survey was distributed via Google Forms (<https://forms.gle/uWEiSfKvXvSqNYB6A>). Participants were recruited through multiple channels within the author's personal and professional environment. The primary recruitment group consisted of peers and acquaintances personally familiar with the application. Additionally, the survey was shared within the Spinner24 startup programme community at W3N in Narva, where the GymBRO project had been presented and won first place — participants from that event were invited both in person and via the programme's Telegram group. A further group of participants was recruited from the author's gym environment, specifically individuals who had been introduced to the application during informal testing sessions. This sampling strategy — purposive convenience sampling across interconnected social and professional networks — is appropriate for an initial prototype evaluation where the goal is to generate qualitative insight and preliminary quantitative indicators rather than statistically generalisable conclusions (Etikan et al., 2016).

Participation was entirely voluntary. All participants were informed of the purpose of the study — evaluation of a fitness application developed as part of a bachelor's thesis — prior to completing the survey. No personally identifiable information was collected.

Participants were informed that their responses would be used solely for academic purposes and that they could withdraw at any point without consequence. These procedures are consistent with the ethical requirements for student research at the University of Tartu Narva College (Üliõpilaste lõputööde koostamise ja vormistamise juhend, 2024).

A total of 21 participants completed the survey between 31 March and 11 May 2026. Participants ranged in age from 18 to 50 years, with the majority falling in the 18–24 range ($n = 13$, 62%). Fitness experience varied across all self-reported levels, from beginners to experienced athletes with two or more years of consistent training.

The survey was designed to generate data directly relevant to the research question: how does an AI assistant embodied as a context-aware virtual avatar companion affect user motivation and engagement in a fitness application? Each section of the survey corresponds to one dimension of this question, allowing findings to be mapped onto the theoretical framework established in Chapter 1.

3.2 Survey Instrument

The survey instrument consisted of five thematic sections with a total of 19 items, combining five-point Likert scale items (15 items) with open-ended qualitative questions (4 items). The full survey is reproduced in Appendix A.

Section 1 — Background (3 items) — collected self-reported demographic and context information: age range, self-assessed fitness experience level (beginner, intermediate, or experienced), and prior use of other fitness applications. This section established baseline characteristics of the sample relevant to interpreting motivation and engagement responses. No personally identifiable information was requested.

Section 2 — Motivation and Engagement (5 Likert items) — assessed the extent to which the GymBRO application affected participants' motivation and sense of agency over their fitness routine. Items were formulated to address the three dimensions of Self-Determination Theory (Ryan & Deci, 2000): autonomy ("Using GymBRO made me feel in control of my own fitness routine"), competence ("Seeing my avatar's progress made me feel that my efforts were making a difference"), and relatedness ("I felt a sense of connection to my avatar companion"). Additional items assessed general motivation increase and behavioural engagement frequency. Each item was rated on a five-point Likert scale from 1 (Strongly Disagree) to 5 (Strongly Agree).

Section 3 — Avatar Companion Experience (3 Likert items + 1 open-ended) — specifically evaluated the living avatar concept. Likert items assessed the meaningfulness of avatar state changes in relation to real user activity, the sense of responsibility toward the avatar (operationalising the Tamagotchi effect), and perceived personalisation compared to a conventional fitness application. The open-ended item invited participants to describe in their own words how the avatar affected their experience, providing qualitative data to complement the quantitative responses.

Section 4 — AI Chat (3 Likert items + 1 open-ended) — evaluated the AI assistant component, with particular focus on context-awareness. Items assessed the perceived relevance of AI responses to the participant's actual fitness situation, the perceived intelligence of the AI regarding the user's context (workouts, nutrition, avatar state), and the degree to which the AI chat motivated concrete action. The open-ended item asked participants to identify what they found most and least useful about the AI chat interaction.

Section 5 — Overall Usability (2 Likert items + 1 open-ended) — provided a holistic evaluation of the application. Items addressed general ease of use and likelihood of recommending the application to others. The open-ended item invited suggestions for improvement, surfacing design insights not captured by the thematic sections.

3.3 Results

A total of 21 participants completed the survey between 31 March and 11 May 2026. Participants were drawn from several interconnected groups: peers and acquaintances of the author, attendees of the Spinner24 programme at W3N in Narva who encountered the project at its competition presentation, individuals reached through the programme's Telegram group, and gym contacts introduced to the application during informal testing. This multi-channel convenience sampling approach is reflected in the demographic distribution: the sample skews strongly toward the 18–24 age group ($n = 13$, 62%), which corresponds to the social and professional environment of the author. Four participants fell in the 25–34 range (19%) and four in the 35–50 range (19%). Older participants — particularly those in the 35–50 group — were, in several cases, less able to engage with the application for extended periods due to professional and family commitments, which may have affected the depth of their experience with the product. As noted in sampling methodology literature, those who participate in convenience samples are often more

available, cooperative, or motivated than those who do not — a dynamic that directly shaped the age distribution observed here (Etikan et al., 2016).

The sample was 67% male ($n = 14$) and 33% female ($n = 7$). Fitness experience was distributed across all self-reported levels: 7 participants (33%) described themselves as irregular or on-and-off, 6 (29%) as experienced (2+ years), 5 (24%) as intermediate (6 months–2 years), and 3 (14%) as beginners. Twelve participants (57%) had prior experience with other fitness applications. Among those who specified an avatar choice, 11 selected the Wolf companion and 7 selected the Bunny; 3 did not record a preference.

Motivation and Engagement

Across the five motivation and engagement items, mean scores ranged from 4.05 to 4.24, yielding a section mean of 4.18. No participant rated any motivation item below 2, and the modal response across the section was 5 (Strongly Agree).

Item	Mean	Median	Min	Max	n
Using GymBRO made me more motivated to exercise	4.05	4.0	3	5	21
I felt in control of my own fitness routine (Autonomy)	4.24	4.0	2	5	21
Avatar progress made my efforts feel meaningful (Competence)	4.24	5.0	3	5	21
I felt a sense of connection to my avatar companion (Relatedness)	4.19	4.0	2	5	21
I opened the app more often out of curiosity about my avatar	4.19	5.0	2	5	21

Table 1. Motivation and Engagement — Descriptive Statistics ($N = 21$). Source: Author.

Avatar Companion Experience

The avatar section produced a section mean of 4.22. The highest-scoring item was "The avatar's visual changes reflected my real activity in a meaningful way" ($M = 4.48$, median = 5.0) — 12 of 21 participants selected the maximum score. The sense of felt responsibility toward the avatar scored somewhat lower ($M = 3.95$), with three participants rating it at 2,

suggesting that the Tamagotchi-style obligation mechanic was experienced differently across users.

Item	Mean	Median	Min	Max	n
Avatar visual changes reflected real activity meaningfully	4.48	5.0	3	5	21
I felt responsible for taking care of my avatar	3.95	4.0	2	5	21
Avatar made experience more personal than a regular fitness app	4.24	4.0	2	5	21

Table 2. Avatar Companion Experience — Descriptive Statistics (N = 21). Source: Author.

Nine participants left open-ended responses about the avatar's effect on their experience. Responses were largely positive and touched on emotional engagement and identity. One participant wrote that they felt more progress "because of connection with my pet, which I was seeing as myself" — a self-projection dynamic consistent with Tamagotchi effect literature (Frude & Jandric, 2015). A female participant described the psychological comfort of avatar interaction: communicating with the avatar felt like talking to "a real hero, albeit a virtual one," rather than interacting with a blank interface. Another noted the avatar motivated them to "be the same or even better" than their companion's current state — an external standard internalised as a personal goal. Critical notes were mostly aesthetic: one respondent found the wolf avatar's humanoid design somewhat unsettling, while overall finding the application positive.

AI Chat

The AI chat section produced the lowest section mean of the four thematic areas (M = 4.02), though all scores remained above 4.0. Context-awareness — "It felt like the AI knew my context" — was the single weakest item in the entire survey (M = 3.95), though it still exceeded the scale midpoint.

Item	Mean	Median	Min	Max	n
AI chat responses felt relevant to my actual fitness situation	4.05	4.0	3	5	21

It felt like the AI "knew" my context (workouts, food, avatar state)	3.95	4.0	3	5	21
AI chat motivated me to take action	4.05	4.0	3	5	21

Table 3. AI Chat — Descriptive Statistics (N = 21). Source: Author.

Open-ended AI feedback showed a consistent pattern: participants found the chat useful for specific queries — workout generation, nutrition questions — but noted that responses sometimes felt generic or overly long. One participant described the experience as "too much text, it was painful." Another observed that responses "at times felt a bit generic and could be more tailored to my specific goals." One participant noted they simply did not use the AI chat much, pointing to a discoverability issue rather than quality dissatisfaction.

Overall Usability

Usability produced the highest section mean of 4.52. The recommendation item reached M

Item	Mean	Median	Min	Max	n
The app was easy to use overall	4.48	5.0	3	5	21
I would recommend GymBRO to a friend who wants to build fitness habits	4.57	5.0	3	5	21

= 4.57, median = 5.0 — 14 of 21 participants selected the maximum score, and no participant rated below 3.*Table 4. Overall Usability — Descriptive Statistics (N = 21). Source: Author.*

Improvement suggestions clustered around three themes: expanding the avatar roster — multiple participants independently requested additional companion types (hamster, dog, cat) and cosmetic customisation options; improving nutrition flexibility (fully customisable meal counts and calorie targets); and refining the AI chat interface (shorter responses, floating button access rather than a dedicated screen). One participant provided a detailed and technically informed list of UX suggestions, including a reference to the Zeigarnik

effect in the context of progress bars — indicating that the sample included participants with design or psychology background.

3.4 Discussion

The survey results support the thesis hypothesis: across all four thematic areas, mean scores remained above 4.0, no item produced a rating below 2, and 14 of 21 participants indicated they would recommend the application to others. The pattern of responses reflects genuine differentiation rather than uniform approval — AI chat consistently scored lower than avatar and usability items, which is the pattern one would expect from honest evaluation of a product with clearly stronger and weaker components.

The SDT framework provides a useful structure for interpreting the motivation results. Autonomy ($M = 4.24$) — the sense of personal control over one's fitness routine — was well-supported. The offline-first architecture, which places all tracking decisions entirely in the user's hands, appears to have contributed to this. Competence ($M = 4.24$, median = 5.0) — the feeling that effort produces visible, meaningful results — was among the strongest individual motivation items and had the highest median in the section, consistent with the living avatar's core design: the companion's visual state is a direct, unambiguous signal of the user's fitness quality. Relatedness ($M = 4.19$) — felt connection to the avatar companion — was supported both quantitatively and through qualitative responses that described the avatar in relational and identity terms. Together, these results suggest the living avatar design addressed all three SDT needs as theorised in Chapter 1.

The avatar section's qualitative data is particularly informative. The self-projection dynamic — participants describing the avatar as a version of themselves — points to a deeper engagement mechanism than simple gamification. This is closer to what Reeves and Nass (1996) described in the Media Equation: users treated the avatar as a genuine social presence, not a decorative UI element. The lower score on felt responsibility ($M = 3.95$) suggests that the obligation dimension of the Tamagotchi effect was not uniformly activated — some users experienced the avatar as a motivating companion, others as a progress indicator. This variation is likely related to frequency and depth of engagement with the application.

The AI chat results highlight the most significant area for improvement. A section mean of 4.02 is positive, but the context-awareness item ($M = 3.95$) being the weakest point in the survey is informative. Users noticed when responses felt generic — and several said so

directly. This is consistent with Aggarwal et al.'s (2023) finding that perceived personalisation is the hardest criterion for AI health assistants to meet. The gap between "the AI has my data" and "the AI really understands my situation" is a design challenge that extends beyond prompt engineering into conversation architecture and response length calibration.

On a personal note, the overall pattern of responses reflects what was observed informally throughout development: the majority of users were genuinely pleased with the experience and engaged with the avatar in ways that went beyond purely functional use, while simultaneously noticing real limitations — occasional performance issues, AI responses that did not always land with full contextual precision, and navigation elements that could be clearer. This combination — enthusiasm about the concept paired with honest criticism of the execution — is more valuable as evaluation data than uniformly positive responses would be. It points toward what needs to be improved rather than suggesting the concept itself is flawed.

The demographic skew toward younger participants, while a sampling limitation, is also informative. The 18–24 group showed comfortable familiarity with AI interaction and avatar-based applications. The smaller group of 35–50 participants, despite limited engagement time due to professional commitments, included some of the most detailed qualitative responses — suggesting the concept has cross-generational appeal when participants have sufficient time to explore the product.

Industry context: The relevance of the living avatar concept extends beyond this prototype. In April 2026, Anthropic's source code for Claude Code was accidentally exposed through a packaging error in version 2.1.88 of the npm package, revealing a fully implemented Tamagotchi-style virtual pet system — named "Buddy" — built into their professional AI coding assistant. Developers analysing the leaked codebase found a complete pet simulation with 18 species, a rarity system, LLM-generated personality, and stat mechanics that respond to the user's coding session activity (InfoQ, 2026; VentureBeat, 2026). The feature was designed to humanise the terminal environment and sustain engagement through emotional companionship rather than purely functional utility — the same mechanism at the core of GymBRO's living avatar. The fact that one of the leading AI research organisations independently arrived at the same design intuition provides external validation for the theoretical basis of this thesis. The domain differs — Claude Buddy lives in a developer terminal, GymBRO's avatar lives in a fitness application — but the underlying mechanism is identical. This convergence suggests the living avatar concept is

not a niche experiment but a direction the broader AI industry is beginning to take seriously. The question for future development is not whether this approach works in principle, but how to implement it correctly: combining contextual intelligence, emotional presence, and meaningful behavioural feedback in a way that feels genuine. GymBRO represents one early answer to that question in the fitness domain.

3.5 Limitations of the Study

The study has several limitations that should be acknowledged. First, the sample relies on purposive convenience sampling from the author's personal network and conference contacts, which introduces potential social desirability bias and limits the generalisability of findings to broader populations. Second, the testing period is short relative to what would be required to assess long-term habit formation effects. Third, the self-report format of the survey means that perceived motivation is measured rather than objective behavioural outcomes such as exercise frequency or duration. These limitations are inherent to prototype evaluation within a bachelor's thesis scope and will be discussed fully in the final version of this chapter alongside the complete dataset.

CONCLUSIONS

GymBRO began as a personal question and became a working product: a fitness application implementing a living avatar concept in which an AI assistant is embodied as a context-aware virtual companion whose visual state is directly driven by the user's real fitness behaviour. The project was developed over six months within the Spinner24 startup programme at Idea Virumaa in Narva, won first place at the programme's competition, and is documented here as both a working prototype and an academic investigation into whether that relational layer produces a measurable effect on user motivation.

The theoretical review in Chapter 1 established that combining relationship-oriented gamification, emotional companion attachment, and personalised AI assistance within an SDT framework provides a well-grounded basis for the living avatar concept. Each component addresses documented limitations of the others: gamification provides behavioural structure, the virtual companion provides emotional anchoring, and the AI provides contextual intelligence.

Chapter 2 described the technical realisation of this concept. The GymBRO application was developed as a cross-platform React and FastAPI system with an offline-first

architecture, integrating an XP-based avatar system, a contextual AI chat powered by GPT-4o-mini, and tracking modules for nutrition, training, and sleep. The system architecture demonstrated that the living avatar concept is technically feasible within a realistic development scope, and the analytics infrastructure provided a foundation for measuring real user behaviour.

The evaluation of the application through user testing is described in Chapter 3. A total of 21 participants completed the survey between March and May 2026, drawn from the author's personal network, the Spinner24 programme community at W3N in Narva, and gym contacts introduced to the application during informal testing. Across all four thematic areas — motivation and engagement, avatar companion experience, AI chat, and overall usability — mean scores remained above 4.0 on a five-point scale. The living avatar concept received the strongest support, with avatar visual feedback scoring highest ($M = 4.48$) and qualitative responses revealing self-projection and emotional companionship dynamics consistent with the Tamagotchi effect. The AI chat was the weakest area ($M = 4.02$), with context-awareness identified as the primary direction for improvement. Fourteen of 21 participants indicated they would recommend the application, providing a direct answer to the research question: the living avatar concept produced measurable positive effects on self-reported motivation, avatar attachment, and perceived application value.

In terms of the stated objectives, the first objective — reviewing existing research — was fulfilled through the systematic theoretical review in Chapter 1, covering 25 academic sources across gamification, virtual companion, motivation, and AI chatbot research. The second objective — designing and developing the application — was fulfilled as documented in Chapter 2. The third objective — evaluating the application through user testing — was addressed through the survey study described in Chapter 3.

Several limitations of this work should be acknowledged. The sample size is small and composed of individuals from the author's personal and professional network, which limits the generalisability of the findings. The testing period was short relative to what would be required to assess long-term motivation effects. The self-reported nature of survey data introduces social desirability bias. These limitations are inherent to prototype evaluation within a bachelor's thesis scope and point toward the research directions described below.

Future development of GymBRO should prioritise four areas. First, cloud synchronisation of local records would enable multi-device use and make the application viable for a broader user population. Second, expanding the avatar system to support multiple companions with more granular emotional states and visual variety would deepen the emotional engagement documented in the Tamagotchi effect literature. Third, deepening the AI personalisation — incorporating longitudinal pattern recognition and proactive recommendations based on the user's historical data — would move the system closer to a genuine AI coaching relationship. Fourth, a longitudinal controlled study with a larger and more diverse sample would provide the empirical evidence needed to make statistically generalisable claims about the living avatar's effect on fitness motivation.

The GymBRO project demonstrates that the living avatar concept is both theoretically grounded and technically achievable, and provides a working prototype as a foundation for future research and development in AI-enhanced gamified fitness applications.

RESÜMEE

Neliubov, D. (2026). GymBRO: tehisintellektil põhineva treeningrakenduse arendamine ja hindamine kontekstitundliku avatari kaaslasega. Bakalaureusetöö, Tartu Ülikool, Narva kolledž.

Käesoleva bakalaureusetöö eesmärk oli kavandada, arendada ja hinnata treeningrakendust, mis ühendab tehisintellektil põhineva vestlusagendi visuaalse avatari kaaslasega, kelle olek peegeldab kasutaja tegelikku treeningkäitumist. Rakendus GymBRO loodi vastusena mobiilsete treeningrakenduste kroonilisele probleemile — kasutajad võtavad need entusiastlikult kasutusele, kuid loobuvad neist nädalate jooksul.

Teoreetiline ülevaade hõlmas nelja omavahel seotud valdkonda: mängustamine tervishoiu- ja treeningrakendustes, virtuaalsed kaaslased ja emotsionaalne kiindumus, motivatsiooniteooriad (eelkõige enesemääratlusteooria) ning tehisintellektil põhinevad vestlusagendid mobiilses tervishoius. Nende valdkondade süntees moodustas „elava avatari“ kontseptsiooni teoreetilise aluse.

Rakendus arendati Agile metoodika abil kuue kuu jooksul (oktoober 2025 – aprill 2026) kasutades React, FastAPI ja OpenAI GPT-4o-mini tehnoloogiaid. Süsteemiarhitektuur on offline-first põhimõttel, kus kohalik andmebaas (IndexedDB) salvestab kasutajaandmed ja server teenindab ainult tehisintellekti päringuid ning analüütikat.

Rakendust hinnati kasutajatesti kaudu, mille raames koguti andmeid struktureeritud küsimustikuga Google Forms'i vahendusel. Küsimustik koosnes 19 küsimusest, mis hõlmasid motivatsiooni, avatari kogemust, tehisintellekti vestlust ja üldist kasutatavust. Tulemused esitatakse lõpptöös pärast andmekogumise lõpetamist.

Märksõnad: mängustamine, virtuaalne kaaslane, treeningrakendus, tehisintellekti assistent, motivatsioon, kasutajate kaasamine, enesemääratlusteooria, elav avatar

REFERENCES

- Aggarwal, A., Tam, C. C., Wu, D., Li, X., & Qiao, S. (2023). Artificial intelligence–based chatbots for promoting health behavioral changes: Systematic review. *Journal of Medical Internet Research*, 25, e40789. <https://doi.org/10.2196/40789>
- Althoff, T., White, R. W., & Horvitz, E. (2016). Influence of Pokémon Go on physical activity: Study and implications. *Journal of Medical Internet Research*, 18(12), e315. <https://doi.org/10.2196/jmir.6759>
- Bitrián, P., Buil, I., & Catalán, S. (2021). Enhancing user engagement: The role of gamification in mobile apps. *Journal of Business Research*, 132, 170–185. <https://doi.org/10.1016/j.jbusres.2021.04.028>
- Deci, E. L., & Ryan, R. M. (1985). *Intrinsic motivation and self-determination in human behavior*. Plenum.
- Deci, E. L., & Ryan, R. M. (2000). The "what" and "why" of goal pursuits: Human needs and the self-determination of behavior. *Psychological Inquiry*, 11(4), 227–268. https://doi.org/10.1207/S15327965PLI1104_01
- Deci, E. L., Koestner, R., & Ryan, R. M. (1999). A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation. *Psychological Bulletin*, 125(6), 627–668. <https://doi.org/10.1037/0033-2909.125.6.627>
- Deterding, S., Dixon, D., Khaled, R., & Nacke, L. (2011). From game design elements to gamefulness: Defining "gamification". *Proceedings of the 15th International Academic MindTrek Conference*, 9–15. <https://doi.org/10.1145/2181037.2181040>
- Duhigg, C. (2012). *The power of habit: Why we do what we do in life and business*. Random House.
- Etikan, I., Musa, S. A., & Alkassim, R. S. (2016). Comparison of convenience sampling and purposive sampling. *American Journal of Theoretical and Applied Statistics*, 5(1), 1–4. <https://doi.org/10.11648/j.ajtas.20160501.11>
- Frude, N., & Jandric, P. (2015). The intimate machine: Social robots and the nature of attachment. *E-Learning and Digital Media*, 12(3–4), 411–424. <https://doi.org/10.1177/2042753015592616>

- Fu, L., Burns, R., Xie, Y., Shen, J., Zhe, S., Estabrooks, P., & Bai, Y. (2026). The development and use of AI chatbots for health behavior change: Scoping review. *Journal of Medical Internet Research*, 28, e79677. <https://doi.org/10.2196/79677>
- Giasiranis, S., & Sofos, A. (2025). Purr-ogrammed love: A narrative review of virtual pets. *Computers in Human Behavior Reports*, 18, Article 100382. <https://doi.org/10.1016/j.chbr.2025.100382>
- Huang, H., Chen, Y., & Zhou, Y. (2024). Motivation crowding effects on the intention for continued use of gamified fitness apps: A mixed-methods approach. *Frontiers in Psychology*, 14, Article 1286463. <https://doi.org/10.3389/fpsyg.2023.1286463>
- Koivisto, J., & Hamari, J. (2019). The rise of motivational information systems: A review of gamification research. *International Journal of Information Management*, 45, 191–210. <https://doi.org/10.1016/j.ijinfomgt.2018.10.013>
- Lister, C., West, J. H., Cannon, B., Sax, T., & Brodegard, D. (2014). Just a fad? Gamification in health and fitness apps. *JMIR Serious Games*, 2(2), e9. <https://doi.org/10.2196/games.3413>
- McFadyen, J., Habicht, J., Dina, L.-M., Harper, R., Hauser, T. U., & Rollwage, M. (2024). AI-enabled conversational agent increases engagement with cognitive-behavioral therapy: A randomized controlled trial [Preprint]. *medRxiv*. <https://doi.org/10.1101/2024.11.01.24316565>
- Middelweerd, A., Mollee, J. S., van der Wal, C. N., Brug, J., & te Velde, S. J. (2014). Apps to promote physical activity among adults: A review and content analysis. *International Journal of Behavioral Nutrition and Physical Activity*, 11, Article 97. <https://doi.org/10.1186/s12966-014-0097-9>
- Reeves, B., & Nass, C. (1996). *The media equation: How people treat computers, television, and new media like real people and places*. Cambridge University Press.
- Ryan, R. M., & Deci, E. L. (2000). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American Psychologist*, 55(1), 68–78. <https://doi.org/10.1037/0003-066X.55.1.68>

- Tartu Ülikool. (2023). Tartu Ülikooli suunis tekstiroboti kasutamiseks õppetöös. <https://ut.ee/et/sisu/tartu-ulikooli-suunis-tekstiroboti-kasutamiseks-oppetoos>
- Teixeira, P. J., Carraça, E. V., Markland, D., Silva, M. N., & Ryan, R. M. (2012). Exercise, physical activity, and self-determination theory: A systematic review. *International Journal of Behavioral Nutrition and Physical Activity*, 9, Article 78. <https://doi.org/10.1186/1479-5868-9-78>
- Xu, J., Beauchamp, M. R., & Rhodes, R. E. (2024). Effect of digital health applications with or without gamification on physical activity: A systematic review and meta-analysis. *eClinicalMedicine*, 75, Article 102789. <https://doi.org/10.1016/j.eclinm.2024.102789>
- Xu, J., Lio, A., Dhaliwal, H., Andrei, I., Balakrishnan, S., Nagani, U., & Sameer, H. (2021). Evaluating the effectiveness of gamification on physical activity: Systematic review and meta-analysis of randomized controlled trials. *Journal of Medical Internet Research*, 24(1), e24260. <https://doi.org/10.2196/24260>
- InfoQ. (2026, April 7). Anthropic accidentally exposes Claude Code source via npm source map file. <https://www.infoq.com/news/2026/04/claude-code-source-leak/>
- VentureBeat. (2026, April 1). Claude Code's source code appears to have leaked: <https://venturebeat.com/technology/claude-codes-source-code-appears-to-have-leaked-heres-what-we-know/>

APPENDICES

Appendix A. User Survey Instrument

The following survey was distributed via Google Forms to participants who had prior experience using the GymBRO application. All responses were anonymous.

Survey link: <https://forms.gle/uWEiSfKvXvSqNYB6A>

Title: GymBRO App — User Experience Survey

Description: This survey is part of a bachelor's thesis research at the University of Tartu. It takes approximately 5–7 minutes to complete. Your responses are anonymous and will be used solely for academic purposes.

SECTION 1 — Background

Q1. What is your age range? [Under 18 / 18–24 / 25–34 / 35 or older]

Q2. What is your sex? [Male/Female]

Q3. How would you describe your fitness experience? [Beginner (less than 6 months) / Intermediate (6 months – 2 years) / Experienced (2+ years)]

Q4. Have you used other fitness apps before GymBRO? [Yes / No]

SECTION 2 — Motivation and Engagement

Rate each statement from 1 (Strongly Disagree) to 5 (Strongly Agree).

Q4. Using GymBRO made me more motivated to exercise.

Q5. Using GymBRO made me feel in control of my own fitness routine.

Q6. Seeing my avatar's progress made me feel that my efforts were making a difference.

Q7. I felt a sense of connection to my avatar companion.

Q8. I opened the app more often because I was curious about my avatar's state.

SECTION 3 — Avatar Companion

Rate each statement from 1 (Strongly Disagree) to 5 (Strongly Agree).

Q9. The avatar's visual changes (growth, states) reflected my real activity in a meaningful way.

Q10. I felt responsible for taking care of my avatar.

Q11. The avatar made the experience feel more personal compared to a regular fitness app.

Q12. (Open-ended) In your own words — how did the avatar affect your experience using the app? (if at all)

SECTION 4 — AI Chat

Rate each statement from 1 (Strongly Disagree) to 5 (Strongly Agree).

Q13. The AI chat responses felt relevant to my actual fitness situation.

Q14. It felt like the AI "knew" my context (my workouts, food, avatar state).

Q15. The AI chat motivated me to take action (exercise, eat better, sleep more).

Q16. (Open-ended) What did you find most useful — or least useful — about the AI chat?

SECTION 5 — Overall

Rate each statement from 1 (Strongly Disagree) to 5 (Strongly Agree).

Q17. The app was easy to use overall.

Q18. I would recommend GymBRO to a friend who wants to build fitness habits.

Q19. (Open-ended) What would you improve or add to GymBRO?

Appendix B. Key System Prompt Structure

The AI system prompt assembled by `generateSystemPrompt()` contains the following sections in order:

- Core identity — defines the pet name, dual role (companion + coach), and the principle that user fitness activity is an act of caring for the pet.
- Voice rules — specifies tone (warm, alive, affectionate, concrete), forbidden phrases ('fitness journey', 'overall well-being', generic wellness slogans), and emoji usage policy (zero to one, never for decoration).

- Emotional modes — four states: Proud, Comfort, Playful, Coach, each with defined behavioural rules.
- Pet-care relationship — explicit connection between training/nutrition/sleep and avatar well-being.
- Response style — length (2–5 sentences), structure (answer first, one next step maximum, one follow-up question maximum).
- Knowledge contract — prohibition on inventing data not present in the current context.
- Activity context — level-based coaching intensity (generated by `buildActivityContext()`).
- 14 dialogue examples — covering identity, praise, comfort, coaching, data awareness, and missing data handling.
- Current context block — live user data: level, XP, streak, training status, schedule, pet care score, today's food, 7-day calorie history, sleep data, user profile.

Appendix C. Source Code Repository

The full source code of the GymBRO application is maintained in a private GitHub repository. The repository contains the React frontend, FastAPI backend, avatar logic, AI integration, and analytics system. The authentication module, developed as shared infrastructure by a co-author, is maintained separately and is not part of this repository. The repository will be made available for review upon request to the thesis supervisor.

LIHTLITSENTS

Lihtlitsents lõputöö reprodutseerimiseks ja üldsusele kättesaadavaks tegemiseks

Mina, Denys Neliubov,

1. annan Tartu Ülikoolile tasuta loa (lihtlitsentsi) minu loodud teose

GYMBRO: DEVELOPMENT AND EVALUATION OF AN AI-POWERED FITNESS APPLICATION WITH A CONTEXT-AWARE AVATAR COMPANION,

mille juhendaja on Nicolai Morozov,

reprodutseerimiseks eesmärgiga seda säilitada, sealhulgas lisada digitaalarhiivi DSpace kuni autoriõiguse kehtivuse lõppemiseni.

2. Annan Tartu Ülikoolile loa teha punktis 1 nimetatud teos üldsusele kättesaadavaks Tartu Ülikooli veebikeskkonna, sealhulgas digitaalarhiivi DSpace kaudu Creative Commons'i litsentsiga CC BY NC ND 4.0, mis lubab autorile viidates teost reprodutseerida, levitada ja üldsusele suunata ning keelab luua tuletatud teost ja kasutada teost ärieesmärgil, kuni autoriõiguse kehtivuse lõppemiseni.

3. Olen teadlik, et punktides 1 ja 2 nimetatud õigused jäävad alles ka autorile.

4. Kinnitan, et lihtlitsentsi andmisega ei riku ma teiste isikute intellektuaalomandi ega isikuandmete kaitse õigusaktidest tulenevaid õigusi.

Denys Neliubov

05.04.2026