Politecnico di Torino

Master of Science in Computer Engineering

A.Y. 2024/2025

Graduation Session DECEMBER 2025

Neurological Consequences of COVID-19

Supervisors

Dr. Sacco ALESSIO

Dr. Marchetto GUIDO

Dr. Vito DE FEO

Dr. Cristina DEL PRETE

Candidate

Otabek FAYZIEV

Abstract

This thesis was developed within a collaborative project between the University of Essex and Politecnico di Torino and focuses on strengthening the digital infrastructure of the Happy Again platform (https://happyagain.essex.ac.uk/). The platform is a web-based research tool designed to investigate the long-term neurocognitive consequences of COVID-19 by collecting behavioural and cognitive markers through online tasks and questionnaires. It enables the assessment of attention, perception, response timing, and cognitive processing in individuals experiencing post-COVID conditions.

The work presented in this thesis involved coordinated development across backend and frontend components to improve system stability, data integrity, and research reproducibility. On the backend, Docker-based containerization was introduced to ensure consistent deployment, the configuration architecture was refactored, APIs and data models were updated, new cognitive and timing indicators (including task-specific `lc_flag` values) were integrated, and robustness of registration and email delivery workflows was enhanced. On the frontend, the administrative data-export module was restructured, unified pipelines for processing task results were implemented, new timing-based and derived metrics were added, and multiple optimizations were applied to improve reliability and performance.

These developments reinforced the platform as a reliable and scalable research environment, enabling more structured and accurate data collection and supporting ongoing investigations into the neurological impact of COVID-19 and future applications in cognitive monitoring and rehabilitation.

To my family, and to the people who became family.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

The long-term neurological and cognitive effects of COVID-19 have become an important focus of research. A significant number of individuals who recovered from the acute phase of the infection continue to report persistent symptoms such as slower reaction times, reduced attention, impaired perceptual integration, and general cognitive fatigue. These symptoms, often categorised under the term *long COVID*, vary substantially across individuals and are difficult to measure objectively. For this reason, digital behavioural platforms capable of collecting consistent, structured, and reproducible data have become essential tools for large-scale studies in this field.

The *Happy Again* platform was developed as part of a collaboration between the University of Essex and Politecnico di Torino to support such research efforts. The system provides an online environment where participants complete cognitive tasks and questionnaires designed to capture behavioural markers related to perception, timing, memory, and COVID-19–related self-reported symptoms. Because all tasks are delivered through a web browser, the platform allows researchers to collect data remotely and at scale, without requiring laboratory conditions. This approach offers a practical way to observe cognitive patterns in diverse populations, including those affected by long COVID.

Although the core platform was functional, several technical limitations affected its stability, the precision of timing-sensitive tasks, and the consistency of the collected datasets. Behavioural experiments are sensitive to even small variations in frontend execution, backend processing, or deployment configuration. Differences between development and production environments, inconsistencies in timestamp structures, heterogeneous encoding rules, and non-uniform data models all contributed to methodological noise and additional preprocessing effort for researchers.

Given these challenges, the main goal of this thesis was to strengthen the technological foundation of the *Happy Again* platform. The work focused on improving backend reliability, refactoring the frontend architecture, standardising data processing, and restructuring deployment workflows. The objective was not to change the scientific content of the cognitive tasks, but to ensure that the platform operates predictably, produces consistent data, and can support future research and extensions.

The contributions developed in this thesis include:

- containerised backend deployment using Docker to ensure reproducible environments;
- refactoring of configuration management and environment handling;
- updates to cognitive-task data models, including unified timing indicators and support for long-COVID classification flags;
- improvements to API endpoints, validation mechanisms, and error handling;
- extensive refactoring of frontend task components and admin dashboard performance;
- unified schema and shared processing logic across all cognitive tasks;
- a standardised timestamp model applied consistently across experiments;
- a restructured data export system with stable column ordering and decoded values;
- improvements to dataset integrity through frontend and backend validation;
- enhanced maintainability and clearer architectural structure across the platform.

These developments improve the reliability of behavioural data collection, reduce the likelihood of incomplete or corrupted submissions, minimise preprocessing requirements for researchers, and establish a more scalable architecture for future studies.

The structure of the thesis is outlined below:

- **Chapter 2** provides background information on long-COVID research, its neurological and cognitive effects, the challenges in measuring these symptoms, and the role of digital cognitive assessment platforms in addressing these challenges.
- **Chapter 3** presents an overview of the structure and functionality of the *Happy Again* platform, including the participant workflow, administrative interface, backend and frontend architecture, database structure, and deployment management.

- **Chapter 4** describes the technical work carried out on the backend components, including infrastructure improvements through Docker containerisation, configuration system refactoring, updates to data models, enhancements to API endpoints, and stability improvements.
- **Chapter 5** focuses on frontend improvements, covering the architecture of the Angular application, changes to participant-facing experimental tasks, optimisation and refactoring of the administrative dashboard, and enhancements to the data submission workflow.
- **Chapter 6** details the data-processing improvements implemented across the platform, including the unification of processing logic, standardisation of timestamps and timing indicators, consistent encoding and decoding of task responses, and enhancements to the administrative export system.
- **Chapter 7** presents the results and impact of these developments, highlighting increased system stability, improved data accuracy and completeness, more efficient research workflows, enhanced reproducibility, and better maintainability for future extensions.
- **Chapter 8** discusses the limitations of the current system, including constraints inherent to browser-based testing, the scope of available cognitive tasks, potential data completeness issues, and areas requiring further harmonisation.
- **Chapter 9** provides the overall conclusion and outlines possible directions for future work, including the expansion of cognitive tasks, improvements to timing measurement accuracy, further automation in deployment and maintenance, integration with external data sources, and enhancements to administrative and analytical workflows.

The improvements introduced in this thesis provide a more stable, predictable, and scientifically reliable platform for long-COVID behavioural research, ensuring that future studies can be conducted with greater methodological confidence.

Chapter 2

Background

2.1 Neurological and Cognitive Effects of COVID-19

Since the beginning of the COVID-19 pandemic, it has become clear that the virus can affect not only the respiratory system but also the central nervous system and cognitive functions. A substantial number of individuals report persistent symptoms after the acute phase of the infection has resolved, including concentration problems, slower thinking, memory difficulties and a general sense of "mental fog". Large-scale online studies have shown measurable deficits in attention, reasoning and problem-solving in people who recovered from COVID-19 compared with non-infected controls, even months after the illness.[1] Other clinical and neuropsychological investigations confirm that long-term neurocognitive symptoms are a frequent component of long COVID.[2]

These cognitive difficulties are typically heterogeneous. Some participants report only mild problems in everyday life, while others describe substantial impairment in work or study. In formal testing, deficits can appear in domains such as sustained attention, executive control, working memory, processing speed and cognitive inhibition.[2] Symptom severity may vary over time, and individuals sometimes report good days and bad days, which makes the condition difficult to characterise with single-point assessments.

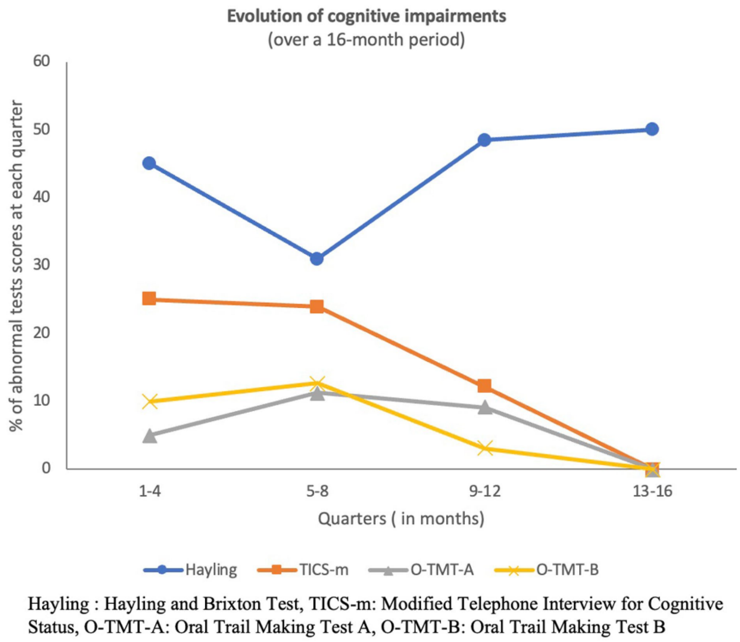

Figure 2.1: Evolution of cognitive impairments in long-COVID patients across a 16-month period, based on multiple neuropsychological tests (adapted from Saucier et al. [3]).

Figure 2.1 shows how the proportion of abnormal scores across several neuropsychological tests changes over a 16-month period in individuals with long COVID. The results demonstrate persistent impairments in executive functions, inhibition, processing speed, and related cognitive domains, even more than a year after infection.

These findings underline the need for accurate and scalable assessment methods, which introduces several methodological challenges discussed in the next section.

2.2 Challenges in Measuring Long-COVID Cognitive Symptoms

Studying the cognitive aspects of long COVID presents several methodological challenges. Traditional neuropsychological assessments require in-person testing, dedicated equipment and trained examiners. They provide high-quality measurements but are time-consuming, expensive and difficult to scale to large populations or to repeated follow-up sessions. This is particularly problematic for long-COVID cohorts, where participants may experience fatigue, mobility issues or limited ability to attend regular laboratory visits.

Another difficulty arises from the partly subjective nature of many complaints. People often describe feeling “slower” or “less focused”, but such impressions depend on self-evaluation and are influenced by stress, sleep, mood and other factors. Objective assessment requires tasks that produce quantifiable behavioural markers such as reaction times, accuracy rates or error patterns. For example, inhibition deficits may manifest as slower or less accurate responses in tasks that require suppressing automatic reactions.[2] However, obtaining reliable behavioural markers outside controlled laboratory environments is technically demanding.

Long-COVID symptoms also tend to fluctuate over days and weeks. A single test session may therefore underestimate or overestimate the actual degree of impairment. For longitudinal research it is useful to collect data repeatedly over longer periods, which is logistically difficult with purely clinic-based approaches. These limitations have motivated a growing interest in remote and digital methods for cognitive assessment.

The participant selection process used in long-COVID studies is illustrated in Figure 2.2.

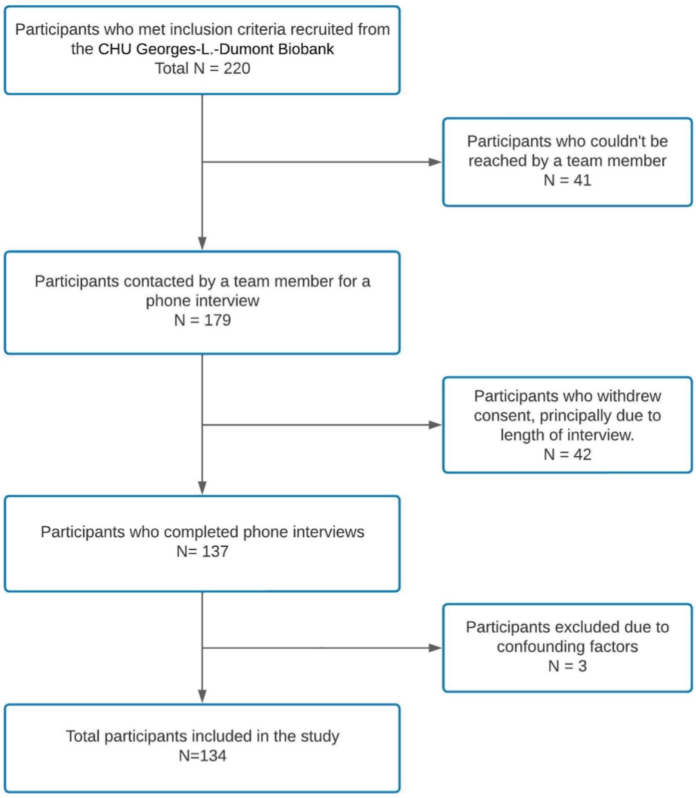

Figure 2.2: Participant recruitment and inclusion flowchart in long-COVID research (adapted from Saucier et al., 2023).

Figure 2.2 illustrates the recruitment and inclusion process commonly reported in long-COVID research. The diagram highlights several methodological challenges: a substantial proportion of participants cannot be reached, many withdraw due to the duration of interviews, and additional cases must be excluded because of confounding factors. As a result, the final sample is significantly smaller than the initially eligible population, which underscores the practical limitations of long-COVID cohort studies.

2.3 Digital Cognitive Assessment Platforms

The increasing availability of web technologies has made it possible to administer cognitive tasks remotely using only a browser. Web-based cognitive platforms allow participants to complete experiments from home, which simplifies recruitment, reduces travel requirements and lowers the cost of data collection. In psychiatric and neurological research, several studies have shown that carefully designed web-based batteries can achieve acceptable validity compared to traditional in-person tests.[4]

Typical digital platforms implement tasks that measure reaction time, accuracy, memory performance or attentional control. Because the procedures are standardised, results from different participants and sessions can be compared more easily. In addition, digital tasks automatically record detailed behavioural traces, including timestamps and trial-level responses, which can later be used to compute derived indicators such as variability, learning effects or fatigue.

Table 2.1 summarises key differences between traditional clinical assessments and web-based cognitive platforms. Digital tools are not a replacement for full clinical neuropsychological examinations, but they are well suited for large-scale and longitudinal studies where frequent testing is required.

Table 2.1: Comparison between traditional clinical assessments and digital cognitive platforms.

Feature	**Clinical neuropsychological tests**	**Digital cognitive tasks**
Testing environment	In-person, controlled lab	Browser-based, remote
Cost	High	Low
Repeated sessions	Difficult	Easy, automated
Sample size	Limited	Large-scale
Timing accuracy	High	Medium–High (depends on implementation)
Behavioural markers	Broad cognitive battery	Reaction times, accuracy, derived indicators
Scalability	Low	High

At the same time, web-based experiments introduce their own technical constraints. Timing precision can be influenced by the browser, operating system and hardware, and must be validated carefully.[4] Network delays or server-side bottlenecks can also introduce noise in the data if task logic and storage routines are not implemented correctly. For reaction-time based research, such issues can reduce the interpretability of results if they are not properly controlled.

Figure 2.3 shows a typical workflow for digital cognitive assessments. Participants access the system through a web interface, complete consent procedures, fill in questionnaires and perform cognitive tasks. Their responses are then transmitted to backend services, where they are stored and prepared for statistical analysis. The architecture of the platform must ensure that this pipeline operates reliably and consistently across sessions and users.

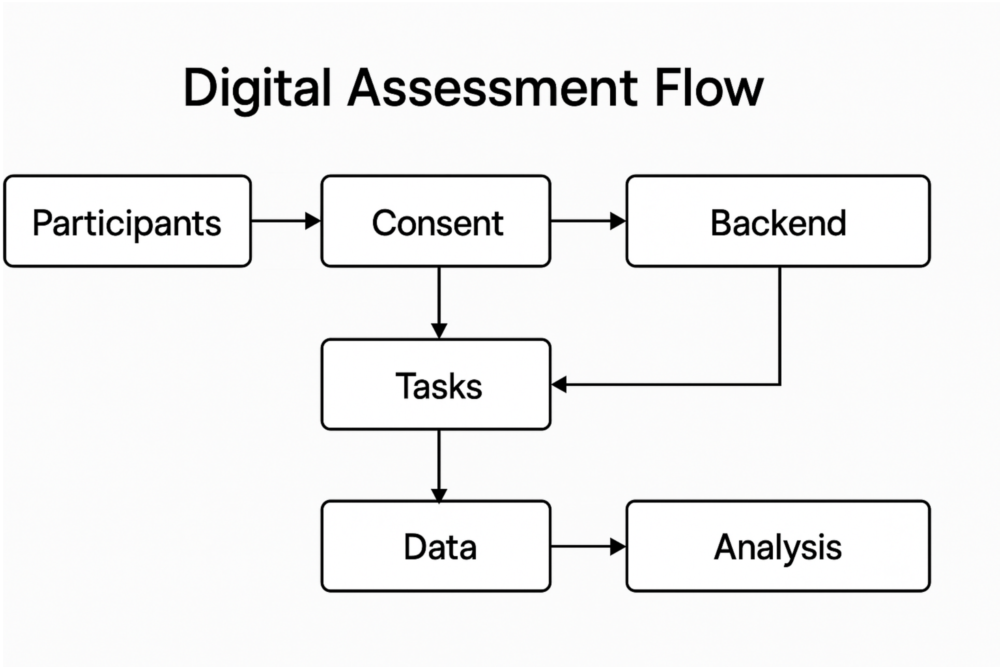

Figure 2.3: Generic workflow of a digital cognitive assessment platform, illustrating data flow from participant access to storage and analysis.

2.4 Motivation for the Happy Again Platform

Within this context, the *Happy Again* project was created as a collaborative effort between the University of Essex and Politecnico di Torino to support research on the long-term neurocognitive consequences of COVID-19. The platform provides a web-based environment in which participants complete cognitive tasks and questionnaires related to their COVID-19 history and current symptoms. By combining behavioural markers (such as reaction times and accuracy) with self-reported information, the system is intended to capture subtle changes in cognitive functioning that may be associated with long COVID.

A browser-based approach is particularly suitable for this population. Participants can take part without travelling to a laboratory or clinic, which reduces barriers for people with fatigue or mobility problems. Researchers, in turn, can collect data from a wider geographic area and monitor cognitive performance over time through repeated sessions. For long-COVID studies, where symptom trajectories can extend over months or years, such longitudinal data are especially valuable.

Figure 2.4 illustrates the rapid growth of scientific publications on long-COVID and related neurocognitive research over the past two decades.

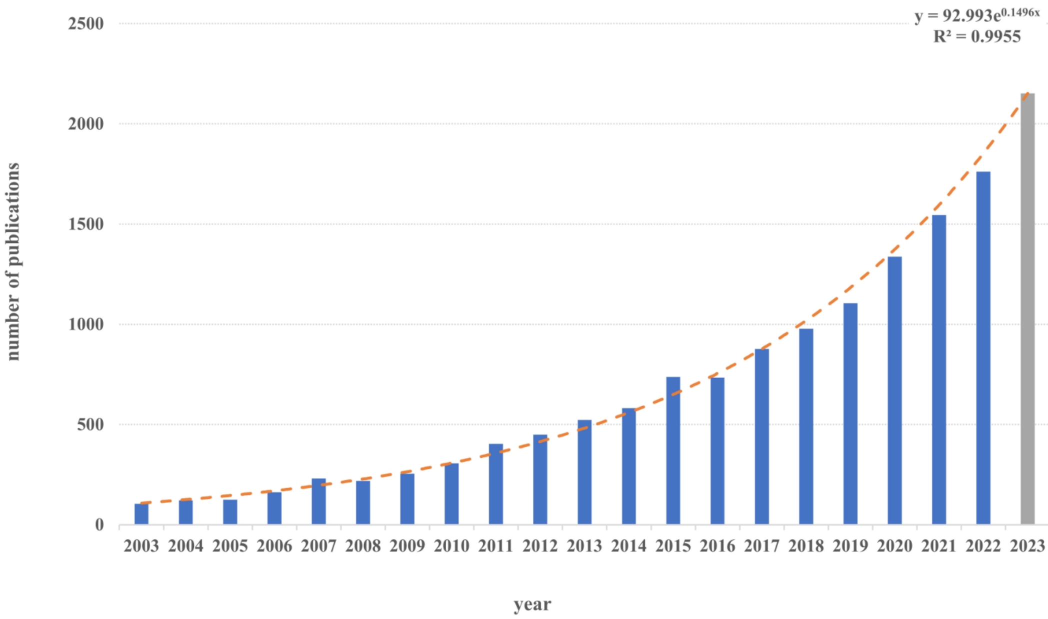

Figure 2.4: Growth of scientific publications on long-COVID and related neurocognitive research (2003–2023). Adapted from [5].

The initial versions of the platform already provided core functionality: user registration, task execution and basic data export. However, as the number of participants and experiments increased, limitations in the technical implementation became more apparent. Inconsistent data formats, missing timestamps, manual deployment procedures and heterogeneous task pipelines all increased the risk of errors and made data preparation more complex. These issues do not only affect convenience; they can have a direct impact on scientific validity if they introduce bias or noise into behavioural measurements.

2.5 Importance of Technical Reliability in Behavioural Research

Cognitive and behavioural experiments are highly sensitive to technical details. A difference of a few tens of milliseconds in a reaction-time task can be meaningful, but similar differences can also be produced by software delays or misconfigured systems.[4] Missing or inaccurately recorded timestamps, inconsistent encoding of responses or partial task logs can undermine the conclusions drawn from the data.

For this reason, the technical reliability of the platform is tightly linked to the quality of the research it supports. Robust deployment workflows help ensure that the same code and configuration are used across environments. Well-defined data models and unified formatting pipelines simplify analysis and reduce the risk of subtle bugs. Clear separation between frontend presentation logic and backend data handling helps to maintain consistency as new tasks and features are introduced.

The work described in the following chapters focuses on strengthening these aspects of the *Happy Again* platform. The goal is to make the system a more dependable tool for long-COVID research by improving the stability of deployments, the structure and completeness of recorded data and the internal processing of task results. This background chapter has outlined the scientific and methodological context in which these technical developments take place. The next chapter presents a detailed overview of the system architecture and components, which serves as a basis for describing the specific backend and frontend improvements.

Data Category	**Description**
Participant identifiers	User ID, attempt number
Timestamps	Cue onset, stimulus timing, response times
Trial metrics	Accuracy, reaction time, response type
Indicators	lc_flag, correctness flags, modality

Table 2.2: Example of structured data fields stored for behavioural experiments.

Chapter 3

System Overview

The Happy Again platform is an online system designed to support research on the long-term neurocognitive effects of COVID-19. It enables the remote execution of cognitive tasks, the collection of participant responses, and the management of behavioural datasets through a unified browser-based environment.

The platform consists of two primary subsystems:

- a **frontend** application (built with Angular), which delivers cognitive tasks, questionnaires, and the administrative interface;

- a **backend** service (FastAPI, Python), which manages authentication, experiment logic, data validation, and storage.

Both components operate together to ensure accurate timing, consistent data formats, and reliable long-term usage. A high-level overview of the system architecture is presented in Figure 3.1.

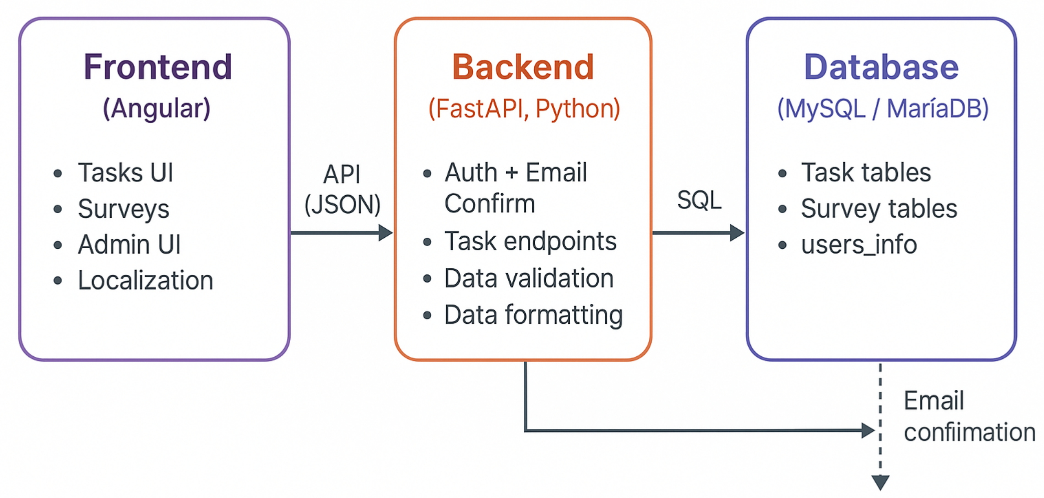

Figure 3.1: High-level architecture of the Happy Again platform.

3.1 Participant Workflow

Participants interact with the platform through a web browser. The typical workflow consists of:

1. **Registration and Login**: creation of an account and identification of the participant.
2. **Questionnaires**: demographic information, COVID-19 history, and symptom surveys.
3. **Cognitive Tasks**: browser-based experiments measuring reaction times, accuracy, attention, and perceptual processing.
4. **Data Submission**: automatic transmission of task results to the backend.

This workflow is illustrated in Figure 3.2.

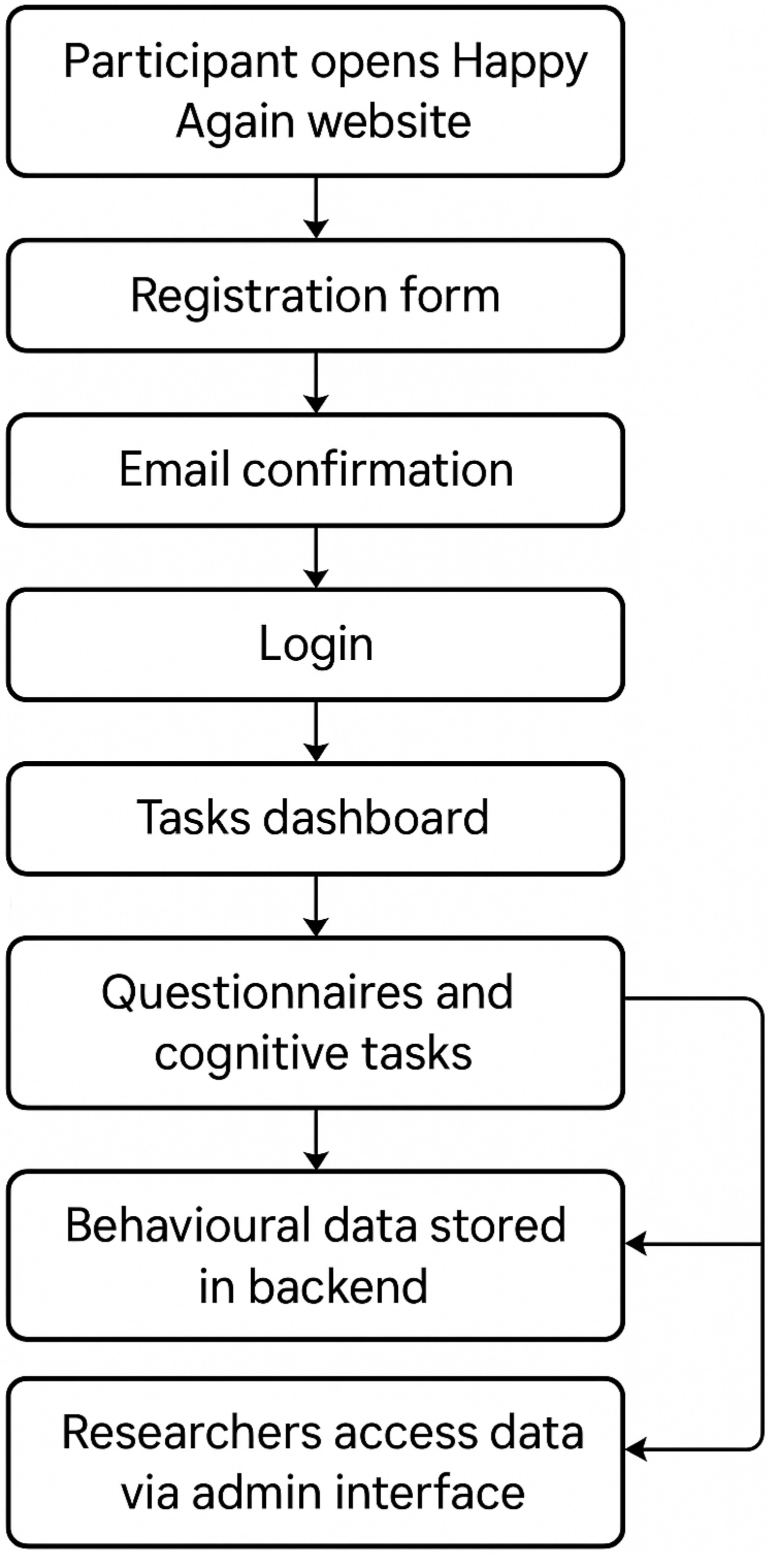

Figure 3.2: Participant workflow within the Happy Again platform.

3.2 Administrative Interface

The platform includes a dedicated administrative interface used by researchers to:

- view participant information,
- inspect completed tasks,
- filter sessions,
- export datasets in structured formats,
- verify timing and behavioural markers.

The interface provides access to encoded and decoded task outputs, helping researchers prepare data for statistical analyses.

3.3 High-Level Backend Overview

The backend coordinates:

- authentication and access control,
- validation of API requests,
- experiment configuration,
- storage of task results,
- communication with the frontend.

The backend does not perform analysis itself but ensures that all required behavioural markers (timestamps, accuracy, reaction times, correctness flags, etc.) are stored consistently.

Only a high-level understanding is presented in this chapter. A full description of backend modifications is provided in Chapter 4.

3.4 High-Level Frontend Overview

The frontend is an Angular application responsible for:

- rendering cognitive tasks and questionnaires,
- controlling stimulus presentation,

- ensuring timing consistency within browser limitations,
- transmitting structured results to the backend,
- providing multi-language support,
- offering administrative tools for researchers.

Accurate timing and predictable task flow are critical requirements, as behavioural data rely on precise measurements.

3.5 Database Overview

The relational database stores:

- participant identifiers,
- timestamps,
- trial-level results,
- performance metrics,
- task-specific indicators (including long-COVID flags),
- metadata such as attempt numbers and session information.

A consistent schema across tasks ensures compatibility during analysis and supports longitudinal tracking.

3.6 Deployment Overview

Deployment relies on a containerized environment to ensure predictable behaviour across development, testing, and production. While the technical details are discussed in Chapter 4, this chapter highlights only that:

- configuration is managed through environment variables,
- containerization improves reproducibility,
- deployments provide stable execution for behavioural experiments.

3.7 Main Components of the Platform

Table 3.1 summarises the key components of the Happy Again system.

Component	Description
Frontend (Angular)	User interface for tasks, questionnaires, and administrator tools.
Backend (FastAPI)	Handles authentication, experiment logic, and data storage.
Relational Database	Stores structured behavioural and questionnaire data.
Admin Dashboard	Provides researcher access to datasets and monitoring tools.
Deployment Environment	Containerized setup ensuring reproducible execution.

Table 3.1: Main components of the Happy Again platform.

This chapter establishes the conceptual and architectural context for the technical improvements described next.

Chapter 4

Backend Improvements

The backend of the Happy Again platform plays a central role in ensuring the stability, accuracy, and reliability of behavioural data collection. It manages user authentication, experiment logic, task configurations, database access, and communication with the frontend. Because cognitive experiments depend on precise timing, consistent data structures, and reproducible task execution, the backend must operate predictably under different environments and workloads.

During the work carried out for this thesis, a large part of the effort focused on improving the backend architecture, resolving inconsistencies across experiments, and strengthening the deployment workflow. These improvements were motivated by practical issues that emerged during active research use of the platform, including unstable deployments, inconsistent API behaviour, and limitations in the existing configuration system. This chapter describes the main backend modifications implemented throughout the project.

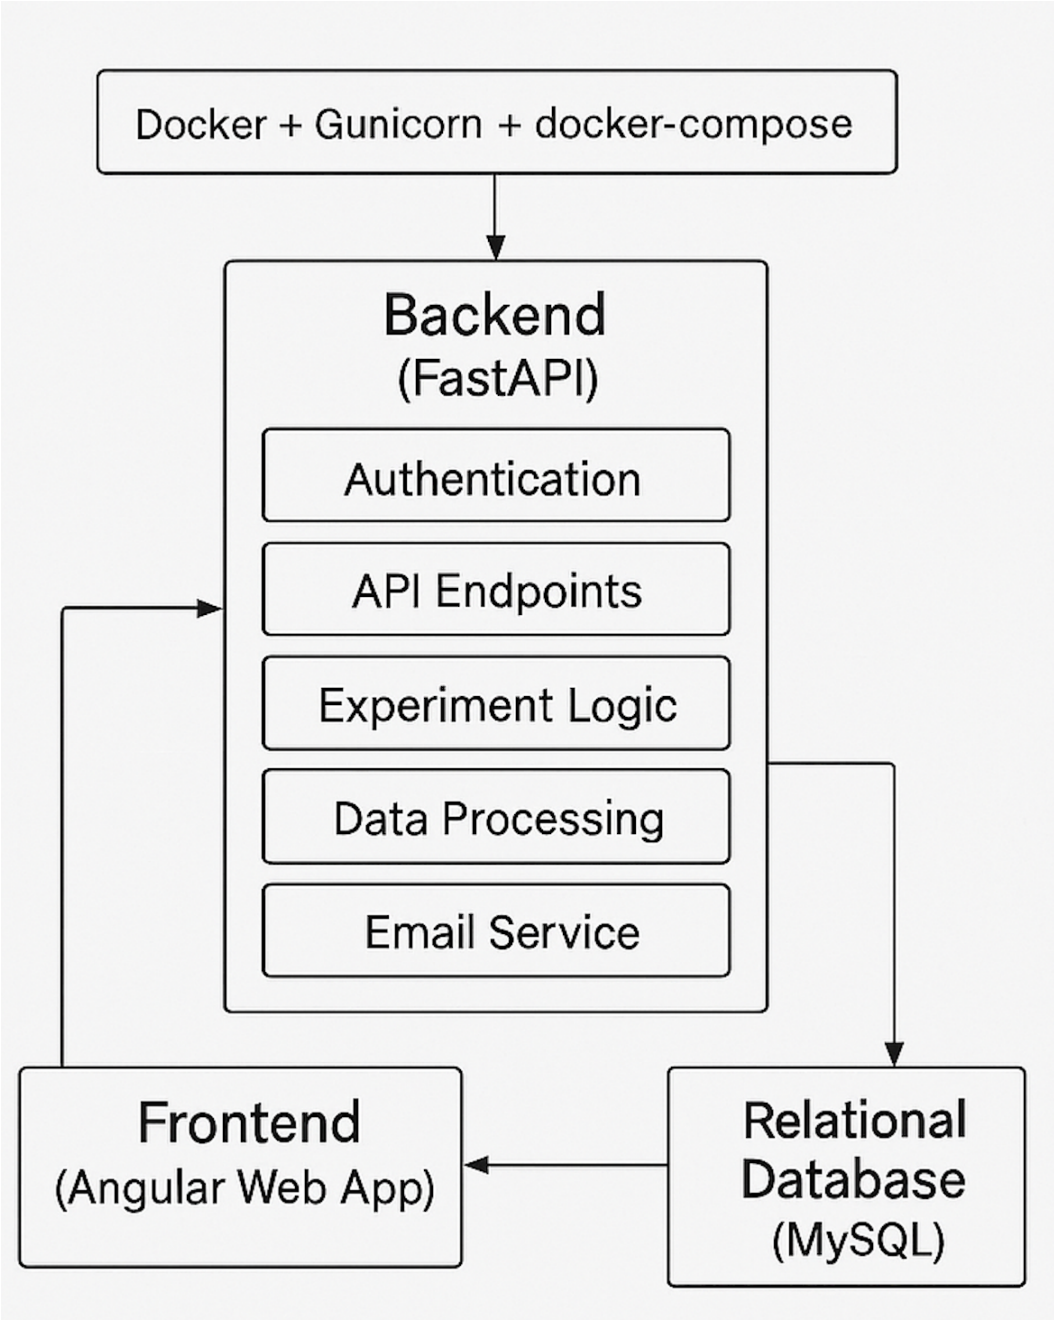

Figure 4.1: High-level backend architecture of the Happy Again platform.

4.1 Infrastructure and Deployment Improvements

Before development began, the backend lacked a fully reproducible deployment process. Configuration files differed between environments, manual setup introduced potential mistakes, and the absence of containerization made it difficult to ensure consistent behaviour across machines.

To address these issues, containerization using **Docker** was introduced. This change provides several benefits:

- **Reproducible environments**: every deployment now uses the same Python version, dependency set, and system configuration.
- **Isolation**: experiments are not affected by local machine differences.
- **Simplified setup**: developers and researchers can run the backend without configuring the system manually.
- **Scalability**: containerized deployments can be extended or moved to cloud infrastructure more easily.

Along with Docker, multiple deployment-related files were added, including:

- `Dockerfile`
- `docker-compose.yml`
- `docker-compose-dev.yml`
- environment templates (`.env.example`)
- a Gunicorn configuration (`gunicorn.conf.py`)
- an initialization script for database setup (`bin/init-db.sh`)

These components allow the entire platform to be launched consistently using predefined commands. They also reduce the time spent resolving environment-specific errors, which previously affected development and testing.

4.2 Configuration System Refactoring

The original backend had a single `config.py` file located at the project root. Over time, this file accumulated unrelated settings, and the platform began using multiple environment-specific patches to manage configuration. This resulted in confusion and occasional inconsistencies between development and production environments.

As part of the refactoring, the configuration was reorganized into a dedicated module:

- `happy_again/config.py`

This change provides a cleaner and more maintainable structure. Configuration values now follow a clear separation based on environment variables, reducing the chance of errors when switching between modes. The new structure also enables tools like Docker and Gunicorn to read settings reliably without relying on fragile path assumptions.

4.3 Updates to Data Models

Each cognitive task in the platform corresponds to a dedicated data model. Over time, the models had diverged: some stored timestamps, some did not; certain tasks used inconsistent naming conventions; several fields were missing for newer analyses; and some indicators could not be derived due to incomplete data.

To address this, multiple models were updated. These included:

- `movement_perception.py`
- `loudness_perception.py`
- `temporal_binding_window.py`
- `questionaire.py`
- `tasks.py`
- `users.py`
- `usersRoles.py`
- `posner_task.py`
- `memory_experiment`
- `word_encoding.py`

The changes implemented across these models include:

- addition of timing indicators, such as reaction-time fields and timestamps
- support for long-COVID classification flags (`lc_flag`)
- consistent naming conventions for fields across tasks
- alignment of data structures with frontend output

- fixes to missing or incorrectly stored values

Model	Main Improvement
movement_perception	Added timestamps, unified naming, lc_flag support
loudness_perception	Standardized structure, added metadata fields
temporal_binding_window	Unified timing markers, corrected missing values
posner_task	Adjusted field naming and consistency rules
word_encoding	Improved encoding scheme and trial merging logic
questionnaire	Unified response format across all survey types

Table 4.1: Summary of updates applied to major backend data models.

These updates ensure that researchers receive complete datasets that follow a uniform structure, making analysis easier and more reliable.

4.4 Enhancements to API Endpoints

Because behavioural experiments depend on precise request handling, several API endpoints were revised to improve validation, consistency, and performance. The main goals of these updates were:

- ensuring that each task receives and stores the correct data format
- preventing incomplete or corrupted submissions
- adding new fields required for analysis
- aligning APIs with updated models
- improving error-handling logic

In several tasks—including target detection, loudness perception, and word categorization—the APIs were updated to incorporate:

- unified data-processing logic
- merged trial structures
- consistent response encoding
- enriched metadata (timestamps, inter-trial intervals, derived variables)

An example of these improvements is the **merge_target_detection_data** functionality, which standardizes how target-detection responses are processed before being stored. This reduces fragmentation and eliminates multiple variations of similar logic that previously existed in the codebase.

4.5 Handling of lc_flag Across Tasks

A recurring requirement from researchers was the ability to mark participants based on criteria related to long-COVID symptoms or questionnaire responses. This classification is used in analysis to compare behavioural performance across groups.

To support this, a unified **lc_flag** mechanism was added across several tasks:

- movement perception
- loudness perception
- word categorization
- flash-beep tasks
- target detection
- COVID questionnaire
- demographic questionnaire

The integration of this flag allows researchers to:

- group participants consistently
- apply filtering during data export
- perform stratified analyses
- trace how behavioural markers differ between flagged and non-flagged participants

By ensuring the flag is present in all relevant models, data extraction becomes more systematic.

4.6 Stability and Error-Handling Improvements

During development, several issues were identified in the backend:

- registration failures due to incomplete flags
- email notifications failing under specific conditions
- timeout errors during task submissions
- inconsistencies in processed data structures
- differences between development and production behaviour

These issues were addressed through:

- fixes to registration logic
- corrections to email-sending and Gmail integration
- improved timeout handling
- normalization of task-processing pipelines
- removal of outdated or unused code paths

These improvements significantly reduce the risk of data loss, which is critical when working with behavioural experiments where each trial can carry scientific value.

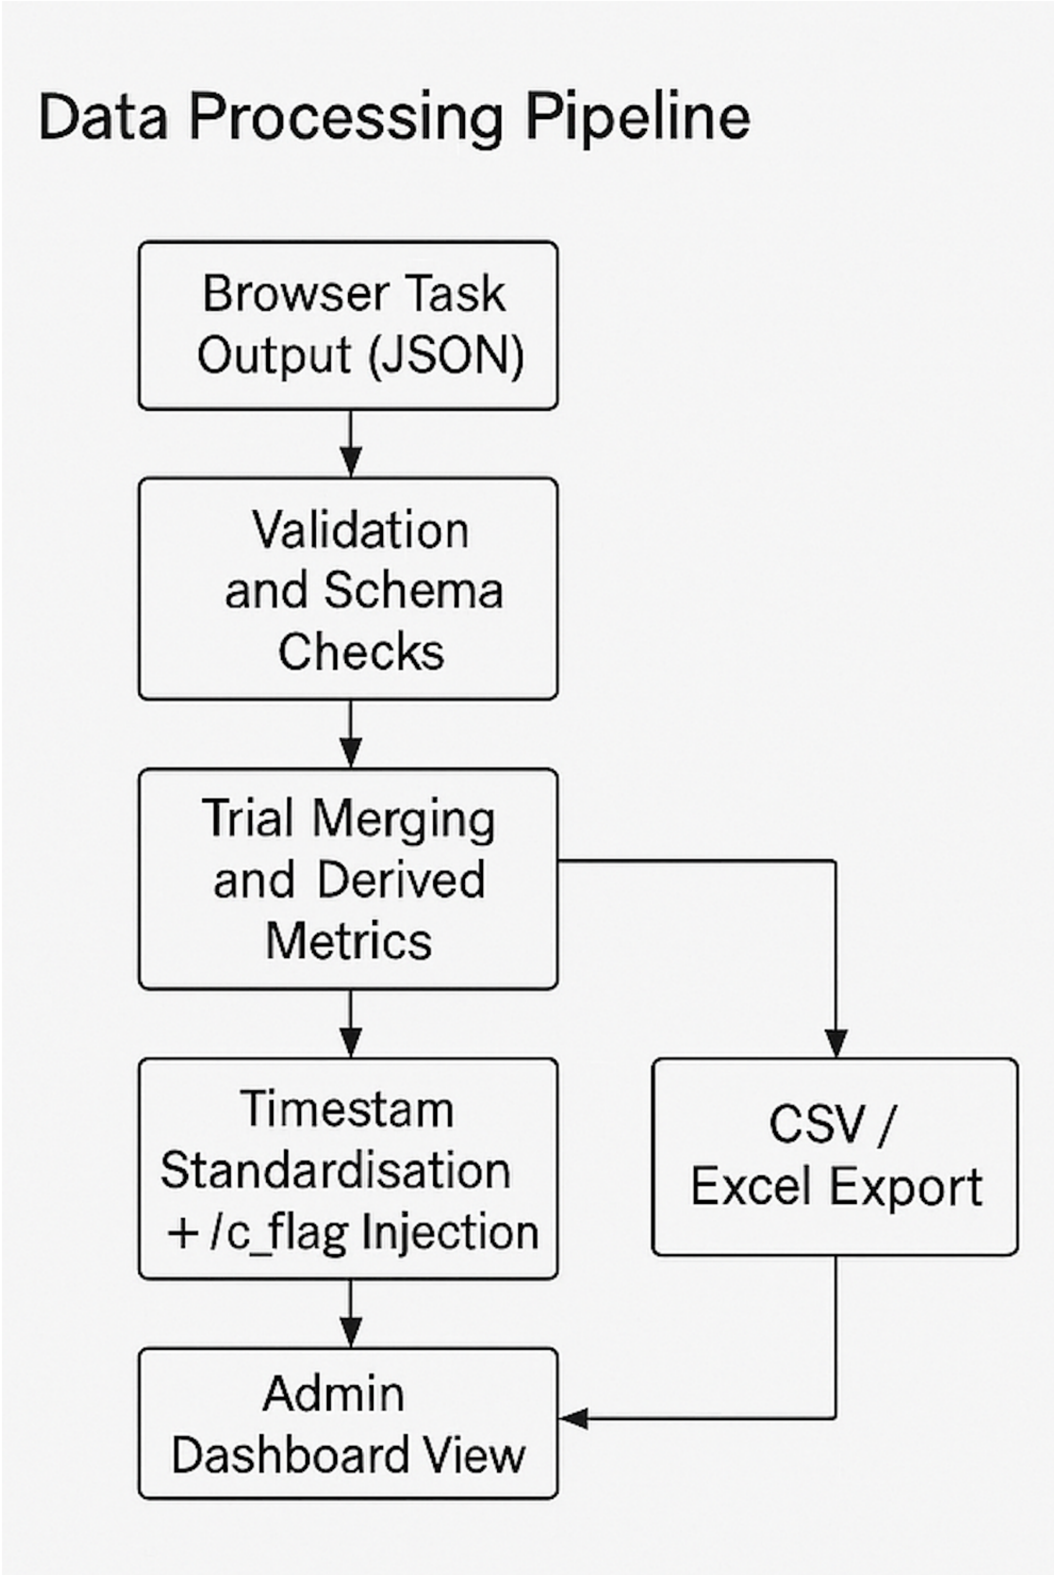

Figure 4.2: Data-processing pipeline used to standardize task results across experiments.

4.7 Summary of Backend Contributions

The backend development carried out during this thesis strengthened the foundation of the Happy Again platform in several ways:

- reproducible deployment through Docker
- unified configuration system
- updated models and APIs for improved data consistency
- integration of long-COVID classification flags
- enhanced error handling and stability
- standardized processing of behavioural tasks
- cleaner and more maintainable codebase

These developments allow researchers to analyse data with greater confidence and ensure that future features can be added without risking architectural instability.

Chapter 5

Frontend Improvements

The frontend of the *Happy Again* platform is responsible for delivering cognitive tasks, questionnaires, and administrative tools directly to participants and researchers through a browser-based interface. Because behavioural experiments rely on precise timing, structured user flows, and predictable execution across devices, the frontend plays a central role in the reliability of the collected data. This chapter describes the architecture of the frontend system and details the improvements implemented during the project in two main areas: (1) participant-facing experimental tasks and (2) the administrative dashboard used by researchers.

5.1 Frontend Architecture Overview

The platform is implemented as a monolithic Angular 11 application organised into a single `AppModule`. The application includes:

- dedicated views for each cognitive task;
- a dynamic questionnaire engine used for multiple survey types;
- authentication and session management components;
- an administrative dashboard for monitoring and dataset exporting;
- a service layer providing API communication, fullscreen handling, device detection, and shared utilities.

Routing is handled through Angular's `RouterModule`, with paths corresponding to each experiment (Flash–Beep, Memory, Target Detection, Movement Perception, Loudness Perception, SPQ, and general questionnaires). Researcher-only routes are protected using authentication guards.

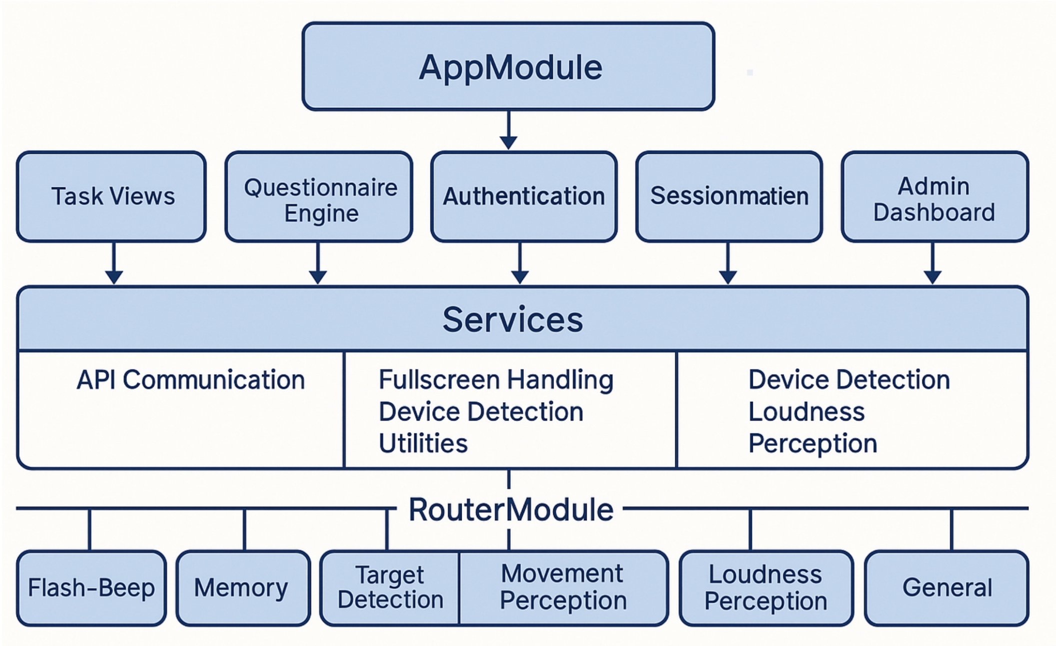

Figure 5.1: High-level frontend architecture of the Happy Again platform.

5.2 Improvements to Task Components

Participant-facing tasks are sensitive elements of the system because scientific validity depends on their timing accuracy and consistency. Prior to this work, several issues existed: duplicated timing logic, inconsistent timestamp formats, missing error handling, and user-interface freezes during long tasks.

The following improvements were introduced:

- **Unified timing logic**: aligned timestamp acquisition using high-resolution timers; structured event processing to minimise delay.
- **Improved feedback and stability**: added loading states, improved fullscreen handling, and safer key event management.
- **Refactoring repeated logic**: consolidated repeated round-handling and delay logic into clearer internal functions.

These improvements ensure that all tasks produce reliable timing indicators and minimise the likelihood of experiment interruptions.

5.3 Major Refactoring of the Admin Area

The administrative dashboard is the most complex part of the frontend, used by researchers to manage users, export datasets, and view analysis outputs. Before the refactoring, several issues impacted performance and usability:

- large datasets caused interface freezes;
- repeated array scans created $O(N^2)$ behaviour;
- long operations provided no visual feedback, causing user confusion;
- dataset exports sometimes produced inconsistent structures.

5.3.1 Algorithmic Optimisation: from $O(N^2)$ to $O(N)$

A major improvement was the introduction of lookup caches using JavaScript `Map` and `Set`. Instead of repeatedly scanning arrays for matching entries, the system builds lookup tables during initialisation. This optimisation transformed slow operations into instant computations.

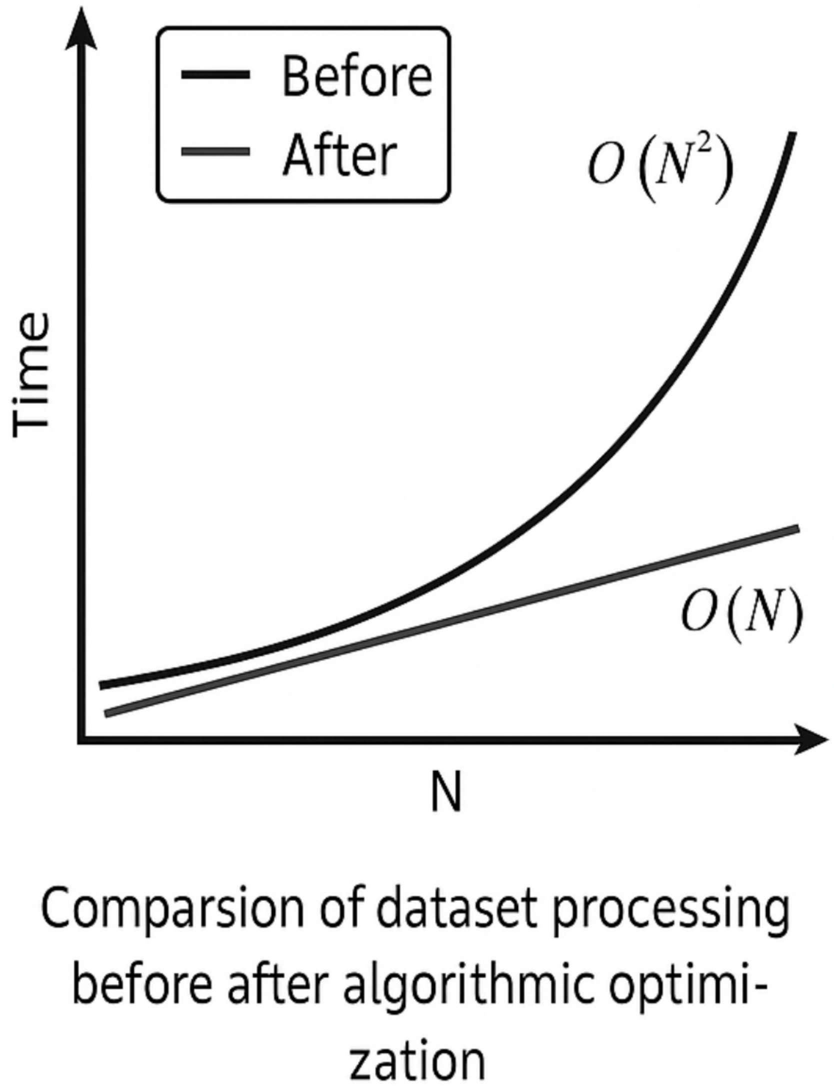

Figure 5.2: Comparison of dataset processing before and after algorithmic optimisation.

5.3.2 Data Handling and Export Improvements

The export logic was restructured to ensure:

- stable column ordering,

- consistent decoding of task responses,
- correct merging of factor-analysis fields,
- reliable handling of missing or malformed values,
- replacement of large conditional blocks with structured processing functions.

5.3.3 User Experience Improvements: Loader Integration

A dedicated visual loader was added to indicate long operations such as dataset preparation, translation-file processing, and factor-analysis merging. Previously, the dashboard appeared frozen during these operations; the new loader provides immediate user feedback and prevents accidental repeated clicks.

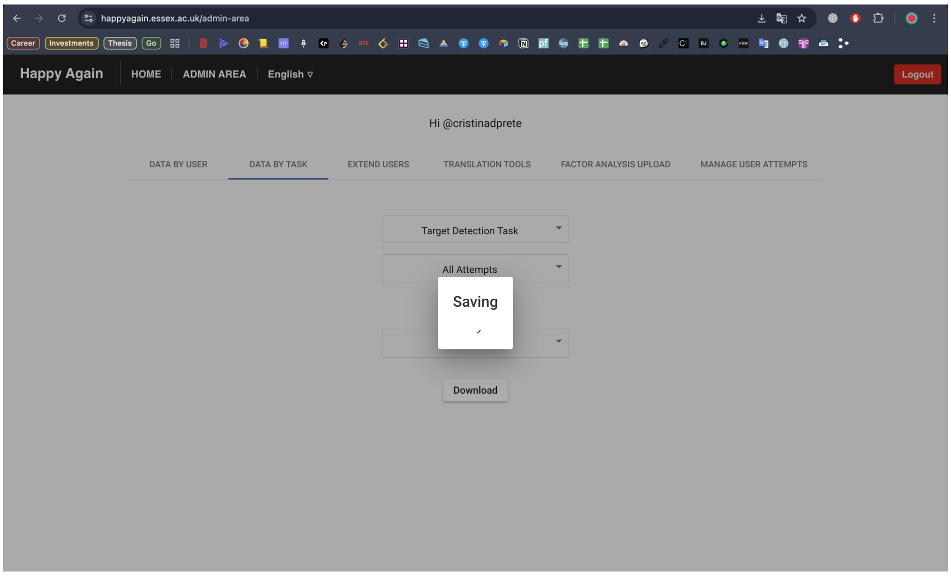

Figure 5.3: Loader indicating dataset export progress in the administrative dashboard.

5.4 Data Submission Workflow Improvements

The workflow for sending experiment results to the backend was standardised to avoid inconsistent payloads or missing timestamps. The updated submission pipeline includes:

1. assembling data into structured in-memory objects;
2. constructing explicit JSON bodies for API requests;

3. attaching session tokens consistently;
4. marking task completion using a dedicated service;
5. providing visual feedback during submission.

This unified approach increases stability during long experimental sessions and ensures that the backend receives complete datasets.

5.5 Summary of Frontend Contributions

The frontend improvements introduced during this thesis significantly enhance stability, performance, and usability. The main contributions include:

- unified timing logic across all experimental tasks,
- improved feedback and error handling,
- major refactoring and performance optimisation of the admin area,
- loader integration for long operations,
- more reliable dataset export workflows,
- updated deployment powered by Docker and CI/CD.

These improvements strengthen the scientific reliability of the platform and ensure that both participants and researchers interact with a consistent and efficient system.

Chapter 6

Data Processing Improvements

High–quality behavioural research depends not only on the correct implementation of cognitive tasks but also on the structure, coherence, and interpretability of the data these tasks produce. Web–based experimental platforms are particularly sensitive to timing variability, inconsistent encoding schemes, and heterogeneous data models. During the early phases of the *Happy Again* project, the data-processing pipeline exhibited several of these limitations: tasks produced non-aligned schemas, timestamp precision differed between modules, derived behavioural metrics were missing, and administrative exports required manual post-processing.

This chapter describes the redesign of the entire data-processing workflow. The objective was to establish a predictable, standardised, and reproducible pipeline capable of supporting large-scale behavioural research, particularly for long-COVID studies where timing accuracy and cross-task consistency are essential.

6.1 Unification of Processing Logic Across Tasks

Before refactoring, each cognitive task implemented its own output model, using different naming conventions, timestamp formats, and local helper functions for encoding. Some tasks included metadata, others returned only raw trial arrays, and several relied on task-specific logic for calculating derived features. This fragmentation complicated analysis and reduced reproducibility.

To address these issues, a unified processing scheme was introduced with five goals:

- enforce a common structural pattern across all tasks;
- provide predictable encoding and decoding rules;

- remove duplicated logic and implicit behavioural assumptions;

- simplify statistical preprocessing;

- reduce cross-participant variability caused by structural inconsistencies.

Each trial now conforms to a shared schema containing timing markers, event metadata, encoded responses, and derived features. This enables automated quality-control and task-agnostic analysis tools.

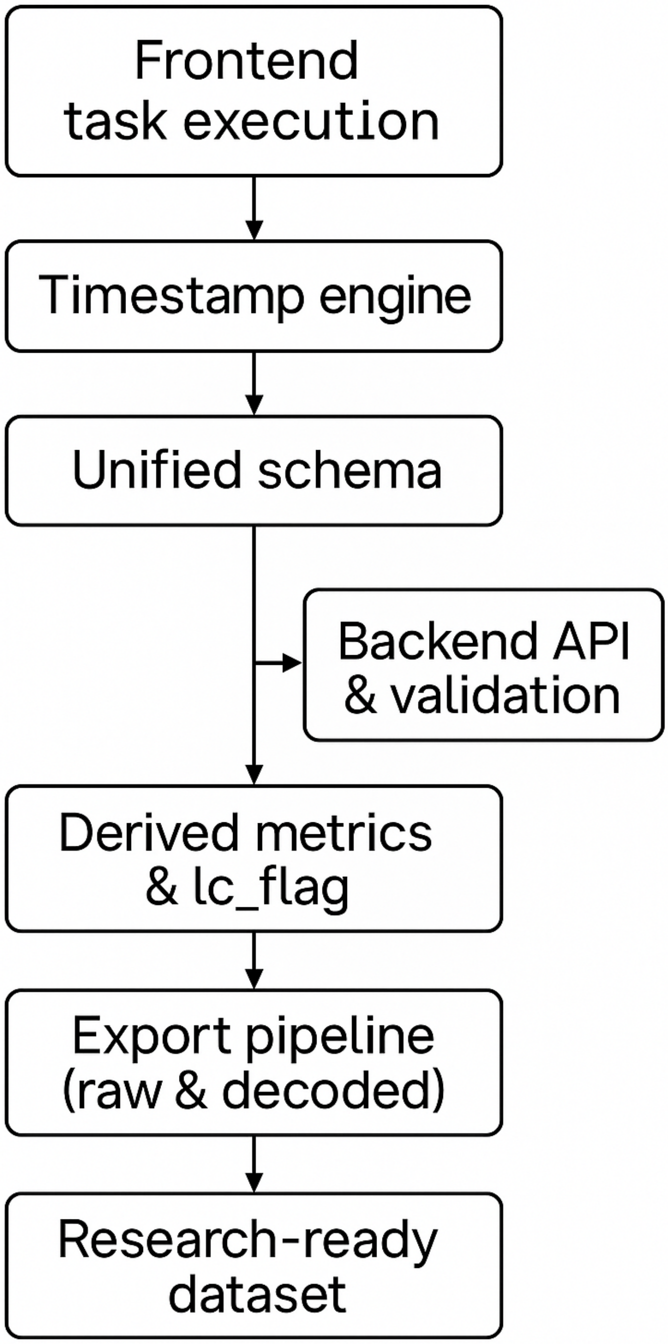

Figure 6.1: Unified data-processing pipeline across tasks, backend services, and export layers.

6.2 Standardisation of Timestamps and Timing Indicators

Timing accuracy is a cornerstone of behavioural research, particularly in reaction-time tasks. Initial analysis revealed several sources of inconsistency:

- mixed usage of `Date.now()` and `performance.now()`;
- missing point-in-time markers such as stimulus onset;
- absence of inter-trial intervals in some tasks;
- inconsistent timestamp precision across modules.

The redesigned timing subsystem introduces:

- high-resolution, monotonic timestamps for all events;
- explicit trial-level markers (start, stimulus onset, response, end);
- unified reaction-time calculation rules;
- inter-trtrial interval tracking across all tasks.

These improvements enhance the precision of latency modelling and enable longitudinal timing analyses.

6.3 Consistent Encoding and Decoding of Task Responses

Prior to the redesign, tasks used heterogeneous response encodings — integers, booleans, and unlabelled strings — without documentation. This increased ambiguity during analysis and required researchers to maintain manual lookup tables.

The improved pipeline introduces:

- a unified encoding scheme for stimuli, responses, and trial outcomes;
- decoding logic integrated directly into the administrative export pipeline;
- shared helper modules used by all tasks and backend services.

Researchers now receive datasets with both encoded and human-readable values, reducing interpretation errors.

6.4 Derived Features and Classification Flags

Many cognitive tasks provide richer insights when raw values are supplemented with derived indicators. Prior to this work, derived metrics were implemented inconsistently or not computed at all.

The new pipeline integrates:

- reaction-time offsets and timing-difference features;
- stimulus–response matching indicators;
- aggregated behavioural markers;
- a unified long-COVID classification flag (`lc_flag`).

The `lc_flag` enables automated participant stratification, facilitating clinical comparisons without additional preprocessing.

Derived Feature	**Purpose**
RT offset	Identifies reaction-time drift across trials
Inter-trial delta	Captures recovery and fatigue effects
Stimulus–response alignment	Detects perceptual integration errors
Trial-type performance index	Aggregates task-specific correctness
Long-COVID flag	Enables clinical grouping

Table 6.1: Examples of derived features added during refactoring.

6.5 Administrative Export Improvements

Exporting datasets is one of the most operationally critical steps for researchers. The previous export system suffered from duplicated code, inconsistent column ordering, and occasional download failures.

The redesigned export pipeline now provides:

- stable, versioned column ordering;
- automatically included derived fields and timestamp structures;
- both raw (encoded) and decoded dataset variants;
- consistent formatting across all tasks;
- elimination of asynchronous race conditions.

6.6 Data Validation and Integrity Measures

Because data is collected on heterogeneous devices, incomplete submissions are inevitable. The updated system reduces this risk through:

- frontend validation before submission;
- backend validation of required fields and types;
- fallback mechanisms for missing timestamps;
- improved error feedback preventing silent failures.

These measures significantly reduce the likelihood of corrupted datasets.

6.7 Summary of Data Processing Contributions

The redesign of the data-processing pipeline resulted in measurable improvements:

- unified schema across tasks;
- standardised timestamp model;
- consistent encoding/decoding;
- integrated derived features;
- stable administrative exports;
- strong validation and error handling.

The platform now provides clean, interpretable data suitable for large-scale behavioural analysis and long-COVID research.

Area	Main Improvement
Task processing	Unified schema and shared helper functions across tasks
Timing	Standardised timestamps, reaction-time rules, and inter-trial intervals
Encoding	Shared encoding scheme with automatic decoding during export
Derived features	Integrated behavioural markers and long-COVID classification flags
Admin export	Stable column ordering, enriched fields, and robust download workflow
Data integrity	Validation on both frontend and backend, safer handling of edge cases

Table 6.2: Summary of data-processing improvements across the Happy Again platform.

Chapter 7

Results and Impact

The improvements introduced during this thesis produced substantial advancements across all subsystems of the *Happy Again* platform. Although the platform itself does not generate scientific findings, its technical design directly determines the quality, completeness, and interpretability of behavioural data collected in long-COVID research. This chapter summarises the measurable outcomes of the work along four main axes: system stability, data quality, workflow efficiency, and scientific reproducibility.

7.1 System Stability and Operational Consistency

Before the redesign, the platform exhibited environment-dependent behaviour. Development, staging, and production environments often behaved differently due to configuration drift, inconsistent dependencies, and manual deployments. Several backend endpoints were also prone to intermittent timeouts, contributing to incomplete submissions.

After introducing Docker-based deployment, unified environment configuration, and improved error-handling mechanisms, the system now demonstrates:

- consistent behaviour across development, testing, and production;
- fewer server-side failures during peak activity;
- stable execution of timing-sensitive tasks;
- reduced probability of incomplete or corrupted submissions.

7.2 Improvements in Data Accuracy and Completeness

Data quality is central to behavioural research. Prior to the improvements, task outputs showed inconsistencies in timestamp formats, missing fields, heterogeneous naming rules, and absence of derived behavioural metrics.

The redesigned data-processing pipeline introduced:

- fully standardised timestamp structures across all tasks;
- consistent naming conventions and field ordering;
- uniform encoding/decoding rules;
- reduced incidence of missing or null fields;
- integrated derived values and behavioural markers.

These improvements substantially reduce preprocessing overhead and improve the validity of reaction-time and longitudinal metrics.

Metric	**Before Improvements**	**After Improvements**
Timestamp consistency	Heterogeneous formats	Fully standardised across tasks
Admin export speed	$O(N^2)$ lookups	$O(N)$ with caching
Deployment stability	Environment-dependent	Fully containerised and reproducible
Incomplete data ratio	High (frequent missing fields)	Strongly reduced due to validation
Task timing precision	Mixed (device-dependent)	Unified high-resolution timing model

Table 7.1: Comparison of system performance metrics before and after refactoring.

7.3 Improvements to Administrative Workflow and Data Export

The administrative dashboard is one of the most frequently used tools by researchers. Before the redesign, export processes suffered from duplicated internal logic, inconsistent column ordering, and occasional failures when handling large datasets.

After restructuring the export pipeline:

- dataset downloads are reliable even for large participant groups;
- column ordering remains stable across all task types;
- newly standardised fields are included automatically;
- decoded datasets provide immediate human-readable interpretation;
- race conditions and asynchronous export failures have been eliminated.

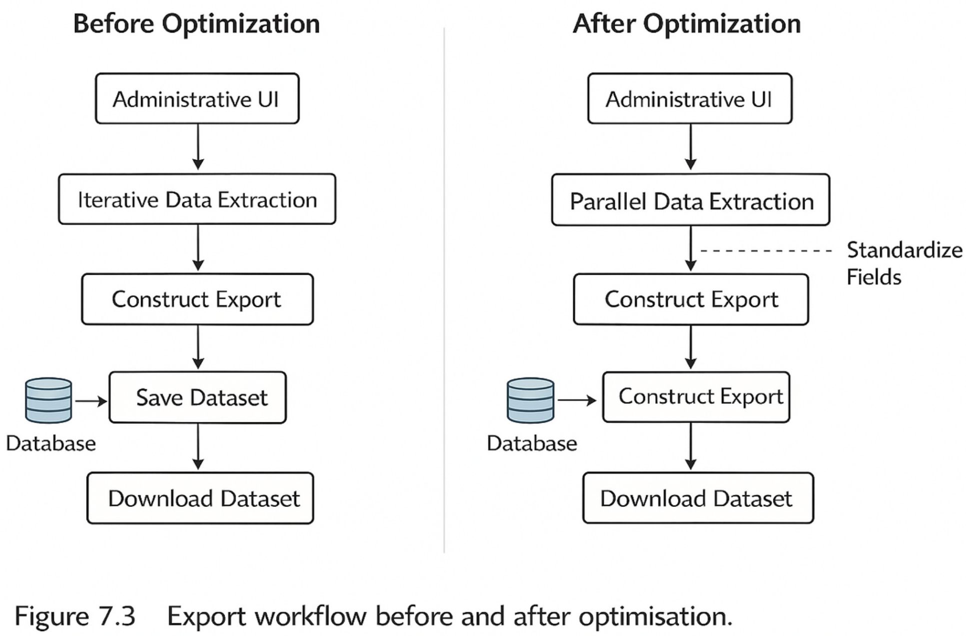

Figure 7.1: Administrative export workflow before and after optimisation.

7.4 Enhanced Reproducibility and Scientific Reliability

Reproducibility is crucial for behavioural research, particularly in studies examining long-term cognitive effects. Prior to the improvements, inconsistencies in internal structures made cross-time and cross-task comparisons difficult.

With the unified processing pipeline and containerised deployment, the platform now ensures:

- stable versioned dataset formats;
- reproducible behavioural markers across time;
- consistent propagation of classification flags (e.g., `lc_flag`);
- reduced methodological noise in longitudinal analyses.

7.5 Maintainability and Extensibility

The platform's architectural redesign improved long-term maintainability:

- redundant logic was removed from both frontend and backend;
- architectural patterns were harmonised across tasks;
- deployment workflow was documented and automated;
- code readability and modularity significantly increased.

These improvements make the platform easier to extend with new tasks, data sources, or analytical modules.

Improvement	**Scientific Impact**
Standardised timestamps	Enables precise reaction-time analysis
Unified task schema	Facilitates cross-task comparisons
Dockerised deployment	Ensures reproducibility across studies
Frontend performance gains	Reduces measurement noise
Export pipeline redesign	Accelerates research workflow
Integrated `lc_flag`	Supports clinical stratification

Table 7.2: Relationship between technical improvements and scientific research benefits.

7.6 Summary of Results and Impact

The combined effect of all improvements is a platform that is:

- significantly more stable during long sessions;

- more predictable across heterogeneous devices;
- more accurate in recording behavioural signals;
- easier for researchers to analyse and interpret;
- easier for developers to maintain and extend;
- better suited for large-scale studies on long-COVID neurological consequences.

Instead of requiring ad-hoc preprocessing or manual corrections, the system now produces structured, interpretable datasets that support robust behavioural science.

Chapter 8

Discussion and Limitations

The improvements carried out in this thesis substantially enhanced the technical reliability, internal consistency, and long-term maintainability of the *Happy Again* platform. These changes directly support the scientific goals of the project by enabling the collection of behavioural data related to long-COVID symptoms in a stable, consistent, and reproducible manner. Nonetheless, several limitations remain—some inherent to browser-based experimental environments, others related to the current architecture, available resources, or the scope of implemented tasks.

8.1 Interpretation of Results in the Context of Behavioural Research

Although the improvements introduced in this thesis significantly reduce technical noise, they cannot eliminate all external sources of variability. Web-based experiments are conducted in uncontrolled environments where participants use heterogeneous devices, network connections, operating systems, display refresh rates, and input peripherals. These differences can influence reaction-time-based tasks, even with optimised timing logic.

While the updated processing pipeline improves internal consistency, absolute timing precision still depends on the participant's hardware and system load. For this reason, behavioural results obtained through the platform—especially latency-sensitive indicators—should be interpreted with awareness of such environmental constraints.

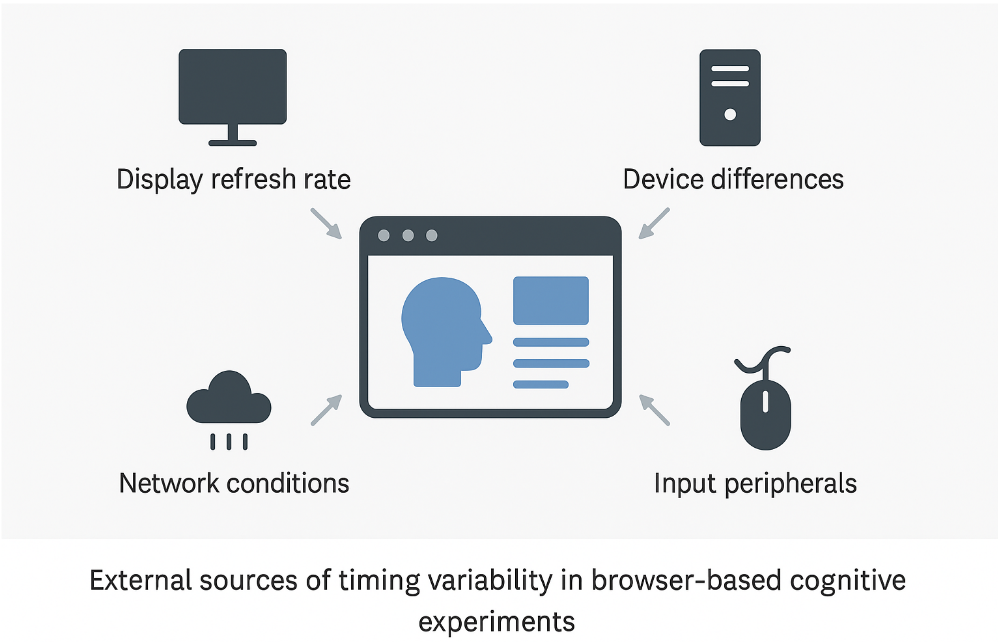

Figure 8.1: External sources of timing variability in browser-based cognitive experiments.

8.2 Technical Limitations

Despite the substantial architectural improvements, several technical limitations persist:

- **Browser timing variability.** Even with unified timestamp logic, tasks remain sensitive to CPU load, rendering delays, and background processes. WebGL-level timing or calibration procedures could further mitigate this but were outside the scope of the thesis.
- **Legacy components.** Several older modules still follow patterns that predate the unified architecture. Although functional, they may require future refactoring to fully align with the updated system design.
- **Partial automation of deployment.** While Docker-based containerisation ensures environment consistency, full CI/CD automation, automated database migrations, and integrated monitoring systems are not yet implemented.

These limitations are summarised in Table 8.1.

Type of Limitation	Description
Environmental variability	Device differences, CPU load, network stability, display refresh-rate variation
Legacy architecture elements	Older modules not fully aligned with the new unified structure
Incomplete automation	Missing CI/CD, automatic migrations, and monitoring tools

Table 8.1: Overview of platform limitations.

8.3 Scientific and Methodological Limitations

From a research-methodology perspective:

- the current set of tasks covers attention, timing, and perceptual processes but does not include domains such as working memory or executive control;
- incomplete datasets may still occur due to browser closures, unstable internet connections, or device-related interruptions;
- browser-based tasks cannot fully replicate laboratory-grade precision, especially for millisecond-level comparisons across individuals.

While the improved validation mechanisms strongly reduce the risk of malformed data, they cannot fully compensate for real-world participant interruptions.

8.4 Impact on Future Research

The improvements made during this thesis create a more predictable and coherent data-collection environment. Researchers will now spend significantly less time on manual preprocessing, data correction, or format conversion.

Table 8.2 summarises the remaining risks and potential mitigation strategies.

Risk	**Potential Mitigation Strategy**
Device-related timing noise	Calibration routines, WebGL timing, refresh-rate detection
Incomplete datasets	Autosave mechanisms, redundancy in submission routines
Legacy code sections	Full modularisation and refactoring passes
Limited task diversity	Integration of new memory, executive-function, and multimodal tasks

Table 8.2: Remaining risks and mitigation strategies.

8.5 Summary

In summary, the platform is now far more stable, coherent, maintainable, and scientifically reliable than before. The remaining limitations are manageable and consistent with the typical constraints of modern web-based cognitive testing. The updated architecture offers a strong foundation for future research and supports the continued expansion of the *Happy Again* platform.

Chapter 9

Conclusion and Future Work

This thesis focused on strengthening the technical foundations of the *Happy Again* platform, a system designed to support large-scale behavioural research on the cognitive consequences of long-COVID. Through extensive refactoring across the backend, frontend, data-processing layer, and deployment workflow, the platform now provides a stable, reproducible, and extensible foundation for future research.

9.1 Conclusion

The key contributions of this work include:

- a unified and reproducible backend deployment using Docker;
- refactored backend models and APIs ensuring consistent data structures;
- major improvements to timing accuracy and event handling in cognitive tasks;
- complete standardisation of timestamps, encoding schemes, and derived fields;
- a fully redesigned administrative export pipeline supporting both raw and decoded datasets;
- performance optimisation in the frontend admin area (from $O(N^2)$ to $O(N)$);
- substantial improvements in data integrity, error handling, and cross-task coherence.

Table 9.1 summarises these contributions by system area.

Together, these improvements transform the platform from a prototype into a robust research tool. The system now produces datasets that are easier to analyse, more consistent over time, and more resilient to environmental variability. These

Area	**Main Contribution**
Backend infrastructure	Containerised deployment with Docker, unified configuration management, and reproducible environments across development, testing, and production.
Backend models and APIs	Refactored data models and endpoints, consistent naming conventions, and strengthened validation of incoming requests.
Frontend task engine	Improved timing accuracy, unified event-handling logic, and reduced UI-related noise during cognitive tasks.
Data-processing pipeline	Standardised timestamps, encoding schemes, and derived fields across tasks, enabling cross-task comparisons and reliable analysis.
Admin tools and exports	Redesigned export workflow with stable column ordering, decoded datasets, and robust handling of large data volumes.
System reliability and integrity	Enhanced error handling, data-integrity checks, and cross-task coherence, reducing technical noise in behavioural datasets.

Table 9.1: Summary of the main technical contributions of this thesis across different areas of the Happy Again platform.

outcomes directly support long-COVID research by reducing technical noise and providing cleaner behavioural indicators.

9.2 Future Work

Although the current platform is significantly improved, several important avenues remain for future development. The most relevant directions are summarised in Table 9.2 and discussed in more detail below.

9.2.1 Expansion of Cognitive Assessments

New tasks targeting memory, executive function, and fatigue-related performance could enrich the behavioural dataset and support broader neuroscientific analyses. The unified processing pipeline introduced in this thesis simplifies the integration of such tasks and ensures that they would follow the same standards for timing, encoding, and export.

9.2.2 Improved Timing Precision

Future work could explore WebGL-based rendering, high-precision timing APIs, or device-specific calibration routines to further reduce timing variability inherent to browser-based testing. These extensions would help bridge the gap between web-based and laboratory-grade temporal precision, particularly in reaction-time-based paradigms.

9.2.3 Infrastructure Automation

Full CI/CD pipelines, automated database migrations, and integrated monitoring would increase deployment consistency and further reduce manual maintenance. As the platform scales to more users and research groups, such automation will become increasingly important for long-term sustainability.

9.2.4 Integration of External Data Sources

Wearable sensors, mobile physiological data, or cognitive fatigue indicators could complement behavioural measures and provide multimodal insights relevant to long-COVID research. The refactored architecture is now better prepared to incorporate such external data streams with additional development.

9.2.5 Enhanced Researcher Tools

Additional analytics features—such as built-in visualisation, quality-control dashboards, or real-time monitoring of participant progress—could significantly streamline the research workflow. These tools would help researchers identify anomalies early and evaluate data quality before formal statistical analysis.

Area	**Main Goal**	**Example Extensions**
Cognitive task set	Broaden the range of behavioural markers relevant to long-COVID	Memory tasks, executive-function tasks, fatigue-related paradigms
Timing precision	Reduce variability in reaction-time and stimulus-onset measurements	WebGL-based rendering, high-precision timing APIs, device calibration procedures
Infrastructure and DevOps	Increase deployment robustness and reduce manual operations	CI/CD pipelines, automated database migrations, monitoring and alerting
Multimodal data integration	Combine behavioural and physiological signals for richer analysis	Wearable devices, mobile sensor data, physiological fatigue indices
Researcher-facing tools	Improve efficiency and transparency of the research workflow	In-platform visualisation, data-quality indicators, progress dashboards

Table 9.2: Overview of proposed directions for future work on the Happy Again platform.

9.3 Final Remarks

The improvements achieved in this thesis establish a solid and extensible foundation for the *Happy Again* platform. By addressing instability, inconsistent data structures, timing irregularities, and performance issues, the platform is now ready for more advanced behavioural research and large-scale deployments. The system remains open for future development, and the modular improvements presented here are designed to support long-term evolution as scientific requirements grow.

Bibliography

[1] Adam Hampshire and et al. «Multidimensional cognitive impairment following COVID-19». In: *Nature* (2021) (cit. on p. 4).

[2] «Neuropsychological manifestations of long COVID». In: *Frontiers in Neurology* (2023) (cit. on pp. 4, 6).

[3] et al. Saucier. «Long-term cognitive dysfunction after COVID-19 infection». In: *Frontiers in Neurology* (2023) (cit. on p. 5).

[4] «Validity of remote cognitive testing». In: *The Lancet Neurology* (2022) (cit. on pp. 7, 8, 10).

[5] Global Health Data Exchange. *Bibliometric Analysis of Long-COVID Research Trends*. Data retrieved from PubMed and Scopus databases. 2023 (cit. on p. 10).