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.

Summary

Happy Again is a digital platform designed to support large-scale behavioural experiments for the study of long-term neurological and cognitive consequences of COVID-19. The purpose of this thesis was to strengthen the technical, architectural, and data-processing foundations of the platform, ensuring that it can be reliably used in real-world research scenarios, across heterogeneous devices, and at population-level scale. Over the course of the project, major improvements were introduced across the backend infrastructure, frontend application, data-processing pipeline, deployment workflow, and administrative tools. This summary presents the motivation, methodology, technical contributions, and scientific impact of the work.

The original version of the *Happy Again* platform suffered from several limitations common to web-based experimental environments. These included inconsistent data formats across tasks, incomplete timing information, duplicate logic in both backend and frontend, performance bottlenecks, and occasional deployment instability. Because behavioural research relies heavily on the precision and reproducibility of collected data, especially in reaction-time-based tasks, these technical issues posed a direct threat to the scientific validity of the project. The aim of this thesis was therefore twofold: first, to redesign the platform to ensure stable and predictable execution, and second, to standardise data structures such that researchers could perform analysis without extensive manual preprocessing.

A first major contribution of this work was the restructuring of the backend system. The platform's API layer was refactored to provide consistent and well-defined endpoints for each cognitive task. The domain models were cleaned and unified, removing duplicated fields and aligning naming conventions across tasks. Data-validation mechanisms were strengthened: input structures are now validated before being written to the database, ensuring that malformed datasets or partially completed trials cannot corrupt research outputs. Furthermore, the entire backend was containerised using Docker, allowing development, testing, and production environments to behave identically. This eliminated a class of deployment-specific bugs that previously caused inconsistent task execution or failed submissions.

The second major area of improvement involved the frontend, which is responsible for delivering cognitive tasks to participants and capturing trial-level behavioural responses. The platform uses an Angular-based architecture that originally contained duplicated timing logic, inconsistent timestamp formats, and fragile event-handling mechanisms. During this thesis, timing logic was unified across all cognitive tasks; high-resolution timestamps were introduced systematically; and stimulus-presentation routines were standardised to minimise delay, drift, and jitter. Performance optimisations reduced user-side freezes, improved responsiveness, and eliminated several sources of UI-related noise in reaction-time measurements. A major refactoring of the administrative dashboard further improved the handling of large datasets, reducing the computational cost of export routines from $O(N^2)$ to $O(N)$ through the introduction of lookup tables and caching mechanisms. This makes the system capable of supporting larger numbers of participants without degradation in performance.

A substantial portion of the work focused on data-processing improvements. Before the thesis, each task generated output in a different structure, often with missing fields, inconsistent naming, or undocumented encoding. This made cross-task comparisons difficult and introduced risk into long-COVID analyses that depend on multiple behavioural indicators. A unified schema was developed and adopted across all tasks. Each trial now includes complete and precise timing information, stimulus identifiers, user responses, derived behavioural markers, and long-COVID classification flags. Encoding and decoding logic was centralised, eliminating ambiguity and reducing preprocessing effort. As a result, datasets exported from the system are now immediately ready for statistical analysis, without requiring manual restructuring.

The deployment and development workflow was also significantly improved. Prior to this thesis, deployments required manual steps, increasing the risk of configuration mismatches. The introduction of Docker-based environment standardisation, the addition of automated startup routines, and the cleaning of environment-variable handling contributed to a more predictable and reliable deployment process. While additional automation such as continuous integration, automated testing, or monitoring pipelines could further enhance stability, the current system already shows markedly reduced operational issues.

Collectively, these improvements produced measurable benefits for system stability, data completeness, reproducibility, and scientific usability. The platform now starts reliably across environments, exhibits fewer backend and frontend errors, and handles high traffic with greater consistency. Datasets collected from participants contain precise, complete, and interpretable behavioural markers. Export workflows in the administrative dashboard are more efficient, less error-prone, and capable of handling larger volumes of data. For researchers, this translates directly into reduced preprocessing time, fewer corrections, and increased confidence in

statistical analyses.

Despite these advances, several limitations remain. Browser-based experiments are inherently affected by hardware variability: differences in display refresh rate, input-device latency, and CPU load introduce small but unavoidable noise into reaction-time measurements. While the improvements in this thesis mitigate these issues through unified timing logic and standardised stimulus routines, they cannot fully eliminate them. Additionally, the current set of cognitive tasks covers a limited range of behavioural domains. Although the platform is now structurally prepared to integrate additional tasks, further development is required to extend its coverage to memory, executive function, or multimodal cognitive assessments.

Nevertheless, the platform now represents a solid technical foundation that can support long-term behavioural research into COVID-19 neurological effects. It is robust enough to handle population-level data collection, flexible enough to incorporate new tasks, and transparent enough to produce datasets suitable for cross-team collaboration and longitudinal analysis. The improvements introduced in this thesis are therefore not only technical upgrades, but essential steps toward enabling scientifically reliable digital experimentation in public-health research.