

From At-Home Menstrual Self-Tracking to Server-Side Analytics and Decision Support: The End-to-End Architecture of Ovufriend 2.0

Lukasz Sosnowski^{1*}

Systems Research Institute, Polish Academy of Sciences
Newelska 6, 01-447 Warsaw, Poland
`sosnowsl@ibspan.waw.pl`

Abstract. Mobile health applications can support long-term health monitoring outside clinical environments, but their effective operation requires the integration of data acquisition, server-side processing, analytical algorithms, and interpretable system outputs. This paper presents the end-to-end architecture of Ovufriend 2.0, a mobile platform developed for at-home menstrual cycle and reproductive health self-tracking.

The system enables users to record daily observations, including cycle events, basal body temperature, bleeding, cervical mucus, symptoms, moods, ovulation tests, medications, and survey-based information. The collected data are processed on the server and transformed through a hierarchy of detector-, cycle-, and user-level concepts. These representations support both individual cycle assessment and longitudinal analysis of recurring patterns.

The analytical layer combines rule-based, fuzzy, granular, and data-driven methods. Selected artificial intelligence techniques are used to estimate or adjust decision thresholds, while dedicated modules support ovulation assessment, cycle characterization, symptom analysis, and detection of selected health-related irregularities. The architecture additionally supports adaptive data enrichment: analytical conditions can trigger targeted surveys or additional user interaction when routinely collected information is insufficient for a more specific assessment.

The paper provides a system-level synthesis of the platform, covering mobile data collection, server-side processing, analytical module integration, adaptive information acquisition, and communication of decision-support results. It also discusses challenges related to incomplete reporting, temporal heterogeneity, interpretability, and responsible use of automated health indications.

Keywords: Mobile health Menstrual health Decision support
Longitudinal health data Adaptive data acquisition

1 Introduction

Mobile health applications make it possible to collect longitudinal health information outside clinical environments and over periods that would be difficult to

* ORCID: 0000-0003-2388-4008

observe during conventional medical visits. Menstrual and fertility tracking is a particularly relevant example, as many potentially informative patterns emerge only through repeated observations collected over consecutive days and menstrual cycles. At the same time, user-generated health data introduce substantial analytical challenges because reporting is voluntary, irregular, and heterogeneous, while individual observations may have limited meaning when considered in isolation [1–3].

Ovufriend was originally developed as a fertility-support platform combining menstrual-cycle self-tracking with computational methods for ovulation and fertile-window analysis. Earlier work focused on such problems as similarity-based fertile-day detection, multivariate ovulation-window detection, and the construction of interpretable cycle-level representations from longitudinal observations [4–7]. The Ovufriend 2.0 project extended this approach beyond fertility prediction toward the identification of health-related irregularities and patterns potentially associated with menstrual and hormonal disorders.

This extension required substantially more than the implementation of additional classifiers. Different health-related problems depend on different combinations of observations and different temporal contexts. Some relevant events can be identified from measurements within a single day, others require aggregation of an entire menstrual cycle, while persistent irregularities can only be evaluated across several cycles or longer periods of user history. Moreover, routinely collected self-tracking data are not sufficient for every analytical problem. Additional contextual information may need to be acquired selectively through questionnaires or other interactions with the user [8, 9].

Consequently, Ovufriend 2.0 was designed as an end-to-end, multi-layer decision-support platform rather than as a collection of independent prediction models. The platform combines a cross-platform mobile application, a central server-side infrastructure, a dedicated analytical backend, persistent longitudinal data storage, a hierarchical representation of health-related concepts, specialized analytical modules, and mechanisms for adaptive data enrichment. The analytical results are subsequently returned to the mobile application, creating a closed information-processing loop between at-home observation and personalized decision support.

A key architectural concept is the transformation of heterogeneous longitudinal observations into a hierarchy of increasingly aggregated representations. Local detector-level concepts describe events identified from detailed observations, cycle-level concepts summarize properties of complete menstrual cycles, and user-level concepts capture persistent characteristics across longer temporal windows [7–9]. The hierarchy allows different analytical modules to reuse a common representation while operating at the temporal level appropriate to a particular health-related problem. The platform additionally supports adaptive information acquisition: preliminary analytical results can activate targeted surveys or other interactions when additional data are required for a more specific assessment [10, 11].

Recent work has also emphasized modular architectures for longitudinal menstrual-cycle monitoring, including the need to accommodate irregularly sampled and incomplete real-world observations [12]. The present paper differs from such monitoring-oriented architectures, as well as from earlier condition-specific Ovufriend studies, by focusing on the integration of mobile data acquisition, multi-level semantic representations, adaptive data enrichment, condition-specific analytical modules, and user-facing feedback within a single operational decision-support architecture.

The resulting architecture was developed and deployed as part of the Ovufriend 2.0 research and development project and integrates application services implemented in PHP with a Java-based analytical backend and PostgreSQL data storage. Communication between the major server-side components is performed through REST APIs, while the server infrastructure is deployed in the Amazon Web Services (AWS) cloud. The mobile application, implemented using the Ionic framework, provides a common user-facing environment for Android and iOS and serves both as a source of longitudinal observations and as a presentation layer for personalized analytical results.

The objective of this paper is to present the end-to-end architecture of Ovufriend 2.0 and to analyze the design decisions required to integrate longitudinal self-tracking, multi-level health-data processing, adaptive information acquisition, and personalized decision support within an operational mobile health platform. The main contributions of the paper are:

- an end-to-end architecture connecting mobile menstrual self-tracking with server-side analytical and decision-support services;
- a multi-level representation integrating daily observations, cycle-level concepts, and longitudinal user-level characteristics;
- an adaptive analytical workflow in which existing observations can trigger selective acquisition of additional condition-specific information;
- a modular integration model separating mobile application logic, application services, persistent data management, and analytical processing;
- design lessons resulting from deploying a decision-support architecture operating on incomplete, irregular, and heterogeneous longitudinal user-generated health data.

The remainder of the paper is organized as follows. Section II introduces the system objectives and main design assumptions. Section III presents the end-to-end architecture of Ovufriend 2.0. Section IV describes the processing of longitudinal health data at daily, cycle, and user levels. Section V discusses the analytical and decision-support modules, while Section VI presents their integration within the operational workflow. Section VII summarizes the main system-level challenges and design lessons, and the final section concludes the paper.

2 System Objectives and Design Assumptions

The architecture of Ovufriend 2.0 was designed around the assumption that meaningful health-related patterns in menstrual data are rarely represented by a single observation. Instead, they emerge from the combination of heterogeneous measurements, symptoms, moods, cycle characteristics, survey responses, and longer-term user history. Consequently, the system must support both detailed temporal data and progressively aggregated representations.

A second assumption concerns the nature of user-generated data. Information collected in a mobile health application is inherently incomplete and irregular. Users differ in reporting frequency, the types of variables they record, and the number of menstrual cycles available for analysis. The analytical architecture therefore cannot depend on uniformly sampled or complete records and must explicitly support missing and partially observed information [1, 2, 13].

The system was also designed to separate routine self-tracking from more specific health-related data acquisition. Standard application use should not require users to complete an extensive set of medical questionnaires. Additional information is acquired selectively when the existing longitudinal history activates a condition-specific analytical process. This assumption led to the integration of boundary-condition evaluators, an intelligent communication mechanism, and reusable survey modules into the decision-support workflow.

Another design objective was the reuse of common analytical representations across multiple health-related modules. Rather than constructing an independent data-processing pipeline for each condition, Ovufriend 2.0 introduces a shared hierarchy of detector-, cycle-, and user-level concepts. Individual analytical modules can then select and combine the concepts relevant to their specific decision problem [8, 9].

From a software-engineering perspective, the architecture was required to separate the user-facing application from the analytical implementation. Mobile and application-level functionality is therefore decoupled from the Java-based analytical backend through REST interfaces. This enables analytical modules to evolve without requiring direct changes to the mobile client.

Finally, Ovufriend 2.0 was designed as a decision-support system rather than a diagnostic system. Its analytical outputs represent detected irregularities, increased vulnerability, or patterns that may justify additional monitoring or medical consultation. Definitive clinical diagnosis remains outside the scope of the platform.

3 End-to-End Architecture of Ovufriend 2.0

Ovufriend 2.0 was designed as an end-to-end information and decision-support environment supporting menstrual health monitoring, fertility awareness, and personalized health-related assessment. The platform integrates user-facing data acquisition, server-side data management, longitudinal health-data processing, specialized analytical modules, and presentation of personalized results to the user.

From an architectural perspective, the system can be divided into three main layers: the mobile application layer, responsible for interaction with the user and acquisition and presentation of health-related information; the backend and data management layer, responsible for persistent storage, validation, preprocessing, and integration of longitudinal data; and the analytics and decision-support layer, which transforms collected observations into higher-level representations and condition-specific assessments.

All major system components communicate with a central backend, which provides a common integration point for mobile clients, longitudinal user records, analytical procedures, and project-related data-processing workflows. This architecture separates user interaction from computational processing while preserving a continuous flow of data from at-home self-tracking to server-side analytics and personalized feedback.

Figure 1 presents the end-to-end architecture of Ovufriend 2.0 and the main data flow between its components. The architecture combines the cross-platform mobile application, application and analytical backends, persistent data storage, and a multi-level analytics and decision-support layer. The figure also illustrates the feedback loop in which analytical results may trigger additional data acquisition and are subsequently returned to the user as personalized decision-support information.




Fig. 1. End-to-end architecture of the Ovufriend 2.0 platform. The system integrates mobile data acquisition, application and analytical backends, multi-level longitudinal data processing, adaptive data enrichment, and personalized decision support.

The architecture follows a closed-loop processing model. Information provided by the user is first collected through the mobile application and transferred to the server-side infrastructure. The analytical backend transforms these observations into higher-level concepts and executes specialized decision-support procedures. Depending on the current analytical state, the system can either generate a result directly or request additional information through targeted surveys and the intelligent communication mechanism. The resulting assessment is then returned to the application and presented to the user.

3.1 Mobile Application Layer

The Ovufriend 2.0 mobile application constitutes the main interaction point between the user and the platform. It was implemented using the Ionic framework, providing a common cross-platform implementation for Android and iOS. Most application functionality is delivered directly within the mobile application, while selected content-oriented components, such as expert articles and educational materials, are provided through WebView-based interfaces.

The original role of the application was to support women attempting to conceive by providing tools for menstrual-cycle self-tracking, fertility awareness, and interpretation of cycle-related observations. Its functionality includes a menstrual calendar, recording of daily symptoms and events, basal body temperature tracking, fertility- and ovulation-related observations, graphical presentation of cycle data, notifications, and personalized information related to the current cycle. During the development of Ovufriend 2.0, this functionality was extended with modules supporting the assessment of selected menstrual and hormonal health-related risks.

Data acquisition is organized primarily through structured forms. Users submit daily observations, cycle-related information, questionnaire responses, and other health-related data through the application interface. The system also supports uploading attachments, including results of medical examinations, which can be transferred to project coordinators as part of the data collection and project-support workflow.

An important characteristic of the system is that its role changed between the research-and-development phase and the subsequent operational use. During system development, participating users enrolled in the project and contributed longitudinal observations, questionnaire data, and selected medical information. Project participants were also provided with medical examinations as part of their participation. The collected data supported the development and evaluation of analytical algorithms and the extension of the mobile application.

After completion of the development phase, the resulting analytical modules became part of the operational platform. Current users can therefore not only record their observations but also receive personalized results generated by the analytical components of the system. The mobile application consequently serves both as a data-acquisition interface and as the presentation layer for server-side decision-support results.

From the architectural perspective, data entered through application forms are submitted to a central backend, where they are stored and integrated with the user’s longitudinal history. More complex processing, including feature construction, temporal aggregation, and execution of health-related analytical procedures, is performed on the server side. The results can subsequently be returned to the mobile application and presented to the user together with other cycle- and fertility-related information.

3.2 Backend and Data Management Layer

The server-side architecture of Ovufriend 2.0 is divided into two cooperating backend components with different responsibilities: an application backend implemented in PHP and an analytical backend implemented in Java. The separation allows application-oriented functionality to evolve independently from computationally intensive health-data processing and decision-support algorithms.

The PHP backend provides the server-side services required by the mobile application. It handles application-level communication, coordinates user requests, and acts as an integration layer between the mobile client and the analytical services. Whenever analytical processing is required, the application backend communicates with the Java-based backend through REST APIs and subsequently makes the resulting information available to the mobile application.

The Java backend constitutes the analytical core of the platform. It is responsible for processing longitudinal menstrual-health data, constructing derived representations, executing analytical procedures, and generating the results used by the decision-support functionality. Its persistent data layer is based on PostgreSQL, which stores structured information required by the analytical processes and supports access to longitudinal user and cycle histories.

Communication between the server-side components is based on REST APIs. This provides a clear service boundary between application functionality and analytical processing and avoids direct dependencies between the mobile client and the internal implementation of the analytical algorithms. Consequently, analytical modules can be modified or extended without requiring corresponding changes in the core mobile application architecture.

The complete server-side environment is deployed in the Amazon Web Services (AWS) cloud. Cloud deployment provides a common infrastructure for the application backend, analytical services, data storage, and communication interfaces. From the perspective of the mobile application, the server-side components form a unified platform, while internally they remain separated according to their application, data-management, and analytical responsibilities.

3.3 Analytics and Decision-Support Layer

The analytics and decision-support layer constitutes the computational core of Ovufriend 2.0. Rather than being implemented as a single predictive model, it consists of a set of cooperating analytical components operating on different

temporal and semantic levels. These components transform heterogeneous longitudinal observations into progressively higher-level representations that can be used to identify fertility-related patterns and potential health irregularities [8, 9].

A central element of this architecture is a hierarchical representation of medical and menstrual-cycle concepts. The lowest level describes local events and observations detected within an ongoing cycle, such as ovulation-related events or other temporally localized phenomena. The second level aggregates these observations into descriptors of a complete menstrual cycle. These cycle-level concepts may represent, for example, abnormal cycle length, characteristics of the luteal phase, symptom recurrence, bleeding patterns, or indicators derived from ovulation analysis. The highest level aggregates information across the user's history and represents persistent characteristics that cannot be reliably inferred from an isolated observation.

The representation combines binary and fuzzy concepts. In particular, cycle-level concepts can express different degrees of occurrence rather than only the presence or absence of a characteristic. This allows complex longitudinal time-series data to be transformed into compact and interpretable intermediate representations. Such representations are subsequently used as inputs to condition-specific analytical modules.

The analytical workflow is adaptive. Initially, the system relies primarily on information routinely collected through menstrual-cycle tracking. Boundary-condition evaluators and analysis activators identify situations in which the available data justify additional processing. If the existing observations indicate a potentially relevant pattern, the system can initiate a data-enrichment process. Additional information is then collected through dedicated surveys or interaction with the intelligent communication module. This mechanism makes it possible to request more specific information only when it becomes relevant to a particular analytical process [10, 11].

Condition-specific modules operate on different combinations of cycle characteristics, symptoms, moods, survey responses, and longitudinal user descriptors. Examples include modules related to ovulation and fertility analysis, luteal-phase abnormalities, PMS/PMDD-related patterns, PCOS-related irregularities, and thyroid-related disorders. Although their internal decision procedures differ, they share the same general processing architecture and reuse the common hierarchy of intermediate concepts.

The resulting outputs represent health-related vulnerability, suspicion, or risk rather than definitive medical diagnoses. When a relevant pattern is identified, the decision-support layer can provide the application layer with information to be presented to the user, including an indication that medical consultation may be advisable. The analytical result is transferred through the server-side integration interfaces and becomes part of the personalized feedback available in the mobile application.

This architecture therefore creates a closed information-processing loop. User-generated observations are transformed into structured longitudinal representa-

tions, analyzed by specialized modules, and returned to the user as personalized decision-support information. At the same time, the system can actively request additional data when the initially available observations are insufficient for a more specific assessment.

4 Longitudinal Health Data Processing

4.1 Daily Observations

The lowest level of the Ovufriend 2.0 data model consists of observations recorded during everyday use of the mobile application. These data form heterogeneous longitudinal time series whose density and completeness depend on the individual user’s tracking behavior [2, 13].

The collected observations include fertility- and ovulation-related measurements, such as basal body temperature, cervical mucus characteristics, cervical position and opening, and ovulation-test results. In parallel, users can record physical symptoms, mood-related observations, bleeding events, and other cycle-related information. The platform also maintains less frequently changing user attributes and longer-term health information.

Individual observations are associated with a specific day and, where applicable, with the corresponding menstrual cycle. They constitute the most detailed temporal representation available in the system and provide input to detector-level analytical procedures.

The system does not require all supported variables to be reported on every day. Consequently, a user’s history may contain irregularly sampled and incomplete sequences. This property is inherent to real-world self-tracking data and is taken into account by subsequent aggregation mechanisms rather than being removed at the data-acquisition stage.

In addition to routinely recorded observations, the dataset can be enriched with information obtained through condition-specific surveys and uploaded medical information. Such additional data are collected when required by the analytical workflow and are integrated with the existing longitudinal user history.

4.2 Cycle-Level Aggregation

Daily observations provide a detailed representation of the user’s menstrual history, but many health-relevant cycle characteristics cannot be inferred reliably from isolated measurements. Ovufriend 2.0 therefore introduces a cycle-level representation in which local events and repeated observations are aggregated into descriptors of the complete menstrual cycle [7, 8].

The aggregation mechanism is based on a hierarchical organization of concepts. Detector-level concepts represent local events identified during an ongoing cycle. These include, for example, events associated with ovulation detection or other characteristics that can be identified from observations available at a particular point in time. Detector-level concepts are binary and indicate whether a specified condition has been detected.

When a menstrual cycle is completed, information collected during the cycle is transformed into cycle-level concepts. These concepts combine lower-level events with aggregated characteristics derived from symptoms, moods, bleeding patterns, fertility observations, and other cycle-related information. Consequently, a multidimensional time series describing an entire menstrual cycle is represented by a smaller set of semantically interpretable attributes.

Unlike detector-level concepts, cycle-level concepts can be fuzzy. Their values express the degree to which a particular characteristic is present in the analyzed cycle [14, 15, 7]. In the implemented system, a four-level representation is used, with values 0, 0.33, 0.66, and 1 corresponding to the absence, low, medium, and high occurrence of a given characteristic, respectively.

An important aspect of this aggregation is normalization with respect to the temporal structure of the cycle. For characteristics based on repeated symptoms, the system can use relative frequency rather than the absolute number of observations. For example, the number of days on which a symptom occurs can be related to the total cycle length. This makes the resulting representation less dependent on whether the analyzed cycle is unusually short or long.

Cycle-level concepts may describe a wide range of characteristics, including cycle length abnormalities, possible absence of ovulation, luteal-phase characteristics, intermenstrual bleeding, recurrent physical symptoms, mood-related patterns, and information derived from targeted surveys. These concepts form a common intermediate representation that can be reused by multiple condition-specific analytical modules.

4.3 User-Level Representation

The highest level of the Ovufriend 2.0 representation describes characteristics of the user that cannot be derived reliably from a single daily observation or an isolated menstrual cycle. Its purpose is to identify persistent or recurrent patterns occurring across the user’s longitudinal history [8, 9].

User-level concepts are constructed primarily from lower-level cycle concepts evaluated within predefined temporal windows. Depending on the analytical context, the system can examine several months of menstrual history and determine whether particular cycle-level characteristics occur repeatedly with sufficient intensity. In the implemented platform, temporal windows of 3, 6, and 12 months are used for selected user-level analyses.

This aggregation makes it possible to distinguish an isolated irregularity from a recurrent characteristic. For example, an abnormally long cycle or a symptom observed once does not necessarily represent a persistent health-related pattern. Repeated occurrence of the corresponding cycle-level concepts within the analyzed history can, however, activate a higher-level concept describing the user’s long-term characteristics.

Whereas many cycle-level concepts use fuzzy values to express the degree of occurrence of a characteristic, user-level concepts can also serve as binary indicators describing whether a specified longitudinal condition has been satisfied.

They are derived from the lower levels of the hierarchy and provide a compact representation of the user’s recent menstrual and health history.

The user-level representation can additionally incorporate information that is not naturally associated with an individual cycle. Examples include selected results of medical examinations, longer-term profile attributes, family-history information, and outcomes of surveys whose validity extends beyond the cycle in which they were completed. Such information can be combined with recurrent cycle-level characteristics when constructing condition-specific analytical inputs.

As a result, Ovufriend 2.0 transforms a heterogeneous collection of daily observations and cycle histories into a progressively more compact representation:

$$\begin{aligned} &\text{daily data} \rightarrow \text{detector-level concepts} \\ &\rightarrow \text{cycle-level concepts} \rightarrow \text{user-level concepts.} \end{aligned} \tag{1}$$

This hierarchy provides the common informational basis for the condition-specific decision-support modules described in the following section.

5 Health Analytics and Decision-Support Modules

5.1 Common Analytical Workflow

Although the analytical modules implemented in Ovufriend 2.0 address different health-related problems, they follow a common decision-support workflow. Each module operates on selected elements of the hierarchical representation described in the previous section and transforms them into a condition-specific assessment.

The processing starts with the identification of a set of medically meaningful concepts relevant to a particular analytical problem. These concepts are defined with the support of medical experts and are mapped onto data available in the platform [8, 9]. Depending on the module, the input may include detector-level events, cycle-level fuzzy concepts, user-level longitudinal characteristics, survey results, and selected medical information.

The concepts are organized into groups representing different aspects of the analyzed condition. Their contribution to the final assessment may differ according to their medical relevance. In the implemented modules, combinations of concept values are used to calculate an intermediate score or ranking representing the degree of consistency between the user’s data and a condition-specific pattern [8, 10].

For selected modules, parameters such as concept-group weights or decision thresholds can additionally be calibrated from reference data using data-driven optimization procedures.

The general processing scheme can be represented as

$$\mathcal{D} \rightarrow \mathcal{C} \rightarrow \mathcal{G} \rightarrow S \rightarrow R, \tag{2}$$

where $\mathcal{D}$ denotes the available longitudinal user data, $\mathcal{C}$ is the set of derived medical concepts, $\mathcal{G}$ represents condition-specific groups of concepts, S is an

intermediate score or analytical representation, and R denotes the resulting risk-related or decision-support assessment.

The final result should not be interpreted as a medical diagnosis. The purpose of the analytical layer is to identify patterns that may justify increased user attention, acquisition of additional information, or consultation with an appropriate medical specialist.

5.2 Condition-Specific Analysis

The common architecture supports multiple condition-specific analytical modules. Their main differences concern the type of concepts used, the temporal scope of analysis, and the rules used to combine individual indicators.

Ovulation- and fertility-related modules operate primarily on measurements and events occurring within a menstrual cycle, including basal body temperature, cervical mucus observations, cervical characteristics, and ovulation-test results. Their outputs provide important intermediate information for higher-level analyses because the estimated ovulation status and cycle phases can subsequently be used by other modules.

Modules related to luteal-phase abnormalities combine information such as luteal-phase length, bleeding patterns, previous reproductive history, and the duration of attempts to conceive. These variables are aggregated into cycle-level concepts and subsequently used to determine whether recurrent abnormalities are present in the user’s history.

PMS/PMDD-related analysis focuses primarily on the temporal distribution and intensity of physical symptoms and moods within the menstrual cycle. In this case, the timing of observations is particularly important because the analytical process must distinguish symptoms concentrated in the premenstrual part of the cycle from symptoms that occur independently of the cycle phase.

PCOS- and thyroid-related modules use a broader range of heterogeneous information. In addition to cycle characteristics, they may use recurrent physical symptoms, mood-related observations, long-term user attributes, survey responses, and selected medical information. Because many of these indicators are non-specific when considered individually, the decision-support process combines multiple weak signals into a higher-level assessment.

Despite these differences, all modules reuse the same longitudinal data infrastructure, concept hierarchy, aggregation mechanisms, and communication interfaces. This reduces duplication between analytical components and allows new condition-specific procedures to be integrated into the existing platform [9, 10].

5.3 Adaptive Data Enrichment and User Feedback

A distinctive feature of Ovufriend 2.0 is that the analytical process is not limited to passive processing of previously collected data. The system can actively request additional information when the available observations indicate that a more detailed assessment may be useful [10, 11].

The process starts with routinely collected self-tracking data. If predefined boundary conditions or analytical activators are satisfied, the system can initiate an additional data-acquisition procedure. The intelligent communication module then interacts with the user and can invite her to complete a condition-specific survey or provide additional information relevant to the current analytical context [10].

The survey mechanism is implemented as a reusable system component. Survey definitions, questions, possible answers, scoring rules, and result categories are maintained independently of the mobile presentation layer. After a survey has been completed, its result is stored in the user’s profile and can be transformed into additional cycle- or user-level concepts. The analytical module can therefore continue its assessment using an enriched representation of the user’s state.

This creates an adaptive closed-loop workflow:

$$\begin{aligned} \text{initial data} &\rightarrow \text{screening} \rightarrow \text{data request} \\ &\rightarrow \text{enrichment} \rightarrow \text{assessment} \rightarrow \text{feedback}. \end{aligned} \quad (3)$$

Depending on the analytical outcome, personalized information can be returned to the mobile application. The system may inform the user about an observed irregularity, suggest further monitoring, request additional information, or indicate that consultation with a medical specialist may be advisable.

The architecture therefore combines prediction with active information acquisition. Instead of requiring every user to provide all possible health-related information in advance, additional data can be requested selectively when they become relevant to a particular decision-support process.

6 Integration and Operational Workflow

6.1 End-to-End Processing Flow

The individual components of Ovufriend 2.0 are integrated into an end-to-end processing workflow connecting at-home self-tracking with server-side analytics and personalized decision support.

The process starts in the mobile application, where the user records cycle-related observations, symptoms, moods, measurements, and other health-related information. Data submitted through application forms are transferred to the application backend, which is implemented in PHP and provides the main server-side interface for the mobile client.

Analytical requests are delegated by the application backend to the Java-based analytical backend through REST APIs. The analytical backend retrieves the required longitudinal information from the PostgreSQL data layer and performs the corresponding processing, including construction or update of detector-, cycle-, and user-level concepts.

Depending on the analytical context, processing can either produce an assessment directly or activate an additional information-acquisition procedure. In the

latter case, the system generates an appropriate interaction that is subsequently presented to the user, for example an invitation to complete a targeted survey. The submitted answers are stored and incorporated into the user’s longitudinal representation, after which the relevant analytical procedure can be executed again using the enriched data.

When the analytical process produces a result, it is exposed through the server-side interfaces to the application backend and subsequently to the mobile application. The user can therefore receive personalized information derived from the complete server-side analytical workflow without direct interaction with the analytical infrastructure.

6.2 Operational Deployment

The server-side components of Ovufriend 2.0 are deployed in the Amazon Web Services (AWS) cloud. The deployment integrates the application backend, Java-based analytical services, persistent data storage, and REST-based communication interfaces within a common operational environment.

A relevant characteristic of the platform is the transition from the research-and-development phase to operational use. During development, participating users contributed longitudinal self-tracking data, questionnaire responses, and selected medical information. These data supported the design, training, and evaluation of the analytical components as well as the development of the user-facing functionality.

After integration of the developed analytical modules with the production platform, the direction of information flow was extended. The mobile application continued to serve as a source of longitudinal observations, but analytical results could also be returned to current users as personalized decision-support information.

This transition required a separation between application logic and analytical processing. The mobile client and the PHP application backend do not depend directly on the internal implementation of individual analytical algorithms. Instead, analytical functionality is exposed through REST interfaces provided by the Java backend. Consequently, individual analytical modules can be maintained, modified, or extended while preserving the overall application architecture.

The resulting platform therefore combines a user-facing fertility application with an independently evolving analytical infrastructure, forming an operational closed-loop system for longitudinal health-data collection, analysis, and feedback.

7 Design Lessons and System-Level Challenges

7.1 Incomplete and Irregular User-Generated Data

A fundamental challenge in the development of Ovufriend 2.0 was that the analytical layer operates on data generated voluntarily by users in everyday con-

ditions. Such data differ substantially from measurements collected under controlled clinical protocols. Users do not necessarily report the same variables every day, and the density of observations may change between cycles and over time [1, 2, 13].

The architecture therefore cannot assume complete and uniformly sampled longitudinal records. Instead, analytical procedures must operate on the information that is actually available and distinguish between the absence of an observation and the absence of a symptom or event.

The hierarchical representation partially addresses this problem by transforming individual observations into more stable cycle- and user-level concepts. Aggregation allows multiple observations to be combined into descriptors such as relative symptom frequency, cycle-length characteristics, or recurrent longitudinal patterns. Consequently, higher-level analytical modules can operate on compact representations without requiring every raw variable to be available at every time point.

At the same time, insufficient data remain meaningful information for the decision-support process. The system may postpone a particular analysis, apply data-quality conditions, or initiate additional data acquisition when the available history is not sufficient for a more specific assessment.

7.2 Multiple Temporal Scales

A second design challenge results from the temporal nature of menstrual health data. Relevant information exists simultaneously at several time scales. Some events can be detected within a single day, others become meaningful only after completion of a menstrual cycle, while persistent health-related patterns require analysis of several cycles or multiple months of history [2, 9].

For this reason, Ovufriend 2.0 does not rely on a single flat feature representation. Detector-level concepts describe local events, cycle-level concepts summarize complete menstrual cycles, and user-level concepts represent recurrent characteristics observed over longer temporal windows.

This separation is particularly important because different decision-support problems require different temporal contexts. Ovulation-related analysis is strongly dependent on local sequences of measurements within an individual cycle, whereas persistent hormonal irregularities may require evidence accumulated over several cycles. Similarly, PMS-related analysis depends not only on symptom occurrence but also on its temporal localization within the cycle.

The use of a common multi-level representation makes these different time scales available to analytical modules without requiring each module to reconstruct the complete longitudinal history independently.

7.3 Passive Tracking versus Active Data Acquisition

Requiring users to provide every potentially relevant medical and contextual variable during routine application use would create a substantial interaction

burden. On the other hand, relying only on standard menstrual self-tracking would leave some analytical problems underspecified.

Ovufriend 2.0 addresses this trade-off by combining routine data collection with adaptive data enrichment [10, 11]. Most users initially provide only information naturally associated with fertility and menstrual self-tracking. Additional questions are introduced selectively when the analytical state indicates that more specific information may improve a condition-related assessment.

The intelligent communication and survey mechanisms make this process part of the analytical workflow. Boundary conditions derived from the existing user history can activate a targeted interaction, after which the newly obtained information becomes part of the user’s analytical representation.

This design reduces unnecessary questionnaires while preserving the possibility of obtaining more specific information when needed. From a decision-support perspective, the system therefore performs not only data analysis but also adaptive information acquisition.

7.4 Decision Support versus Medical Diagnosis

An important system-level requirement is the separation between automated health-related assessment and medical diagnosis. The analytical modules of Ovufriend 2.0 are intended to identify patterns, irregularities, or increased vulnerability that may justify additional attention, rather than to establish definitive clinical diagnoses [11].

This distinction influences both the analytical architecture and the user-facing workflow. Outputs of condition-specific modules are interpreted as decision-support information and can lead to further monitoring, additional data collection, or a recommendation to consult an appropriate medical specialist.

Such an approach is particularly relevant when the input is based largely on self-reported observations. Individual symptoms may be non-specific, incomplete, or shared by multiple conditions. Combining them into higher-level concepts improves their analytical usability, but does not remove the need for professional medical assessment when a clinically relevant abnormality is suspected.

The role of the platform is therefore to support the transition from long-term at-home observation to professional healthcare by surfacing patterns that may justify further attention and presenting them in a form that can support timely follow-up.

Taken together, these challenges indicate that the key architectural problem is not only the construction of accurate analytical models, but the organization of an end-to-end process capable of operating on irregular longitudinal data, reasoning across multiple temporal scales, acquiring additional information selectively, and returning results to the user in an appropriate decision-support form.

8 Conclusions

This paper presented the end-to-end architecture of Ovufriend 2.0, a mobile health platform that integrates longitudinal menstrual self-tracking with server-side analytics and personalized decision-support functionality. The system was designed to process heterogeneous and irregular user-generated health data collected over multiple temporal scales and to transform them into progressively higher-level analytical representations.

A central architectural feature is the separation of the mobile application, application backend, analytical backend, and persistent data layer, connected through REST-based interfaces. This separation allows user-facing functionality and analytical services to evolve independently while remaining part of a common operational workflow. The hierarchical representation of detector-, cycle-, and user-level concepts provides a reusable basis for multiple condition-specific analytical modules and supports reasoning over both local events and longer-term health patterns.

Another important element is the adaptive data-enrichment mechanism. Instead of requiring users to provide all potentially relevant health information in advance, the system can identify situations in which additional data may be useful and selectively activate targeted surveys or other interactions. The resulting information becomes part of the longitudinal user representation and can subsequently be used by the analytical modules. This creates a closed-loop decision-support process connecting routine self-tracking, server-side analysis, additional information acquisition, and personalized feedback.

The experience gained during the development and operational integration of Ovufriend 2.0 shows that the main challenge in real-world mobile health analytics is not limited to the construction of individual predictive models. Equally important is the design of an architecture capable of handling incomplete longitudinal records, combining information across different temporal scales, selectively acquiring missing contextual information, and delivering analytical results back to the user in an appropriate form.

Ovufriend 2.0 is intended to support awareness and early identification of potentially relevant health-related patterns rather than to replace medical diagnosis. Future development can therefore focus on further personalization of data-acquisition strategies, improved estimation of data quality and uncertainty, extension of the concept hierarchy, and integration of additional analytical modules while preserving the modular architecture presented in this paper.

Acknowledgments

This work was co-financed by the EU Smart Growth Operational Programme 2014–2020 under the project “Developing innovative solutions in the domain of detection of frequent intimate and hormonal health disorders in women of procreative age based on artificial intelligence and machine learning – Ovufriend 2.0”, POIR.01.01.01-00-0826/20.

References

1. Schantz, J.S., Fernandez, C.S.P., Jukic, A.M.Z.: Menstrual cycle tracking applications and the potential for epidemiological research: A comprehensive review of the literature. *Current Epidemiology Reports* **8**(1), 9–19 (2021). DOI 10.1007/s40471-020-00260-3
2. Li, K., Urteaga, I., Wiggins, C.H., Druet, A., Shea, A., Vitzthum, V.J., Elhadad, N.: Characterizing physiological and symptomatic variation in menstrual cycles using self-tracked mobile-health data. *npj Digital Medicine* **3**, 79 (2020). DOI 10.1038/s41746-020-0269-8
3. Karasneh, R.A., Al-Azzam, S.I., Alzoubi, K.H., Muflih, S.M., Hawamdeh, S.S.: Smartphone applications for period tracking: Rating and behavioral change among women users. *Obstetrics and Gynecology International* **2020**, Article 2192,387 (2020). DOI 10.1155/2020/2192387
4. Sosnowski, L., Chaber, W., Milobedzki, L., Penza, T., Sosnowska, J., Zaleski, K., Fedorowicz, J., Szymusik, I., Slezak, D.: Similarity-based detection of fertile days at ovufriend. In: 2018 IEEE International Conference on Big Data (Big Data), pp. 2663–2668 (2018). DOI 10.1109/BigData.2018.8622602
5. Fedorowicz, J., Sosnowski, L., Slezak, D., Szymusik, I., Chaber, W., Milobedzki, L., Penza, T., Sosnowska, J., Wójcicka, K., Zaleski, K.: Multivariate ovulation window detection at ovufriend. In: *Rough Sets, Lecture Notes in Computer Science*, vol. 11499, pp. 395–408. Springer (2019). DOI 10.1007/978-3-030-22815-6_31
6. Sosnowski, L., Szymusik, I., Penza, T.: Network of fuzzy comparators for ovulation window prediction. In: *Information Processing and Management of Uncertainty in Knowledge-Based Systems, Communications in Computer and Information Science*, vol. 1239, pp. 800–813. Springer (2020). DOI 10.1007/978-3-030-50153-2_59
7. Sosnowski, L., Penza, T.: Generating fuzzy linguistic summaries for menstrual cycles. In: *Proceedings of the 2020 Federated Conference on Computer Science and Information Systems, Annals of Computer Science and Information Systems*, vol. 21, pp. 119–128 (2020). DOI 10.15439/2020F202
8. Sosnowski, L., Wróblewski, J.: Toward automatic assessment of a risk of women’s health disorders based on ontology decision models and menstrual cycle analysis. In: 2021 IEEE International Conference on Big Data (Big Data), pp. 5544–5552 (2021). DOI 10.1109/BigData52589.2021.9671481
9. Sosnowski, L., Żuławińska, J., Dutta, S., Szymusik, I., Zyguła, A., Bambul-Mazurek, E.: Artificial intelligence in personalized healthcare analysis for womens’ menstrual health disorders. In: *Proceedings of the 17th Conference on Computer Science and Intelligence Systems, Annals of Computer Science and Information Systems*, vol. 30, pp. 751–760 (2022). DOI 10.15439/2022F59
10. Sosnowski, L., Dutta, S., Szymusik, I., Chaber, W., Kasprowicz, P.: Interactive, personalized decision support in analyzing women’s menstrual disorders. In: *Communication Papers of the 18th Conference on Computer Science and Intelligence Systems, Annals of Computer Science and Information Systems*, vol. 37, pp. 279–286 (2023). DOI 10.15439/2023F1683
11. Dutta, S., Skowron, A., Sosnowski, L.: Medical decision support in the light of interactive granular computing: Lessons from the ovufriend project. *International Journal of Approximate Reasoning* **165**, 109,103 (2024). DOI 10.1016/j.ijar.2023.109103
12. Cedro, T.B., Poludniewski, G., Glinkowski, W.M.: A modular digital health architecture for longitudinal menstrual cycle monitoring: System design and for-

- mative usability evaluation. *Applied Sciences* **16**(13), 6469 (2026). DOI 10.3390/app16136469
13. Li, K., Urteaga, I., Shea, A., Vitzthum, V.J., Wiggins, C.H., Elhadad, N.: A predictive model for next cycle start date that accounts for adherence in menstrual self-tracking. *Journal of the American Medical Informatics Association* **29**(1), 3–11 (2022). DOI 10.1093/jamia/ocab182
 14. Zadeh, L.A.: Toward a theory of fuzzy information granulation and its centrality in human reasoning and fuzzy logic. *Fuzzy Sets and Systems* **90**(2), 111–127 (1997). DOI 10.1016/S0165-0114(97)00077-8
 15. Kacprzyk, J., Yager, R.R., Merigó, J.M.: Towards human-centric aggregation via ordered weighted aggregation operators and linguistic data summaries: A new perspective on zadeh’s inspirations. *IEEE Computational Intelligence Magazine* **14**(1), 16–30 (2019). DOI 10.1109/MCI.2018.2881641