Overview of the SMART-BEAR Technical Infrastructure

Vadim Peretokin[1], Ioannis Basdekis[2], Ioannis Kouris[3], Jonatan Maggesi[4], Mario Sicuranza[5], Qiqi Su[6], Alberto Acebes[7], Anca Bucur[1], Vinod Jaswanth Roy Mukkala[4], Konstantin Pozdniakov[6], Christos Kloukinas[6], Dimitrios D. Koutsouris[3], Elefteria Iliadou[8], Ioannis Leontsinis[9], Luigi Gallo[5], Giuseppe De Pietro[5] and George Spanoudakis[2]

[1]Philips Research, Eindhoven, The Netherlands
[2]SPHYNX Technology Solutions AG, Zug, CH, Switzerland
[3]Biomedical Engineering Laboratory, School of Electrical and Computer Engineering, National Technical University of Athens, Athens, Greece
[4]Computer Science Department, Università degli Studi di Milano, Milan, Italy
[5]Institute for High-Performance Computing and Networking, National Research Council of Italy, ICAR - CNR, Naples, Italy
[6]Department of Computer Science, City, University of London, London, U.K.
[7]Atos Research and Innovation. Madrid, Spain
[8]1st Otolaryngology University Department, National and Kapodistrian University of Athens, Athens, Greece
[9]1st Cardiology Clinic, Medical School, National and Kapodistrian University of Athens, Athens Greece

Keywords: Cloud, AI, Semantic Interoperability, HL7 FHIR, Healthcare, GDPR, Evidence-based, Ageing, Hearing Loss, Cardiovascular Disease, Balance Disorder.

Abstract: This paper describes a cloud-based platform that offers evidence-based, personalised interventions powered by Artificial Intelligence to help support efficient remote monitoring and clinician-driven guidance to people over 65 who suffer or are at risk of hearing loss, cardiovascular diseases, cognitive impairments, balance disorders, and mental health issues. This platform has been developed within the SMART-BEAR integrated project to power its large-scale clinical pilots and comprises a standards-based data harmonisation and management layer, a security component, a Big Data Analytics system, a Clinical Decision Support tool, and a dashboard component for efficient data collection across the pilot sites.

1 INTRODUCTION

Providing support for healthy living to an ageing population is a central challenge for EU societies; ageing has a significant social and financial impact due to a higher incidence of health-related issues such as hearing loss, cardiovascular diseases, cognitive diseases, and balance disorders. The current focus in elderly care is to develop solutions for the prevention and effective treatment of age-related ailments.

In this paper, we present the cloud-based data harmonisation and management platform implemented in the SMART-BEAR project [1] that aims to facilitate evidence-based personalised support for elderly patients in their home environment.

The EU-funded SMART-BEAR project develops a platform integrating a variety of sensors and mobile instruments that actively collect data of enrolled patients during their daily life, harmonise these data and analyse them to provide effective recommendations and personalised interventions. The SMART-BEAR solution will be tested in large-scale pilots involving approximately 5 000 senior EU citizens in Portugal, Spain, France, Italy, Romania, and Greece. Prior to the full-scale study, in September 2021, a first small-scale pilot, namely the "Pilot of the Pilots" (PoP) has already started in Madeira, Portugal, targeting to enrol 100 patients by June 2022.

[1] SMART-BEAR: Smart Big Data Platform to Offer Evidence-based Personalised Support for Healthy and Independent Living at Home. https://www.smart-bear.eu/

Peretokin, V., Basdekis, I., Kouris, I., Maggesi, J., Sicuranza, M., Su, Q., Acebes, A., Bucur, A., Mukkala, V., Pozdniakov, K., Kloukinas, C., Koutsouris, D., Iliadou, E., Leontsinis, I., Gallo, L., De Pietro, G. and Spanoudakis, G.
Overview of the SMART-BEAR Technical Infrastructure.
DOI: 10.5220/0011082700003188
In *Proceedings of the 8th International Conference on Information and Communication Technologies for Ageing Well and e-Health (ICT4AWE 2022)*, pages 117-125
ISBN: 978-989-758-566-1; ISSN: 2184-4984

The main components of the SMART-BEAR cloud platform include databases and its underlying information model, clinical repository interfaces, a Big Data Engine, a data collection dashboard, and analytics. The analytics and personalisation components leverage the SMART-BEAR FHIR-based Information Model. Separately from the cloud platform, a smartphone application collects information on the patient's basic physiological, medical and behavioural parameters (such as steps walked daily or weekly, weight in Kgs, or blood pressure). This information is collected by means of smart devices that are provided for one year (12 months) to each participant according to their comorbidities and needs.

In recent years, there has been an increased interest in e-health monitoring systems situated at homes, leading to the creation of Health Smart Homes. Such technologies can facilitate monitoring patients' activities and enable efficient, decentralised healthcare services at home. This type of monitoring may improve the quality of care for the elderly population and increase their well-being in a non-obtrusive way. This approach allows for greater independence and empowerment, maintaining good health longer, preventing social isolation for individuals, and delaying their placement in institutions such as nursing homes and hospitals. The recent advancements in the IoT technology, the improvements with respect to user-friendliness, and the significant cost reduction need to be considered as well. This current wide use (compared to previous periods) was enabled by major advances in wireless technology and computing power, leading to a plethora of diverse and specialised Medical IoT (MIoT) that can generate and transmit data via open protocols – and later, to be picked up and analysed.

It is not just the increase in the supply of affordable MIoT monitoring medical and well-being measurements that is changing the landscape in personalised medicine and consumer health. The data, generated at a rapid rate, along with the devices themselves, are creating a connected infrastructure of medical devices, software applications and health systems and services, that is transforming healthcare delivery. Nowadays, the evolution of e-health systems equipped with Big Data Analytics (BDA) capabilities permits the provision of good quality decision support, enhancing care delivery. The efficient information exchange and data reusability, together with the utilisation of data mining and ML analytics help to convert information into knowledge (Dash, Shakyawar, Sharma, & Kaushik, 2019).

With all the progress achieved in this domain, challenges remain in how to use the information and the derived knowledge productively and in a way that can be evaluated systematically, as the scientific community does not have a commonly accepted way of capturing it, while industry traditionally invests in technological solutions that can be commercially exploited. The HL7 (Health Level Seven) FHIR (Fast Healthcare Interoperability Resources) standard, a well-known specification that can be used for the representation of clinical data, provides the underlying basis for our data harmonisation solution. Accompanied by well-defined semantics captured using widely adopted ontologies such as LOINC and SNOMED-CT for the semantic representation of data, the standardised data representation helps streamline the development of analytics and decision models with the potential to provide accurate, personalised interventions via decision support tools. These tools digest the harmonised information and facilitate decisions that are vetted by health professionals to ensure patient safety.

Along with the production of knowledge, the dimension of data protection needs to be adequately addressed. Processing of sensitive personal data must be compliant with all relevant legal requirements and privacy obligations laid down by national legislation in addition to those imposed by the General Data Protection Regulation (GDPR), a legal framework that fundamentally transformed how personal data must be managed lawfully. In this context, it is not enough just to have in place organisational procedures along with IT-supported processes for exercising certain GDPR rights. Vulnerabilities do happen, even within the best organised and best coded IT systems. Therefore, mechanisms to ensure the security and privacy of the data held, the integrity of any platform storing and managing them (integrity, confidentiality, and availability of data at rest, in transit and processing for data flows), in a continuous security and privacy assurance approach, are of paramount importance. On this axis, and given the legal obligations imposed by the GDPR and the state-of-the-art guidelines (e.g., encryption guidelines of NIST), data minimisation, pseudo/anonymisation, transparency in processing personal data, and audits support are among the appropriate technical (and organisational) measures that must be taken into account, preferably at early stages, to ensure that all legal requirements are met.

Last but not least, Big Data Analytics (BDA) systems for healthcare decision-making must not only focus on the production of ML knowledge but also convey it in an easy-to-use way, accompanied by information that aims to make AI algorithms more understandable by people (AI explainability). Diachronically, e-health systems do not appear to be rated satisfactorily in terms of their usability (Population ageing in Europe Facts, implications and policies) (Basdekis, Sakkalis, & Stephanidis, 2011), while understanding AI remains an open question (Liao, Gruen, & Miller, 2020). These systems, in particular, are being used within a high-stress environment, by non-technical end-users, and perhaps with time constraints that made the situation even worse. Thus, the acceptance and usability by the involved end-users of such functionality is a critical factor for its success and a key requirement in the SMART-BEAR project.

2 THE SMART-BEAR ARCHITECTURE

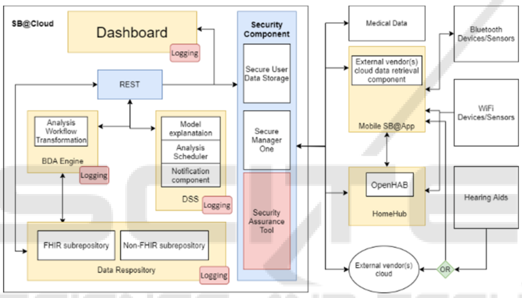

Figure 1: Overall SMART-BEAR architecture.

The architecture consists of three main systems: a mobile phone application (Mobile SB@App), the SB@HomeHub and the SB@Cloud. The overall architecture is depicted in Figure 1. The mobile application allows collecting the data from all portable devices connected to a person's mobile phone (such as heart rate and steps measurements). The HomeHub accumulates data from different sensors of home-based devices (such as movement sensors) directly, as well as from external vendor clouds.

Finally, SB@Cloud is the core system responsible for the secure storage and analysis of the data collected. The main components of the cloud platform include FHIR-compliant/non FHIR-compliant data repository, its underlying information model, the clinical repository interfaces, the Big Data Engine including the synthetic data generation component, the analytics, Decision Support System (DSS) and the Dashboard – a user interface component. The analytics and personalisation components leverage the FHIR-based Information Model. DSS provides the functionality for interventions, reasoning behind the decisions proposed, analysis scheduling and notifications. The clinical repository interfaces allow the accumulation of the data from Electronic Health Records (EHR) external to the project's infrastructure. On the backend level, all components are interconnected via RESTful interfaces. User interaction with the project infrastructure is provided via the dashboard. The communications between components, as well as authentication at the dashboard, are secured according to GDPR via the security component, which assures all security mechanisms are working correctly. The security component also enables interoperability with external platforms that use the FHIR standard to represent medical/usage data. A secure, privacy-preserving machine-to-machine bridge with two platforms, developed within the Smart4Health [2] and Holobalance [3] EU-funded projects, is currently being tested in the PoP (Pilot of Pilots).

3 THE SMART-BEAR CLOUD COMPONENTS

3.1 The Database Implementation

The clinical repository component of SB@Cloud is utilising a combination of FHIR-compliant and non-FHIR databases. All the data that represent medical entities are stored in the FHIR database while data related to non-medical entities are stored in the non-FHIR database of the Cloud Backend. Those contain elements that are not mapped to FHIR models (such as dashboard user settings) and intermediate results of the analytics models which some of them relay data back to the FHIR database. The non-FHIR database is also used to store data transmitted by the HomeHub which is placed in the patient's home and monitors the usage of light sources, temperature and humidity and motion inside the home. The interventions, the notifications and the alerts that are generated by the DSS are stored in the non-FHIR database.

[2] Smart4Health: Citizen-centred EU-HER exchange for personalised health. https://smart4health.eu/

[3] Holobalance: Holograms for personalised virtual coaching and motivation in an ageing population with balance disorders. https://holobalance.eu/

3.2 Data Model Specification Compliant with FHIR

HL7 FHIR is the latest standard from HL7, an international standards development organisation that has been publishing healthcare interoperability standards since 1989. FHIR takes the best of, and builds upon the lessons learnt from the different directions taken previously by HL V2 and HL7 V3 and while applying well-known, modern technologies such as REST and JSON. The standard focuses on implementers first, provides out-of-the-box tooling, is published for free and is free to use.

The standard is spreading rapidly in the international context and the most important international organisations that provide solutions to specific problems in healthcare, like Integrating the Healthcare Enterprise (IHE) (Integrating the Healthcare Enterprise, n.d.), that is an initiative by healthcare professionals and industry to improve the way computer systems share health information and Personal Connected Health Alliance (PCHA) (Personal Connected Health Alliance, n.d.), that is a membership-based Healthcare Information and Management Systems Society (HIMSS) Innovation Company that works for advancing patient/consumer-centred health, wellness and disease prevention by means of the Continua Design Guidelines. The IHE and PCHA are updating their technical specifications to include FHIR.

FHIR was chosen as the standard for clinical data within SMART-BEAR for its speed and ease of implementation, and the fact that modern web technologies such as REST and JSON are an especially good fit for mobile applications, which the project makes use of. FHIR is used nationally in The Netherlands as part of the MedMij project (MedMij, 2022) and Estonia makes use of it in their national electronic health record system, among other countries. Countries such as the Netherlands, Switzerland, Belgium have defined national core profiles for FHIR that standardise clinical information relevant for the countries within. In this context, using a standard that is gaining adoption in Europe enables us to be open-minded about the possibilities of future data exchange.

The nature of the data treated in the project, in accordance with the rules of the FHIR standard, led to the necessity of using a specific Implementation Guide (IG). For this reason, an analysis of the IGs published on the FHIR registries was carried out. Among these, particular attention was paid the Personal Health Device (PHD) (HL7, Personal Health Device Implementation Guide, n.d.) and International Patient Summary (IPS) (HL7, International Patient Summary Implementation Guide, s.d.) IGs.

The PHD IG adapts FHIR resources to convey measurements and supporting data from PHDs to different kind of systems, like platforms for electronic medical records, clinical decision support, etc. The interest for this IG was captured considering that it is based on the Continua Design Guidelines and upon the ISO/IEEE 11073 PHD Domain Information Model (DIM) (Huang, Wang, & Wang, 2020). Regardless, considering that in SMART-BEAR many health data are not acquired by PHDs, but by collecting patients' answers to specific questionnaires, this IG was not considered adequate for the SMART-BEAR project.

The IPS IG defines the rules to produce a document containing the essential healthcare information about a subject of care, designed for supporting unplanned, cross-border care, although it is not limited to it. Although this IG provides an important contribution to identify a minimal, specialty-agnostic, condition-independent, clinically relevant dataset for a patient, it was not considered relevant for the SMART-BEAR project.

For these reasons, in compliance with the FHIR standard and in line with the choices adopted in many European projects, the approach taken for the definition of the information model for the project is to define a dedicated SMART-BEAR IG by profiling a set of identified FHIR resources and individuating the terminologies from international standard code systems as well as internal value sets. The tool chosen for modelling the FHIR information model is SUSHI (FSH School, n.d.), considering that it integrates well with the IG publisher which is an official tool provided by HL7.

Currently, the published IG (implementation guide) consists of 84 profiles (of type Observation, Condition, Questionnaires, Bundle, Patient, DeviceUseStatement, FamilyMemberHistory, MedicationStatement, ResearchSubject), 2 extensions, 33 value Sets, and 133 examples.

3.3 The Clinical Data Repository

The SMART-BEAR (SB) Clinical Data Repository (CDR) is based on the Health Data Hub which is built around the HL7 FHIR standard, structuring and disposing of clinical information using this standard as specification. Therefore, the SB CDR repository stores and serves clinical information in HL7 standardised, safe and scalable way. This allows Big Data Analytics (BDA) and Decision Support System (DSS) developers to focus on having the algorithms or applications that best suit the SMART-BEAR pilots requirements, enabling them to build a common set of solutions and products smoothly connected using standardised data. Medical terminology not fully covered by FHIR will be annotated using SNOMED-CT [4]. The interoperability with some different clinical terminologies (ICD9, LOINC) used across the healthcare industry will be reached by adapting the Atos Terminology Server (ATS). ATS will be customised and implemented in the second phase, after the finalisation of the PoP, and it will provide a RESTful API. This API will allow for safe access to clinical information via interaction with the FHIR database for terminology purposes.

3.4 The Security Component

Data protection is considered a critical issue, especially when dealing with special categories of personal data (Art 9, GDPR). In this context, the SB@Cloud, by virtue of its design, supports privacy. In particular, the Security Component (SB@SC) provides mechanisms that handle data minimisation, authentication and other security and privacy aspects by performing pseudonymisation and resource identifier re-associations (Basdekis, Pozdniakov, Prasinos, & Koloutsou, 2019). This component supports RBAC (Role-based access control) authentication and authorisation of all RESTful API endpoints (Token-based access via encrypted HTTPS connections) to protect the transmission of any (sensitive or not) data, and it also introduces services to cope with the management of privacy-related requests to demonstrate compliance with the GDPR. The way the data is stored in 2 different repositories (i.e., personal data and PII stored encrypted in a separate one, while all pseudonymised medical and usage data in the CDR), allows the analysis of the fully anonymised data to continue after the end of the project, provided that all personal data will be deleted. Thus, after the completion of the SB project, data kept in SB@SC will no longer be needed to conduct the research (e.g., analytics, interventions), and consequently will be erased and not further used for any data process.

In parallel, SB@SC component is also responsible for monitoring, testing, and assessing the security and privacy of all operations of the platform. It will also audit critical components and processes of the infrastructure while leveraging monitoring mechanisms developed in the context of the project to provide an evidence-based, certifiable view of the security posture of the whole platform, along with accountability provisions for changes that occur in said posture and the analysis of their cascading effects. Several built-in security assessments addressing the Confidentiality – Integrity – Availability (CIA) principles among which are custom metrics with respect to the platform's components will be utilised, leveraging an evidence-based approach to provide security and privacy assurance assessments with certifiable results.

3.5 The SMART-BEAR Information Model

As described above, data in SMART-BEAR is partitioned across two databases – clinical data in the FHIR database, and non-clinical or private information that is not exposed to analytics in a non-FHIR one.

FHIR is a platform specification (see section 3.2); it is intended to be constrained for a specific use case and so we profile various resources in the FHIR database for our needs. We naturally use the Patient resource to store basic demographic information such as name, date of birth, and ethnicity, while the bulk of the clinical data resides in Conditions and Observations that are tied to the Encounter resource.

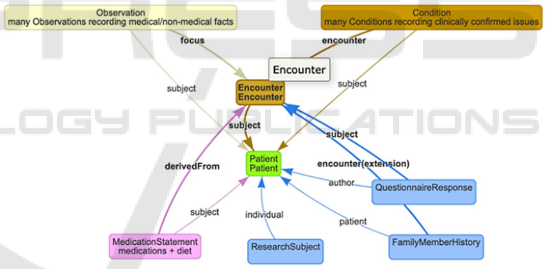

Figure 2: Information model of FHIR resources in use.

Given that we have clinicians performing patient assessments, a single assessment is represented by an instance of an Encounter resource. This Encounter resource is key to the information model as all other resources either link to or from it, creating a graph by which you can reach all the relevant nodes (resources). Any concerning clinical issues noted during an assessment are stored in a Condition resource. Issues or observations of lesser importance are stored as FHIR Observations, which also house 'negations' – issues that a clinician has verified that the patient does *not* have. This fine but important difference between a lack of data (unknown value) and a refuting observation (known negative) allows us to build more accurate analytics algorithms.

[4] https://www.snomed.org

Most Observations follow a simple 'key-value' pattern, where Observation.code identifies the type of measurement and Observation.value[x] records the measurement value. In case of Conditions, Condition.code records the type of condition. As an example, in case of the patient having anxiety Condition.code will be populated with 197480006 |Anxiety disorder| from SNOMED – and should they not be affected by anxiety, Observation.code will have the same terminology code but Observation.valueCodeableConcept will be populated with 260385009 |Negative|.

Where possible, we align with FHIR Vital Signs standard profiles – for example blood pressure, where we record systolic/diastolic measurements using Observation.component, we record the arm (left/right) as bodySite and the patient's position (standing/supine) is part of the LOINC code. Our Implementation Guide (IG) contains a wealth of Condition, and positive/negative Observation examples to assist users in understanding of the FHIR DB.

Specialised resources are used where appropriate; FamilyMemberHistory for example is used to record the family history of hearing loss and ResearchSubject is used to record the source of referral to our clinical study. MedicationStatement records both the list of medications the patient is taking using the WHO ATC value set, and the diet they are prescribed.

A significant part of the data acquired by clinical assessments comes in a form of over 20 Questionnaires; these are internationally recognised, standard data collection points whose outcome scores will be used for analytical purposes.

Previously mentioned Conditions and Observations rely on over 120+ terminology mappings, with most codes coming from SNOMED, to link the semantic meaning within. Codes from LOINC and MESH complement the rest of the mappings. Special care was taken not to create custom codes unless absolutely necessary to avoid the creation of new medical knowledge - just 4 new codes have been introduced that did not have equivalents in any of the searched code systems. Several food/diet-related concepts not available in the SNOMED international core but available in the Australian edition were also made use of for this reason. We verified that this does not impose any additional licensing constraints SNOMED-wise.

Following the theme of avoiding introducing new codes as much as possible, two SNOMED post-coordinated expressions were crafted to accurately represent very specific concepts: "number of non-scheduled visits due to volume overload in subjects with heart failure" as:

4525004 |emergency department patient visit| :362981000 |qualifier value| = 260299005 |number|, 42752001 |due to| = 21639008 |hypervolemia|

and "number of Visits to the ER due to HTN peak" as:

4525004 |emergency department patient visit| :362981000 |qualifier value| = 260299005 |number|, 42752001 |due to| = 38341003 |hypertension|

When it came to recording the patient's ethnicity, and this is an interesting subject because data recorded here really depends on where you happen to be in the world, we chose to re-use the FHIR extension and value set as published by the German Corona Consensus Data Set (GECCO) project, which itself is partially based on WHO ISARIC eCRF valueset. Re-use of existing knowledge like these bolsters long-term interoperability.

Analytics are a crucial aspect of the system, giving it the necessary intelligence for the task at hand. They are driven by the BDA Engine, which has several requirements placed upon it – raw data processing, incremental updates, and scalability.

As mentioned before, clinical data in the system is stored in a FHIR repository. While a FHIR interface has the advantages that make it an excellent choice for clinical data, where it makes compromises is the area of bulk data processing. For this reason, the BDA engine requires a capability to convert and flatten the hierarchical format of FHIR to a relational one that lends itself better to bulk data processing. It should be possible to do this conversion incrementally as new data is arriving in the clinical repository in order to run analytics continuously, and it also needs to be able to scale with large volumes of data.

3.6 The BDA Engine

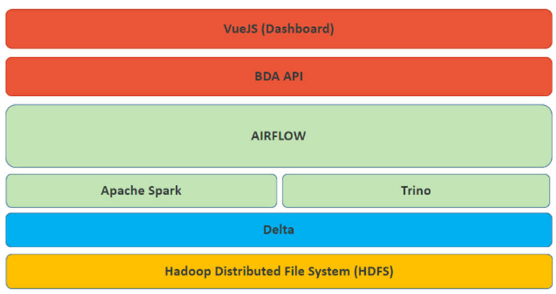

Figure 3: The BDA Engine architecture.

The BDA Engine mainly addresses the functionalities required for processing Data Analysis Workflows (DAWs) and providing/storing the execution results. The BDA Engine exposes a set of APIs to compute and to get raw data to perform analyses. In terms of Machine Learning, a preliminary extraction of data analytics - that will be carried on the pre-processed datasets - are going to indicate variables or combinations of variables for the feature selection approaches. All ML methods and techniques are data-driven, and the "best" method will be decided after its application.

The preliminary extraction of data analytics is performed by the following subcomponents featured in the BDA Engine architecture: Delta Lake [5], Spark[6], Trino[7], Airflow[8]. The components are described here below by following a bottom-up approach, the layer at the bottom being the closest to the data repositories. The architecture is shown in Figure 2 and it is an extended version of the one presented in (Anisetti, et al., 2021).

Delta Lake is an ACID table storage layer over cloud object stores and is the closest component to the repositories. Delta Lake enables to build a Lakehouse Architecture on top of existing storage systems such as Amazon S3, Azure Data Lake Storage (ADLS), Google Cloud Storage (GCS), and Hadoop Distributed File System (HDFS)[9] (Armbrust, et al., 2020). In the case of SMART BEAR, the adopted standard is HDFS.

Spark and Trino are the components collocated on the third layer from the bottom and provide the capability to access data and perform queries on the datasets. Spark is a multi-language engine for executing data engineering, data science, and machine learning on single-node machines or clusters. Spark was chosen because it is capable of processing tasks encompassing custom analytics on large data volumes, and in addition, it features many bindings with other commonly used Data Science and Machine Learning libraries. Spark is also capable to work both on batch and streaming data. Trino is the component providing the capability to access and perform highly parallel and distributed queries on data from multiple systems. Trino was chosen because it provides the BDA Engine with the capability of managing On-Line Analytical Processing (OLAP) queries and data warehousing tasks, and because it can operate on many data sources in addition to data that are stored on HDFS.

Airflow is the component providing the capability to programmatically author, schedule, and monitor workflows that are written in Python programming language. Airflow is the fourth layer from the bottom.

3.7 Decision Support System

The DSS is designed to assist the clinicians in the initial assessment of every patient in terms of the optimal assessments that must be performed to assess the patient and then provide them with the optimal combination of the devices to monitor their health during the pilot study. This component is designed to evolve throughout the project, as it will continuously be trained by the data that will be digested into the platform. The initial version of DSS available for the PoP has adopted the rules and the medical guidelines that have been provided by the clinicians to have a ground truth system based on the most updated medical knowledge. For each of the monitoring conditions of the SMART-BEAR project (Hearing Loss, Cardiovascular Diseases, Mild Cognitive Impairment, Mild Depression, Balance Disorders, and Frailty) the medical teams are providing the rules-based scenarios and the relevant interventions that should be provided to the participants. The rules-based algorithms are taking into consideration the personalised thresholds that are set for each patient individually. For example, for CVDs, optimal and extreme cut-off values are set for the blood pressure, which trigger the generation of notifications and alert to the patient and the clinical care team.

Starting with the PoP, the data that will be collected feed the models of the BDA engine and output of the analytics will be combined with the measured parameters to identify the degree of satisfaction of the patients and to what degree are the personalised thresholds requiring modifications. If the results of the analytics provide insights that lead to a new intervention, the DSS is capable to be extended to support all the new interventions that will be provided by the clinicians. It must be noted that any new intervention must first be validated by the clinicians before it is included in the interventions provided to the patient.

[5] https://delta.io/
[6] https://spark.apache.org/
[7] https://trino.io/
[8] https://airflow.apache.org/
[9] https://hadoop.apache.org/docs/r1.2.1/hdfs_design.html

3.8 Dashboard

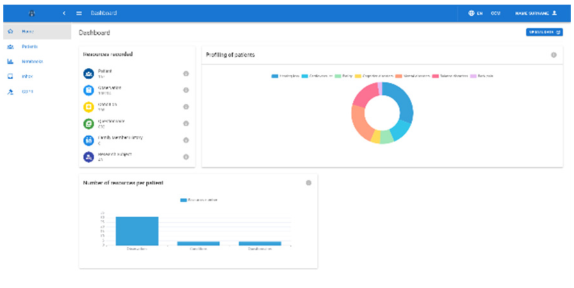

Figure 4: Dashboard homepage.

The SMART-BEAR Dashboard is a component aimed at providing clinicians with a user-friendly graphical user interface. The Dashboard home page is shown in Figure 3. The Dashboard can be used by the clinicians to create and manage a patient, considering his/her devices and medications, conduct the first visit and the check-ups, perform analytics on data and create interventions to be delivered. All the data collected are stored in the FHIR and non-FHIR repositories depending on their clinical value: the first collection is made during the Baseline Assessment from a patient and it concerns the medical history, the physical examinations, and the questionnaire responses, and on the basis of the information provided the dashboard visualises suggestions about the eligibility of the prospective participants of SMART-BEAR pilot studies. A profiling functionality is also featured that suggests the clinician the specific tabs and questionnaires to be activated based on the conditions detected. Although a patient's profile is eventually chosen by a clinician, the profiling functionality redirects the users to the clinical tools and the devices that are required to match a patient's profile consistently with the SMART-BEAR protocol. After a patient is created and deemed eligible, the Dashboard shows specific tabs that enable the patient management and contain information concerning the demographic data, including the living situation and ethnic group, the participation in synergies, and the type and status of provided devices. The patient management tab is shown in Figure 4. Another functionality is featured that is the visualisation of the notifications delivered.

The analytics and intervention mechanisms are in the development phase, and they will make it possible for the clinicians to perform analytics on collected data targeting all patients or only a specific subgroup defined by parameters to monitor in a determined condition in the future. Based on the outcome of the analytics and with support from the DSS, the Dashboard visualises suggestions for clinicians on the interventions to launch, although the final choice is made by a clinician, who will also be able to monitor the intervention outcome. Examples of analytics to be made available in the Dashboard are discussed in (Bellandi, et al., 2021).

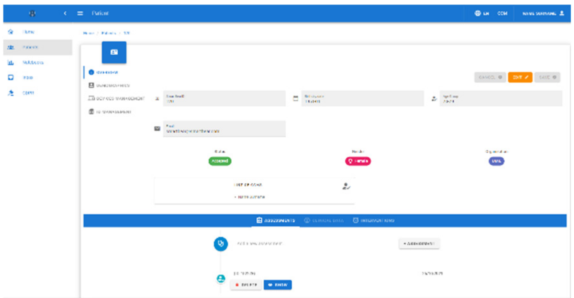

Figure 5: Patient management page.

4 FURTHER WORK

SB@Cloud will operate for at least three years during which the whole solution will be tested and validated through five large-scale pilots involving 5.100 elderly living at home in Greece, Italy, France, Spain, Portugal, and Romania to demonstrate its efficacy, extensibility, sustainability, and cost-effectiveness. During this period, the analysis of the collected data, driven by high level big data analytics and decision models, expected to generate evidence (i.e., metrics, observational evidence base) useful for offering personalised health care and medicine in clinical practice. Still, analysis upon anonymous data can be continued even after the project's lifecycle as the pseudonymisation mechanism in place allows this type of management. To support this notion, SMART-BEAR aims to develop a data sharing and valorisation model (DSVM). This model will identify ways, at a technical and organisational level, for extending the data collected in SMART-BEAR by integrating new data providers and open sources and use the outcomes of data analysis to improve the platform performance, enhance further the personalisation of its relation with its end users, develop new services, and monetise data intensive services out of the platform.

5 CONCLUSIONS

In this paper we have presented an overview of the cloud-enabled standards-based integrated system developed in the SMART-BEAR project, which is able to record assessments for, monitor, and deliver clinician-vetted interventions to senior citizens to assist in monitoring, to empower the patients and to support healthy living at home. The system is supported by an underlying semantic interoperability solution based on widely-adopted standards, such as HL7 FHIR, and advanced analytics. The platform will be leveraged during the SMART-BEAR Pilot of Pilots and further refined to support the planned large-scale pilots in all the participating countries.

ACKNOWLEDGEMENTS

This work was supported by the European Commission's Horizon 2020 research and innovation program under the SMART-BEAR project, grant agreement No 857172.

REFERENCES

Anisetti, M., Ardagna, C., Braghin, C., Damiani, E., Polimeno, A., & Balestrucci, A. (2021). Dynamic and Scalable Enforcement of Access Control Policies for Big Data. *Proceedings of the 13th International Conference on Management of Digital EcoSystems.*

Armbrust, M., Das, T., Sun, L., Yavuz, B., Zhu, S., Murthy, M., . . . Zaharia, M. (2020). Delta lake: high-performance ACID table storage over cloud object stores. *Proceedings of the VLDB Endowment, 13*(12), 3411-3424.

Basdekis, I., Pozdniakov, K., Prasinos, M., & Koloutsou, K. (2019). Evidence Based Public Health Policy Making: Tool Support. *2019 IEEE World Congress on Services (SERVICES)*, (pp. 272-277). Milan, Italy.

Basdekis, I., Sakkalis, V., & Stephanidis, C. (2011). Towards an accessible personal health record. *International Conference on Wireless Mobile Communication and Healthcare* (pp. 61-68). Springer, Berlin, Heidelberg.

Bellandi, V., Basdekis, I., Ceravolo, P., Cesari, M., Damiani, E., Iliadou, E., . . . Maghool, S. (2021). Engineering Continuous Monitoring of Intrinsic Capacity for Elderly People. *2021 IEEE International Conference on Digital Health (ICDH)* (pp. 166-171). IEEE.

Broekhuis, M., van Velsen, L., Peute, L., & Halim, M. (2021). Conceptualizing Usability for the eHealth Context: Content Analysis of Usability Problems of eHealth Applications. *JMIR Formative Research, 5*(7), e18198.

Dash, S., Shakyawar, S., Sharma, M., & Kaushik, S. (2019). Big data in healthcare: management, analysis and future prospects. *Journal of Big Data, 6*(1), 1-25.

European Commission. (2015). *The 2015 Ageing Report Economic and budgetary projections for the 28 EU Member States (2013-2060), Technical Report.* Brussels.

FSH School. (n.d.). *SUSHI*. Retrieved March 2022, from https://fshschool.org

HL7. (n.d.). *International Patient Summary Implementation Guide*. Retrieved March 2022, from https://hl7.org/fhir/uv/ips

HL7. (n.d.). *Personal Health Device Implementation Guide*. Retrieved March 2022, from http://hl7.org/fhir/uv/phd/2019May

Huang, Z., Wang, Y., & Wang, L. (2020). ISO/IEEE 11073 Treadmill Interoperability Framework and its Test Method: Design and Implementation. *JMIR medical informatics, 8*(12), e22000.

Integrating the Healthcare Enterprise. (n.d.). Retrieved March 2022, from Integrating the Healthcare Enterprise: https://www.ihe.net/

Liao, Q., Gruen, D., & Miller, S. (2020). Questioning the AI: informing design practices for explainable AI user experiences. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, (pp. 1-15).

MedMij. (2022, March). Retrieved from The MedMij Project: https://medmij.nl

Personal Connected Health Alliance. (n.d.). Retrieved March 2022, from Personal Connected Health Alliance: https://www.pchalliance.org

Population ageing in Europe Facts, implications and policies. (n.d.). Retrieved from https://ec.europa.eu/research/social-sciences/pdf/policy_reviews/kina26426enc.pdf

Population structure and ageing. Electronic. (2017). Retrieved from http://ec.europa.eu/eurostat/statisticsexplained/index.php/Population_structure_and_ageing

World Health Organisation. (2015). *World report on Ageing and Health.* Retrieved from http://apps.who.int/iris/bitstream/handle/10665/186463/9789240694811_eng.pdf.