CrowdHEALTH: An eHealth Big Data Driven Platform towards Public Health Policies

Argyro Mavrogiorgou[1], Athanasios Kiourtis[1], Ilias Maglogiannis[1], Dimosthenis Kyriazis[1], Antonio De Nigro[2], Vicent Blanes-Selva[3], Juan M. García-Gómez[3], Andreas Menychtas[4], Maroje Sorić[5], Gregor Jurak[5], Mitja Luštrek[6], Anton Gradišek[6], Thanos Kosmidis[7], Sokratis Nifakos[8], Konstantinos Perakis[9], Dimitrios Miltiadou[9], and Parisis Gallos[10]

[1]*Department of Digital Systems, University of Piraeus, Piraeus, Greece*
[2]*Engineering Ingegneria Informatica SpA - R&D laboratory, Roma, Italy*
[3]*Instituto de Tecnologías de la Información y Comunicaciones - BDSLab, Universitat Politècnica de València, Valencia, Spain*
[4]*BioAssist S.A., Athens, Greece*
[5]*University of Ljubljana, Ljubljana, Slovenia*
[6]*Department of Intelligent Systems, Jozef Stefan Institute, Ljubljana, Slovenia*
[7]*Care Across Ltd, London, United Kingdom*
[8]*Department Health Informatics Center, Karolinska Institutet, Stockholm, Sweden*
[9]*Singular Logic EU projects department, Kifisia, Greece*
[10]*European Federation for Medical informatics, Lausanne, Switzerland*
{margy, kiourtis, imaglo, dimos}@unipi.gr, antonio.denigro@eng.it, viblasel@upv.es, juanmig@ibime.upv.es, amenychtas@bioassist.gr, {Maroje.Soric, Gregor.Jurak}@fsp.uni-lj.si, {mitja.lustrek, anton.gradisek}@ijs.si, thanos.kosmidis@careacross.com, Sokratis.nifakos@ki.se, kperakis@ep.singularlogic.eu, dmiltiad@gmail.com, parisgallos@yahoo.com

Abstract: In today's interconnected world, more health data is available than ever before, resulting into a rich digital information environment that is characterized by the multitude of data sources providing information that has not yet reached its full potential in eHealth. CrowdHEALTH introduces a new paradigm of Health Records, the Holistic Health Records (HHRs), which offer the ability to include all this existing health data. To achieve that, CrowdHEALTH seamlessly integrates big data technologies across the complete data path, providing its results to the health ecosystem stakeholders, as well as to policy makers towards a "health in all policies" approach. This paper describes the CrowdHEALTH architecture, summarizing all the mechanisms and tools that have been developed and integrated in the context of CrowdHEALTH. The latter, along with the experimentation with several use cases that provide diverse data from different sources, have provided useful insights towards the successful and wide adaptation of the CrowdHEALTH platform in the healthcare domain.

1 INTRODUCTION

Information and data sharing across heterogeneous healthcare systems, focusing on the management of healthcare, nowadays have become the backbone of modern delivery of sustainable healthcare services and platforms (Ganguly, 2009). From routine patient care to record keeping to requisite regulatory compliance details, the healthcare industry generates enormous amounts of directionless data, which on its own does not hold any tangible value (Cohen, 2019). At the same time, due to the explosion of all the available information and communications technology (ICT) services, there are several sensors and applications supporting personalized care. However, all these services and data are heterogeneous and operate independently, resulting into the limited exploitation of their emerging added-value (The, 2019). Due to this inadequate integration of the technology, as well as the large amount of data being generated by the existing data sources, it is getting increasingly common for important events to be missed, such as the early identification of development of diseases or the creation of inefficient

policies. On top of all these, today's health records (i.e. electronic - EHRs and personal - PHRs) are far from being what the citizens consider as of value to their health. This is consistent with the beliefs of 80% of the public regarding health as more than being disease-free (Edelman, 2011) and includes a variety of everyday living aspects, such as the environment, the active and fit lifestyle, the nutrition, the mental and emotional health. Capturing this information, as well as linking it with other data in EHRs and PHRs would be of benefit for learning about outcomes of prevention strategies and health policies, diseases, and efficiency of patient pathway management. All these highlight the opportunity of exploiting all the existing amounts of healthcare data for achieving effective and targeted policy making, development of personalized medicine, and health promotion.

All these confirm the fact that nowadays there exist a plethora of independent and heterogeneous services, while health records are of limited value since their data exploitation is limited as well. This has resulted into ineffective and untargeted health policies, fragmented health strategies and inefficient personalized healthcare, highlighting the need for a holistic approach to enable public health policies and strategies and efficient medicine, health support and disease prevention. In order to address this gap, CrowdHEALTH platform (CrowdHEALTH, 2019) envisions to incorporate technologies for a paradigm shift from independent and heterogeneous services and data sources, from limited data exploitation and from health records that partially address the policy domain, to complete integrated data views through the Holistic Health Records (HHRs) (Kiourtis, 2019a). This is achieved based on the actual data exploitation emerging from collective knowledge (from HHRs clusters), and effective and targeted health policies based on a set of health analytics tools. Therefore, CrowdHEALTH explores mechanisms that can be clustered across the main areas of the holistic data services exploiting user knowledge, and the efficient policy making across domains.

The rest of this paper is organized as follows. Section 2 describes the overall CrowdHEALTH architecture capturing all its components, in combination with the interactions among them, so as to achieve the integration of all the heterogeneous existing health data towards the creation of successful public health policies and strategies. Section 3 outlines the chosen use cases for evaluating the applicability of the CrowdHEALTH platform in different eHealth scenarios, while Section 4 depicts all the users that are involved in the platform. Finally, Section 5 states the conclusions and future work.

2 CROWDHEALTH ARCHITECTURE

The CrowdHEALTH platform aims to deliver an integrated ICT platform that provides decision support to public health authorities for policy creation, co-creation, and evaluation, through the exploitation of collective knowledge that emerges from multiple information sources and its combination with situational awareness artefacts. The platform incorporates big data management mechanisms addressing the complete data path: from acquisition, and cleaning, to data integration, modelling, analysis, information extraction and interpretation. What is more, CrowdHEALTH provides various services to policy makers, enabling them to utilize causal and risk stratification mechanisms - combined with forecasting and simulation tools, in order to develop multi-modal targeted policies in terms of time scales (i.e. long- / short- term), location properties (i.e. area, regional, national, international), population segmentation (e.g. patients of a specific disease, overnight workers, etc.), and evolving risks (e.g. epidemics).

In order to offer all the aforementioned capabilities, the overall architecture of the CrowdHEALTH platform consists of three (3) main pillars: (i) the Data & Structures, (ii) the Health Analytics, and (iii) the Policies. Fig.1 illustrates the final version of the CrowdHEALTH platform that is an updated version of the architecture proposed in (kBioAssist, 2017), reflecting all the components that have been implemented in the context of the CrowdHEALTH platform.

2.1 Data & Structures

In the context of Data & Structures, the whole pillar is divided into three (3) sub-pillars: (i) Data ingestion, (ii) Data integration, (iii) Data processing (Fig. 1).

Data ingestion: The CrowdHEALTH platform is able to take as an input either live data coming from streaming data sources (i.e. unknown sources) or data at rest that are already in diverse healthcare data stores (i.e. known sources). For the known sources, since these sources are fully trustful and reliable, the nature of their data does not have to be checked. However, with regards to the unknown sources, since their nature and as a result their derived data may be anomalous either from a technical point of view (e.g. malfunctioning of a component) or from a security point of view (e.g. malicious), these sources in combination with their produced data are given as an input into the Trust & Reputation evaluation

component. The latter retrieves by the existing trust evaluation and reaction models' datastore the required trust and reputation ratings, in order to rank the input unknown data sources into the trustfulness list that finally decides whether these sources can be connected into the CrowdHEALTH platform or not. In sequel, for both the input known and unknown sources, in order for their data to be anonymized, the Data anonymization component takes as an input this data so as to completely anonymize it, achieving the required data disclosure and privacy requirements. It should be noted that there exist cases that the part of data anonymization may take place either within or outside the different organizations of the healthcare data providers, in order to enable and achieve data protection and privacy policy. Depending on the type of the data source that the anonymized data has been derived from (i.e. either unknown or known sources), the flow of the data has two (2) different options. In the first option of the unknown sources, the anonymized data is being sent into the Plug'n'play sources component. In this component, different technologies are being provided for easing the connection between the new streaming data sources and the CrowdHEALTH platform, identifying the sources' Application Programming Interfaces (APIs), and finally gathering all their data. Sequentially, all this data is being sent to the Sources reliability component that combines and evaluates (i) the reliability of the collected data, and (ii) the reliability

of the data sources that produced all the collected data, so as to estimate the reliability levels of each connected data source, and keep only the reliable data that comes from exclusively reliable sources. Thus, all this data is sent to the Gateway component so as to be transferred into the remaining flow of the architecture. In the second option of the known sources, the anonymized data is sent immediately to the Gateway component, where both connectivity and communication issues are solved at the same time, so as to gather the data from the connected sources.

Data integration: In order for the collected data to be firstly managed and transformed into an interoperable format, the Gateway sends the collected data in the form of raw data to the Data conversion component. The latter implements different functionalities in order to make it interoperable both structure and terminology wised, translating finally the collected data of the Gateway into the HHR FHIR format (Kiourtis, 2019a) that has been decided to be used for the interoperability purposes, based on a relevant research that was conducted in the past (Kiourtis, 2019b). In more details, the Data conversion component initially transforms the raw data into HHR format using the HHR model (Kiourtis, 2019a) that is being produced by the HHR manager component. In short, the HHR model has been constructed in the context of the CrowdHEALTH platform, for representing in a consistent way all the data required by the underlying
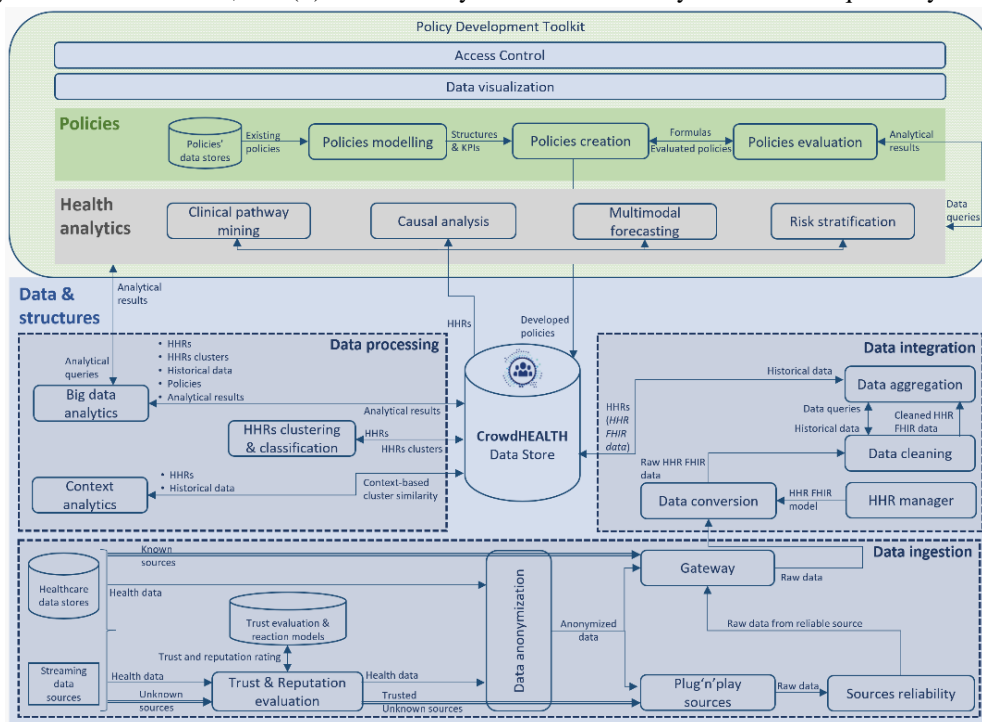

Figure 1: CrowdHEALTH architecture

data sources. It implements an XML language, specifically designed for the HHR model, which allows to specify in a machine-interpretable way the structure of HHR types and map them to the structure of the corresponding FHIR resources (HL7 FHIR, 2018). As soon as all the acquired data is transformed into HHR FHIR format, since it is important to have a certain confidence about the created information "fresh-ness" and appropriateness, the generated HHR FHIR data is sent to the Data cleaning component so as to be cleaned. To achieve that, the latter performs specific data queries to the Data aggregation component. This component, in turn, submits these queries into the CrowdHEALTH Data Store, so as to retrieve historical data and send it back to the Data cleaning component for performing data cleaning actions based on the defined cleaning rules and the patterns created by the acquired historical data. Thus, all the gathered data is fully cleaned, being sent in the form of cleaned HHR FHIR data to the Data aggregation component. In sequel, this component aggregates all the input HHR FHIR data into the corresponding HHRs, storing them finally into the CrowdHEALTH Data Store. It should be noted that the HHRs are not stored as raw HHR documents into the Data Store, but instead, they are translated into tuples that are stored in the data tables of the relational schema of the Data Store, which was designed with respect to the E-R definition of the HHR model.

Data processing: Having constructed and stored into the CrowdHEALTH Data Store all the aforementioned information, the HHRs clustering & classification component is triggered taking as an input the stored HHRs, in order to capture the correlations among the similar HHRs that are identified, and produce the corresponding HHRs clusters. These clusters are then stored into the CrowdHEALTH Data Store. On top of this, the Context analytics component retrieves both the HHRs and the historical data that is stored into the CrowdHEALTH Data Store so as to identify cluster similarities based on the health contexts obtained from this data. Again, this information is stored into the CrowdHEALTH Data Store for future usage. Thus, upon all this stored data (i.e. HHRs, HHRs clusters, historical data), the Big data analytics component performs real-time big data analytics, enabling correlations and extraction of situational factors between biosignals, physical activities, medical data patterns, clinical assessment, and laboratory exams. This component is able to process millions of events per second allowing the exploitation of (often-critical) medical data from different sources as things happen.

2.2 Health Analytics

Since the data from the underlying sources have been successfully imported, transformed, and stored into the CrowdHEALTH HHR-compatible format through the Data & Structures pillar, all this data can then be exploited by the Health Analytics pillar of the CrowdHEALTH platform. In more detail, in the context of Health Analytics, analytical techniques are utilized for carrying out Clinical pathway mining, Causal analysis, Multimodal forecasting, as well as Risk stratification upon all the gathered data. Each one of these components works independently, acquiring as an input from the CrowdHEALTH Data Store all the stored information that they need, which was originally arrived in the platform in the form of constructed HHRs. Additionally, each one of these components exploits the Big data analytics component so as to perform its queries upon the required data and retrieve the corresponding health analytical results. In short, Causal analysis allows the identification of the properties that affect the performance of policies and care plans, while Clinical pathway mining supports data analysis so as to identify similarities or differences in treatment among groups of patients, indicating major effective factors that affect several treatments and establishing a supporting framework for improving the treatment of patients with different diseases. Multimodal forecasting estimates the applicability and effectiveness of health policies, their variations and combinations to particular population segments, considering social information and spatiotemporal properties. Finally, the Risk stratification informs about population-level health risk, identifying what proportions are of low, medium and high risk.

2.3 Policies

On top of the Data & Structures and Health Analytics pillars, there exists the Policies pillar that is mainly responsible for exploiting the results by using the developed Policy Development Toolkit. The latter represents the component that integrates several sub-components to enable policy makers to create, update and validate policy models. In this context, initially the Policies modelling component is triggered, which collects as an input all the existing policy models, formulates new policy models' structures based on the policy makers' inputs, and sends the constructed structures and key performance indicators (KPIs) to the Policies creation component in order to create the corresponding policies. In sequel, as soon as the policies have been created, their evaluation takes

place through the Policies evaluation component, which takes as an input (i) the constructed formulas of the created policies from the Policies creation component, and (ii) the analytical results from the corresponding Health analytics tool (i.e. Clinical pathway mining, Causal analysis, Multimodal forecasting, Risk stratification) that it was decided to be used by the user. Based on this input, the Policies evaluation component outputs to the Policies creation component the evaluated policies so as to conclude to the final context of its created policies and store them into the CrowdHEALTH Data Store.

It should be mentioned, that all this information can be provided to different user groups in the ecosystem (e.g. healthcare providers, policy makers, healthcare professionals, nutrition experts, etc.) through the Data visualization component. In short, this component enables the interaction of all the users with the platform through analytical queries, while processing the results and visualizing them in an adaptive way. The Data visualization component is integrated into the Policy Development Toolkit in order to provide the required enhanced visualizations towards the end users. What is more, the Access control component is also integrated into the Policy Development Toolkit in order to give access to authorised members, providing them with different capabilities based on their privileges and rights. Thus, all the members are allowed to interact and exploit the corresponding capabilities of the Policy Development Toolkit and the Data visualization component.

3   CROWDHEALTH USE CASES

Based on the architecture described in Section 2, CrowdHEALTH aims to design, develop and showcase a novel data integration and health analytics framework for exploiting heterogeneous health data, which leverages the proper understanding for successfully creating and evaluating public health policies. Thus, CrowdHEALTH aggregates healthcare data aiming to track the same patients in different sources, and create a holistic overview of their health conditions, leveraging it to population-based analysis. For that purpose, in this Section six (6) different representative use case scenarios of the CrowdHEALTH platform are described, which consist of both private and public healthcare stakeholders that have a different health scope across Europe. More specifically, these use cases refer to the organizations of University Hospital of La Fe (HULAFE) (HULAFE, 2019), Karolinska Institutet (KI) (Karolinska Institutet, 2019), University of Ljubljana (ULJ) (University of Ljubljana, 2019), CareAcross (CRA) (CareAcross, 2019), BioAssist (BIO) (BioAssist, 2019), and German Research Centre for Artificial Intelligence (DFKI) (DFKI, 2019). Through these use cases, the main scope is to integrate the research and development work of the CrowdHEALTH platform, verifying its purposes, and collecting useful feedback about its developed concepts and technologies. Since the use cases aim to verify the applicability of the whole CrowdHEALTH platform, their collected data has followed all the steps and mechanisms provided by the Data & Structures and Policies pillars, whereas depending on each different use case's scope and requirements, the corresponding health analytical technique was implemented. All this information is described for each different use case in the following paragraphs.

HULAFE use case: This use case has been chosen for the identification of overweight and obese patients in the Health Department Valencia-La Fe through the implementation of CrowdHEALTH. In more detail, HULAFE offers data that is related to patients' demographic information, hospitalization episodes, emergency room episodes, hospital at home episodes, and morbidity. By implementing the whole CrowdHEALTH data and policies process upon the HULAFE collected data, as well as the clinical pathway mining and risk stratification techniques, it has become feasible to understand and characterize which data is crucial to drive effective policies in obesity and overweight fields, whilst improving accuracy of the identification of overweight and obese patients. In the same notion, it has been achieved the improvement of the management and the detection of obesity, including the systematic detection of obese and overweight people and the


Figure 2: CrowdHEALTH use case scenarios

detection of bad nutrition and activity habits to promote better habits on these citizens. What is more, it has been achieved the detection of groups of citizens with greater propensity for obesity to guide public health policies, whereas broaden the knowledge of health professionals through a catalog of physical activity resources and professionals in order to improve the prescribing of physical activity.

BIO use case: This use case has been chosen for monitoring disease progression and healthcare expenditure for improved chronic disease management of patients that have enrolled in the BioAssist platform. More specifically, BIO offers data related to biosignals relevant to patients' conditions, being acquired from pulse oximeters, blood pressure meters, glucometers, spirometers, weighing scales, and physical activity trackers. By implementing the whole CrowdHEALTH data and policies process upon the BIO collected data, as well as the risk stratification technique, CrowdHEALTH bestows added value to patient monitoring technologies, transforming these into tools that support evaluation assessment with regards to attributes of a population that are currently difficult to examine, and providing a link between public health authorities and patients. By applying the CrowdHEALTH technologies within this use case, it is achieved to enhance patients' quality of life, encourage proactive care, and offer efficient support in potentially dangerous situations. Extending this scenario by exploiting the data analysis capabilities provided by CrowdHEALTH, collected data has the potential to equip policy makers with a tool that allows them to measure the impact of relevant policies, in terms of actual results on populations health and quality of life.

CRA use case: This use case has been chosen for evaluating the impact of online coaching and medical education on cancer patient behavior that have enrolled in the CareAcross web platform. More specifically, CRA offers data related to patients' diagnosis, treatment, comorbidities, health behaviors and side-effects. By implementing the whole CrowdHEALTH data and policies process upon the CRA collected data, as well as the causal analysis technique, all this data is analyzed in order to identify potential causal relationships between specific data points. Furthermore, it enables predictions of future behaviors since a patient with specific diagnosis is less likely to report a specific side-effect. Such analyses are very important for patients, for healthcare professionals, but also for public policy makers. This is because the nature of oncology and cancer care services is mostly confined to the clinic.

On the other hand, patients have increased and prolonged support needs. This means that, while there are no specific policies established for the provision of medical information and online coaching, such an approach may be quite helpful. This is not restricted only to the benefit of individual patients, but it may also be fruitful towards the improvement of resource allocation in the healthcare system.

ULJ use case: This use case has been chosen for analysing the current state of physical fitness and weight status of children, analysing its development over time, predicting future levels of fitness and somatic development, through the implementation of CrowdHEALTH. More particularly, ULJ offers data related to cohort, physical activity, sedentariness, sleep, resting heart-rate, socio-economic status, and parental physical activity of school children. Thus, it provides data on physical fitness and physical activity to supplement the data on nutritional status of children and enable the construction of obesity risk assessment and developmental prediction models of somatic and physical fitness development. By implementing the whole CrowdHEALTH data and policies process upon the ULJ collected data, as well as the clinical pathway mining, risk stratification, and causal analysis techniques, ULJ use case obtains a basis for the implementation of policies that enable linking school and health data for early interventions monitoring and evaluation. What is more, the individual growth trends, physical fitness and nutritional development trends, adult stature prediction, adult weight prediction, adult physical fitness prediction, adult obesity-related health risks prediction for all the students are visualized, for easing the monitoring of the physical fitness, physical activity and obesity among school children.

DFKI use case: This use case has been chosen for understanding and characterizing influences of people's nutritional habits, and differences in physical activity upon their overall health and quality of life, through the implementation of CrowdHEALTH. More particularly, DFKI offers citizens' physical and activity data provided by their personal activity trackers. By implementing the whole CrowdHEALTH data and policies process upon the DFKI collected data, all this data can be clustered based upon their common nutritional and physical habits, thus finding relevant correlations among it, and among the habits of the citizens.

KI use case: This use case has been chosen for monitoring patients with chronic kidney diseases (CKD) and cardiovascular diseases through the implementation of CrowdHEALTH. In more detail, KI offers patients' demographic data, drug usage

data, and practitioners' consultation data with regards to these diseases. By implementing the whole CrowdHEALTH data and policies process upon the KI collected data, as well as the clinical pathway mining, risk stratification, and causal analysis techniques, all this data can be combined in order to determine the prevalence of CKD, and ascertain its clinical consequences in terms of comorbid complications and healthcare resource utilization, to determine healthcare- and socioeconomic-related risk factors for progression of CKD, and finally to establish the safety and effectiveness of common drugs in individuals with CKD and the connection to cardiovascular diseases.

4 CROWDHEALTH USERS

Based on the use cases described in Section 3, various users can offer their data and exploit the results of the CrowdHEALTH platform. Thus, all the information provided by the CrowdHEALTH platform can be exploited by different types of user groups that may exist in a healthcare ecosystem. These users may represent either final or intermediate users, depending on whether they have access to the final output information of the platform or they have access to all the information that is being produced, exchanged, and managed through the whole data flow of the platform. All the types of users that are getting involved into the CrowdHEALTH platform and the interactions that they have with the platform's (3) main pillars are depicted in Fig. 3, including both final and intermediate users.

With regards to the final users, these include the healthcare professionals, the healthcare providers, the policy makers, and the citizens. The most crucial and central stakeholder among them is the citizens, since the whole CrowdHEALTH architecture, both initiation and existence are based upon the medical data that is provided by them. Apart from this, the citizens may provide to the platform their experimentation results based upon the policies that were exported by the platform and given to them. Apart from the citizens, a major role in the platform is played by the policy makers. The latter are able to provide to the platform their existing policies - not only the existing health policy models that they currently aim to create, but also the health policies that other policy makers have modelled and created. As a result, they contribute to the successful completion of the policies creation process, being able to get as an output from the platform the final developed policies. Furthermore, another major role in the platform is played by healthcare providers that can offer the citizens' medical data into the platform, whilst they are offering the ability to receive the developed health policies based upon their requested policy requirements. Finally, another major role in the platform refers to the healthcare professionals that, as in the case of the healthcare providers, can enter the citizens' medical data into the platform. Moreover, they are able to view analytics outcomes that they have requested based on their analytics requirements through the visualization component of the platform and receive the developed health policies based upon their requested policy requirements.
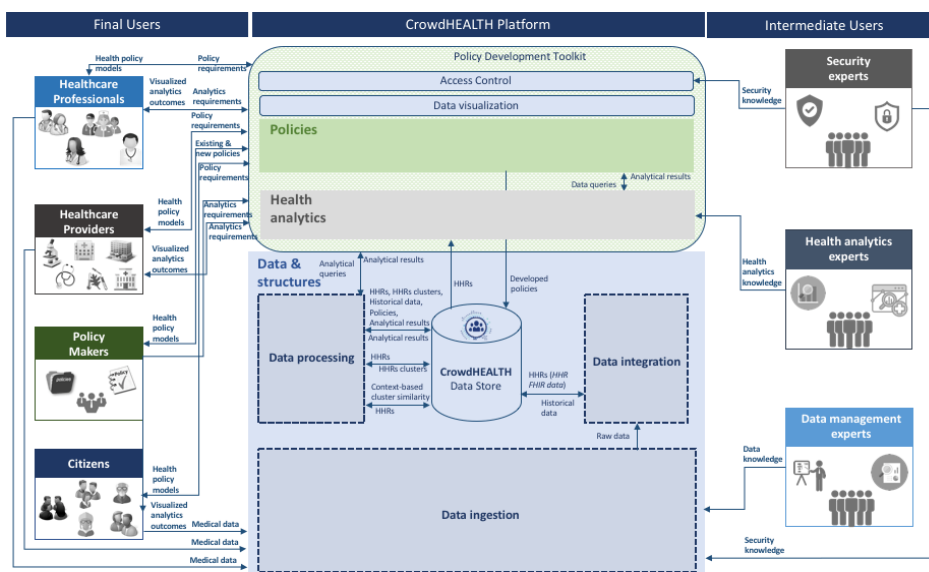

Figure 3: User groups' interaction with CrowdHEALTH architecture

Regarding the intermediate users, these include the security experts, the health analytics experts, and the data management experts. The security experts provide their expertise upon all the security aspects that are developed into the CrowdHEALTH platform of the Policies and the Data & Structures pillars, and are of crucial importance. Apart from the security experts, the health analytics experts are highly involved into the platform, providing their health analytics knowledge upon all the developed health analytical tools of the Health Analytics pillar. In the same notion, the data management experts are responsible for providing their data knowledge upon all the data management procedures that occur within the Data & Structures pillar of the platform.

5 CONCLUSIONS

Patients' data coming from multiple information sources constitutes a computable collection of fine-grained longitudinal phenotypic profiles that may facilitate cohort-wide investigations and knowledge discovery on an unprecedented scale, which is the prerequisite for patient-centered care (Chawla, 2013). To this end, in this paper a complete patient-centered eHealth platform was presented, the CrowdHEALTH platform, being able to capture all the existing health determinants in new structures, the HHRs, while creating groupings of them (i.e. clusters). As a result, it provided the ground for the discovery of deep knowledge about population segments and provision of insight for different segments and users according to various criteria (e.g. location, medication status, emerging risks, etc.), by creating and evaluating the corresponding health policies. This opens the opportunities for successfully achieving personalised medicine, and disease prevention (Chawla, 2013), (Cirillo, 2019).

The applicability of the proposed platform was evaluated through different use case scenarios in terms of collecting and processing data from real-world data sources, being heterogeneous, and having various data formats, analysis needs, information to be included in the HHRs, target groups (e.g. people suffering from chronic diseases, children and youth), and environments (e.g. care centers, social networks, public environments, and living labs). Thus, by currently exploiting the 2 million records and 700.000 streams of lifestyle activities and nutrition data, as well as engaging more than 200.000 users that come from these use case scenarios, the platform is expected to be able to exploit more than 7.5 million measurements from 1 million people across Europe.

REFERENCES

BioAssist (2019). BioAssist platform, https://bioassist.gr/?lang=en

CareAcross (2019). CareAcross platform, https://www.careacross.com/

Chawla, N. V., et al. (2013). Bringing big data to personalized healthcare: a patient-centered framework. Journal of general internal medicine, 28(3), 660-665.

Cirillo, D., et al (2019). Big data analytics for personalized medicine. Current opinion in biotechnology, 58, 161-167.

Cohen, L. (2019). More, More, More: Removing data inefficiencies will remove healthcare roadblocks, https://www.medicaleconomics.com/news/more-more-more-removing-data-inefficiencies-will-remove-healthcare-roadblocks

CrowdHEALTH (2019). CrowdHEALTH EU project, https://crowdhealth.eu/

DFKI (2019). DFKI, https://www.dfki.de/en/web/

Edelman (2011). Health Barometer 2011: Global Findings, http://healthbarometer.edelman.com/2011/10/health-barometer-2011-global-findings

Ganguly, S., et al., (2009). Sharing information and data across heterogeneous e-health systems. Telemedicine and e-Health, 15(5), 454-464.

HL7 FHIR (2018). Resource Index, https://www.hl7.org/fhir/resourcelist.html

HULAFE (2019). Health Department of Valencia-La Fe, http://www.lafe.san.gva.es/home

Karolinska Institutet (2019). Karolinska Institutet, https://ki.se/

kBioAssist, S. (2017). CrowdHEALTH: Holistic Health Records and Big Data Analytics for Health Policy Making and Personalized Health. Informatics Empowers Healthcare Transformation, 238, 19.

Kiourtis, A., et. al (2019a). Holistic Health Records towards Personalized Healthcare. IEEE International Symposium on Circuits and Systems. (under review)

Kiourtis, A., et. al (2019b). Aggregating the syntactic and semantic similarity of healthcare data towards their transformation to HL7 FHIR through ontology matching. International Journal of Medical Informatics 132 (2019): 104002.

LeanXcale (2019). Any load & scale database, https://www.leanxcale.com/

The, L. P. H. (2019). Next generation public health: towards precision and fairness. Public health, 4(5), e209.

University of Ljubljana (2019). University of Ljubljana, https://www.uni-lj.si/eng/