Towards a Model-Driven Platform for Evidence based Public Health Policy Making

Marios Prasinos, George Spanoudakis
City, University of London
London, United Kingdom
{Marios.Prasinos.1@city.ac.uk,
G.E.Spanoudakis@city.ac.uk}

Dimitrios Koutsouris
Biomedical Engineering Laboratory,
National Technical University of Athens
Athens, Greece
dkoutsou@biomed.ntua.gr

Abstract— **The effective management of various health conditions depends on and requires appropriate public health policies (PHP). Such policies are important for several aspects of healthcare provision, including: (a) screening for prevention of disease; (b) early diagnosis and treatment; (c) long-term management of chronic diseases and disabilities; and (d) setting-up standards. Although it is widely recognised that the PHP life cycle (i.e., the analysis, action plan design, execution, monitoring and evaluation of public health policies) should be evidenced based, current support for it is mainly in the form of guidelines, and is not supported by data analytics and decision making tools tailored to it. In this paper, we present a novel model driven approach to PHP life cycle management and an integrated platform for realising this life cycle. Our approach is based on PHP decision making models. Such models steer the PHP decision making process by defining the data that need to be collected and the ways in which these data should be analysed in order to produce the evidence required for PHP making. Our work is part of a new research programme on public health policy making for the management of hearing loss, called EVOTION, that is funded by the European Union.**

Keywords — public health policy; model driven data analytics; ontologies

I. INTRODUCTION

The effective management of various health conditions depends on and requires appropriate public health policies (PHP) as it has been acknowledged by the World Health Organisation [1], governmental institutions and patient associations [2]. Public health policy affects the affordability and, hence, access to health care services (e.g., health check-ups, health care device adjustments, provision of related rehabilitation services), medication and supportive devices. PHP may also have a significant effect on: (a) screening for prevention of disease; (b) early diagnosis and treatment; (c) long-term management of chronic diseases and disabilities; and (d) setting-up standards.

It is widely recognised that the PHP life cycle (i.e., the analysis, action plan design, execution, monitoring and evaluation of public health policies) should be evidenced based. This is because the management of health conditions and their consequences at a public health policy making level can benefit from the analysis of heterogeneous data, including health care device usage (if applicable), physiological, cognitive, clinical and medication, personal, behavioural, life style data, occupational and environmental data. The analysis of these types of data using big data analytics techniques can enable the investigation of whether particular health conditions have comorbidities and reveal contextual factors, social, behavioural and economic, life cycle and other factors affecting them. The outcomes of such analysis can enable the stratification of related risks and effects to the patients, and – through correlation with other economic, social and physical constraints – help developing a holistic systemic perspective of interventions regarding the management of health conditions and the broader support, social and occupational inclusion and well-being of the patients. Evidence arising from big data analytics (BDA) can also help exploring missing, under or over-estimated value of specific medical interventions and analysing their effectiveness (i.e., understanding the trade-offs between their cost and benefits). Despite such potential benefits, current support for PHP making is mainly in the form of guidelines, and is not supported by data analytics and decision making tools based on such analytics.

In this paper, we present a novel model driven approach to PHP life cycle management and an integrated platform for realising this life cycle. Our approach is based on PHP decision making (PHPDM) models. Such models steer the PHP making process by defining the data that need to be collected and the ways in which they should be analysed in order to produce the evidence required for PHP making. Our work is part of a new research programme on public health policy making for the management of hearing loss, called EVOTION, that is funded by the European Union. On-going reforms of PHP in this area related to the management of different types of hearing loss and the spark of social debate that they have caused demonstrate the importance of PHP making in this area (see, for example: https://www.actiononhearingloss.org.uk/get-involved/campaign.aspx).

The rest of this paper is structured as follows. Section II describes related work. Section III provides an overview of our approach and the architecture of platform that we are developing to realise it. Section IV presents the ontology based scheme for specifying PHPDM models. Section V presents an example PHPDM model and how it can be executed to realise our approach. Finally, Section VI presents concluding remarks and directions for future work.

II. Related Work

A recent review of big data applications in biomedical research and health care [3] has pointed out several applications, including, for example: real-time risk monitoring, using online social media combined with epidemiological information as a data source for facilitating public health surveillance [4] and ambulatory cardiovascular care [5].

The opportunities, which arise from using big data analytics to reduce the costs of health care [6] are also unprecedented. Examples of healthcare services that can benefit in this respect include high-cost patients, readmissions, triage, decompensation (when a patient's condition worsens), adverse events, and treatment optimisation for diseases affecting multiple organs [6]. The use of big data analytics in healthcare may also enable the generation and dissemination of new knowledge, the translation of personalised medicine initiatives into clinical practice and the transformation of healthcare by delivering information directly to patients, empowering them to play a more active role [7].

According to a systematic review of studies of decision-making by health care managers and policy-makers, researchers could better inform health care management and policy-making by making several changes to how they produce and update systematic reviews and by adapting existing reviews that are relevant to local health care issues [8]. Also, research in this area becomes important as policy-makers and managers increasingly require access to high-quality evidence syntheses that include research and non-research based evidence, in the form of qualitative and/or quantitative research findings [9]. `Recent applications of big data research in health policy making include breast screening decision making [10] and public health strategies for alcohol harm reduction in the UK [11].

A key to success in learning from big health care data is to remain focused on gaining actionable insights into the best ways to treat the patients in the care system that generated the data [12].

III. Approach

To realise the approach that we outlined in Sect. I, we are developing a platform (referred to as the *EVOTION platform* in the following) whose overall architecture is shown in Figure 1. This platform uses different types of data to inform the PHPDM (PHPDM) process. Such data include: (a) retrospective and prospective patient data, including medical, clinical and medication, personal and occupational data; (b) prospective real time patient data including medical devices usage, cognitive, behavioural and life style, and environment data (e.g., location of patient, noise environment); and (c) dynamic web and social media data (e.g., feedback on proposed or implemented policies that may be useful in predicting or evaluating the perceptions of the public about PHPs).

The operation of the EVOTION platform is driven by PHPDM models. These models specify:

(i) The generic issues that need to be addressed by PHPs and the alternative decisions that may be made to address them.

(ii) The evidence that can support or provide counter indicators for decisions. Evidence may be related to a wide spectrum of factors. Considering the use of medical devices, for example, evidence can be collected to explore whether the difficulties faced by different types of medical device users depend on their condition, their cognitive capabilities, their life style and behaviour, other comorbidities that they may have and/or their overall compliance with medical device usage guidelines given to them by clinicians. Evidence may also be required to explore whether such difficulties can be alleviated by the number of follow up treatments, the time periods between such treatments etc.

(iii) The big data analytic (BDA) processes (e.g., the specific types of statistical analysis or data mining analysis) that should be followed for collecting and analysing the evidence.

(iv) The criteria that should be used to determine if the available evidence is sufficient for making decisions. Criteria determine the extent of the evidence that would be deemed sufficient for supporting a decision, and thresholds that would make the evidence conclusive. Criteria may, for example, determine the combination of the factors referred to in (ii) above that would be a good predictor of the difficulties faced by medical device users.

(v) The processes to be followed for making specific types of health policies. Such processes may, for example, determine who are the stakeholders whose views should be considered and recorded prior to reaching a decision, who has responsibility for making the final decision, and whether a decision should be continually or periodically reviewed upon the acquisition of new evidence.
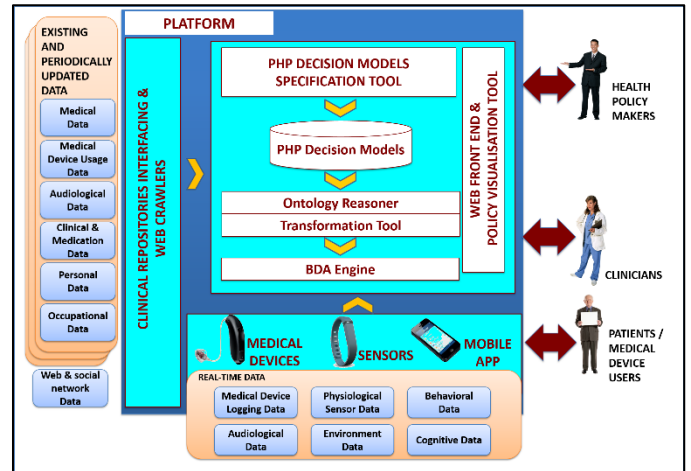
Figure 1 The Overall Evidence Based PHP Decision Making Platform

Our view is that PHPDM models are essential for realising evidence based, scalable, fully dynamic, repeatable and accountable policy making. This is because PHPDM models covering the aspects identified above could:

- be automatically transformed into executable BDA processes whose execution would provide the basic evidence required for making a decision and exploring its consequences;
- drive the collaborative stakeholder decision making processes; and

- provide a structure for organising the alternatives, arguments and rationale for making decisions in a way that makes them traceable and accountable.

PHPDM models can be: (a) repeatedly executed in the same or different policy making settings (e.g., for making policy on the very same issues in different regions); and (b) specified parametrically to make their customisation easy in case that this would be required in different policy making settings.

As shown in Figure 1, the EVOTION platform incorporates a tool supporting the specification of PHPDM models into some high-level language, and their verification and transformation into executable BDA processes that would be passed as input to the BDA engine to drive its execution and generate the evidence required for PHPDM. Also, to enable the data collection processes and policy-making process, the EVOTION platform incorporates and integrates:

(a) Existing repositories of medical data.
(b) Enhanced medical devices enabling the capture and provision of medical device usage related data (e.g., rating of device ease or difficulty of use in different listening conditions, frequency and type adjustments of controls).
(c) Sensors supporting the collection real time contextual patient physiological data (e.g., heart rate, blood pressure, skin conductance)
(d) A mobile application with components supporting the acquisition and transmission of behavioural (e.g., recording of patient daily activities such as participation in conversations, watching TV), contextual (e.g., patient's location), cognitive (e.g., verbal reaction time) data as well as the notification and acceptance/rejection of decisions by the patient and/or their carers (decision selection component); and the execution n of periodic audiological and cognitive to collect the related data (audiological and cognitive test components).

IV. SPECIFICATION OF PHPDM MODELS

The specification of PHPDM models is ontology based. More specifically, PHPDM models are specified as instances of an ontology on an ontology that we have introduced for this purpose. The advocation of an ontology based approach for the specification of PHPDM models has been due to the ability of ontologies to provide an axiomatic foundation of the specification language and the realisation of corresponding reasoning processes that can aid decision making. The specific ontology that we have introduced for specifying PHPDM models has been based on three sub-ontologies, namely: (i) an ontology for governmental policy making (i.e., the G2G Ontology [13]); (ii) an ontology for data mining (i.e., the OntoDM-core ontology [14]); and (iii) an ontology for statistics (i.e., the STATO ontology [15]).

A view of the main classes of the ontology for specifying PHPDM models is presented in Figure 2. According to this ontology, public health policies are introduced to address or explore different ISSUEs. Examples of issues that can be the subject of a public health policy can be the introduction of a screening programme, the provision of certain types of medical devices to different types of patients or the introduction of a medical training programme for clinical professionals. In the policy ex-

ample that we refer to in Section V, the issue of policy exploration is the provision of binaural or single HA to hearing loss patients with different characteristics.
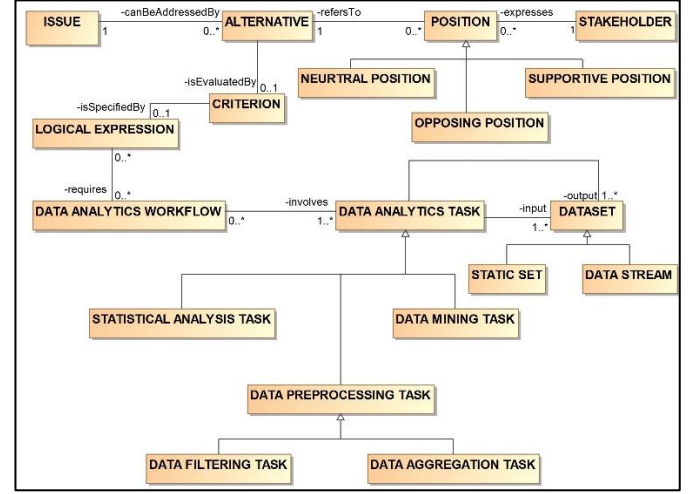

Figure 2 Main Classes of the Ontology Framework

Issues can be addressed by ALTERNATIVE decisions. An alternative (decision) presents a possible way of addressing the issue that is the subject of the policy. Alternatives reflect the key decisions that may be made in PHPDM process. They also need to be explored on the basis of evidence arising from the analysis of data. To express this, the ontology associates alternatives with a CRITERION that determines the circumstances under which the evidence arising from data analytics would support the alternative. A CRITERION is specified by a LOGICAL EXPRESSION over the outcomes of DATA ANALYTICS WORKFLOWS. Each DATA ANALYTICS WORKFLOW involves one or more DATA ANALYTICS TASKs which determine the way in which the available data are to be analysed in order to produce the evidence required by the criterion of the alternative.

A DATA ANALYTICS TASK takes as input one or many DATASETs and has output one or more other DATASETs. A DATASET can be STATIC data SET (i.e., a set that does not change frequently) or a DATA STREAM (i.e., a data set that changes continually). With regards to the processing that it performs upon its input data set(s), data analytics tasks can be distinguished into STATISTICAL ANALYSIS TASKs (i.e., tasks that carry out some statistical analysis upon the data), DATA MINING TASKs (i.e., tasks that carry out some data mining analysis upon the data), or a DATA PREPROCESSING TASK. The latter tasks are further distinguished into DATA FILTERING TASKs and DATA AGGREGATION TASK, i.e., tasks that filter data or aggregate data, respectively. Statistical analysis tasks are modelled using the STATO ontology and data mining tasks are specified using the Onto-DM ontology.

Initially, a PHPDM model includes a specification of tasks and the input data sets that they will be applied to. When a PHPDM model is executed, the description of it is expanded by including the data sets, which are produced by the data analytics workflow, as we explain in Section V.

The alternatives specified in a PHPDM model are related to POSITIONs taken by the different participants of the process (i.e., the STAKEHOLDERs). A position expressed by a stakeholder can be a SUPPORTIVE position (i.e., a position that supports the advocation of the alternative), an OPPOSING position (i.e., a position that is negative to the advocation of the alternative) or a NEUTRAL position (i.e., a decision indicating that the stakeholder neither supports nor objects to the alternative). A stakeholder may express SUPPORTIVE, OPPOSING or NEUTRAL positions for one or more alternatives but cannot express two different POSITIONS for the same alternative.

As discussed above, we use the OntoDM-core ontology [14], to specify in a formal manner the key data mining processes that should be applied to data in order to produce evidence. For this purpose, we have extended the OntoDM-core ontology by introducing classes to representing the different data mining algorithms of WEKA [16]. Each algorithm is represented by a distinct class, which has the appropriate properties for the configuration of the WEKA execution required for the algorithm.
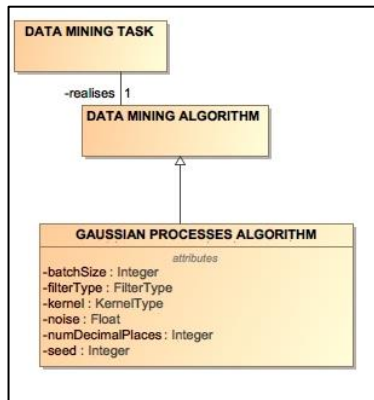

Figure 3 Example of Data Mining Algorithm Specification

An example of such classes, which has been introduced to represent the Gaussian Processes algorithm is shown in Figure 3. This class represents the properties that may be used to determine how the Gaussian Processes algorithm is to be applied, namely:

(i) the *batchSize* – i.e., the preferred number of instances to process if batch prediction is being performed

(ii) the *filterType* – i.e., whether the data should be transformed (the available options/values for this are: Normalize training data, Standardize training data, No normalization/standardization).

(iii) the *kernel* – i.e., the similarity function over pairs of data points in raw representation to use (available options in WEKA are: NormalizedPolyKernel, PolyKernel, PrecomputedKernelMatrixKernel, Puk, PBFKernel, StringKernel)

(iv) the *noise* – the level of Gaussian Noise (added to the diagonal of the Covariance Matrix, after the target has been normalized/standardized/left unchanged)

(v) the *numDecimalPlaces* – i.e., the number of decimal places to be used for the output of numbers in the model

(vi) the *seed* – i.e., the random number seed to be used

For the development of our ontology framework we merged the above three sub-ontologies using Protégé [17], based on the ontology merging approach described [18].

V. EXAMPLE AND IMPLEMENTATION

In this section, we present an example of specifying a PHPDM model using the ontology introduced in Section IV to realise a BDA process and explore alternative PHP issues.

Our example is related to policy regarding the provision of HAs and has been inspired by a policy statement of the Hearing Loss Association of America (HLAA) regarding the provision of binaural or monaural HAs[19] (HLAA recommends the provision of binaural HAs). Such a PHP decision could be based on evidence about the potential benefits from the provision of binaural or monaural HAs and, in particular, whether one of these two alternatives would affect the average daily usage of HAs by different types of patients with hearing loss. To enable the assessment of these alternatives, a BDA process can be set up to establish whether the average daily usage of HAs depends on characteristics of users including their gender, age and type of hearing loss. Exploring this issue could be based on a process, which as shown in Figure 4, involves the following steps:

(i) The specification of PHPDM model as an instance of the ontology introduced in Section IV to identify: the (policy) issue, the alternatives for addressing it, the data and the data analytics tasks that will be used to produce evidence from these data in order to explore each alternative, and the criteria for selecting amongst the different alternatives.

(ii) The execution of the data analytics tasks of the model specified in (1) and the recording of its outcomes as instances of the PHPDM model.

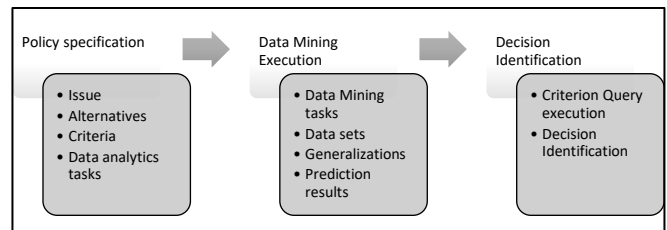(iii) The querying the PHPDM model to identify alternatives with satisfied criteria

Figure 4 Implementation Overview

In the PHPDM model used in the above process:

- The Issue is Binaural Fitting.
- The possible Alternatives are: (ALT 1) binaural HAs to become the normal fitting practice, (ALT 2) no change in policies (i.e., use of monaural HA), and (ALT 3) other means of support to be considered (in case that none of alternatives (ALT 1) and (ALT 2) are clearly supported by the evidence generated by the BDA process.
- The Criteria are:
  ♦ *Criterion 1*: Predicted average daily usage of single HA users to be less than predicted average daily usage of two HA users [(HA1+HA2)/2] with a minimum difference of 20% AND for binaural fitting users with issue delta: predicted average daily usage of single use period less than predicted average daily usage of two HAs use period AND NOT predicted average daily usage for all cases less than 2 hours (concerning ALT 1)
  ♦ *Criterion 2*: Predicted average daily usage of single HA users to be greater than 120% of predicted average

daily usage of two HA users [(HA1+HA2)/2] AND for binaural fitting users with issue delta: predicted average daily usage of single use period greater than 120% of predicted average daily usage of two HA use period AND NOT predicted average daily usage for all cases less than 2 hours (concerning ALT 2)

♦ *Criterion 3*: Predicted average daily usage for all cases less than 2 hours (concerning ALT 3)

Furthermore, there are two data analytics tasks in the model: (1) to predict and compare the average daily usage of single HA users with a single HA vs those with two HAs (DAT_1), and (2) to predict and compare the average daily usage of binaural fitting users from the period of using one HA vs the period of using both HAs (DAT_2). Table I shows how we express the instances of the above-mentioned data analytics tasks in OWL.

TABLE I.　INSTANCES IN OWL

```
Individual: DAT_1
        Types:
            data_analytics_task
Individual: DAT_2
        Types:
            data_analytics_task
```

Also, for the preparation of the mandatory data mining tasks the PHPDM model defines the following pre-processing tasks:

- a task to prepare dataset of single HA users
- a task to prepare dataset of two HA users
- a task to prepare dataset of the single period of usage of two HA users
- a task to prepare dataset of the two hearing-aid period of usage of two HA users

Additionally, we need a test set to get the predictions results from the execution of the generalisations. To achieve this, the PHPDM model includes also a task for preparing the test dataset with all the available combinations of the attributes.

The data mining tasks for the first data analytics task are: (i) to predict the average daily usage of single HA users and (ii) to predict the average daily usage of two HA users and for the second one are: i) to predict the average daily usage of the single period of usage of two HA users and ii) to predict the average daily usage of the two-HA period of usage of two HA users. For each of the above data mining tasks we also insert an instance of DM-dataset which contains the data produced by the pre-processing tasks described in the previous section.

We used the following regression algorithms of WEKA: (a) Gaussian Processes, (b) Linear Regression, (c) Multilayer Perception, (d) Support Vector Machine for Regression, (e) K-nearest neighbours, (f) K*, (g) Locally weighted learning, (h)Additive Regression, (i) Bagging, (j) Cross-Validation Parameter Selection, (k) Multischeme, (l) Random Committee, (m) Randomisable Filtered Classifier, (n) Random Subspace, (o) Regression by Discretisation, (p) Stacking, (q) Vote, (r) Weighted Instances Handler Wrapper, (s) Input Mapped Classifier, (t) Decision Table, (u) M5Rules, (v) Zero R, (w) Decision Stump, (x) M5P, (y) Random Forest, (z) Random Tree, and (aa) Rep Tree.

For each one of the 27 different algorithms, we added four different instances of data mining algorithm execution: one for each data mining task-dataset.

For each of the 108 (i.e., 4x27) algorithm executions, we added an instance of *Generalisation*, an object property of type *is_specified_output_of* to associate the instances with the data mining execution instances, and a data property with the path, where each model is stored. Along with the algorithm execution, we also performed 10-fold cross validation. For each instance of the Generalisation class, we also inserted a data property of type *has_output_details* which contains the textual output details produced by WEKA and an instance of the root mean squared error class, to store its value separately. To define which generalisation should be selected for execution for each data mining task, we need to query the ontology to get the algorithm executions with the minimum root mean squared error.

TABLE II.　SPARQL QUERY

```
SELECT DISTINCT ?generalization ?dm_task WHERE {
?task rdf:type evotion:data_analytics_task.
?isabout rdfs:label "is about"@en.
?dm_task ?isabout ?task.
?algorithm_execution rdf:type  ontodm-core:OntoDM_000033.
?is_concretization_of rdfs:label "is_concretiza-
tion_of"@en.
?algorithm_execution ?is_concretization_of ?algorithm.
?is_specified_output_of rdfs:label "is_specified_out-
put_of"@en.
?generalization ?is_specified_output_of ?algorithm_execu-
tion.
?error ?is_about ?generalization.
?error ?has_value ?min_error_value{
    SELECT  (MIN( DISTINCT ?error_value) as ?min_er-
    ror_value) ?error WHERE {
    ?algorithm_execution rdf:type  ontodm-core:On-
    toDM_000033.
    ?is_concretization_of rdfs:label "is_concretiza-
    tion_of"@en.
    ?algorithm_execution ?is_concretization_of ?algo-
    rithm.
    ?is_specified_output_of rdfs:label "is_specified_out-
    put_of"@en.
    ?generalization ?is_specified_output_of ?algo-
    rithm_execution.
    ?is_about rdfs:label "is about"@en.
    ?error ?is_about ?generalization.
    ?has_value rdfs:label "has value"@en.
    ?error ?has_value ?error_value
    } GROUP BY ?error
}}
```

Table II shows the query (in SPARQL code) which returns, for each data mining task, the produced generalisation with the minimum root mean squared error. After executing the above query, we select which generalisations should be executed. The selected generalisations are executed for: (i) Single HA users, (ii) Two HA users, (iii) One HA period of Binaural fitting users, and (iv) Two HA period of Binaural fitting users. From the generalisation executions, the following result datasets are produced:

(i) result dataset of generalisation execution (Single HA users)
(ii) result dataset of generalisation execution (Two HA users)
(iii) result dataset of generalisation execution (One HA period of binaural fitting users)
(iv) result dataset of generalisation execution (Two HA period of binaural fitting users)

After executing the generalisations, we finally have the data needed to figure out whether the criteria defined are fulfilled.

For each *Criterion*, we proceed to a query execution and insert an instance of the Value class, with an object property containing a Boolean ("true" if the criterion is fulfilled, "false" otherwise). Finally, we proceed to the identification of the proposed policy decisions by extracting the alternative with Value: "true".

The data mining algorithm execution and the generalisation execution process is summarised in Figure 5.
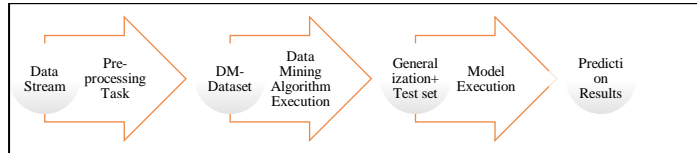

Figure 5 Implementation Tasks Execution

To test the above model, we used data from users of monaural HA and compared them to data from users with binaural HAs. The data that we used were synthetic. More specifically, we generated data for 1000 monaural HA users and 1000 binaural HA users. Also, we generated data from 1000 HA users, who initially had one HA and after some time got a second HA. In generating data, we assumed that the average daily usage data for each patient was from 3 months of the monaural aid use period and 3 months from the binaural HAs period. The generation of the synthetic data set did not consider the HA user acclimatisation period, as we considered data coming 8 weeks after issuing the HA/second HA. For all the constructed data, we generated the gender, the age, as well as the type of hearing loss randomly. To produce better data mining results and to validate our proof of concept we also made the following assumptions when generating the test data:

We generated the average daily usage according to the following conditions:

- if the type of hearing is low, the average daily usage is a random number integer between 0 and 19199 seconds,
- if the type of hearing is moderate, the average daily usage is a random number integer between 19200 and 38399 seconds, and
- if the type of hearing is severe, the average daily usage is a random number integer between 38400 and 57600 seconds.

Additionally, to manipulate the produced decision of our proof of concept, we assumed that for the single aid users the random numbers generated as described above were reduced by 30%. The same assumption was made also for the data produced for the single use period of the binaural fitting users.

Finally, we generated the test dataset, which includes all the combinations of the independent variables, to perform the generalisation executions and get the prediction results.

VI. CONCLUSIONS

In this paper, we have introduced an approach and a platform supporting public health policy decision making based on evidence produced by big data analytics processes. This approach is based on specifying public health policy decision making as instances of an ontology introduced for this purpose and executing them to support the public health policy decision making

process. To demonstrate the use of our approach, we have used an example related to policy regarding provision of monaural or binaural HAs.

Our approach and the platform supporting it is under development. In particular, we are working on developing the scheme for transforming health policy decision making models into BDA processes onto specific BDA platforms (e.g. SPARK) and an experimental evaluation of our approach using real, non-synthetic data. We are also working on the full axiomatisation of the ontology for the specification of PHPDM models.

REFERENCES

[1] World Health Organization, "Key Policy Issues in Long-Term Care," World Heal. Organ., pp. 139–190, 2003.

[2] M. I. Wallhagen, "Access to care for hearing loss: Policies and stakeholders," J. Gerontol. Nurs., vol. 40, no. 3, pp. 15–19, 2014.

[3] J. Luo , et al., "Big Data Application in Biomedical Research and Health Care: A Literature Review.," Biomed. Inform. Insights, vol. 8, pp. 1–10, 2016.

[4] S. I. Hay, et al., "Big Data Opportunities for Global Infectious Disease Surveillance," PLoS Med., vol. 10, no. 4, 2013.

[5] J. V Tu, et.al., "The Cardiovascular Health in Ambulatory Care Research Team (CANHEART): using big data to measure and improve cardiovascular health and healthcare services," Circ Cardiovasc Qual Outcomes, vol. 8, no. 2, pp. 204–212, 2015.

[6] D. W. Bates, et al., "Big data in health care: Using analytics to identify and manage high-risk and high-cost patients," Health Aff., vol. 33, no. 7, pp. 1123–1131, 2014.

[7] T. B. Murdoch and A. S. Detsky, "The inevitable application of big data to health care," Jama, vol. 309, no. 13, pp. 1351–1352, 2013.

[8] J. Lavis, et al., "Towards systematic reviews that inform health care management and policy-making," J. Health Serv. Res. Policy, vol. 10, no. July, pp. 35–48, 2005.

[9] N. Mays, et al., "Systematically reviewing qualitative and quantitative evidence to inform management and policy-making in the health field," J. Health Serv. Res. Policy, vol. 10, no. suppl 1, pp. 6–20, 2005.

[10] L. Parker, "Including values in evidence-based policy making for breast screening: an empirically grounded tool to assist expert decision makers," Health Policy (New. York)., 2017, in press.

[11] A. Brennan, et al., "Developing policy analytics for public health strategy and decisions-the Sheffield alcohol policy model framework," Ann. Oper. Res., vol. 236, no. 1, pp. 149–176, 2016.

[12] S. Schneeweiss, "Learning from Big Health Care Data," N. Engl. J. Med., vol. 370, no. 23, pp. 2161–2163, 2014.

[13] E. N. Loukis, "An ontology for G2G collaboration in public policy making, implementation and evaluation," Artif. Intell. Law, vol. 15, no. 1, pp. 19–48, 2007.

[14] P. Panov, et al., Ontology of core data mining entities, vol. 28, no. 5–6. 2014.

[15] "STATO Ontology." [Online]. Available: http://stato-ontology.org.

[16] M. Hall, et al., "The WEKA data mining software," SIGKDD Explor. Newsl., vol. 11, no. 1, p. 10, 2009.

[17] N. F. Noy, et al., "Protégé-2000: An Open-Source Ontology-Development and Knowledge-Acquisition Environment," p. 2003, 2003.

[18] P. Lambrix, et al., "Evaluation of ontology development tools for bioinformatics," Bioinformatics, vol. 19, no. 12, pp. 1564–1571, 2003.

[19] HLAA, "Dispensing Hearing Aids." [Online]. Available: http://www.hearingloss.org/sites/default/files/docs/HLAA_POLICYSTA TEMENT_Dispensing_Hearing_Aids.pdf.