IEEE Access
Multidisciplinary : Rapid Review : Open Access Journal

Received November 19, 2019, accepted December 18, 2019, date of publication January 8, 2020, date of current version January 16, 2020.

Digital Object Identifier 10.1109/ACCESS.2020.2964802

Health Intervention Evaluation Using Semantic Explainability and Causal Reasoning

JON HAËL BRENAS[1,2] AND ARASH SHABAN-NEJAD[1]

[1]The University of Tennessee Health Science Center – Oak-Ridge National Laboratory (UTHSC-ORNL) Center for Biomedical Informatics, Department of Pediatrics, College of Medicine, Memphis, TN 38103 USA

[2]Big Data Institute - Nuffield Department of Medicine, University of Oxford, Oxford OX3 7LF, U.K.

Corresponding author: Arash Shaban-Nejad (ashabann@uthsc.edu)

This work was supported by the University of Tennessee Health Science Center.

ABSTRACT As serious public health problems require complex responses, health interventions often involve multiple components implemented by groups including policy experts, social workers, and health practitioners. The success or failure of an intervention depends on many different factors, ranging from available resources to characteristics of the targeted public health issue and community to the complex mechanics relating cause and effects of the actions performed. In this paper, we present a novel formal methodology to evaluate public health interventions, policies, and programs. Our method uses the theory of change (TOC) approach along with logic models that define the intervention under consideration to generate a causal diagram and an ontology-based inference model for causal description. The resulting causal diagram will then be compared to existing knowledge and data to determine whether the intervention is coherent, internally consistent and its goals are achievable in the allotted time with the resources provided. The contextual knowledge and semantics provided by the ontology will generate a more explainable, understandable, and trustworthy approach to compare and assess different interventions based on their shared goals. Depending upon the quality and quantity of data available we perform a mix of qualitative and quantitative evaluation of the interventions. This study uses smoking cessation interventions to showcase the proposed methodology in action.

INDEX TERMS Causal graphs, intervention evaluation, logic models, ontologies, explainable AI, public health program evaluation.

I. INTRODUCTION

The Center for Disease Control and Prevention (CDC) identifies three main reasons why evaluating public health programs and interventions is crucial: accountability, identifying successful methods, and improving future programs [1]. Another key benefit of intervention evaluation is the ability to decide whether an intervention is effective, or finding the most effective intervention among several programs/proposals. This type of evaluation needs to be performed before interventions take place whereas most of the existing evaluation methods are retrospective and, thus, not suited to such an endeavor.

Health intervention evaluation may be ‘‘Formative evaluation’’, to assess the feasibility, appropriateness of a program, ‘‘Process/implementation evaluation’’, to control if an intervention’s activity has been developed as planned, ‘‘Outcome/effectiveness evaluation’’, to measure program effects in the target population, and ‘‘Impact evaluation’’, to evaluate intervention effectiveness in attaining its goals [2]. The existing literature [3] identifies the shortage of clear evaluation and automatic procedure to guide and assist ‘‘researchers, grant/journal referees and reviewers in the design, conduct or assessment of process evaluation, which means that process evaluation may be planned in an ad hoc fashion’’.

Public health organizations and health care providers often make decisions on the effectiveness of interventions and policies on the basis of their functionalities, time, budget, and human resource availability. Most of the times complex intelligent decision support systems used during the evaluation process lack proper explainability. Explainable Artificial Intelligence (XAI) [4] tends to overcome the limitations of traditional AI systems to explain their predictions, decisions,

The associate editor coordinating the review of this manuscript and approving it for publication was Xiao-Sheng Si.

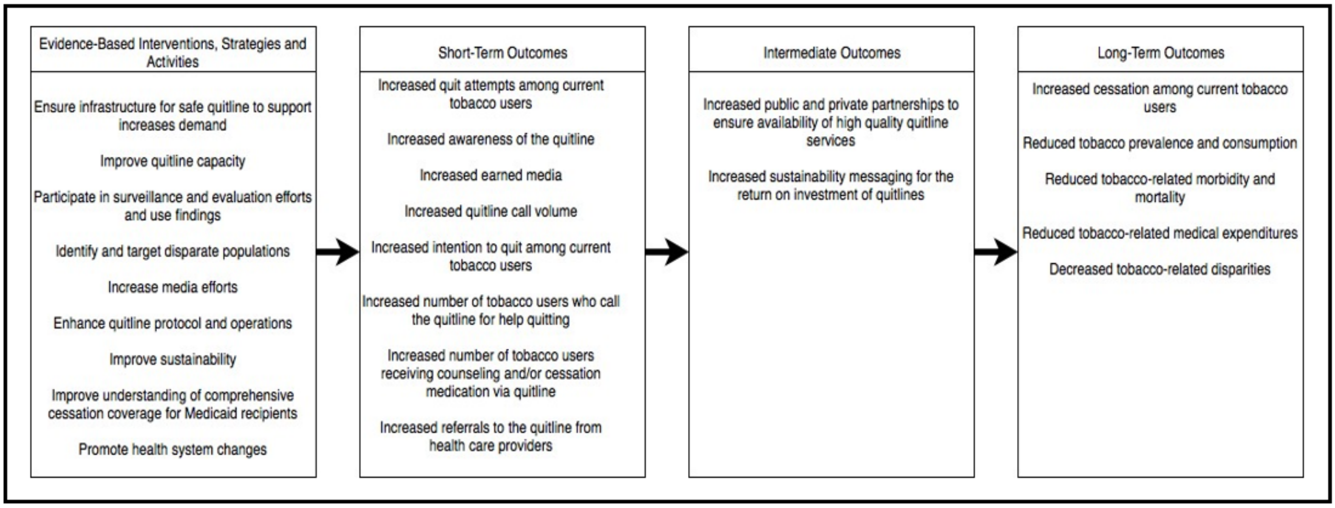

FIGURE 1. An example of a logic model for smoking cessation, adapted from the U.S. Center for Disease Control and Prevention [21].

and actions to the users. Through a constant learning process that increases the prediction precision, XAI aims to generate more explainable, understandable, and trustworthy models. While this might not yield the best of all possible prediction and evaluation, it often improves itself through multiple iterative learning processes. One important aspect of the learning process is learning the semantic associations between different concepts, determinants, and indicators using formal ontologies in the field. Ontologies capture the users' conceptualization and knowledge of a specific domain. They have been successfully used in a wide range of clinical [5], [6], public health [7], [8], and global health [9], [10] applications.

We designed the POLicy EVAlUation & Logical Testing (POLE.VAULT) Framework [11], [12] that aims to facilitate the ontology-based evaluation of public health interventions and policies. It uses Root Cause Analysis (RCA) [13], a technique for detecting and describing the causality path for a problem and recommending remedial actions. Using contextual knowledge captured in ontologies, POLE.VAULT analyzes public health programs and policies and discusses how the context influences the outcome. Ontologies provide reusable modules based on consensus knowledge, which enables knowledge-based interventions to "be responsive to the local context and potentially more effective while still allowing meaningful evaluation in controlled designs" [14]. Here, context covers "any elements which are external to the intervention, but which may impede or strengthen the effects of an intervention" [3]. Building up on top of the POLE.VAULT framework, in this paper, we propose a method to evaluate public health interventions prospectively. The proposed method can be interesting for public health researchers and decision-makers who need to choose between several possible interventions and to prove that the data support the theory behind a selected intervention. Through the mediation of ontologies, many different kinds of data both qualitative (e.g., access to public parks) as well as quantitative (e.g., average income) can be used. Similarly, data coming from different domains (e.g., public health, economy, law enforcement, etc.) can be aggregated to produce a wider and more accurate picture of the situation (including the health status) of a given community.

A complete evaluation also requires a quantitative analysis of the resources available and the existing stressors to help direct the choice. This paper focuses on a more qualitative approach that can be expanded to include quantitative data later on. As a result, the algorithms and methods that we present do not yet form a complete solution to the problem.

To illustrate our ideas, we will use the example of smoking cessation programs. Smoking has been identified as a serious health issue for a relatively long time, and many different interventions, from efforts involving communities [15] to approaches involving law and tax changes [16], have been designed to help individuals quit smoking. In Section II, we will introduce the concepts of the theory of change [17], logic models [18], causal diagrams [19] and ontologies [20] as well as the characteristics of the targeted population. Section III explains how to use the proposed method to evaluate the interventions. Section IV looks at two methods to obtain the causal diagrams that are needed for the evaluation. Section V shows the proposed method in action through an example scenario. Finally, Section VI discusses prospects for future work and conclude.

II. THEORY OF CHANGE, LOGIC MODELS, CAUSAL DIAGRAMS AND ONTOLOGIES

In order to represent the interventions that we intend to evaluate, we use logic models [18]. Figure 1 and Figure 2 show some examples of logic models used in the context of smoking cessation in the United States [21] and Scotland [22], respectively. A logic model gives a visual representation of an intervention. Most logic models, in particular in public health, follow the same archetype where inputs (i.e., resources) are

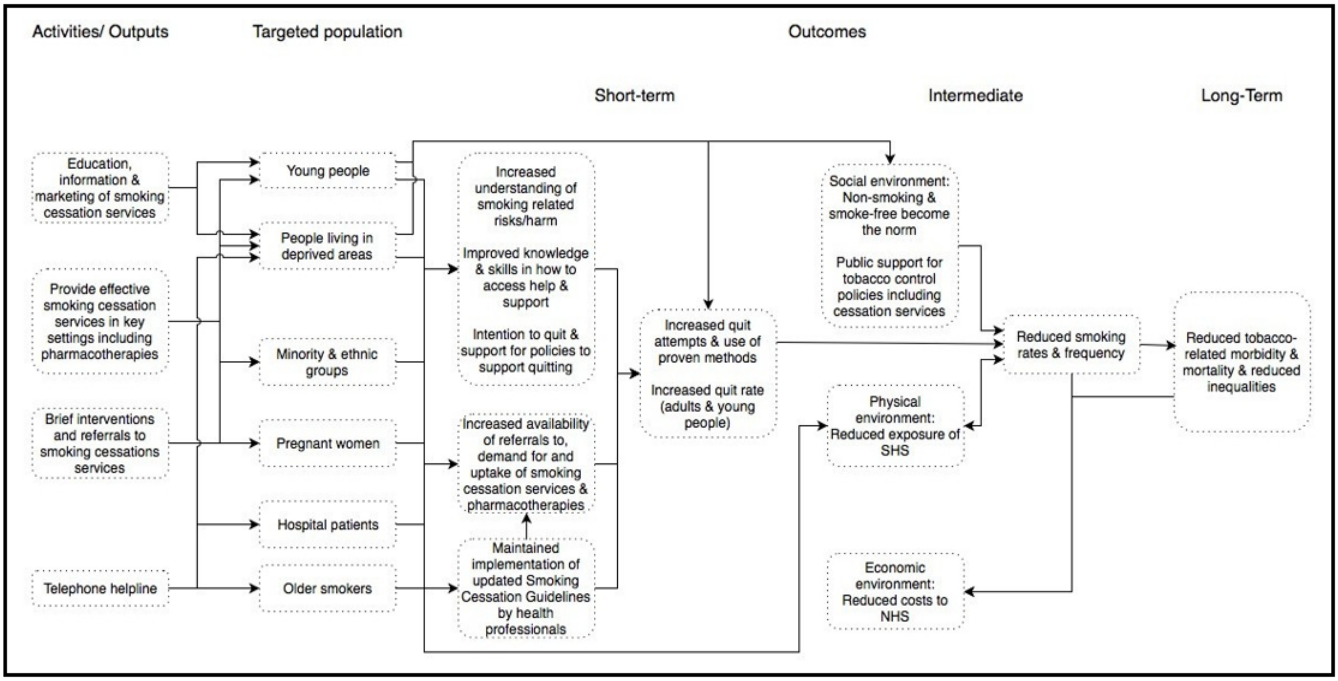

FIGURE 2. An example of a logic model for smoking cessation, adapted from the Scottish National Health Service [22].

connected to activities (i.e., the actions) that will take place during the intervention. The activities are connected to outputs (i.e., immediate or short-term expected consequences of the activities), which are then linked to outcomes, impacts (i.e. the more long-term effect of the intervention). As there is no universal standard, different organizations may come up with slightly different logic models for one intervention. For example, Figure 1 emphasizes on the environment of the intervention as a parameter, while Figure 2 is more focused on giving a framework for intervention development instead of one intervention.

The Theory of Change (TOC) [17] supports the logic model by defining long-term goals and then "maps backward to identify changes that need to happen earlier (preconditions). The identified changes are mapped graphically in causal pathways of outcomes, showing each outcome in logical relationship to all the others". TOC is a helpful tool to test hypotheses and assumptions on the best interventions to reach a desirable goal and identifies measurable indicators of success for evaluating interventions [17].

Behind the development and implementation of an intervention, there is an implied assumption that the activities will result in progress towards the objectives. This entails the existence of causal relationships between the activities, their immediate outputs and the expected impacts of the intervention. We use causal diagrams [19] to represent these causal relationships. Figure 3 shows an example of such a causal graph.

Graphs representing causal pathways have been extensively used [19], [23]–[28] to visualize the links among

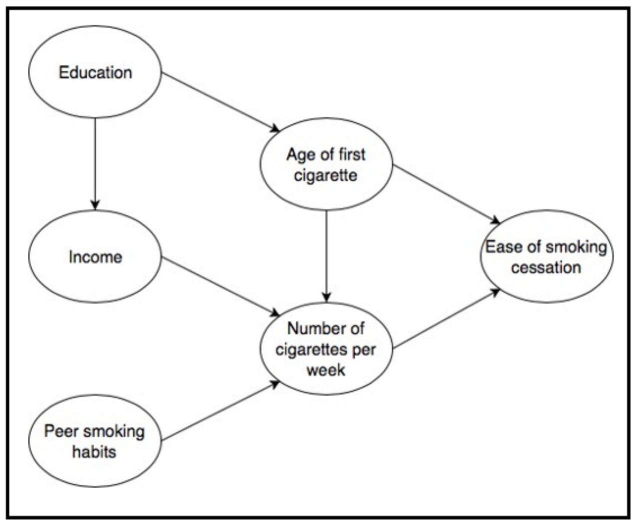

FIGURE 3. An example of a causal graph.

variables of interest. Causal graphs, created from a merger of graphical probability theory with path diagrams, provides a powerful yet intuitive tool to show and "deduce the statistical associations implied by causal relations" [24]. Recent advances in causal inference and discovery [23] rely heavily on advanced mathematical theories that address problems such as "confounding control, policy analysis, mediation, missing data and the integration of data from diverse studies" [26]. Knowledge representation and Semantic Web tools and techniques have been also used to construct and infer causal relationships in multiple health domains [19], [29], [30].

Causal networks from observational and interventional data, including the ones involving multiple cause-effect relationships, can be inferred using different statistical and computational approaches [23], [26], [28], [29], [31]. However, most of the existing causal inference approaches are only applicable for a small number of public health applications, due to the complexity and a high number of parameters (many uncertain) in the field. One of the main novelties of our proposed approach is the disambiguation of parameters and providing of semantic explainability for causal links and relations between different elements of an intervention or between different interventions. Our method provides a semantic backbone to explain and explore questions throughout all three levels (association, intervention, and counterfactuals) of causal hierarchy [32]. One of the open problems [28] in the field of causal discovery is how to determine the semantic and syntactic properties (e.g., soundness, consistency) for existing causal search algorithms. Again, our proposed ontology-driven approach provides semantics that is transparent to the end-users and rich enough to support automatic consistency and soundness assessment of causal pathways through logical reasoning. Moreover, it offers flexibility to accommodate a variety of intervention design, including partial or multi-level interventions. By utilizing detailed axioms, this method also expands the set of criteria that may be used for the identification of a causal graph.

Defining the correct causal graph is difficult because causality is stronger than simply correlation. By considering the notion of temporality, e.g., in Figure 3, "*Age of the first cigarette*" predates the current "*Number of cigarettes per week*" one sees that the causality can only go one way. In a similar way, there are connections that can, intuitively, go in only one direction, e.g., changing "*Number of cigarettes per week*" will not result in an increase or decrease of someone's "*Income*". It is also worth observing that this information is qualitative and does not explain how two values are connected. Indeed, it is possible that the sign of causal relation might change. For instance, increasing the number of helplines may initially increase the ability to find help. However, if the number of helplines grows so large that it causes confusion, it may reduce that ability instead.

Each intervention is tailored for a given population, designed with specific intent and knowledge by different persons with different backgrounds and ideas. This means that two interventions, even if they use the same processes, are likely to use slightly different vocabularies which makes it harder to federate and reuse their data and results. For instance, "*Quitline*" in Figure 1 and "*telephone helpline*" in Figure 2 correspond to the same general concept but do not use the same lexicon. In order to overcome this issue, we use ontologies that define a uniform common lexicon that can be used as a "lingua franca" when dealing with the various interventions. One of the first tasks in intervention evaluation is thus to "translate" the intervention in the language of the ontology to make it possible to compare it with the existing resources. Many ontologies exist in the biomedical domain [33]. Several efforts [34], [35] are underway to develop ontologies specifically to describe human behavior change and the interventions that cause them.

III. EVALUATING INTERVENTIONS

In order to evaluate the interventions, we extracted the processes from the logic models and compared them to the causal diagrams representing the same processes. To be more precise, the idea is to look for the causal paths that are identified by logic models in the causal graphs.

Our proposed algorithm contains several different steps. The first step is to express the information in both logic models and the causal diagrams using the ontologies. Several challenges exist during this process. Logic models often use terms that cannot be interpreted automatically because they use words that are specific to a particular intervention (e.g., "*Quitline*" in Figure 1) or because they can cover many different ideas (e.g., "*high-quality quitline services*" in Figure 1) that may not be regrouped into one concept in the hierarchy of the ontology. Mapping logic models to ontologies often require some computational work (e.g., to compute the list of "*tobacco-related medical expenditures*" from the ontology). A solution to tackle these problems is to update the ontology regularly with new concepts, relations, and axioms but this may create new challenges. When adding new concepts and relations, if they are not globally accepted by the community, the interoperability of the method will be reduced, which is highly detrimental. Adding new axioms can be even more damaging as it could lead to inconsistent ontologies that would thus lose all their meaning. We refer the interested readers to our previous works [9], [10] on how to manage changes and maintain interoperability.

The next step is to identify the causal paths implied by the logic models. In most cases, this step only consists in following the arrows in the logic models. However, in some cases, e.g., in Figure 1, the arrows are not specific enough and choices have to be made on which concepts on the left of the arrow should be connected to which concepts on the right. The consequences of the choice depend on how one tries to connect the paths found in the logic models to the causal graph. We can assume for now that every concept on the left, that we will call "cause (of the arrow)" for ease of understanding, is connected to every concept on the right, that we will call "consequence (of the arrow)". Some of these connections will likely be erroneous and can be removed later if such a problem is detected. Another approach would be to infer the causal relations from the logic models using a heuristic or an algorithm but that would likely use the causal diagram and thus affect the results of the evaluation.

Once the causal paths have been determined, one needs to check whether they appear in the causal graph. If one assumes that all concepts on the left need to be connected to every concept on the right, then it is unlikely to find all these paths in the causal graph. A better solution would be to check whether there exists a path connecting each cause to at least one consequence and at least one cause to each consequence. If that is the case, every concept that appears in the logic model can be expected to have an impact and is thus meaningful. In such a situation, the intervention seems sensible.

In order to improve this method, one may want to slightly modify the notion of the causal graph and make a difference between causes that positively or negatively affect an outcome of interest [36]. The algorithm then asks to identify how the different concepts are supposed to evolve, i.e., whether we want them to increase or decrease. The algorithm then checks that the existing paths in the causal graph match this requirement.

Going slightly further towards a quantitative evaluation of the interventions, we can use the existing knowledge to estimate the relative strength of every connection in the causal diagram. To be able to associate explicit values to the components of the causal diagram, we must be able to find a mapping between the causal diagram and the existing data. Semantic rules, e.g. Positional-Slotted Object-Applicative Rule Markup Language (PSOA RuleML) rules [37], can be used to define how to interpret the data into the ontology. For example, assuming a dataset contains a database with a table "*Policy results*" for policies and how youths and merchants complied with them, one can use a rule of the form shown in Figure 4 that extracts the youth compliance data from the table. Semantic web services, for instance, Semantic Automated Discovery and Integration (SADI) web services [38], can be then created using these rules to automate access to the data. Rules can be a very powerful tool to connect multiple sources and to increase the semantic information associated with the data but their exact interaction with causal reasoning still needs to be investigated.

Once the existing data is available, statistical analysis is used to weigh each arrow. As noted previously, some of the arrows may have negative weight. For instance, an increase in "*Number of cigarettes per week*" could negatively impact "*ease of smoking cessation*" in the causal diagram of Figure 3. Therefore, the existence of a path from the inputs to the targeted outcomes may not necessarily be an indication of the likelihood of success. Rather than focusing only on the existence of a path, the algorithm will now focus on the product of the weights of the arrows along such a path. We may decide that an intervention is deemed likely to be successful if the paths from the inputs to the outcomes are all positive, i.e., the products of the weights along each path are positive. Another option could be to require the paths to have a value higher than a given threshold. However, such an algorithm requires greater care because the weights of different arrows may not have comparable value.

IV. GENERATING CAUSAL DIAGRAMS

One of the key challenges of using causal graphs is that they usually do not readily exist. This means that we need first to generate causal graphs with which to compare the logic models of the interventions. In order to create those causal diagrams, we propose two different methods.

```
Forall ?p ?yc ?mc (
       policy_result(?p, ?yc, ?mc) :-
       policy_youth_compliance(?p, ?yc)
)
```

FIGURE 4. An example of a rule.

The first method makes use of the ontology and its contextual knowledge. We assume the existence of enough data to accurately represent the domain on which the interventions take place. In the case of smoking cessation, this does not seem unrealistic because of the well-documented history of smoking, its health effects and the different approaches to smoking cessation. From the data, using a statistical approach to causality [27], it is quite feasible to generate a graph by analyzing, for every pair of concepts, if there is a causality relationship between them. Obviously, as the data should be diverse and comprehensive enough to be representative, this requires a lot of computation before the evaluation can take place. Furthermore, there is a risk that the resulting causal graph will be extremely complex and highly connected which would severely impede the intervention evaluation. In particular, it might create so many paths in the causal graph that it is almost always possible to find one matching the logic model of the intervention under consideration.

The second method uses the same idea as the evaluation, i.e., using the inherent causal knowledge in the logic models. For this to work, we assume that there exists at least one successful intervention from which knowledge can be extracted. The problems that arise with this method are the same as those for the evaluation. First, contrary to what happened in the first method, the logic models are probably not expressed in the same language as the ontology, and we must parse it to generate the right conceptualization. The correct connections may also be difficult to infer given that the tasks, outcomes, and goals are usually fairly general which may yield an inconsistent or even wrong model. Third, the intervention might be successful despite some inconsistency or miss-match between its components, which results in misleading connections in the causal graph. As more data becomes available, it would be possible to refine the causal graph by extracting knowledge from more than one logic model and filtering out conflictual information. Finally, this method does not weigh the connections between concepts or their strengths. This is not a major problem for a purely qualitative approach but will be a hurdle when a more thorough quantitative approach is followed.

V. APPLICATION

In this section, we will demonstrate the applicability of our method through a simple example. We use a fragment from a logic model for smoking cessation (Figure 5) used in

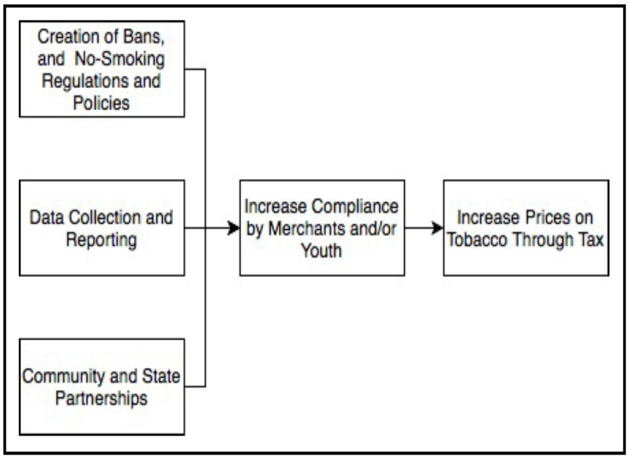

FIGURE 5. A fragment of a logic model for smoking cessation adapted from the Maryland Department of Health and Mental Hygiene [39].

Maryland [39]. Because no real causal diagram exists and we do not have access to sufficient data to infer a realistic causal graph, we will use a dummy one created solely for this example.

The first step is to express the domain knowledge in a formal, machine-readable ontology language. This is already an arduous task. We then map every box of the logic model to an ontological element (e.g., an ontological concept). Figure 6 shows an example of such mapping. Many concepts need to be created because they either belong to slightly different domains or are too specific to have been previously defined. For instance, there is no concept expressing exactly "*Data Collection and Reporting*" so one needs to merge different concepts, which may come from different ontologies. In this case, we express "*Data Collection and Reporting*" as the union of "*Data Collection*" and "*Reporting*".

In many cases, several ontologies define a given concept, with different levels of granularities and expressivity. For instance, both the National Cancer Institute Thesaurus (NCIT) [40] and SNOMED CT [41] define a concept for "Community". Using one definition instead of the other means reducing interoperability because only part of the existing applications will be consistent with the choice made. Hence, we state that both definitions are equivalent and use both definitions at the same time. It can create other problems too if there are axioms in either ontologies that are in conflict with the other ontology, but that is usually not the case. Another possible problem is that the chosen definition for a concept in an ontology may not be the one that we are looking for. For instance, NCIT defines a concept for "*Partnership*"

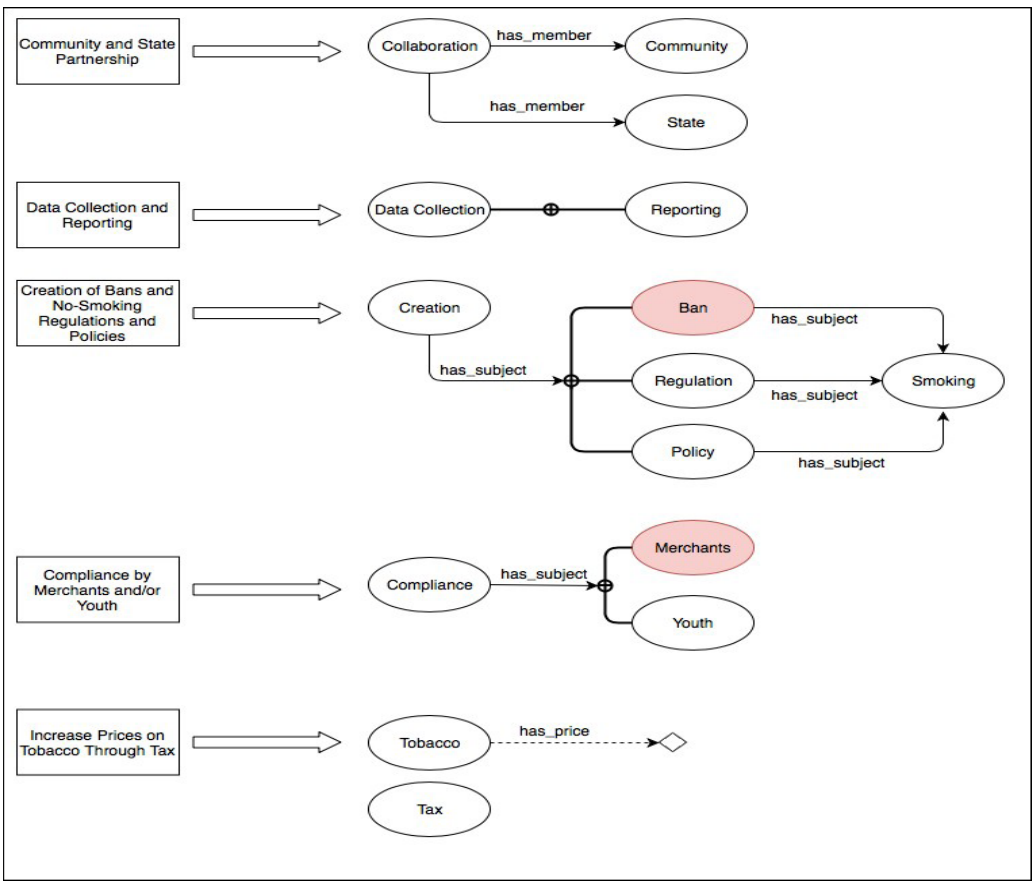

FIGURE 6. A possible mapping of the boxes from the logic model to an ontology.

but its definition is 'A business enterprise entered into for profit which is owned by more than one person, each of whom is a "partner."' [41]. When that is the case, one needs to find another concept with a meaning closer to what they need to express: in this case we chose "*Collaboration*" from SNOMED CT. A similar problem exists for "*Youth*" where the definition is 'The time of life between childhood and maturity. This period overlaps with adolescence.' when what we are looking for is a person whose state, or age category, is "*Youth*". We didn't make that distinction in Figure 6 to try to make the mapping easier to understand and represent. Furthermore, some concepts are not currently defined in the relevant ontologies for many different reasons. In many cases, for instance, "Compliance", the concept that we are interested in is not central to any biomedical domain and thus not defined. In some cases, for example, "*Merchant*" there exist other concepts (e.g., "*Vegetable seller*" in SNOMED CT) that are more specific. This is an indication that adding a more general concept might be a good idea. In other cases, for instance, "*Ban*", the missing concept can be seen as a sub-category of an existing concept (e.g., of "*Regulation*" in NCIT) that can be easily added to the ontology. A final case is when a concept is not the best way to represent an idea in the ontology. For instance, "*Price*" is a value and is thus better represented by a datatype relation. Similarly, an "*Increase*" represents a modification of a data value and is thus not a concept. On the other hand, there might be cases where concepts exist that correspond to some part of a logic model, but we would rather not use them. For instance, using "*Smoke-Free Policy*" from MESH [42] to represent the non-smoking policies may be a good idea but it creates an unbalance between "*Smoke-Free Policy*" and "*Regulation has_subject Smoking*" if a similar concept does not exist to describe it.

We now generate a causal diagram from the logic model. It will have roughly the same shape and directionality as the logic model. As much as possible, we need to try to avoid using newly defined concepts because it is unlikely, they will be used by other applications which would reduce interoperability.

In this situation, some of the qualifiers that were used previously, e.g., "*Creation*", do not matter and can be removed. We obtain the causal diagram in Figure 7. The logic model contains the information that the "*Compliance by Merchants and/or Youth*" affects the "*Price of Tobacco*" through "*Taxation*". This is a piece of information that should appear in the causal diagram but the logic model does not explain how "*Taxation*" is related to the "*Compliance*"s. Here, an ontology-driven causal diagram can help us improve interpretability, transparency, and explainability of our model. As the causal diagram uses values, ideally every node of the causal diagram should correspond to a data relation in the ontology. Such is the case for "*Tobacco Price*" that corresponds to "*Tobacco has_price*". This is not always possible, and indicators may have to be found when a concept is used instead. However, as we are only interested in the qualitative

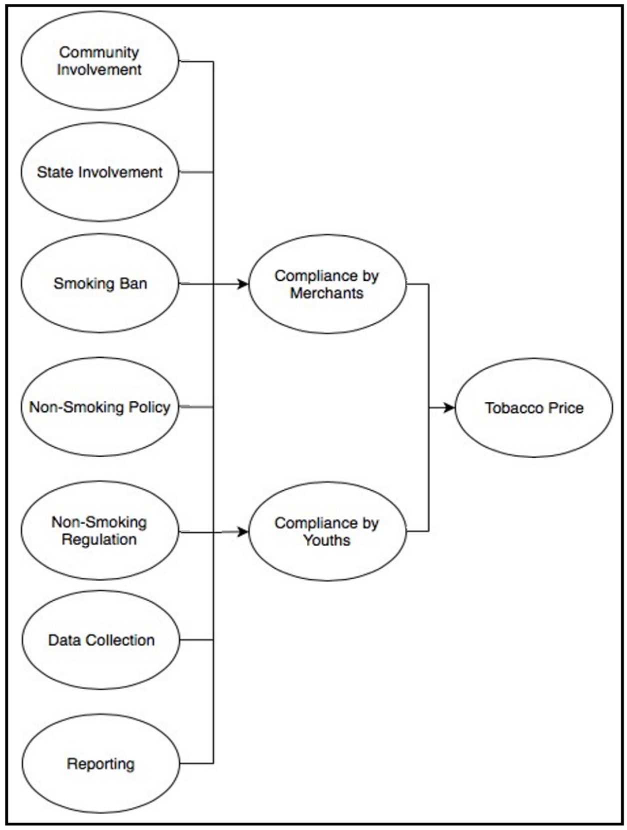

FIGURE 7. A possible causal diagram resulting from the logic model.

evaluation of the interventions, not being able to assign values to the concepts is not a big problem.

In order to evaluate the intervention represented by the logic model in Figure 5, we use a dummy causal diagram, shown in Figure 8, that represents the concepts present in the logic model. The goal is to check the existence of the causal paths implied by the logic model. For instance, the logic model implies that there is a causal path between a "*Smoking Ban*" and "*Compliance by Merchants*". Indeed, a direct path can be found in Figure 8. In a similar way, the logic model implies the existence of a causal path between "*State Involvement*" and "*Compliance by Merchants*" that is not directly in the causal diagram but still exists. Using the underlying semantics provided by ontologies, one can infer indirect paths within a causal diagram. On the other hand, the logic models contain a path from "*Data Collection*" to "*Compliance by Youths*" that does not exist in the causal diagram. As the causal diagram does not contain any edge from "*Compliance by Youths*" to "*Tobacco Price*", the evaluation would validate the paths from "*Community Involvement*", "*State Involvement*" and "*Smoking Ban*" but not the others. In order to obtain a more quantitative approach, it is possible to use existing data (e.g., generated through statistical analysis) to populate a causal diagram before comparing it to the one created from the logic model. Many different approaches can be used to weight the connections in the populated causal diagram. For example, one could use a linear regression and time series

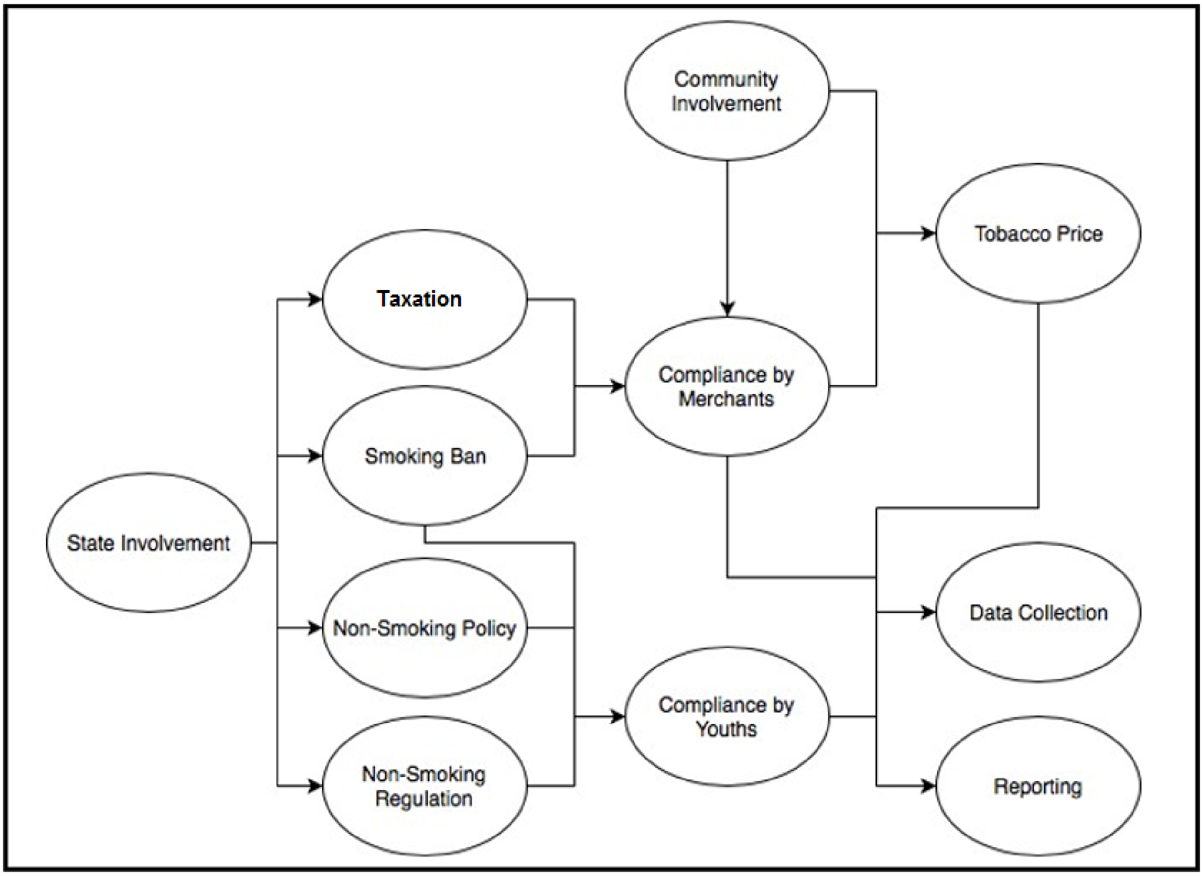

FIGURE 8. A possible causal diagram to compare to the logic model.

models between the various data points. By using the slope of the regression, one gets an approximation of how a change in one of the two parameters affects the evolution of the other. It is, obviously, not perfect and only works if the linear regression is a good approximation of the relationship between the two parameters. Another key difficulty is that it requires huge amounts of data that may not be available or exist. As we are, for now, mostly interested in the theory, we can assume that we have enough data for the method to function properly. That may not necessarily be the case in actual applications.

Using the weighted causal diagram, the paths found in the qualitative evaluation can now be calculated and weighted too. In order to obtain the weight of a path, we multiply the weights of all the edges that it contains. It is not a fully accurate estimate of the relationship between the input resources and the outcomes because several concepts may have all sorts of causes outside of the path that are not taken into account in the evaluation. It should, however, give an idea of what can be expected.

It is worth pointing out that our goal here is not to determine the best way to weight the edges of the causal diagram. The algorithm computing the weights is a key part of any application of our method but, given that our focus is on the theory and we are not collecting any data for the computation of weights, we only use random numbers as an illustration of our method in this example. A complete study of the possible choices for the weights and how to extend them from edges to paths would be of great interest but lies outside of the scope of this paper.

Let us use the modified causal diagram in Figure 9. The values it contains, as discussed above, are purely arbitrary and do not reflect any actual knowledge. In a real-world application, the weights must be computed or inferred from the existing data. One can see that some of the paths that had been identified previously in the qualitative evaluation appear not to be propitious. For instance, the path "*Smoking Ban*" to "*Compliance by Merchants*" to "*Tobacco Price*" has a weight of -3.4 and would thus likely not have positive results. On the other hand, the path "*State Involvement*" to "*Taxation*" to "*Compliance by Merchants*" to "*Tobacco Price*" has a weight of 2.04 and is thus expected to yield the intended results. Some inputs can yield several different weights that can yield similar results, e.g., "*Community Involvement*" is at the start of two paths ending in "*Tobacco Price*" both having positive weights (2.7 for the direct one,

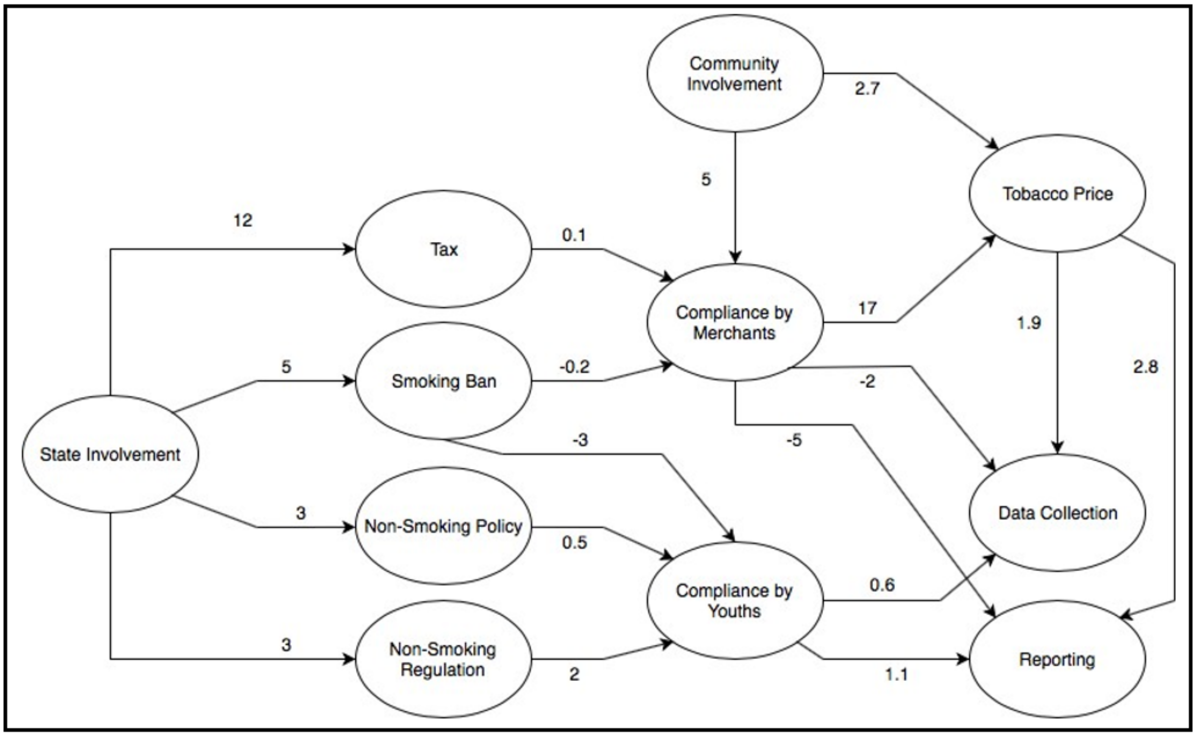

FIGURE 9. A causal diagram with weighted edges.

8.5 for the indirect one), or different results, e.g., "*State Involvement*" also starts a path through "*Smoking Ban*" with a negative weight.

VI. DISCUSSION AND CONCLUSION

In this paper, we proposed a formal approach to evaluate public health interventions. It rests on the idea of studying the causal relationships behind a target intervention, which is expressed in its logic model. This causal pathway is then compared to causal diagrams constructed from real-world evidence data. Several hurdles needed to be cleared for our method to be fully functional. In order to be able to effectively compare the information coming from the logic model of an intervention with the causal diagrams, they must be expressed using similar concepts. This problem can be addressed by using formal ontologies to build a common semantic model and language. Moreover, full causal diagrams are rarely readily available in the real world, even for domains in which sufficient knowledge exists. This issue can be palliated by using existing data to generate a causal diagram according to our best knowledge. To facilitate reuse and integration, we advocate the use of logic rules and semantic web services to create a framework for data access.

Most current intervention assessment methods are either based on survey data or based on a set of pre-defined assumptions without taking into account the context and its dynamics. The contextual knowledge captured in ontologies enables us to analyze how context influences outcomes. Using semantic inference and causal reasoning within an explainable AI framework we can determine what elements in a disease causal pathway (e.g. behaviors, actions, or conditions) must be changed to stop the recurrence of similar, undesired outcomes.

Using the formal consensus knowledge and the logical reasoners enable researchers and policymakers to evaluate (or forecast) various outcomes of one specific intervention (or series of interventions) through logical inference and deductive querying (e.g. what are the potential consequences of a smoke cessation intervention on the prevalence of lung cancer in a defined community in the next decade?)

We have proposed several approaches to intervention evaluation. The purely qualitative approach uses the casual graph and checks for the existence of the causal paths that are present in the logic model. The more quantitative approach additionally checks whether these paths affect the outcome in the right direction. This creates the need for more elaborate causal diagrams that reflect the relative importance of connections. This can be done by using the existing data. To the best of our knowledge, this work is the first study that investigates automatic evaluation and assessment of health interventions using semantic and causal reasoning and inference.

We hope that our intervention evaluation can be useful to all the actors involved in public health intervention design and implementation. We, however, acknowledge that this method is only the first step in creating an automated intervention evaluation tool. Considering that in the real-world the activities defined within a logic model does not usually occur in a straight-line sequence we used formal ontologies to model the way we believe things will happen, rather than how it actually happens. The underlying theory of change provides a "forecast that shows what conditions we believe must exist for other conditions to come into being" [43] (e.g. if x, then y), and goes through by several phases of evaluation and recalibration throughout an intervention evaluation life cycle. Each application requires the creation of a specific causal diagram and each causal diagram requires the acquisition and sharing of knowledge. Furthermore, every additional piece of data can be used to refine the evaluation and make it more accurate. One area for future research is evaluating the way the causal diagram models reality. Our method relies on the assumption that paths are enough to describe the domain, but this is not necessarily the case. The more quantitative approach also raises additional problems as one needs to figure out how to assign weights to the edges, whether they are constants or functions, how to weight the paths from the weights of the edges, etc. More theoretical work is underway to make the proposed model fully functional.

REFERENCES

[1] Centers for Disease Control and Prevention. *Program Performance and Evaluation Office*, Accessed: Dec. 28, 2019. [Online]. Available: https://www.cdc.gov/eval/index.htm

[2] Centers for Disease Control and Prevention. *Types of Evaluation*. Accessed: Dec. 28, 2019. [Online]. Available: https://www.cdc.gov/std/Program/pupestd/Types%20of%20Evaluation.pdf

[3] G. Moore, S. Audrey, M. Barker, L. Bond, C. Bonell, C. Cooper, W. Hardeman, L. Moore, A. O'Cathain, T. Tinati, D. Wight, and J. Baird, "Process evaluation in complex public health intervention studies: The need for guidance," *J. Epidemiol. Community Health*, vol. 68, no. 2, pp. 101–102, Feb. 2014.

[4] D. Gunning and D. Aha, "DARPA's explainable artificial intelligence (XAI) program," *AI Mag.*, vol. 40, no. 2, pp. 44–58, Jun. 2019.

[5] J. H. Brenas, E. K. Shin, and A. Shaban-Nejad, "Adverse childhood experiences ontology for mental health surveillance, research, and evaluation: Advanced knowledge representation and semantic Web techniques," *JMIR Mental Health*, vol. 6, no. 5, Apr. 2019, Art. no. e13498.

[6] A. Shaban-Nejad, H. Mamiya, A. Riazanov, A. J. Forster, C. J. O. Baker, R. Tamblyn, and D. L. Buckeridge, "From cues to nudge: A knowledge-based framework for surveillance of healthcare-associated infections," *J. Med. Syst.*, vol. 40, no. 1, p. 23, Jan. 2016.

[7] A. Shaban-Nejad, D. L. Buckeridge, and L. Dube, "COPE: Childhood obesity prevention [knowledge] enterprise," in *Artificial Intelligence in Medicine* (Lecture Notes in Computer Science), vol. 6747. Berlin, Germany: Springer, 2011, pp. 225–229.

[8] A. Shaban-Nejad, M. Lavigne, A. Okhmatovskaia, and D. L. Buckeridge, "PopHR: A knowledge-based platform to support integration, analysis, and visualization of population health data," *Ann. New York Acad. Sci.*, vol. 1387, no. 1, pp. 44–53, Jan. 2017.

[9] M. S. Al Manir, J. H. Brenas, C. J. Baker, and A. Shaban-Nejad, "A surveillance infrastructure for malaria analytics: Provisioning data access and preservation of interoperability," *JMIR Public Health Surveill.*, vol. 4, no. 2, May 2018, Art. no. e10218.

[10] J. H. Brenas, M. S. Al-Manir, C. J. O. Baker, and A. Shaban-Nejad, "A malaria analytics framework to support evolution and interoperability of global health surveillance systems," *IEEE Access*, vol. 5, pp. 21605–21619, 2017.

[11] A. Shaban-Nejad, A. Okhmatovskaia, E. K. Shin, R. L. Davis, and D. L. Buckeridge, "POLE.VAULT: A semantic framework for health policy evaluation and logical testing," *Stud. Health Technol. Inform.*, vol. 245, pp. 1335–1335, Jul. 2017, doi: 10.3233/978-1-61499-830-3-1335.

[12] A. Shaban-Nejad, A. Okhmatovskaia, E. K. Shin, R. L. Davis, B. E. Franklin, and D. L. Buckeridge, "A semantic framework for logical cross-validation, evaluation and impact analyses of population health interventions," *Stud. Health Technol. Inform.*, vol. 235, pp. 481–485, Apr. 2017, doi: 10.3233/978-1-61499-753-5-481.

[13] T. O. A. Lehtinen, M. V. Mäntylä, and J. Vanhanen, "Development and evaluation of a lightweight root cause analysis method (ARCA method)—Field studies at four software companies," *Inf. Softw. Technol.*, vol. 53, no. 10, pp. 1045–1061, Oct. 2011.

[14] P. Hawe, A. Shiell, and T. Riley, "Complex interventions: How 'out of control' can a randomised controlled trial be?" *BMJ*, vol. 328, no. 7455, pp. 1561–1563, 2004.

[15] J. O. Andrews, S. D. Newman, J. Heath, L. B. Williams, and M. S. Tingen, "Community-based participatory research and smoking cessation interventions: A review of the evidence," *Nursing Clinics North Amer.*, vol. 47, no. 1, pp. 81–96, Mar. 2012.

[16] P. Bader, D. Boisclair, and R. Ferrence, "Effects of tobacco taxation and pricing on smoking behavior in high risk populations: A knowledge synthesis," *Int. J. Environ. Res. Public Health*, vol. 8, no. 11, pp. 4118–4139, Oct. 2011.

[17] D. H. Taplin, H. Clark, E. Collins, and D. C. Colby, "Theory of change," ActKnowledge, New York, NY, USA, Tech. Rep., Apr. 2013. Accessed: Dec. 28, 2019. [Online]. Available: http://www.theoryofchange.org/wp-content/uploads/toco_library/pdf/ToC-Tech-Papers.pdf

[18] J. D. Wright, Ed., *International Encyclopedia of the Social & Behavioral Sciences*, 2nd ed. Oxford, U.K.: Elsevier, 2015, pp. 299–305.

[19] J. Pearl, *Causality: Models, Reasoning, and Inference*. 2nd ed. New York, NY, USA: Cambridge Univ. Press, 2009.

[20] L. Liu and T. M. Özsu, Eds., *Encyclopedia of Database Systems*. New York, NY, USA: Springer-Verlag, 2009.

[21] *National Tobacco Quitline Logic Model*. Accessed: Dec. 28, 2019. [Online]. Available: https://www.cdc.gov/tobacco/about/foa/quitline-capacity/pdfs/logic-model.pdf

[22] *Smoking Cessation*. Accessed: Dec. 28, 2019. [Online]. Available: http://www.healthscotland.com/OFHI/tobacco/logicmodels/lm_03.html

[23] F. Eberhardt, "Introduction to the foundations of causal discovery," *Int. J. Data Sci. Anal.*, vol. 3, no. 2, pp. 81–91, Mar. 2017.

[24] S. Greenland and J. Pearl "Causal diagrams," in *Wiley StatsRef: Statistics Reference Online*. Hoboken, NJ, USA: Wiley, Nov. 2017, doi: 10.1002/9781118445112.stat03732.pub2.

[25] A. Marx and J. Vreeken, "Telling cause from effect by local and global regression," *Knowl. Inf. Syst.*, vol. 60, no. 3, pp. 1277–1305, Sep. 2019.

[26] J. Pearl, "The mathematics of causal inference," in *Proc. Joint Stat. Meetings Conf.* Alexandria, VA, USA: American Statistical Association, 2013, pp. 2515–2529.

[27] J. Peters, D. Janzing, and B. Schölkopf, *Elements of Causal Inference: Foundations and Learning Algorithms* (Adaptive Computation and Machine Learning Series). Cambridge, MA, USA: MIT Press, 2017.

[28] P. Spirtes and K. Zhang, "Causal discovery and inference: Concepts and recent methodological advances," *Appl. Inf.*, vol. 3, no. 1, pp. 1–28, Feb. 2016.

[29] M. Steyvers, J. B. Tenenbaum, E.-J. Wagenmakers, and B. Blum, "Inferring causal networks from observations and interventions," *Cogn. Sci.*, vol. 27, no. 3, pp. 453–489, May 2003.

[30] K. Kozaki, Y. Yamagata, R. Mizoguchi, T. Imai, and K. Ohe, "Disease compass- a navigation system for disease knowledge based on ontology and linked data techniques," *J. Biomed. Semantics*, vol. 8, p. 22, Jun. 2017.

[31] V. K. Raghu, A. Poon, and P. V. Benos, "Evaluation of causal structure learning methods on mixed data types," *Proc. Mach. Learn. Res.*, vol. 92, pp. 48–65, Aug. 2018.

[32] J. Pearl, "The seven tools of causal inference, with reflections on machine learning," *Commun. ACM*, vol. 62, no. 3, pp. 54–60, Feb. 2019.

[33] *NCBO BioPortal*. Accessed: Dec. 28, 2019. [Online]. Available: https://bioportal.bioontology.org/

[34] E. Norris, A. Finnerty, M. M. Marques, R. West, J. Thomas, P. M. Aonghusa, M. Johnston, M. P. Kelly, and S. Michie, "The human behaviour-change project: Developing a behaviour change intervention ontology," in *Proc. 8th Int. Conf. Biomed. Ontol. (ICBO)*, vol. 2137, 2018.

[35] K. R. Larsen, S. Michie, E. B. Hekler, B. Gibson, D. Spruijt-Metz, D. Ahern, H. Cole-Lewis, R. J. B. Ellis, B. Hesse, R. P. Moser, and J. Yi, "Behavior change interventions: The potential of ontologies for advancing science and practice," *J. Behav. Med.*, vol. 40, no. 1, pp. 6–22, Feb. 2017.

[36] A. Okhmatovskaia, A. Shaban-Nejad, M. Lavigne, and D. Buckeridge, ''Addressing the challenge of encoding causal epidemiological knowledge in formal ontologies: A practical perspective,'' *Stud. Health Technol. Inform.*, vol. 205, pp. 1125–1129, Aug. 2014.
[37] H. Boley, ''A RIF-style semantics for RuleML-integrated positional-slotted, object-applicative rules,'' in *Rule-Based Reasoning, Programming, and Applications* (Lecture Notes in Computer Science), vol. 6826. Berlin, Germany: Springer, 2011, pp. 194–211.
[38] M. D. Wilkinson, B. Vandervalk, and L. McCarthy, ''The semantic automated discovery and integration (SADI) Web service design-pattern, API and reference implementation,'' *J. Biomed. Semantics*, vol. 2, no. 1, p. 8, 2011.
[39] M. Kanamori, S. Randolph, O. Carter-Pokras, and J. Wallen, ''Maryland tobacco use prevention and cessation logic model,'' School Public Health, Univ. Maryland, College Park, MD, USA, 2008. Accessed: Dec. 28, 2019. [Online]. Available: https://mdquit.org/sites/default/files/pdf_files/Logic-Model.pdf
[40] *National Cancer Institute Thesaurus*. Accessed: Dec. 28, 2019. [Online]. Available: https://ncithesaurus.nci.nih.gov/ncitbrowser
[41] *SNOMED CT*. Accessed: Dec. 28, 2019. [Online]. Available: https://www.nlm.nih.gov/healthit/snomedct/index.html
[42] National Library of Medicine. *Medical Subject Headings*. Accessed: Dec. 28, 2019. [Online]. Available: https://www.nlm.nih.gov/mesh/meshhome.html
[43] E. Breuer, L. Lee, M. De Silva, and C. Lund, ''Using theory of change to design and evaluate public health interventions: A systematic review,'' *Implement. Sci.*, vol. 11, no. 1, p. 63, May 2016.

JON HAËL BRENAS received the double master's degree in mathematical modeling and digital imagery specialized in modeling, computation, and simulation from the ENSIMAG, Grenoble, France, and in computer engineering from the Politecnico di Torino, Torino, Italy, and the Ph.D. degree from Université Grenoble Alpes, Grenoble. He currently holds a senior postdoctoral position with the Big Data Institute - Nuffield Department of Medicine, University of Oxford, U.K. From 2017 to 2019, he was a Postdoctoral Fellow with the UTHSC-OAK-Ridge National Laboratory (ORNL) Center for Biomedical Informatics, and with the Department of Pediatrics, The University of Tennessee Health Science Center (UTHSC). His research interests include in graphs and graph transformations application to biomedical and health informatics.

ARASH SHABAN-NEJAD received the Ph.D. and M.Sc. degrees in computer science from Concordia University, Montreal, and the M.P.H. degree from the University of California at Berkeley. His additional training was accrued at the Harvard School of Public Health. He is currently an Assistant Professor with the UTHSC-OAK-Ridge National Laboratory (ORNL) Center for Biomedical Informatics, and with the Department of Pediatrics, The University of Tennessee Health Science Center (UTHSC). He is also an Adjunct Faculty with the Bredesen Center for Interdisciplinary Research and Graduate Education, The University of Tennessee, Knoxville. Prior to UTHSC, he was a Postdoctoral Fellow with the McGill Clinical and Health Informatics Group, McGill University. His primary research interests are population health intelligence, epidemiologic surveillance, and big-data semantic analytics using tools and techniques from artificial intelligence, knowledge representation, and semantic web.