



# The Artificial Facilitator: Guiding Participants in Developing Causal Maps Using Voice-Activated Technologies

Thrishma Reddy<sup>1</sup>, Philippe J. Giabbanelli<sup>2</sup>(✉), and Vijay K. Mago<sup>1</sup>

<sup>1</sup> Department of Computer Science, Lakehead University, Thunder Bay, ON, Canada  
{tshivara,vmago}@lakeheadu.ca

<sup>2</sup> Computer Science Department, Furman University, Greenville, SC, USA  
giabbanelli@gmail.com

**Abstract.** Complex problems often require coordinated actions from stakeholders. Agreeing on a course of action can be challenging as stakeholders have different views or ‘mental models’ of how a problem is shaped by many interacting causes. Participatory modeling allows to externalize mental models in forms such as causal maps. Participants can be guided by a trained facilitator (with limitations of costs and availability) or use a free software (with limited guidance). Neither solution easily copes with large causal maps, for instance by preventing redundant concepts. In this paper, we leveraged voice-activated virtual assistants to create causal models at any time, without costs, and by avoiding redundant concepts. Our three case studies demonstrated that our artificial facilitator could create causal maps similar to previous studies. However, it is limited by current technologies to identify concepts when the user speaks (i.e. entities), and its design had to follow pre-specified rules in the absence of sufficient data to generate rules by discriminative machine-learned methods.

**Keywords:** Amazon Alexa · Causal maps · Mental models · Participatory modeling · Virtual assistant

## 1 Introduction

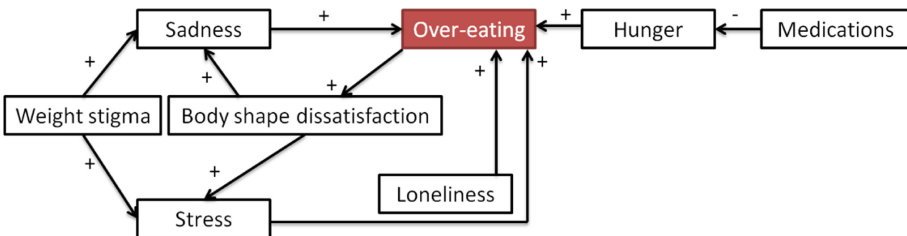
Problems as diverse as ecological management or obesity are often called *complex*, or ‘wicked’. While the complexity sciences provide many definitions and tools to measure complexity<sup>1</sup>, complex problems often share at least two traits which are central to this paper. First, they are *multifactorial*. The traditional reductionist approach trying to fix the ‘root’ cause does not lend itself well to a complex problem [1], and may even cause harm through unintended consequences [2]. Rather, the emphasis is often on mapping [3] and navigating [4] the

Research funded by MITACS Globalink Research Award, Canada.

<sup>1</sup> For an overview of the complexity sciences, see the map at [http://www.art-sciencefactory.com/complexity-map\\_feb09.html](http://www.art-sciencefactory.com/complexity-map_feb09.html).

*complex system of interactions* between factors that contribute to, and/or are impacted by, a problem of interest. Second, dissemination and implementation research emphasizes that solutions to complex problems often require *coordinated actions between stakeholders* from multiple sectors (i.e., a multiactor view [5]). For instance, actions regarding population obesity involve sectors as varied as food production, the built environment (e.g., to promote walkable cities and access to fresh food), mental and physical well-being [6]. Coordinated actions should produce a coherent policy, which implies that stakeholders work together at least by sharing a mission [7].

It can be challenging to assess whether stakeholders share a mission when operating in a complex system of interactions. They may have different views or ‘mental models’ on how the factors interact, which may lead to very different takes on interventions. In the case of ecological management, one stakeholder may ignore the pressure of fishing and instead focus on the environment (e.g., enough nutrition for the fish, not too many predatory birds) while another may acknowledge that fishing reduces the fish population but downplay its importance [8]. Stakeholders may also have the same views but express them differently, for example by naming factors in different ways depending on their fields, which can create a communication gap [9,10]. Consequently, complex problems involving multiple stakeholders often involve *participatory modeling*, which allows to externalize [11] and hence compare [8] the mental models of stakeholders. There are various approaches to participatory modeling, depending on whether the objective is to be able to *simulate* a system [12,13] (e.g., to quantitatively assess *how much* effect an intervention would have) or only capture its *structure* [14] (e.g., to qualitatively assess *what* an intervention would affect). In the example of obesity, qualitative approaches may be realized by systems dynamics or agent-based modelling [15] while qualitative approaches may generate ‘systems maps’ or ‘diagrams’ [16]. The creation of systems maps is particularly important either as an endpoint (for qualitative analysis of stakeholders’ mental models), or as a step toward the creation of quantitative models [14] (e.g., starting with a Causal Loop Diagram to produce a Systems Dynamics model). Causal maps are a widely used form of systems maps, in which concepts are represented as nodes and their causal connections are captured through directed edges (Fig. 1).



**Fig. 1.** Sample causal map where “over-eating” is the problem of interest [17].

Participants interested in developing causal models have often done it with the support of a trained facilitator, who elicits concepts and causal relations [18–22]. Alternatively, tech-savvy participants may receive training and independently develop causal models using software such as `cMap` (common in education research), `MentalModeler` (most used in socio-ecological systems), or `Vensim` (typical in health and systems engineering). However, both approaches have limitations. A trained facilitator can provide ample guidance, but may be costly or unavailable. A software may be free and available anytime, but it does not guide the participant through the process of building a causal map. In addition, both approaches rely on a visual inspection of the map as it is built, which does not easily scale as participants start to have many concepts and/or interrelationships. For example, a participant may add a concept that is actually a synonym of a concept already present. To notice this redundant concept, the facilitator and/or participant would need to manually look at all other concepts, which becomes prohibitive as the number of concepts increases.

There is thus a need for an approach to causal model building that can be available at any time, without costs, and scales easily. In this paper, we address this need by leveraging voice-activated virtual assistants (`Amazon Alexa`) to design and implement a *virtual facilitator*. Our solution guides participants in developing a model through a conversation (like a human facilitator), but is available at anytime without cost (as a software) and continuously examines the map to avoid typical issues such as synonymy of concepts.

The remainder of this paper is organized as follows. In Sect. 2, we provide background information on the process to create a causal map, and we briefly discuss recent uses of conversational agents built on `Amazon Alexa`, `Microsoft's Cortana`, and `Apple Siri`. In Sect. 3, we present the process that our artificial facilitator follows, and we cover its implementation in Sect. 4. Several examples are offered in Sect. 5, where a participant interacts with our technology to develop a model. Videos of the interaction are provided as supplementary online material. Finally, Sect. 5 contextualizes the implications of this work for the development of causal models and participatory modeling in general.

## 2 Background

### 2.1 Why Do We Create Causal Maps?

A causal map is a *conceptual model*. In Modeling and Simulation (M&S), conceptual models are the first stage of model development before quantifying nodes and relationships (mathematical model [23]) and possibly implementing the model as code (computational model). Conceptual models serve multiple objectives such as identifying key elements and aspects (thus delineating the boundaries of a system) or externalizing hypotheses through a transparent list of expected relations [14]. These objectives may be sufficient to warrant the development of a conceptual model as a final product. In this case, the concep-

tual map is often analyzed using network theory<sup>2</sup>. A common type of analysis is the identification of clusters or communities to divide a complex system into broad themes, as exemplified by the Foresight Obesity Map [25,26], maps for the Provincial Health Services Authority [27], or the recent work of Allender, McGlashan and colleagues [28,29]. Other analyses may include the centrality, to identify leverage points in a system [30,31]; an inventory of loops, to better characterize and possibly change the dynamics of the system [32–34]; an exploration of disjoint paths between factors, to capture how a policy impacts an outcome in multiple ways [4,33]; or a comparison of maps, to understand how different are the mental models of participants [10,35].

Map-like artifacts may be constructed solely from data, for instance as Structured Equation Models (SEM) or Fuzzy Cognitive Maps (FCMs) [36]. Alternatively, traces produced by an analyst in exploring the data can be structured in a map [37,38], or the literature on a topic can be synthesized into a map [39]. It would be overly reductive to categorize such data-driven maps as ‘objective’ compared to participant-driven maps being deemed ‘subjective’. Data can also have “biases, ambiguities, and inaccuracies” [40] and the inference process to build a map may not be perfect. Our focus is on participatory modeling (PM), in which participants drive the development of causal maps. Participatory modeling serves a different (and sometimes complementary) purpose than data-driven modeling. As detailed elsewhere [17], data-driven modeling may strive for accuracy with respect to the data whereas PM aims to be transparent and representative of the participants’ mental models. PM can thus be employed in ‘soft’ situations that lack data and rely on human expertise [41], to support decision-making processes [42], or to understand what actions would be acceptable to various stakeholders [43].

The elicitation process consists of externalizing the mental model of a participant or group into a map. The elicitation process is first and foremost a *facilitation* process: we want to support participants in expressing their perspectives, rather than judge whether what they think is ‘right’ given our own ideas. Research in cognitive sciences has long been concerned with how humans store mental models, or their “conceptualization of the world” [44]. This storage takes place in *semantic memory*, which provides functional *relationships* between objects. As we previously summarized, “if mental models are published and shared in the form of maps, it owes to the fact that we seek to capture semantic memory whose structure is network-based” [8]. On one extreme, freeform approaches such as Rich Pictures pose no constraints on the creation of maps [45], which simplifies the process for participants but limits the analytical possibilities. At the other extremes, concept maps and mind maps have a very structured process that lists concepts (e.g., via brainstorming), group them, link them, and

---

<sup>2</sup> A conceptual model is an aggregate model in which *factors* or *concepts* are connected. This is different from a ‘social network’, which is an individual model in which nodes represent individuals rather than factors. Although the methods are often similar (e.g., centrality, community detection), the application of network science to social networks is often presented as ‘social network analysis’ [24].

label the links. However, this process precludes the presence of some structures (e.g., mind maps are trees so they cannot contain cycles) which are important to characterize the dynamics of a system. Causal maps occupy an intermediate position: the development process is more guided than rich pictures, less restricted than concept maps and mind maps, and any network structure can be produced by participants<sup>3</sup>.

## 2.2 How Do We Create Good Causal Maps?

The process to produce a map as shown in Fig. 1 is relatively simple: participants create concept nodes, and link them by indicating the causal relationship to be an increase (+) or a decrease (-) [48, 49]. However, at least three issues may arise if the facilitator does not provide further guidance<sup>4</sup>. First, participants need to choose node labels that have an unambiguous quantification: having ‘more’ or ‘less’ of this concept should be a straightforward notion. For instance, labeling a concept as ‘weather’ does not work, since having more or less weather is undefined. However, having more or less rain would be defined. A facilitator thus regularly ensures that labels are quantifiable, or prompts for clarifications that would change the label. Second, users may forget about concepts that they already have, and add one with a similar name. Facilitators thus continuously monitor the maps to either avoid creating a redundant concept, or merge them once they are discovered. Given the tremendous potential for (subtle) variations in language, discovering equivalent concepts is a difficult problem, particularly as the number of concepts increases [9, 10]. Third, case studies have shown that cognitive limitations make it difficult for participants to think of structures such as loops and disjoint paths [50, 51]. In particular, Ross observed how peculiar it was that “those who set policy think only acyclically, especially since the cyclical nature of causal chains in the real world has been amply demonstrated” [52]. Without paying particular attention to loops, participants may produce star diagrams with the one central problem at the core, and every other factor directly connecting to it. Facilitators may thus prompt participants extensively for relationships, to minimize the risk of missing loops or additional paths [27, 33].

<sup>3</sup> There are at least two limitations to this representation. First, networks or graphs only represent binary relationships. However, participants may think of non-binary relationships, for instance when three concepts are directly involved together. While we have long been aware that cognitive structures could generally be represented by relations between any number of concept (e.g. using a hypergraph), it has been common practice to limit the structure to a graph [46]. Second, the network is only used to represent what is true (i.e. the existence of a causal connection between two factors) rather than what is false. As noted by Johnson-Laird, mental models also include counterexamples, which are important in decision-making processes [47].

<sup>4</sup> Some of these issues are also addressed in our tutorials at <https://www.youtube.com/watch?v=OdKJW8tNDcM> and <https://www.youtube.com/watch?v=D-2Q2IHclo4>.

### 2.3 Smart Conversational Agents

The term ‘conversational agent’ may be used loosely for any system that can carry on a conversation with a human. However, there are significant differences across systems. Unlike chatbots, smart conversational agents are not limited to performing simple conversations. And unlike embodied conversational agents, they do not provide computer-generated characters to mimic the movements or facial expressions of a virtual interlocutor. Smart conversational agents are at the confluence of speech processing, natural language processing (NLP), and artificial intelligence (AI). As detailed by Williams and colleagues [53], voice-activated devices such as Amazon Alexa or Apple Siri start by converting what a user said (i.e., an audio utterance) into text using automatic speech recognition. Words are then processed through spoken language understanding (SLU) and passed onto a dialog state tracker (DST), which results in identifying an appropriate response. The words in the response are prepared by natural language generation (NLG), and turned into audio by text-to-speech (TTS).

Smart conversational agents can be designed in many ways, as shown in the recent review by Laranjo *et al.* applied to healthcare [54]. A conversation may not be oriented toward the completion of a specific *task*, but takes place for its own sake. The flow of the discussion may be *controlled* by the system and/or the user. *Interactions* can be via spoken language and/or written language. Finally, the *dialogue management* may take the user through a sequence of pre-determined steps (i.e., a finite-state system), elicit an input and parse it using a template to decide the dialogue-flow (i.e., a frame-based system), or take an agent-based approach to focus on beliefs and desires. In the specific healthcare context reviewed by Laranjo *et al.*, agent-based approaches were uncommon (1 study) while finite (6 studies) and frame-based systems (7 studies) were equally common [54]. However, when interactions rely on voice and a task has to be accomplished, then the frame-based design is so common that the system may be presented as a *slot-based dialog system* [55].

## 3 Process in an Artificial Facilitation

As described in Sect. 2.2, the process needs to (i) obtain concept labels that are quantifiable and distinct from labels already used, and (ii) help participants provide relationships to minimize the risk of missing essential structures such as loops. To help participants track relationships, a map building process can be conceptualized as a *graph traversal*: we want to elicit/visit all of the concepts (i.e. nodes) that pertain to the user’s mental model, and we move from a concept to another using a relationship. Unlike a *graph exploration* in which we typically come back to the first node, the map building process ends on an arbitrary node.

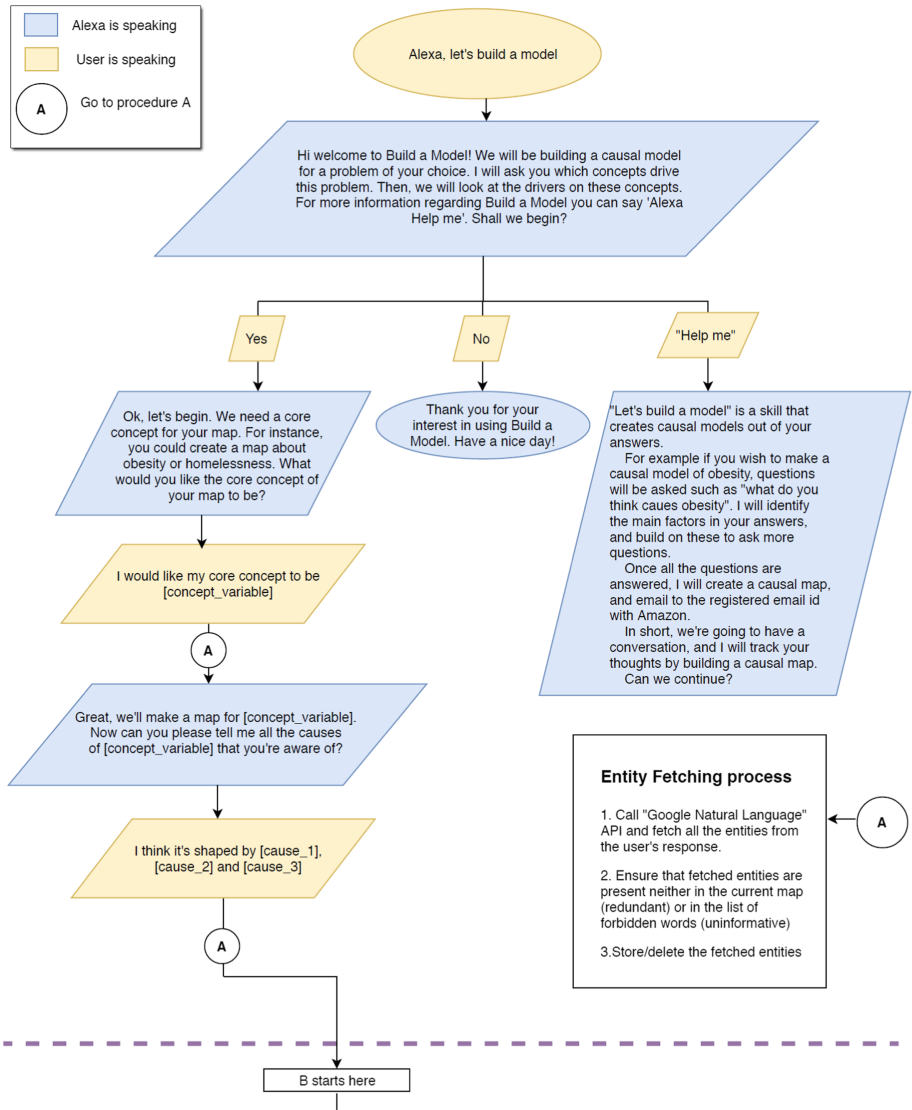
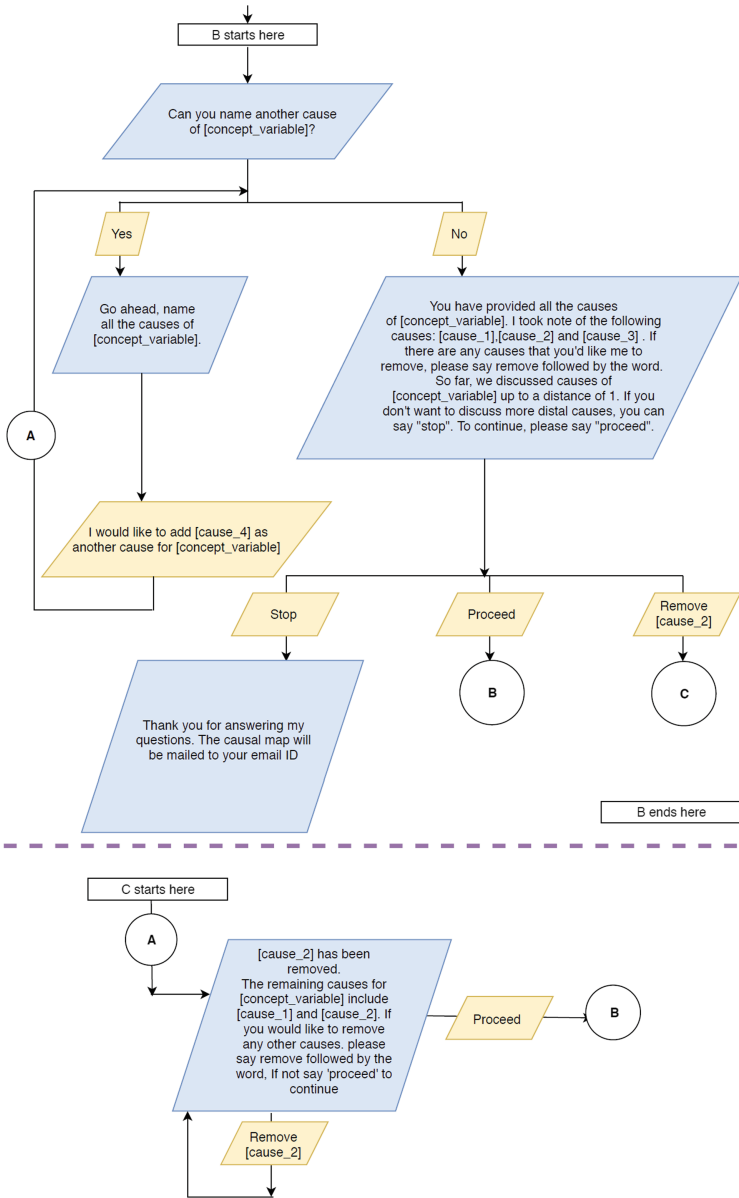


Fig. 2. Process to start a model and provide the first causes.



**Fig. 3.** Continuation of the process, showing how to get additional causes, get another layer of causes, or removing a causal edge.



**Box 1.** Sample conversation between a **user** and our facilitator using a depth-first search.

“Let’s discuss obesity. What is one cause of obesity?”

- Obesity is caused by unhealthy foods.
- What do you think would be one reason for having unhealthy foods?
- They’re cheaper than healthy ones.
- And why are they cheaper?
- Because of government subsidies for food production.
- What’s one driver of agricultural subsidies?
- I think that’s going off topic.
- Ok. Let’s go back. You said unhealthy foods are cheaper than healthy ones, because of agricultural subsidies. Any other reason?
- Not that I can think of.
- That’s fine. Let’s go back. You said people have unhealthy foods because of prices. Any other cause?”

Two typical approaches to a graph traversal are a depth-first search (DFS) and a breadth-first search (BFS). Starting from a root, a DFS follows *one* unexplored node, and from there visits another unexplored node, thus going as far as possible. When it cannot go further, it backtracks until it can branch in a new direction. This approach potentially undesirable in a facilitated process for at least three reasons. First, it can take participants on tangents, quickly going away from the main topic until they realize that factors are no longer relevant to the problem space. As a result, the map may be imbalanced, and a high cognitive load is placed on the individual who needs to frequently think of the problem’s boundaries. Second, often going *back* to a node may feel less natural than going *forward*, possibly coming across as ‘jumping’ between ideas. Third, a DFS requires that the user only provides *one* new concept each time, and may thus ask many times about the same node. This is more cumbersome than providing all known concepts at once, at moving on. These points are illustrated through an example of a DFS-based conversation in Box 1.

Starting from a root, a BFS asks for *all* connected nodes. Intuitively, it acquires the complete layer of connected concepts at distance 1 from the starting one. Then, it goes through all of these concepts and acquires all of their neighbors, thus completing the layer at distance 2. By going through entire layers at a time, it avoids taking participants on tangents. By asking whether participants want to continue when an entire layer is done, it asks for a conscious monitoring of the problem boundary at specific moments instead of offloading this responsibility onto the user at every question. By going through layers, it only goes forward (i.e. uses a queue) instead of backward (i.e. using a stack in the DFS). Finally, by asking for all connected concepts at once, users have the natural opportunity to share all of their thoughts instead of restricting themselves to a

single new concept. For these reasons, our artificial facilitator uses a breadth-first search. The functioning of a BFS is illustrated via a conversation in Box 2.

**Box 2.** Sample conversation between a **user** and our facilitator using a breadth-first search.

“Let’s discuss obesity. Can you share all of the contributors to obesity?”

- Eating unhealthy foods, lack of physical activity.

- Alright, these are the direct drivers. Do you want to go further and share what do you think drives these?

- Sure.

- What do you think would be the reasons for eating unhealthy foods?

- Could be a coping mechanism for stress or depression, an affordability issue because they’re cheap, or a personal taste preference.

- And what about physical activity?

- Could be obesity itself, because it creates some barriers. Or a lack of access to facilities, or a fear of engaging in physical activity.

- So we’ve now looked at indirect drivers. Do you think it’d be relevant to discuss their causes?”

Note that, while the BFS is meant to cover more concepts, the appearance of previous concepts can create loops. As illustrated in Box 2, we have a loop from obesity to a lack of physical activity, which itself contributes to obesity.

As shown in Figs. 2 and 3, our process utilizes the layer-by-layer approach of the BFS. It also closely monitors the names of concepts, as shown in Fig. 2 (inset A). We actively prevent the creation of similar concepts, informing the user that they are already present in the map under a possibly different name. We also attempt to avoid the use of concepts that cannot be quantified, thus promoting more operational definitions of concepts. The technology used to realize these objectives is detailed in the next section.

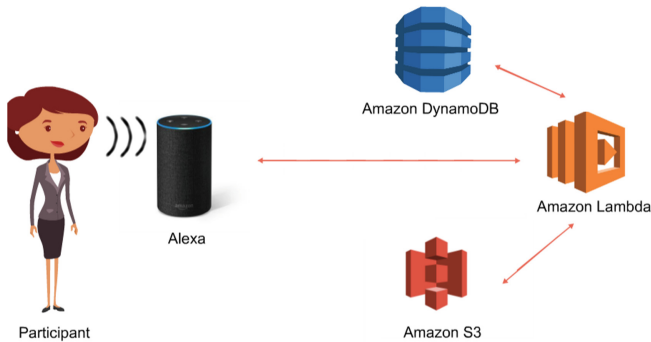
## 4 Implementation of Our Artificial Facilitator

Our implementation is task-oriented as we seek to guide a participant in externalizing their mental model. The virtual facilitator controls the flow of the conversation by asking questions. Interactions in the deployed version are exclusively vocal, but developers in Amazon Alexa also have access to a console that takes written input (for testing only). Dialogue-management uses a frame-based system. All of these technical choices were briefly discussed in Sect. 2.3.

Our code is provided at <https://github.com/datalab-science/causalMapBuilder>. Our implementation involves several technologies, shown in a high-level view in Fig. 4 and detailed in Table 1. We use Amazon Alexa as it provides automatic speech recognition and text-to-speech, in addition to working on three out of four smart speakers [56]. We interact with the Alexa Skills Kit (ASK) through

**Table 1.** List of technologies and versions.

Technology	Version
Amazon web services	Accessed october 2018
Alexa skills kit (ASK)	
Dynamo DB	
Amazon S3	
Google natural language API	1.2 GA release
Python	3.6
NetworkX	2.2
NLTK	3.3
Ubuntu	18.04.1 LTS

**Fig. 4.** High-level view of the prototype.

a program written in the Python language, stored on Amazon S3 (Amazon Simple Storage Service), which is invoked by Amazon lambda functions when objects are created or when intents are triggered through user interactions. The complete conversation log generated during a session with a user is stored in Amazon Dynamo DB, which is Amazon’s fully-managed solution for NoSQL databases. The NetworkX library for Python serves to store and visualize the map. When the discussion ends, the visualization is emailed to the user together with a file containing a list of edges.

Google Natural Language API is queried extensively to find *entities*. Consider that the artificial facilitator asks “what causes obesity?” and the user responds “I believe that obesity is caused by an excess in eating and not enough exercise”. Google Natural Language API will extract the entities ‘obesity’, ‘excess’, ‘eating’, and ‘exercise’. Since an answer often includes a repetition of the subject, we automatically ignore user-provided entities that were part of the question. In this example, ‘obesity’ would be ignored, thus there are only three new concepts: ‘eating’, ‘excess’, and ‘exercise’. As detailed in Sect. 3, we must ensure that the concepts are not already used. When a new concept node is created, we use

WordNet (accessed via the NLTK library in Python) to retrieve all cognitive synonyms (i.e., *synsets*). If the user later mentions an entity that belongs to these synsets, the artificial facilitator points out that it already exists under a different name.

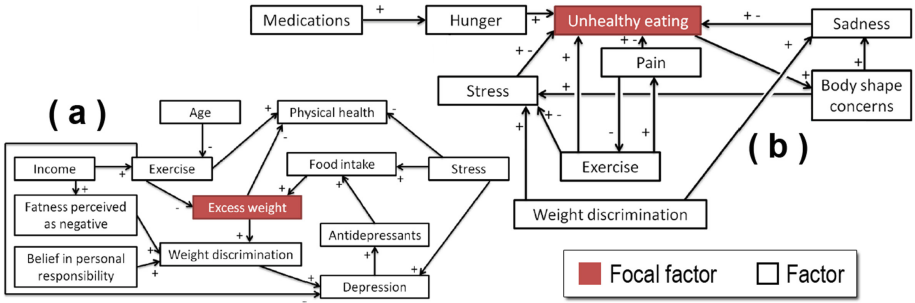
A causal map is not supposed to have unquantifiable concepts, but users may lose track of this requirement. If Google Natural Language API identifies an entity which is unquantifiable, then our application can use it in nonsensical questions. For instance, ‘excess’ was identified as an entity although it is unquantifiable. The application may continue by asking “what causes excess?”. We tested the application with 8 subjects over two months to identify such problematic entities. Since we cannot manually identify *all* such entities, we use the ones we identified as seeds to automatically fetch all similar entities, thus constituting a large dictionary of entities to ignore. The creation of this dictionary takes three steps performed using WordNet:

- (1) We have a set of entities, identified during testing as both (i) fetched by the Google Natural Language API and (ii) unquantifiable. For instance, consider {lack, bunch}.
- (2) For each word, we retrieve all its *hypernyms*, which are words with a broader meaning (e.g., color is a hypernym of red). Here, {lack, bunch} is transformed into {need, agglomeration, collection, cluster, gathering}.
- (3) For each hypernym, we retrieve all its *hyponyms*, which are more specific words (e.g., hyponyms of color would include red, blue, and green). In this example, {need, agglomeration, collection, cluster} would be expanded into a large set including {lack, necessity, urge, . . . , bunch, pair, trio, hive, crowd, agglomeration, batch, block, ensemble, . . . , population}.

Amazon Alexa development features were altered during the development of the artificial facilitator. Our initial implementation relied extensively on an intent (i.e., a template) known as *AMAZON.LITERAL*, which allowed for free-form speech input instead of a defined list of possible values. This slot was deprecated on October 22, 2018. Consequently, the implementation presented here relies on custom slots.

## 5 Case Study: Creating Obesity-Related Maps

We used three case studies to test our system. In the first two case studies, we verified whether a participant could (re)create a previously developed causal map when using our artificial facilitator (Fig. 5). Leveraging the broad variety of languages and accents supported by Alexa, we set the device to Indian English for these two cases, as it is the language spoken by our participant. In the third case, the device was set to American English, and we tested additional features such as detecting redundant concepts or allowing the user to correct the map. All case studies were performed using an Amazon Echo Dot Device version 618571720. We recorded the discussion and the resulting map



**Fig. 5.** Two previously published causal maps from health behaviors research [17]. Each map is centered on a different problem or ‘focal factor’. The original study added an intervention to these maps as part of a virtual trial.

that our artificial facilitator emailed to the participant. To provide full disclosure, our three recordings can be viewed at <https://www.youtube.com/playlist?list=PL7UTR3EL44zrkwrcDkiSwV-7kL0Nv6fQ5>

Our first two case studies demonstrated that the *structure* of the maps could correctly be created using our artificial facilitator. We observed three issues due to the automatic detection of entities. First, it can lead to significantly shorter concept labels (<https://www.youtube.com/watch?v=57tq0w4OEPw&t=324s>). The original map stated that weight discrimination was driven by excess weight, fatness perceived as negative, and a belief in personal responsibility. Our automatic process resulted in weight discrimination being driven by weight, fatness, and responsibility. This loses some nuances: it is not fatness in itself that leads to discrimination, but the societal belief that fatness is an undesirable trait. The problem is aggravated when concepts that should be different are shortened such that they are indistinguishable. For instance, ‘cardiovascular diseases’ and ‘metabolic diseases’ are very different medical situations. However, entity recognition sees both as ‘diseases’ and thus conflates them, which results in structural errors for the map. Second, entity recognition is a bottleneck of the application in terms of time: users can have to silently wait for several seconds before entities have been processed. These awkward silences disrupt the flow of the discussion. Finally, accents can lead to very different performances in terms of entity recognition. Results are not only different between Indian and American participants, but also among Americans (e.g., from the South or the Midwest). As noted by Rachael Tatman, the training dataset for smart speakers results in working “best for white, highly educated, upper-middle-class Americans, probably from the West Coast, because that’s the group that’s had access to the technology from the very beginning” [57].

The third case study demonstrated that additional features of our artificial facilitator worked as specified. For instance, the participant stated that over-eating was caused by over-indulgence, but these two concepts are considered interchangeable per WordNet. Consequently, the artificial facilitator informed

the user (<https://www.youtube.com/watch?v=U2mYkSLE9NE&t=40s>). We also confirmed that users were able to remove causes when they have been incorrectly captured (<https://www.youtube.com/watch?v=U2mYkSLE9NE&t=213s>). Finally, we verified that the virtual facilitator did repeat questions when prompted by the user (<https://www.youtube.com/watch?v=U2mYkSLE9NE&t=95s>).

## 6 Discussion

In collaborative modeling, participants externalize their mental models into various artifacts such as causal maps. This externalization can be guided by a trained facilitator, but there may be associated costs, and availability is limited. Alternatively, free software can be used at any time to create causal maps but they do not guide participants. In addition, neither facilitators nor current software can easily cope with larger causal maps, for instance, to avoid the creation of redundant concepts. To address these limitations, we designed an artificial facilitator that leverages voice-activated technologies. We implemented the prototype via Amazon Alexa, and demonstrated its features through three case studies.

As our system constitutes the first use of voice-activated technologies to build causal maps in participatory modeling, we are at the early stage of a multi-year process. There are several opportunities to improve the system or address additional research questions in the short- and medium-term. In the short-term, our prototype faces two limitations. First, we used hand-crafted rules, which is more in line with early spoken dialog systems than with current ones. Other approaches use generative methods (e.g., Bayesian networks) which often involve hand-crafted parameters, or discriminative methods where parameters are inferred by machine learning from the data. As stated by Henderson, “discriminative machine-learned methods are now the state-of-the-art in dialog state tracking” [55]. However, machine learning requires data to learn from. There is currently no corpus of model building involving a facilitator and one participant. Such sessions are often conducted with *many* participants, and the recordings are not released as the consent forms generally include an anonymity clause. Designing a better artificial facilitator will thus start by assembling a large set of recordings between a facilitator and a participant, for instance by modeling a system in which participants would be comfortable in publicly sharing their perspectives.

Second, our approach extensively relies on Alexa followed by Google Natural Language API to identify entities. Our prototype struggled with creating causal maps with specialized terms (e.g., from the medical domain) as Alexa could not identify them in speech and/or the API would not see them as relevant entities. The API may improve over time, and it may also be assisted with ontologies to identify (i) which specialized terms may be used, and (ii) which term is likely to be used following another one. Similarly, improvements in the API would reduce the processing time which currently results in many awkward seconds of silence. We note that improvements in the API or in Alexa Skill Kit

will automatically benefit the quality of our application, without changes in our design or implementation.

In the medium-term, research may explore how an artificial facilitator can provide guidance in aspects that are necessary yet challenging for trained facilitators. The structure of causal maps is normally analyzed *after they have been built*, for instance by identifying leverage points via centrality [30,31] or inventorying loops that drive the dynamics of the map [32–34]. However, a large map of a complex system that contains no loops may already be identified as problematic, suggesting that some causal edges are potentially missing. Consequently, the artificial facilitator can leverage network algorithms to analyze the structure of the map as it is built, thus informing participants of potential issues and approaches to address them. The artificial facilitator can also build on natural language processing in many ways that go beyond the identification of entities. Causal maps sometimes start with a brainstorming process, in which many concepts are generated and then *grouped*. Our artificial facilitator can use the semantic relatedness of concepts to inform the user about potential themes, which may result in combining several overly-detailed concepts into a more abstract category.

## 7 Conclusion

We successfully used Alexa to develop a voice-activated assistant that guides a user in creating a causal map. We addressed the challenge of finding appropriate concept names. In future work, we will automatically inform the user when concepts related to a theme may be used instead of narrowly defined concepts, and we will monitor the structure of the map as it is being built to support users in identifying loops.

### Supplementary Material

Our code is available at <https://github.com/datalab-science/causalMapBuilder>. Our three case studies (Sect. 5) as well as a video overview can be accessed at <https://www.youtube.com/playlist?list=PL7UTR3EL44zrkwrcDkiSwV-7kL0Nv6fQ5>

**Acknowledgments.** The authors are indebted to Mitacs Canada for providing the financial support which allowed TR to perform this research at Furman University, while mentored by PJG (local advisor) and VKM (home advisor). Publication costs are supported by an NSERC Discovery Grant for VKM.

**Contributions.** TR implemented the software and produced the videos demonstrating its functioning. PJG wrote the manuscript and designed the workflow of the virtual facilitator. TR was advised by PJG and VKM, who jointly initiated the study. All authors read and approved of this manuscript.

## References

1. Heitman, K.: Reductionism at the dawn of population health. In: El-Sayed, A.M., Galea, S. (eds.) *Systems Science and Population Health*, chap. 2, pp. 9–24. Oxford University Press (2017)
2. Fink, D.S., Keyes, K.M.: Wrong answers: when simple interpretations create complex problems. In: El-Sayed, A.M., Galea, S. (eds.) *Systems Science and Population Health*, chap. 3, pp. 25–36. Oxford University Press (2017)
3. Giabbanelli, P.J.: Analyzing the complexity of behavioural factors influencing weight in adults. In: Giabbanelli, P.J., Mago, V.K., Papageorgiou, E.I. (eds.) *Advanced Data Analytics in Health. SIST*, vol. 93, pp. 163–181. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-77911-9\\_10](https://doi.org/10.1007/978-3-319-77911-9_10)
4. Giabbanelli, P.J., Baniukiewicz, M.: Navigating complex systems for policymaking using simple software tools. In: Giabbanelli, P.J., Mago, V.K., Papageorgiou, E.I. (eds.) *Advanced Data Analytics in Health. SIST*, vol. 93, pp. 21–40. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-77911-9\\_2](https://doi.org/10.1007/978-3-319-77911-9_2)
5. Riley, B., et al.: Systems thinking and dissemination and implementation research. In: Brownson, R.C., Colditz, G.A., Proctor, E.K. (eds.) *Dissemination and Implementation Research in Health: Translating Science to Practice*, chap. 9, p. 143. Oxford University Press (2017)
6. Verigin, T., Giabbanelli, P.J., Davidsen, P.I.: Supporting a systems approach to healthy weight interventions in British Columbia by modeling weight and well-being. In: *Proceedings of the 49th Annual Simulation Symposium, Society for Computer Simulation International*, p. 9 (2016)
7. Dubé, L., et al.: From policy coherence to 21st century convergence: a whole-of-society paradigm of human and economic development. *Ann. N. Y. Acad. Sci.* **1331**(1), 201–215 (2014)
8. Lavin, E.A., et al.: Should we simulate mental models to assess whether they agree? In: *Proceedings of the Annual Simulation Symposium, Society for Computer Simulation International*, p. 6 (2018)
9. Gupta, V.K., Giabbanelli, P.J., Tawfik, A.A.: An online environment to compare students' and expert solutions to ill-structured problems. In: Zaphiris, P., Ioannou, A. (eds.) *LCT 2018. LNCS*, vol. 10925, pp. 286–307. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-91152-6\\_23](https://doi.org/10.1007/978-3-319-91152-6_23)
10. Giabbanelli, P.J., Tawfik, A.A.: Overcoming the PBL assessment challenge: design and development of the incremental thesaurus for assessing causal maps (ITACM). *Technol. Knowl. Learn.* **24**, 161–168 (2019)
11. Voinov, A., et al.: Tools and methods in participatory modeling: selecting the right tool for the job. *Environ. Model. Softw.* **109**, 232–255 (2018)
12. Giabbanelli, P.J., Crutzen, R.: Using agent-based models to develop public policy about food behaviours: future directions and recommendations. *Comput. Math. Methods Med.* **2017** (2017). <https://www.hindawi.com/journals/cmmm/2017/5742629/abs/>
13. Maglio, P.P., Mabry, P.L.: Agent-based models and systems science approaches to public health. *Am. J. Prev. Med.* **40**(3), 392–394 (2011)
14. Giabbanelli, P.J., Jackson, P.J.: Using visual analytics to support the integration of expert knowledge in the design of medical models and simulations. *Procedia Comput. Sci.* **51**, 755–764 (2015)
15. Xue, H., et al.: Applications of systems modelling in obesity research. *Obes. Rev.* **19**(9), 1293–1308 (2018)



16. de Pinho, H.: Generation of systems maps: mapping complex systems of population health. In: El-Sayed, A.M., Galea, S. (eds.) *Systems Science and Population Health*, chap. 6, pp. 61–76. Oxford University Press (2017)
17. Giabbanelli, P.J., Crutzen, R.: Creating groups with similar expected behavioural response in randomized controlled trials: a fuzzy cognitive map approach. *BMC Med. Res. Methodol.* **14**(1), 130 (2014)
18. Lich, K.H., et al.: Extending systems thinking in planning and evaluation using group concept mapping and system dynamics to tackle complex problems. *Eval. Program Plan.* **60**, 254–264 (2017)
19. Frerichs, L., et al.: Mind maps and network analysis to evaluate conceptualization of complex issues: a case example evaluating systems science workshops for childhood obesity prevention. *Eval. Program Plan.* **68**, 135–147 (2018)
20. Gray, S., et al.: The structure and function of angler mental models about fish population ecology: the influence of specialization and target species. *J. Outdoor Recreat. Tour.* **12**, 1–13 (2015)
21. Befus, D.R., et al.: A qualitative, systems thinking approach to study self-management in women with migraine. *Nurs. Res.* **67**(5), 395–403 (2018)
22. Yourkavitch, J., et al.: Interactions among poverty, gender, and health systems affect womens participation in services to prevent HIV transmission from mother to child: a causal loop analysis. *PloS One* **13**(5), e0197239 (2018)
23. Giabbanelli, P.J., Torsney-Weir, T., Mago, V.K.: A fuzzy cognitive map of the psychosocial determinants of obesity. *Appl. Soft Comput.* **12**(12), 3711–3724 (2012)
24. Zhang, S., de la Haye, K., Ji, M., An, R.: Applications of social network analysis to obesity: a systematic review. *Obes. Rev.* **19**(7), 976–988 (2018)
25. Finegood, D.T., Merth, T.D., Rutter, H.: Implications of the foresight obesity system map for solutions to childhood obesity. *Obesity* **18**(S1), S13–S16 (2010)
26. Jebb, S., Kopelman, P., Butland, B.: Executive summary: foresight tackling obesity: future choices project. *Obesity Reviews* **8**, vi–ix (2007)
27. Drasic, L., Giabbanelli, P.J.: Exploring the interactions between physical well-being, and obesity. *Can. J. Diabetes* **39**, S12–S13 (2015)
28. Allender, S., et al.: A community based systems diagram of obesity causes. *PLoS One* **10**(7), e0129683 (2015)
29. McGlashan, J., et al.: Comparing complex perspectives on obesity drivers: action-driven communities and evidence-oriented experts. *Obes. Sci. Pract.* **4**, 575–581 (2018)
30. McGlashan, J., et al.: Quantifying a systems map: network analysis of a childhood obesity causal loop diagram. *PloS One* **11**(10), e0165459 (2016)
31. Knapp, E.A., et al.: A network approach to understanding obesogenic environments for children in pennsylvania. *Connections* (02261766) **38**(1) (2018). <https://doi.org/10.21307/connections-2018-001>
32. Owen, B., et al.: Understanding a successful obesity prevention initiative in children under 5 from a systems perspective. *PloS One* **13**(3), e0195141 (2018)
33. Giabbanelli, P., et al.: developing technology to support policymakers in taking a systems science approach to obesity and well-being. *Obes. Rev.* **17**, 194–195 (2016)
34. Rwashana, A.S., et al.: Advancing the application of systems thinking in health: understanding the dynamics of neonatal mortality in uganda. *Health Res. Policy Syst.* **12**(1), 36 (2014)

35. Giabbanelli, P.J., Tawfik, A.A., Gupta, V.K.: Learning analytics to support teachers' assessment of problem solving: a novel application for machine learning and graph algorithms. In: Ifenthaler, D., Mah, D.-K., Yau, J.Y.-K. (eds.) *Utilizing Learning Analytics to Support Study Success*, pp. 175–199. Springer, Cham (2019). [https://doi.org/10.1007/978-3-319-64792-0\\_11](https://doi.org/10.1007/978-3-319-64792-0_11)
36. Singh, M., et al.: Building a cardiovascular disease predictive model using structural equation model & fuzzy cognitive map. In: 2016 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE), pp. 1377–1382. IEEE (2016)
37. Pillutla, V.S., Giabbanelli, P.J.: Iterative generation of insight from text collections through mutually reinforcing visualizations and fuzzy cognitive maps. *Appl. Soft Comput.* **76**, 459–472 (2019)
38. Pratt, S.F., et al.: Detecting unfolding crises with visual analytics and conceptual maps emerging phenomena and big data. In: 2013 IEEE International Conference on Intelligence and Security Informatics (ISI), pp. 200–205. IEEE (2013)
39. Ozawa, S., Paina, L., Qiu, M.: Exploring pathways for building trust in vaccination and strengthening health system resilience. *BMC Health Serv. Res.* **16**(7), 639 (2016)
40. Lukoianova, T., Rubin, V.L.: Veracity roadmap: is big data objective, truthful and credible? *Adv. Classif. Res. Online* **24**, 4 (2014)
41. Jetter, A.J.: Fuzzy cognitive maps for engineering and technology management: what works in practice? In: *PICMET* (2006)
42. Jordan, R., et al.: Twelve questions for the participatory modeling community. *Earth's Future* **6**(8), 1046–1057 (2018)
43. Giabbanelli, P.J., Adams, J., Pillutla, V.S.: Feasibility and framing of interventions based on public support: leveraging text analytics for policymakers. In: Meiselwitz, G. (ed.) *SCSM 2016. LNCS*, vol. 9742, pp. 188–200. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-39910-2\\_18](https://doi.org/10.1007/978-3-319-39910-2_18)
44. Binder, J.R., Desai, R.H.: The neurobiology of semantic memory. *Trends Cogn. Sci.* **15**(11), 527–536 (2011)
45. Lewis, P.: Rich picture building in the soft systems methodology. *Eur. J. Inf. Syst.* **1**(5), 351–360 (1992)
46. Sommerfeld, E., Sobik, F.: Operations on cognitive structures - their modeling on the basis of graph theory. In: Albert, D. (ed.) *Knowledge Structures*, pp. 151–196. Springer, Heidelberg (1994). [https://doi.org/10.1007/978-3-642-52064-8\\_5](https://doi.org/10.1007/978-3-642-52064-8_5)
47. Johnson-Laird, P.N.: Mental models and human reasoning. *Proc. Natl. Acad. Sci.* **107**(43), 18243–18250 (2010)
48. Mago, V.K., et al.: Analyzing the impact of social factors on homelessness: a fuzzy cognitive map approach. *BMC Med. Inform. Decis. Mak.* **13**(1), 94 (2013)
49. Mago, V.K., et al.: Fuzzy cognitive maps and cellular automata: an evolutionary approach for social systems modelling. *Appl. Soft Comput.* **12**(12), 3771–3784 (2012)
50. Axelrod, R.: Decision for neoimperialism: the deliberations of the British Eastern Committee in 1918. In: *The Cognitive Maps of Political Elites, Structure of Decisions*, chap. 4, pp. 77–95 (1976)
51. Axelrod, R.: Results. In: *The Cognitive Maps of Political Elites, Structure of Decisions*, chap. 9, pp. 221–250 (1976)
52. Ross, S.: Complexity and the presidency. In: *The Cognitive Maps of Political Elites, Structure of Decisions*, chap. 5, pp. 96–112 (1976)
53. Williams, J., Raux, A., Henderson, M.: The dialog state tracking challenge series: a review. *Dialogue Discourse* **7**(3), 4–33 (2016)

54. Laranjo, L., et al.: Conversational agents in healthcare: a systematic review. *J. Am. Med. Inform. Assoc.* **25**(9), 1248–1258 (2018)
55. Henderson, M.: Machine learning for dialog state tracking: a review. In: *Proceedings of The First International Workshop on Machine Learning in Spoken Language Processing* (2015)
56. Canalys: Amazon reclaims top spot in smart speaker market in Q3 2018 (2018)
57. Harwell, D.: The accent gap. *The Washington Post* (2018)