

SignalKG: Towards Reasoning about the Underlying Causes of Sensor Observations

Anj Simmons, Rajesh Vasa and Antonio Giardina

Applied Artificial Intelligence Institute, Deakin University, Geelong, Australia

Abstract

This paper demonstrates our vision for knowledge graphs that assist machines to reason about the cause of signals observed by sensors. We show how the approach allows for constructing smarter surveillance systems that reason about the most likely cause (e.g., an attacker breaking a window) of a signal rather than acting directly on the received signal without consideration for how it was produced.

Keywords

knowledge graph, ontology, sensor, surveillance

1. Introduction


Standards such as the Semantic Sensor Network (SSN/SOSA) ontology [1] allow capturing the semantics of sensor observations, and emerging standards for smart buildings [2, 3] and smart cities allow capturing the semantics of the environments in which sensors operate. However, reasoning about the underlying cause of sensor observations requires not only knowledge of the sensors and their environment, but also an understanding of the signals they detect and the possible causes of these signals. For example, inferring that the sound of breaking glass may be due to a broken window requires knowledge of the fact that glass windows produce a distinct sound when broken and that this sound propagates as sound waves through the air to a sensor such as a human ear or microphone. This paper proposes a signal knowledge graph (SignalKG) to support machines to reason about the underlying cause of sensor observations. Sensors and their environment are represented using existing standards, and then linked to SignalKG. To reason about the cause of sensor observations, we automatically generate a Bayesian network based on information in the knowledge graph, and use this to infer the posterior probability of causes given the sensor data.

Figure 1 presents an ontology describing the high level concepts. A category of *entities* (e.g., humans) perform *actions* (e.g., walking) that act on a type of object or place (e.g., in hallways), which in turn create a type of *signal* (e.g., the sound of footsteps). A *sensor* observes a particular type of signal, which usually reduces in strength with distance and can be distorted by surrounding objects (e.g., a wall between a sound source and the receiver attenuates the sound signal). The sensor may implement a *classifier* to detect the presence of the signal (e.g.,

ISWC'22: The 21st International Semantic Web Conference, October 23–27, 2022, Hangzhou, China

✉ a.simmons@deakin.edu.au (A. Simmons); rajesh.vasa@deakin.edu.au (R. Vasa); antonio.giardina@deakin.edu.au (A. Giardina)

ORCID 0000-0001-8402-2853 (A. Simmons); 0000-0003-4805-1467 (R. Vasa); 0000-0003-1047-6339 (A. Giardina)

 © 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)




Figure 1: High level overview of SignalKG ontology

an acoustic detector may make use of a binary machine learning classifier to detect if the sound of footsteps is present or not). A formal RDF/OWL specification of the SignalKG ontology is available online¹ as well as an *interactive demonstration*² of how SignalKG can be applied and used to reason about the underlying causes of sensor observations.

2. Related Work

Past research has utilised ontologies for the purpose of threat and situation awareness using description logic [4] and rule based reasoning [5]. However, these approaches assume that threats can be classified into classes according to deterministic rules, whereas in reality threats may be probabilistic in nature and there may be multiple possible explanations for a given set of observations. To overcome this limitation, reasoning methods have been proposed that combine deterministic rules with a Bayesian network for reasoning probabilistically [6]. However, the structure of the Bayesian network and associated probabilities need to be manually specified. In contrast, we seek to express this knowledge in a reusable and extensible form. There have been attempts to extend OWL to support representing/reasoning about uncertainty [7]. However, general approaches do not directly specify concepts for reasoning about sensor signals.

3. Demonstration

Figure 2 demonstrates how concepts in the SignalKG ontology can be applied to construct a knowledge graph of audio signals. Attackers and employees are entities that are capable of producing an action, such as breaking a window or dropping a tray of glasses in a dining room. Actions create signals, in this case, the sound of breaking glass or the sound of dropped glass. The building asset/room type on/in which an action can occur is represented using the RealEstateCore ontology [2]. To group similar signals together, we use the Simple Knowledge Organization System (SKOS) [8] ‘broader’ property to represent a signal hierarchy, for example, the sound of breaking glass is similar to the sound of dropped glass and thus are grouped under the same category. The knowledge graph also includes information about how signals propagate, for example, that sound intensity reduces with distance (according to an inverse square law) and is attenuated by walls. To model the case in which signals are classified on

¹SignalKG ontology: <https://signalkg.visualmodel.org/skg>

²Interactive demonstration available at: <https://signalkg.visualmodel.org>

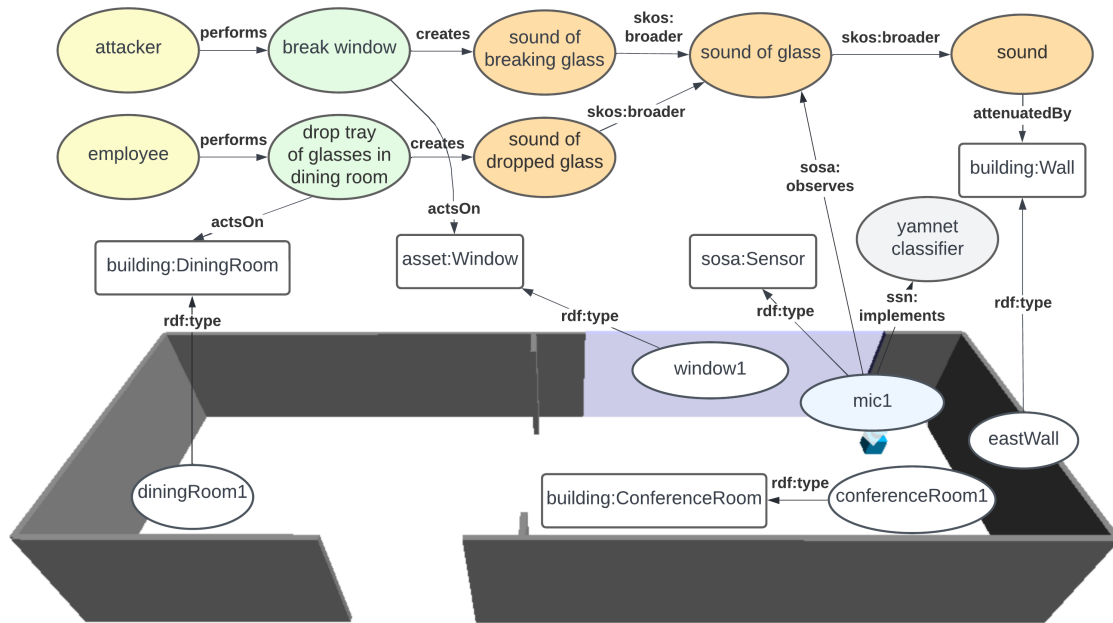


Figure 2: Knowledge graph of audio signals constructed using SignalKG ontology showing link to building rooms/assets and sensors

the sensing device, we allow for describing different classification models, such as YAMNet, an audio classifier that recognises 521 classes of sounds (e.g., glass).

Sensor observations can be linked to the signal knowledge graph via the property (signal) that the sensor observes. Our interactive demonstration includes a simulator to generate sample sensor observations (represented using the SSN/SOSA ontology) for hypothetical scenarios that could occur. The goal is to infer what took place given only the sensor observations, knowledge of the building and sensor placement, and our understanding of possible underlying causes of observed signals (specified using SignalKG).

To support reasoning probabilistically about causes of sensor observations, the knowledge graph also includes probabilities, such as the prior probability that an entity will be present, and the probability that an entity (if present) will perform an action. As our goal is to infer a cause given observations, it lends itself to Bayesian reasoning. Rather than manually specifying a Bayesian network for a particular scenario, we automatically generate one based on the information in the knowledge graph. Encoding all the information needed to reason about signals in the knowledge graph itself helps facilitate reuse and extension.

Nodes in the Bayesian network are generated for each entity, action at a location, signal emitted/received, and sensor. While our simple example results in a 1:1 mapping (shown in Figure 3), in more complex scenarios there may be vastly more nodes due to all the possible permutations of entity, action, location, signal and sensor. If multiple types of signals are present (audio, vision, social, etc.) then these all appear as part of the same generated Bayesian network, e.g., an attacker breaking a window will create both audio (sound of breaking glass) and vision (suspicious behaviour in a video feed) based signals. Prior probabilities for entities and actions need to be specified in the knowledge graph. The probability that a signal emitted by an action

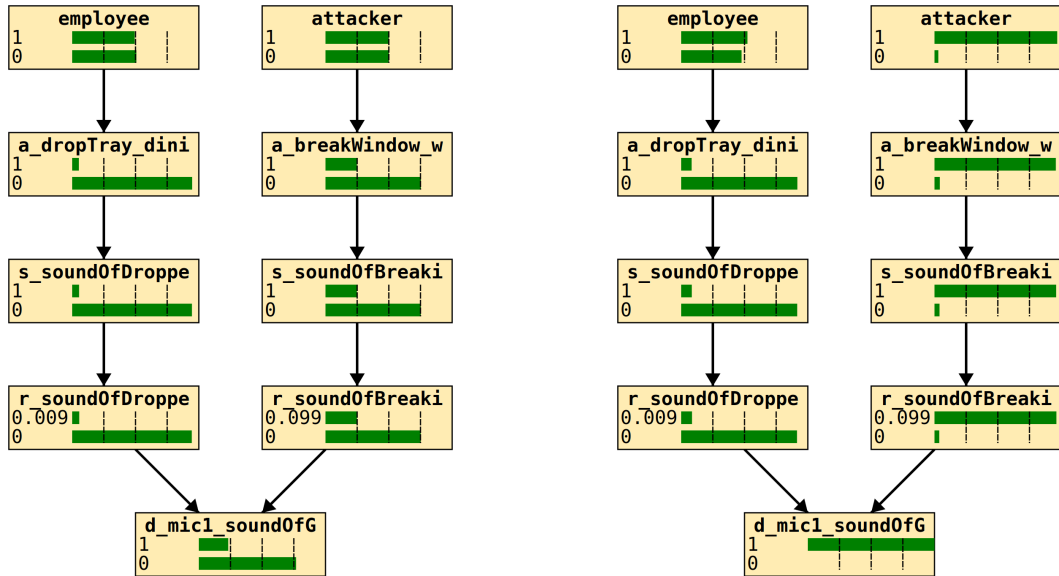


Figure 3: Bayesian network before (left) and after (right) conditioning on sensor observations. Green bars indicate the probability of each node value. (Key: a=action, s=signal emitted due to action, r=received signal strength at location of sensor, d=detected signal after classification). Visualisation created with jsbayes-viz [9].

will be detected by a sensor is calculated based on the distance from the location of the action to the location of the sensor, how the signal intensity reduces with distance (e.g., the knowledge graph may specify an inverse square law for sound signals), any barriers between the source and the sensor that may attenuate/block the signal (e.g., the knowledge graph may specify sound is attenuated by walls), and the sensitivity of the classifier used by the sensor to detect presence of a signal.

Once we have generated a Bayesian network, we can condition it on sensor observations to infer the posterior probability of the underlying cause. For the demonstration, we estimate the posterior probability via likelihood weight sampling, implemented by [9], drawing 20,000 samples (the number of samples to draw is a trade-off between accuracy and computation time). In the example, prior to conditioning on observations, there is a 50% chance of an attacker being present. After conditioning on the observation that the microphone has detected the sound of glass, the posterior probability of an attacker increases to 97%.

4. Next Steps

Even for the simple example of detecting building intrusions, the space of possible causes is large (e.g., an attacker could impersonate an employee then ask someone to let them in, or suspicious sounds could be due to a movie playing in the background). Furthermore, signals are not independent (as assumed by our preliminary prototype), but rather occur in sequences (e.g., the sound of footsteps, followed by a weapon detected in video footage, followed by a scream) that could help more reliably distinguish between possible causes. Also, more realistic models

of signal propagation are needed, which may require continuous probability distributions rather than a conditional probability table over a discrete set of values as in our example Bayesian network. To support efficient reasoning about these more complex scenarios, we plan to explore generation of probabilistic programs in place of the discrete Bayesian networks used in this paper.

A practical barrier to uptake of our approach is the need to specify prior probabilities for each action that an entity can perform. Ordinary behaviour could potentially be learned from data. However, modelling prior probabilities of actions intruders perform is more difficult, as attacks are rare events (limited data to learn from) and an adversary will adjust their actions to avoid detection. In future work, we plan to include the goals of the intruder as part of the knowledge graph, then use a game theoretic approach to determine probable actions they will take rather than manually specifying prior probabilities for each action.

Acknowledgments

This research was funded by National Intelligence Postdoctoral Grant NIPG-2021-006.

References

- [1] A. Haller, K. Janowicz, S. J. Cox, M. Lefrançois, K. Taylor, D. Le Phuoc, J. Lieberman, R. García-Castro, R. Atkinson, C. Stadler, The modular SSN ontology: A joint W3C and OGC standard specifying the semantics of sensors, observations, sampling, and actuation, *Semantic Web* 10 (2018) 9–32. doi:10.3233/SW-180320.
- [2] K. Hammar, E. O. Wallin, P. Karlberg, D. Hälleberg, The RealEstateCore Ontology, in: *The Semantic Web – ISWC 2019*, 2019, pp. 130–145. doi:10.1007/978-3-030-30796-7_9.
- [3] M. H. Rasmussen, M. Lefrançois, G. F. Schneider, P. Pauwels, BOT: the building topology ontology of the W3C linked building data group, *Semantic Web* 12 (2021) 143–161. doi:10.3233/SW-200385.
- [4] J. Roy, A. B. Guyard, Supporting threat analysis through description logic reasoning, in: *2012 IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support*, IEEE, 2012, pp. 308–315. doi:10.1109/CogSIMA.2012.6188401.
- [5] F. P. Pai, L. J. Yang, Y. C. Chung, Multi-layer ontology based information fusion for situation awareness, *Applied Intelligence* 46 (2017) 285–307. doi:10.1007/s10489-016-0834-7.
- [6] H. Yao, C. Han, F. Xu, Reasoning Methods of Unmanned Underwater Vehicle Situation Awareness Based on Ontology and Bayesian Network, *Complexity* 2022 (2022) 7143974. doi:10.1155/2022/7143974.
- [7] R. N. Carvalho, K. B. Laskey, P. C. Costa, PR-OWL – a language for defining probabilistic ontologies, *International Journal of Approximate Reasoning* 91 (2017) 56–79. doi:10.1016/j.ijar.2017.08.011.
- [8] A. Miles, S. Bechhofer, SKOS simple knowledge organization system reference, W3C recommendation (2009). URL: <https://www.w3.org/TR/skos-reference/>.
- [9] J. Vang, jsbayes-viz, 2016. URL: <https://github.com/vangj/jsbayes-viz/>.