Towards an Ensemble Approach for Sensor Data Sensemaking

Athanasios Tsitsipas[a]
Institute of Information Resource Management, Ulm University, Ulm, Germany

Keywords: Sensor Data Sensemaking, Hybrid, Stream Reasoning, Abduction, Shapeoids.

Abstract: In a world of uncertainty and incompleteness, one must "make sense" of found observations. Cyber-physical systems output large quantities of data, opening massive opportunities and challenges for scalable techniques to gain exciting insights. One intriguing challenge is the process of Sensor Data Sensemaking. The research presents an approach to handle this process by bringing together the strands of data and knowledge in a single architecture in an interpretable and expressive way. Differently from other works, the use of interpretable patterns from streaming data is in the spotlight. In addition, background knowledge over these patterns gasps the intention to give meaning to these patterns with several possible explanations. A hybrid implementation realises the approach following big data processing models.

1 INTRODUCTION

Imagine an infrastructure with various sensors deployed, witnessing a flux in the data from a specific sensor. There could be multiple applications and use cases treating this "phenomenon" as an anomaly or just the typical operation of this sensor. The importance of the flux is interchangeable in any of the two cases. A malfunction, or a complete failure of this sensor, can lead to internal issues or shortcomings for the system in place. Indirect Sensing (IS) is a compelling approach to treating such cases. Indirect Sensing (IS) is where single or different sensor compositions deliver the same information when the property in need is no longer attainable to direct sense. Therefore, other alternative sources may continuously provide streaming data that might contain the missing information implicitly. However, making sense (Cook, 2007) of sensor data to tackle this uncertainty and incompleteness in the environment is necessary for gaining exciting insights.

One approach that arises naturally for such use cases is classification as part of pattern recognition assigning observations to various classes. For example, in the context of IS, the authors in (Laput et al., 2017) use the streaming data as an input to a supervised learning algorithm and thus create a model that approximates the real-world's "image" to an understandable format for a machine. To enhance the robustness of such a model requires a substantial amount of training data transforming fundamental values that need to be categorised into features to calculate distances between them later. Such efforts fall under catch-phrases like Machine Learning or Deep Learning. However, both follow the same intuition to explore the input data for prediction, retrodiction, and imputation tasks. While this approach is prodigious for making comparisons between such items (i. e., data) and clustering them accordingly, they share to a degree the natural incentive of "making sense". However, recent studies have shown that making sense of sensor data should be more than just a classification task (Teijeiro and Félix, 2018; Evans et al., 2021).

Sensor Data Sensemaking (SDS) is a process in Human-Computer Interaction (HCI). Participants in field studies interpret and understand their environments and the behaviour of complex systems by "reading off" historical data using situated knowledge (Kurze et al., 2020). The human brain can distinguish or observe similarities in features from data as patterns (Neisser, 2014); use them as evidence which ought to explain by guessing the underlying process that caused such observations. Hence, the main research question is *how to handle IS under a machine-based SDS process in a human-centred manner*. The research assumes that making sense of sensor data in IS requires a theory under explicit symbol manipulation, most commonly referred to and understood as methods based on formal logic.

Additionally, it is similar to current trends in

[a] https://orcid.org/0000-0001-9122-6408

Tsitsipas, A.
Towards an Ensemble Approach for Sensor Data Sensemaking.
DOI: 10.5220/0010898800003122
In *Proceedings of the 11th International Conference on Pattern Recognition Applications and Methods (ICPRAM 2022)*, pages 323-329
ISBN: 978-989-758-549-4; ISSN: 2184-4313

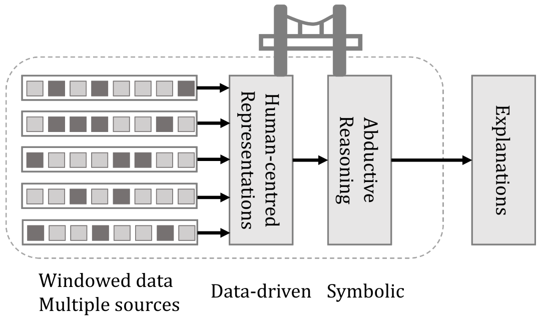

Figure 1: The ensemble approach bridges data-driven human-centred representations from streaming data and symbolic abductive reasoning.

modern Artificial Intelligence (AI), under its subfield Neuro-Symbolic AI, which focuses on merging the neural and symbolic regimes in AI research (Saker et al., 2021). The process of SDS is following modern applications of data work (Fischer et al., 2016; Fischer et al., 2017), which consists of three essential aspects laying down the proposed approach: (i) as part of human cognition requires a representation of objects, actions, numbers and space (Spelke and Kinzler, 2007), which persist and evolve, (ii) a key component in human commonsense is the construction of an explanation for the observed phenomenon (Marcus and Davis, 2019; Lake et al., 2017), (iii) finally, the research postulates that making sense of sensor data series is possible for someone to fathom it by speculating the presence of an unobserved cause, which accounts for and characterises the view of the sensor data (McCarthy, 2006; Inoue et al., 2009; Teijeiro and Félix, 2018). The occurred phenomenon consists of unobserved entities (i. e., environmental properties) that causally interact with the mechanics of the sensors to produce a dimensional footprint translated to the real-valued sensor reaction received as input. Thus, the approach to tackle the problem of IS using low-level processing to extract representations and high-level processes as innate abductive reasoning (Schurz, 2016) becomes ubiquitous to the underlying research areas in Structural Pattern Recognition, Logic and Stream Reasoning, following the field of Knowledge Representation and Reasoning (KRR).

The novelty and advance of the approach weights in an amalgamation of the following contributions. The first contribution is to develop and extract human-centred representations for mapping the time series sensor data to symbolic lexical constructs that constitute the primary objects for the following step. The second contribution is a declarative reasoning method via abduction, employing commonsense and domain knowledge to exploit the above lexical constructs from the sensor data as explanations for their occurrence. The approach uses probabilistic graphical models, which seamlessly bridge the worlds of probability and logic-based programming. Finally, architecture and implementation are imperative for an integrated stream reasoning framework for the SDS process facilitating the IS approach, forming a sensor fusion and a reasoning pipeline. Figure 1 illustrates a high-level view of the hybrid approach for the SDS process.

Section 2 presents a brief overview of the related work. Next, an enumeration of the focused research questions and the approach addressing them reside in Section 3 and Section 4, respectively. Finally, a brief overview of the main contributions in Section 5 and a short conclusion in Section 6 finalise the paper.

2 RELATED WORK

Sensemaking is an active area in HCI and has seen many studies on how people perceive and understand complex organisations of information (Fischer et al., 2016; Kurze et al., 2020). It actively involves the cognitive state of mind of Situational Awareness to analyse and assess dynamic problem environments. Automated methods for sensemaking support have been surveyed in (Llinas, 2014). The survey includes conceptual sensemaking models and analytic tool suites (ATS), primarily visualisation-based operations. The current research employs various techniques and methods to emulate the sensemaking process in HCI as part of an integrated framework. Observing structural primitives in time series data from a graph lends from Structural Pattern Recognition (Olszewski, 2001) and examing the hypothesis synthesis as an abductive reasoning process (Glass, 2019) formulates the final approach. In addition, the research includes a realisation of the sensemaking process over streaming sensor data. Therefore, each part contains individual challenges tackled in the respective research area. The reader has some pointers to the author's previous work (Tsitsipas and Schubert, 2021a; Tsitsipas et al., 2021).

Stream reasoning (Della Valle et al., 2009) is a relatively new field from almost the last decade. It aims to bridge the gap between stream processing and reasoning. Different communities cover aspects of it when the scale moves solely on each side. Data Stream Management Systems (DSMSs) (Golab and Özsu, 2003) cover the element of managing low-level, high-throughput data using continuous queries, but hardly ever they consider incomplete information. Complex Event Processing (CEP) (Cugola and Margara, 2012) approaches express high-level events in compositions of single incoming events. They use probabilistic methods (Tran and Davis, 2008; Wasserkrug et al., 2010; Skarlatidis et al., 2015) to handle uncertainty in incoming data. The research follows these systems, especially where Markov Logic Networks has a central role. However, these approaches usually operate over a static dataset and do not realise a dynamic scenario for inference over dynamic data. The issue is mainly because a Markov Logic Network requires all the available evidence before the inference. The research work applies a partial solution using the method of soft evidence (Jain and Beetz, 2010). Recently, a survey examined how Complex Event Recognition could jump to the Big Data era (Giatrakos et al., 2020). Although, in line with the work, their interests are around event forecasting and inductive logic programming (ILP).

Furthermore, in KRR, novel methods and approaches for representation and reasoning over streaming input exist but lack the scalability over the velocity of data. KRR yields the research behind non-monotonic logic and abductive reasoning as guiding methods for commonsense reasoning (Davis, 2017). Many authors are working on visual understanding, utilising commonsense knowledge for abductive high-level explanations (Suchan et al., 2019; Le-Phuoc et al., 2021). However, although they use neural models for finding the required representations, they do not use domain-independent semantics on the model-based processing or the data processing steps. Therefore, generality is not guaranteed by their approaches. Moreover, the authors in (Kate and Mooney, 2009) examine the application of abduction in Markov Logic Networks, with a recent application to root cause analysis in IT infrastructures (Schoenfisch et al., 2018). However, the application does not foresee a dynamic scenario; instead, it includes statically generated evidence for the failure events of the IT components.

Finally, an architecture for bridging the sense-reasoning gap for stream reasoning developed in the context of unmanned aerial vehicles (UAVs), named DyKnow(Heintz et al., 2010) and extended in a robot operating system for adaptively applied stream processing (De Leng and Heintz, 2016). The main separation factor between this pioneering work in the area and the current research work is the absence of generality expressing the SDS process. DyKnow operates on raw data using threshold-based static queries to extract the low-level abstractions for high-level cognitive functions for agent systems. There is a lack of human-centred data representations and the general concept of sensemaking (including reasoning processes, e. g., abduction).

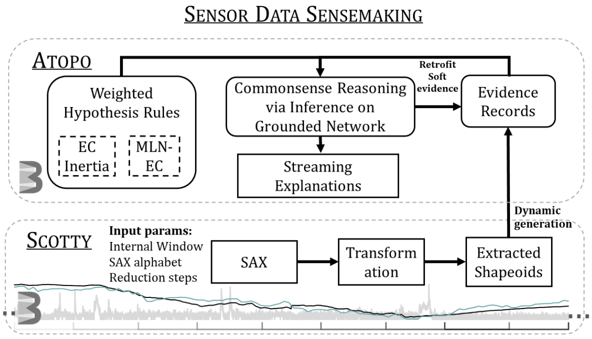

Figure 2: A diagram of the significant dataflows in the hybrid architecture realising the SDS process.

3 RESEARCH QUESTIONS

The work to alleviate the challenges formulated in the introduction poses the following research questions to separate the overall approach logically:

RQ1: Which symbolic descriptions (human-centred) from low-level sensor data the SDS process requires, and how to extract them?

RQ2: How to employ commonsense and domain knowledge to exploit the sensor data dimensional footprint of an unobserved cause as an explanation for its occurrence?

RQ3: How to scale the SDS process in an advanced stream reasoning framework?

4 RESEARCH APPROACH

The approach is driven by developing an expressive SDS artefact within the realm of stream reasoning. As such, it explores the interplay between a data-driven human-centred primitive pattern extraction and a knowledge-driven inference method. They support both a quantitative view on symbolic representations from streaming data and a qualitative view on complex interrelations of properties and features of an IS task. A principal approach conciliates the gap between sensing and reasoning in a streaming setting in research work. Given the magnitude and complexity of the problem, addressing each issue and challenge in the respective field is essential to finally move towards the hybrid architecture of the approach.

Lexical Sensor Observation (LSO). An investigation of techniques and methods from time series representation and more concretely focus on symbolic representations accommodates the approach. The study of subsymbolic and neural processes are not foreseeable in work. Whereas they contain powerful tools for exploring large amounts of data, the resulting feature space is in a numerical representation that needs further interpretation to translate to human-centred encodings. The emphasis in work is on how a human can observe and describe a time series sequence by locking to some points of interest to perceive a general shape or pattern (Agrawal et al., 1995; Bakshi and Stephanopoulos, 1995). The exciting area of shape-based extraction (Siddiqui et al., 2020) acts as an inspiration to describe a time series of sensor perturbations in natural language constructs. The extracted low-level representations from the sensor data series act as interpretable LSOs and can compose more complex shapes expressing their relationships via sequential operators.

Declarative Hypothesis Model (DHM). Formal logic is the only epistemically sufficient representation (Davis, 2017) to encode commonsense and domain knowledge in a declarative and expressive manner. The process of sensemaking is naturally an abductive reasoning process (non-monotonic). Initially, there is inherent uncertainty in the sensor data, the pattern extraction, and the different conjectures (i. e., DHM) created during the SDS process. A prompt hypothesis should be retracted when more data is available or a non-perfect logical rule as a hypothesis. In addition, a formalism to guide the process of reasoning is required, as the semantics and the denotations of logic needed for the hypothesis encoding should be domain-independent.

From Hypotheses to Explanations (HE). The research approach uses as commonsense knowledge the various encoded hypotheses qua alternatives for an IS task. They are expressed adequately in the form of logical language. The evaluation of the hypotheses rules is implemented as an approximately valid inference concerning the logic. The inference techniques in the respective logic are necessary to evaluate the correctness of the SDS process.

A Hybrid Architecture. As a final step, the research work provides a realisation of the approach to streaming sensor data. A novel and unified data processing programming model implements the SDS process capable of high-level abduction over streaming sensor data. A qualitative evaluation of the overall research outcome against a set of requirements for principal approaches attempting to bridge the gap between sensing and reasoning (Heintz, 2009) is presumptive for its resilience.

5 CONTRIBUTIONS

The following contributions realise the approach and thus provide evidence to counter the constituted questions. Figure 2 acts as an accompanying illustration, with more elaborate descriptions of the contribution's internal steps.

5.1 Human-centred Representation for Data-driven Patterns

The process of SDS requires low-level symbolic representations of sensor data as patterns that are human-understandable (e. g., in natural language). The approach offers the time series representation in a "morphable" manner by proposing lexical shape-like primitives, named *shapeoids* (Tsitsipas and Schubert, 2021b). It utilises the Symbolic Aggregate Approximation (SAX) (Lin et al., 2003) method. On top of SAX, developing a string-based algorithm for extracting this set of primitive data-driven lexical constructs constitutes the creation of the *shapeoids*. A novel framework named SCOTTY (Tsitsipas et al., 2021) encompasses this functionality implemented in Java and based on an open-source implementation[1] of SAX. For the sake of conciseness, the reader may regard the *shapeoids*: (i) An ANGLE is a gradual and continuous increase or decrease, (ii) A HOP describes a distinctive phase shift leading to an overall visible change, (iii) The HORN is a temporary effect which fades quickly, as the end of the pattern returns almost to the initial point, and (iv) The FLAT is an almost sturdy line with small variations in the curve. The *shapeoids* offer a comprehensive view of the time series data. SCOTTY can process around a million single data points, in not more than six seconds, with a time and space complexity of *O(N)*.

Additionally, its evaluation has shown interesting capabilities in time series representation in pattern recognition. During the evaluation, a declarative classifier was built, based on SCOTTY, competing with state-of-the-art algorithms in the area. The evaluation uses the Cylinder-Bell-Funnel dataset from the UCR archive (Dau and Keogh et al., 2018). For the SDS process, SCOTTY should run over a sliding window with a parameterised configuration on the window size to overcome the limitation of the underlying method of SAX, which performs in an internal step a statistical computation, requiring bounded data.

[1] https://github.com/jMotif/SAX

5.2 Expressive High-level Abductive Reasoning

Rooted in Markov Logic Networks (MLN) (Richardson and Domingos, 2006), the framework named ATOPO (Tsitsipas and Schubert, 2021a) is designed for integration as a reasoning engine for the hybrid architecture of the SDS. It leverages the available commonsense knowledge to express the alternative hypotheses for the task of IS. It employs axioms from the Event Calculus (EC) (Kowalski and Sergot, 1989; Mueller, 2014) to enforce the persistence of objects whose value change over time by keeping them as hard-constrained rules in the resulting Markov Network. MLN combines the expressiveness of first-order logic and uses concepts from probability theory to tackle uncertainty. The knowledge engineering process for the SDS consists of definite clauses in first-order logic with assigned weights in front of them indicating their probability degree. The manual assignment of weight is part of the sensemaking process. It is an aspect of the commonsense reasoning stating if they support (positive weight) or penalise (negative weight) worlds in which the corresponding hypotheses are satisfied instead of classical logic, where all the statements are hard-constrained. The EC inertia laws must always remain as hard constraints, because otherwise during the inference process, the hypotheses rules' probabilities will eventually converge to be equiprobable (Skarlatidis et al., 2015).

The implementation of commonsense reasoning is a valid inference (Davis, 2017) within MLN, calculating the joint probability distribution of the grounded Markov Network, using approximation techniques because the direct computation of the formula is intractable for large networks. ATOPO uses MC-SAT to output the marginal probabilities. As mentioned above, the knowledge base contains factual predicates about the ground truth and the hypotheses rules encoded with domain-independent semantics. A hypothesis is a clausal rule with the EC predicate *InitiatedAt* as the head. Its body contains *Happens* predicates with the recognised *shapeoid* from a sensor and other contextual constraints (e. g., sensor type, suitable location). The queried fluent in the EC predicate *HoldsAt*, is a possible quantification over verifying an explanation for LSO findings from the raw sensor data. Finally, ATOPO uses an open-source implementation of Markov Logic Networks, named LoMRF (Skarlatidis and Michelioudakis, 2014) and implemented in Scala, which is also its development language. To support the non-monotonic semantics of EC, LoMRF performs circumscription via predicate completion (Skarlatidis et al., 2015).

The evaluation of ATOPO in a situational awareness use case with IS, utilises a real-world dataset (Birnbach et al., 2019). The authors collected various sensor data while performing different activities. In the context of its evaluation ATOPO ran in a static setting, requiring manual encodings for the internal representations in the knowledge base and the evidence for recognising an opened window from the data using five in total surrounding sensors (e. g., temperature, air quality). The evaluation showed promising results even if selecting a rigorous setting to calculate the performance measures, referencing the timestamps of the ground truth (i. e., when the window's status is open or closed, based on a contact sensor). While engaging the SDS process, the visible dimensional footprint of the action (e. g., opening the window) eventuates with a delay as it takes some time until the opened window affects, for example, the temperature sensor sufficiently.

5.3 An Integrated Framework for SDS

The final contribution is the integrated framework implemented on top of Apache Beam[2]. It provides a unified programming model, offering common elements of data processing frameworks (e. g., windowing, transform functions), supporting batch and stream processing and running on many execution engines. Hence, the solution's potential is scalable to clustered data processing platforms (e. g., Apache Flink, Apache Spark). The architecture enables the SDS in a setting under streaming use cases. Apache Beam allows SCOTTY to run over multivariate sensor data in a streaming setting. Additionally, it retrofits ATOPO in a reasoning pipeline, executing deterministic runs of inference on the grounded Markov Networks over a sliding window. It supports incremental reasoning, as the previous states and inference results propagate as evidence and maximum a posteriori elements, respectively, to the following sliding window. Finally, the research work provides evidence for the compliance of the hybrid implementation by complying with the requirements of a knowledge processing middleware (Heintz, 2009).

6 CONCLUSION

The paper presents an "ensemble" approach for SDS. It extracts human-centred data-driven primitives from time series data, representing people's mental mod-

[2] https://beam.apache.org/

els in an SDS process. A knowledge-driven method based on theoretical formalism supports the reasoning steps for expressing alternative hypotheses that explain the extracted observations in an IS setting. The individual parts of the work have been published in various conferences on Artificial Intelligence, Big Data and Pattern Recognition. In addition, the solution includes a prototype implementation with a hybrid architecture under modern data processing frameworks. SDS constitutes a significant challenge to solve, as many state-of-the-art research fields have the same incentive under different settings: how to bridge data with knowledge. The outcome of the work comes to the fore to explain the occurrence of primitive structures in time series data by assuming an underlying reality that triggered them.

Furthermore, when submitting this paper to the conference, the open position of evaluating the primitive structures for the SDS process in other domains, where a sensor is not physical but software-based virtual (Martin et al., 2021), was submitted and accepted in parallel to another conference (Tsitsipas et al., 2022). Finally, the current research work opens the field on the sparsity of other different primitive data structures for domain-dependent scenarios and how the SDS process can be realised in a neuro-symbolic approach while maintaining a human-centred manner.

ACKNOWLEDGEMENTS

Germany's Federal Ministry of Education and Research (BMBF) has partially funded the research under the project HorME.

REFERENCES

Agrawal, R., Psaila, G., Wimmers, E. L., and Zaït, M. (1995). Querying shapes of histories. In *Proceedings of the 21th International Conference on Very Large Data Bases*, pages 502–514.

Bakshi, B. R. and Stephanopoulos, G. (1995). Reasoning in time: Modeling, analysis, and pattern recognition of temporal process trends. *Advances in Chemical Engineering*, 22(C):485–548.

Birnbach, S., Eberz, S., and Martinovic, I. (2019). Peeves: Physical event verification in smart homes. In *Proceedings of the 2019 ACM Conference on Computer and Communications Security*. ACM.

Cook, D. J. (2007). Making sense of sensor data. *IEEE Pervasive Computing*, 6(2):105–108.

Cugola, G. and Margara, A. (2012). Processing flows of information: From data stream to complex event processing. *ACM Computing Surveys (CSUR)*, 44(3):1–62.

Dau, H. A. and Keogh et al., E. (2018). The ucr time series classification archive.

Davis, E. (2017). Logical formalizations of commonsense reasoning: a survey. *Journal of Artificial Intelligence Research*, 59:651–723.

De Leng, D. and Heintz, F. (2016). Dyknow: A dynamically reconfigurable stream reasoning framework as an extension to the robot operating system. In *2016 IEEE International Conference on Simulation, Modeling, and Programming for Autonomous Robots (SIMPAR)*, pages 55–60. IEEE.

Della Valle, E., Ceri, S., Van Harmelen, F., and Fensel, D. (2009). It's a streaming world! reasoning upon rapidly changing information. *IEEE Intelligent Systems*, 24(6):83–89.

Evans, R., Hernández-Orallo, J., Welbl, J., Kohli, P., and Sergot, M. (2021). Making sense of sensory input. *Artificial Intelligence*, 293:103438.

Fischer, J. E., Crabtree, A., Colley, J. A., Rodden, T., and Costanza, E. (2017). Data work: how energy advisors and clients make iot data accountable. *Computer Supported Cooperative Work (CSCW)*, 26(4):597–626.

Fischer, J. E., Crabtree, A., Rodden, T., Colley, J. A., Costanza, E., Jewell, M. O., and Ramchurn, S. D. (2016). " just whack it on until it gets hot" working with iot data in the home. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 5933–5944.

Giatrakos, N., Alevizos, E., Artikis, A., Deligiannakis, A., and Garofalakis, M. (2020). Complex event recognition in the big data era: a survey. *The VLDB Journal*, 29(1):313–352.

Glass, D. H. (2019). Competing hypotheses and abductive inference. *Annals of Mathematics and Artificial Intelligence*, pages 1–18.

Golab, L. and Özsu, M. T. (2003). Issues in data stream management. *ACM Sigmod Record*, 32(2):5–14.

Heintz, F. (2009). *DyKnow: A stream-based knowledge processing middleware framework*. PhD thesis, Linköping University Electronic Press.

Heintz, F., Kvarnström, J., and Doherty, P. (2010). Bridging the sense-reasoning gap: Dyknow–stream-based middleware for knowledge processing. *Advanced Engineering Informatics*, 24(1):14–26.

Inoue, K., Furukawa, K., Kobayashi, I., and Nabeshima, H. (2009). Discovering rules by meta-level abduction. In *International Conference on Inductive Logic Programming*, pages 49–64. Springer.

Jain, D. and Beetz, M. (2010). Soft evidential update via markov chain monte carlo inference. In *Annual Conference on Artificial Intelligence*, pages 280–290. Springer.

Kate, R. J. and Mooney, R. J. (2009). Probabilistic abduction using markov logic networks. In *In: IJCAI-09 Workshop on Plan, Activity, and Intent Recognition*.

Kowalski, R. and Sergot, M. (1989). A logic-based calculus of events. In *Foundations of knowledge base management*, pages 23–55. Springer.

Kurze, A., Bischof, A., Totzauer, S., Storz, M., Eibl, M., Brereton, M., and Berger, A. (2020). Guess the data: Data work to understand how people make sense of and use simple sensor data from homes. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–12.

Lake, B. M., Ullman, T. D., Tenenbaum, J. B., and Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and brain sciences*, 40.

Laput, G., Zhang, Y., and Harrison, C. (2017). Synthetic sensors: Towards general-purpose sensing. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 3986–3999.

Le-Phuoc, D., Eiter, T., and Le-Tuan, A. (2021). A scalable reasoning and learning approach for neural-symbolic stream fusion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 4996–5005.

Lin, J., Keogh, E., Lonardi, S., and Chiu, B. (2003). A symbolic representation of time series, with implications for streaming algorithms. In *Proceedings of the 8th ACM SIGMOD workshop on Research issues in data mining and knowledge discovery*, pages 2–11.

Llinas, J. (2014). A survey of automated methods for sensemaking support. In *Next-Generation Analyst II*, volume 9122, page 912206. International Society for Optics and Photonics.

Marcus, G. and Davis, E. (2019). *Rebooting AI: Building artificial intelligence we can trust*. Vintage.

Martin, D., Kühl, N., and Satzger, G. (2021). Virtual sensors. *Business & Information Systems Engineering*, 63(3):315–323.

McCarthy, J. (2006). Challenges to machine learning: Relations between reality and appearance. In *International Conference on Inductive Logic Programming*, pages 2–9. Springer.

Mueller, E. T. (2014). *Commonsense reasoning: an event calculus based approach*. Morgan Kaufmann.

Neisser, U. (2014). *Cognitive psychology: Classic edition*. Psychology Press.

Olszewski, R. T. (2001). *Generalized feature extraction for structural pattern recognition in time-series data*. Carnegie Mellon University.

Richardson, M. and Domingos, P. (2006). Markov logic networks. *Machine learning*, 62(1-2):107–136.

Saker, M. K., Zhou, L., Eberhart, A., and Hitzler, P. (2021). Neuro-symbolic artificial intelligence. *AI Communications*, pages 1–13.

Schoenfisch, J., Meilicke, C., von Stülpnagel, J., Ortmann, J., and Stuckenschmidt, H. (2018). Root cause analysis in it infrastructures using ontologies and abduction in markov logic networks. *Information Systems*, 74:103–116.

Schurz, G. (2016). Common cause abduction: The formation of theoretical concepts and models in science. *Logic Journal of the IGPL*, 24(4):494–509.

Siddiqui, T., Luh, P., Wang, Z., Karahalios, K., and Parameswaran, A. (2020). Shapesearch: A flexible and efficient system for shape-based exploration of trendlines. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, pages 51–65.

Skarlatidis, A. and Michelioudakis, E. (2014). Logical Markov Random Fields (LoMRF): an open-source implementation of Markov Logic Networks.

Skarlatidis, A., Paliouras, G., Artikis, A., and Vouros, G. A. (2015). Probabilistic event calculus for event recognition. *ACM Transactions on Computational Logic (TOCL)*, 16(2):1–37.

Spelke, E. S. and Kinzler, K. D. (2007). Core knowledge. *Developmental science*, 10(1):89–96.

Suchan, J., Bhatt, M., and Varadarajan, S. (2019). Out of sight but not out of mind: An answer set programming based online abduction framework for visual sensemaking in autonomous driving. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 1879–1885. ijcai.org.

Teijeiro, T. and Félix, P. (2018). On the adoption of abductive reasoning for time series interpretation. *Artificial Intelligence*, 262:163–188.

Tran, S. D. and Davis, L. S. (2008). Event modeling and recognition using markov logic networks. In *European Conference on Computer Vision*, pages 610–623. Springer.

Tsitsipas, A., Eisenhart, G., Seybold, D., and Wesner, S. (2022). Scalable shapeoid recognition on multivariate data streams with apache beam. In *Proceedings of the Computing Conference 2022, forthcoming*. Springer.

Tsitsipas, A., Schiessle, P., and Schubert, L. (2021). Scotty: Fast a priori structure-based extraction from time series. In *2021 IEEE International Conference on Big Data (BigData 2021), forthcoming*.

Tsitsipas, A. and Schubert, L. (2021a). Modelling and reasoning for indirect sensing over discrete-time via markov logic networks. In *Proceedings of the Twelfth International Workshop Modelling and Reasoning in Context (MRC 2021)*. CEUR-WS.org.

Tsitsipas, A. and Schubert, L. (2021b). On group theory and interpretable time series primitives. In *Proceedings of the Seventeenth International Conference on Advanced Data Mining and Applications (ADMA 2021), forthcoming*. Springer.

Wasserkrug, S., Gal, A., Etzion, O., and Turchin, Y. (2010). Efficient processing of uncertain events in rule-based systems. *IEEE Transactions on Knowledge and Data Engineering*, 24(1):45–58.