



HAL
open science

Towards Stream-based Reasoning and Machine Learning for IoT Applications

Markus Endler, Jean-Pierre Briot, Vitor P. Almeida, Francisco Silva E Silva,
Edward H Haeusler

► **To cite this version:**

Markus Endler, Jean-Pierre Briot, Vitor P. Almeida, Francisco Silva E Silva, Edward H Haeusler.
Towards Stream-based Reasoning and Machine Learning for IoT Applications. 2017. hal-01433729v1

HAL Id: hal-01433729

<https://hal.sorbonne-universite.fr/hal-01433729v1>

Preprint submitted on 27 Jan 2017 (v1), last revised 13 Feb 2018 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Towards Stream-based Reasoning and Machine Learning for IoT Applications

Markus Endler

Department of Informatics, PUC-Rio
Rio de Janeiro, Brazil
Email: endler@inf.puc-rio.br

Jean-Pierre Briot

Laboratoire d'Informatique de Paris 6
UPMC-CNRS, France
& Department of Informatics, PUC-Rio
Rio de Janeiro, Brazil
Email: Jean-Pierre.Briot@lip6.fr

Vitor P. de Almeida

Department of Informatics, PUC-Rio
Rio de Janeiro, Brazil
Email: valmeida@inf.puc-rio.br

Francisco Silva e Silva

LSDi, Universidade Federal do Maranhão
São Luis, Brazil
Email: fssilva@lsdi.ufma.br

Edward H. Haeusler

Department of Informatics, PUC-Rio
Rio de Janeiro, Brazil
Email: hermann@inf.puc-rio.br

Abstract—As distributed IoT applications become larger and more complex, the pure processing of raw sensor and actuation data streams becomes impractical. Instead, data streams must be fused into tangible facts and these pieces of information must be combined with a background knowledge to infer new pieces of knowledge. And since many IoT applications require almost real-time reactivity to stimulus of the environment, such information inference process has to be performed in a continuous, on-line manner. This paper proposes a new semantic model for data stream processing and real-time reasoning based on the concepts of *Semantic Stream* and *Fact Stream*, as a natural extensions of Complex Event Processing (CEP) and RDF (graph-based knowledge model). The main advantages of our approach are that: (a) it considers time as a key relation between pieces of information; (b) the processing of streams can be implemented using CEP and that (c) it is general enough to be applied to any Data Stream Management System (DSMS). Last, we will present challenges and prospects on using machine learning and induction algorithms to learn abstractions and reasoning rules from a continuous data stream.

Keywords—*Internet of Things (IoT); sensors; data*

streams; complex event processing (CEP); semantic reasoning; inference; machine learning.

I. INTRODUCTION

Several complex IoT applications, such as manufacturing industry, transportation systems, and healthcare, put hard real time requirements on the acquisition and processing of sensor data for identifying situations and extracting information from systems' operations and its environment. These typically require on-line processing of continuous streams of sensor data (Data Stream Processing), sensor fusion techniques, pattern recognition and timely and autonomous systems control. There are already many well-established algorithms for clustering, classification, regression and sequence-analysis.

However, so far in current IoT systems, sensing and actuation is mostly done at the *bare bones* data level, whereas many IoT applications demand higher level situation awareness of – and reasoning about – the systems' states and the physical envi-

ronment where they operate. For this to be possible, it is necessary to have comprehensive semantic models for data stream analysis and actuation. Semantic models are formally defined concepts and relations on which reasoning engines can operate to derive new bits of information and knowledge about a system and its environment. The main problem is that current semantic models (designed for the Semantic Web) are not suitable for efficient and real-time reasoning. Current data analysis for IoT systems is either done off-line or lacks any semantic-based reasoning.

For example, consider a production plant in the near future, where several – mobile or stationary – robots operate in a product assembly and interact with each other to hand over parts and tools of the assembly line. Suddenly, there is a short power outage and the assembly line stops for a few seconds, so that some robots go back to their consistent initial states, while others continue their activity (e.g. on battery power) and only stop when their sensors notice that the production line is not advancing. In this case, the robots have to “understand” what has happened, and have to “know” which of the machinery (and robots) are in which state when activity is resumed, as well as the assembly stage of items being produced. And like magic, only a few seconds after energy is back, the robots synchronize with each other, identify missed steps in the assembly process of each item, and resume cooperating again. Such knowledge and understanding is only possible because all robots have not only a semantic model of their own state, but also situational awareness, i.e. a comprehensive model of the production process as a whole and their role in the entire process. The semantic model furthermore describes possible localized and global problems of the entire production process, as well as individual and specific actuation plans for some situations. As all possible situations cannot be represented in a model, the robots have to classify features, combine situational patterns and combine parts of specific action plans. In the aforementioned IoT scenario, the robots would be capable of such fast recovery of the manufacturing process because their situational understanding (i.e. semantic-centered inference/reasoning process) is executed very fast, with almost no delay, as soon as each robot’s operational capability is back.

With the goal of finding a suitable semantic model for IoT, this paper proposes a novel approach for real-time symbolic reasoning based on the concepts of *Semantic Stream and Fact Stream*, as natural extensions of Complex Event Processing (CEP) [18] and RDF (graph-based knowledge model) [10]. The main advantages of our approach are that: (a) it uses the timestamp and co-location information to correlate actions/events happening at different real-world entities (i.e. objects and subjects), (b) the online processing of semantic streams can be implemented using conventional CEP technology and semantic reasoning approaches, but (c) using ontology-based reasoning over a knowledge base it is possible not only to deduce future or indirect events that would not be detected through CEP, but also to generate new CEP rules for the stream analysis; (d) the approach is generic enough to be applied to many Data Stream Management Systems (DSMS). This research is being carried out in the scope of the ESMOCYP cooperation project between PUC-Rio, Federal University of Maranhão and University of Stuttgart. We are currently developing a prototype of the semantic stream reasoning using, ContextNet, our distributed and scalable middleware for the Internet of Mobile Things [21]. It is a mobile-cloud architecture where several interconnected CEP agents can be deployed both in a cloud/cluster [5], as well as on Android mobile devices [20].

The paper is structured as follows. In Section II, we explain the basic concepts of Complex Event Processing and list some common approaches for modeling knowledge and performing reasoning. Section III explains the two steps of semantic stream reasoning. In Section IV, we present a scenario to explain how our reasoning process would be performed using temperature and accelerometer sensors embedded into vehicles, houses and in the street. Then we present related work and in Section VI discuss the benefits of our approach and prospects, with in Section VII, a specific first analysis of how machine learning and induction algorithms could help in automatically or semi-automatically constructing or extracting stream analysis patterns and rules. Section VIII then concludes the paper.

II. FUNDAMENTALS

A. Complex Event Processing

Complex Event Processing (CEP) [18] provides a rich set of concepts and operators for processing events, which include the CQL-like (Continuous Query Language) [4] queries, rules, primitive functions (aggregation, filtering, transformation, etc.) and production of derived events. A CEP workflow continuously processes incoming events, analyses and manipulates them, and outputs derived events that are delivered to event consumers (*e.g.* monitoring applications). These output usually represent notifications about detected situations of interest to the the applications.

The processing of events is described by CEP rules, which are Event-Condition-Actions that combine continuous query primitives with context operators (*e.g.* temporal, logical, quantifiers) on received events, checking for correlations among these events, and generating complex (or composite) events that summarize the correlation of the input events. For example, a split rule takes an input event and creates a set of events, while a filter rule only outputs events that satisfy a given criteria. Rules can also operate on a collection of events, for example, an aggregate rule outputs a single event by executing a function on the grouped events, while a join transformation tries to correlate events from various data streams. Another important concept in CEP is that of sliding time and event windows. A time window is a temporal context that subdivides the stream of events into intervals, where CEP rules and operators are applied only to the events within each window. CEP supports three sorts of windows: Landmark, Sliding, and Fading, the latter being a sliding window where a decay factor λ is applied to the events according to their age. *i.e.* more recent events have higher importance than older events. Most CEP systems have the concept of Event Processing Agents (EPAs), which are software modules that implement one transformation within the event processing workflow. The type of an EPA is defined by the rules it implements, such as filtering, counting or specific event pattern detection. Note that rules are hand written by experts. We will address in Section VII, a preliminary analysis of how machine learning and induction algorithms could help in, automatically or semi-automatically, constructing rules as well

as extracting patterns.

B. Knowledge Representation and Reasoning Approaches

There are plenty of Semantic Models that represent knowledge about a system and its environment, but almost all of them have problems of scale (*i.e.* the reasoning has high computational complexity), and thus are not suitable for real-time reasoning. The main semantic approaches are (see, *e.g.*, for a survey and comparison in [19]):

- **Frame Based Models:** A frame is an artificial intelligence data structure used to divide knowledge into substructures by representing "stereotyped situations." They are used in artificial intelligence Frame languages.
- **Conceptual Graphs:** are a logical formalism that includes classes, relations, individuals and quantifiers. This formalism is based on semantic networks, but it has direct translation to the language of first order predicate logic, from which it takes its semantics.
- **Description Logic:** are logics serving primarily for formal description of concepts and roles (relations). These logics were created from the attempts to formalize semantic networks and frame based systems. Semantically they are found on predicate logic,
- **Ontologies:** An ontology is a semantic/concept network that contains a body of knowledge describing some domain, typically common sense knowledge relating concepts.
- **Semantic Web: RDF, RDFS and OWL:** RDF (Resource Description Framework) is a framework for representing information about resources in a graph model, where information is represented by triples subject–predicate–object. RDFS (RDF Schema) extends RDF vocabulary to allow describing taxonomies of classes and properties. It also extends definitions for some of the elements of RDF, for example it sets the domain and range of properties

and relates the RDF classes and properties into taxonomies using the RDFS vocabulary. Web Ontology Language (OWL) brings the expressive and reasoning power of Description Logic (DL) to the semantic web. It is divided into two levels: OWL Lite and OWL DL, which differ in their expressive power and the deduction complexity. The limitation with OWL Lite and OWL DL is that reasoning is hardly implemented in an efficient way, and it also suffers from lack of scalability.

III. GENERAL IDEA

The general idea of our semantic model and reasoning approach is to define two-level CEP transformations, each of which transforms one event flow/stream into a semantically richer one: 1) from annotated preprocessed events to RDF triples; and 2) from RDF triples to a stream of facts. Initially, sensor data received from smart objects are pre-processed so as to identify: a) the entity type and instance from the received UUID; and b) what is happening to the entity, e.g., if it is doing some action, experiencing a state change or any other transformation. This 2nd type of pre-processing may be performed, e.g., through CEP (by matching a sequence of data onto a pre-defined temporal pattern identifying a specific pattern of action). This leads to a stream of semantically annotated data with pairs (*subject, predicate*) or (*object, predicate*). The entity type/instance and predicate identification is performed by CEP agent close to the sensors, cf. Figure 1), that in the specific case of our IoT middleware typically execute on mobile devices. Therefore, we named them Mobile Event Processing Agents (Mobile EPA),

Then, in the first stream processing stage, our approach transforms the stream of annotated data into a stream of RDF statements, and in the second stage, we transform the stream of RDF-triples into semantically richer facts, i.e combining RDF statements. The details of each of these stages are explained in the following.

A. Mapping Data Events to Semantic Events

Our reasoning approach dictates that each simple annotated event (actually, a data object

with member attributes) represents an action-based *predicate* (i.e. the event is the outcome of an action) and has at least one of the other two remaining RDF elements: the *subject* or the *object*. If the event has the ID of the subject and the object then we have the complete RDF triple (*Subject, Predicate, Object*), but otherwise, the missing third RDF element of the triple may be inferred from the shared context (i.e. the temporal and spatial correlation) of both elements, the subject and the object when these are received in separate events. For example, if we consider RDF statement (*ball, kicking, in the front-yard*), then the event instances represent the predicate *kick*. It further carries the ID of either the ball (e.g. when the ball carries an accelerometer sensor), or else the ID of the yard (e.g. the GPS-position or the street number of the yard (e.g., lawn sensors detect some kicking object). And the shared context is defined by the same location (co-location) of the events and the synchronicity of the events that the sensors on the yard ground and the sensor in the ball detect the hitting of the ball with the lawn (the kick). This contextual correlation is performed by CEP rules called "Context mappers", that analyze the streams of events and match Subjects, Objects and Predicates.

Figure 2 shows how Context mappers analyze each pair of events in the sliding time window (e.g. 60 s.) of Data Event Stream and try to identify common contexts, based on time proximity or any other data attribute.

B. Mapping Semantic Events to Knowledge Facts

The mapping from Semantic Events (i.e. RDF triples) to Facts is achieved by Semantic Event (SEv) rules. These are CEP rules that look out to find causality and temporal patterns in several Semantic Event sub-Streams, where each stream comprises the Semantic events of a given context. This "context-specific splitting" is possible in most CEP engines by the concept of a stream partition (*a.k.a.* context). Then, depending on the SEv rule, it might consume, filter out, modify or even insert new RDF triples in some SEv streams, a feature that is supported by CEP. This manipulation is achieved by querying the Knowledge base about all the concepts and relations pertaining to the sub-streams analyzed. For example, the inference might




Fig. 1. Semantic annotation from raw sensor data




Fig. 2. Mapping Data Events to Semantic Events

deduce that the "kicking ball with a given ID" has "Bob" as its owner, and that the "yard where the ball is kicking" is the one where Bob lives. By this, the new piece of knowledge may be derived such as "someone is kicking Bob's ball on his house's yard". And maybe with the context information "Bob has finished his homework", it is possible to deduce - with high probability - that "Bob is kicking his ball in his house's yard".

The Knowledge base is organized as nested contexts [9], which allows a much more efficient checking of concepts and relations when compared to single-layer (or flat) ontologies. For example, the ontology of the Knowledge Base may be organized as the following nested contexts: Spatial nested contexts: "Green Way district" \supset "house at 10 Rodeo Dr." \supset its yard \supset its lawn; Temporal nested contexts: "Bob's leisure time" \supset "Thursday" \supset afternoon \supset "Bob's homework finished"; Containment nested context, such as, "Bob's toys" \supset Balls \supset basket ball with ID, etc.

Figure 3 shows how Semantic Event rules analyze all RDF triples in the sliding time window (e.g. 180 s.) of sub streams of Semantic Events, trying to find event patterns, filtering, manipulating or adding RDF triples into "their" main context sub-stream or also of sub-streams of semantically related contexts, such as, "the front yard" and the "street in front of the yard".

C. Deriving Situations

Using the Facts of the stream and checking them against the Semantic Graph (Ontology) of the knowledge base, complex situations may be identified such as "Bob is playing basketball in the front yard, but should be notified that a strong storm is approaching his house's yard.". Moreover, some of the complex facts may be used for expanding, reinforcing or removing some the knowledge about a subject, an object or a place. For example, after the Bob's pen has finishes writing QED on the page with the exercises of his Math's




Fig. 3. Mapping Semantic Events to Knowledge Facts

homework notebook, the latter has been closed, and his Bob house's main door has been opened and closed, sensing that someone left the house, then the Knowledge Base will be expanded with the facts that (Bob, finished, Math homework), (Bob, left, house) and (Bob, stepped into, yard).

IV. AN EXAMPLE OF REASONING OVER DATA STREAMS

In this section we show how the aforementioned two-phase reasoning could be done with off-the-shelf components and current wireless WPAN technologies, such as Bluetooth Low Energy (BLE). Consider a scenario where smart ambient sensors are everywhere: in houses, offices, public transportation, in the streets and in private cars, and that these smart devices include a temperature and an accelerometer sensor, have a unique UUID and Bluetooth Low Energy interface. Now consider a user, Silva, lives in Rio de Janeiro and carries a smartphone running our Internet of Things middleware ContextNet [11], [21]. This middleware uses the smartphone as the bridge between Bluetooth-enabled smart devices/objects/sensors and IoT application servers executing in a cloud. The mobile middleware (Mobile Hub) periodically issues a BLE scan, discovers nearby BLE devices, connects

to them, subscribes to the smart device's sensors and writes commands to the smart objects that have some actuator. Figure 4 shows the Mobile Hub with four Texas Instruments *SensorTags* that have 6 different sensors, including temperature and accelerometer. Assume that it is summertime and that some IoT application needs to know if Silva is in his office, if he is walking on the street or if he is in a bus or car. Whenever Silva's smartphone encounters a BLE smart sensor, it is possible to deduce if he is in an air conditioned space or not, and whether he is in movement or not (due to the smart device's accelerometer). Moreover, if the location of each deployed smart device is previously registered, it is further possible to deduce if Silva is in his office or elsewhere. And this can be deduced even without the use of GPS, either because of its signal is not available (indoors), or because Silva decides to keep it off to save the smartphone's battery.

In this case, it would be possible to deduce the RDF triple (*Silva, rides, BusLine435*) from following simpler semantic events in the stream (*SilvasPortable, connectedTo, sensorX*) and (*sensorX, in, BusKKZ8674*) and (*sensorX, Abs(Accelerator) > 10*) and (*sensorX, temp = 20*) and the fact that *BusKKZ8674* operates the




Fig. 4. ContextNet Mobile Hub with four SensorTags

BusLine 435. Moreover, it would be possible to deduce that Silva is moving in the traffic, but that he is in an air-conditioned bus, which may be very important during Rio's summertime, when outdoor temperature can reach more than 45 degrees Celsius.

V. RELATED WORK

In an early work, Adi *et al* [1] presents abstractions that describe semantic relationships between events, object and tasks. These are defined as generalizations and associations and through attributes that may reference events. Their abstractions are suitable for specification but cannot be computed efficiently. On the other hand, the work [3] describes a system (ETALIS) that can perform reasoning over streaming events with respect to background knowledge, similar to our Knowledge Base. It implements two languages for specification of event patterns: the rule based ETALIS Language for Events (ELA), and Event Processing SPARQL. ETALIS can evaluate domain knowledge on-the-fly, thereby proving semantic relations among events and reasoning about them. Their semantic relations among events are time-based, but don't have the synchronicity requirement. Another difference is that they do not generate a RDF Stream which they check against a knowledge base. Thus, their inference is much simpler than the one proposed in our project.

Tachmazidis *et al* [23] propose a reasoning method over RDF triples based on defeasible logic (i.e., a non-monotonic logic) which can be implemented in a massively parallel way. They used

Hadoop, an open-source implementation of the MapReduce paradigm, and a stratified rule set for a more efficient processing of the knowledge base. Unlike our proposal, they do not handle Stream Processing and do not apply their method to reasoning for time-critical systems, such as CPS. Moreover, their choice for defeasible logic limits the sorts of knowledge that can be inferred by their system, as opposed to temporal logic, which shall have highly parallelizable implementations.

The following projects CityPulse [15] Star-City [17] and FIESTA-IoT [2] also present research toward the use of Semantic Stream reasoning. All of these projects use the knowledge base in order to deduce new context/facts. Also, they use a single-layer (or flat) ontology model, which differs from our ontology model that is organized as nested contexts. Moreover, none of these projects focus on the problem of delivering real-time reasoning.

The FIESTA-IoT project [2] integrates several other projects and one of them is the CityPulse project [15]. The main goal of these projects is to achieve semantic interoperability at different levels (hardware, data, model, query, reasoning and application levels). The StarCity project has a similar idea, but it is aimed at using semantics to provide interoperability at the data level.

On the other hand, the work by Teymourian *et al* [24] has the same focus as our work. They use a similar idea and combine the use of SCEP rules (semantic web plus CEP) with a semantic knowledge base to deliver real-time reasoning. The difference is that our work uses an ontology

model organized as nested context to represent context information, rather than a flat ontology model. As a result, it is more efficient on query processing, because when we execute a query, the query will be processed only using a sub-set of the knowledge base (a partition of the knowledge base). Furthermore, another difference is that we plan to insert new SCEP rules on-the-fly, based on new facts generated by the reasoning over the knowledge base. Consequently, it will give the application a more efficient approach to adapt to different situations. For example, in a monitoring application, we only need a CEP rule that triggers an action based on an altitude situation only if the monitored person is in a high altitude, until then this rule does not need to be there.

VI. DISCUSSION

Combining symbolic reasoning based on ontologies with Complex Event Processing has several advantages. Firstly, it allows to leverage CEP's efficient processing of dense flows of simple events, not just over raw sensor events but also over RDF triples. Secondly, CEP's ability to produce complex events is also necessary for the iterative generation of higher level information from lower level bits of information.

On the other hand, while CEP is appropriate for processing data that is carried by the incoming events, it is incapable of detecting domain-specific relationships between events that are produced by distinct entities/objects that apparently have no relation with each other, or when this relationship cannot be directly encoded by the (meta-)information carried by the events. Symbolic reasoning using ontologies, on the other hand, can very well model these "indirect" relationships among the monitored entities and/or their corresponding events. And hence, by using the results of a query over a domain-specific ontology during a CEP-based continuous processing, it becomes possible to generate new sorts of events (i.e., fact events), which are produced independently by the Semantic Event reasoners in response to the consumption of some RDF triples. These Fact events, which in some sense embody some semantic knowledge that was forked off the knowledge base, can in turn be further processed by other CEP engines, and may be used to predict events

that actually did not yet happen, but which are a natural consequence of initial events that have been detected by CEP.

This makes us consider the Semantic Web reasoners as a special kind of CEP engines, which have access to knowledge base, consume RDF events and eventually produce fact events that are passed on to other CEP engines in the Event Processing Network. See Figure 2 for a general architecture.

VII. TOWARDS AUTOMATED RULES AND PATTERNS INDUCTION

In this section, we briefly discuss prospects for using machine learning and induction techniques to extract useful information from the data stream. At first, let us mention that although there are known and proven techniques for extracting information from data, from raw data to structured data and knowledge, most of them have been designed as off-line techniques and with the assumption of all data present in the working memory. Therefore, there is a great challenge in adapting when possible current techniques to on-line stream data processing with huge volumes of data or designing new techniques. A good review of the issues (continuous data streams flow, unbounded memory requirements, mining changes, avoiding overfitting...) can be found in [12]. Another good analysis could be found is [22].

Let's start with the raw data produced by the sensors. Unsupervised algorithms, such as sparse autoencoders, may be used to automatically extract higher level features [16]. The basic idea is to use an autoencoder, a neural network with a hidden encoding layer and a decoding output layer identical to the input layer. We add an additional sparsity activation constraint, in order to enforce specialization of each neuron as a specific feature detector. Training an autoencoder, sometimes named self-supervised, relies on traditional supervised learning on learning identity, as the autoencoder learns to reconstruct its input data on its output. Once trained, to extract features from an input, one just needs to feed forward the input data and gather the activations. One may use successive levels (stacks) of autoencoders in order to extract more abstract features. Although standard training is off-line, one may make use it incrementally, with successive

rounds of batch training. Compared to traditional dimension reduction algorithms, such as Principal Component Analysis (PCA), this approach has 2 advantages: the number of features may be arbitrary and not necessarily smaller than the initial dimension of the input and feature extraction is non linear.

An interesting end to end approach has been proposed by Ganz et al. in [13]. The first step, named SensorSAX (as for Sensor Symbolic Aggregate Aproximation), is the discretization of data into qualitative attributes, encoded in some alphabet words. Then a clusterization step is applied, using a k-means non supervised clusterization algorithm, by considering time as one of the criteria, to form patterns, which are proto-concepts (not yet named concepts). Temporal relations between these proto-concepts are extracted by constructing a Markov model, a statistical predictive model of temporal occurrences of proto-concepts. Three kinds of temporal relations are considered: `occursAfter`, `occursBefore` and `occursSame`. The last step, at this stage still manual, consists in labeling, i.e. naming proto-concepts into named/symbolic concepts (e.g., `coldTemperature`). The authors are also considering the possibility of automatic labeling, derived from the labels of the sensors and a common sense ontology.

When starting from RDF triples, one may consider various knowledge extraction methods based on ontologies (mostly based on OWL), designed for the Semantic Web. One objective is induction, to be able to construct more abstract knowledge (concepts/hypotheses) from the facts. Various algorithms exist and aim at both generalizing examples into concepts, while specializing them in order to uncover counter examples. Inductive Logic Programming (ILP) is a seminal formalism but there are many variant (see for instance the DL-Learner framework [8] which includes various ones), as well as related techniques like decision trees construction and also exploratory approaches based on genetic algorithms.

An interesting proposal in [6], offers inductive reasoning as well as deductive reasoning on RDF data streams. Deductive reasoning is performed on queries constrained by concepts expressed in OWL. C-SPARQL [7] is the query language used.

It is an extension for continuous queries on RDF streams of the SPARQL RDF query language. Inductive reasoning is performed on a subset of data in order to be practically computable. The user defines statistical units (entities, e.g., persons) as well as a population of these entities (e.g., at a specific institution or location) on which he wishes to make inductive queries. The inductive engine periodically updates data matrices representing relations between RDF triples corresponding to the population of the statistical units considered (actually, there are two kinds of matrices, one long term stable and one short term representing the trends) and conducts a multivariate analysis of the data matrices. The trained model could then be used to predict relationships between entities at query time.

Last, it is also important to be able to extract temporal relations. A promising proposal is by Georgala et al. [14] to efficiently extract all possible temporal relations (along seminal Allen's taxonomy and algebra of temporal intervals) from time stamped RDF streams.

In summary, we could see that there are various interesting directions for introducing automated machine learning and knowledge extraction techniques into our framework. One important issue is the dynamicity of the data produced. That is, we need to find good trade-offs depending on the usage, the nature of the data and the computing and communicating resources available, between the demand for higher level knowledge, the cost for extracting it and the risk of it being obsolete because of the constant stream of data. For instance, in the scope of our two stage process, we believe that machine learning could be effective as a way to consolidate into the knowledge base the facts which occur very frequently (see Figure 3). An example of such learning (in that case, inductive) is the identification of two temperature settings, inside a air-conditioned bus and outside, that will be extracted from repeated facts of passengers entering and exiting air-conditioned buses.

Therefore, one needs to carefully examine what exact machine learning techniques we will insert into our framework, and at which stage.

VIII. CONCLUSION AND FUTURE WORK

This paper presented a real-time reasoning approach based on semantic events and fact streams for IoT systems. The reasoning approach is based on the assumptions that all objects, people, buildings, places, vehicles, environments, etc. will have many embedded tiny sensors that will emit simple events whenever some action is performed with/to it by an actor, and that each event will carry the items' unique UUID and an accurate timestamp. By enforcing the restriction that predicates in a RDF triple must be action-based, such as "kick", "put", "grab" etc., rather than state-based, such as "has", "is" "belongs to", etc., we are of course limiting the amount of information that the data/event streams are capable to express. However, we believe that the action-based predicates are the really important ones for reasoning in IoT applications. All the state predicates, on the other hand, should instead be represented by the nested context-based ontology in the Knowledge Base. We are aware that this is only a first and initial step towards adding semantics to real-time reasoning over data streams, and that much more theoretical and practical research is required to validate our approach, evaluate it under a broader perspective and show its feasibility for large-scale and distributed IoT applications. However, we are confident that it is a promising first step. As next steps, we will finish the development of the Context Mappers and Semantic Event Rules using Esper's EPL (Event Processing Language) and deploy them on our mobile IoT middleware. In parallel, we will model a simple scenario and the main entities and their relationships, as described in Section IV and represent it as nested contexts.

Acknowledgement: Our ESMOCYP (Efficient Semantic Models and Fault-tolerant Middleware for CYber-Physical Systems) Project is being supported by CAPES (Brazil) and DAAD (Germany) cooperation program PROBRAL (Process No 8148/2015-05) and by a CAPES PVE fellowship to J.-P. Briot.

REFERENCES

- [1] A. Adi, D. Botzera, and O. Etzion. Semantic Event Model and its Implication on Situation Detection. In *European Conference on Information Systems (ECIS 2000)*, 2000.
- [2] Yahya Al-Hazmi and Thomas Magedanz. Towards semantic monitoring data collection and representation in federated infrastructures. In *IEEE International conference on future internet of things and cloud (FICloud)*, pages 17–24, 2015.
- [3] Darko Anicic, Sebastian Rudolph, Paul Fodor, , and Nenad Stojanovic. Stream Reasoning and Complex Event Processing in ETALIS. *Semantic Web*, (1):1–5, 2009.
- [4] Arvind Arasu, Shivnath Babu, and Jennifer Widom. The CQL continuous query language: Semantic foundations and query execution. *The VLDB Journal*, 15(2):121–142, June 2006.
- [5] G. Baptista, F. Carvalho, S. Colcher, and M. Endler. A middleware for data-centric and dynamic distributed complex event processing for iot real-time analytics in the cloud. In *34th Brazilian Symposium on Computer Networks and Distributed Systems (SBRC'2016)*, Salvador, Brazil, June 2016.
- [6] Davide Barbieri, Daniele Braga, Stefano Ceri, Emanuele Della Valle, Yi Huang, Volker Tresp, Achim Rettinger, and Hendrik Wermser. Deductive and inductive stream reasoning for semantic social media analytics. *IEEE Intelligent Systems*, 25:32–41, November–December 2010.
- [7] Davide Francesco Barbieri, Daniele Braga, Stefano Ceri, Emanuele Della Valle, and Michael Grossniklaus. C-sparql: Sparql for continuous querying. In *18th International World Wide Web Conference (WWW'09)*, pages 1061–1062, 2009.
- [8] Lorenz Bühmann, Jens Lehmann, and Patrick Westphal. DL-learner-a framework for inductive learning on the semantic web. *Web Semantics: Science, Services and Agents on the World Wide Web*, 39(C):15–24, August 2016.
- [9] Isabel Cafezeiro, José Viterbo, Alexandre Rademaker, Edward Hermann Haeusler, and Markus Endler. Specifying ubiquitous systems through the algebra of contextualized ontologies. *The Knowledge Engineering Review*, 29(02):171–185, 2014.
- [10] K. Selçuk Candan, Huan Liu, and Reshma Suvarna. Resource description framework: Metadata and its applications. *SIGKDD Explor. Newsl.*, 3(1):6–19, July 2001.
- [11] M. Endler, G. Baptista, L. D. Silva, R. Vasconcelos, M. Malcher, V. Pantoja, V. Pinheiro, and J. Viterbo. Contextnet: Context reasoning and sharing middleware for large-scale pervasive collaboration and social networking. In *ACM/IFIP/USENIX 12th International Middleware Conference Workshop on Posters and Demos Track*, pages 2:1–2:2. ACM, 2011.
- [12] Mohamed Medhat Gaber, Arkady Zaslavsky, and Shonali Krishnaswamy. Mining data streams: A review. *SIGMOD Record*, 34(2):18–26, June 2005.
- [13] Frieder Ganz, Payam Barnaghi, and Francois Carrez. Automated semantic knowledge acquisition from sensor data. *IEEE Systems Journal*, 10(3):1214–1225, September 2016.
- [14] Kleanthi Georgala, Mohamed Ahmed Sherif, and Axel-Cyrille Ngonga Ngomo. An efficient approach for the

- generation of Allen relations. In *22nd European Conference on Artificial Intelligence (ECAI'2016)*, Den Haag, The Netherlands, August–September 2016.
- [15] Maria Giatsoglou, Despoina Chatzakou, Vasiliki Gkatziki, Athena Vakali, and Leonidas Anthopoulos. City-Pulse: A platform prototype for smart city social data mining. *Journal of the Knowledge Economy*, 7(2):344–372, 2016.
- [16] Quoc V. Le, Marc’Aurelio Ranzato, Rajat Monga, Matthieu Devin, Kai Chen, Greg S. Corrado, Jeff Dean, and Andrew Y. Ng. Building high-level features using large scale unsupervised learning. In *29th International Conference on Machine Learning*, Edinburgh, U.K., 2012.
- [17] F. Lécué, S. Tallevi-Diotallevi, J. Hayes, R. Tucker, V. Bicer, M. L. Sbodio, and P. Tommasi. Star-City: semantic traffic analytics and reasoning for city. In *19th international conference on Intelligent User Interfaces*, pages 179–188. ACM, 2014.
- [18] David C Luckham. *The Power of Events: An Introduction to Complex Event Processing in Distributed Enterprise Systems*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2001.
- [19] Natalya Fridman Noy and Carole D. Hafner. The state of the art in ontology design – a survey and comparative review. *AI Magazine*, 18(3), 1997.
- [20] L. Talavera Rios, M. Endler, and S. Colcher. An energy-aware iot gateway, with continuous processing of sensor data. In *34th Brazilian Symposium on Computer Networks and Distributed Systems (SBRC'2016)*, Salvador, Brazil, June 2016.
- [21] L. Talavera Rios, M. Endler, I. Vasconcelos, R. Vasconcelos, M. Cunha, and F. Silva e Silva. The mobile hub concept: Enabling applications for the internet of mobile things. In *12th IEEE Workshop on Managing Ubiquitous Communications and Services (MUCS 2015)*, pages 123–128, St. Louis, MI, USA, March 2015.
- [22] Xiang Su, Ekaterina Gilman, Peter Wetz, Jukka Riekk, Yifei Zuo, and Teemu Leppänen. Stream reasoning for the Internet of Things: Challenges and gap analysis. In *6th International Conference on Web Intelligence, Mining and Semantics (WIMS'16)*, Nimes, France, June 2016.
- [23] I. Tachmazidis, G. Antoniou, G. Flouris, S. Kotoulas, and L. McCluskey. Large-scale parallel stratified defeasible reasoning. In *ECAI-12*, 2012.
- [24] Kia Teymourian, Malte Rohde, and Adrian Paschke. Fusion of background knowledge and streams of events. In *6th ACM International Conference on Distributed Event-Based Systems, DEBS'12*, pages 302–313. ACM, 2012.