Provenance Logic: Enabling Multi-Event Based Trust in Mobile Sensing

Xinlei Wang, Hao Fu, Chao Xu, Prasant Mohapatra

Department of Computer Science
University of California, Davis, CA 95616
{xlwang, haofu, chaoxu, pmohapatra}@ucdavis.edu

***Abstract*—With the proliferation of sensor-embedded mobile computing devices, mobile sensing is becoming a popular paradigm to collect information from participating mobile users. Unlike the well-calibrated and well-tested sensor networks, mobile sensing relies on participants with unknown reliability. Data collected from mobile users may be untrustworthy. There are various solutions proposed in the literature for assessing the trustworthiness of the sensing data that describe an individual event or observation. In addition to single-event based trust models, we propose the concept of *Provenance Logic*, to reason about the logical relations between multiple events by jointly recognizing and linking events from successive sensing observations. We propose an approach that combines logical reasoning and statistical learning techniques. To the best of our knowledge, our work is the first attempt for trust evaluation based on the logical relation among multiple events in the mobile sensing context. We motivate and illustrate our approach with a use case of traffic monitoring mobile sensing. Performance validation has shown that improved trust assessment can be achieved efficiently and effectively on top of single-event based analysis.**

I. Introduction

We have seen the massive prevalence of mobile computing devices such as smartphones and tablet computers in recent years. These devices usually come with multiple embedded sensors, such as camera, microphone, GPS, accelerometer, digital compass and gyroscope. Such advancements have made mobile sensing become an extremely popular data collection paradigm, where people use their personal mobile devices to perform sensory data collection tasks. A notable number of mobile sensing applications have emerged for collecting specific data such as traffic [1], noise pollution [2], cyclist experiences [3], and consumer pricing information [4]. Unlike traditional well-calibrated and well-tested wireless sensor networks (WSN), mobile sensing relies on participants of unknown trustworthiness. Therefore, a challenge in mobile sensing applications lies in the unreliability of data sources. Mishandling of uncertainty in the data may result in a wrong perception of the situation. Therefore, we need schemes to analyze the trust of sensing data streams about events in order to combine data from multiple sources and derive reliable conclusions from them.

Prior literature have studied trust or credibility assessment of information from potentially unreliable sources in social networks [5], [6], traditional WSN [7]–[9] and mobile sensing networks [10], [11]. Provenance, i.e., the derivation history, of the sensing data has been considered a key element to reason about trust of data. Particularly, researchers have looked at provenance data from three different dimensions: node provenance (who processed the data), spatio-temporal provenance (where and when the data were originated and processed) and social provenance (the social relationship of the nodes who processed the data). With provenance information available, various algorithms have been proposed to evaluate the trustworthiness of sensing data describing an individual sensing event.

Mobile sensing applications with different purposes may work under completely different system models and may have very different requirements. In this work, we focus on mobile sensing applications that collect data that describe events that are likely to be in progress in a target environment, with a purpose of performing situation assessment or hazard detection. A category of such applications would be traffic sensing applications, where users share their real-time observations about traffic events/situations like accidents, traffic jam, road hazards, etc. One such application is Waze [12], a GPS navigation application available on the iOS App Store. In such an application scenario, it is common that a sequence of events in the physical world is reported within a short period of time by either the same set or different sets of participants. For example, observations at time t_1 claim that "a major accident happened at location l" and subsequent observations at time t_2 claim "a heavy traffic jam is observed at location l". Events observed within a short period of time or at the same location often possess logical relations in terms of time, location and other contextual factors. With the observation description and the associated provenance data, we can correlate against different events and identify the logical support or conflicts among them, which allows us to further reason about the trust of the sensing data in addition to only looking at single events.

In this work, we propose the concept of *Provenance Logic*, to reason about the logical relations between multiple events with the available provenance information, which provides the foundation for another level of information trust evaluation on top of single-event based trust evaluation. In other words, we are trying to answer the question that, by jointly recognizing and linking events from successive sensing observations, how one can determine the probability that a given observation is true with a higher accuracy than only looking at individual events. To the best of our knowledge, our work is the first attempt to evaluate information trust in mobile sensing based on the logical relation among multiple events.

With the information in the provenance data associated with the sensing reports, we aim to reason about the logical relations between events from three dimensions: *spatial*, *temporal*, *contextual* as well as the inter-correlation between them. Violation of provenance logic, e.g., different events indicate

that the same entity was at different locations at the same time, will degrade our trust towards the sensing data. It is also possible that mutually supporting provenance makes us trust the sensing data more. Event Calculus [13] is adopted and extended in order to represent and analyze the logical relation among multiple events. We then transform such logical relation into a probabilistic graphic representation based on Markov Logic Network (MLN) [14], which enables us to infer the probability of one event given the probability of others.

The main contributions of this paper can be summarized as:

1) An event logic module to extract sensing observations that possess logical relations and translate them with their provenance information to first order predicates.
2) A logical semantic model extending Event Calculus to represent the provenance logic of sensing observations from three dimensions: temporal logic, spatial logic and contextual logic.
3) An MLN based mechanism to perform logic-based trust reasoning for a sensing observation given the other related observations and their pre-computed trust.

The rest of the paper is organized as follows. We highlight the related work in Section II. In Section III, we describe our system model and give a brief overview of how our entire framework works. We then present our proposed framework in Section IV. The performance evaluations based on simulation experiments and a case study are presented in Section V. We give a discussion and talk about our future work in Section VI. Finally, Section VII concludes the paper.

II. Related Work

To assess the trust of reports in mobile sensing, a relevant body of work in machine learning and data mining communities performs trust analysis. Yin et. al. introduced *TruthFinder* as an unsupervised fact-finder for trust analysis on a providers-facts network [15]. Wang et. al. proposed a maximum likelihood estimator to directly and optimally quantify the accuracy of conclusions obtained from credibility analysis in mobile sensing [10].

In the networking community, trust analysis of sensing data in both traditional WSN and mobile sensing has been studied. Raya et. al. [16] developed the notion of data trust in ephemeral networks. They evaluate data reports with corresponding trust levels using several decision logics, namely weighted voting, Bayesian inference, and Dempster Shafer Theory. Several recent efforts [7]–[9], [17] studied the information trust analysis based on node-level provenance in multi-hop sensor networks. The ARTSense [11] is a scheme where data trust is evaluated anonymously in mobile sensing applications. None of these trust analysis schemes has considered the possible logical relation among multiple events or observations. To the best of our knowledge, our work is the first attempt to look at the data trust analysis problem at a multi-event level in the mobile sensing context.

To model and analyze the ordering, concurrency and other logical relations between situations and events, A variety of logical formalisms, e.g., Situation Calculus [18], Event Calculus [13] and Event Logic [19] have been proposed and studied in the artificial intelligence community. They have been widely explored in the field of event recognition from surveillance videos. However, none of the techniques has been applied in analyzing trust of data in the field of mobile sensing. In this work, we employ Event Calculus as the basis of our provenance logical reasoning.

Probabilistic graphical models, e.g., Dynamic Bayesian Networks (DBN) [20], Markov Random Field (MRF) [21] and Markov Logic Networks (MLN) [14] have been applied to a variety of event recognition tasks where uncertainty exists. To analyze the amount of uncertainty, i.e., trustworthiness, in the sensing reports, we adopt MLN as our trust computation framework because it is a new and powerful approach that combines both uncertainty and logic language to tackle the relation among multiple events.

III. Preliminaries

A. System Overview

We consider a mobile sensing application model where a group of participants make individual observations. Each observation describes an event or situation in the physical world at the observation time. We assume there is a finite domain of states that a participant could report.

The provenance information is attached to the sensing reports as meta-data, where each piece of provenance information is a property-value annotation. The provenance annotations include *time*, *location* and a list of *contextual provenance* (e.g. weather condition) required by the specific application. We assume that sensing reports with the same observation and similar provenance information are grouped and consolidated as a *claim*. Each claim's *initial trust* can be evaluated based on an existing single-event trust analysis scheme. Such an initial trust is denoted as $T(0)$.

A sliding window of N claims $C_1, C_2, \ldots, C_N$ are recorded and ordered by their reporting time. We assume each of the N claims has been assessed by our scheme at least once and assigned with an updated trust value $T(x)$, where x denotes the number of rounds of trust updates that have been done by our system. We name this sliding window as the *evidence window* since each claim is considered a potential evidence.

Suppose we received a claim C' based on some newly arrived sensing reports, we first obtain its initial trust $T'(0)$. After that, $C_1, C_2, \ldots, C_N, C'$, together with their provenance annotations and current trust values are fed into our system as the input. The goal of our system is to (1) find the claims in $C_1, C_2, \ldots, C_N$ that have logical relations with C' based on the logic relation definitions in our knowledge base; (2) construct an MLN based on the logical relations and predefined rules; (3) update each of the related claims' (including C') trust based on the MLN probability inference.

After the above steps, the original C_1 is replaced with C' to maintain the size of the evidence window. Not all of the claims in the evidence window are necessarily related to a newly arrived claim. Therefore, only those that have a logical relation with the new claim will get a trust update. However, for clarity purpose, we increase x in the $T(x)$ notation for every claim after a round of trust update. Ultimately, when a claim exits the evidence window, its trust value should be updated as $T(N+1)$. We expect $T(N+1)$ to be more accurate than the corresponding $T(0)$ of the same claim. The size of the evidence window N can be adjusted based on the real-time

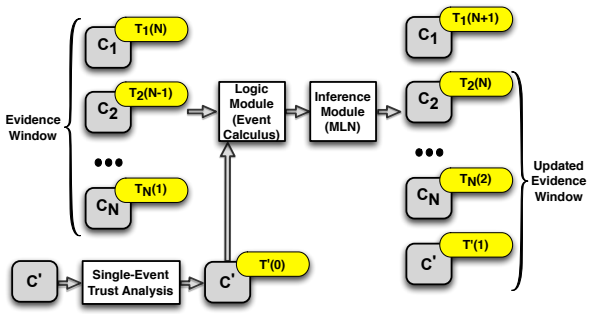

Fig. 1: An overview of multi-event trust analysis

needs of the application.

Figure 1 provides an illustration of a round of multi-event based trust update by our system. The *Logic Module* and *MLN Module* are the main focus of our work.

B. Event Calculus

The Event Calculus, originally introduced by Kowalski and Sergot [13], is a first-order logical language for representing and reasoning about events and their effects. The basic components of the Event Calculus are *time points*, *events* and *fluents*. A fluent is a state whose value may vary over time. When an event occurs, it may change the value of a fluent. Table I summarizes the domain-independent predicates of the Event Calculus. The core domain-independent axioms define whether a fluent holds or not at a specific time-point. Moreover, the axioms incorporate the common sense *law of inertia*, according to which fluents persist over time, unless they are affected by the occurrence of some event. A knowledge base of Event Calculus axioms is defined by a set of first-order logic formulas.

TABLE I: The Event Calculus predicates

Predicates	Meaning
$initially_P(F)$	Fluent F holds at time-point 0
$initially_N(F)$	Fluent F does not hold at time-point 0
$holdsAt(F,T)$	Fluent F holds at time-point T
$happens(E,T)$	Event E occurs at time-point T
$initiates(E,F,T)$	Event E initiates fluent F at time-point T
$terminates(E,F,T)$	Event E terminates fluent F at time-point T
$clipped(F,T_0,T_1)$	Fluent F is terminated at some time-point in the interval $[T_0,T_1)$
$declipped(F,T_0,T_1)$	Fluent F is initiated at some time-point in the interval $[T_0,T_1)$

C. Markov Logic Network

A knowledge base of Event Calculus axioms is defined by a set of first-order logic formulas. Each formula imposes a hard constraint over the set of possible worlds. This, however, does not handle uncertainty adequately. We employ the framework of Markov Logic Networks (MLN) [14] in order to soften the constraints and perform probabilistic inference.

In contrast to the Event Logic, all worlds in MLN are possible with a certain probability. An MLN consists of a set of first-order logic formulas and a weight value ($\in \mathbb{R}$) associated with each formula. Formulas with weights impose soft-constraints over the possible worlds, which enables imperfect knowledge in the domain to be captured and utilized for inference. The stronger the constraint represented by a formula F_i, the higher the value of its weight w_i.

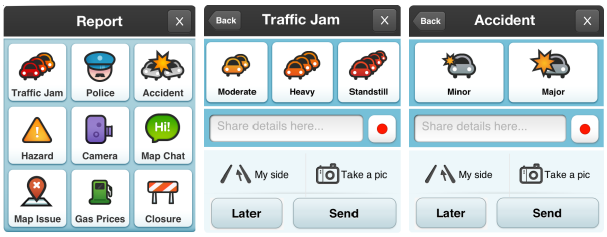

Fig. 2: Reporting interface of Waze [12]

Formally, in our case, the knowledge base K consists of Event Calculus formulas converted into clausal form. Together with the associated weights and a finite domain of constants C where C contains all the distinct values of events, fluents and time. K is transformed into a ground Markov network $M_{K,C}$, which contains one binary node for each possible grounding of each predicate appeared in K. The value of the node is 1 if the ground atom is true, and 0 otherwise. The predicates of a grounded formula form a clique in M_K, C. Each clique is associated with a weight w_i of the corresponding formula. The ground MLN defines a probability distribution over the possible worlds.

A Markov network is a model for the joint distribution of a set of variables $X = (X_1, X_2, \cdots, X_n) \in \mathcal{X}$. The probability distribution over possible worlds x is given by:

$$P(X = x) = \frac{1}{Z}\exp\left(\sum_i w_i n_i(x)\right) \quad (1)$$

where $n_i(x)$ is the number of true groundings appeared in each formula F_i in x and w_i is the weight associated with the formula F_i. Z is known as the *partition function* and is given by $Z = \sum_{x\in\mathcal{X}} \exp\left(\sum_i w_i n_i(x)\right)$.

D. Use Case: Traffic Sensing

To better motivate and explain the system model as well as our entire scheme, we take a traffic sensing application like Waze [12] as an illustrating example. Figure 2 shows the interface of Waze on which users report their observations. For illustration purpose, we consider a simplified case derived based on Waze: the events/situations that a user can report include traffic jam (major, minor), police, accident (moderate, heavy, stand-still), and road hazards. The provenance information of the observation report includes time, location and one contextual property: traveling speed.

Table II shows a mapping of the sensing data (observation and contextual provenance) to the Event Calculus elements (events and fluents). The events and traffic fluents can be extracted from the observations reported and the vehicle fluents can be derived from the traveling speed property in the contextual provenance.

TABLE II: Events and fluents in traffic sensing use case

Events	Fluents	
	Traffic	Vehicles
Minor Accident Major Accident Police Hazard	Moderate Jam Heavy Jam Standstill	High Speed Low Speed Stopped

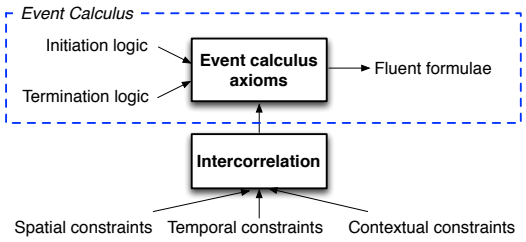

Fig. 3: Adding provenance logical analysis to event calculus

IV. The Scheme

A. Provenance Logic

To demonstrate how logical support and conflict may arise in a traffic sensing scenario, let us assume that the following four claims are received:

C_1: `A major accident is at Location X at 2:00pm (reported by 10 drivers).`
C_2: `A heavy traffic jam is at Location X at 2:15pm (reported by 30 drivers).`
C_3: `A police is at Location X at 2:30pm (reported by driver Alice with traveling speed of 70 mph).`
C_4: `A police is at Location Y (50 miles away from X) at 2:45pm (reported by Alice with traveling speed of 70 mph).`

Clearly we can see the C_1 and C_2 support each other since a major accident is likely to cause a heavy traffic jam. C_3 seems to be reasonable at first glance since police usually appears after an accident. However, there is also an obvious *contextual conflict* between C_2 and C_3 since a car is unlikely to be traveling at 70 mph when there is a heavy traffic jam. In C_4, we can observe that *spatio-temporal constraints* are violated when comparing against C_3 because it is a car could hardly travel 50 miles in 15 mins.

The Event Calculus defines the logical semantics to model the causal relation between events and fluents by integrating the initiation and termination predicates for a particular fluent of interest into formulas.

From the example given above, one can see the necessity to have a way to analyze the logical support/conflict between multiple claims with the available provenance information, in addition to the causal relation modeled by the Event Calculus. First of all, besides direct causal relations, other events or fluents which do not directly initiate or terminate a particular fluent may also impact the validity of the fluent. We define that the fluent for which we intend to design an Event Calculus formula is the *primary fluent* of the formula. The events which have direct causal relations with the primary fluent are defined as the *primary events*. The rest of the events and fluents which may potentially affect the validity of the primary fluent and thus are introduced in the fluent's formula are defined as *secondary events* and *secondary fluents*.

To enable the *Provenance Logic* reasoning of secondary events and fluents to Event Calculus, we introduce the following three Provenance Logic dimensions:

1) **Spatial constraint**: The location where a secondary event happens or a secondary fluent holds should not impair the primary fluent's validity.
2) **Temporal constraint**: The time when a secondary event happens or a secondary fluent holds should not impair the primary fluent's validity.
3) **Contextual constraint**: None of primary/secondary event's or secondary fluent's provenance properties should impair the primary fluent's validity.

To make our ideas generic and clear, we defined the above types of constraint separately. However, to construct a fluent formula, it is very common that more than one of them need to be considered collectively in order to formulate an adequate logic constraint. The conflict between C_3 and C_4 in our above example illustrates a case where both spatial and temporal provenance are jointly considered. Figure 3 gives an illustration of how provenance-based logical reasoning is added to the Event Calculus.

B. Location and Context Aware Event Calculus

The Event Calculus only defines a way to model temporal relations between events and their fluents. However, the trust of reported events cannot be fully judged with only the temporal logic. Since we want to reason about the support and conflicts between events in terms of spatial, temporal and contextual relations, we need to add spacial and contextual semantics to the current Event Calculus. Now we discuss an approach to express spacial and contextual elements in the Event Calculus.

Location-awareness: An extension of Event Calculus, spatio-Temporal Extended Event Language (STEEL) [22], was proposed to address the issue of representing spatial information in the Event Calculus. In STEEL, a new variable L is introduced in every predicate to represent location, e.g., $initially_P(F, L)$ can be interpreted as fluent F holds at time-point 0 at spatial region L and $happens(E, < T, L >)$ can be interpreted as event E occurs at time T in spatial region L.

Introducing an L variable in each predicate certainly makes the Event Calculus capable of modeling various complicated spatial relations between events. However, when we perform logical reasoning of various situations (e.g., construction of the ground Markov network), we need to have every predicates grounded. In this case, an additional variable leads to clauses with a large number of disjunctions and a combinatorial explosion of the number of clauses that are generated.

Considering the fact that generally multiple events/fluents possess logical relations because they happen at the same location, if there is no secondary events or fluents at another location having any logical impact on the primary fluent, we do not necessarily introduce an additional variable to every predicates in the fluent formula. We define such a location of interest as the *primary location*. If there is a location constraint triggered by a secondary event/fluent at a *secondary location*, what really matters is the distance between the secondary event/fluent and the primary fluent. Hence, instead of introducing a new variable in each predicates, we introduce the *distance* variable in two predicates as in Table III.

One thing to be noted is, adding the above predicates to a formula generally does not introduce a complete secondary spatial or temporal constraint. They need to be coupled with additional predicates $D < D'$ or $|T_s - T_p| < T'$ which describe the time and distance bounds between primary and secondary locations and time points.

TABLE III: Predicates with a distance variable

Predicates	Meaning
$holdsAt(F, T, D)$	Fluent F holds at time-point T at a location D distance away
$happens(E, T, D)$	Event E occurs at time-point T at a location D distance away

Context-awareness: Contextual information in the provenance data is often expressed in the form of property-value pairs, e.g., *traveling speed* = 70 mph. For each contextual property, we can derive a requirement on it for the primary fluent's validity. Each contextual property can then be translated to a contextual fluent predicate in the Event Calculus axioms. However, since applications and the validity of fluents usually have different requirements for the contextual conditions, it is easy to image that the way the contextual fluent is defined may vary when the application is different or even when the primary fluent is different. There is no way to define a generic contextual property to contextual fluent translation process. For example, if we want to introduce a traveling speed constraint for a formula of the *heavy traffic jam* fluent at time T, we can add a predicate $holdsAt(LowSpeed, T)$ where $LowSpeed$ defines a contextual fluent where the traveling speed of vehicles is low. However, when we translate the *traveling speed* property when the primary fluent is *Standstill*, we might add a different predicate $holdsAt(Stopped, T)$ since moderate traffic jam and heavy traffic jam obviously put different constraints to the contextual condition of vehicles' traveling speeds.

We have now described the way we can add location and context awareness to the Event Calculus. We will next define a set of Event Calculus formulas related to the *Heavy Traffic* fluent in the traffic sensing use case, as an illustrative example of how the formulas are actually defined (we ignored the distance attribute here for clarity).

$$holdsAt(HJ, T_1) \leftarrow initially_P(HJ) \wedge \neg clipped(HJ, 0, T_1) \quad (2)$$

$$holdsAt(HJ, T_2) \leftarrow happens(MA, T_1) \wedge initiates(MA, HJ, T_1) \wedge \neg clipped(HJ, T_1, T_2) \wedge T_1 < T_2 \quad (3)$$

$$holdsAt(HJ, T_2) \leftarrow happens(MI, T_1) \wedge initiates(MI, HJ, T_1) \wedge \neg clipped(HJ, T_1, T_2) \wedge T_1 < T_2 \quad (4)$$

$$holdsAt(HJ, T_2) \leftarrow happens(HZ, T_1) \wedge initiates(HZ, HJ, T_1) \wedge \neg clipped(HJ, T_1, T_2) \wedge T_1 < T_2 \quad (5)$$

$$clipped(HJ, T_1, T_2) \leftarrow happens(P, T) \wedge terminates(P, HJ, T) \wedge T_1 < T < T_2 \quad (6)$$

Formulas 2 - 6 are derived based on the Event Calculus only. They cover the initiation and termination of the *Heavy Jam* (HJ) fluent. Formula 2 specifies the condition when there is a heavy traffic jam at the beginning point in time of the analysis, i.e., time of the first claim in the evidence window and nothing has terminated the HJ fluent. Formula 3 means HJ holds if a *Major Accident* (MA) happens and nothing has terminated the HJ fluent. Similarly, Formula 4 and 5 describe the logic condition when HJ is caused by a *Minor Accident* (MI) and *Hazard* (HZ) respectively. Formula 6 defines the termination logic of a HJ fluent, where the occurrence of a $Police$ event at the same location will clear the HJ fluent.

Subsequently, we demonstrate how the spatio-temporal and contextual constraints could be formulated and added into the knowledge base. First of all, let us look at the simpler contextual constraint in the use case. We would like to add a constraint that "no car could be traveling at a high speed if there is a heavy traffic jam". This can be translated to a first order logic formula as follows:

$$\neg holdsAt(HS, T) \leftarrow holdsAt(HJ, T) \quad (7)$$

and

$$\neg holdsAt(HJ, T) \leftarrow holdsAt(HS, T) \quad (8)$$

where HS denotes the *High Speed* fluent of a vehicle traveling at a speed higher than a pre-defined speed threshold. Next, we begin to take the distance attribute into consideration and add the spatio-temporal constraint that "a car appeared at T_2 can neither be traveling (with high speed or low speed) or stopped at time T_1 at another location that is more than D distance away if a feasible speed (a pre-defined speed constant S) could enable it travel D distance from T_1 to T_2". By translating this to a formula, we have:

$$holdsAt(HS, T_2) \leftrightarrow \neg\{\{holdsAt(HS, T_1, D) \vee holdsAt(LS, T_1, D) \vee holdsAt(ST, T_1, D)\} \wedge (T_2 - T_1) * S < D\} \quad (9)$$

where LS denotes the *Low Speed* fluent of a vehicle traveling at a speed lower than the speed threshold, and ST denotes the *Stopped* fluent where a vehicle is stopped when reporting an observation. One thing to be noted is, this formula is only applicable when there is the same vehicle reporting multiple observations. To simplify the formula so that the ground Markov network does not become over-complicated, we do not introduce a separate $Vehicle$ entities in the axioms. Since it is easy to give a pre-judgment of whether a same vehicle reported multiple observations at multiple locations, we only consider Formula 9 in a ground Markov network when such conflicts arise and we make the *vehicle* entity implicit.

Due to space limitations, we only listed the above subset of the formulas for our traffic sensing use case. These formulas set the rules when the *Heavy Jam* traffic fluent and *High Speed* vehicle fluent are the primary fluents. The logical formulas for other fluents and events can be easily generalized.

C. MLN-based Trust Reasoning

In Section III-C, we introduced the basic MLN concepts. In our mobile sensing context, we aim to evaluate the trust (i.e. probability) of a target claim given the other claims in the current evidence window as well as all of their trust values computed based on single-event trust analysis. All the claims are mapped to one or more predicates in MLN based on the

claim payload and provenance information.

The set of random variables in $M_{K,C}$ can be divided into two subsets. One is the set of input ground predicates derived from the other claims in the current evidence window and preprocessed spatio-temporal constraints, which is referred to as the evidence random variables $X \in \mathcal{X}$. The other subset correspond to groundings of query predicates, as well as groundings of any other hidden/unobserved predicates, which is referred to as the query random variables $Y \in \mathcal{Y}$. The joint probability distribution of a possible assignment of $Y = y$, conditioned over a given assignment of $X = x$, is defined as follows:

$$P(Y = y|X = x) = \frac{1}{Z(x)} \exp\left(\sum_i w_i n_i(x, y)\right) \quad (10)$$

where $n_i(x, y)$ is the number of true groundings of the F_i in x and y. $Z(x)$ normalizes over all possible assignments $y' \in \mathcal{Y}$ of query/hidden variables given the assignment x, that is, $Z(x) = \sum_{y' \in \mathcal{Y}} \exp\left(\sum_i w_i n_i(x, y')\right)$.

Exact inference by computing Equation 10 directly is intractable in all but the smallest domains. In order to perform a specific inference task, it is not necessary in general to ground the whole network, as parts of it could have no influence on the computation of the desired probability. Grounding only the needed part of the network can allow significant savings both in memory and in time to run the inference. Inference in the partial ground network can be done by Gibbs sampling.

The inference with MLN is supposed to be done based on past evidence whose value (1 or 0) has been determined. In our context, all the *evidence predicates* derived from the claims in the evidence window are supposed to be under trust evaluation, too. Hence, each evident predicate does not have a confirmed value. Instead, each is only associated with a trust value in the range of [0, 1], which is obtained from either the initial single-event based trust analysis or the multi-event based trust analysis from the previous evidence window. This trust value is defined as the *prior trust*. This makes it difficult to determine whether a grounding of a given formula is true or false. To solve this problem, we first classify the value of each evidence predicate to 1 or 0 by comparing their prior trust with a *neutral trust threshold* (0.5). A predicate with a prior trust higher than the neutral trust threshold is considered true. With this classification, all the evidence predicates are fed into the MLN to obtain a probability value of the any particular target predicate. This probability can then be used to update the prior trust of a target predicate as follows:

$$T(x + 1) = T(x) + \lambda \cdot (p - 0.5) \quad (11)$$

where $T(x)$ is the prior trust, $T(x+1)$ is the updated trust, p is the *inferred probability* obtain from MLN, and λ is the *logical sensitive parameter* which controls how much influence our logical analysis has on the resulting trust value. The rationale behind Eqn. 11 is that the trust of a claim (or predicate) should increase if there is more support from other evidence claims such that the MLN inference yields a probability higher than 0.5 and decrease otherwise.

V. Evaluation

In this section, we aim to study the performance of our scheme and understand how the dynamics and system parameters would affect the trust analysis results via simulation experiments. We tested two dimensions of performance: effectiveness and efficiency. We assume the Event Calculus knowledge base has been pre-established given a specific application scenario and real time computation of the logic module only involves mapping the claims to events and fluents of the knowledge base. In terms of effectiveness, since we have not seen another scheme which tries to add multi-event logical analysis to the information trust analysis in the mobile sensing context, we do not have a reference system to compare with. Therefore, we first concentrate on studying the impact of the dynamics and system parameters (e.g., the available relevant claims in the evidence window and the weight assigned to the formulas in the knowledge base) on the effectiveness of trust analysis. There is no expensive computation before the MLN inference. Therefore, for efficiency, we focus on the delay caused by the MLN inference by running test cases with different complexity of the knowledge base and different number of available evidence claims.

TABLE IV: Default simulation settings

Total number of formulas in knowledge base	100
Number formulas relevant to target event	10
Number of evidence claims relevant to target event	10
Weight of relevant formulas	0.5

In our simulation experiments, we define a complete Event Calculus knowledge base for our traffic sensing use case in accordant with the events and fluents defined in Table II. All the simulations are under a closed-world assumption, i.e., all relevant events are defined in the knowledge base. For each of our experiments, we set one target claim (e.g., heavy jam) and create different amounts of other claims as evidence claims (e.g., major accident, hazard, high speed, etc.). We use the open-source statistical relational AI framework Alchemy [23] for executing MLN inference and our testings are carried out on a Linux machine (Ubuntu 12.04 64-bit) with 2.26 GHz Intel Core i5 processor and 6 GB 1333 MHz DDR3 memory. Each of the data points in our evaluation results are obtained based on 30 executions. Table IV lists our default simulation settings. Each of the parameters is varied in our experiments to study their impact on the performance.

A. Inferred Probability

The first performance metric we evaluate is the effectiveness of MLN inference in our proposed framework, i.e., the inferred probability p in Eqn 11. It is a determinative factor for the trust adjustment. To understand if our scheme can really reduce uncertainties, we need to investigate how the system dynamics affect the inferred probability.

For a particular target claim, there could be both supporting and conflicting evidence claims in the evidence window. For the simulated experiments of inferred probability, we change the ratio of the supporting and conflicting claims in order to get a thorough understanding of its impact. Given the varied supporting claim ratio, we then vary two other important factors: the number relevant evidence claims (denoted as EN) and the weight of the relevant formulas (denoted as WT). The results for the former are shown in Figure 4 and the results for the later are shown in Figure 5.

As expected, the supporting claim ratio determines if the inferred probability is higher or lower than 0.5, that is, if the

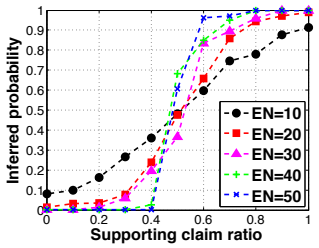

Fig. 4: Impact of supporting evidence ratio and number of evidence claims on inferred probability

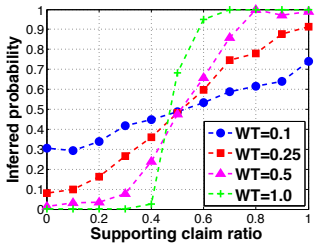

Fig. 5: Impact of supporting evidence ratio and formula weight on inferred probability

target claim gets positive or negative trust adjustment based on Eqn. 11. In addition, both EN and WT have significant influence on the actual inferred probability given a certain supporting claim ration. Both Figure 4 and Figure 5 show similar impact of EN and WT: higher EN or WT leads to more diverged distribution of the inferred probability for different supporting claim ratio. This is determined by the nature of the MLN calculation. As we explained, the higher the value of a formula's weight, the stronger the constraint the formula imposes over the possible worlds. It is thus anticipated that higher WT pushes the inferred probability to either 1 or 0 when there is a quantity difference between the supporting claims and the conflicting claims. Furthermore, each evidence claim represents the same level of constraint given a certain weight to its associated formula. With the same supporting claim ratio, the quantity difference between the supporting claims and the conflicting claims becomes larger when EN is larger. Hence, when WT is fixed, more supporting claims also push the inferred probability towards 1 and more conflicting claims push the inferred probability to 0. This explains the sensitivity of the inferred probability to both EN and WT. With the knowledge specific to the domain under consideration, proper weights can be carefully assigned and thus the inferred probability can provide useful constraints and effectively reduce uncertainties and ambiguities in the trust analysis of the sensing claims.

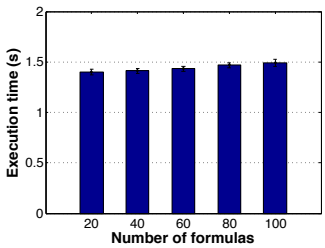

Fig. 6: Execution time against total number of formulas

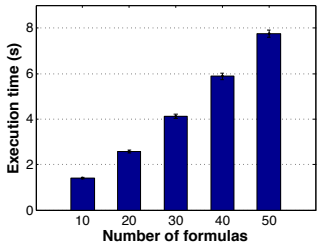

Fig. 7: Execution time against number of relevant formulas

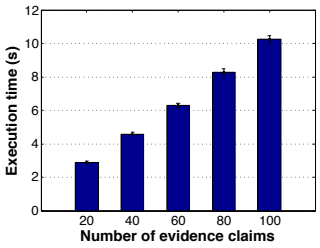

Fig. 8: Execution time against number of evidence claims

B. Execution Time

Other than effectiveness evaluation, we also measure the execution time of MLN inference with Alchemy by varying the total number of formulas in the knowledge base, the number of relevant formulas in the knowledge base as well as the number of relevant claims in the evidence window.

Figure 6 shows our execution time measurement result when we vary the total number of formulas. In this experiment, we only increase the number of formulas that are irrelevant to the target claim in order to see if they would affect the execution time. From the results we can see that there is only a slight increment of the execution time. Figure 7 shows the changes of the execution time when we change the number of formulas that are relevant to the target claim, i.e., those formulas that defined the spatial, temporal or contextual relations between the target claim (mapped to an event/fluent predicate) and other events and/or fluents. We can observe a much more dramatic increment of the execution time when we increase the number of relevant formulas. Finally, Figure 8 shows the changes of the execution time when we change the number of evidence claims, i.e., the claims which have spatial, temporal or contextual relations with the target claim in the evidence window. Similar to Figure 7, the execution delay increases quickly with the number of evidence claims.

From the above results, we can conclude that the major factors that determine the delay of the MLN inference are the number of relevant formulas and relevant evidence claims. The irrelevant formulas in the knowledge base do not incur significant delays. This is a great benefit from Gibbs sampling, i.e., inference in a partial ground Markov network. To clearly demonstrate the increased delay in Figure 7 and Figure 8, we introduced "forged" formulas and evidence claims. In fact, for one particular target claim, the number of relevant formulas or evidence claims is less than 10 in most cases. Hence, the MLN inference is generally rather efficient.

VI. Discussion and Future Work

Our proposed framework is designed with an objective of serving as an add-on to single-event based trust analysis schemes. Trust adjustments are made for each claim reported

by participating users based on the level of support or conflict obtained from other relevant observations. The scale of such adjustments can be controlled depending on the requirements of the specific application. It, however, does not prevent our scheme to be used independently without a single-event based trust analysis scheme. Our scheme is general enough to be applied to cases where logical relations among events is the only determinative factor.

As a validation for the proposed framework, our performance testing included evaluation for both efficiency and effectiveness. As a first attempt to solve the multi-event based trust problem in mobile sensing, we kept our evaluation general via simulated scenarios. We did not explicitly compare our result with a particular existing single-event based scheme as our approach can be integrated on top of any of those schemes. Without a real application scenario, there is no easy way to evaluate the exact amount of improvement in terms of trust analysis accuracy our approach can achieve on top of a single-event based scheme. However, from our effectiveness evaluation, we can see that our scheme lowers the trust values the claims that have more conflicts with other claims and increases the trust values of claims that have more support from other claims as expected. The magnitude of such adjustments is determined by the number of evidence claims and the strength of the pre-defined logic constraints (i.e., formula weights) between the claimed events. The actual system designer has the control over the sensitivity of the trust adjustments to these factors. This achieves our design objectives of dynamically enhancing the accuracy of trust analysis by jointly recognizing and linking events from successive sensing observations.

In the performance evaluation, we varied the formula weight settings of MLN and tested the impact of different settings on the trust evaluation. In fact, MLN has the capability of learning the weights of formulas when a training data set is provided. In this work, we focused on the inference with MLN and did not consider its weight learning capability. As a part of our future work, we will look into the different weight learning mechanisms of MLN and carry out more measurements with formula weights learned from real datasets.

VII. Conclusion

Data collected with the mobile sensing paradigm must be assessed in terms of reliability or trustworthiness. We believe that logical knowledge specific to an application domain under consideration can provide useful constraints to reduce uncertainties and ambiguities in the trust evaluation. In this paper, on top of the traditional single-event based trust analysis, we propose a novel trust evaluation framework based on the mutual logical support or conflict between successive sensing claims. Our approach integrates logical reasoning via an extended Even Calculus and the Markov Logic Network. With an example traffic sensing use case, we illustrated how to construct an Event Calculus knowledge base based on commonsense knowledge of the domain and make it location and context aware. In addition, we describe how the Markov Logic Network is used in our scheme to infer the trust of a particular event with other events as evidence. Performance evaluation results have shown that our scheme can serve as an efficient and effective add-on to single-event based trust analysis.

VIII. Acknowledgments

This work was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-09-2-0053. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official polices, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copy right notation here on.

References

[1] B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. Miu, E. Shih, H. Balakrishnan, and S. Madden, "CarTel: a distributed mobile sensor computing system," in *Proceedings of ACM SenSys*, pp. 125–138, 2006.

[2] R. Rana, C. Chou, S. Kanhere, N. Bulusu, and W. Hu, "Ear-phone: an end-to-end participatory urban noise mapping system," in *Proceedings of ACM/IEEE IPSN*, pp. 105–116, 2010.

[3] S. Eisenman, E. Miluzzo, N. Lane, R. Peterson, G. Ahn, and A. Campbell, "BikeNet: A mobile sensing system for cyclist experience mapping," *ACM Transactions on Sensor Networks*, vol. 6, 2009.

[4] L. Deng and L. Cox, "Livecompare: grocery bargain hunting through participatory sensing," in *Proceedings of ACM HotMobile*, pp. 1–6, 2009.

[5] C. Castillo, M. Mendoza, and B. Poblete, "Information credibility on twitter," in *Proceedings of the ACM WWW*, pp. 675–684, ACM, 2011.

[6] M. Gupta and P. Zhao, "Evaluating Event Credibility on Twitter," in *Proceedings of SIAM International Conference on Data Mining*, 2012.

[7] H.-S. Lim, Y.-S. Moon, and E. Bertino, "Provenance-based trustworthiness assessment in sensor networks," in *Proceedings of International Workshop on Data Management for Sensor Networks*, pp. 2–7, ACM, 2010.

[8] X. Wang, K. Govindan, and P. Mohapatra, "Collusion-resilient quality of information evaluation based on information provenance," in *Proceedings of the IEEE SECON*, 2011.

[9] X. Wang, J.-H. Cho, K. Chan, M. Chang, A. Swami, and P. Mohapatra, "Trust and Independence Aware Decision Fusion in Distributed Networks," in *Proceedings of IEEE PERCOM Workshops*, 2013.

[10] D. Wang, L. Kaplan, H. Le, and T. Abdelzaher, "On truth discovery in social sensing: a maximum likelihood estimation approach," in *Proceedings of ACM/IEEE IPSN*, pp. 233–244, ACM, 2012.

[11] X. Wang, W. Cheng, P. Mohapatra, and T. F. Abdelzaher, "ARTSense: anonymous reputation and trust in participatory sensing," in *Proceedings of IEEE INFOCOM*, 2013.

[12] "Waze." http://www.waze.com/.

[13] R. Kowalski and M. Sergot, "A logic-based calculus of events," *New generation computing*, vol. 4, no. 1, pp. 67–95, 1986.

[14] M. Richardson and P. Domingos, "Markov logic networks," *Machine learning*, vol. 62, no. 1, pp. 107–136, 2006.

[15] X. Yin, J. Han, and P. S. Yu, "Truth discovery with multiple conflicting information providers on the web," *IEEE Transactions on Knowledge and Data Engineering*, vol. 20, no. 6, pp. 796–808, 2008.

[16] M. Raya, P. Papadimitratos, V. Gligor, and J. Hubaux, "On data-centric trust establishment in ephemeral ad hoc networks," in *Proceedings of IEEE INFOCOM*, pp. 1238–1246, 2008.

[17] X. Wang, K. Govindan, and P. Mohapatra, "Provenance-Based Information Trustworthiness Evaluation in Multi-Hop Networks," in *Proceedings of IEEE GLOBECOM*, vol. 10, pp. 1–5, 2010.

[18] J. McCarthy and P. J. Hayes, "Some Philosophical Problems from the Standpoint of Artificial Intelligence," *Machine Intelligence*, vol. 4, 1969.

[19] W. Brendel, A. Fern, and S. Todorovic, "Probabilistic event logic for interval-based event recognition," in *Proceedings of IEEE CVPR*, pp. 3329–3336, 2011.

[20] K. P. Murphy, "Dynamic bayesian networks," *Probabilistic Graphical Models, M. Jordan*, 2002.

[21] R. Kindermann, J. L. Snell, *et al.*, *Markov random fields and their applications*. American Mathematical Society Providence, RI, 1980.

[22] H. Chaudet, "Extending the event calculus for tracking epidemic spread," *Artificial Intelligence in Medicine*, vol. 38, no. 2, pp. 137–156, 2006.

[23] S. Kok, M. Sumner, M. Richardson, P. Singla, H. Poon, D. Lowd, J. Wang, and P. Domingos, "The alchemy system for statistical relational AI," tech. rep., Department of Computer Science and Engineering, University of Washington, Seattle, WA, 2009.