Handling Sequential Observations in Intelligent Surveillance

Jianbing Ma, Weiru Liu, and Paul Miller

School of Electronics, Electrical Engineering and Computer Science,
Queen's University Belfast, Belfast BT7 1NN, UK
{jma03,w.liu}@qub.ac.uk, p.miller@ecit.qub.ac.uk

Abstract. Demand for intelligent surveillance in public transport systems is growing due to the increased threats of terrorist attack, vandalism and litigation. The aim of intelligent surveillance is in-time reaction to information received from various monitoring devices, especially CCTV systems. However, video analytic algorithms can only provide static assertions, whilst in reality, many related events happen in sequence and hence should be modeled sequentially. Moreover, analytic algorithms are error-prone, hence how to correct the sequential analytic results based on new evidence (external information or later sensing discovery) becomes an interesting issue. In this paper, we introduce a high-level sequential observation modeling framework which can support revision and update on new evidence. This framework adapts the situation calculus to deal with uncertainty from analytic results. The output of the framework can serve as a foundation for event composition. We demonstrate the significance and usefulness of our framework with a case study of a bus surveillance project.

Keywords: Intelligent Surveillance; Active System; Situation Calculus; Belief Change; Sequential Observation.

1 Introduction

Recently, more and more attention has been paid by governments and transport operators to protect vehicles and passengers with surveillance cameras, e.g., the Florida School Bus Surveillance project [2], the First Glasgow Bus Surveillance [22], Federal Intelligent Transportation System Program in the US [18], Washington rail corridor surveillance [17], Airport Corridor Surveillance in the UK [16], etc. These applications require deployment of large-scale CCTV systems giving rise to unique problems. For example, in a reasonable sized provincial city there may be several hundred buses, each of which has 12-14 cameras, giving a total of several thousand cameras. The large amount of cameras for monitoring passengers/vehicles makes it almost impossible to detect possible incidents manually without delay. For this reason, video analytics and event reasoning are being introduced to CCTV systems in order to ensure in-time reaction.

The aim of video analytics is single-event recognition. Recently, however, developers have realized that it is necessary to manage the events generated by video analysis software. For instance, to prevent anti-social behaviors on public transport systems, one has to make a decision based on a sequence of detected events. Ideally, this would be straightforward if the recognized events are correct and certain. Unfortunately, in reality, imperfection and mistakes frequently occur in practical applications. For example, in the case of a person entering the bus doorway, the person may be classified as male with a certainty of 85% by the classification analytics, rather than with 100% certainty. Even worse, the analytic algorithm may classify a person as female at one time instant and male at a later time. In addition to inaccurate analytics, a large amount of mistakes are caused by the unreliability of the data sources. For example, in the classification example above, the camera may have been tampered with, illumination could be poor, or the classifier training set may be unrepresentative. Any, or all of these, can result in imperfection and errors.

In this paper, we introduce a sequential observation modeling framework which is able to deal with uncertainty, and can support revision and update when new evidence is received, thereby removing the influence of past errors. This framework, which operates at a higher level than analytic algorithms, deploys a situation calculus foundation with the ability to deal with uncertainty in analytic results. It is also able to handle belief revision and update properly. We demonstrate the significance and usefulness of our framework with a case study of a bus surveillance project [12, 7, 10, 15]. Our approach provides a sound framework for surveillance applications, such as CCTV for buses, airports, etc. The output of the framework, i.e., primitive events, can be used as a starting point of event composition.

Usually in situation calculus, there are sensing and non-sensing actions, or epistemic and ontic actions [4]. A sensing/epistemic action senses a property of the domain and does not change the environment. A non-sensing/ontic action is an action done by the agent which changes the environment. A major difference between real-world situations, such as those encountered in surveillance applications, and situation calculus approaches is that in the former, even the result of an ontic action, e.g., a passenger changes their position from standing to seated, is observed by cameras and analyzed by video analytic algorithms (and hence sometimes we call the agent of interest an *observable*). Therefore, situation calculus should be significantly adapted to make it suitable for intelligent surveillance purposes.

In intelligent surveillance applications the results of both epistemic and ontic actions are provided by video analytics, therefore, we must differentiate between both kinds of actions, since they respond differently upon new evidence being obtained. The properties sensed by epistemic actions are generally invariable, e.g., the gender of a person, etc., whilst properties related to ontic actions are those that can be changed at will, e.g., the position of a person, etc. We also allow external information to be handled in this framework. External information, when received and used, can be seen as a kind of epistemic or ontic action, according to the information properties. For instance, if a piece of information tells us a person is a male, then it can be seen as an epistemic action; if it tells us a person is standing, it can be seen as an ontic action. In addition, a property related to an epistemic action will be called an **epistemic** or an **invariable** property, and similarly, a property related to an ontic action is called an **ontic** or a **variable** property. In summary, if a property indicates an intrinsic character of an observable of interest, and hence is invariable, then it is an epistemic property. But we also need to point out that this property could be mis-classified or even intentionally disguised, which seemingly makes it variable. For example, although a person is a male, he could be wrongly classified as a female. He could even disguise himself as a female if he wants to. However, this superficial variability should not cause any confusion. Instead, ontic properties are usually *external* properties between an observable of interest and the environment, and hence are variable, e.g., a passenger can move from the drivers cabin area to the saloon area on a bus. This differentiation between epistemic and ontic properties also applies to properties obtained by video analytics. For instance, if the gender of a person is in fact estimated from the captured video, it is still called an epistemic property since gender is an intrinsic character of a person.

The rest of the paper is organized as follows. Section 2 provides the preliminaries on the situation calculus. In Section 3, formal approaches to deal with uncertain observations are presented, including the ways in which to handle epistemic and ontic actions. Section 4 shows how belief revision is adequately handled in our framework. We then provide a case study, which is a simplified bus surveillance scenario, in Section 5. Finally, we conclude the paper in Section 6.

2 Preliminaries

Situation calculus, introduced by John McCarthy [13, 14], has been applied widely to model and reason about actions and changes in dynamic systems. It was reinterpreted in [19] as *basic action theories* which are comprised of a set of foundational axioms defining the space of situations, unique-name axioms for actions, action preconditions and effects axioms, and the initial situation axioms [5]. The well known frame problem is solved by a set of special action effects axioms called *successor state axioms*.

Since actions carried out by agents cause constant changes of the agents' beliefs, developing strategies of managing belief changes triggered by actions is an important issue. The problem of *iterated belief change* within the framework of situation calculus has been investigated widely, e.g., [20, 21, 24, 11]. In [24], a new framework extending previous approaches was proposed, in which a plausibility value is attached to every situation. This way, the framework is able to deal with nested beliefs, belief introspection, mistaken beliefs, and it can also handle belief revision and update together in a seamless way. The framework in [24] is based on an extension of action theory [19] stemming from situation calculus [13, 14]. Here we introduce the notion of situation calculus from [24] which includes a belief operator [20, 21].

According to [24], the situation calculus is a predicate calculus language for representing dynamically changing domains. A situation represents a snapshot of the domain. There is a set of initial situations corresponding to what the agents believe the domain might be initially. The actual initial state of the domain is represented by a distinguished initial situation constant, S_0, which may or may not be among the set of initial situations believed by an agent. The term $do(a, s)$ denotes the unique situation that results from the agent performing action a in situation s.

Predicates and functions whose values may change from situation to situation (and whose last argument is a situation) are called *fluents*. For instance, we use the fluent

$\mathsf{InR1}(s)$ to represent that the agent is in room $R1$ in situation s. The effects of actions on fluents are defined using successor state axioms [19], which provide a succinct representation for both effect axioms and frame axioms [13, 14]. For instance, if there are two rooms ($R1, R2$) and an action Leave takes the agent from the current room to the other room. Then, the successor state axiom for $\mathsf{InR1}$ is [24]:

$$\mathsf{InR1}\big(do(a,s)\big) \equiv \big(\big(\neg\mathsf{InR1}(s) \wedge a = \mathsf{Leave}\big) \vee \big(\mathsf{InR1}(s) \wedge a \neq \mathsf{Leave}\big)\big).$$

This axiom says that the agent will be in Room 1 after doing action a in s iff either it is in Room 2 and leaves for Room 1 or is currently in Room 1 and does not leave.

Levesque [6] introduced a predicate, $\mathsf{SF}(a,s)$, to describe the result of performing the binary-valued epistemic action a. $\mathsf{SF}(a,s)$ holds (returns *true*) iff the sensor associated with a returns the sensing value 1 in situation s. Each epistemic action senses some property of the domain. The property sensed by an action is associated with the action using a *guarded sensed fluent axiom* [3]. For example, the following two axioms

$$\mathsf{InR1}(s) \rightarrow \big(\mathsf{SF}(\mathsf{SenseLight}, s) \equiv \mathsf{Light1}(s)\big)$$

$$\neg\mathsf{InR1}(s) \rightarrow \big(\mathsf{SF}(\mathsf{SenseLight}, s) \equiv \mathsf{Light2}(s)\big)$$

can be used to specify that $\mathsf{SenseLight}$ senses whether the light is on in the room the agent is currently located.

In this paper, we adopt the following conventions about *guarded action theories* Σ consisting of: (A) successor state axioms for each fluent, and guarded sensed fluent axioms for each action; (B) unique names axioms for actions, and domain-independent foundational axioms; and (C) initial state axioms which describe the initial state of the domain and the initial beliefs of agents. A *domain-dependent fluent* means a fluent other than the probability fluent p, and a *domain-dependent formula* is one that only mentions domain-dependent fluents. However, since this is a paper focusing on applications, we will not introduce the axioms here. Interested readers can refer to [24, 11]. We further assume that there is only one agent acting in a chosen domain, although the framework is capable of accommodating multiple agents.

3 The Revised Situation Calculus Framework

In this section, we extend the situation calculus to include a probability operator to account for iterated belief changes and deal with uncertainty.

Usually in situation calculus, the result of all actions are accurate. In recent work, e.g., [1, 23, 24, 11], noisy epistemic actions have been proposed and studied. However, in intelligent surveillance applications, not only can epistemic actions be noisy, but ontic actions can also be subject to noise (recall an epistemic action senses some property of the domain, but leaves the environment unchanged, while an ontic action changes the actual environment). That is, in these applications, the results of ontic actions are also reported by video analytic algorithms, which may (and in fact usually, if not always) present uncertain results. For example, if a passenger takes a seat, then from a normal situation calculus point of view, its status certainly changes from "standing" to "seated".

However, if the scenario is analyzed by an algorithm, due to the imperfection of the algorithm, it can only give 90% degree of certainty that the passenger is seated, leaving the remaining 10% still standing. That is, ontic actions can also bring uncertainty, or noise. In fact, in a few scenarios (e.g., light changes suddenly outside the window), the analytic results could be very inaccurate. It may conclude that with 60% degree of certainty the passenger is seated and 40% degree of certainty the passenger is standing. These cases cannot be handled in classical situation calculus. Hence, we need to adapt the situation calculus to deal with such cases.

For convenience, we denote SA the set of all epistemic actions and hence for each action a, $a \in SA$ means that a is an epistemic action while $a \notin SA$ means that a is a ontic action. Since an ontic action a can bring up more than one possible result (e.g., “seated” or ”standing”), the corresponding $do(a, s)$ may also give rise to more than one successive situation. An epistemic action a can also bring up more than one possible result if it is not accurate.

Example 1 *Let a situation $s = M \wedge S$ (the passenger is male and standing), then a noisy (inaccurate) epistemic action presents a result as the passenger is male with probability 0.4 and female with probability 0.6, then two successive situations $s^1 = M \wedge S$ with probability value 0.4 and $s^2 = \neg M \wedge S$ with probability 0.6 should be expected.*

Hence, subsequently in this paper, we assume $do(a, s)$ is a set of situations instead of a single situation. Moreover, in this paper, we assume $\mathsf{SF}(a, s)$ (different from the definition in [6] where $\mathsf{SF}(a, s)$ returns a boolean value) gives the tuple-valued sensing result $(x_1, \cdots, x_k)$ where each x_i stands for the probability that the epistemic action returns result X_i. For instance, in the above example, $\mathsf{SF}(a, s) = (0.4, 0.6)$ means that the passenger is male with probability 0.4 and female with probability 0.6 when a is an epistemic action returning the gender of the passenger. For convenience, we also write $\mathsf{SF}(X, a, s)$ to denote the probability that the epistemic action a returns X, e.g., $\mathsf{SF}(M, a, s) = 0.4, \mathsf{SF}(\neg M, a, s) = 0.6$. Similarly, we write $\mathsf{NSF}(a, s)$ (NS is short for Non-Sensing) to present the tuple-valued ontic action result $(x_1, \cdots, x_t)$ where each x_i stands for the probability that the ontic action returns result X_i. For instance, if a is an ontic action changing the behavior of a passenger, then $\mathsf{NSF}(a, s) = (0.2, 0.8)$ means that there is a probability 0.2 such that a passenger is standing and a probability 0.8 such that it is seated (¬S). Similarly, we also denote $\mathsf{NSF}(X, a, s)$ the probability that the ontic action a returns X.

In this paper, for simplicity, we assume that all actions, regardless of whether they are epistemic or ontic, can only provide two possible results. In fact, if they could return more than two possible results, no essential changes are needed for the framework, but a more cumbersome description of the scenarios, e.g., $\mathsf{SF}(a, s)$ will be a n-tuple value where $n > 2$ and the number of successive situations will become greater, etc..

In this paper, we use ordinals as time points to indicate the sequence of situations. More precisely, all the initial situations will have a subscript 0, denoted as s_0^i (where i indicates the i-th possible situation), and the successive situation of a situation s_n will be s_{n+1}. Let $\mathcal{S}_n$ denote the set of all situations with subscript n, i.e., the set of situations in the n-th run. Note that if s and s' are both in $\mathcal{S}_n$, then we should have

$\mathsf{SF}(a,s) = \mathsf{SF}(a,s')$ (resp. $\mathsf{NSF}(a,s) = \mathsf{NSF}(a,s')$) since the action is taken in the real world, it should return only one result, no matter what we *think* the real-world situation might be (e.g., s or s'). From this sense, we can write $\mathsf{SF}_n(a)$ or $\mathsf{NSF}_n(a)$ to denote the action result for situations in the n-th run. Furthermore, we use s/X to denote a situation that the property corresponding to X in s is changed to X, e.g., for $s = M \wedge S$, we have $s/M = s$ and $s/\neg M = \neg M \wedge S$.

The belief set of $\mathcal{S}_n$ is defined as follows. Let $\phi[s]$ denote that ϕ is assessed in s. For example M (the passenger is a male) is assessed in $s = M \wedge S$, hence $M[s]$ holds. Let $p_n(\phi) = \sum_{s:s\in\mathcal{S}_n\wedge\phi[s]} p(s)$ indicate the total probability of ϕ in $\mathcal{S}_n$.

Definition 1 $Bel_n(\phi) \stackrel{def}{=} p_n(\phi) > p_n(\neg\phi)$.

That is, ϕ is believed in the n-th run if it is more probable than its negation. Since this definition is not closed under deduction, i.e., $Bel_n(\phi) \wedge Bel_n(\psi) \nrightarrow Bel_n(\phi \wedge \psi)$, we usually only consider probabilities (and hence beliefs) on atoms (e.g., $Male$, $Stand$, etc.), while probabilities (and hence beliefs) of other formulae are computed from probabilities of atoms (with independence assumptions) [9].

Based on the above notations and definitions, we can define a probability function p for each situation s to measure how possible an agent considers s is. The p functions for initial situations are provided with a normalization condition that the sum of probabilities of all initial situations is 1. This is expressed as follows:

Axiom 1 *(Initial State Axiom)* $\sum_{s:\mathsf{Init}(s)} p(s) = 1$.

Probabilities of successor situations are defined as follows.

If a is an epistemic action and $\mathsf{SF}(a, s_n) = (t, 1-t)$, or $\mathsf{SF}(X, a, s_n) = t$ and $\mathsf{SF}(\neg X, a, s_n) = 1-t$, then in general it induces two successive situations for s_n, i.e., $s_{n+1} = s_n/X$ and $s'_{n+1} = s_n/\neg X$ with probability $p(s_n)t$ and $p(s_n)(1-t)$ respectively. Note that if $t = 0$ or $t = 1$, it in fact only induces one successive situation (situations with probability 0 will be ignored).

However, it is not always reasonable to simply change the current situation to the successive situations as stated above. In some scenarios, we must keep the probabilities of the beliefs induced by the current situations. For instance, assume that at $\mathcal{S}_k$, a passenger is classified as a male with probability 0.8, and at $\mathcal{S}_{k+1}$, this passenger is classified as a male with probability 0.6, then do we need to change the probability to 0.6? The answer is no. In real-world applications such as intelligent surveillance, we observe that if a video analytic algorithm is used to continuously check the gender of a person based on a video, then the probability of that person being a male will fluctuate. Hence in practice, if at some time point, it is classified as a male with probability 0.9, and later with probability 0.85, we can just keep the probability 0.9. A more persuasive scenario is that at some time point we have external information (e.g., an analyst views the person on a monitor) which indicates that the passenger is 100% a male, but later the algorithm still classifies it as a male with probability 0.85, then it is obvious that we do not need to change the probability from 1 to 0.85.

An exception to the above statement, is that the change of probability leads to a change of beliefs. For example, if at $\mathcal{S}_k$, a passenger is classified as male with probability 0.8 (hence $Bel_k(M)$ holds), but at $\mathcal{S}_{k+1}$, it is classified as female with probability

0.75 (hence $Bel_{k+1}(\neg M)$ holds), then this major change should not be ignored. It may indicate that an interesting event has happened. In real systems, such belief changes with respect to an invariable property, may, by themselves, justify alerting an analyst.

In short, the probabilities of the current belief are only changed when the sensing result overwhelms the current belief. Here by overwhelm we mean either the belief induced by the sensing result is the same as the current belief but with a greater probability or the belief becomes different.

More precisely, let a be an epistemic action and $\mathsf{SF}(a, s_n) = (t, 1-t)$. Without loss of generality, we assume that $t > 0.5$[1] and hence $Bel_{n+1}(X)$ holds. For each situation $s_n \in \mathcal{S}_n$, it induces two successive situations, i.e., $s_{n+1} = s_n/X$ and $s'_{n+1} = s_n/\neg X$. The probabilities of these two situations are defined as follows:

- if $Bel_n(\neg X)$ holds, then $s_{n+1} = s_n/X$ and $s'_{n+1} = s_n/\neg X$ are assigned with probability $p(s_n)t$ and $p(s_n)(1-t)$ respectively,
- if $Bel_n(X)$ holds, then $s_{n+1} = s_n/X$ and $s'_{n+1} = s_n/\neg X$ are assigned with probability $p(s_n)max(p_n(X), t)$ and $p(s_n)(1 - max(p_n(X), t))$ respectively.

For this we have the following result which shows that the assignment of probabilities satisfy the statements we argued before.

Proposition 1 *For any epistemic property X and $n > 0$, if both $Bel_n(X)$ and $Bel_{n+1}(X)$ hold, then $p_{n+1}(X) = max(p_n(X), \mathsf{SF}(X, a, s_n))$. If both $Bel_n(\neg X)$ and $Bel_{n+1}(X)$ hold, then $p_{n+1}(X) = \mathsf{SF}(X, a, s_n)$.*

There might be some equivalent situations in $\mathcal{S}_{n+1}$ (in terms of all fluents). For convenience, they can be merged together.

Example 2 *Assume that a video analyzer detects a passenger on board but it does not know whether the passenger is male or female. So it considers two possible situations S_0^1 and S_0^2 at the beginning where*

$$S_0^1 = \mathsf{Male} \wedge \mathsf{Stand}, S_0^2 = \mathsf{Female} \wedge \mathsf{Stand}$$

The video analytic algorithm gives S_0^1 with probability 0.8 and S_0^2 with probability 0.2. The bottom-half of Fig. 1 illustrates these two situations.

After some seconds, the camera does a second detection from which the video analytic algorithm asserts that the passenger is male with probability 0.9 and female with probability 0.1. In Fig. 1, Sensing Gender *is abbreviated as* SG. *Hence each situation induces two successive situations (in Fig. 1, 0.18=0.2*0.9, etc.) and then equivalent situations are merged together, which finally forms two situations S_1^1 and S_1^2 in Fig. 1.*

It seems that we can let each initial situation in Fig. 1 only induce one successive situation and simply change the probabilities of the successive situations to 0.9 and 0.1, respectively. The reason why we do not follow this way is that the latter approach is not applicable on some occasions. For instance, if there is only one initial situation S_0^1, then from the second detection, probability 0.1 should be assigned to a successive situation that the passenger is a female but no initial situations can induce such a successive situation.

[1] In practice, we can always change $(0.5, 0.5)$ to $(0.5+\epsilon, 0.5-\epsilon)$ for a small positive real ϵ. It does not make much difference.

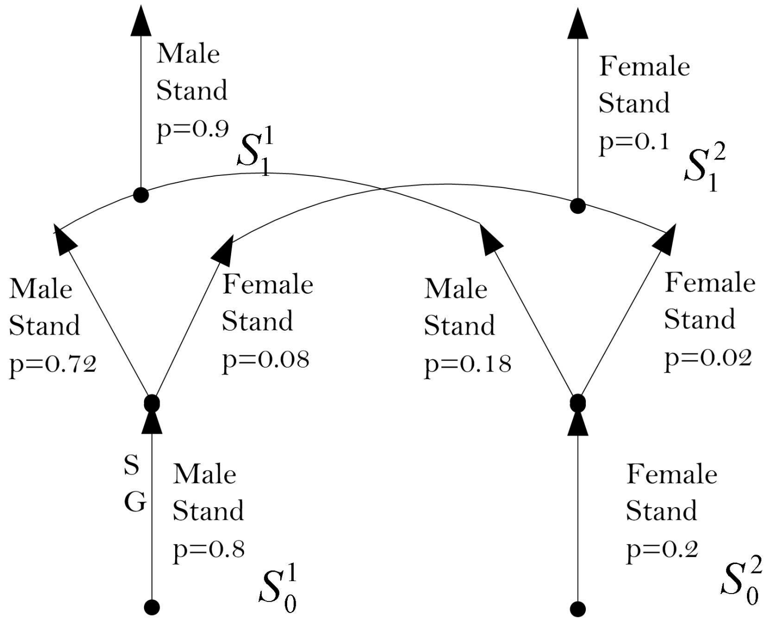

Fig. 1. Situations after Sensing the Gender of a Passenger

If a is an ontic action and $\mathsf{NSF}(a, s) = (t, 1-t)$, then for each situation $s_n \in \mathcal{S}_n$, it induces two successive situations $s_{n+1} = s_n/X$ and $s'_{n+1} = s_n/Y$ with probability $p(s_n)t$ and $p(s_n)(1-t)$ respectively. Similarly, if there are equivalent induced situations, then we can merge them.

We have the following result.

Proposition 2 *For any ontic property X and $n > 0$, $p_{n+1}(X) = \mathsf{NSF}(X, a, s_n)$.*

Example 3 *Assume we have two possible initial situations S_0^1 and S_0^2 at the beginning where*

$$S_0^1 = \mathsf{Male} \wedge \mathsf{Stand}, S_0^2 = \mathsf{Male} \wedge \neg\mathsf{Stand}$$

The video analytic algorithm gives S_0^1 with probability 0.8 and S_0^2 with probability 0.2. The bottom-half of Fig. 2 illustrates these two situations (Note that this figure is similar to Fig. 1 except that the action is an ontic action).

After some seconds, the camera does a second detection from which the video analytic algorithm asserts that the passenger takes a seat with probability 0.9 and it is standing with probability 0.1. In Fig. 2, Sensing Position *is abbreviated as* SP. *Using the above method, finally we get two situations S_1^1 and S_1^2 in Fig. 2.*

For any epistemic or ontic actions, the revised probabilities always sum up to 1.

Proposition 3 *For any $n > 0$, $\sum_{s \in \mathcal{S}_n} p(s) = 1$.*

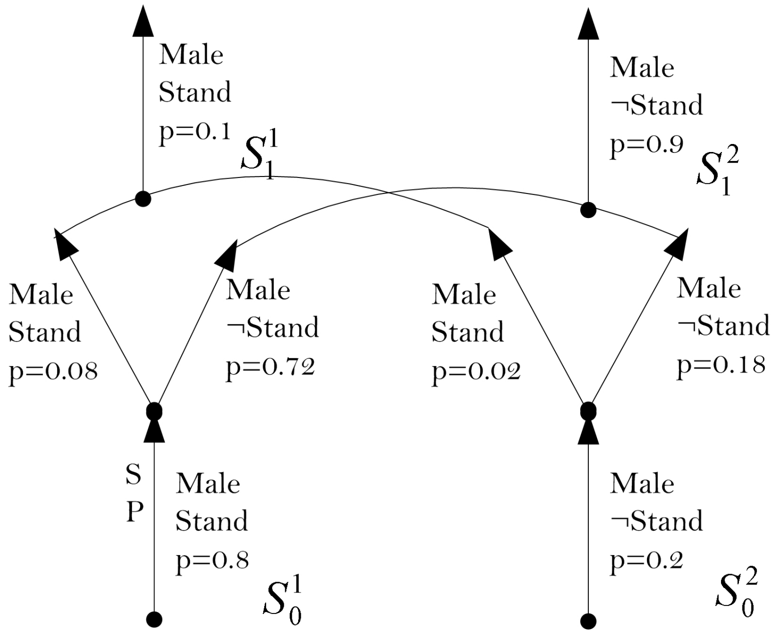

Fig. 2. Situations after Sensing the Position of a Passenger

4 Belief Revision

Belief revision studies how an agent's beliefs can be changed based on some new information if the new information must be believed. Any property of interest (no matter epistemic properties or ontic properties) could be revised when obtaining certain new information on that property of the observable. Studying belief revision in situation calculus is a natural course for managing an agent's beliefs. In the following, we assume that for each formula ϕ to be revised, there is a corresponding action that obtains information on that property.

Definition 2 *(Uniform formula, adapted from [24, 11]) A formula is* uniform *if it contains no unbound variables.*

Definition 3 *A uniform formula ϕ is called obtainable from an action A with regard to a situation s, denoted: $(A, s) \to \phi$, if*

$$\begin{cases} SF(\phi, A, s) > SF(\neg\phi, A, s), & A \textit{ is an epistemic action,} \\ NSF(\phi, A, s) > NSF(\neg\phi, A, s), & \textit{otherwise.} \end{cases}$$

A is called a revision *action for ϕ w.r.t. Σ, if for any s, $(A, s) \to \phi$.*

Note that here the meaning of *revision* is in fact extended to updating as it also handles changes of ontic properties as ontic actions changes the environment[2].

Now by abuse of notation, we use $Bel(\phi, s)$ to denote $Bel_n(\phi)$ where s is a situation in the n-th run.

Theorem 1 *Let ϕ be a domain-dependent, uniform formula, and A be a revision action for ϕ w.r.t. Σ, then we have:*
$\Sigma \models \big[\forall s, \phi[s] \rightarrow Bel\big(\phi, do(A, s)\big)\big] \wedge \big[\forall s, \neg\phi[s] \rightarrow Bel\big(\neg\phi, do(A, s)\big)\big]$.

This theorem proves that revision (as well as updating) in our framework is handled adequately. That is, if new information indicates that ϕ holds, then the agent will believe that ϕ holds after performing A. Conversely, if new information shows that ϕ does not hold, then the agent will believe $\neg\phi$ after performing A. This theorem is also consistent with the framework in [20, 21, 24, 11].

Theorem 2 *Let A be a revision action for domain-dependent, uniform formula ϕ w.r.t. Σ, then the following sentence is satisfiable:*
$\Sigma \cup \{Bel(\neg\phi, S_0), Bel(\phi, do(A, S_0)), \neg Bel(FALSE, do(A, S_0))\}$.

This theorem shows that even if the agent believes $\neg\phi$ in S_0, it will believe ϕ after performing A when action A provides that ϕ is true, and still maintains consistent beliefs ($\neg Bel(FALSE, do(A, S_0))$).

5 Example

The advantage of the methods proposed in this paper is that it can tolerate the existence of errors and correct errors, hence keeps a well established track of video analytics. Error correction can be done by either internal inspections or external inferences. In this section, we use a surveillance example to illustrate this advantage.

Example 4 *Now we are going to model a simplified scenario that a passenger boards a bus. We use Fig. 3 to illustrate the situation pedigree. Multiple passengers can be modeled by multiple situation pedigrees.*

Similar to the previous examples, we use SP *to denote Sensing Position and* SG *to denote Sensing Gender. In addition to the internal actions* SP *and* SG*, we also allow external instructions as external actions into the system, e.g.,* $P(M) = 1$ *(resp.* $P(F) = 1$*) in Fig. 3 which means that the external instructions suggesting the passenger is definitely a male (resp. a female).*

Now we give the explanations of the process depicted by Figure 3.

The two initial situations are:

$$S_0^1 = \mathsf{Male} \wedge \mathsf{Stand}, S_0^2 = \mathsf{Female} \wedge \mathsf{Stand}$$

That is, initially the video analytics tell us that the passenger is standing but it does not know accurately whether it is male or female, only providing probability values 0.8 for male and 0.2 for female.

[2] Revision receives (and accepts) information about a static world while updating means that the world itself has been changed [8].

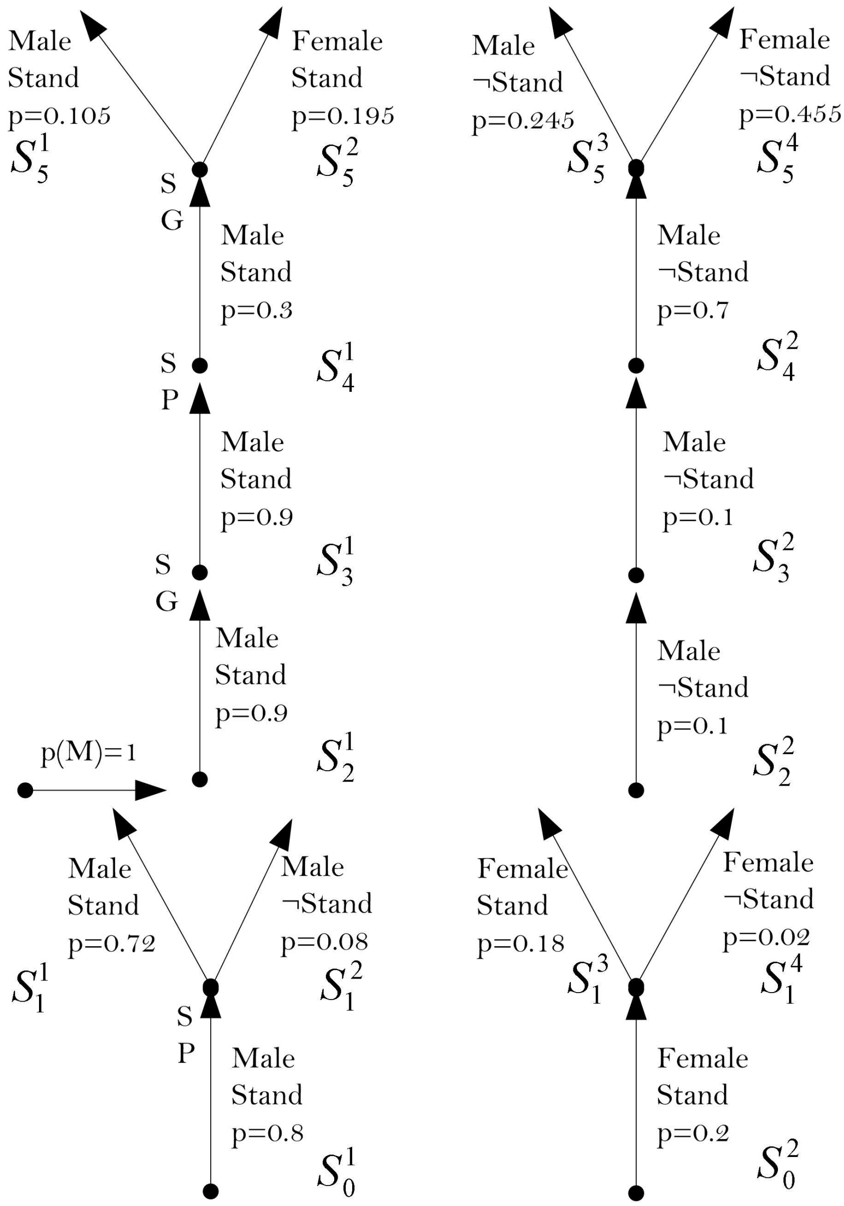

Fig. 3. Surveillance Example

*After a while, the sensor re-examines the position of the passenger (*SP *in the bottom of Fig. 3) and tells us it is now standing with probability 0.9 and seated with probability 0.1. Hence our method gives four possible successive situations* S_1^1, S_1^2, S_1^3*, and* S_1^4 *as shown in Fig. 3.*

Now the monitor in the control room provides a piece of information that this passenger is definitely a male ($p(M) = 1$ *in Fig. 3). Hence we get two successive situations* S_2^1 *and* S_2^2*. We can see that the possibility of the passenger being a female is eliminated.*

*Then the sensor re-examines the gender of the passenger (*SG *in the middle of Fig. 3), and tells us it is a male with probability 0.85 and female with probability 0.15. However, since it does not change the belief that this person is a male and the probability of this person being a male (0.85) is less than the one in the current situation (where the probability of the person being a male is 1), according to our procedure, we do not need to change the probability, hence the two successive situations* S_3^1 *and* S_3^2 *are just the same to their predecessors. Note that here new information from* SG *shows the passenger is male, and after performing* SG*, the proposition* the passenger is male *is believed. It verifies Theorem 1.*

*After that the sensor checks the position of the passenger (*SP *in the top half of Fig. 3). This time the passenger might have taken a seat, hence the video analytics tell us that the passenger is standing with probability 0.3 and seated (*$\neg Stand$*) with probability 0.7. Then two successive situations are induced as* S_4^1 *and* S_4^2 *in Fig. 3.*

Finally, the passenger accidentally removes some of her disguise and the video analytics tell us that it is a female with probability 0.65 *and a male with probability* 0.35*. Then we obtain four successive situations* S_5^1, S_5^2, S_5^3 *and* S_5^4*. Since there is a big change in epistemic properties (*$Male \to Female$*), an alert is triggered and reported to the control room. Also note that here Theorem 1 is verified.*

This example clearly shows how the beliefs are smoothly maintained or changed with uncertain internal (video analytics) and external information, whilst video analytics just tell what is what *at each time point, without continuity.*

6 Conclusion

In this paper, we have proposed a framework to deal with uncertain observations. This framework is based on a revised version of situation calculus. It allows external instructions as well as internal actions. It is able to tackle with uncertain epistemic actions and uncertain ontic actions. Early errors can be corrected in this framework by revision and updating.

In the literature, probabilistic methods (e.g. dynamic bayesian networks, [26], etc.) and other AI techniques (e.g. Bilattice reasoning, [25], etc.) have been applied in computer vision/video surveillance. Comparing to these approaches, our framework considers and easily handles external information. In addition, the ability of belief revision and updating makes it easy to correct past mistakes.

For future work, we are implementing this framework as a part of an on-going intelligent surveillance project (CSIT) to enhance the power of event reasoning, especially for error correction. The full implementation includes a thorough set of ontic actions, epistemic actions for a set of properties of interest. It also can be naturally extended to allow for multiple agents (passengers). In addition, the situations can serve as a foundation for event inference proposed in [12]. Another interesting issue is to study the properties that the framework satisfy.

Acknowledgement

This research work is sponsored by the EPSRC projects EP/G034303/1 and EP/H049606/1 (the CSIT project).

References

1. F. Bacchus, J. Halpern, and H. Levesque. Reasoning about noisy sensors and effectors in the situation calculus. *Artificial Intelligence*, 111(1-2):171–208, 1998.
2. Bsia. Florida school bus surveillance. In *http://www.bsia.co.uk/LY8VIM18989_action;displaystudy_sectorid;LYCQYL79312_caseid;NFLEN064798*.
3. G. De Giacomo and H. Levesque. Progression using regression and sensors. In *Procs. of IJCAI*, pages 160–165, 1999.
4. A. Herzig, J. Lang, and P. Marquis. Action representation and partially observable planning using epistemic logic. In *Procs. of IJCAI*, pages 1067–1072, 2003.
5. G. Lakemeyer and H. Levesque. A semantic characterization of a useful fragment of the situation calculus with knowledge. *Artificial Intelligence*, 175(1):142–164, 2011.
6. H. Levesque. What is planning in the presence of sensing? In *Procs. of AAAI*, pages 1139–1146, 1996.
7. W. Liu, P. Miller, J. Ma, and W. Yan. Challenges of distributed intelligent surveillance system with heterogenous information. In *Procs. of QRASA*, pages 69–74, Pasadena, California, 2009.
8. J. Ma, W. Liu, and S. Benferhat. A belief revision framework for revising epistemic states with partial epistemic states. In *Procs. of AAAI*, pages 333–338, 2010.
9. J. Ma, W. Liu, and A. Hunter. Inducing probability distributions from knowledge bases with (in)dependence relations. In *Procs. of AAAI*, pages 339–344, 2010.
10. J. Ma, W. Liu, and P. Miller. Event modelling and reasoning with uncertain information for distributed sensor networks. In *Procs. of SUM*, pages 236–249. Springer, 2010.
11. J. Ma, W. Liu, and P. Miller. Belief change with noisy sensing in the situation calculus. In *Procs. of UAI*, 2011.
12. J. Ma, W. Liu, P. Miller, and W. Yan. Event composition with imperfect information for bus surveillance. In *Procs. of AVSS*, pages 382–387. IEEE Press, 2009.
13. J. McCarthy. *Situations, Actions and Causal Laws*. Stanford University, Stanford, CA, 1963.
14. J. McCarthy and P. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In *Machine Intelligence*, volume 4, pages 463–502. Edinburgh University Press, 1969.
15. P. Miller, W. Liu, F. Fowler, H. Zhou, J. Shen, J. Ma, J. Zhang, W. Yan, K. McLaughlin, and S. Sezer. Intelligent sensor information system for public transport: To safely go .... In *Procs. of AVSS*, 2010.
16. ECIT Queen's University of Belfast. Airport corridor surveillance. *http://www.csit.qub.ac.uk/Research/ResearchGroups/IntelligentSurveillanceSystems*, 2010.
17. US Defense of the Homeland. Washington rail corridor surveillance. *http://preview.govtsecurity.com/news/Washington-rail-corridor-surveillance/*, 2006.

18. US Department of Transportation. Rita - its research program. *http://www.its.dot.gov/ITS_ROOT2010/its_program/ ITSfederal_program.htm*, 2010.
19. R. Reiter. The frame problem in the situation calculus: A simple solution (sometimes) and a completeness result for goal regression. In *Aritificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*, pages 359–380. Academic Press, 1991.
20. R. Scherl and H. Levesque. The frame problem and knowledge-producing actions. In *Procs. of AAAI*, pages 689–695, 1993.
21. R. Scherl and H. Levesque. Knowledge, action, and the frame problem. *Artificial Intelligence*, 144(1-2):1–39, 2003.
22. Gardiner Security. Glasgow transforms bus security with ip video surveillance. In *http://www.ipusergroup.com/doc-upload/Gardiner-Glasgowbuses.pdf*.
23. S. Shapiro. Belief change with noisy sensing and introspection. In *Procs. of NRAC*, pages 84–89, 2005.
24. S. Shapiro, M. Pagnucco, Y. Lespérance, and H. Levesque. Iterated belief change in the situation calculus. *Artificial Intelligence*, 175(1):165–192, 2011.
25. V. D. Shet, J. Neumann, V. Ramesh, and L. S. Davis. Bilattice-based logical reasoning for human detection. In *IEEE Conf. on Computer Vision and Pattern Recognition, CVPR'07*, pages 1–8, 2007.
26. T. Wang, Q. Diao, Y. Zhang, G. Song, C. Lai, and G. Bradski. A dynamic bayesian network approach to multi-cue based visual tracking. In *Procs. of Pattern Recognition, ICPR*, pages 167–170, 2004.