2009-09-01

Using Dempster Shafer Theory of Evidence for Situation Inference

Susan McKeever

Technological University Dublin, susan.mckeever@tudublin.ie

Juan Ye

University of St Andrews

Lorcan Coyle

University of Limburg

See next page for additional authors

Follow this and additional works at: <https://arrow.tudublin.ie/scschcomcon>

Recommended Citation

McKeever S, Ye, J, & Coyle, L. (2009). Using Dempster Shafer Theory of Evidence for Situation Inference, *Proceedings of EuroSSC 2009*, Sept., London, UK. doi:10.1007/978-3-642-04471-7_12

This Conference Paper is brought to you for free and open access by the School of Computer Sciences at ARROW@TU Dublin. It has been accepted for inclusion in Conference papers by an authorized administrator of ARROW@TU Dublin. For more information, please contact arrow.admin@tudublin.ie, aisling.coyne@tudublin.ie, gerard.connolly@tudublin.ie.

Authors

Susan McKeever, Juan Ye, Lorcan Coyle, and Simon Dobson

Using Dempster-Shafer Theory of Evidence for Situation Inference [★]

Susan McKeever^{★★a}, Juan Ye^a, Lorcan Coyle^b, and Simon Dobson^a

^aSystems Research Group, School of Computer Science and Informatics, University College Dublin, Ireland, ^bLero, University of Limerick, Ireland

Abstract. In the domain of ubiquitous computing, the ability to identify the occurrence of situations is a core function of being 'context-aware'. Given the uncertain nature of sensor information and inference rules, reasoning techniques that cater for uncertainty hold promise for enabling the inference process. In our work, we apply the Dempster Shafer theory of evidence to determine situation occurrence based on uncertain sensor data and inference rules. We also describe a set of evidential operations for sensor mass functions using context quality and evidence accumulation for temporal situation detection. We demonstrate how our approach enables situation inference with uncertain information using a case study based on a published smart home activity data set.

1 Introduction

In the domain of ubiquitous computing, a context-aware system must be able to perceive the state of entities of interest in the environment, such as users. We term these states of interest 'situations', where a situation is a human understandable description of an entity state such as 'user at lunch'. The ability to infer situations, i.e., 'what situation(s) is occurring' is a critical function for a context-aware system, acting as a driver of adaptive behaviour at the application level. Situation inference is reliant on disparate sensor-based information, and ~~this inference process~~ is complicated by the imperfections associated with sensor information, such as problems of noise, breakdown, network delays and user error [3]. Furthermore, observations from multiple sensors can lead to conflicts, such as a user detected in two different locations simultaneously. Therefore, inference mechanisms that treat sensor information as evidence of fact rather than **certain information** are of particular interest in our work.

Bayesian methods, including Bayesian networks [2,10,15] and Hidden Markov Models [1,13] have been used to infer situations in context-aware systems. These

[★] This work is partially supported by Enterprise Ireland under grant number CFTD 2005 INF 217a, "Platform for user-Centred design and evaluation of context-aware services" and by Science Foundation Ireland under grant numbers 07/CE/1147 "Clarity, the centre for sensor web technologies", 03/CE2/I303-1 "Lero, the Irish Software Engineering Research Centre", and 05/RFP/CMS0062 "Towards a semantics of pervasive computing".

^{★★} susan.mckeever@ucd.ie

methods demonstrate that with sufficient training data, higher level context such as activities can be recognised from lower level sensor information. Posterior probabilities of situation occurrence are calculated and imperfections of sensor data and inference rules are absorbed invisibly into these probability calculations. *Dempster-Shafer theory* (DS Theory), which is a generalised form of Bayesian theory, is a tool for representing and combining evidence. It offers an alternative to other Bayesian methods when training data is not easily available. It explicitly quantifies ignorance in the face of uncertain or missing data. It does not rely on training data and it offers a range of operations that can be used for propagating evidence from sensor up to situation level in a scrutable manner. ~~In our work, we~~ use DS theory to incorporate sensor uncertainty into sensor evidence, and to fuse this evidence in order to infer situations. The novelties of our approach are multi-fold: (1) it explicitly caters for quantified sensor uncertainty and uncertain inference rules; (2) domain knowledge can be applied in order to minimise or remove dependance on training data; (3) our approach supports the recognition of temporal situations where causal contexts are distributed over time; (4) the inference process from sensor to situation level is scrutable. We demonstrate our evidence-based situation inference process using sample data from a publicly available home activity data set [13].

This remainder of this paper is organised as follows: Section 2 describes related work by other researchers; Section 3 describes the basic concepts of DS Theory; Section 4 describes situation inference diagrams and the use of DS theory to infer situations. We provide a demonstration of our inference approach in a case study in Section 5. Conclusions and further work are described in Section 6

2 Related work

The application of uncertain reasoning techniques is an active research area within the domain of context-aware computing. In particular, Bayesian methods and fuzzy logic have been used to determine higher level contexts from lower level uncertain sensor information. Van Kasteren et al [13] use Hidden Markov Models to successfully determine a person's activity in the home, where inference of high level states (activities), and temporal activity patterns are learned from training data. Bayesian networks as used by [2,10,15] are used to determine higher level context states (such as activities in a meeting room) from lower level sensor data. These approaches require training data which can be difficult to obtain in particular environments, such as the difficulties of smart home data collection in real life environments, as noted by Tapia et al [12]. Our approach using DS theory has limited or no reliance on training data, relying more heavily on domain knowledge. Fuzzy sets and fuzzy logic as used by [5,7,10] is used to quantify and reason with imprecise context concepts such as describing temperature of a room as 'warm' or 'cold'. We incorporate context fuzziness into our evidence based approach by including fuzzy membership in our sensor mass functions, as described in Section 4.2.

Wu [14] uses DS Theory as a sensor fusion model in context-aware systems. Sensor evidence is supplied via DS mass functions and fused using Dempster's rule of combination. Wu's work does not include the propagation of evidence to support higher level context inference. Uncertainty in sensor mass functions is represented as static weighting of sensors. Wu does introduce a dynamic weighting for evidence sources but this requires the availability of ground truth soon after evidence fusion. Closest to our work is Hong et al's [4] work on activity recognition in smart homes. Similar to our work, they use evidence theory to propagate sensor evidence up to activity to support activity recognition. They use the basic DS theory functions of sensor mass functions, Dempster's combination rule and sensor discounting to process and fuse evidence. They supplement these functions with additional evidential operations to move evidence from sensor level up to activity level. Our work differs from Hong's et al's work by the inclusion of context quality information (such as fuzziness and precision) into sensor mass functions, the processing of evidence for time-distributed situations, and differing evidential operations for evidence fusion.

3 Basic Concepts of Dempster-Shafer Theory

Dempster-Shafer theory is a mathematical theory of evidence [11] which is used to combine separate pieces of information (evidence) to calculate the probability of an event. In a DS Theory reasoning scheme, the set of possible hypotheses are collectively called the *frame of discernment*. This frame Θ represents the set of choices $\{h_1, h_2, \dots, h_n\}$ available to the reasoning scheme, where sources (such as sensors) assign belief or evidence across the frame hypotheses. Let 2^Θ denote the set of all subsets of Θ to which a source of evidence can apply its belief. Then the function $m : 2^\Theta \rightarrow [0, 1]$ is called a mass function that defines how belief is distributed across the frame, if the function satisfies the following conditions:

$$m(\phi) = 0 \text{ and } \sum_{A \subseteq \Theta} m(A) = 1$$

Based on this, belief from an evidence source cannot be assigned to an empty or null hypothesis, and belief from the evidence source across the possible hypotheses (including combinations of hypotheses) must sum to 1. The least informative evidence (ignorance) is the assignment of mass by a source to a hypothesis containing all the elements; i.e., $\{h_1, h_2, \dots, h_n\}$. A crucial part of the process of assessing evidence is the ability to combine evidence from multiple sources. In DS theory, the combination of evidence from two different independent sources is accomplished by Dempster's combination rule:

$$m_{12}(A) = \frac{\sum_{\forall X, Y: X \cap Y = A} m_1(X) \cdot m_2(Y)}{1 - \sum_{\forall X, Y: X \cap Y = \phi} m_1(X) \cdot m_2(Y)} \quad (1)$$

where $m_{12}(A)$ is the combined belief for a given hypothesis A. The numerator in equation 1 represents evidence for hypotheses whose intersection is the exact

hypothesis of interest, A . The denominator, $1 - K$ is a normalisation factor, where K is a *conflict* factor representing all combined evidence that does not match the hypothesis of interest, A . The value of conflict, K , when combining evidence is indicative of the level of disagreement amongst the sources of their belief in hypothesis A .

In Section 4.2, we explain how we apply the DS Theory concepts of *mass functions*, *frames of discernment* and *evidence combination* to situation inference.

4 Situation inference

In this section, we explain how we use Situation Inference Diagrams to capture the inference routes from sensor information to situations. We then explain how basic DS theory (mass functions, evidence combination) and additional evidential operations (context quality in mass functions, evidence transfer, evidence accumulation) are applied to situation inference.

4.1 Situation Inference Diagrams

In order to infer situations using DS Theory, we need to define how evidence is propagated across layers of context, using a multi-layered hierarchy consisting of sensors, abstracted context and situations [8]. We illustrate this hierarchy using a Directed Acyclic Graph (DAG) as shown in Figure 1. Sensors are the root nodes at the base of the diagram. Sensor readings are abstracted or mapped to more human understandable context values. For example, a sensor reading of a user's location may be generated as a set of coordinates '12.4, 10, 5.6, 14:24:08, ID24' and translated using a building layout to a context value of 'Peter hasLocation MeetingRoom5 at 14:24'. Moving up the hierarchy, each context value will be mapped to one or more situations, indicating that the occurrence of a particular context value 'is evidence' of the situation occurring; i.e., an inference rule. Depending upon the complexity and range of situations in the system, higher level situations may also be inferred from lower level situations. Each solid directed arrow in the graph is interpreted as 'is evidence of'. The full notation for the DAG is shown in Figure 2.

Quantified uncertainty is incorporated into the situation inference DAG at sensor and inference rule level. At sensor level, if sensor reliability is quantifiable, it is included as a sensor discount, as described further in section 4.2. Uncertainty can also be quantified against inference rules. For example, in the home data set that we have examined, a user 'sometimes' uses the microwave when preparing breakfast and this is quantified as 40% of the time, by examining sample occurrences of the 'prepare breakfast' situation in the data. Therefore, a certainty of 0.4 is applied along the edge or inference rule from the context value 'microwave used' to the situation of 'prepare breakfast'.

Situations may be inferred from evidence that does not occur at exactly the same time and we term these 'time distributed situations'. The situation

Fig. 1. Situation Inference Diagram

Fig. 2. Situation Inference Diagram Notation

of a user 'preparing breakfast' may be detected by context values such as the 'plate cupboard used', 'grocery cupboard used', and 'fridge used' over a period of time. We represent evidence that accumulates over time by a time period enclosed in '< >' brackets within the time-distributed situation node on the DAG. This number indicates the typical duration of the situation. Where the actual *sequence* of evidence occurrence is also relevant, the duration is enclosed by '> >' brackets.

Some situations may be detectable by the occurrence of one of a set of situations. For example, a 'busy in kitchen' situation may be declared when either situation 'prepare breakfast' or 'prepare dinner' is occurring. This is indicated on the DAG as 'is a type of'.

4.2 Applying DS Theory to Situation Inference

Once the situation inference hierarchy is determined from sensor to situation level, we can then apply evidential operations from DS theory as described in Section 3, and extensions to DS Theory, to produce and propagate evidence for situation inference. In order to infer situation occurrence, we assess sensor readings or evidence at specific points in time t , where each point is separated by a time gap, Δt . At each point in time, the situations with the greatest belief (evidential support) are believed to be occurring. To achieve this, a number of evidential operations need to occur from sensor level upwards. (1) At time t , each sensor system defines its belief across the context values for the sensor via a *mass function*. (2) Quality information associated with each sensor is used to modify the belief distribution from its mass function. (3) The belief associated with each context value is mapped upwards along the inference paths to its associated situations using *compatibility relations* and *evidence propagation* between frames of discernment. This mapping **processes** continues up the hierarchy of situations. (4) For each situation, belief from the relevant sources of evidence is fused in order to determine total belief in the situation at time t , using an appropriate *evidence combination rule*. (5) For situations that have a time duration, belief from sources is fused with belief from pre-existing evidence that was captured within the time duration of the situation.

When these steps have been executed, at time t , a separate level of belief is available for each situation. A decision step selects those situations of highest belief, taking into account which situation can occur at the same time.

Sensor Mass functions Mass functions, as explained in Section 3, define how belief from an evidence source (such as a sensor) is spread across a set of choices in a frame of discernment. In our approach, a mass function is required for each sensor, in order to distribute belief for the sensor across the range of possible context values for that sensor; i.e., the frame of discernment for that sensor. In our activity sensor example, context value 'active' may be defined as any keyboard or mouse activity on the computer within the last 10 seconds. If the sensor at time t has detected a reading 5 seconds ago, the mass function will assign belief

across the frame of discernment $\{active, inactive, \theta\}$ as $\{1, 0, 0\}$ respectively. The mass assigned to θ represents ignorance, where belief is uncertain and cannot be sub-divided amongst other elements in the frame.

Using Context Quality in Sensor Mass Functions In the real world, absolutely reliable sources are rarely found. ‘Discount factors’ are an idea that dates back to Shafer’s evidence theory work [11] and expanded by Lowrance et al [6]. They can be used to account for the uncertainty due to unreliable evidence sources. However, a variety of additional sensor and context quality metrics such as precision [16] and fuzziness [10] are defined to quantify the imperfections of sensor information. Precision indicates the range within which a sensor reading is correct, for a given accuracy. Fuzziness is used to quantify imprecise context. If such quality information is available and quantifiable, we propose that it can be incorporated into sensor mass functions in order to provide a more realistic distribution of belief. We illustrate this through a set of examples:

(1) *Using precision in mass functions:* A location system, Ubisense¹, generates coordinate readings for tags worn for users. The coordinate readings are then mapped to meaningful locations (context values) within the building based on a building map. The measured precisions of Ubisense readings are 3.30 and 2.22 meters along the x- and y- axes respectively, as established using training data in experimental work [16]. When these precisions are used during context abstraction, they form an area around the coordinate reading, and the user/tag may be located anywhere within the area. If this area intersects more than one meaningful space, the proportions of the square in each space can be calculated. Therefore, a reading which is translated to two context values ‘desk1, 0.3’ and ‘desk2, 0.7’, means that the sensor is 30% confident that the user is located at desk1 and 70% confident at desk2. The mass function for this sensor based on this coordinate reading will generate belief of $\{0.3, 0.7, \dots, 0, 0\}$ for a frame of discernment $\{desk1, desk2, \dots, desk_n, \theta\}$.

(2) *Using context fuzziness in mass functions:* Context values can have imprecise meanings such as ‘active’ for our activity sensor. The level of imprecision or fuzziness is represented by a membership value, where the fuzzy membership function applies a numerical value from 0 to 1 to each element of the fuzzy set [17]. For our activity sensor, we observe that any keyboard or mouse activity within the last 10 seconds is definitely active. After 10 seconds, the level or membership of ‘active’ reduces with time, falling to 0 after 60 seconds. The mass function for the activity sensor will include the fuzzy membership values in its calculation of belief assignment to the context values, generating belief distributions such as $\{0.2, 0.8, 0\}$ for the frame $\{active, inactive, \theta\}$, when a reading is more inactive than active.

(3) *Sensor discounting:* When a quantified measure of sensor reliability is available, it is incorporated into the mass function via a sensor discounting function [11,6]. When a source’s evidence is discounted, the remaining evidence is applied to the combination of all options in the frame (i.e., ignorance, θ).

¹ Ubisense is a networked location system: www.ubisense.net

For our Ubisense system, we measured an accuracy of 70% for the precisions described in example (1), meaning that 30% of readings are believed to be incorrect. Using Shafer discounting function, the belief is discounted by 0.7. The remaining 0.3 is attributed to ignorance. Using the Ubisense mass function example from (1) of belief $\{0.3, 0.7, \dots, 0, 0\}$ for a frame $\{desk1, desk2, \dots, desk_n, \theta\}$, the application of the discount factor will alter the belief distribution to $\{0.21, 0.49, \dots, 0, 0.3\}$ where discounting of evidence has resulted in an quantified uncertainty of 0.3.

Reliability of a sensor may be a straightforward measure of physical sensor accuracy as supplied by the sensor manufacturer. However, it may also incorporate additional sources of error such as errors in using a sensor. For example, if a user is known to forget to wear their Ubisense locator tag 10% of the time, this will change the discount factor on the Ubisense sensor system from 0.7 to 0.63.

Evidence Transfer Evidence propagation through layers of the hierarchy from context value upwards is achieved using compatibility relations and evidential mapping. Compatibility relations [6] define the mappings between the compatible beliefs between two frames of discernment; i.e., the elements from the two frames that can be true simultaneously. We can use compatibility relations to define paths for transferring belief from one layer of the situation hierarchy to the next. For example, a fridge sensor can generate belief across two context values $\{FridgeUsed, \neg FridgeUsed, \Theta\}$. The use of the fridge is indicative of the 'get drink' situation, which has a frame of discernment $\{GetDrink, \neg GetDrink, \Theta\}$. 'Fridge used' is compatible with **with** 'get drink' (i.e. they are both true simultaneously) and so on for the remaining elements in both frames. Having defined the evidence paths using compatibility relations, we use evidence *propagation* [4] to propagate evidence from one frame to another. We apply propagation along the paths defined by the compatibility relations, enabling mass of compatible elements to be transferred; e.g., mass of 'fridge used' can be propagated to 'get drink'.

Evidence Combination When each context value has propagated its evidence to situation level, the evidence is combined to produce a final distribution of belief over the choices in the frame of discernment. The basic formalism for evidence combination from two sources is provided in Dempster's rule of combination as described in equation 1 of Section 3. Variations on this combination rule have been introduced in the literature to deal with alternative combination scenarios, such as the use of evidence averaging for unreliable sources when source discounting is used [11] and Murphy's averaged combination rule [9]. Murphy observed that a single piece of evidence can force certainty or overrule a majority when Dempster's rule of combination is used. For scenarios where this may occur (e.g. binary sensors where a single sensor fails to fire), this will distort the evidence, allowing a single sensor to negate evidence from other sources. Murphy's approach is to average the evidence *prior* to combining, thus ruling out the dominance of a single sensor. The averaged evidence is then combined

using Dempster’s combination rule, applied $n - 1$ times where n is the number evidence sources. We propose to use both averaging as proposed by Shafer and Murphy’s alternative combination rule in our work, as demonstrated in section 5, and compare the results in our worked example.

Evidence accumulation over time In order to accumulate evidence for a situation with a specified duration d , we extend the lifetime of evidence to endure over the duration of the situation, combining later evidence with earlier evidence as if it had occurred at the same time, t . At the first occurrence of evidence of a situation (such as ‘grocery cupboard’ detected for the ‘prepare breakfast’ situation), evidence is captured for this time, t . Further evidence for the situation that occurs during the situation (i.e., between time t and $t + d$), is fused with earlier evidence to provide an overall belief in the situation occurrence.

5 Case Study

In this section, we provide a demonstration of our evidence-based situation inferencing approach using a sample of data from Van Kasteren’s home activity data set [13]. We explain each of the steps from processing of sensor evidence up to the fusion of evidence into situation beliefs.

5.1 Data set description

The data set was recorded over 28 days in a house where a 26 year old man lives. 14 sensors were places throughout the house. Each sensor generates binary output only, outputting a value of 1 when fired. The data set is annotated with 7 situations or activities, such as ‘go to bed’, ‘take shower’, and ‘prepare dinner’. None of these activities occur at the same time according to the annotation. In our example, we focus on the three situations that occur in the kitchen: ‘prepare breakfast’, ‘prepare dinner’, ‘get drink’. Our situation inference DAG for these situations (Figure 3) contains all kitchen based sensors as root nodes. Because of the binary nature of the sensors, context values for these sensors are very simple; e.g., the grocery cupboard sensor firing indicates ‘grocery cupboard used’ context. No indication of sensor performance is provided in the data set, so sensor discounting or quality in mass sensor functions cannot be applied. Domain knowledge in this environment for mapping evidence to situations could, in theory, be available from users (‘what do you typically do to prepare breakfast?’) or from examining small amounts of training data.

5.2 Experimental approach

We used a combination of inherent domain knowledge combined with examination of 5 occurrences of each of the three situations to determine the inference paths for our DAG. This highlighted some uncertainty in the inference rules,

Fig. 3. Situation Inference DAG for kitchen situations: get drink, prepare breakfast, prepare dinner

such as the occasional use of the microwave in breakfast preparation (20% of the time), or getting cups from the cups cupboard when preparing dinner (40% of the time). Each situation is a time-distributed situation so we estimate situation durations from a combination of data observation and domain knowledge. For instance, the 'prepare dinner' activity lasted an average of 38 minutes for the 5 occurrences we observed, with a standard deviation of 24 minutes so we applied a duration of 62 minutes to capture 'prepare dinner' evidence. We could have used alternative duration calculations such as mean duration, minimum duration or a duration range, depending upon the nature of the situation. From domain knowledge, we included time as evidence to capture the fact that breakfast occurs in the morning and dinner occurs in the evening. 'Get drink' can occur at any time of the day.

We selected 5 consecutive time slices from a single day when the 'prepare breakfast' activity was annotated, using time slices of 1 minute (as also used by [13]). Evidence for the five time slices is shown in Table 1, showing the kitchen sensor events that occurred over times 9:49 to 9:53 for one day in the data set. Sensor events outside the kitchen such as 'hall bedroom' indicating that the hall bedroom sensor is still firing (open) were discarded as our domain knowledge suggests that only kitchen based sensors have a direct evidential bearing on kitchen situation occurrence. The evidential calculations for these time slices results in total belief for each of the three situations at times 9:49 to 9:53 as shown in Table 2.

The process for inferring situation occurrence using sensor evidence is performed by using the following steps for each time slice:

		Cumulative evidence over situation durations (mins)		
Time	Sensor events	Drink (2)	Prepare Breakfast (15)	Prepare Dinner (62)
9.49	fridge, plates, microwave	fridge	fridge, plates, microwave (0.2)	fridge (0.8), plates
9.50	groceries	fridge	fridge, plate, microwave (0.2), groceries	fridge (0.8), plates, grocery (0.8)
9.51	none	none	fridge, plate, microwave (0.2), groceries	fridge (0.8), plates, grocery (0.8)
9.52	none	none	fridge, plates, microwave (0.2), groceries	fridge (0.8), plates, grocery (0.8)
9.53	microwave, groceries	none	fridge, plates, microwave (0.2), groceries	fridge (0.8), plates, grocery (0.8)

Table 1. Evidence and cumulative evidence by time slice for each situation

1. Use sensor mass functions to obtain context value beliefs.
2. Propagate the belief from the context values up to the relevant situations, using compatibility relations and evidence propagation as explained in Section 4.2.
3. Obtain total belief for each situation by combining evidence for the situation. We will use both basic evidence averaging and Murphy’s alternative combination rule, as described in Section 4.2.
4. Select the situation with the highest belief (given that no situations in the data set are co-occurring).

For illustration, we explain these steps for the first time slice (9:49):

(1) Use sensor mass functions for context values

Fridge, plates cupboard and microwave sensors are firing. No sensor discounting is used. The three sensor mass functions assign mass to context values as follows:

$$\begin{aligned} &\{FridgeUsed = 1, \neg FridgeUsed = 0\} \\ &\{PlateUsed = 1, \neg PlateUsed = 0\} \\ &\{MicrowaveUsed = 1, \neg MicrowaveUsed = 0\} \end{aligned}$$

In addition, the kitchen sensors that did *not* fire are evidential, generating masses as follows:

$$\begin{aligned} &\{CupUsed = 0, \neg CupUsed = 1\} \\ &\{GroceriesUsed = 0, \neg GroceriesUsed = 1\} \\ &\{FreezerUsed = 0, \neg FreezerUsed = 1\} \\ &\{PansUsed = 0, \neg PansUsed = 1\} \end{aligned}$$

Finally, we are using time as evidence. At 9:49am, the time sensor mass gives: $\{Morning = 1, Evening = 0\}$

(2) Propagate the belief from the context values up to the relevant situations

For each situation, propagate context value belief to situations of which that context is evidence, as denoted on the situation inference DAG.

Get drink

$\{FridgeUsed = 1, \neg FridgeUsed = 0\} \rightarrow \{GetDrink = 1, \neg GetDrink = 0\}$
 $\{CupUsed = 0, \neg CupUsed = 1\} \rightarrow \{GetDrink = 0, \neg GetDrink = 0.8, \theta = 0.2\}$

Prepare Breakfast:

$\{FridgeUsed = 1, FridgeNotUsed = 0\} \rightarrow \{Breakfast = 1, \neg Breakfast = 0\}$
 $\{MicrowaveUsed = 1, \neg MicroUsed = 0\} \rightarrow \{Breakfast = 0.2, \neg Breakfast = 0, \theta = 0.8\}$
 $\{PlateUsed = 1, \neg PlatesUsed = 0\} \rightarrow \{Breakfast = 1, \neg Breakfast = 0\}$
 $\{GroceriesUsed = 0, \neg GroceriesUsed = 1\} \rightarrow \{Breakfast = 0, \neg Breakfast = 1\}$
 $\{Morning = 1, Evening = 0\} \rightarrow \{Breakfast = 1, \neg Breakfast = 0\}$

Prepare Dinner

$\{FridgeUsed = 1, \neg FridgeUsed = 0\} \rightarrow \{Dinner = 0.8, \neg Dinner = 0, \theta = 0.2\}$
 $\{PlateUsed = 1, \neg PlateUsed = 0\} \rightarrow \{Dinner = 1, \neg Dinner = 0\}$
 $\{GroceriesUsed = 0, \neg GroceriesUsed = 1\} \rightarrow \{Dinner = 0, \neg Dinner = 0.8, \theta = 0.2\}$
 $\{PansUsed = 0, \neg PansUsed = 1\} \rightarrow \{Dinner = 0, \neg Dinner = 0.8, \theta = 0.2\}$
 $\{CupsUsed = 0, \neg CupsUsed = 1\} \rightarrow \{Dinner = 0, \neg Dinner = 0.4, \theta = 0.6\}$
 $\{FreezerUsed = 0, \neg FreezerUsed = 1\} \rightarrow \{Dinner = 0, \neg Dinner = 1\}$
 $\{Morning = 1, Evening = 0\} \rightarrow \{Dinner = 0, \neg Dinner = 1\}$

Because this is the first time slice, we do not have any cumulative evidence. For the next time slice at 9:50, the sensor events from 9:49 will be assessed as 'still happening' where they contribute to the time-distributed situations, as set out in Table 1.

Timeslice	Averaging			Combination Rule					
	Drink	B'fast	Dinner	Drink		B'fast		Dinner	
				Belief	Conflict	Belief	Conflict	Belief	Conflict
9.49	0.5	0.64	0.26	0.63	43%	0.98	22%	0.02	27%
9.50	0.5	0.84	0.37	0.63	43%	1	0	0.2	40%
9.51	0	0.84	0.37	0	n/a	1	0	0.2	40%
9.52	0	0.84	0.37	0	n/a	1	0	0.2	40%
9.53	0	0.84	0.37	0	n/a	1	0	0.2	40%

Table 2. Total belief in situations by timeslice using evidence averaging and Murphy's combination rule for 'get drink' (Drink), 'prepare breakfast' (B'fast) and 'prepare dinner' (Dinner) situations

(3) Obtain total belief in each situation via evidence combination

We calculate total belief in each situation using two different methods: Simple evidence averaging and Murphy's version of the Dempster's combination rule. In the averaging approach, this involves averaging the belief for each separate situation. At 9:49, this gives belief of 0.5, 0.64, 0.26 for situations 'get drink', 'prepare breakfast', 'prepare dinner' respectively as shown in Table 2. For the

combination rule, we combine the averaged situation belief $n-1$ times using the combination rule in equation 1, where n is the total number of evidence sources for the situation. Total belief is shown for the three situations for our time slices in Table 2. We then repeat the steps for the remaining time slices, using new and cumulative evidence from Table 1.

Looking at Table 2, the situation 'prepare breakfast' is deemed to be occurring at 9:49. As evidence increases for 'prepare breakfast' and 'prepare dinner' (at 9:50) due to the grocery cupboard sensor firing, the belief in these two situations increases. After 9:50, no further sensors fire until 9:53, but the existing evidence endures for time slices 9:51 and 9:52. At 9:53 the microwave and grocery cupboard sensors fire again. However, they do not change the belief because they fired in previous time slices within the system duration so their evidence is still active. For the 'get drink' situation, the situation duration is 2 minutes, so belief drops to zero at 9:51, 2 minutes after the fridge sensor contributed evidence to 'get drink' occurring. For evidence combination, the averaging of evidence provides the same highest belief answer as the combination rule, but the evidence does not converge to the same extent as our combination rule. This is because the combination rule normalises out conflicting evidence and because uncertainty is re-distributed to the two other elements in the frame of discernment. In our example, the three assessed situations cannot be co-occurring (according to the annotations in the published data set), therefore the selection of 'which situation is occurring' is simply based on the highest belief at time t . In a more complex environment where multiple situations may be co-occurring, we will need to develop heuristics for belief thresholds in order to decide which situations are occurring. We anticipate that use of the conflict metric may be useful in developing decision making heuristics for determining situation occurrence.

6 Conclusion and Future Work

In this paper, we presented an approach to inferring situation occurrence using the Dempster Shafer theory of evidence. Our approach incorporates context quality information into sensor evidence, propagates sensor evidence up to situation level and obtains belief for situations via evidence fusion. We also provided a mechanism to accumulate evidence for time-distributed situations. We demonstrated our approach in a case study, using a sample of time-distributed evidence from a publicly available smart home data set. Our approach enables situation inference with uncertain information, with limited or no need for training data.

Our inferred situation for our proof of concept demonstration matched the annotated situation ('prepare breakfast') for the sample data used. The next stage of our work is to establish inference results for the full smart home data set using our DS theory based approach. As part of our evaluation, we will compare our results against published results for the data set that use alternative uncertain reasoning approaches. We will also test our approach against an intelligent office activity data set that we are collecting in-house. The data set tracks office-based users using a variety of sensors. We expect that it will provide us

with richer sensor quality information which we can use to test the impact on situation inference of using context quality in sensor mass functions.

References

1. B. Clarkson, A. Pentland, and K. Mase. Recognizing user context via wearable sensors. *Wearable Computers, IEEE International Symposium*, page 69, 2000.
2. T. Gu, H. K. Pung, and D. Q. Zhang. A bayesian approach for dealing with uncertain contexts. In *Advances in Pervasive Computing*, pages 136–144, 2004.
3. K. Hemricksen, J. Indulska, and A. Rakotonirainy. Modeling context information in pervasive computing systems. In *Proc' of 1st International Conference in Pervasive Computing*, pages 167–180. Springer, 2002.
4. X. Hong, C. Nugent, M. Mulvenna, S. McClean, B. Scotney, and S. Devlin. Evidential fusion of sensor data for activity recognition in smart homes. *Pervasive and Mobile Computing*, In Press, Corrected Proof, 2008.
5. P. Korpipaa, J. Mantyjarvi, J. Kela, H. Keranen, and E. Malm. Managing context information in mobile devices. *Pervasive Computing, IEEE*, 2(3):42–51, 2003.
6. J. D. Lowrance, T. D. Garvey, and T. M. Strat. A framework for evidential-reasoning systems. *Readings in uncertain reasoning*, pages 611–618, 1990.
7. J. Mantyjarvi and T. Seppanen. Adapting applications in handheld devices using fuzzy context information. *Interacting with Computers*, 15(4):521–538, 2003.
8. S. McKeever, J. Ye, L. Coyle, and S. Dobson. A multilayered uncertainty model for context aware systems. In *Adjunct Proc' of the international conference on Pervasive Computing: Late Breaking Result*, pages 1–4, May 2008.
9. C. K. Murphy. Combining belief functions when evidence conflicts. *Decis. Support Syst.*, 29(1):1–9, 2000.
10. A. Ranganathan, J. Al-Muhtadi, and R. Campbell. Reasoning about uncertain contexts in pervasive computing environments. *Pervasive Computing, IEEE*, 3(2):62–70, 2004.
11. G. Shafer. *A mathematical theory of evidence*. Princeton Unverisity Press, 1976.
12. E. Tapia, S. Intille, and K. Larson. Activity recognition in the home using simple and ubiquitous sensors. In *Pervasive Computing*, volume 3001 of *LNCS*, pages 158–175, 2004.
13. E. VanKasteren, Noulas and Krose. *Accurate activity recognition in a home setting*. In *Proc' of 10th Internationall Conference on Ubiquitous Computing*, South Korea, Sept 2008.
14. H. Wu. *Sensor fusion using dempster shafer theory*. PhD dissertation, Carnegie Mellon University, 2003.
15. J. Ye, L. Coyle, S. Dobson, and P. Nixon. Representing and manipulating situation hierarchies using situation lattices. *Revue d'Intelligence Artificielle*, 22(5), 2008.
16. J. Ye, S. McKeever, L. Coyle, S. Neely, and S. Dobson. Resolving uncertainty in context integration and abstraction: context integration and abstraction. In *ICPS '08: Proceedings of the 5th international conference on Pervasive services*, pages 131–140, New York, 2008. ACM.
17. L. Zadeh. Fuzzy sets. *Information Control*, 8:338–353, 1965.