Information Fusion 29 (2016) 40–56

Contents lists available at ScienceDirect

Information Fusion

journal homepage: www.elsevier.com/locate/inffus

Multi-Hypotheses Tracking using the Dempster–Shafer Theory, application to ambiguous road context

Dominique Gruyer[a], Sébastien Demmel[b,*], Valentin Magnier[a], Rachid Belaroussi[a]

[a] COSYS-LIVIC (IFSTTAR), 77 rue des Chantiers, Versaillles 78000, France
[b] Queensland University of Technology (QUT), Centre for Accident Research and Road Safety—Queensland, 130 Victoria Park Road, Kelvin Grove QLD 4059, Australia

ARTICLE INFO

Article history:
Received 27 October 2014
Revised 31 May 2015
Accepted 6 October 2015
Available online 22 October 2015

Keywords:
Tracking
Association
Ambiguity
Dempster–Shafer Theory

ABSTRACT

This paper presents a Multi-Hypotheses Tracking (MHT) approach that allows solving ambiguities that arise with previous methods of associating targets and tracks within a highly volatile vehicular environment. The previous approach based on the Dempster–Shafer Theory assumes that associations between tracks and targets are unique; this was shown to allow the formation of ghost tracks when there was too much ambiguity or conflict for the system to take a meaningful decision. The MHT algorithm described in this paper removes this uniqueness condition, allowing the system to include ambiguity and even to prevent making any decision if available data are poor. We provide a general introduction to the Dempster–Shafer Theory and present the previously used approach. Then, we explain our MHT mechanism and provide evidence of its increased performance in reducing the amount of ghost tracks and false positive processed by the tracking system.

1. Introduction

Many existing or in development Intelligent Transportation Systems (ITS) applications perform tasks that provide some degree of perception of the road's environment, and use this perception to achieve their goal, for example the detection of impending collisions between vehicles. To understand their environment, they rely on sensors that gather information about their surroundings at a given frequency. Because of hardware/software limitations, or because one uses multiple different sensors, the information about the vehicle's environment can be highly asynchronous. For the ITS applications to perform their task properly, a mechanism is needed to reconstruct the evolution of the scene over time, taking into account those gaps and also imperfections in the known data arising from, for example, sensors defects; this mechanism is known as tracking.

Tracking an object such as a vehicle on the road is a three step process, with the stages: (1) synchronisation, (2) association and (3) fusion. The synchronisation task is to predict the evolution of the known objects to the current timestamp k, knowing information on their behaviour at time $k-1$. The classical way to predict this evolution is to use a Kalman Filter estimator [1,2]. Predicted objects are called *tracks*, while observations from the sensor(s) are called *targets*. The association step consist in finding which tracks correspond to which targets before they can be fused together in a last step to obtain a more accurate description of the scene at time k. In this paper, we will focus on the second association step, which is the most complex of the three.

Many different association methods exist (a summary of their advantages and disadvantages are given in Table 1), some fairly straightforward; for example the nearest neighbour method that simply considers the distance between tracks and targets and associate the objects that fall closest to each others. The distance can be computed using the Euclidean distance, the Mahalanobis distance, etc. Unfortunately, this method is inappropriate for complex problems [3]. Some more complex methods are based on probabilistic approaches: probabilistic data association (PDA), joint probabilistic data association (JPDA and JPDAM), nearest neighbour JPDA, etc. [4,5]. Overall, these methods compute association probabilities between the tracks and targets, but cannot manage the appearance or disappearance of tracks.

The Multi-hypotheses Filter (MHF) [6] can manage new tracks. It looks for the probabilities associated with three specific hypotheses for targets, whether (1) they associate with known tracks, (2) they associate with new tracks, and (3) they are false positives. Unfortunately, the MHF is relatively computationally heavy.

An alternative approach growing in popularity is uses the Dempster–Shafer Theory [7,8], also known as the Belief Theory. It

* Corresponding author. Tel.: +61 7 3138 7783.
E-mail addresses: dominique.gruyer@ifsttar.fr (D. Gruyer), sebastien.demmel@qut.edu.au (S. Demmel), valentin.magnier@ifsttar.fr (V. Magnier), rachid.belaroussi@ifsttar.fr (R. Belaroussi).

http://dx.doi.org/10.1016/j.inffus.2015.10.001
1566-2535/

Table 1
Association methods and their characteristics.

Method	Manage appear/disappear	Multi-targets	Process load
Nearest neighbour	No	No	Low
PDA	No	No	Low
JPDA	No	Yes	Average
NNJPDA	No	Yes	Low
MHF	Yes	Yes	Heavy
Dempster–Shafer	Yes	Yes	Average

is advantageous in an automotive context because it can handle imprecision and incertitude in a more suitable way than probabilistic theories, as well as manage ignorance and conflicts. A framework for association using Dempster–Shafer was proposed in [3,9], but this approach has some flaws. Notably, it will not use all the available information because it is forced to make a choice when associating tracks and target, loosing all the potentially useful information contained within ambiguities and conflicts. This leads to the formation of false positives, dubbed “ghost tracks”. In this paper, we propose a solution to this latter problem using a Multi-Hypotheses Tracking (MHT) algorithm.

In the remainder of this paper we will at first introduce the Dempster–Shafer Theory in general terms, notably the notion of belief functions (Section 2). Then, after having led out the formulation of our association problem (Section 3), we will present the principles for generating basic belief assignments (BBA) from sensors data (Section 4) and their combination so that the system can take a decision on which tracks to associate with each targets (Section 5). Then, we will present our MHT algorithm (Section 6) and provide a demonstration of its performance compared to the PDAF and the classical Dempster–Shafer tracking (Section 7). The chosen examples for this comparison focus on pedestrians walking in front of a laserscanner; this situation is particularly complex due to the location of the scanning plane at legs level, creating multiple ambiguities. Another example more linked to the automotive context is also proposed through the study of a car overtaking manoeuvre observed by a laser scanner sensor.

2. Belief functions

Belief functions were introduced by Dempster [7] and further refined by Shafer [8], taking the name of the Dempster–Shafer Theory. It is also sometimes referred to as the Belief Theory. A further extension was undertaken by Smets [10,11] to create the Transferable Belief Model (TBM). Let us at first define Ω, the *universal set* that represents the various possible states for the system under consideration, i.e. the “*frame of discernment*”. The possible states are the simple (singletons) acceptable propositions H_i, so that:

$$\Omega = \{H_1, H_2, \dots, H_n\} \tag{1}$$

$$\{H_i\} \cap \{H_j\} = \varnothing, \quad \forall i \neq j \tag{2}$$

From this universal set, we can define the *power-set* 2^Ω that is the set of all subsets of Ω, including the empty set $\varnothing$. The power-set includes all the combinations based on the hypotheses from the universal set. Proposition A can be a singleton hypothesis or a complex hypothesis, which includes more than one hypotheses. $\varnothing_\Omega$ represents impossible propositions (conflicts) and Ω the total ignorance, since it includes all existing hypotheses.

$$2^\Omega = \{A / A \subseteq \Omega\} = \{\varnothing_\Omega, H_1, H_2, \dots, H_n, H_1 \cup H_2, \dots, \Omega\} \tag{3}$$

In [3,12,13], the authors used an extended universal set Θ, the *extended open world*, by creating a new hypothesis labelled $*$ which represents any new hypothesis that is not initially modelled in Ω. This approach allows discriminating between conflict and new hypotheses, which is not possible in the general approach.

$$\Theta = \{H_1, H_2, \dots, H_n, *\} \tag{4}$$

$$2^\Theta = \{\varnothing_\Theta, H_1, H_2, \dots, H_n, H_1 \cup H_2, \dots, *, \Omega\} \tag{5}$$

The Dempster–Shafer Theory allows to evaluate the likelihood of a proposition A through its belief mass $m^\Omega(A)$, the mass of elementary probability on the said proposition A, a function defined as:

$$\begin{aligned} m^\Omega : 2^\Omega &\rightarrow [0, 1] \\ A &\rightarrow m^\Omega(A) \end{aligned} \tag{6}$$

The set of belief masses constitutes the *basic belief assignment* (BBA), which verifies:

$$m^\Omega(\varnothing_\Omega) = 0 \tag{7}$$

$$\sum_{A \in 2^\Omega} m^\Omega(A) = 1 \tag{8}$$

$m^\Omega(A)$ is the degree of belief assigned to proposition A, more precisely it expresses the proportion of all relevant and available evidence that supports the claim that the actual state belongs to A but to no particular subset of A. $m^\Omega(\Omega)$ represents the mass of ignorance. If A is a complex hypothesis (i.e. not a singleton one), it means that given the current state of knowledge on the system, no mass could be assigned to a more specific proposition. This represents a partial ignorance of the system's state. Total ignorance is represented by the following masses set: $m^\Omega(\Omega) = 1$ and $m^\Omega(A) = 0, \quad \forall A \neq \Omega$. The mass $m^\Omega(\varnothing_\Omega)$ is the mass of conflict (the BBA is labelled as *normal*if $m^\Omega(\varnothing_\Omega) = 0$), and propositions A that have a non-null mass ($m^\Omega(A) > 0$) are called *focal elements*. When focal elements are composed only of singleton hypotheses, the masses are linkable to probabilities, creating a set of Bayesian masses. If the extended open world Θ is used, then, $m^\Theta(\varnothing_\Theta) + m^\Theta(*) = m^\Omega(\varnothing_\Omega)$. See Section 4 for the details on how masses are assigned within our approach.

After the sets of mass assignments have been obtained, the problem becomes how to combine two independent sets of mass assignments, in other words, how to combine evidence from difference sources (such as different sensors)? There are a number of different rules to do so, the principal one being the Dempster–Shafer (DSR) rule. Let us consider S information sources $\forall A \subset \Omega$, which resulting BBA is $m^\Omega_{1,\dots,S}$, called the *joint mass*. Their respective focal elements are $B_1, \dots, B_S$. The DSR is computed so that the final mass of conflict is null ($m^\Omega_{1,\dots,S}(\varnothing_\Omega) = 0$). The joint mass is given by:

$$m^\Omega_{1,\dots,S}(A) = \frac{1}{K} \sum_{B_1 \cap \dots \cap B_S = A} m^\Omega_1(B_1) \dots m^\Omega_S(B_S) \tag{9}$$

where K is a normalisation constant measuring the amount of conflict between the sets, and given by:

$$K = 1 - \sum_{B_1 \cap \dots \cap B_S = \varnothing_\Omega} m^\Omega_1(B_1) \dots m^\Omega_S(B_S) \tag{10}$$

The normalisation is necessary as it distributes the mass of conflict on all the other masses, to maintain the sum at its expected value of 1. If the information sources are in agreement, K tends toward 1; on the other hand, if they are in total conflict, K tends toward 0, making coefficient $\frac{1}{K}$ very large. Use of that rule has come under serious criticism when significant conflict in the information is encountered [14,15], for the normalisation process “destroy” any information that we had on conflicts. In fact, conflict is a kind of information in itself, and the origin of this conflict becomes an issue, specially when the DSR makes a conscious assumption to *ignore* conflict. A solution to this problem will be outlined in Section 5.