Available online at www.sciencedirect.com
ScienceDirect
Information Fusion 9 (2008) 259–277
INFORMATION FUSION
www.elsevier.com/locate/inffus

A Bayesian approach to fusing uncertain, imprecise and conflicting information

Simon Maskell

QinetiQ, St. Andrews Road, Malvern, Worcestershire, WR14 3PS, UK

Received 13 July 2006; received in revised form 16 February 2007; accepted 16 February 2007
Available online 25 April 2007

Abstract

The Dezert–Smarandache theory (DSmT) and transferable belief model (TBM) both address concerns with the Bayesian methodology as applied to applications involving the fusion of uncertain, imprecise and conflicting information. In this paper, we revisit these concerns regarding the Bayesian methodology in the light of recent developments in the context of the DSmT and TBM. We show that, by exploiting recent advances in the Bayesian research arena, one can devise and analyse Bayesian models that have the same emergent properties as DSmT and TBM. Specifically, we define Bayesian models that articulate uncertainty over the value of probabilities (including multimodal distributions that result from conflicting information) and we use a minimum expected cost criterion to facilitate making decisions that involve hypotheses that are not mutually exclusive. We outline our motivation for using the Bayesian methodology and also show that the DSmT and TBM models are computationally expedient approaches to achieving the same endpoint. Our aim is to provide a conduit between these two communities such that an objective view can be shared by advocates of all the techniques.

Keywords: Information fusion; Bayesian; Uncertainty; Imprecision; Conflicting information; Transferable belief model; Dezert–Smarandache theory; Dempster–Shafer theory

1. Introduction

In information fusion applications, it is the representation of uncertainty that is the key enabler to extracting information from multi-sensor data (both co-modal data from multiple sensors of the same type and cross-modal data from sensors of different types). The development of all information fusion algorithms is critically dependent on using an appropriate method to represent uncertainty. A number of different paradigms have been developed for representing uncertainty and so performing data and information fusion, which are now briefly discussed:

- Fuzzy logic [1] represents belief through the definition of a mapping between quantities of interest and belief functions.
- Bayesian probability theory [2] articulates belief through the assignment of probability mass to mutually exclusive hypotheses.
- Dempster–Shafer theory (DST) [3] generalises Bayesian theory to consider upper and lower bounds on probabilities.
- The transferable belief model (TBM) [4] and Dezert–Smarandache theory (DSmT) [5] are further generalisations (over DST) of Bayesian theory. The TBM and DSmT represent uncertainty over the assignment of probability to mutually exclusive hypotheses by instead assigning probability to a power set of mutually exclusive hypotheses.
- Recently, a further generalisation, involving assignment of mass to a hyper-power set of hypotheses has been proposed [6].

Advocates of Bayesian theory make reference to a proof that Bayesian inference is the only way to consistently

E-mail address: s.maskell@signal.qinetiq.com

1566-2535/$ - see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.inffus.2007.02.003

manipulate belief relating to a set of hypotheses [7]. Conversely, advocates of DST, the TBM and DSmT motivate their approaches by the fact that given a set of hypotheses, Bayesian inference is unable to satisfactorily manipulate uncertain, imprecise and conflicting information [3–5]. This paper aims to act as a conduit between these two extreme viewpoints and the associated information fusion research communities. The hope is that this paper acts as a catalyst for the cross-fertilisation of ideas between these communities. The paper is intended to complement related work that has considered how one can subsume DST, the TBM and DSmT into a Bayesian approach [8] and approaches based on robust Bayesian inference [9]; this paper differs in that we explicitly consider how to devise Bayesian models that have the same emergent properties as analysis with DST, the TBM and DSmT.

The approach that is adopted is to accept that an initial application of Bayesian theory to fusion problems involving uncertain, imprecise and conflicting information is unable to satisfactorily manipulate such information. However, rather than attempt to redefine the method for manipulating belief on a given set of hypotheses, we choose to change the model definition and so the definition of the hypotheses. We show that, by exploiting recent advances in the Bayesian analysis of complex data (e.g. the recent development of, for example, particle filters [10] and Markov chain Monte-Carlo algorithms [11]), one can devise a rigorous Bayesian approach to fusing uncertain, imprecise and conflicting information. Furthermore, this approach has the same emergent properties as the TBM and DSmT, which can therefore be regarded as computationally efficient (although approximate) implementation strategies of this Bayesian approach.[1]

It should be noted that, as identified by the Bayesian community [12], model design is a critical component of a fusion system. Strong advocates of Bayesian inference will advocate the Bayesian methodology on the basis that this model design is made explicit. While making this explicit is useful, the problem of understanding how to design fusion systems remains whether model design is an implicit or explicit part of this process!

This paper is a rejection of the hypothesis that a Bayesian approach cannot solve certain problems involving the fusion of uncertain, imprecise and conflicting information. However, the author accepts that, while this paper demonstrates that an axiomatically consistent and robust Bayesian approach can be devised for such problems, specific system level constraints may dictate that approximations (such as those employed in the TBM and DSmT) should be used. The conclusions from any comparison is highly specific to the application being considered. So, this paper does not attempt to consider such comparisons, but aims to demonstrate that Bayesian approaches can and should be included in such comparisons in the future.

The paper begins in Section 2 with a description of how this Bayesian approach is devised. Section 3 considers several examples of how this approach is capable fusing uncertain, imprecise and conflicting information. Finally, Section 4 concludes.

2. Bayesian approach

2.1. Belief

Suppose an event has an outcome, x, that is one of a number of mutually exclusive hypotheses, $x \in X$. Furthermore, suppose one of these hypotheses is true, while the others are all false.

From a Bayesian (not frequentist) perspective, probability quantifies belief. To avoid confusion with belief functions, the term probability will be used from this point hence where appropriate. The probability associated with a hypothesis, $p(x)$, is a number that represents which of the mutually exclusive hypotheses we believe to be true. This probability is always non-negative and sums to unity across the hypotheses[2]:

$$p(x) \geqslant 0 \tag{1}$$

$$\sum_{x \in X} p(x) = 1 \tag{2}$$

Unfortunately, the true event is often very complex and cannot be modeled exactly. In such scenarios one must consider a model, which is an approximation to the real world. This approximation is chosen to be high enough fidelity that it captures the complexity of the event in terms of the parameters of interest but low enough fidelity that the probability can be calculated. It is this model complexity that is the key to the development of a Bayesian approach to fusing uncertain, imprecise and conflicting information (as shown in Section 3.2).

This model is the *prior*; it articulates the anticipated outcome of the event before any measurements are received. The choice of prior makes explicit all relevant knowledge of the system under consideration. Implicit consideration of prior knowledge as part of (for example) maximum likelihood modeling, is often equivalent to a specific explicit model of prior knowledge. However, there is a danger with implicit prior knowledge modeling that one unintentionally can introduce strong prior knowledge implicitly, as a result of parameterisation for example; one cannot be simultaneously ignorant of all parameterisations of a variable[3].

[1] The implication is that since TBM and DSmT approximate the only consistent way to manipulate beliefs, there will be scenarios where these approximations degrade performance significantly. Conversely, there will be scenarios where these approximations do not impact performance and are vital in facilitating real-time processing. Understanding which class of scenarios includes a given scenario remains an open research question.

[2] Open and closed worlds will be considered shortly.

[3] As a simple example, consider a point in a 2D plane. If one assumes all cartesian position of the point are equally likely, this puts a non-uniform prior on points when defined in polar co-ordinates. So, an uninformative prior on one parameterisation is not uninformative in another parameterisation.