

Multiple Model Tracking by Imprecise Markov Trees

Alessandro Antonucci and Alessio Benavoli and Marco Zaffalon

IDSIA, Switzerland

{[alessandro](mailto:alessandro@idsia.ch),[alessio](mailto:alessio@idsia.ch),[zaffalon](mailto:zaffalon@idsia.ch)}@idsia.ch

Gert de Cooman and Filip Hermans

SYSTeMS, Ghent University, Belgium

{[gert.decooman](mailto:gert.decooman@ugent.be),[filip.hermans](mailto:filip.hermans@ugent.be)}@ugent.be

Abstract – *We present a new procedure for tracking manoeuvring objects by hidden Markov chains. It leads to more reliable modelling of the transitions between hidden states compared to similar approaches proposed within the Bayesian framework: we adopt convex sets of probability mass functions rather than single ‘precise probability’ specifications, in order to provide a more realistic and cautious model of the manoeuvre dynamics. In general, the downside of such increased freedom in the modelling phase is a higher inferential complexity. However, the simple topology of hidden Markov chains allows for efficient tracking of the object through a recently developed belief propagation algorithm. Furthermore, the imprecise specification of the transitions can produce so-called indecision, meaning that more than one model may be suggested by our method as a possible explanation of the target kinematics. In summary, our approach leads to a multiple-model estimator whose performance, investigated through extensive numerical tests, turns out to be more accurate and robust than that of Bayesian ones.*

Keywords: tracking, Kalman filtering, imprecise Markov tree, hidden Markov chain, imprecise probability.

1 Introduction

An effective and robust tracking of highly manoeuvrable objects, especially if no accurate object/background representations are available, represents a challenging problem for a variety of important applications.

The most common approach to this problem consists in defining generative manoeuvring object models based on explicit on-line estimation of the manoeuvring kinematics. Formally, this corresponds to considering a *hidden Markov model* (HMM), whose hidden variables describe the model switching. The likelihoods of the models are usually estimated by some filtering process (e.g., Kalman filter) based on measurements

recorded by sensors, while a domain expert usually assesses the probabilities for the transitions between the hidden variables. Yet, it might seem questionable to enforce a very precise quantification of the transition probabilities between two hidden, and hence not directly observable, variables.

The present paper should be regarded as a first step towards a more robust, or more cautious, approach to modelling this generative process. In fact, the tracker we propose is based on an *imprecise-probability* HMM, whose hidden layer is quantified by convex sets of transition probability mass functions. This provides a perhaps more realistic, and certainly more cautious, model of the knowledge of the expert, who is for instance allowed to only give lower and upper bounds for the transition probabilities.

Such higher expressive power has computational drawbacks: inferences based on so-called imprecise-probability models are generally speaking more time-consuming than in their precise-probability counterparts.¹ Nevertheless, for the case of HMMs with their particularly simple topology, a linear-time algorithm for *exact* updating has been proposed recently [7]. This implies that we can estimate the right kinematic model of an object efficiently also under an imprecise probability quantification of the transitions probabilities between hidden variables.

For such imprecise probability models, the results of the inferences may be partially indeterminate, meaning that for a particular set of measurements the estimator may output more than a single model as a possible explanation of the object manoeuvre. We will show that our approach leads to a fast *multiple-model estimator*, which ensures higher accuracy and greater robustness than the ones proposed in the Bayesian framework (see [2] for a survey of multiple-model techniques in that framework).

¹As an example, inferences on polytree-shaped Bayesian nets can be efficiently computed [9], while the same problem for the corresponding imprecise-probability models, which are called credal nets, has been proved to be NP-hard [6].

The paper is organised as follows. In Section 2, we review the basics of the problem and the basic notation. Section 3 outlines general ideas underlying imprecise probability models, as well as particular features of a belief propagation algorithm, that will be employed for inferences, and was originally presented in [7]. A detailed description of our tracker is reported in Section 4. Results of simulations and a comparison with Bayesian approaches can be found in Section 5. Section 6 presents conclusions and an outlook on further research.

2 Bayesian model estimation

We consider tracking a manoeuvring object by a set of models \mathcal{M} . Both the set of models and the time window of the sampling process are discretised. We denote the sampling instants by integers $t \in \{1, \dots, \tau\}$, while the variables M_t and Z_t respectively refer to the model of the manoeuvre and the related measurement at instant t . The sequence $\{M_t\}_{t=1}^\tau$ should be regarded as a set of *hidden* variables, whose direct observation is not possible. In contrast, the actual values of the *manifest* variables $\{Z_t\}_{t=1}^\tau$, denoted by $\{z_t\}_{t=1}^\tau$, are available and correspond to the measurements recorded by the sensors. Model estimation is generally intended to mean the identification of the model that, given the particular set of measurements (z_1, \dots, z_τ) provided by the sensors, is most likely to correspond to the true value $m_\tau^* \in \mathcal{M}$ of M_τ . Such a problem is generally achieved by defining a probabilistic model over the whole set of variables, and hence by solving the following inferential task:

$$m_\tau^* := \arg \max_{m_\tau \in \mathcal{M}} P(M_\tau = m_\tau | z_1, \dots, z_\tau). \quad (1)$$

The solution of (1) requires the specification of a probabilistic model over the whole set of variables, both hidden and manifest. This specification task can be simplified by some structural assumptions involving conditional independence relations between these variables. The variables $\{M_t\}_{t=1}^\tau$ are usually regarded as a generative sequence satisfying the Markov assumption, i.e., the state at each time t depends only on the state at the previous time $t-1$. Similarly, each manifest variable Z_t depends only on the hidden variable M_t that it is supposed to measure, for each $t = 1, \dots, \tau$. These dependence/independence assumptions are depicted by the topology in Figure 1.

A model of this kind is called *hidden Markov model* [10], and it represents one of the most popular paradigms for modelling generative sequences. By exploiting the conditional independence relations captured in Figure 1, a global probabilistic model $P(M_1, Z_1, \dots, M_\tau, Z_\tau)$ can be specified in terms of ‘local’ conditional probabilistic models for the single variables and their immediate predecessors. A domain expert is usually asked to provide a quantification for the




Figure 1: Conditional dependence relations between hidden and observed (or manifest) variables.

hidden variables, by giving the marginal probabilities for M_1 and the conditional probabilities for the transitions $M_t \rightarrow M_{t+1}$, which are assumed *stationary*, i.e., independent of t . For the observable layer, the probabilities of the measurements $P(z_t | M_t)$ are usually quantified through some sequential filtering technique. Once the model has been quantified, algorithms based on belief propagation (e.g., [12]) can solve (1) and, thus, detect the model switching: when the measurements indicate that a previous best explanation or model for the trajectory must be replaced by a new one.

Most of the Bayesian algorithms proposed in the literature conform to the above setup (see [2] for a survey). Nevertheless, if the object moves according to the same model for a sufficiently long time window, the probability assigned to this model by Bayesian tracker generally increases with time. This leads to good accuracy for this part of the manoeuvre, but also causes a sudden worsening of the performance when the kinematics changes (after model switching; such behaviour is for instance evident in Table 3).

The most critical issue for this approach to modelling seems to be the probability quantification for the hidden layer by the domain expert. A wrong assessment of transition probabilities between hidden states might lead to completely wrong predictions, even if the filtering process (corresponding to the quantification of the observable layer) is accurate. And this is especially problematic because domain expertise is more often than not qualitative: it might therefore seem unnatural and perhaps unwise to force an expert to give a very precise quantification of the transition probabilities for variables that, by their very definition, are not directly observable.

For this reason, we describe in the next section how probabilistic graphical models like the ones described above, can be also quantified using sets of probability mass functions. This allows for more a reliable and robust representation of the domain expert knowledge, without unduly compromising the computational efficiency of the inferential process.

3 Imprecise-probabilistic graphical models

The Bayesian tracking treatment we have just described, rests on the assumption that all probability assessments are precisely known. If such is not the case, we need to robustify the results or at least, give experts the opportunity to quantify their knowledge in a less precise but more realistic way.

A very general way to cast off the yoke which precise probabilities sometimes bring along, involves calculations with closed convex sets of probability mass functions, also known as *credal sets*. Instead of specifying one single (precise) probability model, the expert can give bounds on the belief model, for example he can make statements like: “I believe that the probability of A is bigger than that of B ”. In general this will result in a set of precise probability models \mathcal{P} which we assume to be convex because of coherence reasons [13].

As Walley has shown in his seminal work [13], specifying a convex set \mathcal{P} of probability mass functions p on a finite dimensional variable X with probability space \mathcal{X} is equivalent to specifying its *lower and upper previsions* (or expectations), defined for any² $g \in \mathbb{R}^{\mathcal{X}}$ by

$$\underline{P}(g) := \min \left\{ \sum_{x \in \mathcal{X}} g(x)p(x) : p \in \mathcal{P} \right\},$$

$$\overline{P}(g) := \max \left\{ \sum_{x \in \mathcal{X}} g(x)p(x) : p \in \mathcal{P} \right\}.$$

These operators are bounded, non-negatively homogeneous and respectively super- and sub-additive. There is a one-to-one relationship between any credal set and the corresponding lower and upper prevision.

For the most part, in the case of a HMM we need to specify conditional mass functions instead of (marginal) mass functions. This is done in essentially the same way as described above, but now for every possible value of the parent variable a new conditional imprecise probabilistic model (either a credal set or a lower/upper prevision) needs to be specified. Given two variables X and Y where X is the parent variable of Y , we need to specify the value of $\overline{P}(g|x)$ for every $g \in \mathbb{R}^{\mathcal{Y}}$ and every possible value $x \in \mathcal{X}$ of X .

Of course, it is not determined yet what it means to be a parent in an HMM tree: we still need to fix and interpret the independence or irrelevance notion, and say how it is related to the graphical model of a tree. We adopt Walley’s epistemic irrelevance notion which states that X is *irrelevant to* Y whenever our belief model (read: lower prevision \underline{P}) about Y does not change when we learn something about X :

$$(\forall g \in \mathbb{R}^{\mathcal{Y}})(\forall x \in \mathcal{X})(\overline{P}(g) = \overline{P}(g|x)).$$

²The lower (upper) prevision does not necessarily need to be defined on the complete space $\mathbb{R}^{\mathcal{X}}$ of all real-valued maps on \mathcal{X} . Technically, it is the so-called lower and upper natural extensions we are talking about.

It is important to realise that the notion of irrelevance is no longer symmetric and will not imply d-separation in trees.

The interpretation of the graphical model. Consider any node Z in the tree, and its parent variable X . Then *conditional on the parent variable X , the non-parent non-descendant variables are assumed to be epistemically irrelevant to the variables associated with Z and its descendants*. This interpretation turns the tree into a *credal tree under epistemic irrelevance*; we also use the term *imprecise Markov tree* (IMT) for it.

In terms of the imprecise probabilistic models on Z and its children, this means that

$$\underline{P}(\cdot|X, S) = \underline{P}(\cdot|X),$$

for any set of nodes S that is a subset of the non-parent non-descendants of Z .

Until recently most of the research on algorithms for imprecise probabilistic graphical models has focused on models where a different notion of independence (called *strong independence*, [5]) was assumed. However, a computationally efficient exact algorithm for updating beliefs on the IMTs under epistemic irrelevance has very recently been developed [7]. The algorithm is in particular able to compute the lower (upper) prevision for a function of a variable in the tree, given some observed variables. Like the algorithms developed for precise graphical models, our algorithm works in a distributed fashion by passing messages along the tree. This leads to computing lower and upper conditional previsions (expectations) with a complexity that is essentially linear in the number of nodes in the tree.

4 Imprecise-probability based tracking

Following the ideas introduced in the previous section, we can easily generalise the Bayesian trackers described in Section 2 to allow for imprecise-probabilistic quantification. In particular, the domain expert is allowed to assess his beliefs about the kinematics by specifying lower and upper estimates for the transition probabilities. Let us denote these lower and upper bounds as $\underline{P}(M_{j+1}|M_j)$ and $\overline{P}(M_{j+1}|M_j)$. Similarly, the marginal probability mass function for M_1 is quantified by the expert’s lower and upper estimates $\underline{P}(M_1)$ and $\overline{P}(M_1)$. We will assume there is a precise-probabilistic quantification for the measurement probabilities.

This defines an *imprecise HMM*, whose topology (see Figure 1) is a special case of the more general tree topology for which the algorithm mentioned in the previous section has been developed. We can therefore efficiently update the beliefs about M_τ after the observation of the measurements (z_1, \dots, z_τ) , and end up with an imprecise-probability model $\underline{P}(\cdot|z_1, \dots, z_\tau)$. We will see that this achieves robust tracking of the target.

In order to do tracking, the criterion (1) must be extended to allow for imprecise- rather than precise-probabilistic posterior models $\underline{P}(\cdot|z_1, \dots, z_\tau)$. In situations of this kind, it cannot generally be excluded that more than one model provides an ‘optimal’ value for M_τ : we are thus led to *multiple-model* estimation, deriving from a condition of *indecision* between two or more models.³

In practise, a model $m_\tau \in \mathcal{M}$ is regarded as an ‘optimal’, or rather undominated, value for M_τ after the observation of the measurements (z_1, \dots, z_τ) , if there are no models $m'_\tau \in \mathcal{M}$ such that

$$\underline{P}(I_{m'_\tau} - I_{m_\tau} | z_1, \dots, z_\tau) > 0, \quad (2)$$

where I_{m_τ} is the indicator function of $m_\tau \in \mathcal{M}$.⁴ If such is the case, i.e., if (2) is satisfied for some m'_τ , then m_τ is rejected as an ‘optimal’ or undominated model. We adopt this optimality criterion, which is called *maximality* [13, Sect. 3.9.2], in order to identify the set $\mathcal{M}_\tau^* \subseteq \mathcal{M}_\tau$ of those models that are not rejected after the iteration of the test for each possible pair of models. In the next section, we show by numerical simulations that our approach generally leads to multiple-model robust estimates, meaning that \mathcal{M}_τ^* might include more than a single model, i.e., $|\mathcal{M}_\tau^*| \geq 2$.

5 Performance evaluation

Models. In order to test the performance of the tracker described in the previous section, we simulate a planar manoeuvre based on models $m^{(j)} \in \mathcal{M}$ corresponding to the following kinematics

$$\mathbf{x}(t+1) = \mathbf{f}_j(\mathbf{x}(t)) + \mathbf{w}_j(t), \quad (3)$$

where $\mathbf{x} := [x, y, v, h]'$, with x, y Cartesian coordinates of the position, v speed modulus, h heading angle, $\mathbf{w}_j(t)$ zero-mean noise with covariance $\mathbf{Q} = \text{diag}\{0, 0, \sigma_v^2 \Delta t, \sigma_h^2 \Delta t\}$, with Δt sampling period, and the components of the nonlinear function $\mathbf{f}_j(\mathbf{x}(t))$ are

$$\begin{bmatrix} x(t) + \frac{2v(t)}{\omega_t} \sin\left(\frac{\omega_t \Delta t}{2}\right) \cos\left(h(t) + \frac{\omega_t \Delta t}{2}\right) \\ y(t) + \frac{2v(t)}{\omega_t} \sin\left(\frac{\omega_t \Delta t}{2}\right) \sin\left(h(t) + \frac{\omega_t \Delta t}{2}\right) \\ v(t) \\ h(t) + \omega_t \Delta t \end{bmatrix}, \quad (4)$$

where $\omega_t := \dot{h}(t)$ is the angular speed. This is the *coordinated-turn model* [2]. Accordingly, the model

³Indecision between two or more options when the available evidence about the phenomenon of interest is not sufficiently detailed is commonly experienced in human reasoning. The fact that, unlike Bayesian approaches, imprecise probability models may produce this kind of indecision should be regarded as a further warrant for robustness.

⁴The difference between two indicators in (2) is in general not an indicator, and therefore the posterior lower expectation does not reduce to a lower probability. This is not a problem, as posterior lower expectations for arbitrary functions can be computed efficiently by the algorithm described in [7].

$m^{(j)}$ is completely specified by the value assigned to ω_t . In fact, the inclusion of the angular speed ω_t in the state vector and its estimation are not convenient for short-duration manoeuvres since there is little time for a reliable estimation of ω_t . For $\omega_t = 0$, (3) describes a motion with constant velocity and constant heading (straight motion). Conversely for $\omega_t \neq 0$ it describes a manoeuvre (*turn*) with constant angular speed ω_t , a *left turn* ($\omega_t > 0$) or a *right turn* ($\omega_t < 0$) depending on the sign of ω_t .

For our simulations, eleven candidate models are considered, i.e., $\mathcal{M} = \{m^{(-5)}, \dots, m^{(0)}, \dots, m^{(+5)}\}$ and the angular speed corresponding to $m^{(\pm j)}$ is assumed to be $\omega^{(\pm j)} = \pm j \cdot 0.15$ rad/s, for each $j = 0, \dots, 5$.

As noted in the previous section, for the probabilistic quantification of the hidden variables, we assume that the domain expert provides only his lower and upper bounds. More specifically, we assume that $\underline{P}(m_1) = \frac{1}{12}$ and $\overline{P}(m_1) = \frac{1}{10}$, for each $m_1 \in \mathcal{M}$, as an imprecise probability representation of indifference (equiprobability) between the models when the manoeuvre begins. For the transition probabilities, we consider the bounds expressed in Table 1 corresponding to an imprecision of about 5% in the estimates assessed by the expert.

Filtering. For model estimation, the filtering task has only the aim of estimating the likelihood of each model. For this purpose, a bank of eleven *extended Kalman filters* (EKF), each of which is based on a specific model tailored to a possible target behaviour (e.g., straight line motion, left turn, right turn etc.), has been selected for the filtering task. Each filter estimates the kinematic variables (e.g., positions, velocities, etc.) of the object using measurements and a model of the object motion. The relation between measurement and state is described by the following measurement equation, which is the same for all the models, and turns out to be

$$\mathbf{z}(t) = \mathbf{h}(\mathbf{x}(t)) + \mathbf{u}(t), \quad (7)$$

where polar coordinates are considered, i.e.,

$$\mathbf{h}(\mathbf{x}(t)) = \begin{bmatrix} \sqrt{x^2(t) + y^2(t)} \\ \angle(x(t) + jy(t)) \end{bmatrix}, \quad (8)$$

and $\mathbf{u}(t)$ is a zero-mean measurement noise with covariance $\mathbf{R} = \text{diag}\{\sigma_r^2, \sigma_\theta^2\}$.

The estimates returned by each filter are then combined according to a multiple-model approach. In particular, the *generalised pseudo-Bayesian 1* (GPB1) [1, 2, 4] filter has been adopted in this work for combining the estimates provided by the single filters. This choice is motivated by the fact that the GPB1 algorithm can be easily adapted to the imprecise HMM we presented in Section 4.

A full cycle of the GPB1 filtering process in both the precise and the IMT model is summarised in Table 2: the reader is referred to [2] for details about

$$\underline{\mathcal{P}}(M_{t+1}|M_t) = \begin{bmatrix} 77.50 & 6.00 & 3.75 & 2.75 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 5.00 & 77.50 & 5.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.00 & 5.00 & 77.50 & 5.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.00 & 0.00 & 5.00 & 77.50 & 5.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.00 & 0.00 & 0.00 & 5.00 & 77.50 & 5.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.00 & 0.00 & 0.00 & 0.00 & 5.00 & 77.50 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 5.00 & 77.50 & 5.00 & 0.00 & 0.00 & 0.00 \\ 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 5.00 & 77.50 & 5.00 & 0.00 & 0.00 \\ 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 5.00 & 77.50 & 5.00 & 0.00 \\ 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 2.75 & 3.75 & 6.00 & 77.50 \end{bmatrix} \quad (5)$$

$$\underline{\mathcal{P}}(M_{t+1}|M_t) = \begin{bmatrix} 82.50 & 11.00 & 8.75 & 7.75 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 10.00 & 82.50 & 10.00 & 5.50 & 4.50 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 4.00 & 10.00 & 82.50 & 10.00 & 4.00 & 3.50 & 3.50 & 0.00 & 0.00 & 0.00 & 0.00 \\ 3.00 & 4.50 & 10.00 & 82.50 & 10.00 & 4.50 & 3.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.00 & 3.50 & 4.00 & 10.00 & 82.50 & 10.00 & 4.00 & 3.50 & 0.00 & 0.00 & 0.00 \\ 0.00 & 0.00 & 3.50 & 4.00 & 10.00 & 82.50 & 10.00 & 4.00 & 3.50 & 0.00 & 0.00 \\ 0.00 & 0.00 & 0.00 & 3.00 & 4.50 & 10.00 & 82.50 & 10.00 & 4.50 & 3.00 & 0.00 \\ 0.00 & 0.00 & 0.00 & 0.00 & 3.00 & 4.50 & 10.00 & 82.50 & 10.00 & 4.50 & 3.00 \\ 0.00 & 0.00 & 0.00 & 0.00 & 3.50 & 3.50 & 4.00 & 10.00 & 82.50 & 10.00 & 4.00 \\ 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 4.50 & 5.50 & 10.00 & 82.50 & 10.00 \\ 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 7.75 & 11.00 & 82.50 \end{bmatrix} \quad (6)$$

Table 1: Lower and upper bounds of the transition probabilities percentages.

the GPB1. The design parameters in the GPB1 filter are the prior model transition probabilities and the model set. The transition probabilities π_{ij} correspond to the hidden layer in the Markov model that describes the model switching and are fixed in the Bayesian approach. Conversely, in the IMT approach, lower and upper bounds are assumed for π_{ij} and the *predicted model probability* and *model probability* in the GPB1 filter in Table 2 are calculated by using the algorithm in [7]. These are the outputs of the IMT filter, while the inputs correspond to the *model likelihoods*.

As a direct consequence of the imprecision in the model, also the posterior beliefs returned by the algorithm may be imprecise. Yet, a multiple-model filter coping with sets of posterior distributions for the model probabilities is not yet available (this is one of the directions of research we want to explore in our future work by combining the results in [3] with the multiple-model approach described in this paper). Thus, for the time being, in order to employ the GPB1 filter, a single ‘representative’ mass function for both the *predicted model probability* and *model probability* need to be chosen. In this paper, by analogy with the approaches in [11], we simply choose the centre of mass of the convex set of probability mass functions associated to the inferences provided by the algorithm.

Simulations. As a first experimental validation, we have tested the tracker on the trajectory depicted in Figure 2, which corresponds to the following sequence of models:

$$M_t = \begin{cases} m^{(0)} & 1 \leq t \leq 8 \\ m^{(+2)} & 9 \leq t \leq 13 \\ m^{(0)} & 14 \leq t \leq 20, \end{cases} \quad (9)$$

with sampling period $\Delta t = 4s$ (and hence a time window of 80s). Monte Carlo simulations have been performed by varying the measurements’ noise realisations. The values of the simulation parameters are $\sigma_r = 50$ m, $\sigma_\theta = 0.0087$ rad, $q_v = 0.01$ m/s, $q_h = 0.001$ rad, and $\mathbf{x}(0) := [10^4 \text{ m}, 10^4 \text{ m}, 5 \cdot 10^2 \text{ m/s}, 0.1 \text{ rad}]'$.

1. **Re-initialisation** (for $j = 1, \dots, |\mathcal{M}|$):

Predicted model probability:

$$p_j(t|t-1) = \sum_{i=1}^{|\mathcal{M}|} \pi_{ij} p_i(t-1)$$

Mixing estimate:

$$\bar{\mathbf{x}}(t-1|t-1) = \sum_{i=1}^{|\mathcal{M}|} p_i(t|t-1) \hat{\mathbf{x}}_i(t-1|t-1)$$

Mixing covariance:

$$\bar{\mathbf{P}}(t-1|t-1) = \sum_{i=1}^{|\mathcal{M}|} p_i(t|t-1) [P_i(t-1|t-1) +$$

$$(\bar{\mathbf{x}}(t-1|t-1) - \hat{\mathbf{x}}_i(t-1|t-1))(\bar{\mathbf{x}}(t-1|t-1) - \hat{\mathbf{x}}_i(t-1|t-1))']$$

2. **Update of the filter bank** (for $j = 1, \dots, |\mathcal{M}|$):

Model linearisation:

$$F_j(t) = \left[\frac{\partial f(\cdot)}{\partial \mathbf{x}} \right]_{\mathbf{x}=\bar{\mathbf{x}}(t-1|t-1)} \quad H_j(t) = \left[\frac{\partial h(\cdot)}{\partial \mathbf{x}} \right]_{\mathbf{x}=\hat{\mathbf{x}}_j(t|t-1)}$$

Predicted state:

$$\hat{\mathbf{x}}_j(t|t-1) = f_j(t-1, \bar{\mathbf{x}}(t-1|t-1))$$

Predicted covariance:

$$\hat{P}_j(t|t-1) = F_j(t-1) \bar{\mathbf{P}}(t-1|t-1) F_j'(t-1) + Q_j(t-1)$$

Measurement residual:

$$\tilde{z}_j(t) = z(t) - h(t, \hat{\mathbf{x}}_j(t|t-1))$$

Residual covariance:

$$S_j(t) = H_j(t) P_j(t|t-1) H_j'(t) + R_j(t)$$

Filter gain:

$$K_j(t) = P_j(t|t-1) H_j'(t) S_j^{-1}(t)$$

Update state:

$$\hat{\mathbf{x}}_j(t|t) = \hat{\mathbf{x}}_j(t|t-1) + K_j(t) \tilde{z}_j(t)$$

Update covariance:

$$P_j(t|t) = P_j(t|t-1) - K_j(t) S_j(t) K_j'(t)$$

3. **Model probability update** (for $j = 1, \dots, |\mathcal{M}|$):

Model likelihood:

$$L_j(t) = \mathcal{N}(\tilde{z}_j(t); 0, S_j(t))$$

Model probability:

$$p_j(k) = \frac{p_j(t|t-1) L_j(t)}{\sum_{i=1}^{|\mathcal{M}|} p_i(t|t-1) L_i(t)}$$

4. **Estimate fusion:**

Overall estimate:

$$\hat{\mathbf{x}}(t|t) = \sum_{i=1}^{|\mathcal{M}|} p_i(t) \hat{\mathbf{x}}_i(t|t)$$

Overall covariance:

$$P(t|t) = \sum_{i=1}^{|\mathcal{M}|} p_i(t) [P_i(t|t) + (\hat{\mathbf{x}}(t|t) - \hat{\mathbf{x}}_i(t|t))(\hat{\mathbf{x}}(t|t) - \hat{\mathbf{x}}_i(t|t))']$$

Table 2: One cycle of the GPB1 filter




Figure 2: The trajectory corresponding to the manoeuvre described by (9).

The aim of the simulations is to compare the imprecise- with the precise-probabilistic approach to the estimation of the right kinematic model. Two metrics are employed to evaluate the performances of the two approaches: *accuracy* and *average number of models*. The first measures the percentage of cases in which the right model is correctly estimated, while the second corresponds to the average number of models returned by the model estimator at each time instant. Clearly, the Bayesian model estimator is based on the criterion in (1) and always returns a single model, while the IMT estimator is based on the criterion in (2) and can therefore return more than one model.

The results are summarised in Table 3. The first thing to observe is that the imprecise-probabilistic approach always provides more accurate tracks. One downside might be the multiple-model estimation, as the IMT model estimator might provide indeterminate tracks. Nonetheless, this feature has been already noted to be a consequence of our cautious modelling of the generative sequence, and the number of possible models is significantly reduced (from eleven to less than four if $t > 3$).

We also observe (see Figure 3) a clear correlation between the number of models detected by the IMT estimator, and the accuracy of the Bayesian one. For manoeuvres that are ‘difficult’ to identify, the imprecise probabilistic approach tends to maintain a good accuracy by increasing the number of models, while the Bayesian estimator simply decreases its accuracy. On the other side, if a manoeuvre is ‘easy’ to identify (e.g., if the object moves according to the same model for many instants), the IMT gradually reduces the number of models while maintaining a high accuracy, whereas the Bayesian estimator tends to increase its accuracy. We also noted that, after model switching (e.g., $t = 9$ and $t = 14$) the accuracy of both estimators suddenly decreases. Yet, the IMT rapidly returns to a high level

of accuracy, while the Bayesian estimator needs much more time to do the same.

Moreover, we observe that, if two models are in \mathcal{M}_τ^* , then also every model whose angular speed is between those of the two models belongs to \mathcal{M}_τ^* . Accordingly, our multiple-model estimation corresponds to the identification of a range for the angular speed of the target.

| τ | IMT estimator accuracy | $ \mathcal{M}_\tau^* $ | Bayesian est. accuracy |
|--------|------------------------|------------------------|------------------------|
| 1 | 94.74 | 9.6 | 3.51 |
| 2 | 34.50 | 5.2 | 24.56 |
| 3 | 91.81 | 4.8 | 46.78 |
| 4 | 100.00 | 3.7 | 68.42 |
| 5 | 100.00 | 3.2 | 87.72 |
| 6 | 100.00 | 2.6 | 93.57 |
| 7 | 100.00 | 2.8 | 95.32 |
| 8 | 100.00 | 2.2 | 95.91 |
| 9 | 28.07 | 2.3 | 0.00 |
| 10 | 98.25 | 3.7 | 18.71 |
| 11 | 100.00 | 3.9 | 61.40 |
| 12 | 100.00 | 3.3 | 81.87 |
| 13 | 100.00 | 3.1 | 83.04 |
| 14 | 69.59 | 3.5 | 12.28 |
| 15 | 100.00 | 3.8 | 57.89 |
| 16 | 100.00 | 3.6 | 91.23 |
| 17 | 100.00 | 2.9 | 97.08 |
| 18 | 100.00 | 2.1 | 99.42 |
| 19 | 100.00 | 1.9 | 98.83 |
| 20 | 100.00 | 1.2 | 99.41 |

Table 3: IMT versus Bayesian estimator. Statistics are computed over 200 Monte Carlo runs of the simulated manoeuvre described in (9). The second and the third columns report respectively the percentage of runs for which the set of models returned by the IMT includes the true model of the manoeuvre, and the average of the number of models in this set. The fourth column reports the percentage of runs for which the model returned by the Bayesian estimator is the true one.

It should be also pointed out that the posterior probability mass functions computed by the Bayesian estimator is usually very strongly peaked around the model m_τ^* . Even by relaxing (1) with a threshold ϵ , such that the optimal models are the models $m_\tau \in \mathcal{M}$ for which

$$|P(m_\tau^*|z_1, \dots, z_\tau) - P(m_n|z_1, \dots, z_\tau)| \leq \epsilon, \quad (10)$$

no significant increase in the number of models is observed unless a very high threshold values (e.g. $\epsilon > 50\%$) is considered.

Regarding computational times, the computation of $P(\cdot|z_1, \dots, z_\tau)$ through the algorithm in [7] takes a time linear in the size of the graph in Figure 1 and hence in the time window τ , while only constant time is needed in the Bayesian case. This means that the IMT estima-




Figure 3: Average number of classes (circles) of the imprecise tracker versus accuracy (triangles) of the Bayesian tracker.

tor is $k \cdot \tau$ times slower than its Bayesian counterpart, with the constant $k \simeq 4$ in our particular problem.

In the filtering section, it has been explained that we use the GPB1 filter only for calculating the likelihoods of each model. However, since the GPB1 filter also provides estimates of the object's state, we can also calculate the position and speed root mean-square errors (RMSE). Figure 4 shows the RMSE of the following algorithms: the GPB1 filter and the IMT based GPB1 (IP-GPB1). The IP-GPB1 is calculated considering the centre of mass of the set of posteriors returned by the IMT as discussed in the filtering section. The choice of the centre of mass, although arbitrary, is very often used in the literature since, in general, it provides good results. In fact, as it is shown in Figure 4, this approach provides a lower RMSE than the classical GPB1. Peaks in the RMSE appear during the manoeuvres for the GPB1, while the RMSE of the IP-GPB1 is almost flat and lower than that of the classical GPB1. This shows that a robust estimation of the kinematic model is also beneficial for the estimation of the state of the object.

6 Conclusions and outlooks

We have presented a new approach to tracking highly manoeuvrable objects by hidden Markov models, with a imprecise-probabilistic quantification of hidden variables. This leads to a computationally fast multiple-model estimator that can possibly produce indeterminate tracks, where more than one model is suggested as a possible explanation of the object kinematics. In comparison with similar approaches proposed in the Bayesian framework, our approach displays good performances both in terms accuracy and robustness.

This work should be regarded as a starting point




Figure 4: Root mean square errors in the estimation of the object speed and position.

for future research focusing on the generalisation of the whole architecture, i.e., model and state estimation, to imprecise probability. In addition, connections with other multiple-model Bayesian estimators (e.g., [8]) should be considered for future work.

Acknowledgements

Research by De Cooman and Hermans has been supported by Flemish BOF-project 01107505. Research by Antonucci, Benavoli and Zaffalon has been partially supported by the Swiss NSF grants n. 200020-116674/1 and n. 200020-121785/1.

References

- [1] G. A. Ackerson and K. S. Fu. On state estimation in switching environments. *IEEE Transactions on Automatic Control*, 15:10–17, 1970.
- [2] Y. Bar-Shalom, X. Li, and T. Kirubarajan. *Estimation with Applications to Tracking and Navigation*. John Wiley & Sons, New York, 2001.
- [3] A. Benavoli, M. Zaffalon, and E. Miranda. Reliable hidden markov model filtering through coherent lower previsions. In *Proceedings of the 12th International Conference on Information Fusion*, Seattle (USA), 2009. Accepted for publication.
- [4] C. B. Chang and M. Athans. State estimation for discrete systems with switching parameters. *IEEE Transactions on Aerospace and Electronic Systems*, 14:418–425, 1978.
- [5] F. G. Cozman. Credal networks. *Artificial Intelligence*, 120:199–233, 2000.
- [6] C. P. de Campos and F. G. Cozman. The inferential complexity of Bayesian and credal networks. In

Proceedings of the International Joint Conference on Artificial Intelligence, pages 1313–1318, Edinburgh, 2005.

- [7] G. de Cooman, F. Hermans, A. Antonucci, and M. Zaffalon. Epistemic irrelevance in credal networks: the case of imprecise Markov trees. In *ISIPTA '09 – Proceedings of the Sixth International Symposium on Imprecise Probability: Theories and Applications*, Durham (UK), 2009. Accepted for publication.
- [8] X. R. Li, Youmin Zhang, and Xiarong Zhi. Multiple-model estimation with variable structure part iv: Design and evaluation of model-group switching algorithm. *Transactions on Aerospace and Electronic Systems*, 35, 1999.
- [9] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA, 1988.
- [10] L. Rabiner. A tutorial on HMM and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, February 1989.
- [11] P. Smets. Decision making in the TBM: the necessity of the pignistic transformation. *Int. J. Approx. Reasoning*, 38(2):133–147, 2005.
- [12] A. J. Viterbi. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*, 13:260–269, 1967.
- [13] P. Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London, 1991.