

2CoBel : An Efficient Belief Function Extension for Two-dimensional Continuous Spaces

Nicola Pellicanò, Sylvie Le Hégarat-Mascle, Emanuel Aldea
SATIE Laboratory, CNRS UMR 8029
Paris-Sud University, Paris-Saclay University, Orsay 91405, France
Email: {nicola.pellicano,sylvie.le-hegarat,emanuel.aldea}@u-psud.fr

Abstract—This paper introduces an innovative approach for handling 2D compound hypotheses within the Belief Function Theory framework. We propose a polygon-based generic representation which relies on polygon clipping operators. This approach allows us to account in the computational cost for the precision of the representation independently of the cardinality of the discernment frame. For the BBA combination and decision making, we propose efficient algorithms which rely on hashes for fast lookup, and on a topological ordering of the focal elements within a directed acyclic graph encoding their interconnections. Additionally, an implementation of the functionalities proposed in this paper is provided as an open source library. Experimental results on a pedestrian localization problem are reported. The experiments show that the solution is accurate and that it fully benefits from the scalability of the 2D search space granularity provided by our representation.

I. INTRODUCTION

Belief Function Theory (BFT) [1] [2] is an increasingly popular framework for the generalization of probability and possibility theory by modeling imprecision and partial ignorance of information, in addition to its uncertainty. BFT is widely used in fundamental tasks which benefit from multimodal information fusion, such as object detection and tracking [3] [4], object construction [5], outdoor localization [6], or autonomous robot mapping and tracking [7] [8]. Several public evidential theory libraries exist [9] [10] [11], but they are limited to 1D representations.

The main limitation, when dealing with such theory, since it copes with compound hypotheses, is the size of the set of hypotheses to handle, which may become intractable when the size of the exclusive hypothesis set increases. Such issue becomes critical especially at higher dimensions, as in the case of 2D space hypotheses. Moreover, different tasks may require different levels of precision for the solution, thus calling for a 2D space discretization which would increase quadratically the representation space size. In such a scenario, straightforward binary-word representation of hypotheses, as the one commonly used in 1D evidential theory, which allow bitwise operations and which are therefore very efficient, are no longer possible when the cardinality becomes greater than a few tens of possible solutions.

For such reasons, some works rely on different approaches to handle the 2D case: by proposing a smart sub-sampling of the 2D space to maintain tractability [12]; by proposing a sparse representation of the set of hypotheses, and by

keeping in memory only the ones which are carrying non-null information [6]. However, such proposals suffer from several problems which harm their use in practice. Sub-sampling based approaches suffer from non-scalability, since the operations defined by the framework are still dependent on the size of the frame. They are also precision-bounded, since they involve a coarse approximation of the space. On the other hand, current proposals for sparse representation still suffer from non-unique definition of compound hypotheses, from high accessory management costs, and from the need of non-unique space approximations.

Following the idea of providing a sparse representation for 2D BFT, and motivated by the great benefit that an efficient representation would carry to high dimensional problems, we propose a new two-dimensional representation which has full scalability properties with respect to the size of the hypothesis space, while allowing a theoretical infinite precision (bounded by the hardware precision limitations). The main contributions of this paper are:

- The proposal of a new polygon-based compound hypothesis representation, which makes use of polygon clipping operators as basis functions.
- The use of a hashable representation for fast lookup, and the proposal of a scale independent decision making algorithm.
- The release of a public library for multidimensional evidential theory, working with generic representations, and including the proposed definition.
- The demonstration of the interest of the proposed representation in pedestrian tracking with real data.

II. SETTINGS AND DEFINITIONS

Let us denote by Ω the *discernment frame*, i.e. the set of mutually exclusive hypotheses representing the solutions. The power set 2^Ω is the set of the Ω subsets, i.e. the disjunctions of the singleton hypotheses in Ω , having cardinality $2^{|\Omega|}$. The *mass function* m , specifying a basic belief assignment (BBA), is defined as $m : 2^\Omega \rightarrow [0, 1]$ such that $\sum_{A \in 2^\Omega} m(A) = 1$. A subset of hypotheses $A \in 2^\Omega$ such that $m(A) > 0$, is a *focal element* of m . A BBA is said to be *consonant* if the focal elements are nested: $\forall (A, B) \in 2^\Omega \times 2^\Omega m(A) > 0, m(B) > 0 \Rightarrow A \subseteq B \vee B \subseteq A$.

III. A GENERIC EXTENSION AND EFFICIENT VARIANTS

A. Focal element representation

Let us consider a 2D discernment frame Ω . We will refer, as reference, to the toy example illustrated in Figure 3. Such an example, inspired by [12], represents a typical localization scenario, where the discernment frame is a bounded region representing the ground plane. An exhaustive representation of Ω discrete hypotheses (as usually) implies a discretization of the area in a grid, where each cell of the grid represents a singleton hypothesis [12] [6]. Focal elements are then described by using a binary word, where a bit equal to 1 means that the cell belongs to the focal set. However, such representation suffers from major drawbacks when used in real world applications. Since there are $2^{|\Omega|}$ potential focal elements, large discernment frames become intractable, when the discretization resolution or the size of the whole area increase. In order to make the representation manageable, [12] proposes to condition the detections acquired from one sensor in its field of view, and to perform a coarsening at a lower spatial resolution of the focal elements, depending on the physical properties of the sensor. While these workarounds help in practice, they do not make the application fully scalable with the size of the scene, and they involve approximations such as the already cited coarsening, or frequent BBA simplification, which aims at maintaining under control the number of focal elements of the BBAs.

Such limitations derive from the fact that the complexity of any basic operator between focal elements (e.g. intersection, union,...) depends on the cardinality of the focal elements themselves. The works in [5] overcome this limitation by proposing a representation of any focal element as a set of rectangular boxes, and then by expressing the basic operators as performed on arrays of rectangles. In this setting the complexity of the basic operators will be a function of the number of boxes, but it will be independent from the cardinality of discernment frame. However, such representation suffers from some practical limitations. First, the representation is not unique. The same focal element may be represented by different sets of boxes, which do not allow for fast focal element comparisons and lookup. Second, the box set representation implies a non-unique approximation of the real focal element shape once edges are not parallel to the axes of reference. Geometric approximation of such focal elements may require a very large set of boxes when precision is a concern. Moreover, subsequent operations involve continuous box fragmentation which may be detrimental both for performance and for memory load. In order to avoid deep fragmentation, in [6] some representation simplification procedures are presented, which in turn increase the cost of BBA management.

We propose to represent the focal elements as generic polygons (or sets of polygons for focal elements having multiple connected components), by exploiting the capabilities of the generic 2D polygon clipping algorithms in the basic operator implementations. A focal element is represented by a set of closed paths, each of them represented by an ordered array of

vertexes (counterclockwise for positive areas, clockwise for holes). We exploit an extension of the Vatti's algorithm for clipping [13] implemented in the Clipper library [14].

The polygons are constrained to be simple, i.e. defined by closed simple paths (no crossing) and with a minimum number of vertexes (no vertex joining two co-linear edges). Under these constraints, the complexity of the basic operators between two polygons having n and m number of vertexes respectively, is $O(nm)$. Such lightweight representation presents also the advantages of uniqueness and precision. The (circular) vector of vertexes of a focal element provides a unique representation. The vertex coordinates use integer values for numerical robustness and correctness. This means that the continuous representation provided by polygons implies an underlining discretization. However, differently from the previous approaches, the coordinates can be rescaled at the desired level of precision (up to $\approx 10^{19}$) without any impact on the speed and memory requirements of the algorithm, being bounded only by the numerical representation limits of the hardware. This implies full scalability of the focal elements with respect to their size. Figure 1a shows an example of a focal element definition in the case of a localization application. The camera detection (red) is represented as a disk focal element, whereas the focal elements which have the shape of ring sectors embed the imprecision of the location and the ill-knowledge of the camera extrinsic parameters; the track (green) represents the location of the target at the previous frame, whereas its dilation is used in order to model the imprecision in its position introduced by time; the gray and blue focal elements belong to two different BBAs representing scene priors, of building and road presence respectively. The disk shaped focal elements are modeled as 64 to 128 vertexes regular polygons.

B. BBAs combination

Numerous combination rules exist in order to relate the information provided by two sources. When the sources m_1 and m_2 are independent, the conjunctive combination rule is the most popular among them:

$$\forall A \in 2^\Omega, m_1 \odot m_2(A) = \sum_{\substack{(B,C) \in \mathcal{F}_1 \times \mathcal{F}_2, \\ B \cap C = A}} m_1(B)m_2(C),$$

where \mathcal{F}_i is the set of focal elements of m_i . In computational terms, the rule involves the construction of a new BBA by performing intersection operations between all pairs of focal elements from the two BBAs. According to the sum in the previous equation, when creating a new focal element from an intersection, one has to check for its existence and add up masses if it already exists. Such necessity is not specific of the conjunctive rule, but it is shared with several other rules.

The above considerations justify the need for a BBA representation which allows for a fast lookup of a focal element in an array. The uniqueness and compactness of the proposed representation allow for an efficient and low collision prone hashing. The sparse set of focal elements of a given BBA can be stored in a hash table, where the circular vector of vertexes




Fig. 1: Toy localization example. (a) BBA definition through its focal elements: camera detection (red), track at $t-1$ (green), building presence mask (gray), road presence prior (blue). (b) Focal elements obtained as a result of performing a conjunctive combination over the defined BBAs. (c) Intersection graph for $BetP$ maximization and the result of graph simplification. The solid lines show the inclusion relationship, while the dashed lines highlight the intersection relationship. X^* is the set with maximum $BetP$ value.

is used to compute the hash. For a given polygon, its hash will be unique given a policy to decide the starting vertex (e.g. the top left). The array hashing function is equivalent to the one implemented in the Boost library's [15] `hash_range` method.

The binary-word representation, in comparison, uses the full word as a unique key. However, the key length (in number of bits) grows linearly with the cardinality of the discernment frame, needing the use of big data structures in order to store it. On the other side, the proposed hash has a fixed length, while having collision resistance property. The box set representation, being not unique, does not allow for direct hashing without the extraction of the minimal set of vertexes on the boundary. A cheap alternative could be to hash the bounding box of the focal element, but this could cause frequent collisions, since it is common to have spatially close focal elements related to the same BBA.

Figure 1b illustrates the result of the conjunctive combination of the sources introduced in Figure 1a. Seven focal elements are produced.

C. Decision making

Once different sources have been combined, the decision is generally taken on singleton hypotheses w by maximizing the

pignistic probability, defined as:

$$\forall w \in \Omega, BetP(w) = \frac{1}{1 - m(\emptyset)} \sum_{B \supseteq w} \frac{m(B)}{|B|}.$$

Even if the search space size is now $|\Omega|$, the decision making process is still dependent on the cardinality of the discernment frame, and thus not scalable, limiting the precision level which can be set for a specific context.

In order to overcome this limitation, we propose a maximization algorithm which is independent from the cardinality of the sets, and which is only related to the number of focal elements in the BBA. The underlying idea is that, since $BetP$ is an additive measure, its maximum value can be located only in areas of the discernment frame which present *maximal intersections*, defined as follows: given a set of focal elements $\mathcal{A} = \{A_1, \dots, A_n\}$, a maximal intersection I_m satisfies:

$$I_m = \bigcap_{A_k \in \tilde{\mathcal{A}}} A_k, \tilde{\mathcal{A}} \subseteq \mathcal{A}, |I_m| > 0 \text{ s.t.}$$

$$\nexists A_s \in \mathcal{A} \setminus \tilde{\mathcal{A}}, |A_s \cap I_m| > 0.$$

Finally the set X^* of hypotheses that maximizes the $BetP$ is researched within the set of maximal intersections \mathcal{I} :

$$X^* = \arg \max_{I_m \in \mathcal{I}} \frac{BetP(I_m)}{|I_m|},$$

where the $BetP$ function for compound hypotheses derives from the generalized formula:

$$\forall A \in 2^\Omega, BetP(A) = \frac{1}{1 - m(\emptyset)} \sum_{B \in \mathcal{F}, B \cap A \neq \emptyset} \frac{|A \cap B|}{|B|} m(B).$$

Consequently to this formulation, the $BetP$ maximization algorithm reduces to the subproblem of maximal intersection search. Let us assign an ordering to the set of focal elements for the given BBA. For optimization reasons explained further, the focal elements are labeled according to decreasing cardinality and the ordering follows the element label. We build a directed acyclic graph (DAG) $G = (V, E)$ where each node $v \in V$ is a focal element and an edge $e \in E$ represents a non empty intersection between two focal elements. The direction of the edges follows the given topological order. Each node is iteratively selected as the root. For each root a depth first search strategy is used to traverse the graph. The graph traversal is performed as follows: given the current node v_i , the intersection between all the nodes of the current path is propagated as I_i ; given an edge $e = (i, j)$, the node A_j is explored if $|I_j| = |I_i \cap A_j| > 0$. Such an operation is equivalent to performing a dynamic graph pruning which is a function of the current path. Once a leaf l is reached (a leaf is a node without any edge which can be further explored), the resulting I_l is a candidate for maximal intersection. However, it could be non-maximal, as its associated set \tilde{A} could be a subset of a maximal intersection which has already been found. So, when a maximal intersection I_m is found, the list p_m of focal sets involving it is stored (using a bit-set representation). Once the new candidate I_l is produced, the p_l list is tested for inclusion against the stored candidates (by an AND operation between the bit-sets). Even if the number of node visits can be very large in the worst case, in practice, the number of operations is much lower, since the dynamic pruning helps to cut out early dead paths.

Moreover, further optimization can be performed by inspecting the inclusion relationships between focal elements. Consider the node v_j as the current root. If $A_j \subseteq A_i$, for some $i < j$, there exists no maximal intersection including A_j and not A_i . This implies that no maximal intersection can be found starting from v_j as root. Thus, only the focal elements not included in others preceding them in topological order are used as root nodes (*root suppression*). This is the reason why we impose the topologically ordered in ascending order of cardinality, since any edge representing inclusion will be directed from the including to the included focal element. This allows us to exploit root suppression as much as possible.

Following the same principle, an *early stopping* criterion can be introduced. Let us consider the algorithm being executed for a root v_r . Given the current node in a path v_i and an edge $e = (v_i, v_j)$, the node v_j is explored only if it is not a subset of any previous root $v_k, k < r$. This derives from the fact that since v_k is no longer reachable, every path containing v_j is non maximal.

Given the mentioned topological order, a *graph simplification* can be applied to reduce the number of edges in the graph.

Given a node v_j having more than one incoming connection with a superset focal element, all the inclusion connections but the one from the highest index in topological order can be removed. The reason behind this is that any path which contains v_j must contain all its including sets, so, given a list of including nodes $\mathcal{V}_I = \{v_{I,1}, \dots, v_{I,k}\}$, a path between $v_{I,n}, n = 1 \dots k$ and v_j must include all the $v_{I,i}, n < i \leq k$, thus only $v_{I,k}$ can have a direct edge to v_j . Such optimization leads to clear performance gains when inclusion chains are present (such as when dealing with consonant BBAs). An inclusion chain of k elements leads to a complete subgraph in the output DAG, with 2^k possible paths. However, after graph simplification, only the edges going from element A_i to A_{i+1} are kept, resulting in a single path including all the nodes.

Figure 1c shows the intersection graph and its simplification for the proposed toy example. Two intersection graphs are present, and X^* is selected as the one at maximum $BetP$. For this example, raw traversal intersection graph performs 42 node visits, while with optimizations 12 are executed. On the other hand, a straightforward $BetP$ maximization by singleton hypothesis exploration would process 1100 locations (included into at least one focal element) with a factor 10 subsampling of the discernment frame.

IV. EXPERIMENTS

We present test results on a tracking application scenario, which makes use of the proposed representation on real data, as well as of our publicly available *2CoBel* library, embedding all the described methodologies, and exploited throughout the entire testing.

A. The *2CoBel* library

2CoBel is an open source¹ evidential framework embedding essential functionalities for generic BBAs definition, combination and decision making. An *Evidence* object defines common operations for a BBA containing any generic type of *FocalElement*. The current supported methods are: mass to Belief Functions conversion (plausibility, belief, commonality), conjunctive and disjunctive rules, vacuous extension and marginalization, conditioning, discounting, (generalized) $BetP$ computation, $BetP$ maximization (with singleton hypothesis enumeration or maximal intersections). Different types of *FocalElement* are supported, each of them defining basic operators (intersection, union, equality, inclusion) : *uni-dimensional* (hashable), representing the 1D focal element as a binary string; *2D bitmap*, providing a bitmap representation as in [12]; *2D box set*, implementing the definition and focal elements simplification operations proposed in [6]; *2D polygon* (hashable), implementing our proposed representation.

The library has full support for cartesian product of discernment frames.

¹Implementation available at: <https://github.com/MOHCANS-project/2CoBel>




Fig. 2: Example of pedestrian tracking steps. (a) Pedestrian detection blob. (b) Focal elements of detection BBA m_{d_0} on the ground plane at $t = 0$. (c) Focal elements of the conjunctive combination $\tilde{m}_{t_0,7}$ between the track and the associated detection at $t = 7$ (16 focal elements). (d) Focal elements of the BBA simplification of $\tilde{m}_{t_0,7}$ with the Joussemme's distance criterion (5 focal elements). (e) Focal elements after dilation of the track BBA $m_{t_0,8}$ by polygon offsetting.

B. Case study: pedestrian tracking

We apply the proposed representation to the problem of tracking pedestrians detected by imprecise sensors, on the ground plane. The belief function framework allows for direct modeling of the imprecision associated with the detections and the tracks and provides a measure for data association between detections and tracks.

We make use of the detector proposed in [16], which performs low level information fusion from multiple cameras in order to provide a dense pedestrian detection map, together with pedestrian height estimations, in a range between 1.4 m and 2 m. The output of the detector allows to project and track detections on the ground plane. We demonstrate the use of the 2D polygon representation provided in the 2CoBel library in order to perform joint multiple target tracking in the *Sparse* sequence presented in [16]. We perform tracking on the provided detections for 20 frames of the *Sparse* sequence, and we measure the localization error of the real tracks (13 pedestrians, 4 standing and 9 moving) with respect to the ground truth.

1) *Discernment frame definition*: The area under analysis is the ground plane region where the field of views of the cameras overlap. The area of the analysis region is 330 m^2 . The algorithm is run at a resolution of 10^{-4} m , so the cardinality of the discernment frame is $|\Omega| = 33 \times 10^9$. While the desired localization precision is 10^{-2} m , the chosen resolution is higher for increased robustness to rounding errors.

2) *BBA construction and assignment*: Given a detection d_i at time t located in (x_i, y_i) , we build a consonant BBA consisting of two focal elements. The first focal element is a disk centered at (x_i, y_i) and with a radius of 20 cm, taking into account the person's head and shoulder occupancy on the ground plane; the second focal element is a ring sector (approximated by a trapezoidal shape), which embeds the height uncertainty (on the direction point towards the camera location) and the camera calibration imprecision. In order to break the symmetry, the two focal elements are not assigned with 0.5 mass each, but with 0.51 for the internal disk and 0.49 for the trapezoid. In the presented case the choice of the mass allocation has a negligible impact on the quantitative results,

while it may become critical when additional sensors/sources are included into the problem.

3) *Data association and combination*: Given a set of tracks at time δ , $\mathcal{T} = \{t_1, \dots, t_k\}$ and a set of detections $\mathcal{D} = \{d_1, \dots, d_h\}$, the data association aims to compute an optimal one-to-one association set $A_t = \{(t_i, d_j), i \in \{1 \dots k\}, j \in \{1 \dots h\}\}$ with respect to some defined cost. A (t_i, \emptyset) association means that the track is into an inactive state (so it keeps propagating until it associates with a new detection or dies), while a (\emptyset, d_j) association means a new track has to be initialized with detection d_j . We make use of the criterion in [17] to define the association cost:

$$C_{t_i, d_j} = -\log(1 - m_{t_i} \odot m_{d_j}(\emptyset)),$$

which expresses the data association task as a conflict minimization problem, which can be solved by the use of the Hungarian algorithm.

The data association task is followed by a conjunctive combination which produces for every (t_i, d_j) the new track:

$$\tilde{m}_{t_i, \delta} = m_{t_i, \delta} \odot m_{d_j} \odot m_p,$$


where m_p corresponds to the prior. It performs a masking operation on the visible region of interest of the camera on the ground plane.

4) *BBA simplification*: A BBA simplification step is essential in tracking applications for two different reasons. First, we want to avoid that the number of focal elements grows without control as the time progresses, because it would mean that the real-time performance of the algorithm would degrade in time, bounding the maximum number of processed frames. Second, we want to avoid an excessive fragmentation of the belief. The BBA simplification aims at reducing the number of focal elements of a given BBA while respecting the least commitment principle. We adopt the method proposed in [12], which chooses iteratively two focal elements to aggregate (by performing a union operation) as the ones which minimize the Joussemme's distance [18] between the original BBA and the one obtained after the aggregation.


The proposed representation allows, conversely from the one in [12] (which simplifies the BBA after each conjunctive




(a)



(b)



(c)



(d)

Fig. 3: Pedestrian tracking. (a) Detection blobs on the image space ($t = 0$) estimated by the detector in [16]. Colors refer the estimated height values from 1.4 m (red) to 2 m (green). (b) Focal elements of the detection BBAs on the ground plane ($t = 0$). (c) Focal elements of track and detection BBAs on the ground plane ($t = 8$). Associated tracks and detections share the same color. (d) Final estimated tracks. Red crosses refer to target locations, while colored sets correspond to regions presenting maximum $BetP$ value.

combination), to perform the simplification on a less frequent time step. In the proposed experiment a target BBA is simplified when it reaches 15 focal elements, by producing a 5 focal elements BBA.

5) *BetP maximization*: At each time step, we run the $BetP$ maximization algorithm presented in Section III-C for each active track $\tilde{m}_{t_i, \delta}$ in order to extract the most probable location of the target. The cardinality of the resulting polygonal set represents the irreducible ambiguity in the target location. The target position is then estimated as the barycenter of the set.

6) *Modeling the imprecision of the tracks prediction*: Given the track $\tilde{m}_{t_i, \delta}$, which represents the result of the conjunctive combination, we need to model the prediction step imprecision. In order to model the track displacement from the current location, a random walk term is added to the track.

Such term boils down to an isotropic dilation of the focal elements. In the proposed representation, this corresponds to applying a scalable polygon offsetting algorithm, having $O(n \log n)$ complexity, where n is the number of vertexes. Polygon offsetting allows a dilation which respects the inclusion relationship of the original focal elements. The result of such step is the predicted track $m_{t_i, \delta+1}$ at time $t + 1$.

7) *Results*: In order to evaluate quantitatively the tracking accuracy, the target predicted locations are compared against an available ground truth. Such ground truth consists into coordinates in the image space where the heads are located. Since the height of such individuals is not known a priori, each location in the image space projects to a segment in the ground plane, allowing for any possible height in the interval of study. One computes the localization error as the distance between the target estimated location, and the ground




Fig. 4: Normalized histogram of the localization error of pedestrian tracking on the *Sparse* sequence.

Resolution	Average Localization error
$10^{-1} m$	31.061 <i>cm</i>
$10^{-2} m$	22.630 <i>cm</i>
$10^{-3} m$	20.111 <i>cm</i>
$10^{-4} m$	20.006 <i>cm</i>
$10^{-5} m$	20.002 <i>cm</i>

TABLE I: Average localization error on the *Sparse* sequence using different discretization resolutions. By using a representation able to deal with finer resolutions, one may achieve a significant performance gain.

truth head location, under the assumption that the height of such head corresponds to the predicted one. Such metric corresponds to computing the distance between the ground truth segment and a height uncertainty segment drawn at the target location. Target locations for inactive track states are estimated by linear regression fit of the estimated target positions at previous states.

Figure 4 shows the results in terms of (normalized) histogram of localization error. The average localization error is $\epsilon = 0.2 m$, which reaches the empiric limit set by the intrinsic uncertainty of head spatial occupation. On the other hand, the average localization error remains steady in time, meaning that the estimated tracks do not tend to drift away from the real ones. The standard deviation of the average localization error in time is $\sigma = 2.3 cm$.

Table I shows the average localization error obtained by the tracking algorithm for different choices of the resolution at which the discernment frame is discretized. When a coarse resolution of 10 *cm* is considered, the performance drops consistently. At this resolution the size of the discernment frame is already large enough to be intractable using methods based on binary representations, as in [12]. Moreover, while for the theoretically desired resolution of 1 *cm* the average localization error consistently drops, the proposed representation allows us to scale at finer resolutions to account for rounding errors, thus providing an additional performance boost.

V. CONCLUSION

This paper proposed a new representation for multi-modal information fusion in bi-dimensional spaces in the BFT domain. Such representation exhibits uniqueness, compactness, space and precision scalability, which make it suitable for intensive tasks constrained to large hypothesis spaces. We make available a public library for the community, in order to ease the reproducibility of such representation for active research. In our experiments, we show the effectiveness of this formulation on multi-target tracking scenarios, where tenths of tracks have to be estimated on a wide region of interest.

In our future work, we are interested to demonstrate the flexibility of the proposed representation by introducing richer BBAs for detections, in order to model the uncertainty of a detection blob centroid location, which require a non-regular polygon shaping tin order to be exploited. Moreover, we will extend the *2CoBel* library, by studying efficient canonical decomposition approaches.

In terms of application perspectives, we are interested in developing a tracking algorithm for dense crowds, by performing cautious fusion of multiple detection sources from a smart camera network. We aim to demonstrate the use of the proposed representation to make such algorithm scale for high density crowds, for which the number of targets to track jointly can be intractable for state-of-the-art tracking frameworks.

ACKNOWLEDGMENT

This work was supported by ANR grant ANR-15-CE39-0005. We gratefully acknowledge the support from Regent’s Park Mosque for providing access to the site during the collection of the data used for illustrating our contribution.

REFERENCES

- [1] A. P. Dempster, “A generalization of bayesian inference.” in *Classic works of the dempster-shafer theory of belief functions*. Springer, 2008, pp. 73–104.
- [2] G. Shafer, *A mathematical theory of evidence*. Princeton university press, 1976, vol. 42.
- [3] R. O. Chavez-Garcia and O. Aycard, “Multiple sensor fusion and classification for moving object detection and tracking,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 2, pp. 525–534, 2016.
- [4] T. Denoeux, N. El Zoghby, V. Cherfaoui, and A. Joulet, “Optimal object association in the dempster–shafer framework,” *IEEE transactions on cybernetics*, vol. 44, no. 12, pp. 2521–2531, 2014.
- [5] W. Rejik, S. Le Hégarat-Masclé, R. Reynaud, A. Kallel, and A. Ben Hamida, “Dynamic object construction using belief function theory,” *Information Sciences*, vol. 345, pp. 129 – 142, 2016.
- [6] S. Zair and S. Le Hégarat-Masclé, “Evidential framework for robust localization using raw gnss data,” *Engineering Applications of Artificial Intelligence*, vol. 61, pp. 126 – 135, 2017.
- [7] G. Tanzmeister, J. Thomas, D. Wollherr, and M. Buss, “Grid-based mapping and tracking in dynamic environments using a uniform evidential environment representation,” in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 2014, pp. 6090–6095.
- [8] M. Kurdej, J. Moras, V. Cherfaoui, and P. Bonnifait, “Controlling remanence in evidential grids using geodata for dynamic scene perception,” *International Journal of Approximate Reasoning*, vol. 55, no. 1, pp. 355–375, 2014.
- [9] M. Kurdej, “BFT - Belief Functions Theory library,” <https://github.com/mkurdej/bft>, 2014, last accessed 2018-03-13.
- [10] T. Reineking, “Dempster-Shafer theory library,” https://pypi.python.org/pypi/py_dempster_shafer/0.7, 2014, last accessed 2018-03-13.

- [11] A. Martin, "Matlab toolbox for belief functions," <http://www.arnaud.martin.free.fr/Doc>, 2014, last accessed 2018-03-13.
- [12] C. André, S. Le Hégarat-Masclé, and R. Reynaud, "Evidential framework for data fusion in a multi-sensor surveillance system," *Engineering Applications of Artificial Intelligence*, vol. 43, pp. 166 – 180, 2015.
- [13] B. R. Vatti, "A generic solution to polygon clipping," *Commun. ACM*, vol. 35, no. 7, pp. 56–63, Jul. 1992. [Online]. Available: <http://doi.acm.org/10.1145/129902.129906>
- [14] A. Johnson, "Clipper - an open source freeware library for clipping and offsetting lines and polygons." <http://www.angusj.com/delphi/clipper.php>, 2014, last accessed 2018-03-13.
- [15] Boost, "Boost C++ Libraries," <http://www.boost.org/>, 2015, last accessed 2018-03-13.
- [16] N. Pellicanò, E. Aldea, and S. Le Hégarat-Masclé, "Geometry-Based Multiple Camera Head Detection in Dense Crowds," in *28th British Machine Vision Conference (BMVC) - 5th Activity Monitoring by Multiple Distributed Sensing Workshop*, Londres, United Kingdom, Sep. 2017. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-01691761>
- [17] B. Ristic and P. Smets, "The tbm global distance measure for the association of uncertain combat id declarations," *Information fusion*, vol. 7, no. 3, pp. 276–284, 2006.
- [18] A.-L. Jousselme, D. Grenier, and É. Bossé, "A new distance between two bodies of evidence," *Information fusion*, vol. 2, no. 2, pp. 91–101, 2001.