# Real-Time Detection and Prediction of Relative Motion of Moving Objects in Autonomous Driving 

Lalintha G. Polpitiya, ${ }^{1}$ Kamal Premaratne ${ }^{2}$<br>${ }^{1}$ The MathWorks, Inc., Natick, Massachusetts, USA, lpolpiti@mathworks.com<br>${ }^{2}$ University of Miami, Coral Gables, Florida, USA, kamal@ miami.edu


#### Abstract

Autonomous driving vehicles must have the ability to identify and predict behaviors of surrounding moving objects (e.g., other vehicles, cyclists, and pedestrians) in real-time. This is especially true in urban environments, where interactions become more complex due to high volumes of traffic. The work in this paper harnesses the Dempster-Shafer (DS) theoretic framework's ability to capture and account for various types of evidence uncertainty to develop a robust event detection and prediction model, which is appropriately calibrated to account for the underlying uncertainty so that it may be employed to arrive at a more informed decision.


## Introduction

Advances in sensor and AI technologies have propelled rapid developments in autonomous driving during the past decade. Numerous research projects (DARPA 2014) aimed at solving the autonomous driving challenges have accelerated these developments. Today, scientists are diligently working at making fully self-driving cars a reality.

Fully driver-less cars, however, require effective methodologies for handling uncertainties and exceptions that may arise from the behavior of nearby vehicles (e.g., sudden turns at intersections and lane changes, mechanical breakdowns, vehicles blocking the roadway, etc.), cyclists and pedestrians (e.g., abrupt movement onto the road), navigating in highly congested traffic situations, bad weather, and other unpredictable and unanticipated situations. With the increased automation required by the next levels of autonomous driving tasks as identified in the Society of Automotive Engineers (SAE) taxonomy (SAE 2018), there is heightened interest in how to manage these uncertain and exception conditions.

The lack of effective methods to handle such situations is considered perhaps the most serious roadblock preventing the development of fully autonomous vehicles (AVs) (Ottley 2019). Current exceptions and uncertainty handling mechanisms rely heavily on the involvement of the human driver (Burns 2018; Kaber 2018). Irrespective of whether a human is involved in the decision-making process, we believe that the information provided should be properly calibrated

[^0]to capture the underlying uncertainty so that the decisionmaker can make a fully informed decision. The DS theoretic (DST) (Dempster 1968; Shafer 1976) models are ideally suited to capture and propagate these underlying uncertainties throughout the decision-making pipeline.

We examine this supposition by developing a DST model to capture the intended relative movements of other vehicles and pedestrians. With the predicted behavior of a pedestrian crossing the road or a vehicle making a turn, together with the uncertainty associated with this prediction in hand, the AV (or the human-in-the-loop) can arrive at a more informed decision as to the action it could take. In turn, with an appropriately calibrated model which captures the underlying uncertainty as the basis on which better more informed decisions are made, this work has the potential to significantly contribute to reduce driver fatigue and the number of traffic accidents, solve the problem of traffic congestion, allow individuals who do not (or cannot) drive more freedom, as well as make commuting easier and safer.

This paper is organized as follows: First, we provide key terminology in common use for autonomous driving and DST basic notions. This is followed by the proposed model for lateral and longitudinal relative motion of surrounding moving objects. The experimental results are provided next, followed by concluding remarks.

## Preliminaries

Levels in Autonomous Driving. The level of driver attention, and how much driver action is needed to perform the driving task, are important factors to consider when attempting to define the level of autonomy. The SAE defines six levels of driving automation (SAE 2018): 0 (no automation), 1 (driver assistance), 2 (partial automation), 3 (conditional automation), 4 (high automation), and 5 (full automation). These levels depend on whether the vehicle can perform all or part of the dynamic driving task (DDT). In levels 1-2, the human driver monitors the driving environment; in levels 35, an automated driving system (ADS) performs all or part of the DDT on a sustained basis.

The Driving Task. The driving task that an AV would have to execute involves the following (SAE 2018): (a) Lateral control: Turning left, right, tracking a turn. (b) Longi-
tudinal control: Maintaining set speed and appropriate gap, detecting preceding vehicle, appropriate braking (acceleration/deceleration to maintain speed). (c) Object and Event detection: Real-time environmental object and event detection, recognition, classification and response preparation, i.e., first to identify the object (e.g: car, cyclist, pedestrian etc), and second to understand the motion of the object and see if we can infer what it will do next. (d) Planning: Informed decision-making aspect which follows the information gathering via perception and behavior prediction. This includes long and short term plans needed to travel to a destination, maneuvers such as lane changes, intersection crossings, and driving around other moving objects or traffic participants. (e) Miscellaneous: Identifying actions such as hand-waving, signaling, and interacting with drivers.

DST Basic Notions. In DS theory, the frame of discernment ( $F o D$ ) refers to the set of all possible mutually exclusive and exhaustive propositions (Shafer 1976). We consider the case where the FoD is finite and we denote it as $\Theta=\left\{\theta_{0}, \theta_{1}, \ldots, \theta_{n-1}\right\}$. Proposition $\left\{\theta_{i}\right\}$, which is referred to as a singleton, represents the lowest level of discernible information. The power set of $\Theta$, denoted by $2^{\Theta}$, form all the propositions of interest in DS theory. A proposition that is not a singleton is referred to as a composite. The set $A \backslash B$ denotes all singletons in $A \subseteq \Theta$ that are not included in $B \subseteq \Theta$, i.e., $A \backslash B=\left\{\theta_{i} \in \Theta \mid \theta_{i} \in A, \theta_{i} \notin B\right\}$. We use $\bar{A}$ to denote $\Theta \backslash A$ and $|A|$ to denote the cardinality of $A$.
Definition 1. The mapping $m: 2^{\Theta} \mapsto[0,1]$ is said to be $a$ basic belief assignment (BBA) or $a$ mass assignment if $m(\emptyset)=0$ and $\sum_{A \subseteq \Theta} m(A)=1$. The belief assigned to $A \subseteq$ $\Theta$ is $B l: 2^{\Theta} \mapsto[0,1]$ where $B l(A)=\sum_{B \subseteq A} m(B)$.

So, the mass captures the 'support' that is strictly allocated to a given proposition and the belief assigned to a proposition takes into account the support for all of its subsets. The mass of a composite proposition (a general focal element) is free to move into its subsets (e.g., into individual singletons), which allows one to model the notion of ignorance. Complete ignorance is captured via the vacuous $B B A$ $1_{\Theta}: m(A)=1$ for $A=\Theta$, and $m(A)=0$ for $A \subset \Theta$. Propositions that possess non-zero mass are referred to as focal elements; the set of all focal elements in an FoD is referred to as its core $\mathfrak{F}$, i.e., $\mathfrak{F}=\{A \subseteq \Theta \mid m(A)>0\}$. Note that $|\mathfrak{F}|$ is the number of focal elements. $\mathcal{E}=\{\Theta, \mathcal{F}, m(\cdot)\}$ is referred to as the body of evidence (BoE).

## Relative Motion of Adjacent Moving Objects

DST Model. We now develop a DST model which captures the uncertainties that are inherent in the evidence generated by vehicles, pedestrians, cyclists, and other moving objects in uncertain road situations. To this end, we consider the lateral relative movement (as Left, Center, or Right) and the longitudinal relative movement (as Away, Stationary, and Toward) relative to the AV. Fig. 1 illustrates these movements when another car, a cyclist, and a pedestrian are present near the AV . These categories will be further divided into slow


Figure 1: Lateral and longitudinal relative motion of moving objects. Lateral motion: Left, Center, and Right; longitudinal motion: Away, Stationary, and Toward.
and fast, giving ten categories of movement as Fast Left (FL), Slow Left (SL), Center (C), Slow Right (SR), and Fast Right (FR) for lateral movement and Fast Away (FA), Slow Away (SA), Stationary (S), Slow Toward (ST), and Fast Toward (FT) for longitudinal movement.

Each moving object in the near vicinity of the AV is allocated a pair of lateral and longitudinal DST BoEs with the following two DST FoDs:

$$
\begin{align*}
& \text { Lateral FoD: } \quad \Theta_{U}=\{F L, S L, C, S R, F R\} ; \\
& \text { Longitudinal FoD: } \Theta_{V}=\{F A, S A, S, S T, F T\} . \tag{1}
\end{align*}
$$

Initialization of the BoEs. Let us link track $T_{i}$ which may span several frames of a video signal, with object $O_{i}$. For $n$ objects of interest, we have the tracks $\left\{T_{0}, T_{1}, \ldots, T_{n-1}\right\}$. We denote the object $O_{i}$ in the previous and current frames as $O_{i}^{\prime}$ and $O_{i}^{\prime \prime}$, respectively. We denote the centroid of $O_{i}^{\prime}$ as $\left\{x^{\prime}, y^{\prime}\right\}$, and the width and height of the relevant bounding box as $w^{\prime}$ and $h^{\prime}$, respectively. Similarly, the centroid of $O_{i}^{\prime \prime}$ is $\left\{x^{\prime \prime}, y^{\prime \prime}\right\}$, and the width and height of the relevant bounding box are $w^{\prime \prime}$ and $h^{\prime \prime}$, respectively. See Fig. 2.

Consider the object $O_{i}$ on track $T_{i}$. For convenience, we will ignore the subscript $i$. The objects $O^{\prime}$ and $O^{\prime \prime}$, the object $O$ (again, we ignore the subscript $i$ ) in the previous and current frames respectively, are endowed with two BoEs, one defined on the lateral FoD $\Theta_{U}$ and the other defined on the longitudinal FoD $\Theta_{V}$ (see (1)). Let us denote these BoEs as $\mathcal{E}_{U}^{\prime}=\left\{\Theta_{U}, \mathfrak{F}_{U}^{\prime}, m_{U}^{\prime}(\cdot)\right\}$ and $\mathcal{E}_{V}^{\prime}=\left\{\Theta_{V}, \mathfrak{F}_{V}^{\prime}, m_{V}^{\prime}(\cdot)\right\}$ for the previous frame and $\mathcal{E}_{U}^{\prime \prime}=\left\{\Theta_{U}, \mathfrak{F}_{U}^{\prime \prime}, m_{U}^{\prime \prime}(\cdot)\right\}$ and $\mathcal{E}_{V}^{\prime \prime}=\left\{\Theta_{V}, \mathfrak{F}_{V}^{\prime \prime}, m_{V}^{\prime \prime}(\cdot)\right\}$ for the current frame.

The starting BoE (the BoE associated with the first frame of track $T_{i}$ ) is taken to be vacuous, i.e.,

$$
m(A)=1_{\Theta_{U}}\left(\text { or } 1_{\left.\Theta_{V}\right)}\right) \equiv \begin{cases}1, & \text { for } A=\Theta_{U}\left(\text { or } \Theta_{V}\right) ;  \tag{2}\\ 0, & \text { for } A \subset \Theta_{U}\left(\text { or } \Theta_{V}\right)\end{cases}
$$

Thus, the belief and plausibility of each proposition (except the full FoD itself) is 0 and 1 , respectively.


Figure 2: Updating the lateral and longitudinal BoEs of object $O_{i}$ with CUE. $\mathcal{E}^{\prime \prime}=\left(\mathcal{E}^{\prime} \triangleleft \mathcal{E}^{T R}\right)$. Previous frame: centroid $\left\{x^{\prime}, y^{\prime}\right\}$, width $w^{\prime}$, height $h^{\prime}$; current frame: centroid $\left\{x^{\prime \prime}, y^{\prime \prime}\right\}$, width $w^{\prime \prime}$, height $h^{\prime \prime}$.

DST Modeling of New Evidence. New evidence is extracted by comparing the relative lateral and longitudinal movement between $O^{\prime}$ and $O^{\prime \prime}$. Let us refer to these BoEs as lateral and longitudinal transmit-BoEs, and identify them as $\mathcal{E}_{U}^{T R}=\left\{\Theta_{U}, \mathfrak{F}_{U}^{T R}, m_{U}^{T R}(\cdot)\right\}$ and $\mathcal{E}_{V}^{T R}=$ $\left\{\Theta_{V}, \mathfrak{F}_{V}^{T R}, m_{V}^{T R}(\cdot)\right\}$, respectively.

The BBAs associated with $\mathcal{E}_{U}^{T R}$ and $\mathcal{E}_{V}^{T R}$ are determined as follows. When an object is moving left or right, the value of $x$ will change; when an object is moving away or toward, the value of $y$ and also the area of the detected object will change. We use these changes to allocate DST masses which quantify the support the new evidence has for each proposition from the two FoDs $\Theta_{U}$ and $\Theta_{V}$.

Algorithm 1 shows how these DST masses are computed. We use two constant parameters in this algorithm: $\Pi, \Gamma: \Pi$ is used to classify the lateral focal element as either $F L$ or $S L$ (or either $F R$ or $S R$ ); and $\Gamma$ is used to classify the longitudinal focal element as either $F A$ or $S A$ (or either $F T$ or $S T$ ). The exact values of $\Pi$ and $\Gamma$ depend on the configurations of the video input. Another parameter used in Algorithm 1 is the detection confidence score $S$ which captures the confidence of the new evidence generated from $O_{i}^{\prime \prime}$ via the AV's pedestrian or vehicle detector.

The steps of Algorithm 1 are as follows: (a) Line \#1 takes as inputs $x^{\prime}, y^{\prime}, x^{\prime \prime}, y^{\prime \prime}, B B A_{U}^{T R}, B B A_{V}^{T R}, \Pi, \Gamma, S$. The centroid of $O_{i}^{\prime}$ is $\left\{x^{\prime}, y^{\prime}\right\}$ and the centroid of $O_{i}^{\prime \prime}$ is $\left\{x^{\prime \prime}, y^{\prime \prime}\right\}$. We represent the BoEs $B B A_{U}^{T R}$ and $B B A_{V}^{T R}$ as $D S$-Vectors (although one could also employ $D S$-Matrix or DS-Tree structures) (Polpitiya et al. 2016; 2017). $\Pi$ and $\Gamma$ values are constants for the entire run; the value $S$ varies with every new detection. (b) Lines \#2-12 determine the lateral focal element and assigning the mass. Mass of the lateral focal element will be $S$. (c) Line $\# 13$ assigns value $1-S$ to $m_{U}^{T R}\left(\Theta_{U}\right)$. (d) Lines \#14-24 determine the longitudinal focal element; its mass is $S$. (e) Line \#25 assigns value $1-S$ to $m_{V}^{T R}\left(\Theta_{V}\right)$. (f) Line \#26 returns $B B A_{U}^{T R}$ and $B B A_{V}^{T R}$

```
Algorithm 1 Determining the masses of \(\mathcal{E}_{U}^{T R}\) and \(\mathcal{E}_{V}^{T R}\)
    procedure ASSIGNMASSES \(\left(x^{\prime}, y^{\prime}\right.\),
    \(\left.x^{\prime \prime}, y^{\prime \prime}, B B A_{U}^{T R}, B B A_{V}^{T R}, \Pi, \Gamma, S\right)\)
        if \(x^{\prime}>\Pi+x^{\prime \prime}\) then
            \(m_{U}^{T R}(F L) \leftarrow S\)
        else if \(x^{\prime}>x^{\prime \prime}\) then
            \(m_{U}^{T R}(S L) \leftarrow S\)
        else if \(x^{\prime}+\Pi<x^{\prime \prime}\) then
            \(m_{U}^{T R}(F R) \leftarrow S\)
        else if \(x^{\prime}<x^{\prime \prime}\) then
            \(m_{U}^{T R}(S R) \leftarrow S\)
        else
            \(m_{U}^{T R}(C) \leftarrow S\)
        end if
        \(m_{U}^{T R}\left(\Theta_{U}\right) \leftarrow 1-S\)
        if \(y^{\prime}>\Gamma+y^{\prime \prime}\) then
                \(m_{V}^{T R}(F A) \leftarrow S\)
        else if \(y^{\prime}>y^{\prime \prime}\) then
                \(m_{V}^{T R}(S A) \leftarrow S\)
        else if \(y^{\prime}+\Gamma<y^{\prime \prime}\) then
                \(m_{V}^{T R}(F T) \leftarrow S\)
        else if \(y^{\prime}<y^{\prime \prime}\) then
            \(m_{V}^{T R}(S T) \leftarrow S\)
        else
            \(m_{V}^{T R}(S) \leftarrow S\)
        end if
        \(m_{V}^{T R}\left(\Theta_{V}\right) \leftarrow 1-S\)
        Return \(B B A_{U}^{T R}, B B A_{V}^{T R}\)
    end procedure
```


## DS-Vectors.

Evidence Updating. For updating the BoEs, we use the conditional update equation (CUE) and conditional fusion equation (CFE):
Definition 2 (CUE). (Premaratne et al. 2009) Suppose the receive-BoE $\mathcal{E}^{\prime}=\left\{\Theta, \mathfrak{F}^{\prime}, m^{\prime}(\cdot)\right\}$ is presented with the evidence provided by the transmit-BoE $\mathcal{E}^{T R}=$ $\left\{\Theta, \mathfrak{F}^{T R}, m^{T R}(\cdot)\right\}$. Then, the CUE-based BBA update of the receive-BoE $\mathcal{E}^{\prime \prime}$ is

$$
m^{\prime \prime}(B)=\alpha m^{\prime}(B)+(1-\alpha) \sum_{A \in \mathfrak{F}^{T R}} \beta(A) m^{T R}(B \mid A)
$$

for all $B \subseteq \Theta$. Here, the CUE parameters $\{\alpha, \beta(\cdot)\}$ are non-negative real and satisfy

$$
\alpha+(1-\alpha) \sum_{A \in \mathfrak{F}^{T R}} \beta(A)=1, \forall k \in \mathbb{N}_{+}
$$

We denote this operation as $\mathcal{E}^{\prime \prime}=\left(\mathcal{E}^{\prime} \triangleleft \mathcal{E}^{T R}\right)$.
Definition 3 (CFE). (Wickramarathne, Premaratne, and Murthi 2012) The CFE that fuses the evidence of $\mathcal{E}^{T R 1}=$ $\left\{\Theta, \mathfrak{F}^{T R 1}, m^{T R 1}(\cdot)\right\}$ and $\mathcal{E}^{T R 2}=\left\{\Theta, \mathfrak{F}^{T R 2}, m^{T R 2}(\cdot)\right\}$ is

$$
\begin{aligned}
m^{T R}(B)= & K^{T R 1} \sum_{A \in \mathfrak{F}^{T R 1}} \beta^{T R 1}(A) m^{T R 1}(B \mid A) \\
& +K^{T R 2} \sum_{A \in \mathfrak{F}^{T R 2}} \beta^{T R 2}(A) m^{T R 2}(B \mid A)
\end{aligned}
$$



Figure 3: Updating the lateral and longitudinal BoEs of object $O_{i} . \mathcal{E}^{T R}=\left(\mathcal{E}^{T R 1} \bowtie \mathcal{E}^{T R 2}\right), \mathcal{E}^{\prime \prime}=\left(\mathcal{E}^{\prime} \triangleleft \mathcal{E}^{T R}\right)$.
for all $B \subseteq \Theta$. The parameters $\left\{K^{T R 1}, K^{T R 2}, \beta(\cdot)\right\}$ are non-negative real and satisfy

$$
K^{T R 1} \sum_{A \in \mathfrak{F}^{T R 1}} \beta^{T R 1}(A)+K^{T R 2} \sum_{A \in \mathfrak{F}^{T R 2}} \beta^{T R 2}(A)=1
$$

We denote this operation as $\mathcal{E}^{T R}=\left(\mathcal{E}^{T R 1} \bowtie \mathcal{E}^{T R 2}\right)$.
The CUE is used to update a given BoE when it confronts a new incoming BoE (see Experiment 1 below). Fig. 2, illustrates the CUE update operation, $\mathcal{E}^{\prime \prime}=\left(\mathcal{E}^{\prime} \triangleleft \mathcal{E}^{T R}\right)$. We determine the CUE parameters $\{\alpha, \beta(\cdot)\}$ via the following strategy: (a) Selecting $\alpha$ : We use a constant value for $\alpha$ so that the incoming evidence is given the same weight throughout the update process. The value of this constant $\alpha$ is a function of the fidelity of the input (e.g., the resolution and frame rate of the video input) and the particular domain of application. (b) Selecting $\beta(\cdot)$ : These are selected according to the receptive update strategy (Premaratne et al. 2009), i.e., $\beta(A)=m^{T R}(A), \forall A \in \mathfrak{F}^{T R}$. In essence, receptive updating 'weighs' the incoming evidence according to the support each focal element receives from $\mathcal{E}^{T R}$. In the special case of probability mass functions, this receptive strategy reduces to a weighted average of probability mass functions corresponding to $\mathcal{E}^{\prime}$ and $\mathcal{E}^{T R}$.

The CFE is used to generate a new BoE by fusing two (or more) incoming BoEs (see Experiment 2 below). Fig. 3, illustrates the CFE fusion operation, $\mathcal{E}^{T R}=$ $\left(\mathcal{E}^{T R 1} \bowtie \mathcal{E}^{T R 2}\right)$. We determine the CFE parameters $\left\{K^{T R 1}, K^{T R 2}, \beta(\cdot)\right\}$ via the following strategy: (a) Selecting $K^{T R 1}$ and $K^{T R 2}$ : We use constant values for $K^{T R 1}$ and $K^{T R 2}$ so that the fusing evidences are given the same weight throughout the fusion process. The values of these constants $K^{T R 1}$ and $K^{T R 2}$ are a function of the properties of the sensor inputs. (b) Selecting $\beta(\cdot)$ : These are also selected according to the receptive update strategy, i.e., $\beta(A)=m^{T R 1}(A), \forall A \in \mathfrak{F}^{T R 1}$. The receptive updating 'weighs' the incoming evidence according to the support each focal element receives from the BoEs $\mathcal{E}^{T R 1}$ and $\mathcal{E}^{T R 2}$.

## Experiments

We conducted experiments of these models of evidence and updating on two AV navigation tasks: pedestrian navigation
and vehicle lane change. These experiments were carried out with the help of the following MATLAB toolboxes: Automated Driving, Computer Vision, and Deep Learning (https://www.mathworks.com/products/automateddriving.html, computer-vision.html, deep-learning.html).

## Experimental Setup.

Video Input. We used a 30 frames/s (FPS) video input with $(640 \times 480)$-pixel $(480$ p) resolution. The video input was obtained using a Canon EOS Rebel T3i Digital Single Lens Reflex (DSLR) camera with a EF-S 18-55mm f/3.5-5.6 IS STM standard zooming lens. Following are the other camera intrinsic parameters: Focal length: [2.3814e $+03,1.8671 \mathrm{e}+03$ ], Principal point: [283.7174, 282.2949], and Radial distortion: [7.1048, -95.8961]. Video inputs were collected in urban environments and highways in Florida during the day time.

Object Identification and Detection. We used the following pre-trained detectors: (a) For detection of vehicles, vehicleDetectorACF ('full-view') and vehicleDetectorFasterRCNN('full-view'). from the Automated Driving Toolbox and the Deep Learning Toolbox. These 'full-view' models have been trained with images that contain unoccluded views from the front, rear, left, and right sides of vehicles. (b) For detection of pedestrians and cyclists,
peopleDetectorACF ('caltech')
from the Computer Vision Toolbox. It has been trained using the Caltech Pedestrian data set which contains approximately 10 hours of $640 \times 48030 \mathrm{~Hz}$ video taken from a vehicle driving through regular traffic in an urban environment and includes 350,000 pedestrian bounding boxes labeled in 250,000 frames (Dollar et al. 2012).

Object Tracking. We implemented a multiple object tracker using the Computer Vision Toolbox. It keeps an array of tracks, where each track is a structure representing a pedestrian or a vehicle in the video. The purpose of the structure is to maintain the state of a tracked moving object. The tracker estimates the state vector and state vector covariance matrix for each track using a Kalman filter. These state vectors are used to predict a track's location in new frames and to calculate the likelihood of each detection being assigned to each track.

Modeling and Updating of Evidence. Both DST and probability models were employed to capture evidence. We updated the belief, plausibility, and probability values for each object as each frame was received.

DST Model. The FoDs associated with the lateral and longitudinal movements are $\Theta_{U}$ and $\Theta_{V}$, respectively. With the initial BoEs taken as vacuous, each BoE was updated using the CUE, for which after calibration and a series of experiments, we used the parameter value $\alpha=0.66$; the receptive update strategy was employed to select the $\beta$ values.

The fusion operations were carried out with the CFE, for which we used the parameter values, $K^{T R 1}=0.5$ and $K^{T R 2}=0.5$; the receptive update strategy was utilized to select the $\beta$ values. Data structures, algorithms, DS-Conditional-One model and DS-Conditional-All model present in (Polpitiya et al. 2016; 2017; Polpitiya 2019) were utilized for DST implementations.

Probability Model. The sample spaces used were the


Figure 4: One frame of a video stream corresponding to a pedestrian crossing the road from left to right.
same, viz., $\Theta_{U}$ and $\Theta_{V}$. Initially, having no evidence to use as a basis, we assumed that the probabilities are uniformly distributed, i.e., a probability of $1.0 / 5=0.2$ was used for each singleton. Probability updating was carried out in the following manner: Consider the object $O_{i}^{\prime}$ in the frame and its corresponding object $O_{i}^{\prime \prime}$ in the current frame. Suppose the lateral and longitudinal probability distributions associated with $O_{i}^{\prime}$ and $O_{i}^{\prime \prime}$ are $P_{U}^{\prime}(\cdot), P_{V}^{\prime}(\cdot)$ and $P_{U}^{\prime \prime}(\cdot), P_{V}^{\prime \prime}(\cdot)$, respectively. We get the lateral and longitudinal new probability distributions, which we receive as new evidence upon comparing the relative positions of $O_{i}^{\prime}$ and $O_{i}^{\prime \prime}$. Suppose these new probability distributions are $P_{U}^{T R}(\cdot)$ and $P_{V}^{T R}(\cdot)$. Using an update mechanism analogous to CUE, we generate the masses of $P_{U}^{T R}(\cdot)$ and $P_{V}^{T R}(\cdot)$ via

$$
\begin{equation*}
P^{\prime \prime}(\cdot)=\alpha P^{\prime}(\cdot)+(1-\alpha) P^{T R}(\cdot), \text { with } \alpha=0.66 \tag{3}
\end{equation*}
$$

We analyzed several scenarios under different driving and road conditions. Results of the scenarios being presented are the following: (a) Experiment 1: Pedestrian road crossing; (c) Experiment 2: Pedestrian road crossing with an additional fictitious sensor input incorporated to the video input.

## Experiment 1

Fig. 4 shows one frame of the analyzed video of a pedestrian, starting from the middle of the road, crossing the road from left to right. The vehicle, whose passenger was videoing, was turning left inside a roundabout.

Fig. 5 shows the quantities $B l(), P l()$, and $P()$ versus the frame number. Consider the 4 segments (1)-(4).
(1) $B l()$ values start from $0, P l()$ values start from 1, and $P()$ values are in between. With incoming evidence, the uncertainty intervals $U n() \equiv P l()-B l()$ are narrowed. Note how the DST model allows us to start with total ignorance (i.e., $U n()=1$ ) instead of a point probability.
(2) Evidence received being of high confidence, $B l()$, $P l()$, and $P()$ values converge. The pedestrian motion prediction is SL (which is the ground-truth).
(3) Here we receive evidence with low confidence, and while $B l(S R)$ and $P(S R)$ values stay equal, the $P l(S R)$


Figure 5: $B l(), P l()$ and $P()$ values when a pedestrian is crossing the road from left to right.
value is at 1 . The ground-truth is contained within $U n(S R)$. The $P(S R)$ value gets lower and moves away from the ground-truth; the other $P()$ values are increasing.
(4) At the end of the object detection, $U n()$ increases while $B l(S R)$ and $P(S R)$ values stay equal. Eventually $P(S R)$ reaches a near zero value which contradicts the ground-truth (the pedestrian is moving to the right and out of the frame). While the other $P()$ values are also increasing, $B l()$ and $P l()$ values correctly represent the uncertainty.

## Experiment 2

This experiment uses the evidence from the same video stream utilized above (denoted by the BoE $\mathcal{E}^{T R 1}$ ) together with the evidence from an extra fictitious sensor input (denoted by the $\operatorname{BoE} \mathcal{E}^{T R 2}$ ). The CFE is first applied to fuse $\mathcal{E}^{T R 1}$ and $\mathcal{E}^{T R 2}$; then the CUE is applied as before the update the evidence (see Fig. 3).

The BoEs for the two sensor streams are generated as follows: (a) $\mathcal{E}^{T R 1}$ : Here we apply Algorithm 1. The constant value $\Pi$ classifies the lateral focal element as either $F L$ or $S L$ (or either $F R$ or $S R$ ) and it depends on the video input configuration. The constant $\Gamma$ classifies the longitudinal focal element. (b) $\mathcal{E}^{T R 2}$ : Here we employ an algorithm similar to Algorithm 1. The constant $\Pi$ is used to discriminate the lateral focal element as either Slow or Fast (not as FL, SL, $C, F R, S R$ ). Focal elements are not discriminated as Right or Left when using the fictitious sensor. This ensures that the BoEs generated are not necessarily consonant.

Fig. 6 shows the quantities $B l()$ and $P l()$ versus the frame number. Results of Experiment 1 are also provided in the figure for comparison purposes. Note that the results are obtained in real-time. Each fusion and updating operation is completed within 33 (ms).


Figure 6: $B l()$ and $P l()$ values when a pedestrian is crossing the road from left to right. Simulation results with two sensor inputs.

## Concluding Remarks

The absence of effective mechanisms to handle uncertain and exception conditions is perhaps the most serious hindrance toward the development of fully autonomous vehicles (Ottley 2019). The current approach of using deep learning methods to improve decision-making must rely on 'experience' captured via training data. However, most traffic accidents are due to exceptions (e.g., bad weather, traffic violations, sudden movement of other moving objects and pedestrians, mechanical breakdowns, etc.). Achieving full autonomy calls for appropriate mechanisms for handling data/evidence uncertainty and exceptions. Currently, exception handling relies on the human driver, and it is not practical to assume that a driver is always receptive to take over when an AV issues a request to intervene. Irrespective of whether it is the AV itself or a human driver who has to take control of the situation, the decision-maker will be better served if a fully informed decision which accounts for the underlying uncertainty can be made.

We explore the utility of DST models as a way to capture the uncertainty in intended relative movements of other vehicles and pedestrians. The calibrated information that such a model generates may allow an AV (or a human-in-the-loop) to make more informed decisions. To ensure real-time operation, we have made use of the data structures and algorithms in (Polpitiya et al. 2016; 2017; Polpitiya 2019).

We believe that the proposed model is a significant step toward addressing the current challenges involving levels 4 5 autonomous driving tasks. DST models are ideal to capture and propagate the uncertainties throughout the decisionmaking pipeline and provide an interval-valued measure of the uncertainty underlying any inference it makes.

## Acknowledgments

This work is based on research supported by the U.S. Office of Naval Research (ONR) via grant \#N00014-10-10140, U.S. National Science Foundation (NSF) via grant \#1343430, and the University of Miami U-LINK Initiative on Interdisciplinary Inquiry.

## References

Burns, C. M. 2018. Automation and the human factors race to catch up. Journal of Cognitive Engineering and Decision Making 12(1):83-85.
DARPA. 2014. The DARPA Grand Challenge: Ten Years Later. https://www.darpa.mil/news-events/2014-03-13.
Dempster, A. P. 1968. A generalization of Bayesian inference. Journal of the Royal Statistical Society. Series B 30(2):205-247.
Dollar, P.; Wojek, C.; Schiele, B.; and Perona, P. 2012. Pedestrian detection: An evaluation of the state of the art. IEEE Transactions on Pattern Analysis and Machine Intelligence 34(4):743-761.
Kaber, D. B. 2018. Issues in human-automation interaction modeling: Presumptive aspects of frameworks of types and levels of automation. Journal of Cognitive Engineering and Decision Making 12(1):7-24.
Ottley, S. 2019. Nissan: Autonomous cars impossible without humans. https://www.drive.com.au/news/nissan-autonomous-cars-impossible-without-humans-120948.
Polpitiya, L. G.; Premaratne, K.; Murthi, M. N.; and Sarkar, D. 2016. A framework for efficient computation of belief theoretic operations. In Proceedings of the Nineteenth International Conference on Information Fusion (FUSION), 1570-1577.
Polpitiya, L. G.; Premaratne, K.; Murthi, M. N.; and Sarkar, D. 2017. Efficient computation of belief theoretic conditionals. In Proceedings of the Tenth International Symposium on Imprecise Probability: Theories and Applications (ISIPTA), 265-276.
Polpitiya, L. G. 2019. A Framework for Efficient Implementation and Effective Visualization of Dempster-Shafer Belief Theoretic Computations for Reasoning Under Uncertainty. Ph.D. Dissertation, University of Miami, Coral Gables, FL.
Premaratne, K.; Murthi, M.; Zhang, J.; Scheutz, M.; and Bauer, P. 2009. A Dempster-Shafer theoretic conditional approach to evidence updating for fusion of hard and soft data. In Proceedings of the Twelfth International Conference on Information Fusion (FUSION), 2122-2129.
SAE. 2018. Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles. Technical report, SAE International.
Shafer, G. 1976. A Mathematical Theory of Evidence. Princeton, NJ: Princeton Univ. Press.
Wickramarathne, T. L.; Premaratne, K.; and Murthi, M. N. 2012. Consensus-Based credibility estimation of soft evidence for robust data fusion. In Proceedings of the Second International Conference on Belief Functions (BELIEF), 301-309.


[^0]:    Copyright (C) 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

