2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
Congress Center Hamburg
Sept 28 - Oct 2, 2015. Hamburg, Germany

Financialized methods for market-based multi-sensor fusion

Jacob Abernethy and Matthew Johnson-Roberson

***Abstract*— Autonomous systems rely on an increasing number of input sensors of various modalities, and the problem of sensor fusion has received attention for many years. Autonomous system architectures are becoming more complex with time, and the number and placement of sensors will be modified regularly, sensors will fail for many reasons, information will arrive asynchronously, and the system will need to adjust to rapidly changing environments. To address these issues we propose a new paradigm for fusing information from multiple sources that draws from the rich of field pertaining to *financial markets*, particularly recent research on prediction market design. Among the many benefits of this financialized approach is that, both in theory and in practice, markets are well-equipped to robustly synthesize information from diverse sources in a decentralized fashion. Our framework poses sensor processing algorithms as profit-seeking market participants, data is incorporated via financial transactions, and the joint estimation is represented as a price equilibrium. We use pedestrian detection as a motivating application. Pedestrian detection is a well studied field and essential to autonomous driving. Real world fusion results are presented on RGB and LIDAR data from the KITTI Vision Benchmark Suite. We demonstrate we can achieve comparable performance to state-of-the-art hand designed fusion techniques using the proposed approach.**

I. INTRODUCTION

Alongside the burst in development of autonomous systems such as self-driving vehicles, intelligent sensor networks, and personal robots we observe increasing demand for new methods to obtain and synthesize data from an agent's environment. Modern platforms are equipped with a wide range of *sensors* that gather and report information on any number of aspects of the system's surroundings. These can include *camera sensors* for collecting image-based data, acoustic sensors like *sonar*, *LIDAR* for active range measurement and radar for penetrative electromagnetic sensing. Typically each of these sensors report data which will at various times be critical to the autonomous system, and often the system will heavily rely on synthesizing multiple sensor streams simultaneously.

The problem of sensor fusion has received attention for many years with now a sizable literature [1, 2, 3, 4]. Much of this work has utilized probabilistic modeling tools to aggregate data sources; Bayesian methods such as Kalman filtering [5] have been particularly popular. Such techniques have found broad utilization in modern sensor fusion systems in that they have the ability to handle various levels of noise across the array of sensors.

Probabilistic aggregation mechanisms require a huge amount of human-led design and fine tuning, and this is both a major benefit in terms of the capabilities of the system yet also a significant weakness when it comes to maintaining global flexibility and robustness. Carefully designed models can excel under the right conditions, but they are typically fragile to even small changes in system design, and they need full retraining after only slight modification to the input types.

To address these open problems in sensor fusion this paper focuses on the development of a data synthesis platform that will rely on tools from the design of *financial markets* to perform efficient aggregation of the various data streams. The approach we advocate shall exploit the obvious benefits of prediction market mechanisms to efficiently combine the information from a range of different sources for the sake of prediction and estimation.

Currently, the field has been dominated by Bayesian methods. De Finetti's classical perspective on probability theory, that probabilities should only be interpreted as prices offered for betting contracts [6], is particularly salient. When we view a system with a diverse array of sensor inputs, where each generates a stream of probabilities/prices based on the available data, then a very natural way to aggregate such prices is through a market. Indeed, it has already been established that several existing Machine Learning techniques can be endowed with financial semantics [7], where various estimation methods correspond to price equilibria under particular agent/utility models.

While the market-oriented view of data fusion shares many similarities with existing Bayesian methods, it has many unique benefits that are quite well suited to the needs of autonomous systems:

- Economic structures are inherently founded on the principle of decentralization and diversely-held information. Whereas statistical methods can often be implemented with a decentralized architecture, they are not fundamentally designed as such.
- The success of market mechanisms in practice is that they are incredibly robust without the need

J. Abernethy is with the Department of Electrical Engineering & Computer Science, University of Michigan, Ann Arbor, MI 48109, USA jabernet@umich.edu.

M. Johnson-Roberson are with the Department of Naval Architecture & Marine Engineering, University of Michigan, Ann Arbor, MI 48109, USA mattjr@umich.edu.

for strong restrictions or assumptions on the market participants. Electoral prediction markets, for example, exhibit impressive predictive power even when most individual traders are uninformed.

- Typical market frameworks are designed with asynchrony in mind, hence prices will respond to new information when it becomes available. This is crucial when we realize that sensors differ in the frequency of measurement as well as lags due to computational and other time costs.
- Markets tend to have the long-term benefit of separating good information from bad; market participants with truly novel and useful information are given rewards in order that their future "bets" will generate a large effect on the market prediction.

II. Prior Work

The field of data fusion has been the source of extensive research and the broad coverage of topics is beyond the scope of this paper. Bar-Shalom et. al present a comprehensive review of recent techniques in [8, 9]. Most relevant to the proposed market approach is the sub field of Decentralized Data Fusion (DDF), including the work of Sukkarieh et. al where DDF was applied to tracking ground targets from four Unmanned Aerial Vehicles [10]. There we see the use of individual robots as independent agents exhibiting a decentralized solution. Additionally, the Decentralized Bayesian approach of Makarenko and Durrant-Whyte [4] addressing the DDF problem with a formal probabilistic framework allowing for the Bayesian integration of multiple decentralized sensors this time within a sensor network context. In a generalization of the Bayesian approach Dempster-Shafer evidence theory has formed the basis for several approaches that attempt to fuse sensor data in a principled manner [11, 12]. Such approaches have utilized the concept of distributed support for a proposition. Newer work by Lanillos et al. [13] has moved beyond strictly hand designed fusion architectures for the DDF problem but must share larger amount of data between agents to perform fusion.

Financialized methods within robotics have appeared in coordination and task allocation in multi-agent systems. Early work by Wellman and Wurman [14] showed the promise of market aware agents in a robotics context. They explore the concept of using market forces to govern intra-team robot coordination in non-competitive environments. This was validated in simulation and later applied to physically embodied robotic systems by Zlot and Stenz [15, 16]. A greater degree of formalization to market-based task allocation was provided by Liu and Shell who expanded previous approaches with bounds on complexity [17]. Such market-based techniques were very well suited to task allocation. Job bidding and allocation has very natural parallels to typical economic principles in market creation.

This paper addresses the problem of integrating multiple sensors using information markets a topic that has received less attention as it requires a higher degree of abstraction to make parallels to the real world economy. It is important to note that this paper is not the first to put forth a market-like framework for sensor data aggregation; Jumadinova and Dasgupta [18, 19, 20] have done some preliminary work looking at how a prediction market can be used to fuse a metal detector, a ground penetrating radar and an infra-red camera for a land mind detection system. Their work, while interesting and enlightening, has only scratched the surface on avenues of exploration, and has not taken advantage of the wealth of research that has emerged in the last 5-10 years within the EconCS community for developing prediction market mechanisms.

III. Market-based Sensor Fusion

The data fusion problem is critical to real-time field robotic systems. The ability to robustly fuse data from multiple sources has been a major driver in the advances in autonomous systems over the last twenty years [21]. The benefits of data fusion apply to many core functions robots undertake frequently. Sensor fusion offers the possibility of improved range and resolution, measurement across a diverse set of sensor modalities, mitigation of imprecision and uncertainty, and improved resolution.

A. The Fusion Market System Architecture

A diagram of the architecture of the system is shown in Figure 1. Let us now describe the ingredients of the proposed Sensor Fusion Market Mechanism, and how each piece will interact in the following section:

1) **The Sensors.** Each sensor will produce a stream of raw data that can be observed by one or several agent models. The sensor data need not be reliable at all times, and the sensors may have different processing lags etc.
2) **A Class of Predictive Models.** Agents in our system will synthesize sensor data according to a predictive model that has been trained using historical data. A model receives a stream of data and generates a sequence of predictions. An initial prototype for the proposed system is discussed in §V.
3) **An abstract currency.** The way in which algorithms will interact with each other is through the use of a kind of "pseudo-currency" that can be used to purchase predictive contracts. Contract rewards will pay out in this currency, and the agents in our system will aim to maximize their long-term profit accordingly.

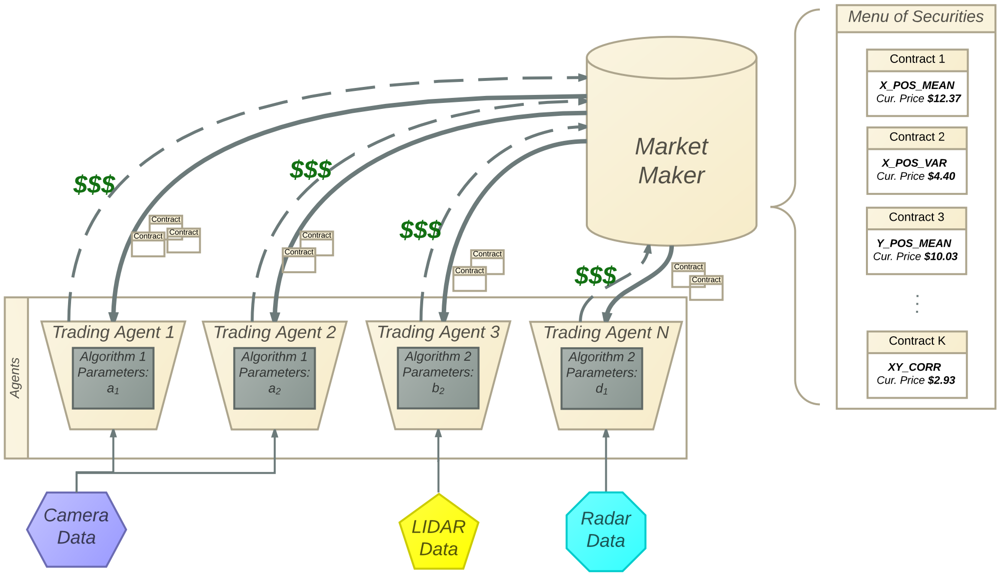

Fig. 1. An overview of the multi-sensor fusion market.

4) **A Set of Contingent Securities (Contracts).** The fusion market will facilitate the trade of abstract prediction contracts characterized as follows:
   a) A *verifiable event or quantity* corresponding to a real or discrete variable associated with a future uncertain state of the environment, such as "there will be an object at position X,Y at time T" and "there is a human standing within a 20 foot range of the vehicle";
   b) A *payout function* that determines the reward, in the market's pseudo-currency, paid to the purchaser of this contract under the range of possible future outcome states;
   c) An *expiration time* specifying a time by which the contract will settle;
   d) A *price*, set according to the market, for investing in this security.
5) **The Agents.** Each agent is characterized by the following attributes:
   a) *A sensor input* that is the unique source of data input;
   b) *a predictive model* used to internally generate predictions;
   c) *available capital* that was obtained via trading, and can be used to make future speculative investments;
   d) a set of *invested securities*, contracts that have been purchased (or sold) that will determine future possible payouts to the agent.

   At a high level, each agent is endowed with the objective to maximize profit over the long term by trading in the fusion market utilizing the (potentially private) information available to it via its specific sensor. At a given point in time the agent will observe sensor data and generate one or several predictions, each corresponding to a prediction contract made available by the market making authority. Thus the agent must query the menu of available securities sold by the market, as well as the posted price of each, and the agent will determine which contracts provide a suitable profit expectation. Of course, the agent must also reason about its available capital and existing investments (financial exposure) to decide what investments to make or whether to abstain from trading.
6) **The Market Authority.** Sitting at the center of the set of agents is the market authority, which we shall call a *Market Maker* (MM). The MM advertises the set of currently-available securities for purchase, and the MM exposes a trading API available to the agents for making purchases (or sales) of such contracts, including both a price-query method as well as a purchase/sale method. The MM's primary job is to adjust prices according to executed trade's, as these prices shall reflect the aggregate estimate of the agents.

IV. Market-driven Sensor Fusion Framework

A. Prediction Markets

Information or *prediction markets*, which play a fundamental role in economics and finance, are used as a central tool in the present work. Prediction markets allow individuals to bet on the outcome of future events, either for the sake of speculation (World Cup gambling) or to hedge risk (farmers buying commodity futures). Many information markets already exist and are available to the public; at sites like `Betfair.com`, individuals can bet on everything ranging from election outcomes to geopolitical events. There has been a recent burst of interest in such markets, not least of which is due to their potential for combining disparate information from many sources. In the words of Hanson et al. [22]: "Rational expectations theory predicts that, in equilibrium, asset prices will reflect all of the information held by market participants. This theorized information aggregation property of prices has lead economists to become increasingly interested in using securities markets to predict future events." In practice, prediction markets have proven impressively accurate as a forecasting tool [23, 24].

Our goal in the present section is to give a broad overview of the mathematical framework for designing prediction market mechanism. In particular we want to describe some of existing literature which is the basis for the design of the centralized market-making authority described in §V.

1) Proper Scoring Rules, to the Design of Prediction Markets: In its simplest incarnation, a proper scoring rule [25] is a two input function $S(\hat{P};x)$, where $\hat{P}$ is a forecast x an "outcome", that must satisfy the following: for distribution P, if x is sampled from P, then for any $\hat{P}$ we must have $\mathbb{E}_{x\sim P}[S(P,x)] \geq \mathbb{E}_{x\sim P}[S(\hat{P},x)]$. That is, the *expected* value of S is maximized by reporting the true distribution P. A popular example is the *logarithmic scoring rule* defined by $S(P,x) := \log P(x)$, and it is an easy exercise to see that this satisfies the desired property. The literature has grown significantly [26, 27] since the work of Savage [25]; see [28] for useful survey.

When x arrives from a large (or infinite) space, we may not need the full distribution P over x but rather some *statistic* of P, say $\mathbb{E}_P[\phi(x)]$, where ϕ is some function of interest taking values in $\mathbb{R}^d$. A proper scoring rule for ϕ is a function $S : \mathbb{R}^d \times \mathcal{X} \to \mathbb{R}$ satisfying the natural property: $\mathbb{E}_{x\sim P}[S(\mu,x)] \geq \mathbb{E}_{x\sim P}[S(\hat{\mu},x)]$, where $\mu = \mathbb{E}_{x\sim P}[\phi(x)]$ and for any $\hat{\mu}$. It turns out that, given any smooth convex function R on $\mathbb{R}^d$, we can construct a proper scoring rule via the Bregman divergence as follows: $S(\hat{\mu},x) = D_R(\phi(x),\hat{\mu})$. Abernethy and Frongillo [29] established that *every* proper scoring rule can be cast in this form, up to additive terms.

Hanson [30] developed the beautiful insight that one can use a scoring rule not only to elicit correct forecasts from a single individual but also to design a *prediction market*. In such a market, traders would have the ability to place bets with a central authority, known as a *market maker*[1](MM). The betting framework proceeds as follows: The market maker publishes a proper scoring rule S and an initial probability estimate P_0. On each round t the current consensus probability P_t is posted and any trader can place a bet by *modifying* the probability to a desired value P_{t+1}. In the end, the true outcome x is publicly revealed, and each trader receives a (potentially negative) profit of $S(P_{t+1},x) - S(P_t,x)$.

Notice two facts about this framework: (a) if a trader at time t knows the true probability P^* then he always maximizes expected profit by setting $P_{t+1} = P^*$ and (b) because of the telescoping sum, if P_T is the final estimated probability then the market maker needs only to pay out a total of $S(P_0,x) - S(P_T,x)$. Hanson referred to this form of prediction market as a *Market Scoring Rule* (MSR), known as the Logarithmic Market Scoring Rule (LMSR) under the log score.

Hanson's prediction market framework, which requires traders to make probability estimates and judges them according to a scoring rule, does not fit into our typical understanding of betting markets, as well as other financial markets, in which parties buy and sell "shares." We can think of a share as an *Arrow-Debreu security* which would involve a payoff of \$1 in the even that a particular state of the world is reached, and \$0 otherwise [31]; these are often called *contingent securities*. A very nice observation, initially discoverd by [32], shows that the market scoring rule idea of Hanson can be converted into an alternative betting language, where traders simply purchase bundles of Arrow-Debreu securities at prices set by the market maker. The corresponding mechanism involves a market formulation based on a "cost function", which we sketch here:

- Before outcome $i \in [n]$ is realized, MM shall sell Arrow-Debreu securities (shares) for all i. MM has a smooth convex $C : \mathbb{R}^n \to \mathbb{R}$ and maintains a "quantity vector" $\mathbf{q} \in \mathbb{R}^n$, initialized to $\mathbf{0}$.
- Traders may purchase share "bundles" $\mathbf{r} \in \mathbb{R}^n_{\geq 0}$ (r_i is quantity of security i). Given current $\mathbf{q}$, the price for $\mathbf{r}$ is $C(\mathbf{q}+\mathbf{r})-C(\mathbf{q})$. After selling $\mathbf{r}$, MM updates $\mathbf{q} \leftarrow \mathbf{q}+\mathbf{r}$.
- At the close of the market, when the outcome i is revealed, the market maker has to make a payout to all winning contracts, which is a total cost of q_i.

The derivative $\nabla C(\mathbf{q})$ is essentially the market estimate of the true distribution on the outcome, since $\nabla_i C(\mathbf{q})$ is the marginal cost of a tiny purchase of contract i, which in equilibrium should be the expected return of the contract; that is, the probability. Indeed, to avoid arbitrage

[1]The term market maker is used generally to describe an agent or firm that is willing to facilitate the market by offering to transact with interested traders. Here, the market maker is also the central authority that manages the market as well.

opportunities the market maker must ensure that $\nabla C(\mathbf{q})$ is always a distribution. Chen and Pennock [32] showed that the LMSR is implementable via this cost-function framework, with $C(\mathbf{q}) := \alpha^{-1} \log\left(\sum_i \exp(\alpha q_i)\right)$.

B. Sensor model and processing

As this paper is not focused on the direct creation of new classification techniques but rather a framework for the effective synthesis of multiple techniques, we draw from the state-of-the-art in object classification literature for autonomous driving applications. Further constraining the exploration of this problem we have chosen a well researched area in this field: pedestrian detection [33, 34]. It was desirable to choose a very important yet constrained problem to deeply explore the fusion of multiple algorithms sensors' hypotheses about the semantic label and physical position of a target like a pedestrian.

In the proposed framework a classification algorithm is wrapped in a trading agent. While the approach is agnostic to the underlying classification algorithm here we describe the state-of-the-art approach used in this paper. Using the multisensor approach of Premebida et al. [35] LIDAR and image data is processed to extract high-level features for classification. The image data is run through a histograms of gradients (HOG) feature extractor [36]. This produces a fairly illumination invariant representation that has been quite popular for object detection. HOG splits an image into a set of patches on a regular grid and then creates a binning of gradient orientations within those patches. These patches are then aggregated in coarser overlapping grids and normalized. These patches are then concatenated into a single feature vector representing the contained region. Given an object of interest (e.g. a human pedestrian) a template is learned discriminatively using a linear support vector machine [37]. This template can then be run against an image to detect occurrences of that object.

To address issues around occlusion and improve the detection of articulated object like pedestrians, Felzenszwalb et al. [38] proposed a deformable parts model. This addresses shortcomings in detecting non-rigid objects with HOG. The algorithm does this by clustering positive examples to generalize template learning for an object. It creates a star model to relate object parts and learn filters for each part individually. Finally, the algorithm uses multiple instance learning to refine bounding box detection.

To complement the image based approach Premebida et al. [35] propose the use of HOG on up-sampled depth maps created using a Velodyne HDL-64E LIDAR. We employ the Premebida et al. [35] technique to perform some preprocessing to allow HOG to be run on the depth image and produce similar results to the complementary RGB image data.

V. EXPERIMENTAL SETUP AND MARKET IMPLEMENTATION

Our experimental setup and implementation of the prediction market mechanisms is as follows:

1) *Image Data.* The KITTI dataset [39] images were all resized to the same dimensions. For each image, and for the RGB sensor as well as the LIDAR, we have a set of bounding boxes for predicted pedestrian locations within the image. You can see an example image in Figure 2.
2) *Sliding Windows.* We divided each image into a set of regions where we generated a binary label corresponding to "is there a pedestrian in this window?". This is computed by calculating if at least 50% of a pedestrian (predicted or ground truth) bounding box is contained within the window. These sliding windows span the entire height of the image and are shifted 50% of their width creating 50% overlap between each window.
3) *Trading Agents.* For each sensor type s, and for each set of parameter values θ used to interpret the sensor's readings, we have a trading agent $T_{s,\theta}$ who will execute trades based on the inputs from this sensor and the parameters θ. For this particular experimental setup, the only parameter we considered was the sensitivity to the size of overlap of the bounding box of a predicted pedestrian in the window, and considered 10 parameters uniformly from 0% to 100%. Each agent was also afforded an amount of initial capital to make trades, and an agent's existing capital affects the size of the bets s/he can make.
4) *Binary contracts.* For each image i and window j the traders can purchase a "contract" $C_{i,j}$ that give a payout of $1 in the event that a pedestrian is detected in this window in the image. The contract prices are set by market demand and are managed by the market maker. The payouts are made after the true labels are revealed, which in our setup occurs directly after the market closes and before the traders proceed to the next image.
5) *Market Maker.* For each image i there is a market maker (MM) who facilitates trading among the set of traders, and the MM offers $C_{i,j}$ for each window j. In order to price $C_{i,j}$, the MM also maintains the total outstanding shares of $C_{i,j}$. The MM sets prices according to the cost function described in §IV-A.1: if the market maker has sold q shares on a particular contract, and a trader wants to purchase r additional shares, the cost is $\frac{1}{\alpha}\log\left(\frac{1+\exp(\alpha(q+r))}{1+\exp(\alpha q)}\right)$; we set $\alpha = 1.0$. This is the LMSR described previously.

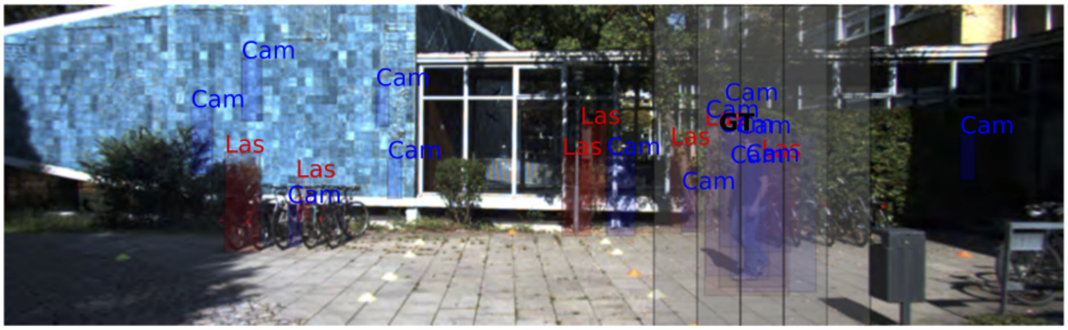

Fig. 2. An example image in the KITTI dataset [39], with predicted detections by LIDAR ("Las") and camera ("Cam") relative to ground truth ("GT"). Active windows are shown as well.

6) *Trading Process.* For each image, corresponding to a single market, we facilitated trading among the set of trading agents as follows. For each of a set of rounds (numbering 1000), we sample an agent at random, compute the current market prices for the set of available contracts, pass these prices to the agent, and the agent determines a bet size for each contract. If the agent's internal belief is p on this contract, corresponding to the probability of a pedestrian in image i window j, and the market price is $\hat{p}$, then the agent's purchase size is proportional to $\max(0, p - \hat{p})$, where the amount also scales with the agent's remaining budget as well as a betting-fraction parameter which we set to 0.1.
7) *Market predictions.* At the end of the trading process, when the market prices have settled, we can query the market for the various prices on the set of available securities. As we have already mentioned, these market prices correspond in some sense to the aggregate belief of all traders. We use these prices as "likelihood scores" when comparing to the other methods (in §VI).

VI. Results

We performed a set of experimental results aimed at demonstrating the feasibility of the proposed prediction market framework. We implemented a market infrastructure to aggregate sensor data on a particular detection task involving recognizing pedestrians. Our results show that when we use the market prices as an aggregation mechanism for our sensor data we obtain predictions that are at least competitive with existing methods. Figure 3 shows that the market predictions are superior for at least some precision/recall tradeoffs.

A. Evaluation Metric

The pedestrian detection performance is evaluated by assessing the number of true positives (TP) and false positives (FP) using VOC-PASCAL metric [40] of 50% overlap criteria:

$$score = \frac{area(B_u \cap B_v)}{area(B_u \cup B_v)}, \tag{1}$$

where $B_u \cap B_v$ denotes the intersection of two bounding boxes and $B_u \cup B_v$ their union. (Note: in this case we calculate the bounding boxes as clamped to sliding windows as described in §V.)

B. Training and Testing

To perform a fair evaluation comparing our results to Premebida et al. [35] we use the exact same validation set which consists of 3741 frames, of the KITTI dataset [39] starting from frame 003739 and ending at 007480. The market agents are "trained" on the first 3738 images where they make bets on the presence or absence of a pedestrian in each of the sliding windows in a image. This "training" process allocates the capital in the market to the agents (the varied modalities and parameter settings). Then once the training rounds of betting are complete the agents are shown the new data, one image at a time (images 003739–007480), and asked to bet on the presence or absence of a pedestrian in a single sub-window. The resulting market price for a security after 1000 betting rounds (see §V.6) can be expressed between $0.0 - 1.0$ and taken as a binary classifier confidence output for the current sliding window. These results are then compared to three existing methods which also output binary classifications on a $0.0 - 1.0$ scale with ground truth evaluated using the bounding box overlap method described in §VI-A. The comparison approaches are the RGB-only, LIDAR, and SVM-fused results in [35]. Each of these methods and the proposed approaches are shown in Figure 3 as Receiver operating characteristic curves.

VII. Conclusions and Future Work

This paper seeks to address one of the most significant problems in robotic perception, our ability to integrate

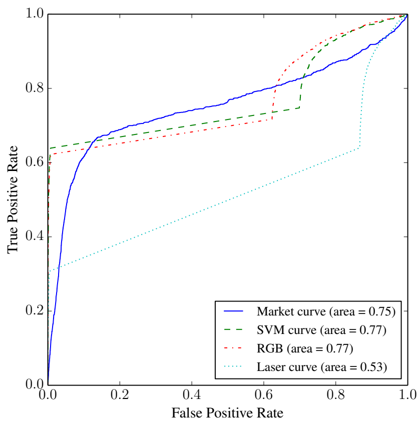

Fig. 3. Receiver operating characteristic curves for the proposed approach and a state-of-the-art SVM-based fusion technique [35]. Note that we are able to achieve similar performance without the need for an elaborately hand-tuned system. The simplicity of the market-based technique is a strong motivator for its adoption.

information from multiple sensing modalities and make decision is the face of sensor noise and failure. We demonstrate a proof-of-concept on a real-world robotics problems. We show it is possible to achieve comparable performance to state-of-the-art hand designed fusion techniques using the proposed approach.

While improvements in sensor quality matched with reduced costs have pushed new frontiers for autonomous systems. Higher-resolution cameras, LIDARs, Flash LIDARs, automotive radars, commodity structured light sensors and so on have all enabled new fields of robotic research. The complement of a new alternative fusion paradigm has great potential to enable advances in perception and classification research.

In future work we hope to broaden the set of techniques here in a number of directions. For example, our market mechanisms were designed to estimate a set of binary outcome probabilities, but a richer approach would provide a full density estimate on the existence of objects (pedestrians) in 3D space.

References

[1] H. B. Mitchell. *Multi-Sensor Data Fusion: An Introduction*. 1st. Springer Publishing Company, Incorporated, 2007.

[2] J. Manyika and H. Durrant-Whyte. *Data fusion and sensor management: a decentralized information theoretic approach*. Electrical and Electronic Engineering. New York: Ellis Harwood, 1994.

[3] H. F. Durrant-Whyte, B. Y. S. Rao, and H. Hu. "Toward a fully decentralized architecture for multi-sensor data fusion". In: *Proceedings of the IEEE International Conference on Robotics and Automation*. Cincinnati, OH, USA, 1990, pp. 1331–1336.

[4] Alexei Makarenko and Hugh Durrant-Whyte. "Decentralized Bayesian algorithms for active sensor networks". In: *Information Fusion* 7.4 (2006), pp. 418–433.

[5] R.E. Kalman. "A new approach to linear filtering and prediction problems". In: *Transactions of the ASME—Journal of Basic Engineering* 82.Series D (Mar. 1960), pp. 35–45.

[6] I. J. Good. "Review: Bruno de Finetti, Theory of probability". In: *Bull. Amer. Math. Soc.* 83.1 (Jan. 1977), pp. 94–97.

[7] Amos Storkey. "Machine learning markets". In: *arXiv preprint arXiv:1106.4509* (2011).

[8] Yaakov Bar-Shalom, Peter K Willett, and Xin Tian. "Tracking and data fusion". In: *A Handbook of Algorithms. Yaakov Bar-Shalom* (2011).

[9] Y. Bar-Shalom. *Tracking and Data Association*. San Diego, CA, USA: Academic Press Professional, Inc., 1987.

[10] Salah Sukkarieh, Eric Nettleton, Jong-Hyuk Kim, Matthew Ridley, Ali Goktogan, and Hugh Durrant-Whyte. "The ANSER Project: Data Fusion Across Multiple Uninhabited Air Vehicles". In: *The International Journal of Robotics Research* 22.7-8 (2003), pp. 505–539.

[11] Huadong Wu, Mel Siegel, Rainer Stiefelhagen, and Lei Yang. "Sensor fusion using Dempster-Shafer theory [for context-aware HCI]". In: *Instrumentation and Measurement Technology Conference, 2002. IMTC/2002. Proceedings of the 19th IEEE*. Vol. 1. IEEE. 2002, pp. 7–12.

[12] Otman Basir and Xiaohong Yuan. "Engine fault diagnosis based on multi-sensor information fusion using Dempster–Shafer evidence theory". In: *Information Fusion* 8.4 (2007), pp. 379 –386.

[13] Pablo Lanillos, Seng Keat Gan, Eva Besada-Portas, Gonzalo Pajares, and Salah Sukkarieh. "Multi-UAV target search using decentralized gradient-based negotiation with expected observation". In: *Information Sciences* 282 (2014), pp. 92 –110.

[14] Michael P. Wellman and Peter R. Wurman. "Market-aware agents for a multiagent world". In: *Robotics and Autonomous Systems* 24.3–4 (1998), pp. 115 –125.

[15] M.B. Dias, Robert Zlot, N. Kalra, and A. Stentz. "Market-Based Multirobot Coordination: A Survey and Analysis". In: *Proceedings of the IEEE* 94.7 (2006), pp. 1257–1270.

[16] Robert Michael Zlot and Anthony (Tony) Stentz. "Market-based Multirobot Coordination for Complex Tasks". In: *International Journal of Robotics*

Research, Special Issue on the 4th International Conference on Field and Service Robotics 25.1 (2006), pp. 73–101.

[17] Lantao Liu and Dylan Shell. "Optimal Market-based Multi-Robot Task Allocation via Strategic Pricing". In: *Proceedings of Robotics: Science and Systems*. Berlin, Germany, 2013.

[18] J. Jumadinova, M.T. Matache, and P. Dasgupta. "A Multi-agent Prediction Market Based on Boolean Network Evolution". In: *Web Intelligence and Intelligent Agent Technology (WI-IAT), 2011 IEEE/WIC/ACM International Conference on*. Vol. 2. 2011, pp. 171–179.

[19] Janyl Jumadinova and Prithviraj Dasgupta. "Multirobot autonomous landmine detection using distributed multisensor information aggregation". In: *SPIE Defense, Security, and Sensing*. 2012, 84070B–84070B–14.

[20] Janyl Jumadinova and Prithviraj Dasgupta. "Prediction Market-Based Information Aggregation for Multi-sensor Information Processing". English. In: *Agent-Mediated Electronic Commerce. Designing Trading Strategies and Mechanisms for Electronic Markets*. Springer, 2013, pp. 75–89.

[21] Hugh Durrant-Whyte and Thomas Henderson. "Multisensor Data Fusion". English. In: *Springer Handbook of Robotics*. Ed. by Bruno Siciliano and Oussama Khatib. Springer Berlin Heidelberg, 2008, pp. 585–610.

[22] R. Hanson, R. Oprea, and D. Porter. "Information aggregation and manipulation in an experimental market". In: *Journal of Economic Behavior & Organization* 60.4 (2006), pp. 449–459.

[23] J. Ledyard, R. Hanson, and T. Ishikida. "An Experimental Test of Combinatorial Information Markets". In: *Journal of Economic Behavior and Organization* 69 (2009), pp. 182–189.

[24] J. Wolfers and E. Zitzewitz. "Prediction Markets". In: *Journal of Economic Perspective* 18.2 (2004), pp. 107–126.

[25] Leonard J. Savage. "Elicitation of Personal Probabilities and Expectations". In: *Journal of the American Statistical Association* 66.336 (1971), pp. 783–801.

[26] J.E. Matheson and R.L. Winkler. "Scoring rules for continuous probability distributions". In: *Management Science* (1976), pp. 1087–1096.

[27] A. Dawid. "The geometry of proper scoring rules". In: *Annals of the Institute of Statistical Mathematics* 59.1 (2007), pp. 77–93.

[28] T. Gneiting and A. Raftery. "Strictly Proper Scoring Rules, Prediction, and Estimation". In: *Journal of the American Statistical Association* 102.477 (2007), pp. 359–378.

[29] Jacob D Abernethy and Rafael M Frongillo. "A Characterization of Scoring Rules for Linear Properties." In: *(COLT) Conference on Learning Theory*. 2012.

[30] R. Hanson. "Combinatorial information market design". In: *Information Systems Frontiers* 5.1 (2003), pp. 107–119.

[31] K. J. Arrow. "The role of securities in the optimal allocation of risk-bearing". In: *Review of Economic Studies* 31.2 (1964), pp. 91–96.

[32] Y. Chen and D.M. Pennock. "A utility framework for bounded-loss market makers". In: *UAI*. 2007.

[33] Pierre Sermanet, Koray Kavukcuoglu, Soumith Chintala, and Yann LeCun. "Pedestrian detection with unsupervised multi-stage feature learning". In: *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*. IEEE. 2013, pp. 3626–3633.

[34] Piotr Dollár, Christian Wojek, Bernt Schiele, and Pietro Perona. "Pedestrian detection: A benchmark". In: *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. IEEE. 2009, pp. 304–311.

[35] C. Premebida, J. Carreira, J. Batista, and U. Nunes. "Pedestrian detection combining RGB and dense LIDAR data". In: *IROS*. IEEE. 2014, pp. 4112–4117.

[36] N. Dalal and B. Triggs. "Histograms of oriented gradients for human detection". In: *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*. Vol. 1. 2005, 886–893 vol. 1.

[37] Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. "LIBLINEAR: A Library for Large Linear Classification". In: *Journal of Machine Learning Research* 9 (2008), pp. 1871–1874.

[38] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. "Object Detection with Discriminatively Trained Part Based Models". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32.9 (2010), pp. 1627–1645.

[39] Andreas Geiger, Philip Lenz, and Raquel Urtasun. "Are we ready for Autonomous Driving? The KITTI Vision Benchmark Suite". In: *Conference on Computer Vision and Pattern Recognition (CVPR)*. 2012, pp. 3354–3361.

[40] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. "The Pascal Visual Object Classes (VOC) Challenge". In: *International Journal of Computer Vision* 88.2 (2010), pp. 303–338.