Identification in Market-Based Multi-Robot Coordination

Amir Danak, and Shie Mannor
Centre for Intelligent Machines,
Department of Electrical and Computer Engineering
McGill University, Montréal, Québec, Canada

Abstract—Since agent-based systems behave more rationally than humans, game theory approaches have been used in many robotic applications. As major marketing solutions, and due to their capabilities in team information collection, auctions have received notable attention in planning and task allocation missions. This paper discusses the problem of information acquisition in market-based multi-agent coordination. Novel applications for identification, an important topic in marketing, is proposed and using partially revealed submitted bids and the concept of risk aversion, we address estimation algorithms to fully recover each agent's private attributes. The proposed learning algorithms are applicable for all teammates, since using less disclosure of auction data, internal communication cost among them becomes negligible. Therefore, each agent's belief about her colleagues' private attributes is updated using the partially revealed information by the group leader.

Index Terms—Agent-Based Systems, Intelligent Systems, Multi-Robot Coordination.

1 Introduction

Teamwork among robots to accomplish certain group tasks such as distributed sensing [1], planetary exploration [2], search and rescue missions [3], [4], observation and tracking of multiple moving targets [5], etc, has been studied by researchers (in several aspects) in the recent years. How to coordinate members to perform each service efficiently as a team constitutes a significant proportion of the aforementioned analyses. Considering the long history of negotiation among humans in sophisticated trading, market-based coordination algorithms have recently been applied to agent-based systems [6]. Based on their straightforward structure and minimal computational requirements, these methods were implemented successfully in a variety of central and distributed task allocation problems.

As a fundamental concise marketing solution, auctions have become a leading approach in multi-robot planning operations. In market-based approaches, decision-making is distributed through auctions. Generally, tasks and resources are modeled as valuable commodities for which agents compete by bidding strategically. Although both first-price and second-price auctions have been contemplated in previous studies, assuming that all agents report their true valuations of each task results in no strategic difference between the two aforementioned mechanisms.

However, in first-price auctions in the marketplace, players shade their bids in order to avoid the so-called winner's curse. In the light of this fact, identification (the process of estimating private attributes of rational agents based on their submitted bids) is demonstrated in this paper. We extend the primary results on the subject to the generic case including more intelligent robots who take their future payoffs in to consideration, which to our knowledge, has not been addressed before. In addition, we present a novel robotic application for a newly published identification algorithm for risk averse players which results in more detailed feature extraction.

Despite its complexity, the optimal bidding strategy has been studied thoroughly for private and common value first price auctions with symmetric or asymmetric players [8], [9], [10], [11]. Assuming a repeated first-price auction game for a multi task allocation process, we show that private values are identifiable for each teammate even if we remove the assumption of bidding truthfully. In each stage game, a target is announced to the team by a central agent and the mission is assigned to the highest bidder.

In addition, we present the solution to the identification problem in the case of partially disclosed information, in which only the winning bid is announced after each task assignment. Due to its negligible internal communication cost, partial disclosure of data can be used to update each agent's individual belief about her teammates' private attributes. This results in collaborative control techniques, the benefits of which are thoroughly studied (e.g. see [12]). We also extend our results to the more general case, where agents consider their expected future profits while deciding how much to bet on a certain task as well as trying to maximize their payoff for the current stage.

Since several factors are combined to define the utility (or cost) function (e.g. target's priority, relative position, or each agent's remaining supplies), neither the central agent nor any teammate is capable of deducing a certain attribute of the agent based on the observation of her proposed bid. Describing agents with fewer remaining energy levels as being less interested in accepting new missions, we explain how to use risk aversion as a divisive tool for encapsulating more information in each agent's submitting bids.

This paper is organized as follows. In the following section, the problem of multi-robot task allocation is formularized and the required concepts are defined for further reference. Section 3 describes the identification procedure for different models, which is followed by simulations. Finally, we summarize the proposed algorithm and suggest some topics for further consideration.

2 Problem Description

We use an exploration task as our benchmark. We assume that the number of agents in the field is n, fixed, and known to all participants. We denote the set of robots in the exploration team by $\mathcal{N}$. Using the terminology of game theory, we will refer to members of $\mathcal{N}$ as agents, or players, throughout the rest of the paper. For each player $i \in \mathcal{N}$ we need to define a private cost variable describing the expense she is charged by performing a certain task.

Suppose that there are multiple tasks occurring consecutively. We denote the set of all tasks by $\mathcal{T}$. With a slight abuse of notation, each task $t \in \mathcal{T}$ can be described using a tuple $< t, v^t, \boldsymbol{p}^t >$, in which the elements denote its time of occurrence, priority and position. Since the initial position of each target $t \in \mathcal{T}$ is unknown, $\boldsymbol{p}^t$ is a random variable. We distinguish random variables from their realizations using boldface letters.

We now define $\boldsymbol{e}_i^t$ as player i's expense incurred by performing task $t \in \mathcal{T}$. Intuitively, $\boldsymbol{e}_i^t$ is proportional to $\boldsymbol{d}_i^t$, the distance between the initial position of the agent i and its proposed target t (and thus is a random variable).

We assume that $\boldsymbol{d}_i^t$ is drawn from a common but unknown distribution function $F(\cdot)$ with its support set on $[\underline{v}, \overline{v}]$. To find a Bayesian Nash equilibrium, it is necessary for $f(\cdot) \triangleq F'(\cdot)$ to be bounded away from zero on $[\underline{v}, \overline{v}]$. Following the majority of literature on the subject, we define $\mathcal{F}$ as the set of twice differentiable cumulative distribution functions with bounded support, and assume $F \in \mathcal{F}$. It is noteworthy that since the valuations are defined based on factors like the distance from the target and its priority, the assumption of a finite upper bound for the support set of the distribution function is justifiable. Other conditions are weak as well.

We define the valuation of target $t \in \mathcal{T}$ for bidder $i \in \mathcal{N}$ by $\boldsymbol{v}_i^t = v^t - \boldsymbol{d}_i^t$. The missions can be described as a repeated game in which, in each stage game (at time t), the true position of the target $\boldsymbol{p}^t = p^t$ is announced to all agents. Therefore, each agent $i \in \mathcal{N}$ realizes its privately known valuation of the mission $\boldsymbol{v}_i^t = v_i^t$.

The expected utility of player $i \in \mathcal{N}$ for bidding b_i^t is described as:

$$\mathbb{E}\mathrm{u}_i(v_i^t, b_i^t) = \mathrm{w}_i(v_i^t - b_i^t)\Pr(b_j^t \leq b_i^t; \forall j \neq i) \qquad (1)$$

where $\Pr(b_j^t \leq b_i^t; \forall j \neq i)$ denotes the probability that player i assigns to submitting the highest bid. For the case where bidders are *ex ante* the same (symmetric bidders), we assume that $\mathrm{w}_i(\cdot) = \mathrm{w}(\cdot)$ for all $i \in \mathcal{N}$.

The aforementioned generic model covers both risk neutrality and risk aversion. For each risk neutral player $i \in \mathcal{N}$, $\mathrm{w}_i(x) = x$. In our problem of interest, risk neutrality means that the value that player i assigns to a certain task $t \in \mathcal{T}$, is the value of performing that task weighted by the probability of awarding it. On the other hand, risk averse bidders are reluctant to accept an agreement with risky payoff.

In other words, risk averse buyers may bid more to increase their chance of winning the contract, although this lowers their expected return, $v_i^t - b_i^t$. This fact is modeled by concave utility functions. In order to guarantee the existence of a symmetric Bayesian Nash equilibrium, in addition to concavity, other technical regularity conditions should be satisfied by $\mathrm{w}_i(\cdot)$. Therefore, we define a certain class of functions $\mathcal{W}$ as follows:

Assumption 1. Utility functions are chosen from $\mathcal{W}$, the set of utility functions $\mathrm{w}(\cdot)$ satisfying:

1. $\mathrm{w} : [0, +\infty) \to [0, +\infty), \mathrm{w}(0) = 0$, and $\mathrm{w}(1) = 1$,
2. $\mathrm{w}(\cdot)$ is continuous in $[0, +\infty)$, and admits three continuous derivatives on $[0, +\infty)$ with $\mathrm{w}'(\cdot) > 0$ in $(0, +\infty)$,
3. $\lim_{x \downarrow 0} \lambda^{(r)}(x)$ is finite for $r \in \{1, 2\}$, where $\lambda^{(r)}(x)$ is the rth derivative of $\lambda(\cdot) \triangleq \mathrm{w}(\cdot)/\mathrm{w}'(\cdot)$.

These regularity assumptions are weak and are satisfied by many von Neumann-Morgenstern utility functions, and requires $\lambda(\cdot)$ to admit two continuous derivatives on $[0, +\infty)$.

The curvature of any function $\mathrm{w}_i(\cdot) \in \mathcal{W}$ models how risk averse player i is. To formulate, $r_{\mathrm{w}}(\cdot) = -\frac{\mathrm{w}''(x)}{\mathrm{w}'(x)}$, and $R_{\mathrm{w}}(x) = x \cdot r_{\mathrm{w}}(x)$ are defined as the coefficients of absolute risk aversion (ARA) and relative risk aversion (RRA), respectively.

Player $i \in \mathcal{N}$ exhibits constant absolute risk aversion (CARA) when $\mathrm{w}_i(x) = 1 - e^{-\alpha x}$, where, $r_{\mathrm{w}_i}(x) = \alpha$ and $\alpha > 0$. In the same way, in CRRA, player i's utility function is of the form $\mathrm{w}_i(x) = x^{1-\beta}$, where $R_{\mathrm{w}_i}(x) = \beta$, is independent of x and $\beta \in (0, 1)$. It is noteworthy that for $\beta = 1$, players tend to be neutral and ignore the risk.

In the following section, we address how to identify each agent's private attributes using the aforementioned descriptions.

3 Learning Algorithm

Obviously, ordering agents to report their true types is the easiest way to learn their private valuations. However, the goal of this section is to confirm that optimal bidding behavior lets every player (and not only the central agent) to update her belief on her teammates' valuations as well. We extend our results by showing that estimations can be also made using even partially revealed information and for more general utility functions. On top of the aforementioned obviation, we close this section by addressing the identification algorithm for risk averse buyers that discloses more knowledge of agents' characteristics.

Generally, the process of learning in auctions can be described in two major steps [13]: in the first step, we *estimate* the distribution $G(\cdot)$ of announced equilibrium bids using methods like kernel estimation, particle filtering, etc. Consequently, knowing that players bid optimally, a relationship between each agent's type, v_i, and her submitted equilibrium bid, b_i, is found using first order conditions, and the distribution of player i's type, $F(\cdot)$, is *recovered* based on the estimated marginal distribution of equilibrium bids.

Our problem is said to be *identifiable* if the recovery step results in a unique model for agents. In general, a model is a set of structures $(\mathrm{w}, F) \in \mathcal{W} \times \mathcal{F}$ and is said to be *globally identified* using knowledge of the bid distribution function $G(\cdot)$, if there does not exist any structure $(\check{\mathrm{w}}, \check{F}) \in \mathcal{W} \times \mathcal{F}$ that leads to the same distribution. In the rest of this section, we discuss the identification problem for risk neutral and risk averse bidders.

3.1 Risk Neutral Agents

We start by modeling each robot as a risk neutral bidder. Since $\mathrm{w}_i(\cdot)$ is the identity function in this case, the recovery step reduces to approximating $F(\cdot)$, based on observing the disclosed information on submitted bids.

The identification problem for risk neutral bidders in first price auctions is addressed in [14]. Assuming that there is no *a priori* knowledge of $F(\cdot)$, they show that for symmetric players with independent private values (IPV model), the distribution of players' valuations is identified.

Based on the symmetry, $\mathrm{w}_i(x) = x$ for all $i \in \mathcal{N}$. It is useful to define $\tilde{\boldsymbol{b}}_{-i} = \max_{j \neq i} \boldsymbol{b}_j$, in describing player i's bidding behavior, as a random variable denoting the highest bid submitted by her opponents. Using the first order condition, we obtain:

$$b_i^t = v_i^t - \frac{G_{\tilde{\boldsymbol{b}}_{-i}|\boldsymbol{b}_i}(b_i^t|b_i^t)}{g_{\tilde{\boldsymbol{b}}_{-i}|\boldsymbol{b}_i}(b_i^t|b_i^t)} \tag{2}$$

where

$$G_{\tilde{\boldsymbol{b}}_{-i}|\boldsymbol{b}_i}(\tilde{b}_{-i}|b_i) = \Pr(\max_{j \neq i} \boldsymbol{b}_j \leq \tilde{b}_{-i}|\boldsymbol{b}_i = b_i)$$

denotes the distribution of the maximum equilibrium bid among player i's opponents conditional on i's own equilibrium bid, and $g_{\tilde{\boldsymbol{b}}_{-i}|\tilde{b}_i}(\tilde{b}_{-i}|b_i)$ denotes the corresponding conditional density.

The term $G_{\tilde{\boldsymbol{b}}_{-i}|\tilde{b}_i}(\tilde{b}_{-i}|b_i)/g_{\tilde{\boldsymbol{b}}_{-i}|\boldsymbol{b}_i}(\tilde{b}_{-i}|b_i)$ indicates how player i shades her bid. In the independent private values model, equation (2) is simplified to

$$b_i^t = v_i^t - \frac{G_{\boldsymbol{b}}(b_i^t)}{(n-1)g_{\boldsymbol{b}}(b_i^t)} \tag{3}$$

where $G_{\boldsymbol{b}}(\cdot)$ denotes the marginal distribution of equilibrium bids and $g_{\boldsymbol{b}}(\cdot)$ is the corresponding density. Note that because of the symmetry, $G_{\boldsymbol{b}_i}(\cdot) = G_{\boldsymbol{b}}(\cdot)$, and therefore equation (3) holds based on the first order statistics of random variable $\tilde{\boldsymbol{b}}_{-i}$.

Let $\hat{G}_{\boldsymbol{b}}(\cdot)$ and $\hat{g}_{\boldsymbol{b}}(\cdot)$ denote the approximated cumulative marginal distribution of equilibrium bids and its associated density, respectively. Suppose $T \in \mathbb{Z}_1$ targets have been announced so far. Using non-parametric empirical estimation, the *a priori* approximation of $\hat{G}_{\boldsymbol{b}}(\cdot)$ will be

$$\hat{G}_{\boldsymbol{b}}(b) = \frac{1}{T} \sum_{t=1}^{T} \sum_{i=1}^{n} \mathbb{1}\{b_i^t \leq b\} \tag{4}$$

in which $\mathbb{1}(\cdot)$ denotes the indicator function. Similarly, the density function of equilibrium bids can be approximated using kernel estimation as follows.

$$\hat{g}_{\boldsymbol{b}}(b) = \frac{1}{nTh_g} \sum_{t=1}^{T} \sum_{i=1}^{n} \mathrm{K}_g\big(\frac{b - b_i^t}{h_g}\big) \tag{5}$$

where $\mathrm{K}_g(\cdot)$ denotes a standard kernel function with compact support, and h_g is an appropriately chosen bandwidth sequence. Using equation (3), the pseudo private value of player $i \in \mathcal{N}$, $\hat{v}_i^t$ corresponding to her submitted bid for target t, b_i^t is defined as

$$\hat{v}_i^t \triangleq b_i^t + \frac{\hat{G}_{\boldsymbol{b}}(b_i^t)}{(n-1)\hat{g}_{\boldsymbol{b}}(b_i^t)} \tag{6}$$

Similarly, using pseudo private values $\{\hat{v}_i^t; i = 1, \ldots, n \text{ and } t = 1, \ldots, T\}$, we can estimate the density function of players' valuations using kernel approximation.

$$\hat{f}(v) = \frac{1}{nTh_f} \sum_{t=1}^{T} \sum_{i=1}^{n} \mathrm{K}_f\big(\frac{v - \hat{v}_i^t}{h_f}\big) \tag{7}$$

We now assume that only the winning bid is revealed after each stage game and try to estimate the distribution of types. It is remarkable that the second step in learning is replaced by estimation of the distribution of the winning bid using auction data. We denote the first order statistics of equilibrium bids by $\bar{\boldsymbol{b}} = \max_{i \in \mathcal{N}} \boldsymbol{b}_i$. For symmetric players with independent private values, we obtain

$$\frac{G_{\bar{\boldsymbol{b}}}(b)}{g_{\bar{\boldsymbol{b}}}(b)} = \frac{n-1}{n}\Big(\frac{G_{\boldsymbol{b}}(b)}{(n-1)g_{\boldsymbol{b}}(b)}\Big) \tag{8}$$

where $G_{\bar{\boldsymbol{b}}}(\cdot)$ denotes the cumulative distribution function of the winning bids and $g_{\bar{\boldsymbol{b}}}(\cdot)$ denotes its associated density function. Therefore, using (8), equation (4) can be rewritten as

$$\hat{v}_i^t \triangleq b_i^t + \frac{n}{n-1} \frac{\hat{G}_{\bar{\boldsymbol{b}}}(b_i^t)}{\hat{g}_{\bar{\boldsymbol{b}}}(b_i^t)} \tag{9}$$

where $\hat{G}_{\bar{\boldsymbol{b}}}(\cdot)$ denotes the approximate distribution function of the winning bid and can be calculated in a manner similar to that of the aforementioned kernel estimation method.

We now consider a more general utility function for each player $i \in \mathcal{N}$. Suppose that while choosing her optimal bidding strategy for stage $t \in \mathcal{T}$ each agent contemplates her future expected payoff (in the following stages $\{t+1, t+2, \cdots\}$) as well. Thus,

$$\begin{aligned} \mathrm{u}^*(v_i, \bar{b}) &= \max_b \{\mathrm{u}(v_i, b, \bar{b})\} \\ &= \max_b \Big\{ \Pr(i \text{ wins}|b)\big(\vartheta(v_i, \bar{b}) - b\big) \\ &\quad + \delta \Pr(i \text{ wins}|b)\ \mathbb{E}(\mathrm{u}^*(v, b)) \\ &\quad + \delta \sum_{j \in \mathcal{N}\setminus i} \Pr(j \text{ wins}|b)\ \mathbb{E}_{\breve{b} \geq b}\big(\mathbb{E}_v(\mathrm{u}^*(v, \breve{b}))\big)\Big\} \end{aligned} \tag{10}$$

where, $\mathrm{u}^*(v_i, \bar{b})$ denotes player i's maximum expected payoff if her private valuation for the target is v_i and the previously announced highest bid is $\bar{b}$. $\vartheta(v_i, \bar{b}) = cv_i + (1-c)v^*$ denotes the value of the target based on player i's private valuation and v^*, her estimate on the last winner's private valuation. $\Pr(i \text{ wins}|b)$ denotes the probability that player i assigns to the event of winning the mission by bidding b, and $\Pr(j \text{ wins}|b)$ denotes the probability that player i assigns to the event that the task is assigned to player j when player i bids b. Based on the symmetry, and using our notation,

$$\begin{cases} \Pr(i \text{ wins}|b) & = G^{n-1}(b) \\ \Pr(j \text{ wins}|b) & = \dfrac{1}{n-1}\big(1 - G^{n-1}(b)\big) \end{cases}$$

In our generic definition, $\delta \in [0,1]$ indicates players' persistency and is called the discount factor. The same results listed before are followed by choosing a small discount factor while $\delta = 1$ indicates that agents devalue their payoff for current stage. Using first order conditions, we solve the iterative equation (10) for the optimal bidding function b for uniformly distributed private values' case. Thus,

$$b(v_i, \bar{b}) = c\frac{n-1}{n}v_i + (1-c)\frac{n}{n-c}\bar{b} \tag{11}$$

Equation (11) reveals how optimal bids and private valuations are related and is used in the recovery phase. The method for estimating the distribution of bids is the same as the aforementioned case for risk neutral bidders with access to partially disclosed information.

3.2 Risk Averse Agents

It is logical to use risk aversion to model exploring agents. If we assume that robots become inactive after consuming their entire energy supply (which is often the case), player i's utility function should be concave to model her reluctance to winning the mission with much effort. We focus on CRRA and CARA models in which the measure of risk aversion is constant. By describing r_{w_i} and R_{w_i} so as to be proportional to the amount of remaining supplies for bidder $i \in \mathcal{N}$, we assume that it is unlikely that agent i wins several consecutive missions.

The inability to identify the structure of $(\mathrm{w}, F) \in \mathcal{W} \times \mathcal{F}$ using data released in a series of second price auctions forces the designer to perform a repeated game of first price auctions. (This is because bidding one's true type is still optimal for risk averse bidders in a second price auction.)

Although a general non-parametric model can not be identified using the bidding data, a semi-parametric solution was presented recently in [15]. In their working paper, several different parameterizations are considered in addition to proving that a more general model can not be identified. We summarize their propositions on general models in proposition 1.

Proposition 1. (See [15]) *In a repeated first price auction with full observation of bidding data,*

1. *No structure* $(\mathrm{w}, F) \in \mathcal{W} \times \mathcal{F}$ *is identified,*
2. *The semi-parametric model* $\mathcal{W} \times \mathcal{F}(\Gamma)$ *is not necessarily identified,*
3. *No structure* $(\mathrm{w}, F) \in \mathcal{W}^{CARA} \times \mathcal{F}$ *or* $\mathcal{W}^{CRRA} \times \mathcal{F}$ *is identified.*

Besides disproving the non-parametric identifiability of auction models with risk aversion, the proposition invalidates two semi-parametric approaches as well. In part (2), the private value distribution is parametrically parameterized as $F(\cdot, \gamma) \in \mathcal{F}(\Gamma)$ for some $\gamma \in \Gamma \subset \mathbb{R}$. In part (3) the utility function $\mathrm{w}(\cdot)$ is parameterized such that $\mathrm{w}(\cdot) = \mathrm{w}(\cdot, \alpha) \in \mathcal{W}^{\text{CARA}} \times \mathcal{F}$ for some $\alpha > 0$ or $\mathrm{w}(\cdot) = \mathrm{w}(\cdot, \beta) \in \mathcal{W}^{\text{CRRA}} \times \mathcal{F}$ for some $\beta \in [0,1]$.

However, as addressed in proposition 1, unilateral parameterization of the model does not suffice for performing the identification process. Therefore, the following impositions are enforced.

Assumption 2. For $\mathcal{I}$ a subset of $\{2, 3, \ldots\}$,

1. $\mathrm{w}(\cdot) = \mathrm{w}(\cdot, \theta) \in \mathcal{W}$ for every $\theta \in \Theta \subset \mathbb{R}$,
2. $F(\cdot|\cdot,\cdot) \in \mathcal{F}(\mathcal{Z} \times \mathcal{I}) \equiv \big\{F(\cdot|\cdot,\cdot) : F(\cdot|z, I) \in \mathcal{F},\ \forall (z, I) \in \mathcal{Z} \times \mathcal{I}\big\}$
3. For some $\alpha \in (0,1]$, $v_\alpha(z, I) = v_\alpha(z, I; \gamma)$ for all $(z, I) \in \mathcal{Z} \times \mathcal{I}$ and some $\gamma \in \Gamma \subset \mathbb{R}$,
4. The function $\phi_\alpha(z, I; \theta, \gamma) \equiv \lambda(v_\alpha(z, I; \gamma) - b_\alpha(z, I); \theta)$ for $(z, I) \in \mathcal{Z} \times \mathcal{I}$ determines uniquely $(\theta, \gamma) \in \Theta \times \Gamma$.

It is noteworthy that a set of variations in observed characteristics $\boldsymbol{z}$ and/or the number of bidders $I \in \mathcal{I}$ is assumed across auctions in order to reveal more information. In other words, private values are considered to be drawn from the conditional distribution $F(\cdot|z, I)$. Straightforwardly, conditions (1) and (2) require that $\mathrm{w}(\cdot)$ and $F(\cdot|z, I), \forall (z, I) \in \mathcal{Z} \times \mathcal{I}$ are both smooth functions.

The intuition behind conditions (3) and (4) is that parameterizing $F(\cdot)$ and $\mathrm{w}(\cdot)$, one is able to obtain more equations relating unknown variables (to be estimated) using α-quantiles of $\boldsymbol{v}$ and $\boldsymbol{b}$. The choice of $\alpha = 1$ results in a parametric estimation of the upper bound of the valuations' support set which is desirable due to its simplicity and fast convergence. We summarize the algorithm below:

1. Non-parametric estimation of the upper bound of submitted bids, $\bar{b}(\cdot,\cdot)$, and $g(\bar{b}(\cdot,\cdot)|\cdot,\cdot)$ at the observed data z_l, I_l,

2. Estimation of $(\theta, \overline{v})$ using nonlinear least square by $(\hat{\theta}_N, \hat{v}_N)$ from

$$g(\overline{b}(z_l, I_l)|z_l, I_l) = \frac{1}{I-1} \frac{1}{\lambda(\overline{v} - \overline{b}(z_l, I_l); \theta)}$$

3. Recovery of pseudo private values $\hat{v}_i$ using

$$\hat{v}_i^t = b_i^t + \lambda^{-1}\Big(\frac{1}{I-1} \frac{G(b_i^t)}{g(b_i^t)}\Big)$$

Using the three aforementioned steps, every team member i can estimate both the parameter θ of the utility function $\mathrm{w}(\cdot, \theta)$ and each player j' private valuation $\{v_j; \forall j \neq i\}$, thus their relative distance from the target and remaining energy by merely observing her submitted bids.

4 Simulation

Suppose a search and rescue operation [3], in which n robots are responsible for reporting characteristics of certain incidents (e.g. a leakage, fire, etc.). There are T missions (possible incidents) and robots are equipped enough to discover targets' characteristics individually.

Tasks are announced either by a central or a field agent who is also responsible for allocating them to team members. Either all submitted bids or only the highest ones are announced by the auctioneer to all agents to update their knowledge of their teammates' attributes.

Assume $n = 10$ and $T = 100$, therefore, a thousand bids are observed in full revelation case. Assume that players' initial distances from each target is identically and independently distributed with log-normal cumulative distribution function. In the same manner, we assume that agents' private values, $\{\boldsymbol{v}_i; i \in \mathcal{N}\}$, are distributed as

$$F(v) = \frac{1}{2}\Big[1 + \mathrm{erf}(\frac{\ln(v) - \mu}{\sigma\sqrt{2}})\Big]$$

with $\mu = 0$, and $\sigma = 1$. Consequently, $\ln(\boldsymbol{v})$ is a zero mean Gaussian random variable with $\sigma^2 = 1$. In order to comply with the characteristics of $\mathcal{F}$, we truncate $F(v)$ at $\underline{v} = 0$ and $\overline{v} = 5$. However, since $\boldsymbol{v}_i$ is proportional to the distance from the target, the scale is irrelevant.

Based on the rationality assumption, each player i submits her equilibrium bid to the central agent. Optimal bids are calculated using simple numerical integration techniques. Using kernel estimation with triweight kernel function

$$\mathrm{K}(x) = \frac{35}{32}\Big((1-x^2)^3\Big)\mathbb{1}(|x| \leq 1)$$

to compute pseudo density function of submitted bids and empirical approximation of the associated distribution function, pseudo private values are calculated using (6).

The release of all auctioning data results in a fairly accurate approximation. First fifty valuations (the data related to first five auctions) and their corresponding estimation using the proposed identification algorithm is drawn in Fig. 1. The identification method is validated by comparison of true and approximated types.

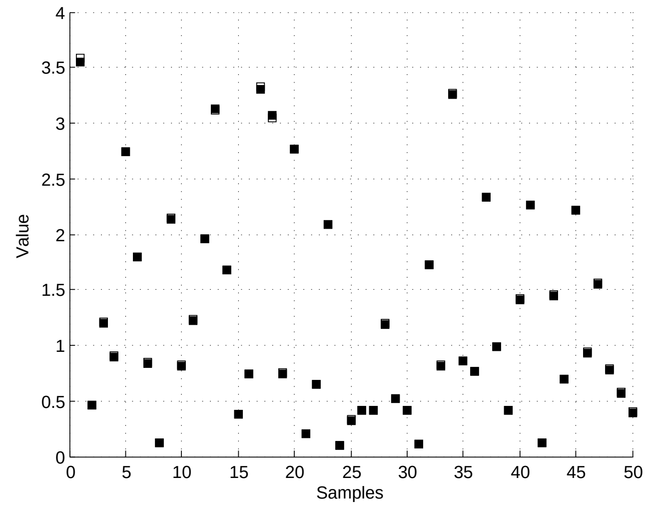

Fig. 1. Identification using fully revealed auctioning data for risk neutral players. First fifty true valuations (respectively, their associated pseudo private values) are shown by open (respectively, solid) squares. As can be seen, agents' private attributes is estimated accurately in the case of risk neutral agents with full observation

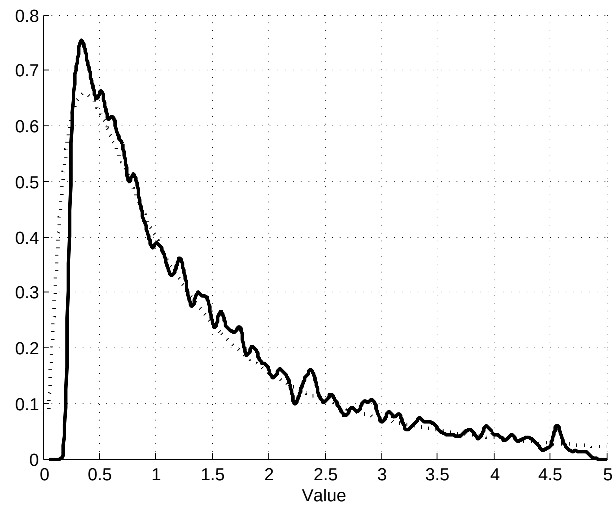

Fig. 2. Estimated private values density function. In the full revelation case, the approximate density function (solid line) is a fairly good estimate of the true density function (dotted line)

Using the same estimation technique for approximating the distribution of bids, pseudo density function of private attributes is approximated using (7). However, due to the lack of precision in kernel estimation in boundary points (valuations v close to $\overline{v}$ or $\underline{v}$), the training data is trimmed before applying the pseudo procedure. The resulting approximated density of private values is drawn and compared with our initial choice in Fig. 2. It is noteworthy that $\hat{f}(v)$ is estimated *non-parametrically*, since there were no initial assumptions on how values were distributed.

As mentioned in Section 3-A, valuations can be identi-

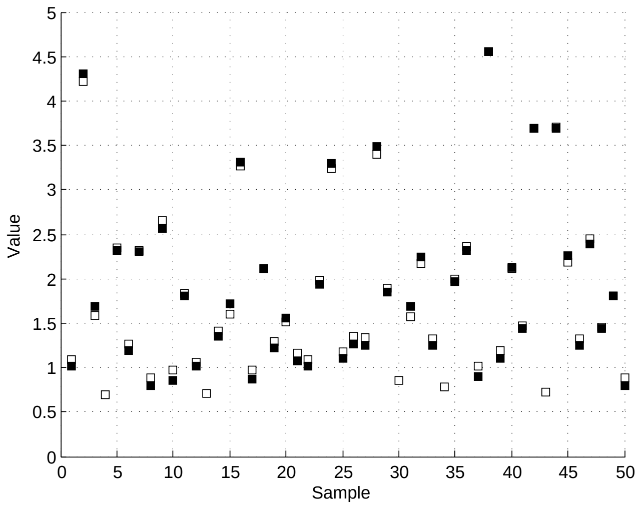

Fig. 3. Identification using partially revealed auctioning data for risk neutral players.Fifty true winning valuations (respectively, their associated pseudo private values) are shown by open (respectively, solid) squares.

fied using only the announcement of highest submitted bid. We choose a training set of 100 samples associated with the winning bids for targets $\{t \in \mathcal{T}\}$. Since agents are assumed to be identical, winners' identities are insignificant.

Using the similar kernel estimation techniques to approximate $g_{\bar{\boldsymbol{b}}}(\cdot)$ and its associated distribution, the order statistics of $\boldsymbol{b}$ is obtained. Therefore, private values are approximated using (9). Fifty samples of the resulting pseudo valuations and their true realizations are drawn in Fig. 3. Noting the isolated open squares and relatively less accurate approximation, the weakness of kernel estimation in calculating very low (or high) values is visible.

5 Conclusion and Further Research

The problem of identification in a generic collaborative task is studied in this paper. Using non-parametric approximation methods, it has been shown that players' individual attributes can be estimated satisfactorily when they send optimal bidding signals to a central agent. The data recovery procedure has been addressed for both cases of full data revelation, in which the entire auctioning data is available, and the case of partial information disclosure, in which after the assignment of each new target, merely the winning bid is announced to the team. We showed that even with partially revealed information (which lowers the cost of communication and enhances the chance of success), a good approximation of agents' attributes is achievable. This is advantageous due to bandwidth deficiency in catastrophic situations.

We extended our results to the more generic case of state identification for more intelligent agents who choose their optimal policy based on its associated long-term profit. Considering future expected payoffs results in a more meaningful cooperation which we showed to be identifiable as well. In addition, by modeling each robot's remaining energy supply as a measure of how much it dares to risk for performing a new task, agents' individual utility functions were modified. Using first-price auction methodology a procedure is addressed to recover agents' two distinctive private attributes, which was impossible for risk neutral agents.

Although, several parametric approximation methods can be used for identifying each robot's individual properties, due to the random nature of robots' initial deployment (relative to targets' location) in many practical collaborative tasks, non-parametric solutions (e.g. the methods presented in this paper) seem to be more appropriate. However, finding an optimal parametric estimation technique (based on the nature of the problem) is an interesting research topic.

References

[1] A.W. Stroupe, M.C. Martin, and T. Balch "Distributed sensor fusion for object position estimation by multi-robot systems," *Proc. Int'l Conf. Robotics and Automation*, vol. 2, pp. 1092-1098, 2001.

[2] W. Burgard, M. Moors, C. Stachniss, F.E. Schneider, "Coordinated Multi-Robot Exploration," *IEEE Trans. on Robotics*, vol. 21, issue 3, pp. 376-386, June 2005.

[3] E. Jones, M.B. Dias, and A. Stentz, "Learning-enhanced Market-based Task Allocation for Disaster Response," *Tech. Report CMU-RI-TR-06-48*, Robotics Institute, Carnegie Mellon University, October 2006.

[4] M. Mazo, Jr., A. Speranzon, K.H. Johansson, and Hu Xiaoming, "Multi-robot tracking of a moving object using directional sensors," *Proc. Int'l Conf. Robotics and Automation*, vol.2, pp. 1103-1108, May 2004.

[5] J.E. Vargas, K. Tvarlaparti, and Z.J. Wu, "Target Tracking with Bayesian Estimation," *Distributed Sensor Networks: A Multiagent Perspective*, edited by V. Lesser, C.L. Ortiz, and M. Tambe, Kluwer Academic Publisher, ISBN 1402074999, pp. 43-58 2003

[6] M.B. Dias, R.M. Zlot, N. Kalra, and A. Stentz, "Market-Based Multirobot Coordination: a Survey and Analysis," *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1257-1270, July 2006.

[7] N. Nisan, "Introduction to Mechanism Design (for Computer Scientists)," *Algorithmic Game Theory*, to be published by Cambridge University Press, Available: http://www.cs.huji.ac.il/∼noam/introMD.pdf, to appear.

[8] P.R. Milgrom, *Putting Auction Theory to Work*, Cambridge University Press, ISBN 0521536723, 2004.

[9] S. Campo, I. Perrigne, and Q. Vuong, "Asymmetry in first-price auctions with affiliated private values," *J. of Applied Econometrics*, vol. 18, issue 2, pp. 179-207, January 2003.

[10] J.J. Laffont, H. Ossard, and Quang Vuong, "Econometrics of First-Price Auctions," *Econometrica*, Vol. 63, No. 4, pp. 953-980, July 1995.

[11] J. Kagel and D. Levin, *Common Value Auctions and the Winner's Curse*, Princeton University Press, 2002.

[12] I. Rekleitis, G. Dudek, and E. Milios, "Multi-robot collaboration for robust exploration," *Annals of Mathematics and Artificial Intelligence*, vol. 31, pp. 7-40, October 2001.

[13] S. Athey, and P.A. Haile, "Nonparametric Approaches to Auctions," *Handbook of Econometrics*, edited by J. Heckman and E. Leamer, vol. 6A, part 14, North-Holland, November 2007 (forthcoming).

[14] E. Guerre, I. Perrigne, and Q. Vuong, "Optimal Nonparametric Estimation of First-price Auctions," *Econometrica*, vol. 68, no. 3, pp. 525-574, May 2000.

[15] S. Campo, E. Guerre, I. Perrigne, and Q. Vuong, "Semiparametric Estimation of First-Price Auctions with Risk Averse Bidders," *Working Paper*, University of Southern California, Available: http://www.qmul.ac.uk/∼tew077/CGPV6.pdf, March 2006.