

IDENTIFYING COORDINATION IN A COGNITIVE RADAR NETWORK - A MULTI-OBJECTIVE INVERSE REINFORCEMENT LEARNING APPROACH

Luke Snow[†], Vikram Krishnamurthy[†], and Brian M. Sadler^{*}

[†] Electrical and Computer Engineering, Cornell University, Ithaca, NY

^{*} DEVCOM Army Research Laboratory, Adelphi, MD

ABSTRACT

Consider a target being tracked by a cognitive radar network. If the target can intercept some radar network emissions, how can it detect coordination among the radars? By 'coordination' we mean that the radar emissions satisfy Pareto optimality with respect to multi-objective optimization over each radar's utility. This paper provides a novel multi-objective inverse reinforcement learning approach which allows for both detection of such Pareto optimal ('coordinating') behavior and subsequent reconstruction of each radar's utility function, given a finite dataset of radar network emissions. The method for accomplishing this is derived from the micro-economic setting of Revealed Preferences, and also applies to more general problems of inverse detection and learning of multi-objective optimizing systems.

Index Terms— Cognitive Radar, Multi-Objective Inverse Reinforcement Learning, Revealed Preferences

1. INTRODUCTION

Cognitive radars [1] use the perception-action cycle of cognition to sense the target, learn relevant information, then optimally adapt their output emissions in response. We consider the case when there is a *network* of cognitive radars which collaborate to optimally track a target. In a coordinating radar network, not only do the individual cognitive radars optimally adapt their output (with respect to an individual utility function) subject to resource constraints, but also the *allocation of resources* between radars is subject to an optimization procedure. The resource to be allocated is often interpreted as the total power available to the radar network at each time step. Such an optimal power allocation strategy has been studied in [2], [3], [4], [5], in which algorithmic game-theoretic methods are employed. Specifically, [2] poses the problem of adaptive power allocation for radar networks as a cooperative game, and provides an iterative cooperative Nash bargaining algorithm which converges quickly to the Pareto optimal equilibrium.

In this work we are interested in the *inverse* problem. Suppose "we" are a target being tracked by a group of cognitive radars, and we have a mechanism for intercepting some of their emissions. Given such a set of measured signals, how can we determine if the radars compose a *coordinating* (in the sense of Pareto optimal power allocation) network? If they are coordinating, can we subsequently reconstruct the individual utility functions which induce each radar's output? This problem is similar to multi-agent inverse reinforcement learning (MAIRL) frameworks [6], [7], [8] in that we aim to reconstruct each radar's utility function. However, it is fundamentally different in that we first aim to *detect coordination* among the radars; MAIRL approaches *a priori* assume some structure, such as coordination, among the agents. Thus, this problem can be considered at the intersection of inverse game theory [9] and MAIRL.

In [10], the framework of [11] was exploited to approach this problem for the single-radar case (i.e. detect if a single radar is a constrained utility maximizer). We extend this work to the multi-radar regime by employing the micro-economic framework of [12]. In [12] a Revealed Preference methodology is applied to the setting of collective consumption models, and a framework is developed for testing consistency with Pareto optimality given a finite dataset of input-output group responses. Remarkably, conditions are given which are *necessary and sufficient* for a finite dataset to be consistent with Pareto optimality. We map our radar network inverse problem into this framework, thus allowing us to test for Pareto optimal power allocation (coordination) given a finite dataset of (incompletely) measured cognitive radar network emissions. We also show how this can be used to reconstruct feasible utility functions for each observed radar.

This paper is organized as follows: In section 2 we provide background on the micro-economic framework of [12]. In section 3 we introduce the radar network protocol and measurement model, and show how this can be mapped into the framework presented in Section 2. In section 4 we provide a numerical example that illustrates the validity of our coordination-testing and utility reconstruction procedures. Finally we conclude in section 5.

^{*}This research was funded by National Science Foundation grant CCF-2112457, Army Research office grant W911NF-21-1-0093, and Air Force Office of Scientific Research grant FA9550-22-1-0016.

2. COORDINATION DETECTION AND UTILITY RECONSTRUCTION

Our framework for coordination detection and utility function reconstruction (inverse reinforcement learning) is the micro-economic setting of Revealed Preferences. Here we present the Revealed Preference framework for the collective consumption setting [12]. We present this framework in its original generality, but in section 3 we adapt it to the radar network power allocation setting.

2.1. Agent Consumption

Consider a group composed of M members, each of which can consume some quantity of N goods. At time $t \in \mathbb{N}$ each member $i \in \{1, \dots, M\}$ has consumption given by the vector $\beta_t^i \in \mathbb{R}_+^N$, and the aggregate group consumption is given by $\beta_t = \sum_{i=1}^M \beta_t^i \in \mathbb{R}_+^N$. It is assumed that the preferences of each member i can be represented by a non-satiated and non-decreasing utility function $U^i(\beta)$, $\beta \in \mathbb{R}_+^N$. If the agents choose their individual consumptions in accordance with group Pareto optimality, then we say the group *coordinates*:

Definition 1. A group of M agents with respective utility functions $\{U^i : \mathbb{R}_+^N \rightarrow \mathbb{R}\}_{i=1}^M$, and subject to price vector α_t , behaves with **coordination** if consumption $\{\beta_t^i\}_{i=1}^M$ satisfies

$$\begin{aligned} \sum_{i=1}^M \mu_t^i U^i(\beta_t^i) &\geq \sum_{i=1}^M \mu_t^i U^i(\zeta^i) \\ \forall \{\zeta^i \in \mathbb{R}_+^N\}_{i=1}^M : \alpha_t'(\sum_i \zeta^i) &\leq \alpha_t'(\sum_i \beta_t^i) \end{aligned} \quad (1)$$

for some set of Pareto weights $\{\mu_t^i\}_{i=1}^M$

The Pareto weights $\{\mu_t^i\}_{i=1}^M$ correspond to the bargaining power, or relative importance, of each individual in the formation of the Pareto optimal solution. We now switch perspectives to an empirical analyst who observes the group consumption and attempts to detect coordination, in the sense of Definition 1.

2.2. Analyst Detection of Coordination

For each $t \in \{1, \dots, T\}$, suppose an analyst observes probe signals $\alpha_t \in \mathbb{R}_+^N$, aggregate consumption $\beta_t = \sum_{i=1}^M \beta_t^i \in \mathbb{R}_+^N$, and "assignable quantities" $\hat{\beta}_t^i \leq \beta_t^i \forall i \in \{1, \dots, M\}$. The assignable quantity $\hat{\beta}_t^i$ represents observed consumption of some subset of the total consumed quantities by individual i at time t , and hence it is elementwise no greater than the actual consumption vector β_t^i . The apparatus by which these assignable quantities are observed, and the amount observed, varies by application. We denote this dataset as

$$\mathcal{D} = \{\alpha_t, \beta_t, \{\hat{\beta}_t^i\}_{i=1}^M, t \in \{1, \dots, T\}\} \quad (2)$$

We emphasize that the true individual consumption vectors β_t^i are hidden. Given this dataset, [12] provides necessary and sufficient conditions for consistency with coordination, in the sense of Definition 1. To formulate these conditions, we need one more definition: For each observation t , *feasible personalized quantities* $q_t^i \in \mathbb{R}_+^N$, $i = 1, \dots, M$ satisfy $q_t^i \geq \hat{\beta}_t^i \forall i$ and $\sum_{i=1}^M q_t^i = \beta_t$.

Now we provide the main result of [12] which states the equivalence between a consistency of a dataset $\mathcal{D}$ with coordination and existence of a non-empty feasible region of a set of inequalities.

Theorem 1. [12] Let $\mathcal{D}$ be a set of observations. The following are equivalent:

1. there exist a set of M concave and continuous utility functions $U^1, \dots, U^M$ such that $\mathcal{D}$ is consistent with coordination, in the sense of Definition 1.
2. there exist feasible personalized quantities q_t^i and numbers $u_s^i > 0, \lambda_j^i > 0$ such that for all $s, t \in \{1, \dots, T\}$:

$$u_s^i - u_t^i \leq \lambda_t^i [\alpha_t' q_s^i - \alpha_t' q_t^i] \quad (3)$$

for each member $i = 1, \dots, M$.

Proof. See Appendix A.1 of [12]. $\square$

Lemma 1. Suppose the conditions of Theorem 1 hold. Explicit monotone and concave utility functions that rationalize the dataset by satisfying the condition of coordination given in Definition 1 are given by

$$\hat{U}^i(\beta) = \min_{t \in \{1, \dots, N\}} \{u_t^i + \lambda_t^i \alpha_t'(\beta - q_t^i)\} \quad (4)$$

Proof. This follows immediately by application of Afriat's Theorem [11]. $\square$

Thus if the feasible region of (3) is nonempty, then monotone concave rationalizing utility functions can be reconstructed using (4).

Implementing this feasibility test (3) in practice presents a problem however, because the set $\{q_t^i\}_{i=1}^M$ is *unobserved*. To overcome this, [12] present an equivalent mixed-integer linear program (MILP) formulation which allows for practical computation of this coordination test.

Theorem 2. [12] Let $\mathcal{D}$ be a set of observations. There exists a set of M concave and continuous utility functions $U^1, \dots, U^M$ such that $\mathcal{D}$ is consistent with coordination if and only if there exist $q_t^i \in \mathbb{R}_+^N$, $\eta_t^i \in \mathbb{R}_+$, and $x_{st}^i \in \{0, 1\}$, $i = 1, \dots, M$ that satisfy

- i) $\sum_{i=1}^M q_t^i = \beta_t$ and $q_t^i \geq \hat{\beta}_t^i$
- ii) $\eta_t^i = \alpha_t' q_t^i$
- iii) $\eta_s^i - \alpha_t' q_t^i < y_s x_{st}^i$
- iv) $x_{su}^i + x_{ut}^i \leq 1 + x_{st}^i$
- v) $\eta_t^i - \alpha_t' q_t^i \leq y_t(1 - x_{st}^i)$

where $y_t = \alpha_t' \beta_t$ is the group consumption at time t .

Proof. See Proposition 2 of [12]. $\square$




Fig. 1: Interaction of our dynamics with the adversary's cognitive radar network. Each cognitive radar is comprised of a Bayesian tracker and a radar controller. Based on the time series $\{\alpha_n, \beta_n, \beta_n^1, \dots, \beta_n^M\}_{n=1}^T$, our goal is to determine if the radar network is coordinating, i.e., if the true responses $\{\beta_n^i, n = 1, \dots, T\}_{i=1}^M$ satisfy coordination as in (1).

3. TESTING FOR COGNITIVE RADAR NETWORK COORDINATION

Here we test for Pareto optimal power allocation among M radars, and subsequently reconstruct each radar's utility function, from a finite dataset of intercepted radar emissions. We take a target-centric view and assume that the target can measure the power of the radar network signals that are being output in response to its maneuvers.

3.1. Waveform Adaptation by Cognitive Radar Network

3.1.1. Cognitive Radar Network Response Model

Here we outline the abstract interaction between "us" (target) and the network of M cognitive radars. Let $n = 1, 2, \dots$ denote discrete time. Our probe signal (modeling acceleration maneuvers) is $\alpha_n \in \mathbb{R}^N$ and the response of radar i is $\beta_n^i \in \mathbb{R}^N$. The Pareto optimal interaction dynamics are as follows:

- i) our state $x_n \sim p_{\alpha_n}(x|x_{n-1}), x_0 \sim \pi_0$
- ii) radar i observation $y_n^i \sim p_{\beta_n^i}(y|x_n)$
- iii) radar i tracker $\pi_n^i = T(\pi_{n-1}^i, y_n^i)$
- iv) radar i output β_n^i , where

$$\{\beta_n^i\}_{i=1}^M \in \operatorname{argmax}_{\{\zeta_i\}_{i=1}^M: \alpha_n(\sum_{i=1}^M \zeta_i) \leq C} \sum_{i=1}^M \mu_n^i U^i(\zeta_i) \quad (5)$$

and C is the radar network power constraint

- v) aggregate radar network output $\beta_n = \sum_{i=1}^M \beta_n^i$

These interaction dynamics can be seen in Figure 1. The interpretation of α_n is as a parameter which influences the state transition dynamics of the target (us). This will soon be specified in terms of the eigenvalue spectrum of the state-noise covariance matrix. The coordinating radar network output $\{\beta_n^i\}_{i=1}^M$ lies on the Pareto Front corresponding to the multi-objective optimization

problem $\operatorname{argmax}_{\{\zeta_i\}_{i=1}^M} \sum_{i=1}^M \mu_n^i U^i(\zeta_i)$, with constraint $\alpha_n(\sum_{i=1}^M \zeta_i) \leq C$ and Pareto weights $\{\mu_n^i\}_{i=1}^M$. At an abstract level, the observations of the i 'th radar, y_n^i , are parameterized by its output signal β_n^i . We will soon specify this interpretation in terms of the eigenvalue spectrum of the measurement noise covariance matrix.

3.1.2. Linear Gaussian Target Model and Radar Tracker

Linear Gaussian dynamics for a target's kinematics [13] and linear Gaussian measurements at each radar are widely assumed as a useful approximation [14]. Thus we will consider the following linear Gaussian state dynamics and measurements, as a special case of the general network response model of the previous subsection:

$$\begin{aligned} x_{n+1} &= Ax_n + w_n(\alpha_n), \quad x_0 \sim \pi_0 \\ y_n^i &= Cx_n + v_n^i(\beta_n^i) \end{aligned} \quad (6)$$

Here we represent the target (our) state as $x_n \in \mathbb{R}^N$, with initial density $\pi_0 \sim \mathcal{N}(\hat{x}_0, \Sigma_0)$. Cognitive radar i obtains measurements $y_n^i \in \mathbb{R}^N$. State and measurement noise processes are given by $w_n \sim \mathcal{N}(0, Q(\alpha_n)), v_n^i \sim \mathcal{N}(0, R(\beta_n^i))$, respectively, and we assume $\{w_n\}, \{v_n^1\}, \dots, \{v_n^M\}$ are independent i.i.d. processes. Note that we implement the dependence of the dynamics on signal α_n and response $\{\beta_n^i\}$ through the state and observation noise covariance matrices Q and R , respectively. A radar's control of its ambiguity function can be mapped to a control of the measurement noise covariance matrix R . However, reducing the measurement noise covariance incurs several costs to the system, such as increased power output. In our setting the cost constraint is interpreted as a bound on the total available power to be allocated across the M radars. The target's probing of the radars is performed via purposeful maneuvers that result in modulation of the state noise covariance matrix Q in (6) by α_n .

3.1.3. Testing for Coordination among Cognitive Radars

We now show how Theorem 1 can be used to determine if the cognitive radar network is coordinating, i.e. the responses satisfy coordination in the sense of Definition (1). What remains is to specify the structure of the process by which we (the target) can acquire information about the radar network signal output. Specifically, suppose:

- i) Our probe α_n , characterizing our maneuvers, is the vector of eigenvalues of the positive definite matrix Q governing the state-noise covariance in (6).
- ii) The response β_n^i of the i 'th radar is the vector of eigenvalues of the positive definite matrix R^{-1} governing the noise covariance in (6).
- iii) $\{\beta_n^i\}_{i=1}^M$ is chosen to be Pareto optimal with respect to individual utilities $\{U^i\}_{i=1}^M$ and Pareto weights $\{\mu_n^i\}_{i=1}^M$, and subject to the constraint $\alpha_n(\sum_{i=1}^M \beta_n^i) \leq 1$, i.e. chosen according to (5).

- iv) We (the target) can observe the aggregate output signal power across all M radars $\beta_n = \sum_{i=1}^M \beta_n^i$ through a broadband omni-directional receiver configuration.
- v) For each radar we observe an associated signal power $\hat{\beta}_n^i$ which is upper-bounded by the true signal power β_n^i , i.e. $\hat{\beta}_n^i \leq \beta_n^i$. This is a reasonable assumption that can arise from a narrowbeam receiver that steers to one radar at a time, and thus only observes some fraction of the power output of each radar.

Thus we obtain the dataset

$$\mathcal{D} = \{\alpha_n, \beta_n, \{\hat{\beta}_n^i\}_{i=1}^M, n = 1, \dots, T\} \quad (7)$$

with $\alpha_n, \beta_n, \hat{\beta}_n^i \in \mathbb{R}^N$. Then we can directly use the MILP from Theorem 2 to test for coordination, and if this MILP has a non-empty feasible region then we can employ Lemma 1 to reconstruct feasible utility functions for each of the M cognitive radars.

4. NUMERICAL EXAMPLE: TRI-RADAR SPECTRAL COORDINATION TEST

Here we demonstrate the use of the MILP of Theorem 2 to test for Pareto optimal power allocation of a network of $M = 3$ radars. For ease of visualization, we choose $N = 2$ so that the probe and response signals $\alpha_n, \beta_n^i \in \mathbb{R}^2$ and thus the utility functions $U^i(\cdot)$ can be displayed on a 2-d contour plot. The elements of our probe signal α_n are chosen randomly and independently over time n as $\alpha_n(1) \sim \text{Unif}(0.1, 1.1)$ and $\alpha_n(2) \sim \text{Unif}(0.1, 1.1)$ where $\text{Unif}(a, b)$ denotes the uniform pdf with support (a, b) . Our probe signal parameterizes the state-noise covariance matrix as $Q_n = \text{diag}[\alpha_n(1), \alpha_n(2)]$ in (6). Suppose the utility functions of the radars are $U^1(\beta) = \det(R^{-1}(\beta)) = \beta(1) \times \beta(2)$, $U^2(\beta) = \text{Tr}(R^{-1}(\beta)) = \beta(1) + \beta(2)$, $U^3(\beta) = \sqrt{\beta(1) \times \beta(2)}$. The distributed response $\{\beta_n^i\}_{i=1}^3$ is then chosen at each time n such that (5) is satisfied, with $M = 3$, $C = 1$, $\mu_n^1 = 0.3, \mu_n^2 = 0.3, \mu_n^3 = 0.4 \forall n \in \{1, \dots, T\}$. This simulation is run for $T = 10$ steps.

We simulate the sensor-level measurement to obtain the dataset (2) in the following way: the aggregate consumption $\beta_n = \sum_{i=1}^3 \beta_n^i$ is directly measured for all $n = 1, \dots, T$, and the assignable quantities are obtained as $\beta_n^i = S \times \beta_n^i$, where $S \sim \text{Unif}(0.1, 1)$. Using this dataset, we have verified¹ that the MILP of Theorem 2 has a feasible solution², indicating that the data is consistent with Pareto optimality, or coordination. Then taking a resultant feasible set $\{q_t^i, i = 1, 2, 3\}_{t=1}^{10}$, we use (4) to reconstruct feasible utility functions which rationalize the data. These reconstructed utility functions

¹Matlab's `intlinprog` function is used to test the linear program

²There exists $\{q_t^i \in \mathbb{R}_+^N, \eta_t^i \in \mathbb{R}_+, x_{st}^i \in \{0, 1\}, i = 1, 2, 3\}_{s,t=1}^{10}$ such that conditions $i - v$ of Theorem 2 hold.




Fig. 2: $U^i(\beta)$ is the true utility function of the i 'th radar, inducing the responses $\{\hat{\beta}_n^i\}_{n=1}^{10}$. Given the dataset $\mathcal{D} = \{\alpha_n, \hat{\beta}_n, \{\hat{\beta}_n^i\}_{i=1}^3\}_{n=1}^{10}$, the MILP 2 has a non-empty feasible region. From this MILP we obtain the set $\{q_t^i\}_{t=1}^{10}$ and reconstruct $\hat{U}^i(\beta)$ via (4) for each $i = 1, 2, 3$.

are shown in Fig. 2, where plots 2a, 2c, 2e are the true utility functions of each radar and plots 2b, 2d, 2f are the reconstructed utility functions for each respective radar, given the dataset $\mathcal{D}$.

Conversely, we have verified that the test of Theorem 2 is capable of rejecting non-cooperative data. Specifically, we tested the MILP 100 times with each β_n^i generated independently (without coordination) as a two-dimensional uniform random variable, $\beta_n^i = [\beta_n^i(1), \beta_n^i(2)]$ where $\beta_n^i(k) \sim \text{Unif}(0, 1), k = 1, 2$. Out of these 100 tests, 5 resulted in a non-empty feasible region, indicating coordination is present, and 95 resulted in an empty feasible region, indicating the absence of coordination. Thus we observe a Type-II error (conclude that the radars are coordinating when in fact they are acting independently) rate of 5%.

5. CONCLUSION

We have utilized the micro-economic framework of Revealed Preferences to test for Pareto optimality, indicating coordination, among a cognitive radar network. We have provided a numerical example which validates the proposed framework and illustrates the reconstructed utility functions of the individual cognitive radars.

6. REFERENCES

- [1] S. Haykin, “Cognitive radar: a way of the future,” *IEEE signal processing magazine*, vol. 23, no. 1, pp. 30–40, 2006.
- [2] C. Shi, S. Salous, F. Wang, and J. Zhou, “Power allocation for target detection in radar networks based on low probability of intercept: A cooperative game theoretical strategy,” *Radio Science*, vol. 52, no. 8, pp. 1030–1045, 2017.
- [3] A. Panoui, S. Lambotharan, and J. A. Chambers, “Game theoretic power allocation technique for a mimo radar network,” in *2014 6th International Symposium on Communications, Control and Signal Processing (ISCCSP)*. IEEE, 2014, pp. 509–512.
- [4] P. Chavali and A. Nehorai, “Scheduling and power allocation in a cognitive radar network for multiple-target tracking,” *IEEE Transactions on Signal Processing*, vol. 60, no. 2, pp. 715–729, 2011.
- [5] G. Bacci, L. Sanguinetti, M. S. Greco, and M. Luise, “A game-theoretic approach for energy-efficient detection in radar sensor networks,” in *2012 IEEE 7th Sensor Array and Multichannel Signal Processing Workshop (SAM)*. IEEE, 2012, pp. 157–160.
- [6] L. Yu, J. Song, and S. Ermon, “Multi-agent adversarial inverse reinforcement learning,” in *International Conference on Machine Learning*. PMLR, 2019, pp. 7194–7201.
- [7] S. Natarajan, G. Kunapuli, K. Judah, P. Tadepalli, K. Kersting, and J. Shavlik, “Multi-agent inverse reinforcement learning,” in *2010 ninth international conference on machine learning and applications*. IEEE, 2010, pp. 395–400.
- [8] X. Lin, S. C. Adams, and P. A. Beling, “Multi-agent inverse reinforcement learning for certain general-sum stochastic games,” *Journal of Artificial Intelligence Research*, vol. 66, pp. 473–502, 2019.
- [9] V. Kuleshov and O. Schrijvers, “Inverse game theory: Learning utilities in succinct games,” in *International Conference on Web and Internet Economics*. Springer, 2015, pp. 413–427.
- [10] V. Krishnamurthy, D. Anglely, R. Evans, and B. Moran, “Identifying cognitive radars-inverse reinforcement learning using revealed preferences,” *IEEE Transactions on Signal Processing*, vol. 68, pp. 4529–4542, 2020.
- [11] S. N. Afriat, “The construction of utility functions from expenditure data,” *International economic review*, vol. 8, no. 1, pp. 67–77, 1967.
- [12] L. Cherchye, B. De Rock, and F. Vermeulen, “The revealed preference approach to collective consumption behaviour: Testing and sharing rule recovery,” *The Review of Economic Studies*, vol. 78, no. 1, pp. 176–198, 2011.
- [13] X. R. Li and V. P. Jilkov, “Survey of maneuvering target tracking. part i. dynamic models,” *IEEE Transactions on aerospace and electronic systems*, vol. 39, no. 4, pp. 1333–1364, 2003.
- [14] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with applications to tracking and navigation: theory algorithms and software*. John Wiley & Sons, 2004.