IEOR 290A – Lecture 36
Inverse Decision Making

1 Problem Framework

Suppose we have an individual (or system) making decisions by optimizing some utility (or cost) function that depends on inputs from the environment. In some problems, we have the following scenario: We get to observe the inputs and the corresponding decisions, and we would like to infer the utility function of the individual. This problem is actually a type of regression problem, but there are additional complications which mean that we have to use a more complex formulation in order to solve the problem.

The first question that needs to be addressed is the model of the individual, and the number of individuals. There are two important cases.

1.1 Utility Maximizing Agent

Suppose that an agent makes decisions by solving the following optimization problem:

$$x_i^* = \arg\max\{J(x, u_i) \mid x \in \mathcal{X}(u_i)\},$$

where $u_i \in \mathbb{R}^q$ are inputs, $x_i^* \in \mathbb{R}^d$ are decisions, $J(x, u_i)$ is the utility function of the agent, and $\mathcal{X}(u_i)$ is a bounded set (that depends on u_i). In this model, we observe (u_i, x_i^*) for $i = 1, \ldots, n$ and would like to infer the function $J(x, u_i)$.

Note that this same model has an alternative interpretation. Suppose we have an expert (or controller) that is operating a system by minimizing a cost function, similar to MPC or LBMPC. And suppose that we observe the control actions and states of the system, but do not know the cost function. This scenario is captured by the above model by noting that the cost function is $-J(x, u_i)$, u_i is the initial condition of the system (corresponding to x_0 in MPC), and x_i^* is the control action chosen (corresponding to u_0).

1.2 Multiple Strategic Agents

Suppose that we have $p \geq 2$ agents, and the k-th agent is making decisions $x^{*,k} \in \mathcal{X}^k$ to maximize their utility function while also taking into account the strategic behavior of each other agent. In this model, we need to also specify our notion of strategic behavior. Perhaps the most well known notion is that of the Nash equilibrium, in which

$$x_i^{*,k} \in \arg\max\{J^k(x_i^{*,1}, \ldots, x_i^{*,k-1}, x^k, x_i^{*,k+1}, \ldots, x_i^{*,p}, u_i) \mid x^k \in \mathcal{X}^k\}.$$

However, there are other notions of rationality such as correlated equilibria. Similar to the case of the utility maximizing agent, we observe $(u_i, x_i^{*,1}, \ldots, x_i^{*,p})$ for $i = 1, \ldots_n$ and would like to infer the utility function of each agent J^k.

2 Key Technical Difficulty

These problems are difficult to solve, and they are special cases of what are known as inverse optimization problems. To specifically understand the difficulty, consider the utility maximizing agent. And assume that we have a parametrization of the utility function, that is we have $\phi(x, u; \beta)$ and a bounded set B such that there exists $\beta_0 \in B$ with $J(x, u) = \phi(x, u; \beta_0)$. Then, one way that the inverse decision making problem can be formulated is as

$$\begin{aligned} \hat{\beta} = \arg\min_{\beta}\ & 0 \\ \text{s.t. }\ & x_i^* = \arg\max_x \{\phi(x, u_i; \beta) \mid x \in \mathcal{X}(u_i)\} \\ & \beta \in B \end{aligned}$$

This feasibility problem is difficult to solve because it has an atypical constraint: The constraint that x_i^* be the maximizer to some optimization problem cannot be directly handled by nonlinear programming techniques.