

Process-Conditioned Investing with Incomplete Information using Maximum Causal Entropy

Brian D. Ziebart

*Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA,
bziebart@cs.cmu.edu*

Abstract. Investing to optimally maximize the growth rate of wealth based on sequences of event outcomes has many information-theoretic interpretations. Namely, the mutual information characterizes the benefit of additional side information being available when making investment decisions [1] in settings where the probabilistic relationships between side information and event outcomes are known. Additionally, the relative variant of the principle of maximum entropy [2] provides the optimal investment allocation in the more general setting where the relationships between side information and event outcomes are only partially known [3]. In this paper, we build upon recent work characterizing the growth rates of investment in settings with inter-dependent side information and event outcome sequences [4]. We consider the extension to settings with inter-dependent event outcomes and side information where the probabilistic relationships between side information and event outcomes are only partially known. We introduce the principle of minimum relative causal entropy to obtain the optimal worst-case investment allocations for this setting. We present efficient algorithms for obtaining these investment allocations using convex optimization techniques and dynamic programming that illustrates a close connection to optimal control theory.

Keywords: Causal entropy, directed information, Kelly criterion, investing, gambling

INTRODUCTION

Though most famous for its applicability to problems in communications, information theory [5] plays a key role in a number of different fields. For investment purposes, Kelly showed that when gambling on independent repeated events, $x \in \mathcal{X}$, the mutual information $I(X; Y) = H(X|Y) - H(X)$ characterizes the increase in the expected growth rate of wealth when side information $y \in \mathcal{Y}$ related to event outcomes is available [1]. The **Kelly criterion**, as the resulting investment allocation is known, and related Kelly-inspired criteria are commonly employed for investment and gambling. Permuter *et al.*'s recent work [4] makes the ideas of Kelly investment applicable to more complex settings with sequentially inter-dependent event outcomes $P(\mathbf{X}|\mathbf{Y}) = \prod_{t=1}^T P(X_t|\mathbf{X}_{1:t-1}, Y_t)$ for event outcome sequences \mathbf{X} and side information sequences \mathbf{Y} using the Marko-Massey theory of directed information [6, 7].

For statistical estimation tasks, the **principle of maximum entropy** [2] prescribes the “least committed” distribution consistent with partial constraints, using Shannon’s information entropy [5], $H(X) = \mathbb{E}_{P(X)}[-\log P(x)]$, to measure the “committedness” of a distribution. The maximum entropy distributions for joint, conditional, and marginal distributions with moment-matching constraints correspond to many state-of-the-art machine learning techniques, including Markov random fields and conditional random fields [8]. The **principle of maximum causal entropy** [9] extends the maximum entropy approach

to settings with information revelation and feedback using directed information theory [6, 7] with applications to single [9] and multi-agent decision prediction [10].

A powerful justification for the maximum entropy approach is that it provides an estimate that minimizes the worst-case log-loss. Since the growth rate of an investment is directly related to the log-loss of event outcome beliefs, this makes the relative variant of maximum entropy applicable for maximizing the worst-case growth rate in the Kelly investment setting with uncertain relationships between side information and event outcomes [3]. This supports the many Kelly criterion practitioners who address model uncertainty in an ad-hoc manner by moderating the allocations suggested by *e.g.*, reducing the difference of the allocation from a neutral allocation by a factor of two—the **half-Kelly criterion**. However, it also offers a more principled treatment for model uncertainty in these investment settings.

In this work, we extend the maximum causal entropy approach [9] to model uncertainty in sequential, inter-dependent investment settings. This contributions can alternately be viewed as extending the maximum entropy approach for uncertain Kelly investment [3] to causal information settings or generalizing the causal gambling work of Permuter *et. al.* [4] to incorporate model uncertainty. Our main result is a relative causal entropy formulation of the sequential maximum growth rate investment task that assures belief in the available odds except when in conflict with partial beliefs relating outcomes and side information. We describe efficient algorithms to calculate the growth-rate optimal investment allocation in this setting using convex optimization and dynamic programming techniques.

BACKGROUND

We first review background information on sequential investment and directed information theory for quantifying expected growth rates of investment in settings where the probability distributions governing event outcomes are known.

Sequential Events, Odds, and Growth Rates

We are interested in temporal sequences of events, $\mathbf{X} = (X_1, X_2, \dots, X_T)$, which are random variables that take on values, $\mathbf{x} = (x_1, x_2, \dots, x_T)$, from a set of event sequences, $\mathcal{X} = (\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_T)$. A probability distribution, $P(\mathbf{X}) \in \Delta_{\mathcal{X}}$, specifies the probability that each event sequence, $\mathbf{x} \in \mathcal{X}$, will occur. The odds available for a particular sequence of outcomes \mathbf{x} is denoted as $o(\mathbf{x})$ and indicates the multiplier by which a wagered amount will grow if outcome sequence \mathbf{x} occurs. This formulation is quite general; it allows both event probabilities and odds to be dependent on the past: $P(\mathbf{X}) = \prod_{t=1}^T P(X_t|X_1, \dots, X_{t-1})$ and $o(\mathbf{X}) = \prod_{t=1}^T o(X_t|X_1, \dots, X_{t-1})$. We assume that at each point in time during the sequence, all assets must be allocated to event outcomes as bet distributions, $b(\mathbf{X}) \in \Delta_{\mathcal{X}}$ over the simplex (denoted as Δ), specifying the portion of capital invested on each of the event outcome sequences $\mathbf{x} \in \mathcal{X}$. Further, we assume that odds belong to a book with fair odds, meaning that the odds-maker does not receive

a “take” from all the bets¹ *i.e.*, $\sum_{\mathbf{x} \in \mathcal{X}} \frac{1}{o(\mathbf{x})} = 1$. The inverse of these odds, $P_{\text{odds}} = \frac{1}{o(\mathbf{x})}$, can be viewed as the odds-marker’s belief in event outcomes.

The expected return when betting according to distribution $b(\mathbf{X})$ is:

$$R(b(\mathbf{X})) = \sum_{\mathbf{x} \in \mathcal{X}} P(\mathbf{x}) o(\mathbf{x}) b(\mathbf{x}). \quad (1)$$

A betting allocation that maximizes this expected return, $b_{\text{MER}}^*(\mathbf{X}) = \text{argmax}_{b(\mathbf{X})} R(b(\mathbf{X}))$, wagers everything on a single sequence of events, $\mathbf{x}_{\text{MER}}^*$. Unfortunately, due to the *gambler’s paradox*, the probability of losing all money, $1 - P(\mathbf{x}_{\text{MER}}^*)$, converges to one over infinite time horizons with uncertain events.

A less risky objective is to maximize the expected growth rate of investment,

$$W(b(\mathbf{X})) = \sum_{\mathbf{x} \in \mathcal{X}} P(\mathbf{x}) \log(o(\mathbf{x}) b(\mathbf{x})) \quad (2)$$

$$= D_{\text{KL}}(P, o^{-1}) - D_{\text{KL}}(P, b), \quad (3)$$

which greatly penalizes the loss of all capital. As shown by Equation 3 (from Cover & Thomas [11]), the growth rate can be equivalently expressed as the differences in **relative entropy** (also known as Kullback-Leibler divergence [12]) $D_{\text{KL}}(P, Q) = \sum_{x \in \mathcal{X}} P(x) \log \frac{P(x)}{Q(x)}$ between event probabilities and the (inverse) available odds and event probabilities and the bet allocations (Equation 3). Thus, the bettor must have a better estimate of P to realize any expected growth.

As a historical note, the growth rate optimizing objective (Equation 2) is equivalent to the log-scaled utility for increased wealth employed by Bernoulli to address the St. Petersburg paradox [13], which describes a game that, despite having infinite expected payoff, does not appear to be worth an arbitrary amount to play.

Theorem 1 (from Kelly [1]) *The optimal allocation of bets for maximizing the expected growth rate (Equation 2) when $P(\mathbf{x})$ is known exactly is*

$$\forall \mathbf{x} \in \mathcal{X}, b_{\text{MEG}}^*(\mathbf{x}) = P(\mathbf{x}), \quad (4)$$

with an expected rate of growth of:

$$W(b_{\text{MEG}}^*(\mathbf{X})) = \left(\sum_{\mathbf{x} \in \mathcal{X}} P(\mathbf{x}) \log o(\mathbf{x}) \right) - H(\mathbf{X}), \quad (5)$$

where $H(\mathbf{X}) = -\sum_{x \in \mathcal{X}} P(x) \log P(x)$ is Shannon’s information entropy [5].

The bet allocations of Theorem 1 are chosen independently from the outcome odds, $o(\mathbf{X})$. That is, growth rate-optimal investors with perfect information should bet in proportion to their beliefs and take no consideration of the payoff odds.

¹ Unfair books are of practical interest and were also investigated by Kelly [1], but they lack the information-theoretic interpretations of fair odds that we investigate in this work.

Side Information and Directed Information Theory

In many settings, bets can be based on available side information $\mathbf{y} = (y_1, y_2, \dots, y_T)$, which take on values from set of symbols $\mathcal{Y} = (\mathcal{Y}_1, \mathcal{Y}_2, \dots, \mathcal{Y}_T)$.

Theorem 2 (from Kelly [1]) *When event outcomes are independent, $\forall i, j, X_i \perp X_j$, and known to follow distribution $P(\mathbf{X}|\mathbf{Y}) = \prod_t P(X_t|Y_t)$, the growth-rate optimal bet allocation is in proportion to the conditional probability distribution:*

$$\forall \mathbf{x} \in \mathcal{X}, b_{MEG}^*(\mathbf{x}|\mathbf{y}) = P(\mathbf{x}|\mathbf{y}).$$

The expected growth rate of the growth-rate optimal betting allocation is:

$$W(b_{MEG}^*(\mathbf{X}|\mathbf{Y})) = \left(\sum_{\mathbf{x} \in \mathcal{X}} P(\mathbf{x}) \log o(\mathbf{x}) \right) - H(\mathbf{X}|\mathbf{Y}). \quad (6)$$

The benefit of having side information is directly related to the mutual information between event outcomes and side information:

$$\begin{aligned} W(b_{MEG}^*(\mathbf{X}|\mathbf{Y})) - W(b_{MEG}^*(\mathbf{X})) &= H(\mathbf{X}|\mathbf{Y}) - H(\mathbf{X}) \\ &= I(\mathbf{X}; \mathbf{Y}). \end{aligned} \quad (7)$$

The independence assumptions of Theorem 2 correspond to the structural relationship model of Figure 1a. More generally, event outcomes can be inter-dependent and the values of revealed side information variables may depend on the outcomes of earlier events, violating that independence assumption. Directed information theory measures are needed to appropriately represent the sequential availability of side information in these settings. The **causally conditioned probability** of \mathbf{X} given \mathbf{Y} from the Marko-Massey theory of directed information [6, 7] is:

$$P(\mathbf{X}^T || \mathbf{Y}^T) = \prod_{t=1}^T P(X_t | \mathbf{Y}_{1:t}, \mathbf{X}_{1:t-1}). \quad (8)$$

It crucially differs from the **conditional probability**, $P(\mathbf{X}|\mathbf{Y}) = \prod_{t=1}^T P(X_t | \mathbf{Y}_{1:T}, \mathbf{X}_{1:t-1})$, in how the \mathbf{Y} variables are conditioned upon. In the causally conditioned probability, each X_t is only conditioned on previously occurring side information variables, $\mathbf{Y}_{1:t}$. In the conditional probability distribution, each X_t is conditioned on all side information, $\mathbf{Y}_{1:T}$.

The **causal entropy** measures the uncertainty corresponding to the causally conditioned probability distribution. It is:

$$H(\mathbf{X}^T || \mathbf{Y}^T) = \mathbb{E}_{P(\mathbf{X}, \mathbf{Y})} [-\log P(\mathbf{X}^T || \mathbf{Y}^T)]. \quad (9)$$

In the sequential side information setting, future side information is not available since side information is revealed only after bet allocations on earlier events have been made. Thus, directed information rather than conditional information measures must be employed.




FIGURE 1. Structural relationships between sequential side information and event outcomes. (a.) The independent relationship structure investigated by Kelly [1]. (b.) A structure with dependence between event outcomes over time. (c.) A more general side information process and event outcome structure investigated by Permuter *et. al.* [4].

Theorem 3 (from Permuter *et al.* [4]) *The maximum expected growth rate investment allocation is a causally conditioned quantity:*

$$\forall \mathbf{x}, b_{MEG}^*(\mathbf{x}|\mathbf{y}) = P(\mathbf{x}|\mathbf{y}),$$

with expected growth rate:

$$W(b_{MEG}^*(\mathbf{X}|\mathbf{Y})) = \left(\sum_{\mathbf{y} \in \mathcal{Y}, \mathbf{x} \in \mathcal{X}} P(\mathbf{y}) P(\mathbf{x}|\mathbf{y}) \log o(\mathbf{x}|\mathbf{y}) \right) - H(\mathbf{X}|\mathbf{Y}). \quad (10)$$

When the odds are independent of side information, $o(\mathbf{x}|\mathbf{y}) = o(\mathbf{x})$, this reduces to:

$$W(b_{MEG}^*(\mathbf{X}|\mathbf{Y})) = \left(\sum_{\mathbf{x} \in \mathcal{X}} P(\mathbf{x}) \log o(\mathbf{x}) \right) - H(\mathbf{X}|\mathbf{Y}). \quad (11)$$

with the increase in growth rate by having causally available side information being the directed information [6, 7] $I(\mathbf{Y} \rightarrow \mathbf{X}) = H(\mathbf{X}|\mathbf{Y}) - H(\mathbf{X})$.

This formulation generalizes the setting considered by Kelly (Figure 1a), allowing dependencies between event outcomes (Figure 1b) and more general side information processes that have side information that is dependent on previous side information and event outcomes, as shown in Figure 1c.

SEQUENTIAL INVESTING WITH INCOMPLETE SIDE INFORMATION

Apart from well-defined games of chance (*e.g.*, dice, cards), investments must typically be made without precise knowledge of the distributions of event outcomes and relations to side information. How should investment allocations be made in these incomplete information settings? Following Grünwald & David [3], who suggest maximizing the worst-case growth rates of investment, we extend maximum worst-case growth rate investment to the sequential side information setting with Theorem 4.

Theorem 4 *Maximizing the worst-case sequential investment growth subject to constraints (Equation 13) is equivalent to minimizing the relative causal entropy—also known as the Kullback-Leibler divergence (**KL divergence**) [12]—between $P(\mathbf{X}||\mathbf{Y})$ and $P_{\text{odds}}(\mathbf{X}||\mathbf{Y}) = \frac{1}{o(\mathbf{x}||\mathbf{y})}$,*

$$D_{\text{KL}}(P(\mathbf{X}||\mathbf{Y}), P_{\text{odds}}(\mathbf{X}||\mathbf{Y})) = \mathbb{E}_{P(\mathbf{X}, \mathbf{Y})} \left[\log \frac{P(\mathbf{X}||\mathbf{Y})}{P_{\text{odds}}(\mathbf{X}||\mathbf{Y})} \right], \quad (12)$$

subject to constraints:

$$\begin{aligned} \forall i, f_i(P(\mathbf{X}||\mathbf{Y})) &= 0 \\ \forall j, g_j(P(\mathbf{X}||\mathbf{Y})) &\leq 0, \end{aligned} \quad (13)$$

with sets of linear functions f_i and g_j .

Proof. We transform the maximum worst-case sequential investment growth rate optimization’s objective to prove equivalence. The original objective function can be viewed as a maximin problem where “nature” adversarially chooses $P(\mathbf{X}||\mathbf{Y})$ to minimize expected growth rate after the belief, $b(\mathbf{X}||\mathbf{Y})$, is selected.

$$\begin{aligned} & \max_{b(\mathbf{X}||\mathbf{Y})} \min_{P(\mathbf{X}||\mathbf{Y})} \sum_{\mathbf{y} \in \mathcal{Y}, \mathbf{x} \in \mathcal{X}} P(\mathbf{y}, \mathbf{x}) \log(b(\mathbf{x}||\mathbf{y})o(\mathbf{x}||\mathbf{y})) && \text{(def. max. worst-case growth)} \\ &= \min_{P(\mathbf{X}||\mathbf{Y})} \max_{b(\mathbf{X}||\mathbf{Y})} \sum_{\mathbf{y} \in \mathcal{Y}, \mathbf{x} \in \mathcal{X}} P(\mathbf{y}, \mathbf{x}) \log(b(\mathbf{x}||\mathbf{y})o(\mathbf{x}||\mathbf{y})) && \text{(Minimax duality)} \\ &= \min_{P(\mathbf{X}||\mathbf{Y})} \sum_{\mathbf{y} \in \mathcal{Y}, \mathbf{x} \in \mathcal{X}} P(\mathbf{y}, \mathbf{x}) \log(P(\mathbf{x}||\mathbf{y})o(\mathbf{x}||\mathbf{y})) && (b(\mathbf{X}||\mathbf{Y}) = P(\mathbf{X}||\mathbf{Y}) \text{ at optimum}) \\ &= \min_{P(\mathbf{X}||\mathbf{Y})} \sum_{\mathbf{y} \in \mathcal{Y}, \mathbf{x} \in \mathcal{X}} P(\mathbf{y}, \mathbf{x}) \log \frac{P(\mathbf{x}||\mathbf{y})}{P_{\text{odds}}(\mathbf{x}||\mathbf{y})} && \text{(def. } P_{\text{odds}}(\mathbf{x}||\mathbf{y})) \\ &= \min_{P(\mathbf{X}||\mathbf{Y})} D_{\text{KL}}(P(\mathbf{X}||\mathbf{Y}), P_{\text{odds}}(\mathbf{X}||\mathbf{Y})) && \text{(def. KL-divergence)} \end{aligned}$$

□

The information-theoretic interpretation is that bet allocations should be as close as possible to the beliefs implied by the odds, except where constraints (Equation 13) guarantee that the bet allocation is a better estimate of outcomes $P(\mathbf{X}||\mathbf{Y})$ than the odds.

The Principle of Minimum Relative Causal Entropy: Duality

Directly optimizing the primal (Equation 12) in terms of causally conditioned probability terms to find the minimum worst-case investment allocation is possible, but it becomes computationally intractable for large time horizons. When knowledge relating event outcomes and side information linearly decomposes over time, more efficient optimization is possible. Building upon our past work [9], we can investigate and interpret the form of the dual optimization.

Definition 1 (Markovian minimum relative causal entropy) *Given Markovian side information dynamics, $P(Y_t|Y_{t-1}, X_{t-1})$, and knowledge constraining the expectations*

of a vector of statistics $\mathbf{f}: \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^K$ to be (component-wise) equivalent² to some $C \in \mathbb{R}^K$, the maximum worst-case sequential investment growth allocation is defined by the following optimization:

$$\begin{aligned} & \min_{P(\mathbf{X}|\mathbf{Y})} D_{KL}(P(\mathbf{X}|\mathbf{Y}), P_{\text{odds}}(\mathbf{X}|\mathbf{Y})) \\ & \text{such that: } \mathbb{E}_{P(\mathbf{X}, \mathbf{Y})} \left[\sum_{t=1}^T \mathbf{f}(x_t, y_t) \right] = C \text{ and} \\ & \forall \mathbf{x} \in \mathcal{X}, \mathbf{y} \in \mathcal{Y} \ P(\mathbf{x}, \mathbf{y}) = P(\mathbf{x}|\mathbf{y}) \prod_{t=1}^T P(y_t | y_{t-1}, x_{t-1}). \end{aligned} \tag{14}$$

By limiting the constraints to additively decompose in terms of statistics over time, a Markovian form for the distribution is obtained in the dual (Theorem 5).

Theorem 5 *The solution distribution for the Markovian minimum relative causal entropy optimization (Definition 1) is:*

$$\begin{aligned} V(y_t) &= \text{relative-softmax}_{x_t \in \mathcal{X}} (Q(x_t, y_t), P_{\text{odds}}(x_t | y_t, x_{t-1})) \\ Q(x_t, y_t) &= \mathbb{E}_{P(Y_{t+1} | x_t, y_t)} [V(y_{t+1}) | x_t, y_t] + \text{Reward}_{\theta}(x_t, y_t), \end{aligned}$$

where: $\text{relative-softmax}_{x \in \mathcal{X}}(f(x), P(x)) = \log \sum_{x \in \mathcal{X}} P(x) e^{f(x)}$ and $\text{Reward}_{\theta}(x, y) = \theta^T \mathbf{f}(x, y)$ with the Lagrange multipliers θ being fit from data to match the constraints of Definition 1. The optimal conditional bet allocation is then: $b(x_t | y_t) = P_{\text{odds}}(x_t | y_t, x_{t-1}) e^{Q(x_t, y_t) - V(y_t)}$.

This recursive definition for the bet allocation matches $P_{\text{odds}}(x_t | y_t)$ when no reward signal is present. It can be interpreted as a softened, relative relaxation of the value iteration algorithm for obtaining the expected future value of each state, $s \in \mathcal{S}$ and action $a \in \mathcal{A}$ for the optimal policy (and the optimal policy itself, $\pi: \mathcal{S} \rightarrow \mathcal{A}$) in a decision process with given dynamics $P(S_{t+1} | a_t, s_t)$:

$$\begin{aligned} V(s_t) &= \max_{a_t \in \mathcal{A}} (Q(a_t, s_t)) \\ Q(a_t, s_t) &= \mathbb{E}_{P(S_{t+1} | a_t, s_t)} [V(s_{t+1}) | a_t, s_t] + \text{Reward}(a_t, s_t). \end{aligned}$$

Under uniform payoff odds, optimizing the relative causal entropy objective of Theorem 5 is equivalent to maximizing the causal entropy (Equation 9). In that special case, this distribution has been employed to forecast sequential decisions in stochastic decision processes based on previously observed decision sequences [9]. Indeed, the proof of Theorem 5 can be obtained by extending the proof for the maximum causal entropy setting using the relative entropy.

² Inequality constraints of linear functions (and, more generally, of convex functions) are also possible. They lead to a similar dual form with half-space constraints on Lagrange parameters.

Algorithm 1 Find the worst-case optimal growth rate betting allocation.

Require: Time horizon T , side information dynamics $P(y_t|x_{t-1}, y_{t-1})$, odds $o(X||Y)$.

Ensure: $b(x_t|y_t, x_{t-1})$ is worst-case growth rate optimal.

$t \leftarrow 0$

while $\mathbb{E}_{P(\mathbf{X}, \mathbf{Y})}[\sum_{t=1}^T \mathbf{f}(x_t, y_t)] - C \neq 0$ **do**

 Compute $P_\theta(X_t|Y_t, X_{t-1})$ from the equations of Theorem 5

 Compute $\mathbb{E}_{P(\mathbf{X}, \mathbf{Y})}[\sum_{t=1}^T \mathbf{f}(x_t, y_t)]$ from $P_\theta(X_t|X_{t-1}, Y_t)$ and $P(y_t|x_{t-1}, y_{t-1})$

 Update parameters: $\theta \leftarrow \theta - \lambda_t(\mathbb{E}_{P(\mathbf{X}, \mathbf{Y})}[\sum_{t=1}^T \mathbf{f}(x_t, y_t)] - C)$

end while

Return $b(X_t|Y_t, X_{t-1}) = P_\theta(X_t|Y_t, X_{t-1})$

The dual parameters θ and corresponding investment allocation are obtained according to gradient descent (Algorithm 1) using a slowly time-decaying learning rate λ_t . In practice, only approximate convergence may be efficiently realized.

CONCLUSIONS

In this work, we investigated the investment setting with side information processes and incomplete knowledge of how outcome events relate to side information. Due to this uncertainty, we pursued maximum worst-case growth rate investing. We showed that this formulation is equivalent to maximizing the relative entropy of a causally conditioned distribution subject to available partial information constraints, and that the corresponding algorithm for determining bet allocations uses a relaxation of standard optimal control procedures and convex optimization.

REFERENCES

1. J. Kelly Jr, *IRE Transactions on Information Theory* **2**, 185–189 (1956).
2. E. T. Jaynes, *Physical Review* **106**, 620–630 (1957).
3. P. D. Grünwald, and A. P. Dawid, *Annals of Statistics* **32**, 1367–1433 (2003).
4. H. H. Permuter, Y.-H. Kim, and T. Weissman, “On Directed Information and Gambling,” in *Proc. IEEE International Symposium on Information Theory*, 2008, pp. 1403–1407.
5. C. E. Shannon, *Bell System Technical Journal* **27** (1948).
6. H. Marko, “The Bidirectional Communication Theory – A Generalization of Information Theory,” in *IEEE Transactions on Communications*, 1973, pp. 1345–1351.
7. J. L. Massey, “Causality, Feedback and Directed Information,” in *Proc. IEEE International Symposium on Information Theory and Its Applications*, 1990, pp. 27–30.
8. J. Lafferty, A. McCallum, and F. Pereira, “Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data,” in *Proc. ICML*, 2001, pp. 282–289.
9. B. D. Ziebart, J. A. Bagnell, and A. K. Dey, “Modeling Interaction via the Principle of Maximum Causal Entropy,” in *Proc. International Conf. on Machine Learning*, 2010.
10. B. D. Ziebart, J. A. Bagnell, and A. K. Dey, “Maximum Causal Entropy Correlated Equilibria for Markov Games,” in *Proc. International Conf. on Autonomous Agents and Multiagent Systems*, 2011.
11. T. Cover, and J. Thomas, *Elements of information theory*, John Wiley and sons, 2006.
12. S. Kullback, and R. A. Leibler, *Annals of Mathematical Statistics* **22**, 49–86 (1951).
13. D. Bernoulli, *Econometrica: Journal of the Econometric Society* pp. 23–36 (1954).