
Game-Theoretic Algorithms for Conditional Moment Matching

Gokul Swamy¹ Sanjiban Choudhury² J. Andrew Bagnell^{2,1} Zhiwei Steven Wu³

Abstract

A variety of problems in econometrics and machine learning, including instrumental variable regression and Bellman residual minimization, can be formulated as satisfying a set of conditional moment restrictions (CMR). We derive a general, game-theoretic strategy for satisfying CMR that scales to nonlinear problems, is amenable to gradient-based optimization, and is able to account for finite sample uncertainty. We recover the approaches of (Dikkala et al., 2020) and (Dai et al., 2018) as special cases of our general framework before detailing various extensions and how to efficiently solve the game defined by CMR.

1. Introduction

Let X , Y , and Z be random variables on (potentially non-finite) sample spaces \mathcal{X} , \mathcal{Y} , and \mathcal{Z} . We are interested in finding a function $h \in \mathcal{H}$ that satisfies a set of *conditional moment restrictions* (Chamberlain, 1987),

$$\forall z \in \mathcal{Z}, \mathbb{E}[Y|z] = \mathbb{E}[h(X)|z], \quad (1)$$

or CMR for short. While this problem might seem a bit abstract, it is at the core of two disparate problems in machine learning: *instrumental variable regression* and *Bellman-residual minimization*. We give a brief introduction to each before presenting a unified method for solving for consistent h from finite samples that elegantly scales to nonlinear problems. Throughout, we assume we optimize over a class \mathcal{H} that is convex, compact, closed under negation, and of finite Rademacher complexity.

Our key insight is that *we can formulate conditional moment matching as a zero-sum game, allowing us to both eliminate double sample issues and explicitly reason about the effects of constraint relaxation*. We call this family of techniques *estimation via relaxation*.

¹Robotics Institute, Carnegie Mellon University ²Aurora Innovation ³Institute for Software Research, Carnegie Mellon University. Correspondence to: Gokul Swamy <gswamy@cmu.edu>.

1.1. Instrumental Variable Regression

Let's assume that X , Y , and Z have the following dependency structure:




Figure 1. The graphical model considered in instrumental variable regression. We are interested in finding h , the causal relationship from X to Y , even though there is an unobserved confounder, U .

We refer to X as the *treatment* and Y as the *response* or *outcome*. Given a dataset of (x, y, z) tuples, we are interested in determining the causal relationship between X and Y , $\mathbb{E}[Y|\text{do}(x)]$. In the above graphical model, this is equivalent to determining h . Because of the presence of an unobserved confounder, U , that affects both X and Y , standard regression (e.g. Ordinary Least Squares or OLS) can produce inconsistent estimates. If we only have observational data and are unable to perform randomized control trials, a canonical technique to recover h is Instrumental Variable Regression (IVR) (Winship & Morgan, 1999). Formally, an instrument Z must satisfy:

1. *Unconfounded Instrument*: $Z \perp\!\!\!\perp U$ – i.e. “independent” randomization from instrument.
2. *Exclusion*: $Z \perp\!\!\!\perp Y|X$ – i.e. no extraneous paths.
3. *Relevance*: $Z \not\perp\!\!\!\perp X$ – i.e. conditioning has an effect.

We note that the unconfounded instrument and exclusion conditions are structural assumptions that cannot be checked from observational data. Assuming access to such a Z and linear relationships between all variables, one can recover $h(x) = \beta x$ by computing $\beta = \mathbb{E}[ZX]^{-1}\mathbb{E}[ZY]$. Notice that the recovered h satisfies the CMR (1). With a bit of linear algebra, one can show this calculation is equivalent to a Two-Stage Ordinary Least Squares (2SLS) procedure, for which one first regresses from Z to X and then regresses from the predicted \hat{X} to Y , returning the latter coefficients.

We focus on the general, nonlinear problem, which can still be formulated in terms of CMRs. We assume that noise

U enters additively to Y ¹, and write out the following equations:

$$X = g(Z, U), \quad (2)$$

$$Y = h(X) + U. \quad (3)$$

Without loss of generality, we assume that $\mathbb{E}[U] = 0$. As is standard in the causal inference literature (Nie et al., 2020), we assume that our collected data satisfies consistency of potential outcomes (i.e. the data generating process did not change since our dataset was collected (Rehkopf et al., 2016)) and overlap (i.e. we have data for all X we care about (Imbens, 2000)). Perhaps the most natural way to tackle this problem would be to minimize some notion of squared error between the two sides of the CMR:

$$\min_{h \in \mathcal{H}} \mathbb{E}_z [\mathbb{E}[Y - h(x)|z]^2]. \quad (4)$$

For nonlinear problems, gradient-based optimization is a common and scalable technique. Unfortunately, differentiating the squared expectation in the preceding expression leads to a “double sample” issue in which one needs two independent samples from $P(X|z)$ to compute each gradient step, which can be quite challenging outside of simulated environments. To resolve this issue, (Dikkala et al., 2020) propose instead solving the following zero-sum game:

$$\min_{h \in \mathcal{H}} \max_f \mathbb{E}[2(Y - h(X))f(Z) - f(Z)^2], \quad (5)$$

which they prove produces consistent estimates of h and does not involve a squared expectation. We show in Sec. 2 that this is an example of a general technique for satisfying CMR.

1.2. Bellman Residual Minimization

Consider a Markov Decision Process (MDP) parameterized by $\langle \mathcal{S}, \mathcal{A}, \mathcal{T}, r, \gamma \rangle$, where \mathcal{S} is the state space, \mathcal{A} is the action space, $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ is the transition operator, $r : \mathcal{S} \times \mathcal{A} \rightarrow [-1, 1]$ is a reward function, and γ is the discount factor. It is a well-known fact that one can find the optimal policy, π , for this MDP by first computing the optimal *value function* via the Bellman Equation,

$$V(s) = \max_{a \in \mathcal{A}} r(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{T}(s, a)} [V(s')], \quad (6)$$

and choosing actions greedily (Puterman, 2014):

$$\pi(s) = \arg \max_{a \in \mathcal{A}} r(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{T}(s, a)} [V(s')]. \quad (7)$$

We can combine these two expressions into one that is satisfied by the optimal (V, π) pair:

$$V(s) = \mathbb{E}_{\substack{a \sim \pi(s) \\ s' \sim \mathcal{T}(s, a)}} [r(s, a) + \gamma V(s')]. \quad (8)$$

¹This is an assumption inherited from standard regression. Without it, one can only bound the treatment effect (Kilbertus et al., 2020).

When one cannot simply enumerate all states and actions to perform either policy or value iteration (Russell & Norvig, 2002), a standard technique is *Bellman residual minimization* – minimizing the difference between the two sides of (8). The updated value function can then be used to compute a more optimal policy via the greedy optimization of (7). We can write (8) in the template of CMR by setting

$$Z = s, a \quad (9)$$

$$X = s, h(X) = V(s) \quad (10)$$

$$Y = r(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{T}(s, a)} [V(s')]. \quad (11)$$

Similar to (4), one could attempt to satisfy the CMR by minimizing the expected squared Bellman error:

$$\min_{V: \mathcal{S} \rightarrow \mathbb{R}} \mathbb{E}[\mathbb{E}[V(s) - (r(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{T}(s, a)} [V(s')])|s, a]^2], \quad (12)$$

which also has double sample issues (Baird, 1995). Recently, (Dai et al., 2018) suggested that one could avoid double sample issues by instead solving the following zero-sum game:²

$$\min_{V: \mathcal{S} \rightarrow \mathbb{R}} \max_{f: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}} \mathbb{E}_{s, a, s'} [2(r(s, a) + V(s') - V(s))f(s, a) - f(s, a)^2], \quad (13)$$

which they derive via an appeal to convex conjugates (Boyd et al., 2004). We present an alternative construction of this objective which helps elucidate the properties of the algorithm proposed by (Dai et al., 2018).

2. Estimation via Relaxation

SETTING	Z	X AND $h(X)$	Y
IVR	INSTRUMENT	TREATMENT	OUTCOME
RL	(s, a)	s AND $V(s)$	$r(s, a) + \gamma \mathbb{E}_{s'} [V(s')]$

Table 1. The work of both (Dikkala et al., 2020) and (Dai et al., 2018) can be seen as examples of a more general template for satisfying conditional moment restrictions from finite samples and for nonlinear problems.

Throughout this section, we use **bold fonts** to designate vectors. To aid us in our quest to satisfy the conditional moment restrictions, we are given access to N samples from the joint distribution of X, Y , and Z . Because we have finite samples and can therefore only estimate conditional expectations up to some tolerance, it is natural to relax the CMRs to

$$\min_{h \in \mathcal{H}, \delta} J(\delta) \quad \text{s.t.} \quad |\mathbb{E}[Y - h(X)|z]| \leq \delta_z \quad \forall z \in \mathcal{Z}, \quad (14)$$

where the δ_z are slack variables and J is some convex function of δ that keeps the slacks from getting too large. We

²For simplicity, we drop the entropy regularization in their objective. Our derivations can be easily extended.

note that (14) is reminiscent of the classical framework of Tikhinov regularization before presenting some interesting properties that result from setting $J(\delta) = \frac{1}{2}\mathbb{E}_z[\delta_z^2]$. The Lagrangian with a $P(z)$ -weighted inner product is

$$L(h, \delta, \lambda) = \sum_{z \in \mathcal{Z}} P(z) \lambda_z (\mathbb{E}[Y - h(X)|z] - \delta_z) + P(z) \frac{1}{2} \delta_z^2, \quad (15)$$

where λ are the Lagrange multipliers.³ Applying the stationarity component of the KKT conditions, one arrives at

$$\nabla_{\delta_z} L(h, \delta, \lambda) = -P(z) \lambda_z + P(z) \delta_z = 0, \quad (16)$$

implying that $\delta_z = \lambda_z$. Plugging this back into the Lagrangian, we can simplify our function to

$$L(h, \lambda) = \sum_{z \in \mathcal{Z}} P(z) \lambda_z \mathbb{E}[Y - h(X)|z] - P(z) \frac{1}{2} \lambda_z^2. \quad (17)$$

We refer to (17) as the *Regularized Lagrangian* or ReLa for short. Now, solving for the optimal Lagrange multipliers via stationarity, we arrive at

$$\nabla_{\lambda_z} L(h, \lambda) = P(z) \mathbb{E}[Y - h(X)|z] - P(z) \lambda_z = 0, \quad (18)$$

which implies the equilibrium λ_z is equal to $\mathbb{E}[Y - h(X)|z]$. Plugging this back into (17) recovers function

$$L(h) = \sum_{z \in \mathcal{Z}} P(z) \mathbb{E}[Y - h(X)|z]^2. \quad (19)$$

Thus, in the population limit, we are optimizing the conditional MSE of (Chen & Pouzo, 2012), leading to consistent estimates.

2.1. Generative Modeling Approach

Perhaps the most immediate way to minimize (19) over $h \in \mathcal{H}$ would be to minimize the empirical MSE,

$$\min_{h \in \mathcal{H}} \frac{1}{N} \sum_i (y_i - h(x_i))^2. \quad (20)$$

Unfortunately, this only gives us a function that matches unconditional moments (i.e. $\mathbb{E}[Y] = \mathbb{E}[h(X)]$). For the IVR setting, this would give us the inconsistent, OLS-like estimates of h we are explicitly trying to avoid. For the MDP setting, this would not produce a valid value function.

Instead, one could learn the distribution $P(X|z) = g(z)$ and pass samples from it to a candidate h , ensuring one is attempting to match the conditional moments. Intuitively, this is the generalization of the 2SLS procedure to nonlinear

³While written above in terms of finite z , our derivation easily extend to infinite sets.

functions. Because the second stage is nonlinear, one cannot simply compute the first moment of the $P(X|z)$ distribution (which is recovered by linearly regressing from X to Z in the 2SLS procedure). To see this, consider \mathcal{H} being the set of quadratic functions of X . Then,

$$\mathbb{E}[h(X)|z] = \sum_{x \in \mathcal{X}} p(x|z) h(x) \quad (21)$$

$$= \sum_{x \in \mathcal{X}} p(x|z) (ax^2 + bx + c) \quad (22)$$

$$= a\mathbb{E}[X^2|z] + b\mathbb{E}[X|z] + c. \quad (23)$$

Here, one needs to have access to $\mathbb{E}[X^2|z]$ to check the CMR. For more complex \mathcal{H} , one therefore needs to learn the entire $P(X|z)$. This kind of approach was first proposed for the IVR setting by (Hartford et al., 2017) and amounts to first learning a $g(z)$ via maximum likelihood estimation and then solving

$$\min_{h \in \mathcal{H}} \frac{1}{N} \sum_i (y_i - \mathbb{E}_{\hat{x} \sim g(z_i)} [h(\hat{x})])^2. \quad (24)$$

Unfortunately, this approach suffers from the well-known “double-sample” issue where multiple independent samples from $g(z)$ are required to compute gradients of h . To see this, note that the gradient w.r.t. h of (24) is

$$\sum_i (y_i - \mathbb{E}_{\hat{x} \sim \hat{g}(z_i)} [h(\hat{x})]) (-\mathbb{E}_{\hat{x} \sim \hat{g}(z_i)} [\frac{\partial}{\partial h} h(\hat{x})]). \quad (25)$$

Approximating this gradient with a single sample,

$$\sum_i (\mathbb{E}_{\hat{x} \sim \hat{g}(z_i)} [(y_i - h(\hat{x})) \frac{\partial}{\partial h} h(\hat{x})]), \quad (26)$$

can produce biased estimates of the gradient that are inconsistent in the limit of infinite data. Additionally, learning an accurate first-stage model might be quite challenging for some problems.

2.2. Game-Theoretic Approach

Ideally, we would like to avoid the added complexity of learning a generative model and the double-sampling required for gradient based-optimization. One gets a two-for-one deal by instead solving the two-player zero-sum game with the ReLa (17) as the payoff. Denoting by $f \in \mathcal{F} \subseteq \{\mathcal{Z} \rightarrow \mathbb{R}\}$ the function that maps z 's to corresponding Lagrange multipliers, we can write this game as:

$$\min_{h \in \mathcal{H}} \max_{f \in \mathcal{F}} \mathbb{E}[2(Y - h(X))f(Z) - f(Z)^2]. \quad (27)$$

Notice the similarity of this expression to the objectives of (Dikkala et al., 2020) (5) and (Dai et al., 2018) (13) and

that there are no squared expectations of h and therefore no double sample issues. Additionally, one does not need to learn a generative model of $P(X|z)$ for these sorts of game-theoretic approaches. Under the assumption that $\mathbb{E}[Y - h(X)|z] \in \mathcal{F}$, finding a Nash equilibrium of this game corresponds to finding a $h \in \mathcal{H}$ that is close in an ℓ_2 sense to satisfying the CMR.

2.3. Takeaways Thus Far

We pause briefly to consider some of the important facts revealed to us by the above derivation:

1. The approaches of (Dikkala et al., 2020) and (Dai et al., 2018) are solving a relaxed CMR problem with a penalty on the ℓ_2 norm of the constraint violation. Thus, their consistency is in an ℓ_2 sense rather than a uniform, ℓ_∞ sense.
2. The choice of J allows us to control how slack is dispersed between the CMRs – derivations like the above allow us to explicitly consider the effects of the kind of regularization we choose.
3. The value of $f(z)$ (regardless of the choice of J) is exactly the amount of conditional moment mismatch the optimization procedure tolerates. If this value is particularly high for values of z where we have many samples, we might need to consider a more expressive \mathcal{H} . Thus, we can explicitly reason about the effects of finite-sample uncertainty in satisfying CMRs.

3. Extensions

We now briefly sketch two extensions to our above setup before diving into how to solve ReLa games efficiently.

3.1. Constraints on δ

A natural question after the above discussion might be how one could include information about sample uncertainty into the optimization procedure, rather than performing an after-the-fact check. Let n_z denote the number of samples we have for z . For simplicity, we assume $n_z > 0, \forall z \in \mathcal{Z}$. A Hoeffding bound tells us that we should expect a sample expectation $\mathbb{E}_{n_z}[\cdot|z]$ to be within $\propto \frac{1}{\sqrt{n_z}}$ of the population expectation. We refer to this sampling error as ϵ_z . Then, in expectation over Z ,

$$\mathbb{E}_Z[\epsilon_z^2] = \sum_{z \in \mathcal{Z}} P(z) \epsilon_z^2 \approx \sum_{z \in \mathcal{Z}} \frac{n_z}{N} \left(\frac{1}{\sqrt{n_z}}\right)^2 = \frac{|\mathcal{Z}|}{N} = \kappa(N). \quad (28)$$

Our relaxed optimization problem then becomes

$$\begin{aligned} \min_{h \in \mathcal{H}, \delta} & 0 \\ \text{s.t.} & |\mathbb{E}[Y - h(X)|z]| \leq \delta_z \quad \forall z \in \mathcal{Z} \\ & \mathbb{E}_Z[\delta_z^2] \leq \kappa(N). \end{aligned} \quad (29)$$

Algorithm 1 No-Regret Conditional Moment Matching

Input: Dataset \mathcal{D} of (x, y, z) tuples. No-regret algorithm over \mathcal{H} , Best-response oracle over \mathcal{F} , Threshold ϵ

Output: Causal effect of X on Y , h

Set $t = 1, h^t \in \mathcal{H}, f^t \in \mathcal{F}, L(h^t, f^t) = 2\epsilon$

while $L(h^t, f^t) > \epsilon$ **do**

$L(h^t, f^t) = \mathbb{E}_{(x,y,z) \sim \mathcal{D}}[2(y - h^t(x))f^t(z) - f^t(z)^2]$

No-regret alg. computes h^t over $L(\cdot, f^t)$ history.

Best-response computes $f^t = \arg \max_{f \in \mathcal{F}} L(h^t, \cdot)$.

$t \leftarrow t + 1$

end while

Return h^t .

This form, reminiscent of Ivanov regularization, takes into account sample uncertainty explicitly and could be solved via standard constrained optimization machinery like Augmented Lagrangians (Hestenes, 1969). We note that in practice, it is far more common to solve the Tikhonov-style problem (14) and scale J based on performance on some holdout data.

3.2. Regularization of h

Consider two solutions $(h_1, f_1), (h_2, f_2)$ of the ReLa Game (27) such that

$$|\mathbb{E}[Y - h_1(X)|z]| \leq f_1(z), \forall z \in \mathcal{Z}, \quad (30)$$

$$|\mathbb{E}[Y - h_2(X)|z]| \leq f_2(z), \forall z \in \mathcal{Z}, \quad (31)$$

$$\mathbb{E}_Z[f_1(z)^2] = \mathbb{E}_Z[f_2(z)^2]. \quad (32)$$

Our game-theoretic perspective wouldn't be able to break ties between h_1 and h_2 , regardless of desirable properties h_1 might have over h_2 like smoothness in X or being the max-margin classifier. This is because as long as two solutions exist within CMR-violation balls of the same size (e.g. in ℓ_2 norm), we consider them equally optimal. Thus, some sort of regularization on h , producing an optimization problem of the form

$$\begin{aligned} \min_{h \in \mathcal{H}, \delta} & J(\delta) + \alpha R(h) \\ \text{s.t.} & |\mathbb{E}[Y - h(X)|z]| \leq \delta_z \quad \forall z \in \mathcal{Z} \end{aligned} \quad (33)$$

with α tuned empirically on hold-out data might be helpful for finding h with additional desirable properties. Concretely, one could regularize to the OLS solution by setting $R(h) = \mathbb{E}[(h(X) - Y)^2]$, encouraging h to also match unconditional moments. Both (Dikkala et al., 2020) and (Dai et al., 2018) consider regularization on h (e.g. in $\|\cdot\|_{\mathcal{H}}^2$ norm or by entropy-regularizing the policy), which we identify as a form of tie-breaking.

3.3. Efficiently Solving the ReLa Game

Solving a two-player game like (27) can be done provably efficiently via a reduction to no-regret online learning, following the classic analysis of (Freund & Schapire, 1997).

For completeness, we provide such a procedure in Algorithm 1 and prove efficiency below.

Proof. Let $L(h, f) = \mathbb{E}[2(Y - h(X))f(Z) - f(Z)^2]$. Our no-regret assumption tells us that

$$\frac{1}{N} \sum_t L(h^t, f^t) - \frac{1}{N} \min_{h \in \mathcal{H}} \sum_t L(h, f^t) \leq \epsilon. \quad (34)$$

for some N that is $\text{poly}(\frac{1}{\epsilon})$. Under a realizability assumption, we can write that

$$\frac{1}{N} \sum_t L(h^t, f^t) \leq \epsilon + \frac{1}{N} \min_{h \in \mathcal{H}} \sum_t L(h, f^t) \leq \epsilon. \quad (35)$$

Then, utilizing the fact that there must be at least one element in an average that is at most the value of the average,

$$\min_t L(h^t, f^t) \leq \frac{1}{N} \sum_t L(h^t, f^t) \leq \epsilon. \quad (36)$$

To complete the proof, we recall that f^t is chosen as the best response to h^t in Algorithm 1, giving us that:

$$\min_t \max_f L(h^t, f) \leq \epsilon \quad (37)$$

Thus, the h^t that minimizes $L(h^t, f^t)$ is half of an ϵ -approximate Nash equilibrium of the ReLa game and can be computed within $\text{poly}(\frac{1}{\epsilon})$ iterations. In our setting, ϵ corresponds to the additional expected CMR violation our recovered h suffers on top of the best element in \mathcal{H} . \square

We note this reduction allows one to plug in *any* no-regret algorithm (e.g. Follow the Regularized Leader (McMahan, 2011) or Multiplicative Weights (Arora et al., 2012)) and the efficiency and approximate CMR satisfaction result to still hold. Practically, one could instantiate this no-regret reduction via a GAN-like optimization procedure with the learning rate for the h player much lower than that of the f player, simulating the no-regret vs. best response iterations of Algorithm 1. While the connection to no-regret online learning is not made explicit, both (Dikkala et al., 2020) and (Dai et al., 2018) follow an approach in this vein to scale their methods to high-dimensional tasks.

4. Discussion

We unify the techniques of (Dikkala et al., 2020) and (Dai et al., 2018) under the umbrella of satisfying relaxed conditional moment restrictions, providing a framework that allows one to reason about sample uncertainty and the effects of regularization on such problems. We further consider explicit constraints on slack variables, the tie-breaking benefits of regularizing h , and provide a reduction to no-regret online learning for efficient solving of Regularized Lagrangian Games.

Moving forward, we would be interested in passing the higher bar of matching not just the first moments of $Y|z$ but the whole distribution, as well as fitting other problems into the framework of conditional moment restrictions.

5. Acknowledgements

We thank Keegan Harris and Logan Stapleton for comments on a draft of this work. ZSW was supported in part by the NSF FAI Award #1939606, a Google Faculty Research Award, a J.P. Morgan Faculty Award, a Facebook Research Award, and a Mozilla Research Grant. GS is supported by NSF Award #1952085 and his family and friends.

References

- Arora, S., Hazan, E., and Kale, S. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.
- Baird, L. Residual algorithms: Reinforcement learning with function approximation. In *Machine Learning Proceedings 1995*, pp. 30–37. Elsevier, 1995.
- Boyd, S., Boyd, S. P., and Vandenberghe, L. *Convex optimization*. Cambridge university press, 2004.
- Chamberlain, G. Asymptotic efficiency in estimation with conditional moment restrictions. *Journal of econometrics*, 34(3):305–334, 1987.
- Chen, X. and Pouzo, D. Estimation of nonparametric conditional moment models with possibly nonsmooth generalized residuals. *Econometrica*, 80(1):277–321, 2012.
- Dai, B., Shaw, A., Li, L., Xiao, L., He, N., Liu, Z., Chen, J., and Song, L. Sbeed: Convergent reinforcement learning with nonlinear function approximation, 2018.
- Dikkala, N., Lewis, G., Mackey, L., and Syrgkanis, V. Minimax estimation of conditional moment models, 2020.
- Freund, Y. and Schapire, R. E. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.
- Hartford, J., Lewis, G., Leyton-Brown, K., and Taddy, M. Deep iv: A flexible approach for counterfactual prediction. In *International Conference on Machine Learning*, pp. 1414–1423. PMLR, 2017.
- Hestenes, M. R. Multiplier and gradient methods. *Journal of optimization theory and applications*, 4(5):303–320, 1969.
- Imbens, G. The role of the propensity score in estimating dose-response functions. *Biometrika*, 87(3):706–710, 2000.

- Kilbertus, N., Kusner, M. J., and Silva, R. A class of algorithms for general instrumental variable models, 2020.
- McMahan, H. B. Follow-the-regularized-leader and mirror descent: Equivalence theorems and ℓ_1 regularization. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011.
- Nie, X., Brunskill, E., and Wager, S. Learning when-to-treat policies, 2020.
- Puterman, M. L. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- Rehkopf, D. H., Glymour, M. M., and Osypuk, T. L. The consistency assumption for causal inference in social epidemiology: when a rose is not a rose. *Current epidemiology reports*, 3(1):63–71, 2016.
- Russell, S. and Norvig, P. *Artificial intelligence: a modern approach*. 2002.
- Winship, C. and Morgan, S. L. The estimation of causal effects from observational data. *Annual review of sociology*, 25(1):659–706, 1999.