
Supplement to Causal Strategic Linear Regression

Yonadav Shavit¹ Benjamin L. Edelman¹ Brian Axelrod²

Algorithm 1 Agent Outcome Maximization

Input: scalar parameters λ_{max}, ϵ , matrix V
Let $n = O(\epsilon^{-1}\lambda_{max}d)$
Sample $x_1 \dots x_n$ with $\omega = 0$.
Let $\hat{\mu} = \frac{1}{n} \sum x_i$
for $i = 1$ **to** d **do**
 Sample $x_1 \dots x_n$ with $\omega = e_i$
 Let $\hat{\omega}_i = \frac{1}{n} \sum x_i$
end for
Let $\hat{\omega} = \frac{\hat{\omega}}{\|\hat{\omega}\|}$
Return $\hat{\omega}$

Theorem 1. *Suppose the covariance matrix Σ has largest eigenvalue $\leq \lambda_{max}$, and suppose the outcome noise η is 1-subgaussian. Then Algorithm 1 learns an approximation of ω_{imp} with squared ℓ_2 error at most ϵ in $d + 1$ rounds with $O(\epsilon^{-1}\lambda_{max}d + 1)$ samples in each round. Furthermore, Algorithm 1 is non-adaptive. It succeeds and achieves the desired error with high probability.*

Proof. First we note that it is straightforward to compute the action that each agent will take. The agent maximizes $\omega^T V(x + Ma) - \frac{1}{2}\|a\|^2$ over $a \in \mathbb{R}^m$. Note that $\nabla_a(\omega^T VMa - \frac{1}{2}\|a\|^2) = M^T V\omega - a$. Thus,

$$\begin{aligned} & \arg \max_a \omega^T V(x + Ma) - \frac{1}{2}\|a\|^2 \\ &= \arg \max_a \omega^T VMa - \frac{1}{2}\|a\|^2 \\ &= M^T V\omega \end{aligned}$$

That is, every agent chooses $x_g = x + MM^T V\omega = x + G\omega$. This means that if the decision-maker publishes ω , the resulting expected agent outcome is $\omega^{*T}(x + G\omega)$. Hence, the optimal value of ω for the decision-maker to choose is

$$\frac{G^T \omega^*}{\|G^T \omega^*\|_2}$$

Imagine for a moment that there were no outcome noise. We could directly make linear measurements of perfect quality by choosing ω 's that form a basis. We do the same thing in Algorithm 1, just in a robust fashion.

We now describe the algorithm.

In the first round we set $\omega = 0$ in order to obtain an estimate for $E[\omega^{*T}x]$. Since ω is a unit vector, the variance of $\omega^{*T}x$ is at most λ_{max} plus a constant (from the 1-subgaussian noise).

¹Harvard School of Engineering and Applied Sciences, Cambridge, MA, USA ²Stanford Computer Science Department, Palo Alto, CA, USA. Correspondence to: Yonadav Shavit <yonadav@g.harvard.edu>.

This means that $O(\epsilon^{-1}d\lambda_{max})$ samples suffice for the empirical estimator of $E[\omega^{*T}x]$ to have no more than $\frac{\epsilon}{2d}$ error with failure probability $\Omega(\frac{1}{2d})$. We call the output of this estimator $\hat{\mu}$ and let $\hat{\mu}_d$ be the r -dimensional vector with $\hat{\mu}$ in every coordinate.

Now we choose $\omega_1 \dots \omega_d$ that form an orthonormal basis of the image of the diagonal matrix V . For each ω we observe the reward $\omega^{*T}(x + G\omega) + \eta$, subtract out $\hat{\mu}$, and plug it into the empirical mean estimator. For each ω_i , let $\hat{\omega}^*$ be the resulting coefficient. After $O(d\lambda_{max})$ samples, each coefficient has at most $\frac{\epsilon}{d}$ error with failure probability at most $\frac{1}{2d}$. Since we have computed $d + 1$ estimators, each one with failure probability at most $\frac{1}{2d}$, a union bound gives us a total failure probability that is sub-constant.

We can now bound the total squared ℓ_2 error between said coefficients and ω_{imp} in the $\omega_1 \dots \omega_d$ basis (noting that the choice of basis does not affect the magnitude of the error). We can break up the error into two components using the triangle inequality: the error due to $\hat{\mu}_d$ and the error in the subsequent rounds. Each coordinate of $\hat{\mu}_d$ has magnitude at most $\frac{\epsilon}{2d}$, so the total magnitude of the error in $\hat{\mu}_d$ is at most $\frac{\epsilon}{2}$. The same argument applies for the error in the coordinate estimates, leading to a total squared ℓ_2 error of at most ϵ .

We note that if our estimate is outside the unit ball, re-normalizing to within the unit ball only decreases the algorithm's error. □

Lemma 1. *Let ω be the linear decision rule and let a be the action taken by agents in response to h . Suppose that Ma and x are uncorrelated. Then the expected squared error of a decision rule ω on the gamed distribution can be decomposed as the sum of the following two positive terms (plus a constant offset):*

1. *The risk of ω on the original distribution*
2. *The squared error of h in predicting the expected impact (on agents' outcomes) of a , weighted by p .*

That is,

$$\begin{aligned} Risk(\omega) &\propto \mathbb{E}_{x \sim P} \left[((V\omega - \omega^*)^T x)^2 \right] \\ &\quad + p \cdot \mathbb{E}_{x \sim P} \left[((V\omega - \omega^*)^T (Ma))^2 \right] \end{aligned} \tag{1}$$

Note that this lemma holds regardless of the choice of agent action model.

Proof.

$$\begin{aligned} Risk(\omega) &= p \cdot \mathbb{E}_{x,a} \left[\left(\omega^T V(x + Ma) - \omega^{*T}(x + Ma) \right)^2 \right] + (1-p) \mathbb{E}_{x,a} \left[\left(\omega^T Vx - \omega^{*T}x \right)^2 \right] \\ &= p \cdot \mathbb{E}_{x,a} \left[\left(\left(\omega^T Vx - \omega^{*T}x \right) + \left(\omega^T VMa - \omega^{*T}Ma \right) \right)^2 \right] + (1-p) \mathbb{E}_{x,a} \left[\left(\omega^T Vx - \omega^{*T}x \right)^2 \right] \\ &= \mathbb{E}_{x,a} \left[\left(\omega^T Vx - \omega^{*T}x \right)^2 \right] + p \cdot \mathbb{E}_{x,a} \left[\left(\omega^T VMa - \omega^{*T}Ma \right)^2 \right] \end{aligned}$$

where the last line follows because Ma and x are uncorrelated. □

Parameter Recovery

In this section we describe how we recover $\hat{\omega}_{opt}$ in L^2 -distance when there exists an ω such that $\Sigma + G\omega$ is full rank. Before we proceed we make a couple of observations. When there is no way to make the above matrix full rank, we cannot hope to recover the optimal $\hat{\omega}_{opt}$. If there is no natural variation in e.g. the last two features, and furthermore no agent can act along those features, it is not possible to disentangle their potential effects on the outcome. This also suggests that the parameter recovery is a more substantive demand for the decision maker than the standard linear regression setting. To discover this additional information, the decision maker can incentivize the agents to take actions that help the decision-maker recover the true outcome-governing parameters.

This motivates the algorithm we present in this section. It operates in two stages. First, it recovers the information necessary in order to identify the decision rule which will provide the most informative agent samples after those agents have gamed. Second, it collects data while incentivizing this action. Finally, it computes an estimate of $\hat{\omega}_{opt}$ using the collected data. We present the complete procedure in Algorithm 2.

Algorithm 2 Recovering the Causal Model

- 1: Let $k_1 = \lambda_{max}(G^T G)$ and $k_2 = \|\Sigma\|^2$
 - 2: Let $\kappa_{min} = \lambda_{min}(\Sigma)$
 - 3: Choose an $\epsilon > 0$
 - 4: Let $n_1 = O(\max(\frac{dk_1}{\kappa_{min}}, \frac{d^2 k_2}{\kappa_{min}}))$
 - 5: Collect samples x_1, \dots, x_{n_1}
 - 6: Let $\hat{\mu} = \frac{1}{n_1} \sum x_i$
 - 7: Let $\hat{\Sigma} = \frac{1}{n_1} \sum x_i x_i^T$
 - 8: Let $n_2 = O(\max(d^2 \|\hat{\mu}\|^2 \text{tr}(\Sigma), d^3 \|G\|^2 \text{tr}(\Sigma)))$
 - 9: **for** $i = 1 \dots d$ **do**
 - 10: $\omega = e_i$
 - 11: Sample $x_1, \dots, x_{n_2}^i$ and subtract $\hat{\mu}$ from each one.
 - 12: Let $\hat{G}_i = \frac{1}{n_2} \sum_{j=1}^{n_2} x_j$
 - 13: **end for**
 - 14: Let $\hat{\omega}_{opt} = \arg \min_{\omega} \hat{\Sigma} + 2\mu\omega^T \hat{G}^T + \hat{G}\omega\omega^T \hat{G}^T$
 - 15: Let $n_3 = O(\frac{d}{\epsilon \kappa_{min}})$
 - 16: Sample x_1, \dots, x_{n_3} with $\omega = \hat{\omega}_{opt}$.
 - 17: Return the output of OLS on x_1, \dots, x_{n_3}
-

The procedure in Algorithm 2 can be summarized as follows:

1. Estimate the first and second moments of the distribution of agents' features.
2. Estimate the Gramian of the action matrix G .
3. Compute the most informative choice of ω .
4. Collect samples under the most informative ω and then return the output of OLS.

Before we proceed to the proof of correctness of Algorithm 2, let us build some intuition for why this procedure of choosing a single ω and collecting samples under said ω makes sense. As we show later, the convergence of OLS for linear regression can be controlled by the minimum eigenvalue of the second moment matrix of the samples. Our algorithm finds the value of ω that, after agents game, maximizes this minimum eigenvalue in expectation. It turns out the minimum eigenvalue of the expected second moment matrix of post-gaming samples is convex with respect to the choice of ω . The convexity of the objective suggests that a priori, when choosing ω s to obtain informative samples, the optimal strategy is choose a single specific ω .

The main difficulty in the rest of the algorithm is achieving the necessary precision in the estimation to be able to set up the above optimization problem to identify such an ω .

Theorem 3. *When $V = I$, the output of Algorithm 2 run with parameter ϵ satisfies $\|\omega - \omega^*\| \leq \epsilon$ with probability greater than $\frac{2}{3}$.*

The proof of Theorem 3 relies on several lemmas. First we bound the L_2 error of OLS as a function of the empirical second moment matrix in Lemma 2. Note that the usual bound for the convergence of OLS is distribution dependent. That is, the expected error is small.

Lemma 2. *Assume $V = I$. Consider samples x_1, \dots, x_n and $y_i = \hat{\omega}_{opt}^T x_i + \eta_i$. Let ω be the output of OLS (x_i, y_i) . Then*

$$\mathbb{E}_{\eta} [\|\omega - \hat{\omega}_{opt}\|^2] \leq \frac{d}{n\kappa_{min}}$$

The above proof is elementary and a slight modification of the standard textbook proof (see for example, (Liang, 2016)).

The proof also requires that the optimization to choose the optimal ω is convex.

Lemma 3. *The minimum eigenvalue of the following matrix is convex with respect to ω for any values of x, G .*

$$\sum_i (x_i + \hat{G}\omega)(x_i + \hat{G}\omega)^T$$

Furthermore, when the following conditions are true, then the minimum eigenvalue of the above is within a constant factor of the optimal value.

$$\mathbb{E}[(x + G\omega)(x + \hat{G}\omega)^T].$$

- $\|\hat{\Sigma} - \Sigma\|^2 \leq \epsilon$
- $\|\mu - \hat{\mu}\|^2 \leq \frac{\lambda_{\max}(G^T G)\epsilon}{d}$
- $\|\hat{G} - G\|^2 \leq \frac{\epsilon}{d\|\mu\|^2}$
- $\|\hat{G} - G\|^2 \leq \frac{\epsilon}{d^2\|G\|^2}$

Finally, the above holds true even for an ω with distance at most $O(\frac{1}{\text{poly}(d)})$ from the optimum.

Finally, we use a minor lemma for recover of a random vector via the empirical mean estimator. Note that we treat the matrix G as a vector.

Lemma 4. *Assume $V = I$. Let g_1^i, \dots, g_n^i be drawn from the distribution $G_i + \xi$ and \hat{G} be the empirical mean estimator computed from said g_j^i 's. Let Σ be the expected second moment matrix of the ξ s. Then*

$$\mathbb{E}_\xi \|G - \hat{G}\|^2 \leq \frac{d^2 \text{tr}(\Sigma)}{n}$$

We proceed with the proof of Theorem 3 below.

Proof. The first step of the algorithm is for recovering an estimate of Σ and μ . Note that n_1 samples suffice to recover $\hat{\Sigma}$ and $\hat{\mu}$ such that:

- $\|\hat{\Sigma} - \Sigma\|^2 \leq \epsilon$
- $\|\mu - \hat{\mu}\|^2 \leq \frac{\lambda_{\max}(G^T G)\epsilon}{d}$

The for loop recovers an estimate of G . Via Lemma 4, the samples suffice to ensure that the following two conditions hold:

- $\|\hat{G} - G\|^2 \leq \frac{\epsilon}{d\|\mu\|^2}$
- $\|\hat{G} - G\|^2 \leq \frac{\epsilon}{d^2\|G\|^2}$

Then the algorithm computes an estimate of the optimal ω . Via Lemma 3, we have that the optimum guarantees the minimum eigenvalue of an approximate solution will be within a constant factor of the optimum.

This ω guarantees that n_3 samples suffice to ensure the recover of ω^* within squared L^2 -distance of $O(\epsilon)$ in expectation.

Finally the expectations can be used with a Markov inequality to ensure the algorithm succeeds with (arbitrarily high) constant probability. \square

Now we prove the lemmas. We begin with Lemma 2. This proof is a slight modification of the textbook proof for the convergence of OLS.

Proof. In this section we derive a bound on the convergence of the least squares estimator when a fixed design matrix X is used. Note this is exactly the case we encounter, since the choice of ω lets us affect the entries of the design matrix. This is a standard, textbook result and not a main contribution of the paper.

In order to state the result more formally we have to introduce some notation. The goal of the procedure is to recover $\hat{\omega}_{opt}$, when given tuples $(x_i, \hat{\omega}_{opt}x_i + \eta)$ where η is 1-subgaussian. We aim to characterize $\|\omega - \hat{\omega}_{opt}\|$ where ω is obtained from ordinary least squares. Let X be the vector with the x_i 's in its columns. Let κ_{min} be the minimum eigenvalue of $\frac{1}{n}X^T X$ (the second moment matrix).

Below all expectations are taken *only* over the random noise. We assume the second moment matrix is full rank.

$$\begin{aligned} \mathbb{E}[\|\omega - \hat{\omega}_{opt}\|^2] &\leq \mathbb{E}\left[\frac{1}{n\kappa_{min}}(\omega - \hat{\omega}_{opt})X^T X(\omega - \hat{\omega}_{opt})\right] \\ &= \mathbb{E}\left[\frac{1}{n\kappa_{min}}\|X(\omega - \hat{\omega}_{opt})\|^2\right] \\ &= \frac{1}{n\kappa_{min}}\mathbb{E}[\|X(X^T X)^{-1}X^T(X\hat{\omega}_{opt} + \eta) - X\hat{\omega}_{opt}\|^2] \\ &= \frac{1}{n\kappa_{min}}\mathbb{E}[\|X(X^T X)^{-1}X^T\eta\|^2] \\ &\leq \frac{d}{n\kappa_{min}} \end{aligned}$$

This motivates our procedure for parameter recovery. We do so in a fashion that attempts to maximize κ_{min} . Note that it is the minimum eigenvalue that determines the convergence rate. This is due to the fact that little variation along a dimension makes it hard to disentangle the features' effect on the outcome via $\hat{\omega}_{opt}$ from the constant-variance noise η . \square

Lemma 3 is somewhat more involved. It is proven in three parts. The first is that the optimization problem is convex. The second is that approximate recovery of S , μ , and G suffice for approximately minimizing the original expression. The third is that an approximate solution suffices.

Proof. In this section we describe how to choose the value of ω that maximizes the value of κ_{min} for the samples we obtain.

To do so, we examine the expectation of the second moment matrix and make several observations. Let Σ denote the expected second moment matrix of x (i.e. $\mathbb{E}[xx^T]$). We have:

$$\mathbb{E}[(x + G\omega)(x + G\omega)^T] = \Sigma + 2\mu\omega^T G^T + G\omega\omega^T G^T$$

1. The minimum eigenvalue of the above expression is concave with respect to ω . This follows due to the following: $x + G\omega$ is a linear operator, the minimum eigenvalue of a Gramian matrix $X^T X$ is concave with respect to X , and the expectation of a concave function is concave (Boyd & Vandenberghe, 2004).
2. Since the agent attempts to maximize their motion in the ω direction, we want to ensure that we move them toward toward the direction that maximizes the minimum eigenvalue of $\mathbb{E}[(x + G\omega)(x + G\omega)^T]$.

However, we do not operate with exact knowledge of G , etc. It turns out that even approximately solving this optimization problem with estimates for G , Σ , μ suffices for our purposes, as long as the ω we obtain from our optimization (using the estimates) results in a high value for the minimum eigenvalue of $\mathbb{E}[(x + G\omega)(x + G\omega)^T]$. Let $\hat{\omega}$ be the maximizing argument for the estimated optimization problem and let ω be the maximizing argument for the original optimization problem. Let Q be the true maximized second moment matrix including gaming, and \hat{Q} be the maximizing second moment matrix with gaming resulting from replacing the true Σ, μ, G with the estimates. In formal terms, we need to show the minimum eigenvalue of the following is large: $\mathbb{E}[(x + G\hat{\omega})(x + G\hat{\omega})^T]$. We note that when $y^T \hat{Q}y$ is within ϵ of $y^T Qy$ for all y in the

unit ball, the minimum eigenvalues may differ by at most ϵ .

$$\begin{aligned} \|y^T \hat{Q}y - y^T Qy\|^2 &= \|y^T (\hat{Q} - Q)y\|^2 \\ &\leq \lambda_{max}^2(\hat{Q} - Q)(y) \|y\|^2 \\ &\leq \|\hat{Q} - Q\|^2 \end{aligned}$$

And now we bound the norm of $\|\hat{\Sigma} - \Sigma\|^2$ assuming the following:

1. $\|\hat{\Sigma} - \Sigma\|^2 \leq \epsilon$
2. $\|\mu - \hat{\mu}\|^2 \leq \frac{\lambda_{max}(G^T G)\epsilon}{d}$
3. $\|\hat{G} - G\|^2 \leq \frac{\epsilon}{d\|\mu\|^2}$
4. $\|\hat{G} - G\|^2 \leq \frac{\epsilon}{d^2\|G\|^2}$

We work out the bound below.

$$\begin{aligned} \|\hat{Q} - Q\|^2 &= \|\Sigma + 2\mu\hat{\omega}_{opt}^T G^T + G\hat{\omega}_{opt}\hat{\omega}_{opt}^T G^T - (\hat{\Sigma} + 2\hat{\mu}\omega^T \hat{G}^T + \hat{G}\hat{\omega}_{opt}\hat{\omega}_{opt}^T \hat{G}^T)\|^2 \\ &\leq \|\Sigma - \hat{\Sigma}\|^2 + 2\|\mu\hat{\omega}_{opt}^T G^T - \hat{\mu}\hat{\omega}_{opt}^T \hat{G}^T\|^2 + \|\dots\|^2 \\ &\leq \epsilon + 2\|\mu\hat{\omega}_{opt}^T G^T + \hat{\mu}\hat{\omega}_{opt}^T \hat{G}^T - \hat{\mu}\hat{\omega}_{opt}^T G^T - \hat{\mu}\hat{\omega}_{opt}^T \hat{G}^T\|^2 + \|\dots\|^2 \\ &\leq \epsilon + d\|\mu - \hat{\mu}\|^2 \|\hat{\omega}_{opt}^T G\|^2 + \|\hat{G} - G\|^2 \|\hat{\mu}\omega^T\|^2 + \dots \\ &\leq \epsilon + \epsilon + \|\hat{G} - G\|^2 \|\hat{\mu}\omega^T + \mu\omega^T - \mu\omega^T\|^2 + \dots \\ &\leq \epsilon + \epsilon + \|\hat{G} - G\|^2 (\|\hat{\mu} - \mu\|^2 + \|\mu\hat{\omega}_{opt}^T\|) + \dots \\ &\leq \epsilon + \epsilon + \|\hat{G} - G\|^2 d\|\mu\|^2 + \dots \\ &\leq 3\epsilon + \|\hat{G}\hat{\omega}_{opt}\hat{\omega}_{opt}^T \hat{G}^T - G\hat{\omega}_{opt}\hat{\omega}_{opt}^T G^T\|^2 \\ &\leq 3\epsilon + \|\hat{G}\hat{\omega}_{opt}\hat{\omega}_{opt}^T \hat{G}^T - \hat{G}\hat{\omega}_{opt}\hat{\omega}_{opt}^T G + \hat{G}\hat{\omega}_{opt}\hat{\omega}_{opt}^T G - G\hat{\omega}_{opt}\hat{\omega}_{opt}^T G^T\|^2 \\ &\leq 3\epsilon + \|(\hat{G} - G)\hat{\omega}_{opt}\hat{\omega}_{opt}^T \hat{G}^T\|^2 + \|(\hat{G} - G)\hat{\omega}_{opt}\hat{\omega}_{opt}^T G^T\|^2 \\ &\leq 3\epsilon + \|(\hat{G} - G)\hat{\omega}_{opt}\hat{\omega}_{opt}^T \hat{G}^T - (\hat{G} - G)\hat{\omega}_{opt}\hat{\omega}_{opt}^T G^T + (\hat{G} - G)\hat{\omega}_{opt}\hat{\omega}_{opt}^T G^T\|^2 + d^2\|\hat{G} - G\|^2 \|G\|^2 \\ &\leq 4\epsilon + \|(\hat{G} - G)\hat{\omega}_{opt}\hat{\omega}_{opt}^T (\hat{G} - G)\|^2 + \|(\hat{G} - G)\hat{\omega}_{opt}\hat{\omega}_{opt}^T G^T\|^2 \\ &\leq 5\epsilon + d^2\|\hat{G} - G\|^4 \\ &\leq 6\epsilon \end{aligned}$$

This means if we find an ϵ -approximate solution to the system with the estimated values, we obtain a κ_{min} within 6ϵ of the optimal. \square

Finally, we present the proof of Lemma 4:

Proof. Recall that when the decision-maker fixes ω , it receives samples of the form $x + G\omega$. We note this can be used to recover the matrix G . In particular, we show how d rounds, each with $O(\frac{d\text{tr}(\Sigma)}{\epsilon})$ samples, suffices to recover the matrix to squared Frobenius norm ϵ . Recall the procedure we propose simply chooses $\omega = e_1, \dots, e_d$, one-hot coordinate vectors in each round. We first bound the error in \hat{G} . coordinate-wise: $\mathbb{E}[\|\hat{G}_{i,j} - G_{i,j}\|^2] \leq \frac{\mathbb{E}[x_i^2]}{n}$. A union bound across coordinates shows that $O(\frac{d^2\text{tr}(\Sigma)}{\epsilon})$ samples suffice to recover G within squared Frobenius norm ϵ . \square

References

Boyd, S. and Vandenberghe, L. *Convex Optimization*. Cambridge university press, 2004.

Liang, P. Cs229t/stat231: Statistical learning theory (lecture notes). <https://web.stanford.edu/class/cs229t/notes.pdf>, April 2016.