
Estimation of Personalized Effects Associated With Causal Pathways

Razieh Nabi

Computer Science Dept.
Johns Hopkins University
Baltimore, MD, 21218

Phyllis Kanki

Immunology and Infectious Diseases Dept.
Harvard T. H. Chan School of Public Health
Boston, MA, 02115

Ilya Shpitser

Computer Science Dept.
Johns Hopkins University
Baltimore, MD, 21218

Abstract

The goal of personalized decision making is to map a unit's characteristics to an action tailored to maximize the expected outcome for that unit. Obtaining high-quality mappings of this type is the goal of the dynamic regime literature. In healthcare settings, optimizing policies with respect to a particular causal pathway may be of interest as well. For example, we may wish to maximize the chemical effect of a drug given data from an observational study where the chemical effect of the drug on the outcome is entangled with the indirect effect mediated by differential adherence. In such cases, we may wish to optimize the direct effect of a drug, while keeping the indirect effect to that of some reference treatment. [15] shows how to combine mediation analysis and dynamic treatment regime ideas to define policies associated with causal pathways and counterfactual responses to these policies. In this paper, we derive a variety of methods for learning high quality policies of this type from data, in a causal model corresponding to a longitudinal setting of practical importance. We illustrate our methods via a dataset of HIV patients undergoing therapy, gathered in the Nigerian PEPFAR program.

1 INTRODUCTION

There has been growing interest in making personalized decisions in different domains to account for inherent heterogeneity among individuals and optimize individual-level experiences. For instance, personalized medicine aims at systematic use of individual patient history including biological information and biomarkers to

improve patient's health care. Personalized actions can be viewed as realizations of decision rules where available information is mapped to the space of possible decisions.

Making good personalized decisions often involves acting in multiple stages. For instance, multiple successive medical interventions may be required for long-term care of patients with chronic diseases. The goal of personalized medicine is to tailor a sequence of decision rules on treatment, known as dynamic treatment regimes, based on patient characteristics seen so far, to maximize the likelihood of a desirable outcome. A number of algorithms have been developed for finding optimal treatment regimes from either observational data, or data from experiments tailored for providing information on regime quality, such as sequential multiple assignment randomized trials (SMARTs) [3, 6]. These algorithms use methods from causal inference, and aim to predict counterfactual outcomes under policies different from those actually followed in the data [11, 5, 2].

A natural extension of these methods is finding treatment regimes that optimizes a *part* of the effect of the treatment on the outcome. We illustrate the utility of this problem with the following example. Patients infected with human immunodeficiency virus (HIV) are typically put on courses of antiretroviral therapy, as a first line of therapy. Although these sort of medications are effective in combating the disease, the full benefit is not realized since patients often do not fully adhere to the medication regimen. One of the causes of poor adherence to the therapy is toxicity of the medication. Hence, viral failure in a patient receiving treatment may be attributed to either poor drug effectiveness or lack of adherence to an otherwise effective treatment plan. In observational studies of HIV patients, treatments are not randomly assigned, and patients have differential adherence. Because of this, finding a policy that optimizes the *overall effect* of the treatment plan on the outcome entangles two very different causal pathways – the chemical pathway asso-

ciated with the active ingredients in the treatment, and the pathway associated with adherence.

We may, instead, consider the problem of finding a set of policies that optimize only the direct chemical effect of the drug, in the counterfactual situations where the indirect effect mediated by adherence can be kept to that of some reference treatment. Policies of this type may be more directly relevant in precision medicine contexts where adherence varies among patients.

[15] defines counterfactual responses to policies that set treatments only with respect to a particular causal pathway, and gives a general identification algorithm for these responses, as a generalization of similar algorithms for standard dynamic treatment regimes [19], and effects associated with causal pathways in mediation analysis [14]. In this paper, we consider algorithms that can be used to find policies that maximize effects along particular causal pathways from data, in a causal model of most immediate relevance in longitudinal settings.

The paper is organized as follows. We fix our notation, and define counterfactual responses to treatment policies and policies associated with pathways in Section 2. In Section 3, we review existing techniques used in learning policies optimizing the *overall effect* of actions. In Section 4, we fix the causal model corresponding to a typical longitudinal study, prove identification of counterfactual policy responses under this model, and show how techniques for maximizing treatment policies, described in Section 4, may be extended to maximize policies associated with causal pathways. In Section 5, we illustrate the methods we propose via an application involving data on treatment of HIV patients in Nigeria. Our conclusions are in Section 6. The Appendix contains the proofs of all claims, a basic description of statistical inference in semi-parametric models, as context for some of our estimation strategies, visualizations of learned policies, and descriptions of the experiments.

2 NOTATIONS AND PRELIMINARIES

Consider a multi-stage decision problem with K pre-specified decision points, indexed by $i = 1, \dots, K$. Let Y denote the final outcome of interest and A_i denote the action made at decision point i with the finite state space of \mathfrak{X}_{A_i} . The set of all actions is denoted by \mathbf{A} . Let W_0 denote the available information prior to the first decision, and W_i denote the information collected between decisions i and $i+1$, ($Y \equiv W_K$). Given A_i , denote \bar{A}_i to be all treatments administered from time 1 to i , similarly for W_i and \bar{W}_i . We combine the treatment and covariate history up to treatment decision A_i into a history vector H_i . The state space of H_i is denoted by \mathfrak{X}_{H_i} .

We are interested in learning policies that map H_i to values of A_i , for all i , that maximize the expected value of the outcome Y . Doing this from observed data entails considering *counterfactual outcomes* Y had A_i been assigned in a different way from what was actually observed. We briefly review graphical causal models, and the potential outcome notation from causal inference, which will be used to define such counterfactuals.

Causal models are sets of distributions defined by restrictions associated with directed acyclic graphs (DAGs). We will use vertices and variables interchangeably – capital letters for a vertex or variable (V), bold capital letter for a set (\mathbf{V}), lowercase letters for values (v), and bold lowercase letters for sets of values (\mathbf{v}). By convention, each graph is defined on a vertex set \mathbf{V} . For a set of values \mathbf{a} of \mathbf{A} , and a subset $\mathbf{A}^\dagger \subseteq \mathbf{A}$, define $\mathbf{a}_{\mathbf{A}^\dagger}$ to be a restriction of \mathbf{a} to elements in \mathbf{A}^\dagger . For a DAG \mathcal{G} , and any $V \in \mathbf{V}$, we define the parents of $V \in \mathbf{V}$ to be the set $\text{pa}_{\mathcal{G}}(V) \equiv \{W \in \mathbf{V} \mid W \rightarrow V\}$, and the children of $V \in \mathbf{V}$ to be the set $\text{ch}_{\mathcal{G}}(V) \equiv \{W \in \mathbf{V} \mid V \rightarrow W\}$.

Causal models of a DAG \mathcal{G} consist of distributions defined on counterfactual random variables of the form $V(\mathbf{a})$ where \mathbf{a} are values of $\text{pa}_{\mathcal{G}}(V)$. These variables represent outcomes of V had all variables in $\text{pa}_{\mathcal{G}}(V)$ been set, possibly contrary to fact, to \mathbf{a} . In this paper we assume Pearl’s functional model for a DAG \mathcal{G} with vertices \mathbf{V} which is the set containing any joint distribution over all potential outcome random variables where the sets of variables $\{\{V(\mathbf{a}_V) \mid \mathbf{a}_V \in \mathfrak{X}_{\text{pa}_{\mathcal{G}}(V)}\} \mid V \in \mathbf{V}\}$ are mutually independent [8]. The *atomic counterfactuals* in the above set model the relationship between $\text{pa}_{\mathcal{G}}(V)$, representing direct causes of V , and V itself. From these, all other counterfactuals may be defined using *recursive substitution*. For any $\mathbf{A} \subseteq \mathbf{V} \setminus \{V\}$,

$$V(\mathbf{a}) \equiv V(\mathbf{a}_{\text{pa}_{\mathcal{G}}(V) \cap \mathbf{A}}, \{\text{pa}_{\mathcal{G}}(V) \setminus \mathbf{A}\}(\mathbf{a})), \quad (1)$$

where $\{\text{pa}_{\mathcal{G}}(V) \setminus \mathbf{A}\}(\mathbf{a})$ is taken to mean the (recursively defined) set of counterfactuals associated with variables in $\text{pa}_{\mathcal{G}}(V)$, had \mathbf{A} been set to \mathbf{a} .

A causal parameter is said to be *identified* in a causal model if it is a function of the observed data distribution $p(\mathbf{V})$. In all causal models of a DAG \mathcal{G} in the literature, all interventional distributions $p(\{\mathbf{V} \setminus \mathbf{A}\}(\mathbf{a}))$ are identified by the *g-formula*:

$$p(\{\mathbf{V} \setminus \mathbf{A}\}(\mathbf{a})) = \prod_{V \in \mathbf{V} \setminus \mathbf{A}} p(V \mid \text{pa}_{\mathcal{G}}(V)) \Big|_{\mathbf{A}=\mathbf{a}} \quad (2)$$

As an example, $Y(\mathbf{a})$ in the DAG in Fig. 1 (a), is defined to be $Y(a, M(a, W), W)$, and its distribution is identified as $\sum_{w,m} p(Y|a, m, w)p(m|a, w)p(w)$. In our sequential decision setting, the relevant counterfactual is $Y(\bar{a}_K)$, i.e. the response of Y had the treatment history $\bar{A}_K = \bar{a}_K$ been administered, possibly contrary to fact. Comparison of $Y(\bar{a}_K)$ and $Y(\bar{a}'_K)$ in expectation,




Figure 1: (a) A simple causal DAG, with a single treatment A , a single outcome Y , a vector W of baseline variables, and a single mediator M . (b) A more complex causal DAG with two treatments A_1, A_2 , an intermediate outcome W_1 , and the final outcome W_2 .

$\mathbb{E}[Y(\bar{a}_K)] - \mathbb{E}[Y(\bar{a}'_K)]$, where \bar{a}_K is the treatment history of interest, and \bar{a}'_K is the reference treatment history, gives the average causal effect of \bar{a}_K on the outcome Y .

Counterfactual Response to Policies

A dynamic treatment regime $\mathbf{f}_A = \{f_{A_1}, \dots, f_{A_K}\}$ is a sequence of decision rules that forms a treatment plan for patients over time. At the i th decision point, the i th rule f_{A_i} maps the available information H_i prior to the i th decision point to a treatment a_i , i.e. $f_{A_i} : \mathfrak{X}_{H_i} \mapsto \mathfrak{X}_{A_i}$. Given a treatment regime \mathbf{f}_A , we define the counterfactual response Y had \mathbf{A} been assigned according to \mathbf{f}_A , or $Y(\mathbf{f}_A)$, as the following generalization of (1)

$$Y(\{f_{A_i}(H_i(\mathbf{f}_A)) | A_i \in \text{pa}_{\mathcal{G}}(Y) \cap \mathbf{A}\}, \{\text{pa}_{\mathcal{G}}(Y) \setminus \mathbf{A}\}(\mathbf{f}_A)). \quad (3)$$

Under a causal model associated with the DAG \mathcal{G} , the distribution $p(Y(\mathbf{f}_A))$, is identified as the following generalization of the g-formula

$$\sum_{V \setminus Y, A, V \in V \setminus \mathbf{A}} \prod p(V | \{f_{A_i}(H_i) | A_i \in \text{pa}_{\mathcal{G}}(V) \cap \mathbf{A}\}, \text{pa}_{\mathcal{G}}(V) \setminus \mathbf{A}) \quad (4)$$

As an example, $Y(a = f_A(W))$ in Fig. 1 (a) is defined as $Y(a = f_A(W), M(a = f_A(W), W), W)$, and its distribution is identified as $\sum_{w,m} p(Y | a = f_A(w), m, w) p(m | a = f_A(w), w) p(w)$. We will call policies corresponding to counterfactuals defined in (3) *overall policies*, to distinguish them from policies associated with causal pathways to be defined later.

Mediation Analysis

An important goal of causal inference is understanding the mechanisms by which the treatment A influences the outcome Y . A common framework for mechanism analysis is *mediation analysis* which seeks to decompose the effect of A on Y into the *direct effect* and the *indirect effect* mediated by a third variable, or more generally into components associated with particular causal pathways. Consider the graph in Fig. 1 (a): the direct effect corresponds to the effect along the edge $A \rightarrow Y$, and indirect effect corresponds to the effect along the path $A \rightarrow M \rightarrow Y$, mediated by M .

Counterfactuals associated with mediation analysis have been defined using a more general type of intervention in a graphical causal model, namely the *edge intervention* [16], which maps a set of directed edges in \mathcal{G} to values of their source vertices. Edge interventions have a natural interpretation in cases where a treatment variable has multiple components that a) influence the outcome in different ways, b) occur or do not occur together in observed data, and c) may in principle be intervened on separately. For instance, smoking leads to poor health outcomes due to two components: smoke inhalation and exposure to nicotine. A smoker would be exposed to both of these components, while a non-smoker to neither. However, one might imagine exposing someone selectively only to nicotine but not smoke inhalation (via a nicotine patch), or only smoke inhalation but not nicotine (via smoking plant matter not derived from tobacco leaves). These types of hypothetical experiments correspond precisely to edge interventions, and have been used to conceptualize direct and indirect effects [12, 7].

We will write the mapping of a set of edges to values of their source vertices as \mathbf{a}_α to mean edges in α are mapped to values in the *multiset* \mathbf{a} (since multiple edges may share the same source vertex, and be assigned to different values). For a subset $\beta \subseteq \alpha$ and an assignment \mathbf{a}_α , denote \mathbf{a}_β to be a restriction of \mathbf{a}_α to edges in β . For simplicity, in the remainder of the paper we assume that if $(AW)_{\rightarrow} \in \alpha$, then for all $V \in \text{ch}_{\mathcal{G}}(A)$, $(AV)_{\rightarrow} \in \alpha$, where $(XY)_{\rightarrow}$ is a directed edge from X to Y .

We will write counterfactual responses to edge interventions as $Y(\mathbf{a}_\alpha)$. An edge intervention that sets a set of edges α to values in the multiset \mathbf{a} is defined via the following generalization of recursive substitution (1):

$$Y(\mathbf{a}_\alpha) \equiv Y(\mathbf{a}_{\{(ZY)_{\rightarrow} \in \alpha\}}, \{\text{pa}_{\mathcal{G}}^{\bar{\alpha}}(Y)\}(\mathbf{a}_\alpha)), \quad (5)$$

where $\text{pa}_{\mathcal{G}}^{\bar{\alpha}}(Y) \equiv \{W \mid (WY)_{\rightarrow} \notin \alpha\}$. For example, in the DAG in Fig. 1 (a), $Y(\mathbf{a}_{\{(AY)_{\rightarrow}, (AM)_{\rightarrow}\}})$ assigning $(AY)_{\rightarrow}$ to a and $(AM)_{\rightarrow}$ to a' is defined as $Y(a, M(a', W), W)$.

Identification of edge interventions in graphical causal models of a DAG corresponds quite closely with identification of regular (node) interventions, as follows. Let $\mathbf{A}_\alpha \equiv \{A \mid (AB)_{\rightarrow} \in \alpha\}$. Consider an edge intervention given by the mapping \mathbf{a}_α . Then, under the functional model of a DAG \mathcal{G} , the joint distribution of counterfactual responses $p(\{\mathbf{V} \setminus \mathbf{A}_\alpha\}(\mathbf{a}_\alpha))$ is identified via the following generalization of (2) called the *edge g-formula*:

$$\prod_{V \in \mathbf{V} \setminus \mathbf{A}_\alpha} p(V | \mathbf{a}_{\{(ZV)_{\rightarrow} \in \alpha\}}, \text{pa}_{\mathcal{G}}^{\bar{\alpha}}(V)). \quad (6)$$

For example, in Fig 1 (a), $p(Y(\mathbf{a}_{\{(AY)_{\rightarrow}, (AM)_{\rightarrow}\}})) = \sum_{W, M} p(Y | a, M, W) p(M | a', W) p(W)$, which is obtained by marginalizing W and M from the edge g-formula.

Counterfactual Responses To Policies Associated With Pathways

We now define counterfactual responses to policies that operate only with respect to particular outgoing edges from \mathbf{A} . These counterfactuals generalize those in the previous two sections.

As an example, we can view Fig. 1 (a) as representing a cross-sectional study of HIV patients of the kind described in [4], where W is a set of baseline characteristics, A is one of a set of possible antiretroviral treatments, M is adherence to treatment, and Y is a binary outcome variable signifying viral failure. In this type of study, we may wish to find $f_A(W)$ that maximizes the expected outcome Y had A been set according to $f_A(W)$ for the purposes of the direct effect of A on Y , and A were set to some reference level a' for the purposes of the effect of A on M . In other words, we may wish to find $f_A(W)$ to maximize the counterfactual mean $\mathbb{E}[Y(f_A(W), M(a', W), W)]$. This would correspond to finding a treatment policy that maximizes the direct (chemical) effect, if it were possible to keep adherence to a level $M(a')$ as if a reference (easy to adhere to) treatment a were given.

We now give a general definition for responses to such path-specific policies. Fix a set of directed edges α , and define $\mathbf{A}_\alpha \equiv \{A \mid (AB)_\rightarrow \in \alpha\}$. As before, we assume if $(AW)_\rightarrow \in \alpha$, then for all $V \in \text{ch}_{\mathcal{G}}(A)$, $(AV)_\rightarrow \in \alpha$. Define $f_\alpha \equiv \{f_{A_i}^{(A_i W)_\rightarrow} : \mathfrak{X}_{H_i} \mapsto \mathfrak{X}_{A_i} \mid (A_i W)_\rightarrow \in \alpha\}$ as the set of policies associated with edges in α . Note that f_α may contain multiple policies for a given treatment variable A .

Define $Y(f_\alpha)$, the counterfactual response of Y to the set of path-specific policies f_α , as the following generalization of (5) and (3):

$$Y(\{f_{A_i}^{(A_i Y)_\rightarrow}(H_i(f_\alpha)) \mid (A_i Y)_\rightarrow \in \alpha\}, \{\text{pa}_{\mathcal{G}}^{\bar{\alpha}}(Y)\}(f_\alpha)) \quad (7)$$

To reformulate our earlier example, if $\tilde{f}_A^{(AM)_\rightarrow}$ assigns A to a constant value a' , and $f_{\{(AY)_\rightarrow, (AM)_\rightarrow\}} \equiv \{f_A^{(AY)_\rightarrow}(W), \tilde{f}_A^{(AM)_\rightarrow}\}$, then $Y(f_{\{(AY)_\rightarrow, (AM)_\rightarrow\}}) \equiv Y(f_A^{(AY)_\rightarrow}(W), M(a', W), W)$.

The joint counterfactual distribution for responses to path-specific policies, $p(\{V(f_\alpha) \mid V \in \mathbf{V} \setminus \mathbf{A}_\alpha\})$, is identified under the functional model, and generalizes (6) and (4) as follows:

$$\prod_{V \in \mathbf{V} \setminus \mathbf{A}_\alpha} p(V \mid \{f_{A_i}^{(A_i V)_\rightarrow}(H_i) \mid (A_i V)_\rightarrow \in \alpha\}, \text{pa}_{\mathcal{G}}^{\bar{\alpha}}(V)) \quad (8)$$

For example, $p(Y(f_A(W), M(a', W), W))$ is identified as $\sum_{W, M} p(Y \mid f_A(W), M, W) p(M \mid a', W) p(W)$ in Fig. 1 (a). We prove a general version of this result in the Appendix. A general identification theory for responses to path-specific policies can be found in [15].

3 LEARNING OPTIMAL OVERALL POLICIES

The goal of this paper is developing methods for learning optimal path-specific policies, in cases where responses to such policies are identified. Before discussing optimal path-specific policies, we first review existing approaches to finding optimal overall policies. Optimality may be quantified in a number of ways. The set of optimal policies is commonly defined as $\mathbf{f}_\mathbf{A}^* \equiv \arg \max_{\mathbf{f}_\mathbf{A}} \mathbb{E}[Y(\mathbf{f}_\mathbf{A})]$.

We will discuss existing methods for finding optimal policies $\mathbf{f}_\mathbf{A}^*$ in the context of a causal model implying *positivity and sequential ignorability* [10]. This model is graphically represented, for two time points, in Fig. 1 (b), where W_0 are baseline factors, W_1 and W_2 are intermediate and final outcomes, A_1, A_2 are treatments. Variables W_0, W_1, W_2 may be confounded by a hidden common cause U . It is well known that in this model, $\mathbb{E}[Y(\mathbf{f}_\mathbf{A})]$ is identified via (4). We now discuss a number of approaches for computing the optimal set $\mathbf{f}_\mathbf{A}^*$ given this identifying formula.

Approaches Based on Backwards Induction

In sequential decision problems, choosing optimal functions $\mathbf{f}_\mathbf{A}$ may appear to be a difficult search problem over a large set of function combinations. However, it is possible to use ideas from dynamic programming to transform the problem of choosing optimal $\mathbf{f}_\mathbf{A}$ into a sequential problem where only a single optimal function is chosen at a time.

Multiple modeling approaches are possible here. A simple approach is to model conditional densities of all outcomes W_i given their past. Assuming all such models were correctly specified, given any particular history H_K up to the last decision point A_K , the optimal $f_{A_K}^*$ is equal to $\mathbb{I}(\mathbb{E}[W_K(A_K = 1) \mid H_K] > \mathbb{E}[W_K(A_K = 0) \mid H_K])$, which under our model is equal to $\mathbb{I}(\mathbb{E}[W_K \mid A_K = 1, H_K] > \mathbb{E}[W_K \mid A_K = 0, H_K])$. Assuming optimal $f_{A_{i+1}}^*, \dots, f_{A_K}^*$, denoted by $f_{\Delta_{i+1}}^*$, were already chosen, the optimal $f_{A_i}^*$ is defined inductively as

$$\mathbb{I} \left[\mathbb{E}[W_K(A_i = 1, f_{\Delta_{i+1}}^*) \mid H_i] > \mathbb{E}[W_K(A_i = 0, f_{\Delta_{i+1}}^*) \mid H_i] \right].$$

Note that under our model, the counterfactual expectations above are identified via a modification of (4) where the outer summation is only with respect to variables W_{i+1}, \dots, W_{K-1} . For many kinds of statistical models for densities appearing as terms in (4), evaluating this sum may be challenging. An alternative strategy that avoids repeated summations is modeling the above expectations directly via *Q-functions* Q_i , which are conditional expectations over *value functions* V_{i+1} , given the

history. These are defined recursively as follows:

$$\begin{aligned} Q_K(H_K, A_K; \gamma_K) &= \mathbb{E}[W_K | A_K, H_K], \\ V_K(H_K) &= \max_{a_K} Q_K(H_K, a_K; \gamma_K), \end{aligned}$$

and for $i = K - 1, \dots, 1$, as

$$\begin{aligned} Q_i(H_i, A_i; \gamma_i) &= \mathbb{E}[V_{i+1}(H_{i+1}) | A_i, H_i], \\ V_i(H_i) &= \max_{a_i} Q_i(H_i, a_i; \gamma_i). \end{aligned}$$

The optimal policy at each stage may be easily derived from Q-functions as $f_{A_i}^*(H_i) = \arg \max_{a_i} Q_i(H_i, a_i; \gamma_i)$. Q-functions are recursively defined regression models where outcomes are value functions, and features are histories up to the current decision point. Thus, parameters $\gamma_i, i = 1, \dots, K$, of all Q-functions may be learned recursively by maximum likelihood methods applied to regression at stage i , given that the value function at stage $i + 1$ was already computed for every row [2, 13].

Value Search

Consider a restricted class of policies \mathcal{F} with elements $\mathbf{f}_A \equiv \{f_{A_i}(H_i); A_i \in \mathbf{A}\}$. It is often of interest to estimate the optimal policy within the class \mathcal{F} , even if the class does not contain the true optimal policy. For example, \mathcal{F} may be the set of clinically interpretable policies. For a sufficiently simple class \mathcal{F} , we can directly search for the optimal $\mathbf{f}_A^{*\mathcal{F}} \equiv \arg \max_{\mathbf{f}_A \in \mathcal{F}} \mathbb{E}[Y(\mathbf{f}_A)]$. This is called policy search or value search.

The expected response to an arbitrary treatment policy $\beta = \mathbb{E}[Y(\mathbf{f}_A)]$, for $\mathbf{f}_A \in \mathcal{F}$ can be estimated in a number of ways. For $\mathbf{A} = \{A\}$, a simple estimator for β that uses only the treatment assignment model $\pi(H; \psi)$ for $p(A = 1|H)$ is the inverse probability weighting (IPW) estimator:

$$\mathbb{E} \left[Y C_{f_A} / \pi_{f_A}(H; \hat{\psi}) \right], \quad (9)$$

where $C_{f_A} \equiv \mathbb{I}(A = f_A(H))$, $\pi_{f_A}(H; \psi) \equiv \pi(H; \psi) f_A(H) + (1 - \pi(H; \psi))(1 - f_A(H))$, the expectation is evaluated empirically, and $\hat{\psi}$ is fit by maximum likelihood. This estimator will not in general yield the optimal policy within \mathcal{F} if π is misspecified. An alternative estimator for β that provides some protection against this is the following:

$$\mathbb{E} \left[\frac{C_{f_A} Y}{\pi_{f_A}(H; \hat{\psi})} - \frac{C_{f_A} - \pi_{f_A}(H; \hat{\psi})}{\pi_{f_A}(H; \hat{\psi})} \mathbb{E}[Y|H, A = f_A(H); \hat{\gamma}] \right]. \quad (10)$$

The above estimator is *doubly-robust* meaning that it is a consistent estimator if either the propensity score model $\pi(H; \psi)$ or the regression model $\mathbb{E}[Y|H, A; \gamma]$ is correctly specified.

Value search methods can in some cases be rephrased as a weighted classification problem in machine learning [20]. In the interest of space, we do not discuss methods based on this observation further here.

G-Estimation

An alternative method for learning policies is to directly model the counterfactual contrast functions known as *optimal blip-to-zero functions*, or the counterfactual deviations in outcome from a reference treatment value (which we take to be $A = 0$), conditional on history, assuming all future decisions are already optimal. Specifically, for each decision point i , we posit a *structural nested mean model (SNMM)* $\gamma_i(H_i, A_i; \psi)$ for the contrast

$$\mathbb{E} \left[Y(\bar{a}_i, f_{A_{i+1}}^*) | H_i \right] - \mathbb{E} \left[Y(\bar{a}_{i-1}, a_i = 0, f_{A_{i+1}}^*) | H_i \right].$$

Note that if the true $\gamma_i(H_i, A_i; \psi)$ were known, the optimal treatment policies are those that maximize the blip function at each stage: $f_{A_i}^* = \arg \max_{a_i} \gamma_i(H_i, A_i; \psi_i)$. In order to estimate ψ using data, let

$$\begin{aligned} U(\psi, \zeta(\psi), \alpha) &= \sum_{i=1}^K \{G_i(\psi) - \mathbb{E}[G_i(\psi) | H_i; \zeta]\} \\ &\quad \times \{d_i(H_i, A_i) - \mathbb{E}[d_i(H_i, A_i) | H_i; \alpha]\}, \quad (11) \end{aligned}$$

where $d_i(H_i, A_i)$ is any function of H_i and A_i , and

$$G_i(\psi) = Y - \gamma_j(H_i, A_i; \psi) + \sum_{k=i+1}^K [\gamma_k(H_k, a_k^*; \psi) - \gamma_k(H_k, a_k; \psi)].$$

Consistent and asymptotically normal (CAN) estimators of ψ can be obtained using the estimating equations $\mathbb{E}[U(\psi, \zeta(\psi), \alpha)] = 0$, as shown in [11]. The estimate obtained in (11) is doubly-robust, meaning that the estimator $\hat{\psi}$ is consistent if either $\mathbb{E}[G_i(\psi) | H_i; \zeta]$ or $p_i(A_i = 1 | H_i; \alpha)$ is correctly specified.

Methods closely related to G-estimation based on *counterfactual regret* were developed in [5]. We do not discuss them here in the interest of space.

4 LEARNING OPTIMAL PATH-SPECIFIC POLICIES

We now consider how the methods for optimizing over all policies translate to optimizing path-specific policies. Some of the generalizations we consider are currently only known for single-stage decision problems.

Consider the generalization of Fig. 1 (b) to the longitudinal setting with mediators, shown (for two time points) in Fig. 2 (a). This causal model corresponds to the setting described in detail in [4], representing an observational longitudinal study of HIV patients. Here, W_0 represents the baseline variables of a patient, A_1, A_2 represent treatment assignments, which were chosen based on observed treatment history according to physician's best judgement, W_1, W_2 are intermediate and final outcomes (such as CD4 count or viral failure), and M_1, M_2 are measures of patient adherence to their treatment regimen. We are interested in finding policies $f_{A_1}(H_1), f_{A_2}(H_2)$ that optimize the effect of A_1, A_2 on W_2 that is either direct or via intermediate outcomes, but not via adherence, and where adherence is kept to that of a reference

treatment a'_1, a'_2 . Specifically, we are interested in choosing f_{A_1}, f_{A_2} to optimize the counterfactual expectation $\mathbb{E}[W_2(f_{A_1}, f_{A_2})]$, which expands via (7) to

$$\mathbb{E} \left[W_2 \left(\begin{array}{c} W_0, f_{A_1}(H_1), M_1(a'_1), \\ W_1(f_{A_1}(H_1), M_1(a'_1)), f_{A_2}(H_2), \\ M_2(a'_1, a'_2, W_1(f_{A_1}(H_1), M_1(a'_1)), M_1(a'_1)) \end{array} \right) \right). \quad (12)$$

The K stage version of the causal model in Fig. 2 (a) is given by a complete DAG on variables $W_0, A_1, M_1, W_1, \dots, A_K, M_K, W_K$, listed in topological order, with a hidden common cause U of W_0, \dots, W_K . Let α be all directed edges out of A_1, \dots, A_K . The general version of (12) is the expectation of W_K taken with respect to the distribution $p(W_K(f_\alpha))$, where f_α sets all edges $(A_i M_j)_{\rightarrow}$ to a'_i , and all other edges in α to a policy $f_{A_i}(H_i)$.

Identifiability of $p(W_K(f_\alpha))$ is given by the following corollary of results in [14], which can be viewed as a generalization of the *collapse of the g-formula* [10] to longitudinal mediation settings.

Theorem 1 *In the above model, with a positive observed data distribution $p(W_K(f_\alpha))$ is identified as*

$$\sum_{H_K, M_K} \prod_{i=1}^K \left\{ p(W_i | \bar{M}_{i-1}, \bar{W}_{i-1}, f_{A_i}(H_i)) \right. \\ \left. p(M_i | \bar{a}'_i, \bar{W}_{i-1}, \bar{M}_{i-1}) \right\} p(W_0) \quad (13)$$

As a consequence (12) is identified as

$$\sum_{\bar{W}_1, \bar{M}_2} \mathbb{E} \left[W_2 | f_{A_1}(H_1), f_{A_2}(H_2), \bar{W}_1, \bar{M}_2 \right] p(M_1 | a'_1, W_0) \times \\ p(W_1 | W_0, f_{A_1}(H_1), M_1) p(M_2 | \bar{W}_1, M_1, a'_1, a'_2) p(W_0).$$

We now discuss a number of strategies for finding optimal path-specific policies identified by (13).

Parametric Backwards Induction

A simple approach for finding optimal path-specific policies is to combine backwards induction with a maximum likelihood estimator for the conditional densities in the identifying functional (13).

Consider Fig. 2 (a), and let a_i and a'_i denote the active and reference levels of a_i , respectively. Beginning at the last stage $K = 2$, and assuming treatment is binary, the optimal decision, $f_{A_2}^*$, for a given patient with history (w_0, a_1, m_1, w_1) sets A_2 to either a_2 or a'_2 to maximize the response W_2 , while keeping M_2 at whatever value it would have attained under a sequence of reference interventions (a'_1, a'_2) :

$$f_{A_2}^* = \mathbb{I} \left(\mathbb{E}[W_2(a_2, M_2(\bar{a}'_2)) | H_2] > \mathbb{E}[W_2(a'_2, M_2(\bar{a}'_2)) | H_2] \right).$$

The optimal decision at the first stage, $f_{A_1}^*$, is given by

$$\mathbb{I} \left\{ \mathbb{E} \left[W_2 \left(\begin{array}{c} a_1, f_{A_2}^*(H_2), M_1(a'_1), W_1(a_1, M_1(a'_1)), \\ M_2(\bar{a}'_2, M_1(a'_1), W_1(a_1, M_1(a'_1))) \end{array} \right) \middle| H_1 \right] > \right. \\ \left. \mathbb{E} \left[W_2 \left(\begin{array}{c} a'_1, f_{A_2}^*(H_2), M_1(a'_1), W_1(a'_1, M_1(a'_1)), \\ M_2(\bar{a}'_2, M_1(a'_1), W_1(a'_1, M_1(a'_1))) \end{array} \right) \middle| H_1 \right] \right\}.$$

Under the causal model we described, the above counterfactuals can be estimated using a modification of (13) with no summations over, but instead conditioning on histories H_1 or H_2 . This approach easily generalizes to any number of decision stages. The difficulty here, as before, is the increasing amount of marginalizations that must be performed as the number of stages grows.

Path-Specific Policies Via Q-Learning

We now describe how to generalize Q-learning to path-specific policies, using the HIV example in Fig. 2 (a) to ground the discussion. Recall that in this example, we wish to set A_1 and A_2 with respect to edges into W_1, W_2 to maximize the outcome, while setting A_1 and A_2 to reference values a'_1, a'_2 for the purposes of edges into M_1, M_2 .

Simply defining Q-functions as expectations over value functions conditional on history does not work in our setting, since mediators behave in a counterfactually different way from either observed variables, or variables we wish to set according to a policy. Moreover, the sequential nature of the problem means the nested counterfactuals needed become quite involved to write down.

An alternative is to define Q-functions not in the observed data distribution, corresponding to Fig. 2 (a), but in a counterfactual distribution where A_1, A_2 behave as observed, *except* for the purposes of edges into M_1, M_2 , in which case they are counterfactually set to a'_1, a'_2 . The graph corresponding to this counterfactual world is shown in Fig. 2 (b), and can be viewed as a generalization of a single world intervention graph [9], where treatment variables are only intervened on for the purposes of certain outgoing edges. Note that descendant variables of M_1, M_2 are marked with a tilde to make clear that these variables are counterfactual and no longer equal to their observed counterparts.

The distribution corresponding to this situation is simply $p(\mathbf{V}(\alpha_{(A_1 M_1)_{\rightarrow}, (A_1 M_2)_{\rightarrow}, (A_2 M_2)_{\rightarrow}}))$, where α sets these edges to a'_1, a'_2 . For K stages, the distribution is $p(\mathbf{V}(\alpha_\alpha))$, where α are all edges of the form $(A_i M_j)_{\rightarrow}$, and α sets each such edge to the reference value a'_i . We have the following claim.

Theorem 2 *In K stage version of the model in Fig. 2(a), $p(\mathbf{V}(\alpha_\alpha))$ is identified as*

$$\bar{p}(\bar{W}_0, \bar{A}_1, \bar{M}_1, \bar{W}_1, \dots, \bar{W}_K, \bar{A}_K, \bar{M}_K) = \\ p(W_0) \prod_{i=1}^K p(W_i | M_i, A_i, H_i) p(A_i | H_i) p(M_i | \bar{a}'_i, H_i \setminus A) \quad (14)$$

We can now define Q-functions as value function expectations on this new distribution, and proceed with backwards induction as before. The only difference between

the previous formulation is how Q-functions parameters are fit. In particular, we must compensate for the fact that \tilde{p} above is not the observed data distribution. Define

$$\begin{aligned}\tilde{Q}_K(\tilde{H}_K, \tilde{A}_K; \gamma_K) &= \mathbb{E}[\tilde{W}_K | \tilde{A}_K, \tilde{H}_K], \\ \tilde{V}_K(\tilde{H}_K) &= \max_{a_K} \tilde{Q}_K(\tilde{H}_K, a_K; \gamma_K),\end{aligned}\quad (15)$$

and for $i = K - 1, \dots, 1$, define

$$\begin{aligned}\tilde{Q}_i(H_i, A_i; \gamma_i) &= \mathbb{E}[\tilde{V}_{i+1}(\tilde{H}_{i+1}) | \tilde{A}_i, \tilde{H}_i], \\ \tilde{V}_i(\tilde{H}_i) &= \max_{a_i} \tilde{Q}_i(\tilde{H}_i, a_i; \gamma_i).\end{aligned}\quad (16)$$

In our example in Fig. 2 (a), $K = 2$, $H_1 \equiv \{W_0\}$, $\tilde{H}_2 \equiv \{W_0, M_1, W_1\}$, and \mathbb{E} denotes expectations with respect to appropriate conditional distributions derived from \tilde{p} . Q-functions defined in this way can be used to obtain the optimal path-specific policy at each stage.

Theorem 3 *Given that each $\tilde{Q}_i, i = 1, \dots, K$ is specified correctly, the optimal treatment at stage i given H_i is equal to: $f_{A_i}^*(H_i) = \arg \max_{a_i} \tilde{Q}_i(H_i, a_i; \gamma_i)$. Since parameters γ are not generally known, they must be estimated from data. This can be done as follows.*

Theorem 4 *Assume models in the set $\{\tilde{Q}_i(\tilde{H}_i, \tilde{A}_i; \gamma_i), p(M_i | A_i, H_i; \phi) | \forall i\}$ are correctly specified. Then the estimation equations*

$$\begin{aligned}\mathbb{E} \left[\frac{\partial \tilde{Q}_K}{\partial \gamma_K} \{W_K - \tilde{Q}_K(A_K, H_K; \gamma_K)\} w_K(H_K; \hat{\phi}_K) \right] &= 0, \text{ and} \\ \mathbb{E} \left[\frac{\partial \tilde{Q}_i}{\partial \gamma_i} \{V_{i+1}(H_{i+1}) - \tilde{Q}_i(H_i, A_i; \gamma_i)\} w_i(H_i; \hat{\phi}_i) \right] &= 0,\end{aligned}$$

are consistent for γ_K and γ_i , where

$$w_i(H_i; \hat{\phi}_i) \equiv \frac{p(M_i | \tilde{A}_i = a', H_i; \hat{\phi}_i)}{p(M_i | \tilde{A}_i, H_i; \hat{\phi}_i)} \forall i = 1, \dots, K.$$

Path-Specific Value Search

For simplicity, we restrict attention to a single-stage decision problem, as shown in Fig. 1 (a), where we are interested in picking a policy that maximizes the counterfactual mean $\beta = \mathbb{E}[Y(A = f(W), M(a'))]$.

Consider a restricted class of path-specific policies \mathcal{F} with elements $f \equiv \{f(W), \tilde{f}\}$ the latter setting A to a constant value a' , regardless of W . Give any estimation strategy for the counterfactual mean under a path-specific policy, we can implement a direct search for the optimal policies within \mathcal{F} as before. By analogy with earlier discussion of value search, we give an IPW estimator for β which generalizes (9):

$$\mathbb{E} \left[\frac{Y \tilde{C}}{\tilde{\pi}(W; \hat{\psi})} \cdot \frac{p(M|A = a', W; \hat{\phi})}{p(M|A = f(W), W; \hat{\phi})} \right], \quad (17)$$




Figure 2: (a) A causal model that generalized Fig. 1 (b) by also considering mediators. (b) A “multiple world intervention graph,” representing the counterfactual situation where adherence levels are kept to a reference treatment level, but the chemical effect of drugs is operating normally.

and an estimator which generalizes (10):

$$\begin{aligned}\mathbb{E} \left[\frac{\tilde{C}}{\tilde{\pi}(W; \hat{\psi})} \frac{f(M|W, A = a'; \hat{\phi})}{f(M|W, f(W); \hat{\phi})} \{Y - \mathbb{E}[Y|f(W), M, W; \hat{\zeta}]\} + \right. \\ \left. \frac{\mathbb{I}(A = a')}{\pi_{a'}(W; \hat{\psi})} \{ \mathbb{E}[Y|f(W), M, W; \hat{\zeta}] - \sum_M \mathbb{E}[Y|f(W), M, W; \hat{\zeta}] p(M|W, A = a'; \hat{\phi}) \} \right],\end{aligned}\quad (18)$$

where the expectation is evaluated empirically, $\tilde{C} \equiv \mathbb{I}(A = f(W))$, $\pi_{a'}(W; \psi) = p(A = a' | W; \psi)$, $\tilde{\pi}(W; \hat{\psi}) \equiv \sum_a \pi_a(W; \psi) \mathbb{I}(a = f(W))$, and $\hat{\psi}, \hat{\phi}, \hat{\zeta}$ are fit by maximum likelihood. We have the following.

Theorem 5 *Under regularity assumptions referenced in the Appendix, the estimator in (17) is consistent and asymptotically normal (CAN) if the models in the set $\{\pi(W; \psi), p(M|W, A; \phi)\}$ are correctly specified, and the estimator in (18) is CAN in the union model, where any two models in the set $\{\pi(W; \psi), \mathbb{E}[Y|A, M, W; \zeta], p(M|W, A; \phi)\}$ are correctly specified.*

This claim, and the estimators above, are extensions of the results in [17].

Single-Stage G-Estimation For Path-Specific Policies

Results in [18] generalized optimal blip-to-zero functions to mediation settings, by positing the following SNMM $\gamma(A, W; \psi)$:

$$\mathbb{E}[Y(A, M(A = 0)) | W] - \mathbb{E}[Y(A = 0, M(A = 0)) | W]. \quad (19)$$

Note that the policy f_A^* maximizing $\mathbb{E}[Y(A = f_A^*(W), M(0))]$ can be directly obtained from γ via

$$f_A^*(W) = \arg \max_a \gamma(A, W; \psi).$$

Since ψ is not generally known, it must be estimated from data. The following is a consistent set of estimating equations for ψ , under assumptions described in [18]:

$$\mathbb{E} \left[\left(\frac{\mathbb{I}(A=1)p(M|A=0,W)}{p(A=1|W)p(M|A=1,W)} \{Y - \mathbb{E}[Y|A=1, M, W]\} - \frac{\mathbb{I}(A=0)}{p(A=0|W)} \{Y - \mathbb{E}[Y|A=1, M, W] + \gamma(1, W; \psi)\} \right) h(W) \right] = 0 \quad (20)$$

for $h(W)$ any $|\psi|$ -dimensional function of W .

Note that unlike SNMMs associated with overall policies, which can be defined for any number of treatments, SNMMs associated with path-specific policies have only been defined for a single treatment. A longitudinal generalization of these models is left to future work.

5 EXPERIMENTS

We now illustrate our methods via a dataset on HIV patients from Nigeria. The data consists of more than 50k treatment-naive HIV-1 infected patients who were enrolled in the Harvard PEPFAR/AIDS Prevention Initiative program prior to Oct 2010. The patients were put on courses of antiretroviral therapy (ART), with five standard first-line regimen, and were followed every six months for at least one year after the ART initiation. Patients stayed on their initial treatment plan unless they were experiencing toxicity within a first-line drug regimen and consequently moved to a second-line regimen. The data has records on demographics and clinical test results such as CD4 counts, viral loads, and toxicity measures at 6-month intervals.

In order to combat chronic diseases such as HIV, patients are required to follow long term use of medications. Full benefits of the medications are realized when patients take their medications as prescribed. Unfortunately, factors such as side effects, caused by drug toxicities, can be developed alongside the course of treatment and lead to poor adherence to the therapy. A primary measure of adherence is provided in the data as average percent adherence that is the total number of days that the patient has drug supply over the total number of days in the time period (6-month intervals).

We restricted our attention to the first year of follow up, including two decision points and picked (log) CD4 count at the end of the first year as the outcome of interest. CD4 count is one of the biomarkers that quantifies how well the immune system is functioning and higher values are more favorable. Drug toxicity and adherence can mediate the effect of treatment on outcome. Since reactions to drug toxicity and adherence behavior vary among patients, we consider the problem of finding policies that optimize the direct chemical effect of the drug

where the indirect effect mediated through drug toxicity and adherence are set to some reference levels. We run an additional experiment where only the effect through adherence is set to a reference level [4]. The latter results are provided in the Appendix for the interest of space.

The causal model can be represented by the DAG in Fig. 2 (a). Treatment decisions at the baseline and the first follow up are respectively denoted by A_1 and A_2 . W_1 is a dichotomized intermediate outcome that indicates whether the CD4 count is above 450 cells/mm³ at the the end of the same six month period. W_2 is the final outcome and denotes the log CD4 count at the end of the first year after ART initiation. W_0 includes all the baseline factors such as sex, age, and test results prior to any treatment initiation. Toxicity and adherence measures during the first six months and the first year after treatment initiation are collected in M_1 and M_2 , respectively. Drug toxicity indicates any lab toxicities (alanine transaminase ≥ 120 UI/L, creatinine ≥ 260 mmol/L, hemoglobin ≤ 8 g/dL), and adherence indicates whether average percent adherence was no less than 95%.

We first illustrate the results on finding optimal policies in a two-stage decision problem using G-formula and Q-learning methods, and then provide the results for finding optimal policies for the single (first) stage decision problem using value search and G-estimation. In the single stage analysis, we consider the log CD4 count at the end of the sixth month as the outcome of interest. All modeling assumptions are described in the Appendix.

Methods For The Two-Stage Decision Problem

Optimal overall polices and path policies are estimated as described in Sections 3 and 4. Expected outcomes under optimal policies we learned, along with 95% confidence intervals obtained by bootstrap, are shown in Table. 1. In the interest of space, we do not consider corrections necessary to preserve the validity of these intervals in cases of model discontinuities, but these are straightforward [1]. Both optimal polices have higher expected outcomes than the observed data, using either method. Path-specific policies led to higher expected outcomes compared to overall policies. This could be explained by the phenomenon of countervailing adherence-mediated effect, described in [4]. In other words, in some cases the more chemically effective drugs are also harder to take. In order to visualize optimal policies we learned as clinically interpretable flowcharts, we viewed policy-recommended decisions as class labels, and history as features, and used a multilabel decision tree classifier, as implemented in the `rpart` R package. The results are shown in Fig. 1 and 2 in the Appendix. Since the classifiers are not perfect predictors of the policies, they are

Table 1: Comparison of population log CD4 counts under different policies (under treatment assignments in the observed data, the value is 5.64 ± 0.01 in the 2-stage and 5.54 ± 0.01 in the 1-stage problem). G-formula and Q-learning are used with 2-stage decision points. Value search and G-estimation are used with 1-stage decision point.

	Path Policies	Overall Policies
G-formula	6.89 (5.76, 7.10)	6.79 (5.68, 6.90)
Q-learning	7.34 (6.10, 7.55)	6.89 (5.82, 7.12)
Value search	5.58 (5.54, 5.62)	5.56 (5.55, 5.58)
G-estimation	5.62 (5.50, 5.67)	5.79 (5.59, 6.04)

to be viewed as interpretable approximations of the true learned policy. Variables involved in the decisions, such as age, gender, CD4, virological status, and so on, are clinically reasonable.

Methods For The Single-Stage Decision Problem

We focus on patients that appear at baseline, and treat the five first-line treatments as a binary decision by grouping the first three (which we denote as group 1) and the last two treatments (which we denote as group 2). We are interested in finding a treatment group assignment at the first decision point that leads to a higher CD4 count at the first follow up, assuming adherence and toxicity are set to a reference point for every patient. The model for this setting is Fig. 1(a), where W is observed history up to the first decision.

We considered policies of the form $\mathbb{I}\{CD4_{m00} < \alpha\}$, where $CD4_{m00}$ denotes the CD4 count right before the first decision point, to decide what treatment group the patient should be assigned to. The normal range for this variable is 500 to 1500 cells/mm³. We searched for an optimal policy in this restricted class by varying α from 100 cells/mm³ to 1000 cells/mm³ by 50 cells/mm³ increments, and estimated the value for each α using (18). Under the modeling assumptions described in the Appendix, the optimal threshold between group 1 and group 2 treatments was chosen to be $\mathbb{I}\{CD4_{m00} < 550 \text{ cells/mm}^3\}$. In other words, if the CD4 count is less than 550, it is better for the patients to receive any of the treatments in the first treatment group. However, if we pick the optimal policy based on the overall effect of treatment on outcome, then value search leads to policies of the form $\mathbb{I}\{CD4_{m00} < 250 \text{ cells/mm}^3\}$. The optimal overall policy decides to switch to group 2 treatments when HIV gets more severe, while the optimal path policy decides on switching when CD4 count is still within a healthy range but at lower values. This implies that if treatment adherence could be kept to that of a reference treatment, initiation of treatments within group 2 could be delayed to lower CD4 values. As before, both policies lead to a

higher than observed log CD4 count, with the path policy yielding slightly higher outcomes than the overall policy.

Finally, we learned optimal path policies via G-estimation. We estimated the parameters ψ for the SNMM using (20). The population log CD4 count under path policies and overall policies learned by G-estimation are given in Table. (1), and a decision tree view of the policies is shown in Fig. 3 in the Appendix. Under our assumptions, the optimal path policy found by G-estimation did not do much better than the outcomes under observed treatments and the overall policies.

6 CONCLUSIONS

In this paper, we generalized ideas in mediation analysis and dynamic treatment regimes to consider the problem of estimation of responses to policies associated with particular causal pathways, and methods for learning policies that optimize these responses. Since validating findings in causal inference is difficult, and conclusions are sensitive to specific modeling assumptions made, we developed multiple approaches for learning optimal policies that rely on orthogonal sets of modeling strategies. In particular, we considered strategies based on backwards induction with either plug-in estimation or Q-learning, value search for restricted classes of policies, and G-estimation of structural nested mean models (SNMMs) generalized to mediation.

We illustrated these methods by finding policies for HIV-positive patients that optimize the direct chemical effect of antiretroviral therapy, where the indirect effect mediated through drug toxicity and adherence are set to a reference level. The results in the experiment section suggest that policies that aim to optimize the direct effect of the treatment have better outcome responses than policies that optimize the overall effect of the treatment, and optimal policies decide on clinically relevant variables, such as age, gender, viral load, and CD4 count.

The estimation methods we described are applicable to sequential decision problems, except for G-estimation which is currently only limited to single-stage decision problems. Generalizing the version of G-estimation to longitudinal mediation problems, and deriving semi-parametric estimators for the version of Q-learning we described are areas for future work.

Acknowledgments

This research was supported in part by the NIH grants R01 AI104459-01A1 and R01 AI127271-01A1. We thank the anonymous reviewers for their insightful comments that greatly improved this manuscript.

References

- [1] B. Chakraborty, E. B. Laber, and Y. Zhao. Inference for optimal dynamic treatment regimes using an adaptive m-out-of-n bootstrap scheme. *Biometrics*, 69(3), 2013.
- [2] B. Chakraborty and E. Moodie. *Statistical Methods for Dynamic Treatment Regimes: Reinforcement Learning, Causal Inference, and Personalized Medicine*. New York: Springer-Verlag, 2013.
- [3] J. W. Lavori and R. Dawson. A design for testing clinical strategies: Biased adaptive within-subject randomization. *Journal of the Royal Statistical Society*, 163:29–38, 2000.
- [4] C. Miles, I. Shpitser, P. Kanki, S. Meloni, and E. T. Tchetgen. Quantifying an adherence path-specific effect of antiretroviral therapy in the nigeria pepfar program. *Journal of the American Statistical Society*, 2017.
- [5] S. Murphy. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B*, 65:331–366, 2003.
- [6] S. A. Murphy. An experimental design for the development of adaptive treatment strategies. *Statistics in Medicine*, 24:1455–1481, 2005.
- [7] J. Pearl. Direct and indirect effects. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence (UAI-01)*, pages 411–420. Morgan Kaufmann, San Francisco, 2001.
- [8] J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2 edition, 2009.
- [9] T. S. Richardson and J. M. Robins. Single world intervention graphs (SWIGs): A unification of the counterfactual and graphical approaches to causality. *Working Paper 128, Center for Statistics and the Social Sciences, Univ. Washington, Seattle, WA.*, 2013.
- [10] J. M. Robins. A new approach to causal inference in mortality studies with sustained exposure periods – application to control of the healthy worker survivor effect. *Mathematical Modeling*, 7:1393–1512, 1986.
- [11] J. M. Robins. Optimal structural nested models for optimal sequential decisions. In *Proceedings of the second Seattle symposium on biostatistics*, pages 189–326, 2004.
- [12] J. M. Robins and S. Greenland. Identifiability and exchangeability of direct and indirect effects. *Epidemiology*, 3:143–155, 1992.
- [13] P. Schulte, A. Tsiatis, E. Laber, and M. Davidian. Q- and A-learning methods for estimating optimal dynamic treatment regimes. *Statistical Science*, 29:640–661, 2014.
- [14] I. Shpitser. Counterfactual graphical models for longitudinal mediation analysis with unobserved confounding. *Cognitive Science (Rumelhart special issue)*, 37:1011–1035, 2013.
- [15] I. Shpitser and E. Sherman. Identification of personalized effects associated with causal pathways. In *proceedings of the 34th Conference on Uncertainty in Artificial Intelligence*, 2018.
- [16] I. Shpitser and E. T. Tchetgen. Causal inference with a graphical hierarchy of interventions. *Ann. Stat.*, 2015.
- [17] E. J. T. Tchetgen and I. Shpitser. Semiparametric theory for causal mediation analysis: efficiency bounds, multiple robustness, and sensitivity analysis. *Ann. Stat.*, 40(3):1816–1845, 2012.
- [18] E. J. T. Tchetgen and I. Shpitser. Estimation of a semiparametric natural direct effect model incorporating baseline covariates. *Biometrika*, 101(4):849–864, 2014.
- [19] J. Tian and J. Pearl. On the identification of causal effects. Technical Report R-290-L, Department of CS, UCLA, 2002.
- [20] Y. Zhao, D. Zeng, A. J. Rush, and M. R. Kosorok. Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association*, 107:1106–1118, 2012.