

Inverse Preference Elicitation for Dynamic Treatment Regimes

Daniel J. Lizotte, Michael Bowling, Susan A. Murphy

danj1@umich.edu

bowling@cs.ualberta.ca samurphy@umich.edu



Abstract

Purpose: In order to provide decision support that allows room for clinician and patient preferences, we assess each potential treatment choice in a dynamic treatment regime, and determine “What health outcome preferences are consistent with this choice of treatment?”

Method: Dynamic treatment regimes can be constructed from data to recommend tailored treatment choices based on both baseline characteristics and on accumulating patient information (response to previous treatments, etc.) so as to optimize the mean value of a specified health outcome. The recommendations can then be used for decision support. This approach is problematic when there is no single outcome that is appropriate to optimize for all patients. For example, maximizing symptom relief and minimizing side-effect burden are both desirable objectives. However, depending on available treatments, the use of one or the other as the health outcome to be optimized will result in different dynamic treatment regimes. For example, a regime that maximizes expected symptom relief will tend to choose more aggressive drugs that are very effective but have a more severe side-effect burden. On the other hand, a regime that minimizes expected side-effect burden will choose drugs that are less effective but have a milder side-effect profile. Our more flexible approach is to consider a set of weighted combinations of conflicting health outcomes that reflect different possible clinician and patient preferences. We present a method that allows us to determine, for each possible treatment choice in a dynamic treatment regime, the range of preferences (say for symptom relief versus side-effect reduction) that are consistent with that treatment choice. This enables tailoring of treatments to both patient characteristics and preferences.

Result: Using data from the CATIE (Clinical Antipsychotic Trials of Intervention Effectiveness) trial, we illustrate how different treatment choices are consistent with different preferences for having better symptom relief versus better weight control.

Conclusion: Our methodology allows us to use dynamic treatment regimes to help inform clinical decision making while respecting patient and clinician preferences for optimizing different health outcomes.

Individualized Treatment Rules

A dynamic treatment regime consists of a sequence of individualized treatment rules. An individualized treatment rule takes as input a patient’s **state**, which consists of measurements about the patient that are relevant for predicting the patient’s **outcome** given different candidate **treatments**. For example, patient **state** may include characteristics such as age, gender, genetic markers, and response to previous treatments.

Given the value of a patient’s **state**, we make a prediction of what the **outcome** of the patient will be given each of the different candidate **treatments**. These predictions can be constructed from data using regression: The **state** and **treatments** are the dependent variables, and the **outcome** is the independent variable. Once we have the regression model, by holding the **state** part of the input fixed and varying the **treatment** we can find which treatment is predicted to result in the best **outcome** for a patient with that **state**.

Note that in order to operationalize this procedure, we must express our satisfaction with any particular patient **outcome** using a single numeric value.

Preference-Dependent Treatment Rules

There are few cases where encoding every patient’s outcome in the same way as a single numeric value captures everything we hope to achieve through treatment. In most settings, clinicians want to provide maximal symptom relief while inducing minimal side-effects and/or other burdens, but the relative importance of one outcome or the other will depend on many factors that we cannot anticipate at the time we construct the treatment rule.

Furthermore, when treating chronic diseases like major depressive disorder and schizophrenia, clinicians now have access to a gamut of possible treatments that provide varying levels of symptom relief and induce varying levels of side-effects. Given this set of candidate treatments, a decision rule constructed with an emphasis on maximizing symptom relief will not select the same treatment as a decision rule constructed with an emphasis on minimizing side-effects.

We propose an approach that first defines a set of outcomes that we might want to optimize, and then examines how the induced treatment rule changes as our **preference** changes from maximizing symptom relief to minimizing side-effect burden.

We begin by identifying two outcomes that are likely to be of concern for all patients, say O_0 and O_1 , both coded so higher is better. We consider all possible outcomes of the form

$$O_\delta = (1 - \delta) \times O_0 + \delta \times O_1$$

for $0 \leq \delta \leq 1$

The quantity δ thus determines the **preference** for optimizing outcome O_0 versus optimizing outcome O_1 : Given any *particular* δ , we can develop a treatment rule that optimizes the resulting single numeric outcome O_δ . By setting $\delta=0$ we produce a treatment rule that only “cares about” O_0 and by setting $\delta=1$ we produce a treatment rule that only “cares about” O_1 . Other values of δ “care about” both outcomes to some degree. In our example, O_0 is a measure of symptoms and O_1 is a measure of side-effect burden.

Inverse Preference Elicitation

If we knew the preference of a clinician or patient *a priori*, expressed as the quantity δ , we could use the outcome O_δ to construct an appropriate treatment rule and recommend a treatment. However, acquiring this preference through elicitation can be burdensome. Therefore, rather than ask “What treatment choice should we recommend given this clinician or patient’s particular value of δ ,” we ask instead, “**What set of preferences, expressed as ranges of δ , that are consistent with each choice of treatment?**”

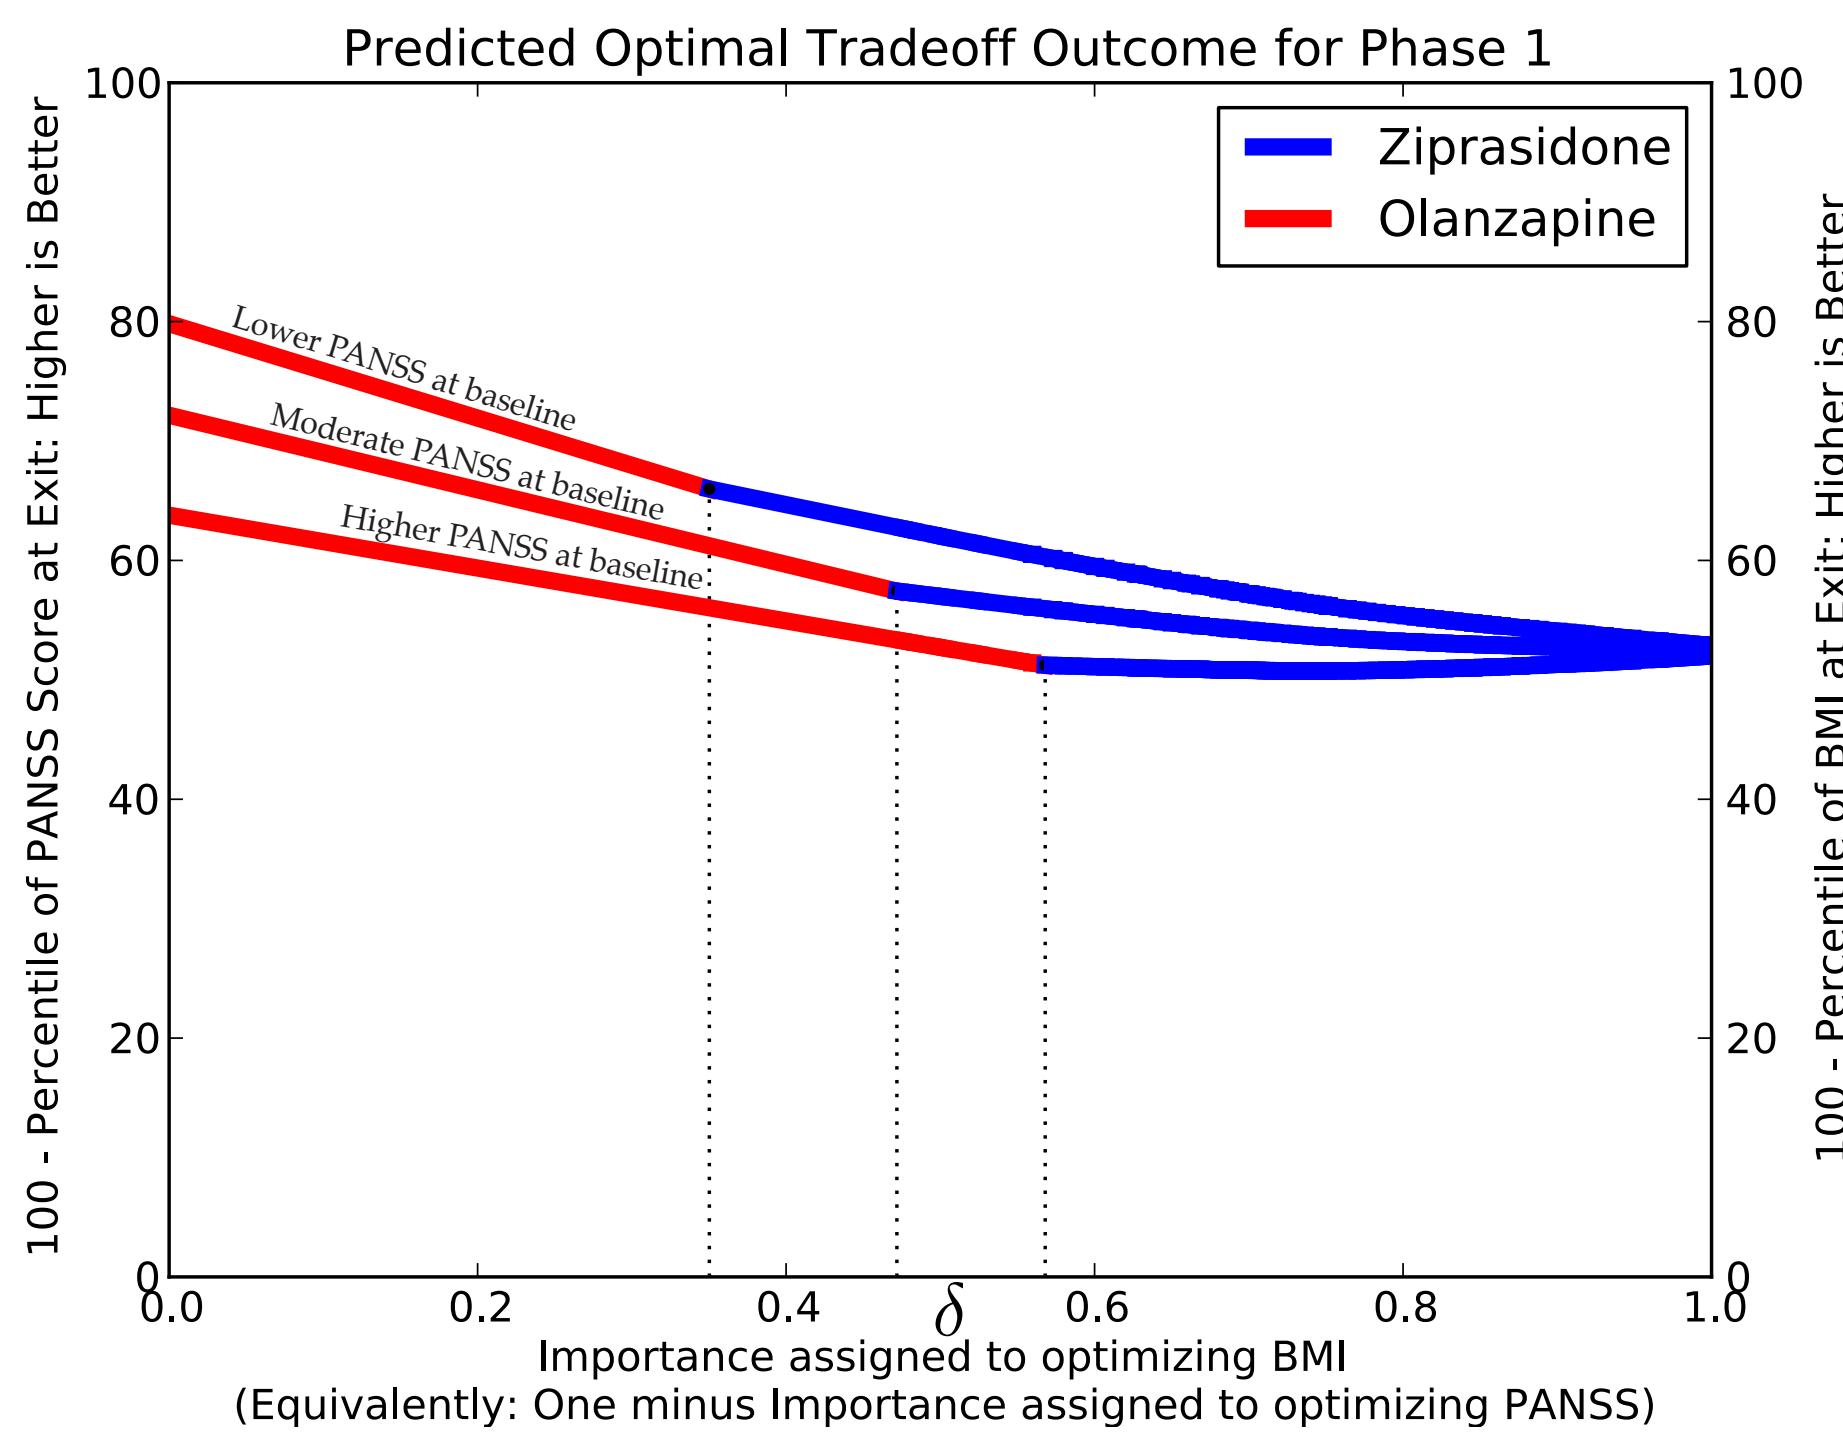
We call this analysis **Inverse Preference Elicitation**. Our goal is to use this approach to more effectively accommodate clinician and patient preferences when developing evidence-based clinical decision support systems.

I.P.E. in Dynamic Treatment Regimes

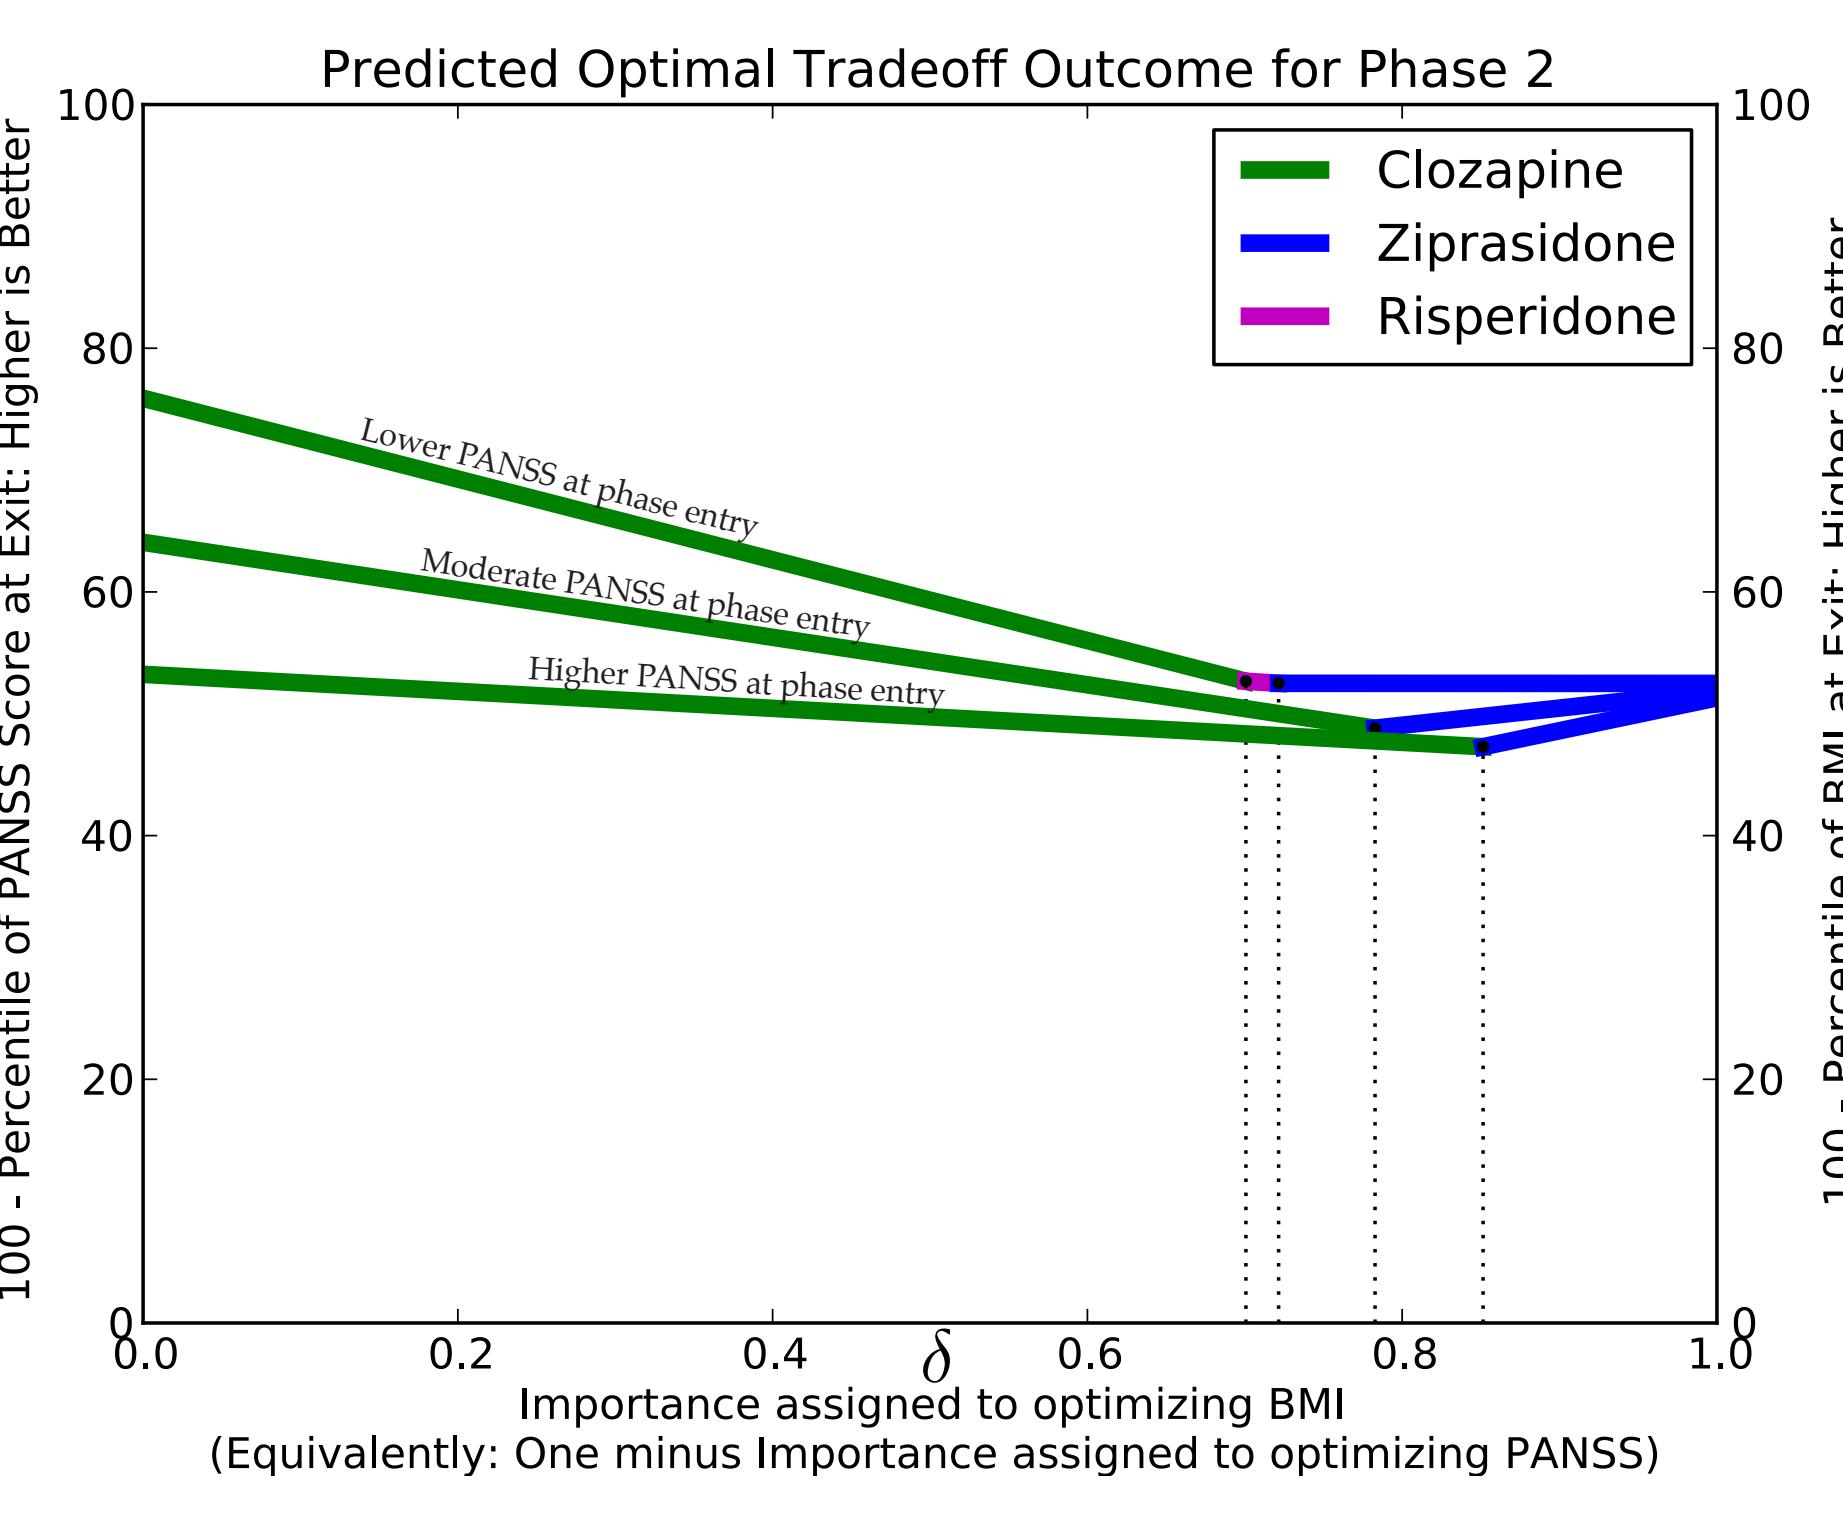
Our method can also be applied in the sequential decision making setting, where treatment choice is based not only on an immediate outcome, but also on *how we will treat this patient in the future*. Finding the optimal treatment in this setting is more complex, and typically involves **dynamic programming**. We have developed computational methods to do inverse preference elicitation in conjunction with **Q-learning**, which is an extension of regression that is used to find optimal treatments in dynamic treatment regimes.

Example Inverse Preference Elicitation Exploratory Analysis of CATIE

We now demonstrate the kind of output that our **Inverse Preference Elicitation** method can provide. As an illustrative example, we have chosen to examine a portion of the Clinical Antipsychotic Trials of Intervention Effectiveness (CATIE) data. Participants in CATIE were randomized more than once over the course of the study, so our analysis is done in the context of a dynamic treatment regime. Our two outcomes of interest are based on the Positive and Negative Syndrome Scale (PANSS), and Body Mass Index (BMI).



At Phase 2, patients in CATIE were randomized to one of five treatments: Clozapine, Olanzapine, Quetiapine, Risperidone, and Ziprasidone. Clozapine appears optimal when considering only PANSS, and Ziprasidone is optimal when considering only BMI. For patients who are less symptomatic at baseline, Risperidone appears optimal for a very small range of δ . Again the δ where more than one treatment is equally optimal shifts depending on the state of the patient, which in this case is their PANSS score at entry to Phase 2.



Summary of Exploratory Analysis

Recommendation given State and Preference	Strong Preference for Symptom Relief over Weight Control	Mild Preference for Symptom Relief over Weight Control	Mild Preference for Weight Control over Symptom Relief	Strong Preference for Weight Control over Symptom Relief
Lower PANSS at Entry to Phase 1	Olanzapine	Olanzapine or Ziprasidone	Ziprasidone	Ziprasidone
Moderate PANSS at Entry to Phase 1	Olanzapine	Olanzapine or Ziprasidone	Ziprasidone	Ziprasidone
Higher PANSS at Entry to Phase 1	Olanzapine	Olanzapine or Ziprasidone	Olanzapine or Ziprasidone	Ziprasidone
Lower PANSS at Entry to Phase 2	Clozapine	Clozapine	Clozapine, Risperidone, or Ziprasidone	Ziprasidone
Moderate PANSS at Entry to Phase 2	Clozapine	Clozapine	Clozapine	Clozapine or Ziprasidone
Higher PANSS at Entry to Phase 2	Clozapine	Clozapine	Clozapine	Clozapine or Ziprasidone

The above table is a possible summarization of the recommendations made by our approach. Note that this output does not include information the certainty with which we can make these recommendations. Incorporating information about measures of uncertainty in the recommendations is an important direction for future work.