

This is a repository copy of *Estimating the comparative effectiveness of feeding interventions in the paediatric intensive care unit: a demonstration of longitudinal targeted maximum likelihood estimation*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/118577/>

Version: Accepted Version

Article:

Kreif, Noemi, Tran, Linh, Grieve, Richard et al. (3 more authors) (2017) Estimating the comparative effectiveness of feeding interventions in the paediatric intensive care unit: a demonstration of longitudinal targeted maximum likelihood estimation. *American Journal of Epidemiology*. pp. 1-34. ISSN: 0002-9262

<https://doi.org/10.1093/aje/kwx213>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Estimating the comparative effectiveness of feeding interventions in the paediatric intensive care unit: a demonstration of longitudinal targeted maximum likelihood estimation

Journal:	<i>American Journal of Epidemiology</i>
Manuscript ID	AJE-00670-2016.R2
Manuscript Type:	Practice of Epidemiology
Key Words:	targeted maximum likelihood estimation, intensive care, time-dependent confounding, machine learning
Note: The following files were submitted by the author for peer review, but cannot be converted to PDF. You must view these files (e.g. movies) online.	
main.tex	

Estimating the comparative effectiveness of feeding interventions
in the paediatric intensive care unit: a demonstration of
longitudinal targeted maximum likelihood estimation

Noémi Kreif ^{*1}, Linh Tran², Richard Grieve^{3,4}, Bianca deStavola^{3,5}, Robert C
Tasker^{6,7}, and Maya Petersen²

¹Centre for Health Economics, University of York, York, UK

²Division of Biostatistics, UC Berkeley School of Public Health, Berkeley, CA, USA

³Centre for Statistical Methodology, London School of Hygiene and Tropical Medicine, London, UK

⁵Department of Medical Statistics, London School of Hygiene and Tropical Medicine, London, UK

⁴Department of Health Services Research and Policy, London School of Hygiene and Tropical
Medicine, London, UK

⁶Department of Anesthesiology, Perioperative and Pain Medicine, Division of Critical Care, Boston
Children’s Hospital, Boston, MA, USA

⁷Department of Neurology, Boston Children’s Hospital, Boston, MA, USA

March 31, 2017

For Peer Review

*corresponding author details: Noemi Kreif, Centre for Health Economics, University of York, Heslington, York, YO10 5DD, UK. Tel: work +44 (0)1904 321401 e-mail: noemi.kreif@york.ac.uk.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Word count: 4591 (text), 200 (abstract)

Running head: A demonstration of longitudinal targeted maximum likelihood estimation

Abbreviations:

CHIP: Control of Hyperglycemia in Paediatric Intensive Care; IPTW: inverse probability of treatment weighting; PICU: pediatric intensive care unit; TMLE: targeted maximum likelihood estimation

Risk Adjusted classification for Congenital Heart Surgery (RACHS-1)

Abstract

Longitudinal data sources offer new opportunities for the evaluation of sequential interventions. To adjust for time-dependent confounding in these settings, longitudinal targeted maximum likelihood based estimation (TMLE), a double-robust method that can be coupled with machine learning, has been proposed. This paper provides a tutorial in applying longitudinal TMLE, in contrast to inverse probability of treatment weighting and g-computation based on iterative conditional expectations. We apply these methods to estimate the causal effect of nutritional interventions on clinical outcomes of critically ill children. We estimate the risk of being discharged alive from the pediatric intensive care unit by a given day, under a range of static and dynamic regimes. We find that before adjustment, patients who follow the static regime 'never feed', are discharged by the end of the 5th day with a probability of 0.88 (95 % CI: 0.87 - 0.90), while for the patients who follow the regime 'feed from day 3', the probability of discharge is 0.64 (95% CI: 0.62 0.66). After adjusting for time-dependent confounding, most of this difference disappears, and the statistical methods provide similar results. TMLE offers a flexible estimation approach, hence we provide a practical guidance in implementation to encourage its wider uptake.

Large, observational databases such as electronic health records are increasingly used to answer questions of comparative effectiveness. The longitudinal structure of these datasets allows studies to estimate the effects of interventions that change over time. Examples include the treatment of chronic diseases such as diabetes and hypertension, where decisions such as when to initiate a treatment, change the dose or introduce a concomitant medication are repeatedly updated over time. For decision makers to compare the consequences of alternative longitudinal interventions, it is essential to carefully define the strategies of interest (1). A static regime for time-varying interventions pre-specifies the full sequence of interventions, irrespective of changing patient characteristics over time (for example: ‘always treat’). Dynamic regimes, or individualized treatment rules, in contrast, define a set of rules as a function of time-varying patient characteristics (2; 3; 4; 5; 6). While sequentially randomized trials provide an ideal design for the evaluation of dynamic regimes (7; 8) such trials are still relatively rare (9), and are impractical for many clinical and health policy questions (10).

Where the only available data for estimating the treatment effects of interest are from observational studies, statistical methods are required to address both baseline and time-dependent confounding. The latter arises in longitudinal settings, when the uptake of treatment may depend on factors that influence the outcome and are also affected by earlier treatments. It is widely recognized that standard regression analysis cannot deal with time-dependent confounding (2; 11). While in the last decades progress has been made in developing appropriate statistical methods for addressing time-varying confounding (see (12) for a review), applications of these approaches have been confined to relatively few clinical areas, such as human immunodeficiency virus infection (13). Further methodological and applied research is required that demonstrates these approaches for handling time-varying

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

confounding, in different contexts.

Inverse probability of treatment weighting (IPTW) (14; 11), a simple and intuitive method for estimating the effect of time-varying treatments, results in unstable and biased estimates, in the presence of data sparsity, even when weights are normalized and extreme weights are truncated (15). A commonly used alternative approach, parametric g-computation (2) requires parametric specification of multiple aspects of the data distribution, including models for the full conditional densities (or probability distributions) of the outcome and of the time-dependent confounders given the past. Similarly, structural nested mean models require a parametric model for the treatment effect, the ‘blip function’ (16). A perennial concern with these approaches is that they are prone to model misspecification, leading to biased estimates of treatment effects.

Targeted learning (17) has been proposed as a general approach for causal inference problems with both time-constant, and time-varying interventions to estimate a range of causal parameters (18; 19; 20; 21; 22). Targeted learning encompasses a semi-parametric, double-robust estimation approach, targeted maximum likelihood estimation (TMLE) (17) for single time point and longitudinal causal effects. TMLE combines estimates of the treatment and outcome mechanism, and provides a consistent estimator of the target parameter if either the treatment or the outcome mechanism are estimated consistently. If both are estimated consistently, TMLE is efficient (19). In order to reduce bias, achieve efficiency, and ensure accurate statistical inference, TMLE is often coupled with machine learning, in particular the Super Learner, a cross-validation based estimator selection approach (23; 24). For a tutorial for single time-point interventions, see for example (25).

This paper demonstrates the application of the longitudinal TMLE estimator based on

iterative conditional expectations (26; 21), and highlights how it is related to the IPTW (11) and g-computation (2; 26) estimators. While prior studies exist demonstrating the application of this estimator in longitudinal settings (26; 22; 27; 28; 29; 30) few observational data applications have discussed the use of the estimator to study the effects of dynamic regimes, or subject responsive adaptive treatment strategies (31).

We apply IPTW, g-computation based on iterative conditional expectations and longitudinal TMLE to an empirical study investigating an unanswered question of high relevance to clinical decision-makers: what is the optimal timing and quantity of caloric intake for critically ill children? We re-analyze the “Control of Hyperglycemia in Paediatric Intensive Care” (CHIP) clinical trial (32), to estimate the effect of alternative treatment regimes on the probability of being discharged from the Paediatric Intensive Care Unit (PICU) by a given day, under a range of clinically relevant treatment regimes. We follow the general targeted estimation roadmap (33) to formulate the hypothetical regimes, define the causal parameters of interest, and discuss how to identify, estimate and interpret these parameters.

METHODS

Research question in the CHIP study

An important objective of critical care medicine is to provide the appropriate level of nutritional support over the course of the patient’s stay. In most critical care settings, the preferred mode of nutritional support is via the nasogastric (enteral) tube, which is often complemented with intravenous (parenteral) feeding. For adult patients admitted to critical care, evidence-based guidelines exist, and recent randomized controlled trials have reported

that nutrition started early after admission to the PICU favorably alters outcomes (34). For critically ill children, guidelines for nutritional support are limited by the lack of available evidence (35). While a recent randomized trial of nutritional support for children admitted to the PICU (36) found that delaying parenteral nutrition led to favorable clinical outcomes, there is no randomized trial evidence to address more complex but important questions, such as what is the optimal timing and total quantity of nutritional support for children admitted to the PICU.

The recently published CHIP randomized trial, undertaken at 13 centers in England, with 1369 participants (recruited between 2008 and 2011, aged 0 to 16 years) found that tight glycemic control in critically ill children has no effect on the primary clinical outcome, the number of ventilator-free days (32). We undertake a secondary analysis of the CHIP dataset to investigate the causal effect of different levels of nutritional support on a clinical outcome. We focus on the subgroup who were admitted to the PICU to undergo cardiac surgery, and were younger than three years of age.

We follow the National Academy of Medicine (formerly Institute of Medicine) guidelines (37) in standardizing individual caloric intake by dividing the individual measures of daily intake by a target level, specific to the patient's gender, age, height and weight. We define a patient as 'fed' on a given day if he or she receives at least 20% of the individualized target. The outcome of interest is being discharged alive from the PICU to other hospital wards by a given day.

The effect of the feeding strategy on a given day on the patient's discharge status is subject to potential confounding from baseline characteristics such as age, sex, weight, height and the Risk Adjusted classification for Congenital Heart Surgery (RACHS-1) risk score, expressing

severity at admission (38), and the randomization arm (tight or standard glycemic control). Younger children, and those with higher risk scores tend to be fed less aggressively, and are also likely to stay longer in the PICU, potentially biasing the effect of feeding compared to no feeding towards a seemingly protective effect. While all patients are mechanically ventilated at baseline, being taken off mechanical ventilation is a strong predictor of discharge on the next days. Patients just taken off mechanical ventilation can, for safety reasons, be fed only through the parenteral (and not the enteral) route, making it less likely that their caloric intake reaches the 20 % threshold. Hence, a lack of adjustment for mechanical ventilation status could make no feeding appear beneficial. Further time-varying potential confounders include renal replacement therapy, infection, and a vasoactive inotrope score (39).

Observed data structure

For each patient, the level of calorific intake (the treatment) was measured daily from study entry (randomization) until the relevant clinical outcome was recorded (discharge from the PICU, or death while in the PICU). We restrict the follow-up data used in this analysis to the first seven days post randomization as the majority of patients were discharged from the PICU by this time point. We denote time by $t = 0, \dots, T + 1$, where $T + 1 = 7$ is the end of follow-up. At each time point, the patient's feeding status is represented by the binary variable A_t , while confounders are denoted by the multi-dimensional variable Z_t . The first measurement of the time-varying confounders and the vector of baseline confounders are jointly denoted with Z_0 . Mechanical ventilation, renal replacement therapy, infection and randomization arm are binary variables, while the RACHS-1 and vasoactive inotrope scores, weight, height and age are continuous.

M_t and Y_t indicate whether by the end of time period t , a patient has died ($M_t = 1$), or has been discharged alive from the PICU ($Y_t = 1$). We use overbars to denote histories, e.g., treatment history is denoted by $\bar{A}_t = (A_0, \dots, A_t)$. We assume that the observed data are n independent and identically distributed copies of $O = (Z_0, A_0, Y_1, M_1, Z_1, A_1, \dots, A_T, Y_{T+1}) \sim P_o$, where P_o is the true underlying distribution from which the data are drawn, and where, for notational convenience, we assume that variables after death or discharge are deterministically equal to their last observed values. The ordering of the elements of O represent their assumed causal ordering. For example, the baseline covariates (Z_0) precede the first instance of feeding (A_0), which precedes whether the patient is discharged by the end of the first day (Y_1).

We denote the history of the confounders, discharge status and death with a single vector $\bar{L}_t = (\bar{Z}_t, \bar{Y}_t, \bar{M}_t)$ that will be referred to as ‘covariates’. Death is treated as a competing event for PICU discharge: if $M_t = 1$, then any subsequent $Y_{t'} = 0$ for $t' > t$.

Formulating the interventions of interest

We consider two types of longitudinal interventions: static and dynamic treatment regimes. Let the vector $\bar{a} = (a_0, a_1, \dots, a_T)$ denote a longitudinal feeding regime, defined up to the last period before the end of follow-up. The elements of this vector, a_t define the feeding intervention, with static and dynamic regimes differing in how a_t is specified. For a static regime, a_t is a pre-specified constant for each t . For example, the static regime ‘never feed’, sets a_t to 0 in each t , resulting in the treatment regime $a = (0, 0, \dots, 0)$, while the static regime ‘feed from day 3’, would be defined as $a = (0, 0, 1, \dots, 1)$.

For dynamic regimes, a_t is set by a decision rule. We define d_t as a function that incor-

porates information available on a subject up to time t , such as some subset of the covariate history, denoted \bar{V}_t . We use $d(\bar{V}_t)$ to denote the vector of interventions required by regime d from time 0 to time t , given the realized covariate history. Here, we specify a dynamic treatment regime in which clinicians are required to feed a patient on each day that he or she is not mechanically ventilated ($Z_{t,1} = 1$.)

Clinical guidelines may not require the intervention to start on the first day, and could allow delaying the start of the intervention. For example, the regime 'feed by the third day' leaves the treatment values to be random for 2 days, then requires feeding from day 3. This regime is denoted by $\bar{a}_{2:T} = (A_0, A_1, 1, \dots, 1)$, where A_0 and A_1 are the observed levels of feeding on the first two days.

Throughout we consider regimes that implicitly only assign a feeding intervention up till time of discharge or death. Thus a subject who followed a regime of interest up to death or discharge, according to our definition, continues to follow this regime up to time $T + 1$. To simplify notation, in the sections that follow, we refer to the counterfactual interventions generally as $d = (d_0, \dots, d_T) \in \mathcal{D}$, where \mathcal{D} is the set of regimes of interest, and note that this notation includes static regimes as special cases of dynamic regimes.

Target causal parameter and identifying assumptions

The counterfactual discharge status at time period t that would have been observed under a given feeding regime d is denoted by Y_t^d . Our causal parameter of interest is the intervention-specific mean outcome, the expected discharge status by a selected time t^* , under a given regime d , where $t^* = 1, \dots, T + 1$:

$$\psi_{d,t^*} = E[Y_{t^*}^d].$$

ψ_{d,t^*} can be interpreted as the counterfactual cumulative risk of discharge by day t^* if all subjects had followed a given regime.

In order to identify ψ_{d,t^*} from the observed data, the following assumptions are required (2):

The **sequential randomization assumption** states that, conditional on the observed treatment and confounder history, the potential outcome is independent of treatment status in each preceding time period,

$$Y_{t^*}^d \perp A_t | \bar{L}_t, \bar{A}_{t-1} = d(\bar{V}_{t-1}),$$

for $t = 0, \dots, t^* - 1$, and $d \in \mathcal{D}$. This assumption requires that a sufficiently rich set of confounders are measured, so that it can be assumed that conditional on observed covariates, and following the regime of interest, the feeding decision at time t is ‘at random’.

The **positivity assumption** requires that for each feeding regime d , in each period t before the final time period of interest t^* , patients must have a positive probability to follow that regime, conditional on having followed it up to that time point, for any combination of observed covariate history:

$$Pr[A_t = d(\bar{V}_t) | \bar{L}_t, \bar{A}_{t-1} = d(\bar{V}_{t-1})] > 0,$$

for $t = 0, \dots, t^* - 1$.

Estimation

This section describes the IPTW estimator, the g-computation and longitudinal TMLE estimators, for interventions starting on the first day. For interventions with a delayed start, we provide small modifications of the estimators in the Web Appendix 2.

Inverse probability of treatment weighting

IPTW estimates the intervention-specific mean of a treatment regime by re-weighting the observed outcomes of the subset of the study sample who followed the regime (11). We denote the probability of a subject following a regime of interest d_t at time t , given her covariate and treatment history with

$$g_t = Pr[A_t = d_t(\bar{V}_t) | \bar{A}_{t-1} = d(\bar{V}_{t-1}), \bar{L}_t],$$

and the cumulative conditional probability of following regime d through time $t^* - 1$ as:

$$g_{0:t^*-1} = \prod_{t=0}^{t^*-1} g_t.$$

The stabilized Horvitz-Thompson IPTW estimator (14; 40) of the cumulative risk of discharge by period t^* under treatment regime d is based on estimating the following quantity:

$$E(Y_{t^*} \times I(\bar{A}_{t^*-1} = d(\bar{V}_{t^*-1})) / g_{0:t^*-1}) / E(Y_{t^*} \times I(\bar{A}_{t^*-1} = d(\bar{V}_{t^*-1})) / g_{0:t^*-1}),$$

where $I(\bar{A}_{t^*-1} = d(\bar{V}_{t^*-1}))$ indicates whether a patient has followed the treatment regime d up to one before the final period of interest. Implementation is based on estimating g_t for $t = 0, \dots, t^* - 1$, plugging in these estimates, and taking the empirical mean of the numerator and denominator.

Drawbacks of the IPTW estimator include reliance on consistent estimation of the treatment mechanism, as well as susceptibility to violations and near violations of the positivity assumption, resulting in unstable estimates (see for example (15)). The next section describes the longitudinal targeted maximum likelihood estimator, a double-robust estimator which can improve on the properties of the IPTW estimator by using information not only

on the treatment mechanism, but also on the outcome-confounders relationship.

The longitudinal targeted maximum likelihood estimator

The conditional expectation representation of the g-computation formula

The longitudinal TMLE (21) uses the identifiability result established by the g-computation formula (2). In short, the g-computation formula expresses the intervention-specific mean as a function of the conditional distributions of the outcome, and the time-varying confounders, given the past among subjects who followed the regime of interest. For discrete-valued confounders, this can be written as follows:

$$E[Y_{t^*}^d] = \sum_{\bar{l}_{t^*-1}} E[Y_{t^*} | \bar{A}_{t^*-1} = d(\bar{v}_{t^*-1}), \bar{L}_{t^*-1} = \bar{l}_{t^*-1}] \prod_{t=0}^{t^*-1} Pr(L_t = l_t | \bar{A}_{t-1} = d(\bar{v}_{t-1}), \bar{L}_{t-1} = \bar{l}_{t-1}),$$

where the summation is taken over all possible values \bar{l}_{t^*-1} of the confounder history.

Intuitively, the g-computation formula estimates the conditional expectation of the outcome under the treatment regime of interest, and averages these expectations over the intervened-on distribution of the confounders, i.e. the distribution that the confounders would take under the treatment regime of interest. Parametric g-computation (2; 41; 42) estimates the components of this formula directly, and makes strong parametric assumptions due to the need to specify conditional densities or probabilities for each of the time-varying confounders (12; 43).

The g-computation formula can be re-written as a series of iterated conditional expectations of the observed outcome (44; 45; 26):

$$E[Y_{t^*}^d] = E \left[\dots E \left[E \left[E \left[Y_{t^*} | \bar{A}_{t^*-1} = d(\bar{V}_{t^*-1}), \bar{L}_{t^*-1} \right] | \bar{A}_{t^*-2} = d(\bar{V}_{t^*-2}), \bar{L}_{t^*-2} \right] | \bar{A}_{t^*-3} = d(\bar{V}_{t^*-3}), \bar{L}_{t^*-3} \right] \dots \right] , \quad (1)$$

where the innermost expectation is the conditional distribution of the outcome, given the full treatment and confounder history, evaluated at the treatment values that would have been assigned according to the intervention of interest d . The second innermost expectation marginalizes over the intervened-on history of L_{t^*-1} , the next one over L_{t^*-2} , and so on, until the last expectation is taken over the empirical distribution of baseline confounders L_0 , where $M_0 = 0$ and $Y_0 = 0$. We first briefly review how to obtain the target parameter using these iterative regressions, then describe how the longitudinal TMLE extends this approach.

Steps of g-computation using sequential regressions

Step 1: Regress the outcome on full treatment and confounder history

First the innermost expectation of Equation 1 is estimated: $E[Y_{t^*} | \bar{L}_{t^*-1}, \bar{A}_{t^*-1} = d(\bar{V}_{t^*-1})]$. We will refer to this quantity as \bar{Q}_{t^*} . This expectation can be estimated by regressing the outcome on past covariates and treatment variables, for example using a logistic regression, and taking predictions at the treatment values corresponding to the intervention of interest.

Step 2: Take the previous predictions as the new outcome, regress on history up to $t^* - 2$.

The predictions from the previous step, \bar{Q}_{t^*} are now taken as the new outcome, and are regressed on confounders and treatment variables up to time period $t^* - 2$. As before,

predictions are generated for treatment values required by the regime \bar{d} , up to time period $t^* - 2$. This expectation, \bar{Q}_{t^*-1} , corresponds to the second innermost expectation in Equation 1. \bar{Q}_{t^*-1} is marginal over the intervened-on distribution of the time varying confounder L_{t^*-1} , but conditional on the time varying confounders up to time period L_{t^*-2} .

Steps 3, Step 4,..., to Step t^* : Iterate Step 2

Step 3 takes the predictions from Step 2, \bar{Q}_{t^*-1} and regresses them on the treatment and confounder history up to $t^* - 3$, then takes predictions as described above, stored as \bar{Q}_{t^*-2} . This step is iterated until the last step, where the expectation is only conditional on the baseline covariates: $\bar{Q}_1 = E[\bar{Q}_2 | \bar{L}_0, \bar{A}_0 = d(V_0)]$.

Step $t^* + 1$: Average over the empirical distribution of the baseline covariates

By averaging \bar{Q}_1 over the empirical distribution of the L_0 , the g-computation estimator for the intervention-specific mean is obtained as $\bar{Q}_0 = E[\bar{Q}_1]$.

Each \bar{Q}_t can be obtained using a regression, for example a linear or logistic regression. This approach offers substantial advantages over the parametric g-computation approach, by avoiding the need to estimate the conditional density of each time-varying confounder. However, estimating these iterative regressions well can be challenging and the approach remains susceptible to bias due to misspecification. Bang and Robins (26) propose a double robust and semiparametric efficient version of this sequential regression estimator based on including an additional, ‘clever’ covariate that uses information from the treatment assignment mechanism. Joffe (46) subsequently suggested moving this clever covariate to a weight, an approach that improved performance in the face of practical positivity violations. The resulting estimator is double robust in the sense that if either the treatment mechanisms or the sequential regressions are estimated consistently, then the estimator is consistent. If both

are estimated consistently, it is efficient in a semiparametric model that makes assumptions, if any, only on the treatment mechanism (26). Van der Laan and Gruber (21) subsequently placed this estimator in the general TMLE framework. The general idea behind this TMLE is that it is a two-step estimator: first, the conditional expectation of the outcome is estimated, then, this estimate is updated using information from the treatment assignment mechanism, targeted in a way that it reduces bias for the parameter of interest. Longitudinal TMLE performs the update step at each stage of the sequential regressions, as we summarize below.

The update step of the TMLE estimator

\bar{Q}_{t*} is defined and estimated as in Step 1 of the g-computation approach. This initial estimate is then updated, by perturbing the initial fit \bar{Q}_{t*} using a parametric submodel, defined as $\text{logit}(\bar{Q}_{t*}^1(\epsilon_{t*})) = \text{logit}(\bar{Q}_{t*}) + \epsilon_{t*}$. We estimate ϵ_{t*} by fitting a logistic regression of Y_{t*} on the intercept, using the prior predicted value of \bar{Q}_{t*} as offset, and weights corresponding to $I(\bar{A}_{t*-1} = d(\bar{V}_{t*-1}))/g_{0:t*-1}$, an indicator of whether a subject has followed the regime of interest up to the previous time period divided by the predicted probability of having done so. The estimated ϵ_{t*} is then used to update the initial estimate, which is stored as \bar{Q}_{t*}^1 , and will be used as the new outcome for the next iteration.

This update is performed after each step of the sequential regressions, described for the g-computation estimator. The regression and update steps are iterated until the last step, in which the updated expectation \bar{Q}_1^1 is only conditional on the baseline covariates. Analogous with the last step of the g-computation estimator, the TMLE estimator for the intervention-specific mean is obtained as $\bar{Q}_0^1 = E[\bar{Q}_1^1]$.

The consistency of the estimator relies on the consistent estimation of either the treatment

mechanism or the iterated conditional regressions while its efficiency relies on consistent estimation of both. In practice, often both components are expected to be misspecified, when fixed, parametric models such as logistic regressions are used. Machine learning or data adaptive approaches are thus advocated to estimate both. (19). We use the Super Learner (47), a machine learning algorithm that uses cross validation to find the optimal weighted convex combination of multiple candidate prediction algorithms, for estimating both the treatment assignment mechanism, as well as the sequential regressions (See the Web Appendix 3 for more details).

Implementation

We implement the IPTW, g-computation and TMLE estimators described above to estimate the cumulative probability of PICU discharge by the end of days 1-7, under a range of pre-specified static treatment regimes: never feed, feed from day 1,2,3,.. to 7, static regimes over limited time periods: feed by day 2,3,.. to 7, and the dynamic regime ‘feed when off ventilation’.

We use the Super Learner to estimate the treatment assignment mechanism and the sequential regressions, and use these models to construct the three estimators. Among the Super Learner candidates, we included an intercept model, a main terms and a logistic regression model with all possible two-way interactions in the linear predictor, a stepwise logistic regression, generalised additive models (48), a Bayesian generalized linear model with main terms in the linear predictor (49), lasso (50), boosting (51) and neural networks (52). We specified a 10-folds cross validation (47). We estimate separate models for the treatment assignment mechanism for each period, while assuming that treatment decisions are

influenced only by treatment and confounder values in the two most recent periods. The regressions to obtain the conditional probability of treatment and the iterative regressions and the update steps of the TMLE are only run among those who remain alive and not discharged. We contrast these estimates to ‘naive’ estimates, taken as the simple proportion of discharge status among those who follow a given regime.

The 95% confidence intervals are based on an estimate of the empirical influence function (53; 54) of the IPTW and TMLE estimators. For the g-computation estimator, no influence function based approach for inference is readily available, and the point estimates are reported without 95% confidence intervals. While the non-parametric bootstrap represents an alternative approach to variance estimation, when Super Learner is used to estimate the sequential regressions without subsequent targeting, bootstrapping can impose substantial computational burden while still failing to provide valid inference. The availability of an influence curve based variance estimator compatible with machine learning approaches is thus an additional attractive feature of the TMLE. The methods are implemented using the ltmle R package Version 0.9-9 (55; 56), which incorporates the Super Learner R package (57). We provide the main R functions used for the analysis in the Web Appendix 4.

RESULTS

706 children were included in the study sample. The number of patients who were still in the PICU at each time point, and among those, the number of patients receiving less than 20% of their daily caloric target (not fed), and those receiving at least 20% (fed) is reported in Table 1. Table 2 shows patient numbers observed to have followed each static regime of interest, up to a given day. Figure 1 contrasts the static regimes ‘never feed’, and ‘feed from

day 3', reporting the naive estimates, not adjusted for any observed confounders (Figure 1 A), and the IPTW (Figure 1 B) , g-computation (Figure 1 C) and TMLE estimates (Figure 1 D). See Web Table 1 in the Appendix, for illustrative calculation of the naive estimates for patients who followed the regime 'feed from day 3'.

The naive estimates indicate a significantly higher probability of being discharged by each day, for the 'never feed' regime, compared to the 'feed from day 3' regime. For example, the probability of discharge by the end of day 5 is 0.88 (95% CI: 0.87 - 0.90) for 'never feed', in contrast to the significantly lower estimate of 0.64 (95% CI: 0.62 - 0.66) for 'feed from day 3'. Adjustment for baseline and time-varying confounders shifts the estimated probability of discharge at each time period downwards, and reduces the difference between the two regimes. Using TMLE we estimate a 0.66 probability of discharge by the end of day 5 for those who were never fed (95 % CI: 0.59 - 0.72), and a 0.53 probability for those who were fed from day 3 (95% CI: 0.48 - 0.59). The TMLE, IPTW and g-computation estimators report similar point estimates, while the TMLE reported narrower 95% CIs compared to the IPTW estimator. For example, the probability of discharge by the end of day 5 for the 'feed from day 3' regime was estimated to be 0.54 (0.47 - 0.60) using IPTW, and 0.59 using g-computation. The smallest estimated cumulative probability of following a given regime, across all regimes considered was > 0.05 , so no weight truncation was used.

Focusing on the cumulative probability of discharge by day 4, Figure 2 contrasts the intervention-specific mean estimates across all regimes distinguishable by this time point, estimated by TMLE. The estimated probability of the static regime "feed from day 1" is 0.42, with the widest 95% CI among all regimes (0.21 - 0.62), which can be explained by the low number of patients following this regime. Regimes requiring starting feeding from or

by the third day, compared to starting on the second day, have lower expected probability of discharge; however, the 95% CIs overlap. As before, the regime ‘never feed’ has the most favorable expected outcomes (TMLE reporting an estimated probability of discharge of 0.63 (95% CI: 0.57 - 0.79); however, the probability of discharge under this regime was not statistically significantly different from the other regimes.

DISCUSSION

We implemented a double-robust approach, targeted maximum likelihood estimation, to contrast the counterfactual probability of being discharged alive from the PICU under a set of static and dynamic longitudinal feeding regimes, in a population of critically ill children. While the unadjusted estimates show a significant difference in discharge probabilities between the treatment regimes ‘start feeding from the 3rd day’, compared to ‘never feed’, after adjustment, most of this difference disappears. TMLE estimators report narrower confidence intervals than IPTW, as predicted by theory (17), while influence curve based confidence intervals for g-computation estimators are not readily available. We found no strong evidence that high levels of caloric intake may lead to adverse health outcomes in critically ill children. While in this paper the three statistical approaches led to similar conclusions, depending on the setting, TMLE may give substantially different results from estimation methods that are not double-robust, or don’t exploit data-adaptive model selection (see e.g. (27) for an application, and (28) for simulation evidence).

An observational analysis of data from a clinical trial enabled us to investigate the impact of alternative longitudinal feeding practices on clinical outcomes. We contributed to the literature of applying longitudinal causal methods in the PICU setting, where, due to the

1
2
3 fast changing prognosis of patients, and subsequently updated treatment decisions, time-
4
5 dependent confounding is an important concern (58). Using clinical judgement on meaningful
6
7 longitudinal treatment regimes, we selected a range of static and dynamic interventions which
8
9 were supported by the data, and asked new causal questions. While data collected in a clinical
10
11 trial can provide advantages, such as regular intervals of follow-up, measurement of a rich
12
13 set of observed and time-varying confounders, and little missing data, the approach taken
14
15 generalizes to settings of observational data.
16
17
18
19

20
21 The paper further provides a demonstration of the application of TMLE for longitudinal
22
23 static and dynamic regimes, and highlights how it builds on alternative approaches such as
24
25 IPTW and g-computation, under the challenging circumstances of a real-world comparative
26
27 effectiveness study: large number of covariates to adjust for, and a medium sized sample.
28
29 Application of the methods to address an unanswered clinical question of high relevance in
30
31 intensive care raised several methodological issues. Beyond static and dynamic treatment
32
33 regimes, we also considered interventions with a delayed start, for example “feed by day
34
35 3”, motivated by clinical practice. The availability of daily measurements of time-varying
36
37 confounders resulted in high dimensionality of observed covariates to adjust for. Informed by
38
39 clinical judgement, we assumed that the decision on whether to feed on a given day is only
40
41 influenced by observed characteristics measured on the given and on the previous day. To
42
43 deal with the challenge of model specification, we used the data-adaptive algorithm, Super
44
45 Learning.
46
47
48
49
50
51
52

53 Each of the methods applied here relies on the assumption that in each period, all time-
54
55 constant and time-varying confounders that can influence treatment assignment and the
56
57 outcome were observed. While the CHIP trial recorded a rich set of covariates, patients’
58
59
60

prognosis changes quickly over time, and the observed time-varying characteristics (mechanical ventilation, renal replacement, inotrope score) may not capture all confounders. In particular, if a clinician expects a patient to be discharged from the PICU soon, she may temporarily decrease, or not initiate enteral feeding, to prevent delay in discharge. Further research using methods to analyse the sensitivity of the parameter estimates to the presence of unobserved confounders is therefore warranted (59).

In summary, this paper illustrates that existing data sources such as well-conducted randomized controlled trials can be exploited, to address important questions of clinical decision making, beyond those originally posed. A wider use of appropriate causal methods could add to the understanding of the advantage of alternative sequencing of time-varying treatments, and provide estimates of the effectiveness and cost-effectiveness of realistic treatment strategies.

ACKNOWLEDGMENTS

Author affiliations:

Centre for Health Economics, University of York, York, UK (Noemi Kreif).

Division of Biostatistics, UC Berkeley School of Public Health, Berkeley, CA (Linh Tran, Maya Petersen)

Centre for Statistical Methodology, London School of Hygiene and Tropical Medicine, London, UK (Richard Grieve and Bianca DeStavola)

Department of Health Services Research and Policy, London School of Hygiene and Tropical Medicine, London, UK (Richard Grieve)

Department of Medical Statistics, London School of Hygiene and Tropical Medicine, London, UK (Bianca DeStavola)

Department of Anesthesiology, Perioperative and Pain Medicine, Division of Critical Care,
Boston Children's Hospital, Boston, MA (Robert C Tasker)

Department of Neurology, Boston Children's Hospital, Boston, MA (Robert C Tasker)

Dr. Noemi Kreif was supported by the UK Medical Research Council (Early Career
Fellowship in the Economics of Health, MR/L012332/1).

The authors thank Dr. Mark van der Laan and Josh Schwab for expert advice, and Dr.
Elizabeth Allen for data access.

Conflict of interest: none declared.

References

[1] Hernán MA. Counterpoint: Epidemiology to guide decision-making: moving away from practice-free research. *American journal of epidemiology*. 2015;182(10):834–839.

[2] Robins J. A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical Modelling*. 1986;7(9):1393–1512.

[3] Robins JM. Information recovery and bias adjustment in proportional hazards regression analysis of randomized trials using surrogate markers. In: *Proceedings of the Biopharmaceutical Section, American Statistical Association*. vol. 24. American Statistical Association; 1993. p. 3.

[4] Murphy SA, Van Der Laan M, Robins JM. Marginal mean models for dynamic regimes. *Journal of the American Statistical Association*. 2001;96(456):1410–1423.

[5] Murphy SA. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2003;65(2):331–355.

[6] Chakraborty B, Murphy SA. Dynamic treatment regimes. *Annual review of statistics and its application*. 2014;1:447.

[7] Murphy SA. An experimental design for the development of adaptive treatment strategies. *Statistics in medicine*. 2005;24(10):1455–1481.

[8] Bembom O, van der Laan MJ. Analyzing sequentially randomized trials based on causal effect models for realistic individualized treatment rules. *Statistics in medicine*. 2008;27(19):3689–3716.

[9] Parmar MK, Carpenter J, Sydes MR. More multiarm randomised trials of superiority are needed. *The Lancet*. 2014;384(9940):283–284.

[10] Hernán MA, Robins JM. Using big data to emulate a target trial when a randomized trial is not available. *American journal of epidemiology*. 2016;183(8):758–764.

[11] Robins JM, Hernan MA, Brumback B. Marginal structural models and causal inference in epidemiology. *Epidemiology*. 2000;11(5):550–560.

[12] Daniel R, Cousens S, De Stavola B, Kenward M, Sterne J. Methods for dealing with time-dependent confounding. *Statistics in medicine*. 2013;32(9):1584–1618.

[13] Cain LE, Robins JM, Lanoy E, Logan R, Costagliola D, Hernán MA. When to start treatment? A systematic approach to the comparison of dynamic regimes using observational data. *The international journal of biostatistics*. 2010;6(2).

[14] Horvitz DG, Thompson DJ. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*. 1952;47(260):663–685.

- [15] Petersen ML, Porter KE, Gruber S, Wang Y, van der Laan MJ. Diagnosing and responding to violations in the positivity assumption. *Statistical methods in medical research*. 2012;21(1):31–54.
- [16] Robins JM, Blevins D, Ritter G, Wulfsohn M. G-estimation of the effect of prophylaxis therapy for *Pneumocystis carinii* pneumonia on the survival of AIDS patients. *Epidemiology*. 1992;p. 319–336.
- [17] van der Laan MJ, Rubin D. Targeted maximum likelihood learning. *The International Journal of Biostatistics*. 2006;2(1):1–38.
- [18] Gruber S, van der Laan MJ. A targeted maximum likelihood estimator of a causal effect on a bounded continuous outcome. *The International Journal of Biostatistics*. 2010;6(1).
- [19] Van der Laan MJ, Rose S. Targeted learning: causal inference for observational and experimental data. Springer-Verlag New York; 2011.
- [20] Zheng W, van der Laan MJ. Targeted maximum likelihood estimation of natural direct effects. *The international journal of biostatistics*. 2012;8(1):1–40.
- [21] van der Laan MJ, Gruber S. Targeted minimum loss based estimation of causal effects of multiple time point interventions. *The international journal of biostatistics*. 2012;8(1):1–39.
- [22] Petersen M, Schwab J, Gruber S, Blaser N, Schomaker M, van der Laan M. Targeted maximum likelihood estimation for dynamic and static longitudinal marginal structural working models. *Journal of causal inference*. 2014;2(2):147–185.
- [23] van der Laan MJ, Dudoit S, Keles S. Asymptotic optimality of likelihood-based cross-validation. *Statistical Applications in Genetics and Molecular Biology*. 2004;3(1):1–23.
- [24] Van der Laan MJ, Polley EC, Hubbard AE. Super learner. *Statistical applications in genetics and molecular biology*. 2007;6(1):1–21.
- [25] Schuler MS, Rose S. Targeted Maximum Likelihood Estimation for Causal Inference in Observational Studies. *American Journal of Epidemiology*. 2016;185(1):65–73.
- [26] Bang H, Robins JM. Doubly robust estimation in missing data and causal inference models. *Biometrics*. 2005;61(4):962–973.
- [27] Decker AL, Hubbard A, Crespi CM, Seto EY, Wang MC. Semiparametric Estimation of the Impacts of Longitudinal Interventions on Adolescent Obesity using Targeted Maximum-Likelihood: Accessible Estimation with the ltmle Package. *Journal of Causal Inference J Causal Infer*. 2014;2(1):95–108.
- [28] Schnitzer ME, van der Laan MJ, Moodie EE, Platt RW. Effect of breastfeeding on gastrointestinal infection in infants: A targeted maximum likelihood approach for clustered longitudinal data. *The annals of applied statistics*. 2014;8(2):703.

[29] Tran L, Yiannoutsos CT, Musick BS, Wools-Kaloustian KK, Siika A, Kimaiyo S, et al. Evaluating the Impact of a HIV Low-Risk Express Care Task-Shifting Program: a case study of the targeted learning roadmap. *Epidemiologic Methods*. 2016;5(1):69–91.

[30] Brown DM, Petersen M, Costello S, Noth EM, Hammond K, Cullen M, et al. Occupational exposure to PM2. 5 and incidence of ischemic heart disease: longitudinal targeted minimum loss-based estimation. *Epidemiology (Cambridge, Mass)*. 2015;26(6):806.

[31] Neugebauer R, Schmittdiel JA, Laan MJ. Targeted learning in real-world comparative effectiveness research with time-varying interventions. *Statistics in medicine*. 2014;33(14):2480–2520.

[32] Macrae D, Grieve R, Allen E, Sadique Z, Morris K, Pappachan J, et al. A randomized trial of hyperglycemic control in pediatric intensive care. *New England Journal of Medicine*. 2014;370(2):107–118.

[33] Petersen ML, van der Laan MJ. Causal models and learning from data: integrating causal modeling and statistical estimation. *Epidemiology (Cambridge, Mass)*. 2014;25(3):418–426.

[34] McClave SA, Martindale RG, Rice TW, Heyland DK. Feeding the critically ill patient. *Critical care medicine*. 2014;42(12):2600–2610.

[35] Mehta NM, Compher C, et al. ASPEN Clinical Guidelines: nutrition support of the critically ill child. *Journal of Parenteral and Enteral Nutrition*. 2009;33(3):260–276.

[36] Fizez T, Kerklaan D, Mesotten D, Verbruggen S, Wouters PJ, Vanhorebeek I, et al. Early versus late parenteral nutrition in critically ill children. *New England Journal of Medicine*. 2016;374(12):1111–1122.

[37] Gerrior S, Juan W, Peter B. An easy approach to calculating estimated energy requirements. *Preventing chronic disease*. 2006;3(4):1–4.

[38] Al-Radi OO, Harrell FE, Caldarone CA, McCrindle BW, Jacobs JP, Williams MG, et al. Case complexity scores in congenital heart surgery: a comparative study of the Aristotle Basic Complexity score and the Risk Adjustment in Congenital Heart Surgery (RACHS-1) system. *The Journal of thoracic and cardiovascular surgery*. 2007;133(4):865–875.

[39] Wernovsky G, Wypij D, Jonas RA, Mayer JE, Hanley FL, Hickey PR, et al. Postoperative course and hemodynamic profile after the arterial switch operation in neonates and infants A comparison of low-flow cardiopulmonary bypass and circulatory arrest. *Circulation*. 1995;92(8):2226–2235.

[40] Hernán MA, Lanoy E, Costagliola D, Robins JM. Comparison of dynamic treatment regimes via inverse probability weighting. *Basic & clinical pharmacology & toxicology*. 2006;98(3):237–242.

[41] Young JG, Cain LE, Robins JM, O'Reilly EJ, Hernán MA. Comparative effectiveness of dynamic treatment regimes: an application of the parametric g-formula. *Statistics in biosciences*. 2011;3(1):119–143.

- [42] Snowden JM, Rose S, Mortimer KM. Implementation of G-computation on a simulated data set: demonstration of a causal inference technique. *American Journal of Epidemiology*. 2011;173(7):731–738.
- [43] Petersen ML. Commentary: Applying a Causal Road Map in Settings with Time-dependent Confounding. *Epidemiology*. 2014;25(6):898–901.
- [44] Robins JM. Analytic methods for estimating HIV-treatment and cofactor effects. In: *Methodological Issues in AIDS Behavioral Research*. Springer; 2002. p. 213–288.
- [45] Robins JM. Commentary on ‘Using inverse weighting and predictive inference to estimate the effects of time-varying treatments on the discrete-time hazard’. *Statistics in Medicine*. 2002;21(12):1663–1680.
- [46] Robins J, Sued M, Lei-Gomez Q, Rotnitzky A. Comment: Performance of double-robust estimators when "inverse probability" weights are highly variable. *Statistical Science*. 2007;22(4):544–559.
- [47] van der Laan MJ, Dudoit S. Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples; 2003. Working Paper 130. Available from: <http://biostats.bepress.com/ucbbiostat/paper130>.
- [48] Hastie T. *gam: Generalized Additive Models*; 2013. R package version 1.09. Available from: <http://CRAN.R-project.org/package=gam>.
- [49] Gelman A, Su YS. *arm: Data Analysis Using Regression and Multilevel/Hierarchical Models*; 2013. R package version 1.6-09. Available from: <http://CRAN.R-project.org/package=arm>.
- [50] Simon N, Friedman J, Hastie T, Tibshirani R. Regularization Paths for Cox’s Proportional Hazards Model via Coordinate Descent. *Journal of Statistical Software*. 2011;39(5):1–13. Available from: <http://www.jstatsoft.org/v39/i05/>.
- [51] Ridgeway G, et al. *gbm: Generalized boosted regression models*. R package version. 2006;1(3).
- [52] Venables WN, Ripley BD. *Modern Applied Statistics with S*. 4th ed. New York, New York: Springer; 2002. ISBN 0-387-95457-0. Available from: <http://www.stats.ox.ac.uk/pub/MASS4>.
- [53] Tsiatis A. *Semiparametric theory and missing data*. New York, New York: Springer; 2007.
- [54] Huber PJ. *Robust statistics*. New York, New York: Springer; 2011.
- [55] Schwab J, Lendle S, Petersen M, van der Laan M. *ltmle: Longitudinal Targeted Maximum Likelihood Estimation*; 2015. R package version 0.9-9,.

[56] R Core Team. R: A Language and Environment for Statistical Computing. Vienna, Austria; 2013. Available from: <http://www.R-project.org/>.

[57] Polley E, van der Laan M. SuperLearner: Super Learner Prediction; 2013. R package version 2.0-10. Available from: <http://CRAN.R-project.org/package=SuperLearner>.

[58] Vansteelandt S, Mertens K, Suetens C, Goetghebeur E. Marginal structural models for partial exposure regimes. *Biostatistics*. 2009;10(1):46–59.

[59] Brumback BA, Hernán MA, Haneuse SJ, Robins JM. Sensitivity analyses for unmeasured confounding assuming a marginal structural model for repeated measures. *Statistics in medicine*. 2004;23(5):749–767.

TABLES AND FIGURES

Table 1: Patient flow in the CHIP study, by treatment status and outcomes, on each hospital day

Hosp day	In PICU	In PICU, fed	In PICU, not fed	Cumulative dead	Cumulative discharged
1	706	28	678	0	0
2	701	260	441	0	5
3	597	387	210	0	109
4	434	340	94	0	272
5	325	278	47	3	378
6	248	222	26	5	453
7	188	169	19	7	511

PICU: pediatric intensive care unit

Table 2: Cumulative number of patients whose data are consistent with each static regime

Hosp day	Feed from..							Never feed
	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7	
1	28	678	678	678	678	678	678	678
2	24	241	442	442	442	442	442	442
3	24	220	254	270	270	270	270	270
4	22	212	237	205	197	197	197	197
5	21	205	232	195	178	173	173	173
6	21	202	226	192	176	165	165	165
7	21	200	223	186	175	164	163	162


Hosp: hospital

Figure 1: Estimated cumulative probabilities of discharge by end of days 1 to 7, for the regimes ‘never feed’ vs. ‘feed from day 3’


Part A: Unadjusted estimates and 95 % CIs; Part B: IPTW estimates and 95 % CIs, Part C: G-computation estimates; Part D: TMLE estimates and 95 % CIs. The x axis displays days, while the y axis displays the estimated counterfactual probability of discharge from PICU, by the end of a given day, for a given regime. The triangle marker stands for the regime ‘never feed’ (corresponding 95 % CIs displayed with dashed lines), while the circle marker displays the regime ‘feed from day 3’ (corresponding 95 % CIs displayed with solid lines). The 95 % CIs of the g-computation estimates are not reported.

Figure 2: Estimated cumulative probabilities of discharge by end of day 4


The x axis displays the regimes compared, while the y axis displays the TMLE estimates and corresponding 95 % CIs of the counterfactual probabilities of discharge from PICU, by the end of day 4, for a given regime. Abbreviations of regimes compared: From Day 1: Feed from day 1, From Day 2: Feed from day 2, By Day 2: Feed by day 2, From Day 3: Feed from day 3,..., Never: never feed, Dynamic: feed when off mechanical ventilation




279x361mm (300 x 300 DPI)




279x361mm (300 x 300 DPI)




279x361mm (300 x 300 DPI)







279x361mm (300 x 300 DPI)





279x361mm (300 x 300 DPI)



Web appendix for the article: "Estimating the comparative effectiveness of feeding interventions in the paediatric intensive care unit: a demonstration of longitudinal targeted maximum likelihood estimation"

Noémi Kreif ^{*1}, Linh Tran², Richard Grieve^{3,4}, Bianca deStavola^{3,5}, Robert C Tasker^{6,7}, and Maya Petersen²

¹Centre for Health Economics, University of York, York, UK

²Division of Biostatistics, UC Berkeley School of Public Health, Berkeley, CA, USA

³Centre for Statistical Methodology, London School of Hygiene and Tropical Medicine, London, UK

⁵Department of Medical Statistics, London School of Hygiene and Tropical Medicine, London, UK

⁴Department of Health Services Research and Policy, London School of Hygiene and Tropical Medicine, London, UK

⁶Department of Anesthesiology, Perioperative and Pain Medicine, Division of Critical Care, Boston Children's Hospital, Boston, MA, USA

⁷Department of Neurology, Boston Children's Hospital, Boston, MA, USA

March 31, 2017

Web Appendix 1

Causal model

We assume the following nonparametric structural equation model, \mathcal{M} :

$$Y_t = f_{Y_t}(\overline{Y}_{t-1}, \overline{M}_{t-1}, \overline{Z}_{t-1}, \overline{A}_{t-1}, U_{Y_t}), \text{ for } t = 1, \dots, T + 1$$

$$M_t = f_{M_t}(\overline{Y}_t, \overline{M}_{t-1}, \overline{Z}_{t-1}, \overline{A}_{t-1}, U_{M_t}), \text{ for } t = 1, \dots, T$$

$$Z_t = f_{Z_t}(\overline{Y}_t, \overline{M}_t, \overline{Z}_{t-1}, \overline{A}_{t-1}, U_{Z_t}), \text{ for } t = 0, \dots, T$$

$$A_t = f_{A_t}(\overline{Y}_t, \overline{M}_t, \overline{Z}_t, \overline{A}_{t-1}, U_{A_t}), \text{ for } t = 0, \dots, T,$$

where $U_t = (U_{Y_t}, U_{M_t}, U_{Z_t} \text{ and } U_{A_t}), t = 0, \dots, T + 1$ are unmeasured exogenous random variables from some underlying probability distribution P_U . This causal model specifies how each of the variables in the data are generated, with randomness arising only from the exogenous variables U . For example, the outcome at a given time period, Y_t is a deterministic function of the full history of treatment and confounder values, and a random error. Y_t and M_t are also functions of previous values of Y and M , encoding the information that after a patient is discharged, she always remains discharged, but if a patient dies in a given time period, she remains dead, and can never be discharged. More generally, after an event of death or discharge, all the processes become degenerate, and for notational convenience we assume that they take the last value observed. For notational convenience the causal model allows for Y_0 , and M_0 , which are both assumed to take value 0 (at baseline no one is dead or discharged), and Z_{-1}, A_{-1} which are assumed to be empty vectors.

*corresponding author details: Noemi Kreif, Centre for Health Economics, University of York, Heslington, York, YO10 5DD, UK. Tel: work +44 (0)1904 321401 e-mail: noemi.kreif@york.ac.uk.

Web Appendix 2

Modification of the estimators for the static regimes with delayed start

For regimes where the intervention starts with a delay, such as ‘feed by day k ’, corresponding to an intervention beginning at $t = k - 1$, the A_t nodes denoting feeding prior to time $k - 1$ are treated as non-intervention nodes or ‘covariates’. As a result, the baseline covariates (measured prior to the first intervention node A_{k-1}) consist of

$$(Z_0, A_0, \dots, Y_{k-1}, M_{k-1}, Z_{k-1}).$$

When estimating $E[Y_{t^*}^d], t^* = k, \dots, T + 1$, there are thus a total of $t^* - (k - 1)$ rather than t^* intervention nodes, and a corresponding number of components to the regimes of interest ($d(\bar{V}_t) = d_{k-1}(\bar{V}_{k-1}), \dots, d_{t^*-1}(\bar{V}_{t^*-1})$).

The IPTW, g-computation, and TMLE estimators are modified accordingly. First, the indicator of following a regime of interest through time $t^* - 1$, $I(\bar{A}_{t^*-1} = d(\bar{V}_{t^*-1}))$, used in the numerator of the weights for the IPTW and TMLE estimators, corresponds to an indicator of following the regime from time $k - 1$ to $t^* - 1$. (In other words, all subjects follow the regime of interest before $k - 1$). Second, the cumulative probability of following the regime of interest, used in the denominator of the weights for the IPTW and TMLE estimators, is now based on a product of time point-specific probabilities of continuing to follow the regime beginning at time $k - 1$:

$$g_{k-1:t^*-1} = \prod_{t=k-1}^{t^*-1} g_t(A_t = d_t(\bar{V}_t) | \bar{A}_{t-1} = d(\bar{V}_{t-1}), \bar{L}_t).$$

Finally, the presence of fewer intervention nodes implies that the longitudinal g-formula can be expressed using $t^* - (k - 1)$ rather than t^* iterated conditional expectations; one conditional expectation is needed for each intervention node. Thus implementation of the sequential regression g-computation and TMLE estimators requires fitting fewer conditional

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

regressions.

For Peer Review

Web Appendix 3

Super Learning estimation of the treatment and outcome mechanism

The Super Learner (1) is a machine learning algorithm that uses cross validation to find the optimal weighted convex combination of multiple candidate prediction algorithms. The algorithms are pre-selected by the analyst, potentially including parametric and non-parametric regression models, as well as a range of machine learning approaches. Asymptotically, the Super Learner algorithm performs as well as the best possible combination of the candidate estimators, assuming that none of the candidates in the library is a correctly specified parametric model; in the latter case it achieves almost parametric rate of converge (see (2) and (3) for details). Beyond its use for prediction (4; 5), it has been used for estimating the propensity score and the outcome model to obtain causal parameters (for example, (6; 7; 8)), and has been shown to reduce bias from model misspecification (9; 10; 11).

Web Table 1: Patient flow and unadjusted estimates: regime ‘feed from day 3’

Hosp day	In PICU (t)	In PICU & follows (t))	Dischg event (t+1)	Death (t+1)	Stops following (t+1)	Cum. follows (t)	Cum. discharge (t)	Prob dischg by t+1 follows t
1	706	678	5	0	236	678	5	0.007
2	701	437	77	0	265	442	82	0.186
3	597	172	23	0	17	254	105	0.413
4	434	132	33	0	5	237	138	0.582
5	325	94	10	0	6	232	148	0.638
6	248	78	21	0	3	226	169	0.748
7	188	54	11	0	0	223	180	0.807

Cumulative discharge is calculated amongst those whole followed the rule. The last column corresponds the unadjusted estimates reported in Figure 1.

Web Appendix 4

Main R functions used in the analysis

```
##### deterministic Q function #####

# objective: set Q to 1 deterministically if any prior L2=1

my.det.Q.fun <- function(data, current.node, nodes, called.from.estimate.g) {
  dnodes <- grep("L2", names(data))
  if (! any(dnodes < current.node)) { # outputs FALSE if there is no death node
    before currnt Anode
    return(NULL)
  }
  dnodes <- dnodes[dnodes < current.node] # only look at death nodes before current
  node
  dnodes.is1 <- data[, dnodes, drop=FALSE] == 1 & !is.na(data[, dnodes, drop=FALSE
    ]) # true if dnode is 1 (and not NA)
  dead <- apply(dnodes.is1, 1, any)
  return(list(is.deterministic=dead, Q.value=0))
}

#### MAIN FUNCTION ACTUALLY CALLING TMLE, it takes different arguments for
  different kinds of interventions

#### It calculates iptw, tmlw and gcomp, for 2 interventions. One is never feed (
  static), this stays fixed, I call it control.
```



```

1
2
3
4 ##### The other one is a "treatment" regime, either static, or with delayed
5     intervention ( v ) needs to be specifed, or dynamic.
6
7
8
9
10 ##### Dynamic intervention is currently a rule based on the presence of mechanical
11     ventilation each day
12
13
14
15
16
17 my.ltmle.contrast <- function(time, treatment, adjusted, sl) {
18     d.time <- pick_data(time)
19     n <- nrow(d.time$d)
20
21
22
23
24
25     abar.static.0 <- rep(0, time)
26
27     if (identical(treatment, "dynamic")) {
28         ### dynamic 1: mech vent #####
29         abar.1 <- as.matrix(d.time$d[, paste0("L4.", 0:(time - 1))])
30         intervene.time <- 1:time
31     } else {
32         set.to.1 <- if (treatment$day <= time) treatment$day:time else NULL
33         if (treatment$delay) {
34             intervene.time <- set.to.1
35         } else {
36             intervene.time <- 1:time
37         }
38
39         ### static or delayed: intervention starts on day
40         abar.1 <- matrix(0, nrow = n, ncol = time)
41         abar.1[, set.to.1] <- 1
42     }
43
44     abar.1.subset <- abar.1[, intervene.time, drop = FALSE]
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

```

```

1
2
3
4
5
6   if (adjusted) {
7       my.qform <- qform.generate(time)
8
9       my.gform.0 <- gform.generate(time)
10
11      my.gform.1 <- my.gform.0[intervene.time]
12
13  } else {
14
15      my.qform <- qform.generate.unadjusted(time) # intercept only for Q
16
17      my.gform.0 <- matrix(0, nrow=n, ncol=time) #set P(A=0)=1; unadjusted estimates
18          are reported by setting the g matrix to 1s
19
20      my.gform.1 <- abar.1.subset
21
22  }
23
24
25
26
27  result.list <- list()
28
29  for (gcomp in c(FALSE, TRUE)) {
30
31      # run ltmle for control (always static)
32
33      result.0 <- ltmle(d.time$d, Anodes=d.time$my.A.nodes, Lnodes=d.time$my.L.nodes,
34          Ynodes=d.time$my.Y.nodes, abar=abar.static.0, SL.library=sl, estimate.time=
35          FALSE, survivalOutcome=TRUE, variance.method='ic', deterministic.Q.function=
36          my.det.Q.fun, gform=my.gform.0, Qform=my.qform, gcomp=gcomp)
37
38
39
40
41
42      # run ltmle for treated
43
44      result.1 <- ltmle(d.time$d, Anodes=d.time$my.A.nodes[intervene.time], Lnodes=d.
45          time$my.L.nodes, Ynodes=d.time$my.Y.nodes, abar=abar.1.subset, SL.library=sl
46          , estimate.time=FALSE, survivalOutcome=TRUE, variance.method='ic',
47          deterministic.Q.function=my.det.Q.fun, gform=my.gform.1, Qform=my.qform[
48          intervene.time], gcomp=gcomp)
49
50
51
52
53
54
55      result.list <- c(result.list, GetAllResults(result.0, result.1, gcomp))
56
57  }
58
59
60

```

```
1
2
3
4
5
6
7
8     return(result.list)
9
10 }
11
12
13 GetConfInt <- function(result, estimator) {
14
15     c(summary(result, estimator)$treatment$estimate, summary(result, estimator)
16
17         $treatment$CI)
18
19 }
20
21
22
23 GetResults <- function(result.0, result.1, estimator) {
24
25     x <- list(GetConfInt(result.0, estimator), GetConfInt(result.1, estimator))
26
27     names(x) <- paste("est", c("ctrol", "tr"), estimator, sep = ".")
28
29     return(x)
30
31 }
32
33
34
35 GetAllResults <- function(result.0, result.1, gcomp) {
36
37     if (gcomp) {
38
39         GetResults(result.0, result.1, "gcomp")
40
41     } else {
42
43         c(GetResults(result.0, result.1, "tmle"), GetResults(result.0, result.1, "iptw
44
45             "))
46
47     }
48
49 }
50
51
52
53
54
55
56
57
58
59
60
```

References

- [1] van der Laan MJ, Dudoit S. Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples; 2003. Working Paper 130. Available from: <http://biostats.bepress.com/ucbbiostat/paper130>.
- [2] Van der Laan MJ, Polley EC, Hubbard AE. Super learner. Statistical applications in genetics and molecular biology. 2007;6(1):1–21.
- [3] Van der Laan MJ, Rose S. Targeted learning: causal inference for observational and experimental data. Springer Science & Business Media; 2011.
- [4] Polley EC, van der Laan MJ. Super learner in prediction; 2010. Working Paper 266. Available from: <http://biostats.bepress.com/ucbbiostat/paper266>.
- [5] Rose S. Mortality risk score prediction in an elderly population using machine learning. American journal of epidemiology. 2013;177(5):443–452.
- [6] Eliseeva E, Hubbard AE, Tager IB. An Application Of Machine Learning Methods To The Derivation Of Exposure-Response Curves For Respiratory Outcomes;. Working Paper 309. Available from: <http://biostats.bepress.com/ucbbiostat/paper309>.
- [7] van der Laan MJ, Luedtke AR. Targeted learning of an optimal dynamic treatment, and statistical inference for its mean outcome; 2014. Working Paper 329. Available from: <http://biostats.bepress.com/ucbbiostat/paper329>.
- [8] Gruber S, Logan RW, Jarrín I, Monge S, Hernán MA. Ensemble learning of inverse probability weights for marginal structural modeling in large observational datasets. Statistics in medicine. 2015;34(1):106–117.
- [9] Porter KE, Gruber S, van der Laan MJ, Sekhon JS. The relative performance of targeted

maximum likelihood estimators. The international journal of biostatistics. 2011;7(1):1–34.

[10] Kreif N, Gruber S, Radice R, Grieve R, Sekhon JS. Evaluating treatment effectiveness under model misspecification: a comparison of targeted maximum likelihood estimation with bias-corrected matching. Statistical methods in medical research. 2016;25(5):2315–2336.

[11] Pirracchio R, Petersen ML, van der Laan M. Improving Propensity Score Estimators’ Robustness to Model Misspecification Using Super Learner. American journal of epidemiology. 2015;181(2):108–119.