UNLEARN

Using AI-based Prognostic Models to Design Efficient, Unbiased Clinical Trials

Statistical principles of clinical trials with Digital Twins

Unlearn® works with pharma, biotech companies, and academic researchers to optimize human clinical trials by using a cutting-edge machine learning platform and data. We generate Digital Twins: disease- and study-specific, longitudinal, patient-level placebo clinical predictions that increase confidence in observing the treatment effect and reduce the required number of enrolled patients while keeping regulatory risk in check. We value statistical rigor and validated innovation. Our research and methods are published in peer reviewed papers and listed at Unlearn.AI.

Introduction

Rising costs and timeline delays continue to burden drug development. A growing resource of historical patient data have highlighted the opportunity for novel clinical trial designs that leverage these historical data sources to make clinical trials more efficient. Unfortunately, using historical data sources to make clinical trials more efficient while maintaining evidentiary standards required of late-stage clinical trials is a difficult task. For example, simply incorporating real world or historical data into the control arm of a clinical trial may increase statistical power, but this often comes at the expense of increased bias and loss of strict type-1 error rate control — even if propensity score-based methods are used to balance observed covariates. As a result, it's generally considered suboptimal to incorporate historical information into pivotal clinical trials that are used to inform regulatory decisions.

Fortunately, recent progress in Artificial Intelligence (AI) and Machine Learning (ML) provides an avenue for using historical data to create more efficient trials without introducing bias. Rather than incorporating data from external sources directly into the trial, we leverage predictions from AI-based prognostic models — called Digital Twins — trained on historical control data to reduce uncertainty in estimated treatment effects. This enables optimally efficient clinical trial designs that require fewer subjects to achieve pre-specified statistical power while rigorously controlling bias and type-I error rates.

Prognostic Covariate Adjustment (PROCOVA™)

It's well-known that regression techniques can be used to estimate treatment effects from Randomized Controlled Trials (RCTs)[1]. Here, we describe how similar methods can be used to leverage AI-based prognostic models trained on historical control patient data to create more efficient clinical trial designs that don't introduce bias.

Analysis of covariance (ANCOVA) is a technique for estimating average treatment effects from RCTs[2]. In an ANCOVA analysis, the treatment effect on a continuous outcome is estimated by regressing the outcome against a treatment indicator variable and other baseline covariates that may be correlated with the outcome. For example, the effect of an experimental treatment on a cognitive outcome in an Alzheimer's disease clinical trial may be estimated while adjusting for each subject's age. Including the covariate in the regression helps to explain some of the variability in outcome among the subjects, thereby leading to a more precise estimate of the treatment effect.

[1] Data Analysis Using Regression and Multilevel/Hierarchical Models

[2] ANCOVA versus change from baseline: more power in randomized studies, more bias in nonrandomized studies

This approach is widely-used and is supported by guidance from both FDA[3] and EMA[4].

Prognostic Covariate Adjustment (PROCOVA™) relies on similar statistical techniques and theory as ANCOVA but, instead of adjusting for simple baseline covariates, we adjust for predicted outcomes calculated from a subject's Digital Twin. Digital Twins are the ideal covariate, because Digital Twins are precise predictions generated from an AI-model optimized to explain the outcome. To implement PROCOVA™, first, an AI-based prognostic model is trained on historical control data. Then, this model is used to predict clinical outcomes from each subject's baseline data in the RCT (Figure 1). Finally, the treatment effect is estimated from the RCT using a regression that adjusts for this predicted clinical outcome. Schuler et al[5] prove this procedure is optimal when a high quality prognostic model is used; it attains the minimum variance among a wide class of estimators (Box 1). **In addition, PROCOVA™ provides unbiased estimates of treatment effects while ensuring strict control of the type-I error rate.**

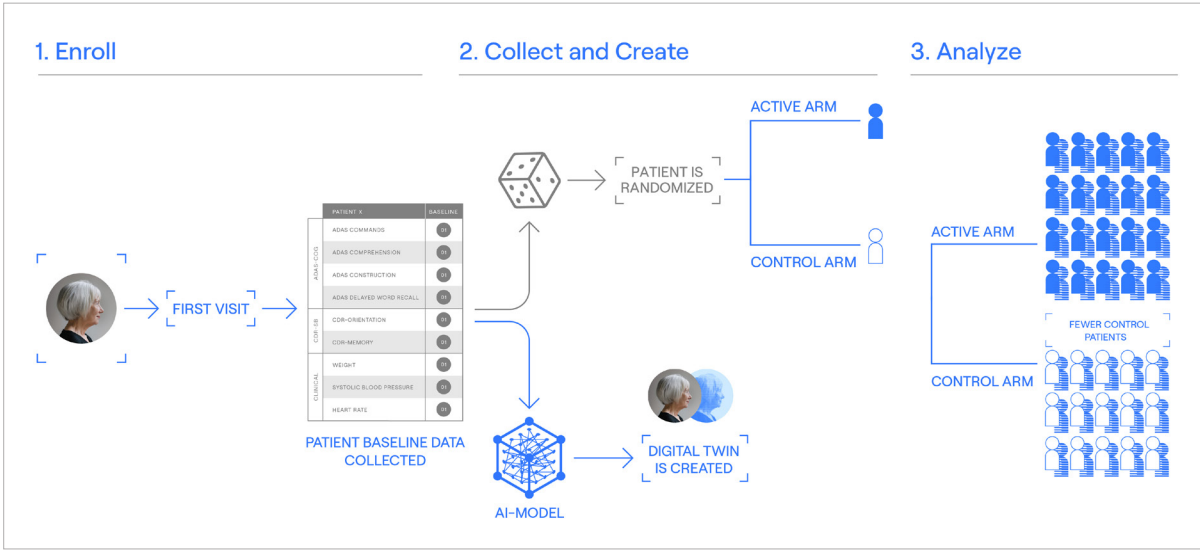

Figure 1. Using Digital Twins in Randomized Controlled Trials. Digital Twins are created for each patient in a clinical trial at first visit. A Digital Twin is a comprehensive, longitudinal clinical prediction representing a potential outcome for a particular patient in the control arm. Each Digital Twin is created from each patient's baseline data — before treatment — using an AI-model trained on a compendium of historical patient data from previously completed clinical trials and observational studies. A particular patient may go on to be randomized to either active treatment or control, such that Digital Twins can be included into randomized controlled trials in a way that maintains blinding.

[3] Adjusting for Covariates in Randomized Clinical Trials for Drugs and Biological Products

[4] Guideline on adjustment for baseline covariates in clinical trials

[5] Schuler et al. Increasing the efficiency of randomized trial estimates via linear adjustment for a prognostic score, 2020.

BOX 1. MATHEMATICAL DESCRIPTION OF PROCOVA™

Consider an RCT with $i=1...N$ patients randomized to active treatment $(W_i=1)$ or control $(W_i=0)$, in which we are interested in the effect of the treatment on $k=1...K$ outcomes. A vector of baseline covariates x_i is measured for each patient i after they enroll in the clinical trial, but before they receive treatment. The effect of the treatment on the k^{th} outcome is estimated by fitting the linear regression[6],

$$Y_{i,k} = \beta_{0,k} + \beta_{W,k} W_i + \beta_{f,k} f_k(X_i) + e_{i,k}$$

in which $\beta_{W,k}$ represents the average treatment effect. In addition to the treatment indicator variable, we've also adjusted for a covariate $f_k(X_i)$ obtained by applying a transformation to the baseline covariates[7].

This regression approach to estimating average treatment effects has two particularly interesting statistical properties:

Property 1. Unbiasedness. The estimate for the average treatment effect will be unbiased for any choice of transformation, $f_k(X_i)$. In fact, one can state further that this analysis will provide valid p-values for testing the null hypothesis $\beta_{W,k}=0$ for any choice of transformation, $f_k(X_i)$.

Property 2. Increased Power. The variance of the treatment effect estimator can be decreased by choosing the transformation so that $f_k(X_i)$ is correlated with $Y_{i,k}$. This choice of transformation also increases the statistical power for detecting the treatment effect.

Property 1 ensures that inferences about treatment effects drawn from these analyses are unbiased and can form the basis of regulatory decisions for the approval of new drug products. Property 2 provides an avenue for designing more powerful, and more efficient, clinical trials by choosing the optimal transformation.

[6] Similar approaches based on generalized linear models or generalized estimating equations can be used for other types of outcomes (e.g., binary, ordinal, survival times, etc).

[7] For simplicity, we'll only consider adjusting for a single covariate but it's easy to generalize this approach to multiple covariates.

Naturally, this leads to the question: what is the optimal transformation?

Schuler et al[8] proves that the optimal transformation is $f_k(X_i) = E[\, Y^0_{ik} \mid X_i]$ in which Y^0_{ik} is the k^{th} potential outcome that would be observed for subject i if she/he were randomized to the control group. Although this expected value isn't known, it can be predicted from a subject's baseline characteristics using AI-based prognostic models trained on historical control data.

The asymptotic variance of the PROCOVA™ estimate for the treatment effect in a trial with n subjects is less than

$$\frac{1}{n}\left[\frac{\sigma_0^2}{\pi_g} + \frac{\sigma_1^2}{\pi_1} - \pi_0\pi_1\left(\frac{\rho_1\sigma_1}{\pi_1} + \frac{\rho_0\sigma_0}{\pi_0}\right)^2\right]$$

where π_g is the fraction of patients randomized to group g, σ_g is the standard deviation of the outcome, and ρ_g is the correlation between the predicted and observed outcomes. Each of these parameters is specified or can be estimated from historical data, enabling the prospective design of clinical trials using PROCOVA™ control analyses.

PROCOVA™ enables more efficient clinical trials

Digital Twin outcomes can be used to explain away the variance in the real outcomes, making it easier to detect treatment effects (see Figure 2). Using PROCOVA™ to adjust for an AI-based prognostic model that obtains a correlation of 0.5 with an outcome can decrease the variance of the estimated effect of a treatment on that outcome by up to 25%. **Therefore, a clinical trial using PROCOVA™ with an accurate prognostic model requires fewer subjects to achieve a desired statistical power.** Reducing the sample size required to conduct a clinical trial has a direct impact on the time and cost of that study; e.g., fewer subjects may translate into fewer sites, faster study completion, and lower costs.

[8] Schuler et al. Increasing the efficiency of randomized trial estimates via linear adjustment for a prognostic score, 2020.

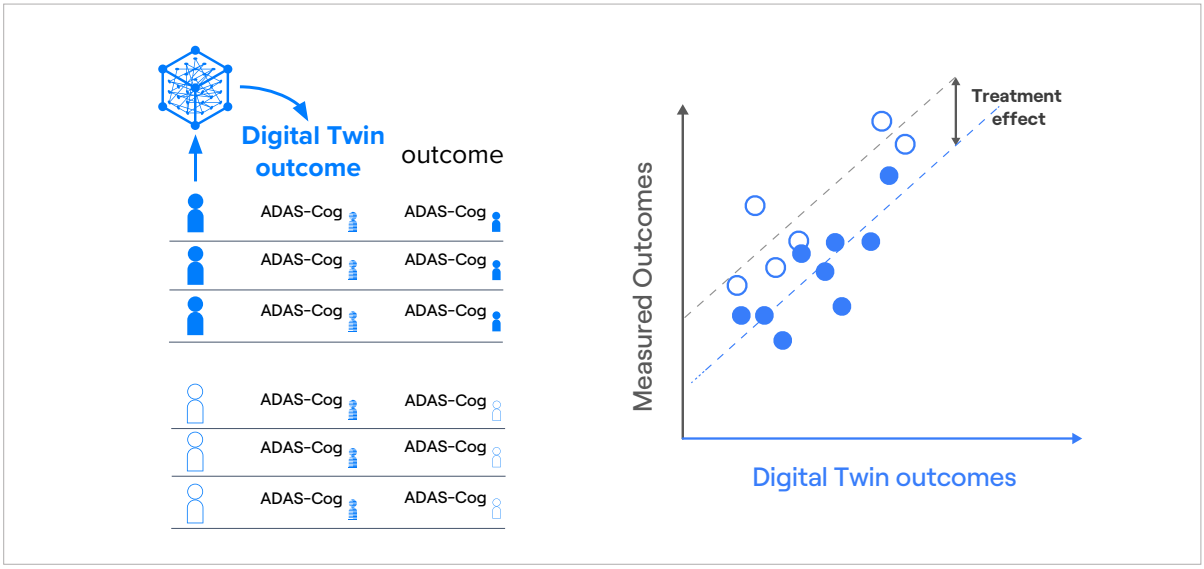

Figure 2: Prognostic Covariate Adjustment (PROCOVA™) powered by Digital Twins increases certainty, and thus power, by robustly explaining variability in the outcome.

We can actually go further if we use Bayesian methods and relax the requirement of strict type-I error rate control. In addition to being used to train a prognostic model, historical data can be used to estimate the accuracy and calibration of that prognostic model for predicting patient outcomes. If this analysis provides evidence for a high-quality prognostic model, more information can be borrowed from that model to decrease required sample sizes even further.

In PROCOVA+™ (Prognostic Covariate Adjustment "plus," which uses a Bayesian framework), prior distributions describing the relationship between predicted and observed outcomes are obtained from analyses of the control arms from previously completed clinical trials and used to decrease variance beyond what's possible with strict type-I error rate control. Although the Bayesian approach doesn't strictly control the type-I error rate, FDA's guidance on Bayesian statistics[9] says "when using prior information, it may be appropriate to control type I error at a less stringent level than when no prior information is used" especially if "the prior information is based on empirical evidence such as data from clinical trials."

Simulations of historical Alzheimer's Disease trials including Digital Twins with PROCOVA+™ demonstrate reductions in required sample size more than twice as large as what is achievable using strictly frequentist methodologies. These are significant improvements over the frequentist design but at the expense of relaxing strict control of the type-I error rate[10].

[9] FDA Guidance on Guidance for the Use of Bayesian Statistics in Medical Device Clinical Trials

[10] Gains in power or reductions in required sample sizes depend on characteristics of the trial design including inclusion/exclusion criteria and the choice of endpoints.

Summary

Unlike pooling and propensity score-based methods used to incorporate external data into clinical trials, PROCOVA™ reduces the chances of missing a true treatment effect without increasing the probability of false positives (see Figure 3). Although regulatory decision making from a clinical trial typically focuses on a single outcome designated as the primary endpoint, there are, in fact, many analyses that need to be conducted in any given clinical trial. To accommodate these diverse analyses, it's necessary to use prognostic models that predict all of the relevant patient characteristics over time. As a result, AI-based methods trained to create Digital Twins — accurate, comprehensive, longitudinal clinical predictions of control outcomes — enable the design of more efficient clinical trial, with higher power and smaller required sample sizes (see Figure 4), that are appropriate for both early and late-stage clinical development.

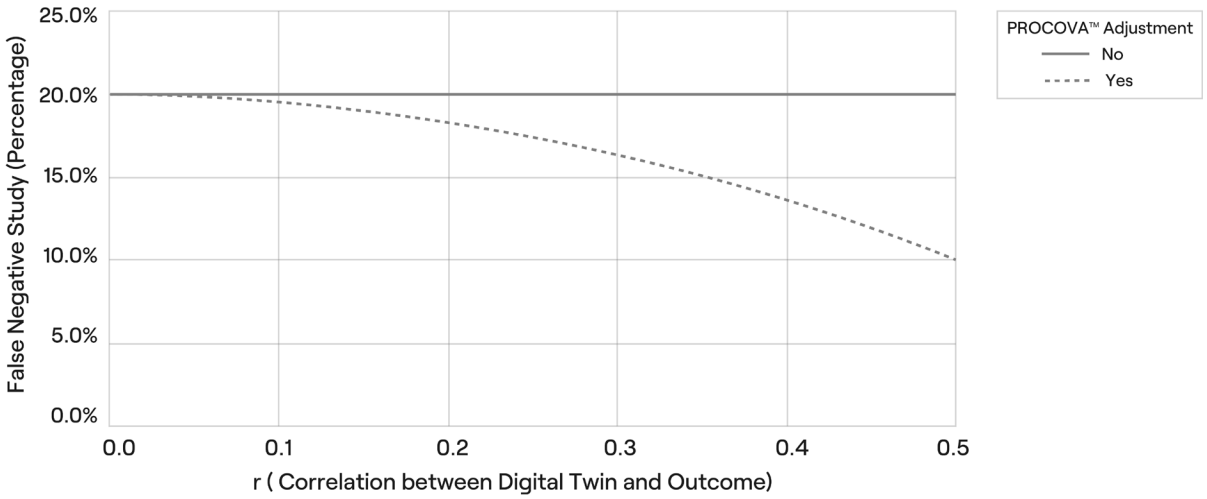

Figure 3: False Negative Studies with PROCOVA™ compared to original design with 80% power. When using PROCOVA™ (Prognostic Covariate Adjustment), as r (the correlation between a patient's actual outcome and their Digital Twin outcome) increases, the likelihood of a false negative study occuring decreases. Without PROCOVA™, i.e. the original clinical trial designed with 80% power, the likelihood of a false negative study occurring remains constant. When an r of 0.5 is achieved, the likelihood of a false negative study is cut in half with PROCOVA™ (10% vs 20%), and lower correlations, such as an r=.35 can still cut the chances of a false negative study by a quarter.

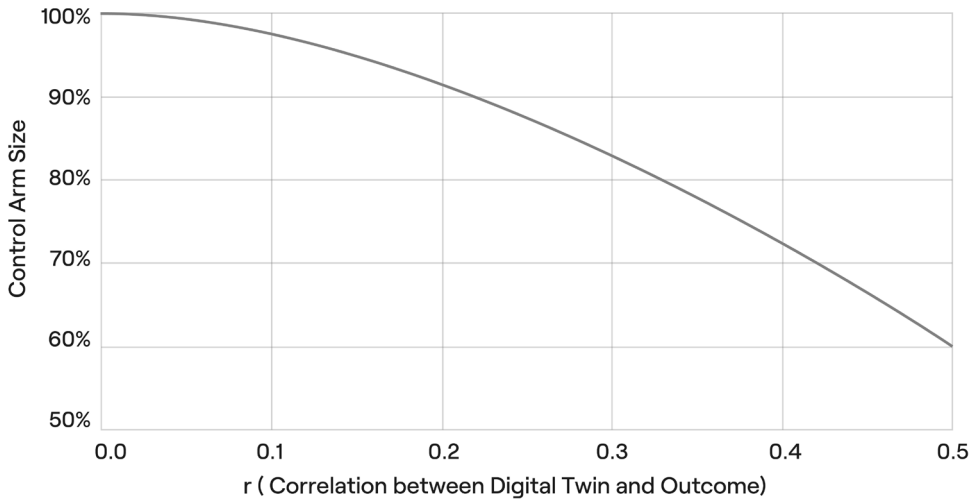

Figure 4: Reduction in placebo sample size with PROCOVA™ retaining 80% power. When using PROCOVA™ (Prognostic Covariate Adjustment), as r (the correlation between a patient's actual outcome and their Digital Twin's outcome) increases, the required control arm size decreases. The number of control arm patients is 60% of that from the original clinical trial with PROCOVA™ when an r of 0.5 is achieved and 80% of the original control arm size with a smaller correlation (e.g., r=.35).

With Prognostic Covariate Adjustment (PROCOVA™), we adjust for predicted clinical outcomes calculated from a subject's Digital Twin. The Digital Twin is the ideal covariate, because Digital Twins are precise clinical predictions generated from an AI-model optimized to explain the outcome. Digital Twins have a higher correlation with the outcome than any other single covariate. PROCOVA™ enables more efficient, higher powered clinical trial designs with fewer patients that don't introduce bias.

UNLEARN

We are seeking partners

to adopt this novel approach to accelerate clinical development in CNS, Immunology, and other indications.

Reach out to us to learn more about how to leverage Digital Twins and corresponding statistical approaches like PROCOVA™ to enable faster, more efficient trials. **It's About Time.**

unlearn.ai | partnerships@unlearn.ai

SCIENTIFIC REPORTS
nature research