Adaptive Designs

Mark van der Laan
Division of Biostatistics, UC Berkeley

September 28 , 2018
Workshop on Study Designs for Implementation Science
UCSF

Joint work with Antoine Chambaz, Wenjing Zheng, Ivana Malenica, Romain Pirrachio

Outline

Outline

Foundations of Statistical Learning

- **Observed data**: Realization of a random variable $O^n = (O_1, \ldots, O_n)$ with a probability distribution (say) P_0^n, indexed by "sample size" n.
- **Model stochastic system of observed data realistically**: Statistical model $\mathcal{M}^n$ is set of possible probability distributions of the data.
- **Define query about stochastic system**: Function Ψ from model $\mathcal{M}^n$ to real line, where $\Psi(P_0^n)$ is the true answer to query about our stochastic system.
- **Estimator**: An a priori-specified algorithm that takes the observed data O^n and returns an estimate ψ_n to the *true answer to query*. Benchmarked by a dissimilarity-measure (e.g., MSE) w.r.t true answer to query.
- **Confidence interval for true answer to query:** Establish approximate sampling probability distribution of the estimator (e.g., based on CLT), and corresponding statistical inference.

Targeted Learning (TL)

is the subfield of statistics concerned with development of estimators P_n^* based on data $O^n \sim P_0^n$ from the stochastic system P_0^n with corresponding estimates $\Psi(P_n^*)$ and **confidence intervals** for true answer $\Psi(P_0^n)$, **based on realistic statistical models** $\mathcal{M}^n$.

By necessity, TL involves highly data adaptive estimation (e.g., machine learning).

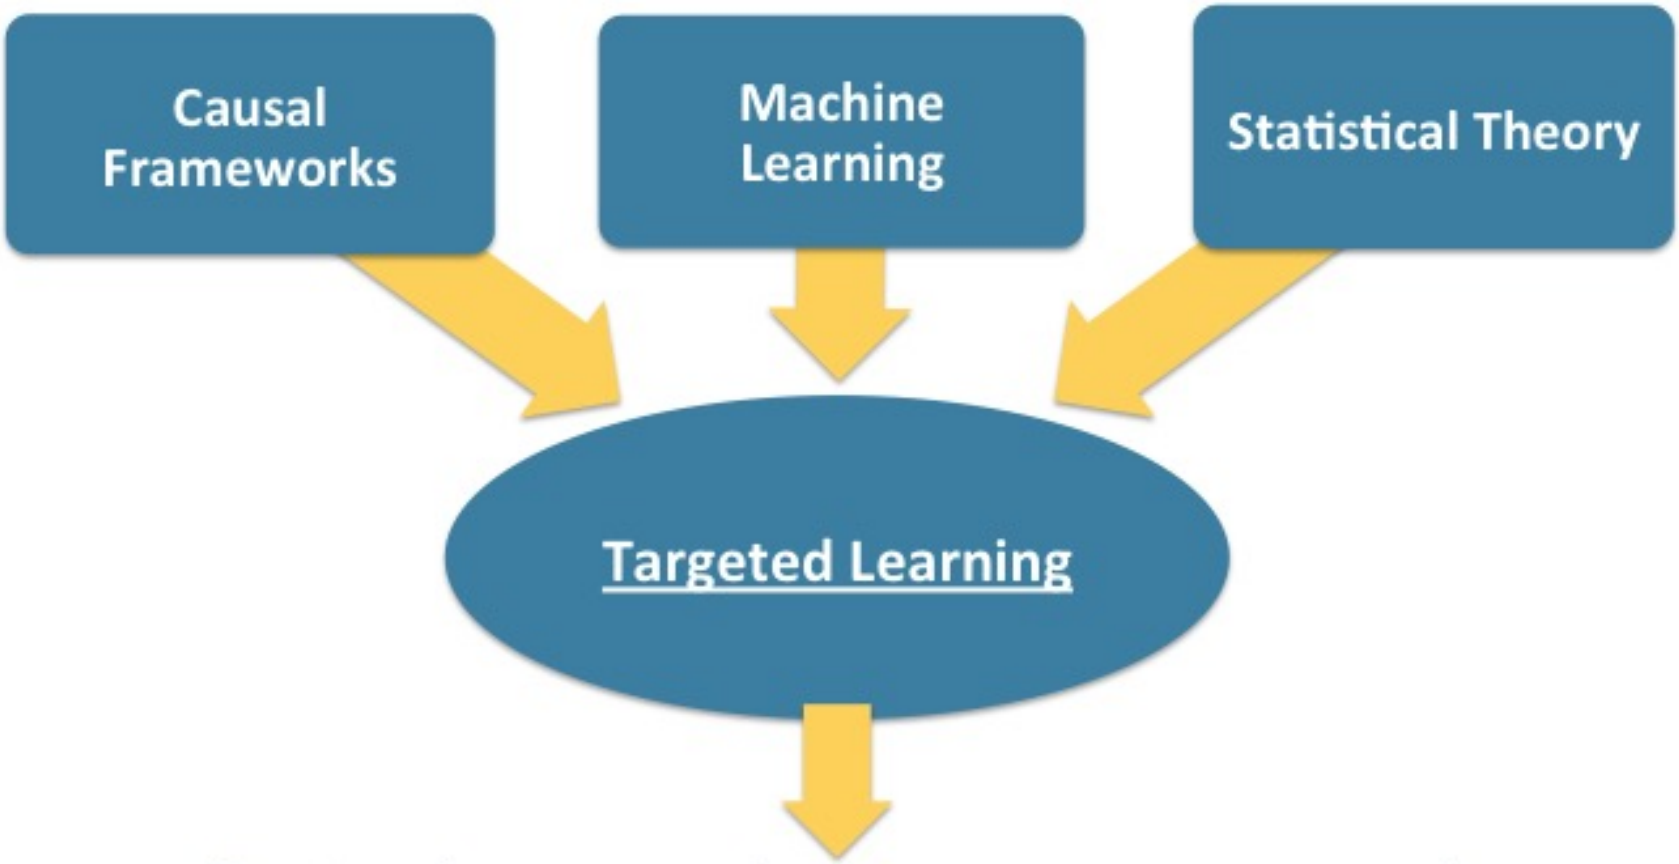

Better (more precise) **answers** to **causal** (actionable) **questions** with **accurate quantification of uncertainty** (signal from noise)

Targeted Learning (targetedlearningbook.com)

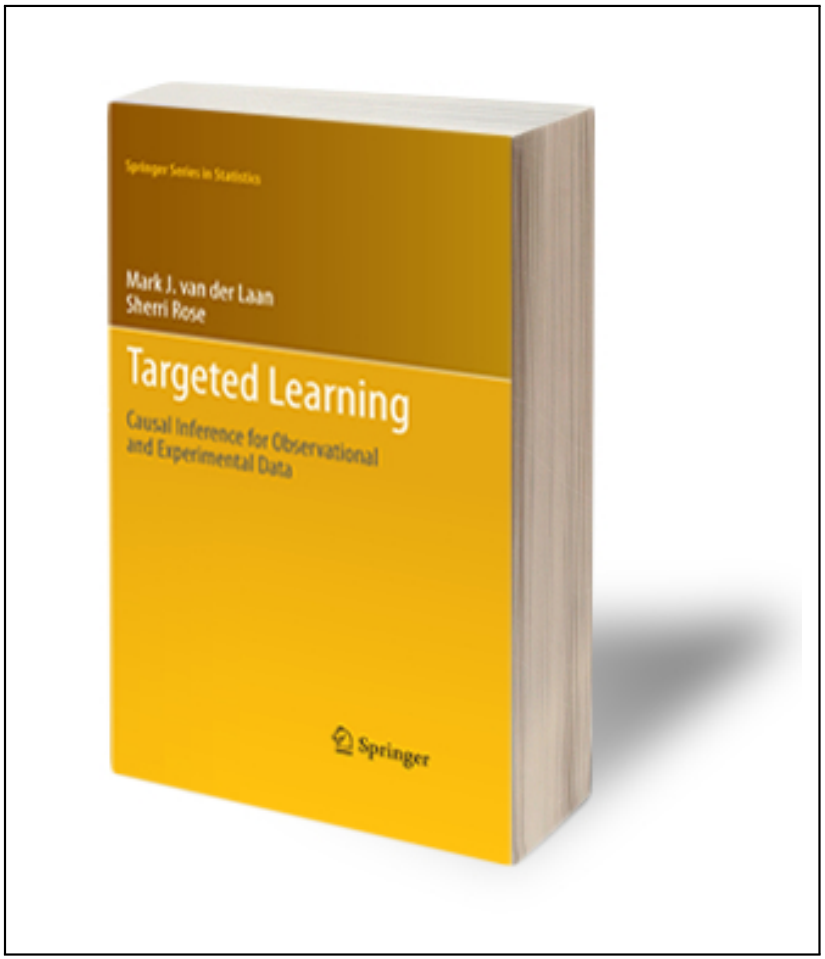

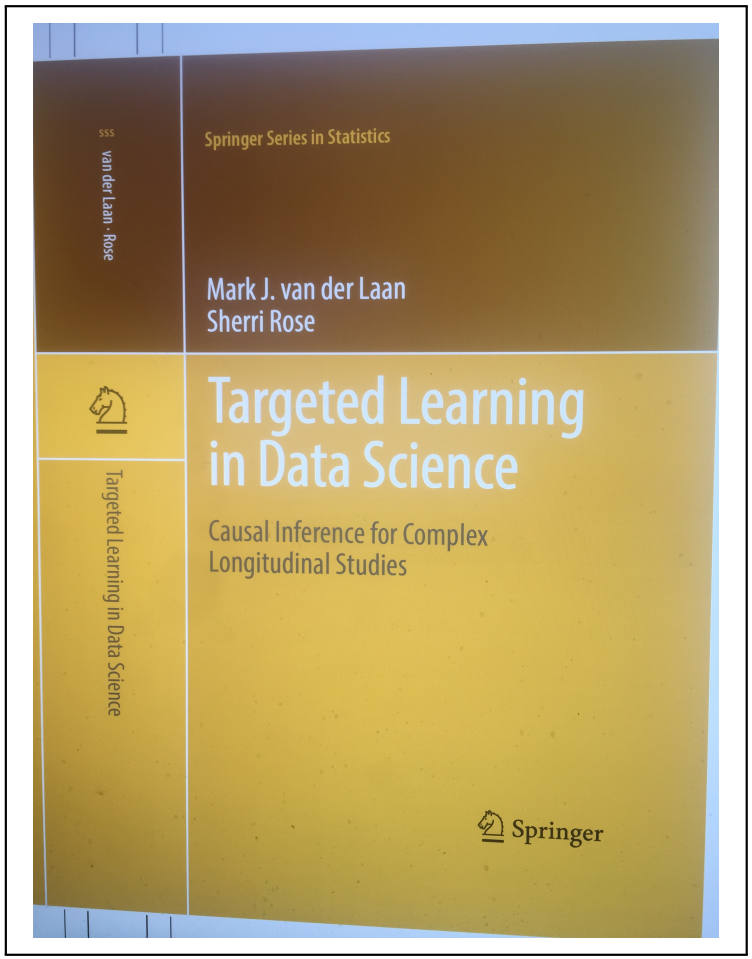

van der Laan & Rose, *Targeted Learning: Causal Inference for Observational and Experimental Data*. New York: Springer, 2011.

Outline

1 Better, cheaper trials

Do corticosteroids reduce mortality for adults with septic shock?

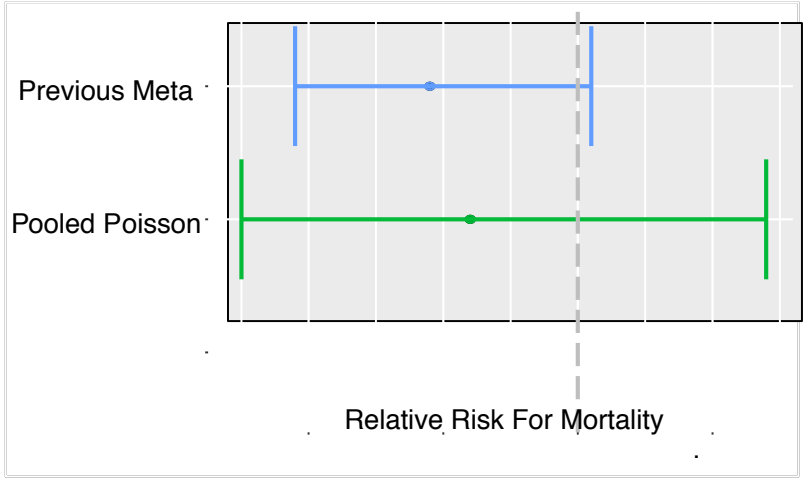

Previous Meta-Analysis of 31 trials: ***No significant benefit***

Pooled analysis of 3 major RCTs (1300 patients) with standard methods: ***No significant benefit***

Pirracchio 2016

Better, cheaper trials

Do corticosteroids reduce mortality for adults with septic shock?

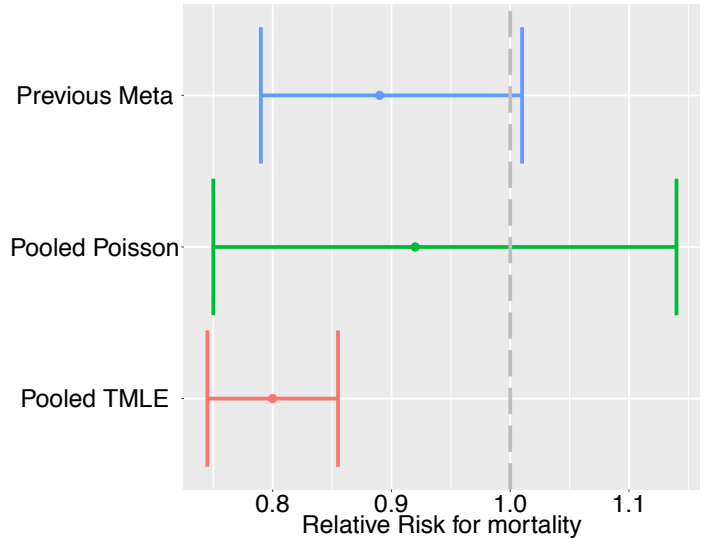

Previous Meta-Analysis of 31 trials: No significant benefit

Pooled analysis of 3 major RCTs (1300 patients) with standard methods: **No significant benefit**

Pooled analysis of 3 major RCTs using Targeted Learning: significant reduction of mortality.

Not just is there an effect, but for whom?

- In Sepsis re-analysis: Targeted Learning showed **all benefit** occurred in a key subgroup
  - Heterogeneity in patient populations one cause of inconsistent results

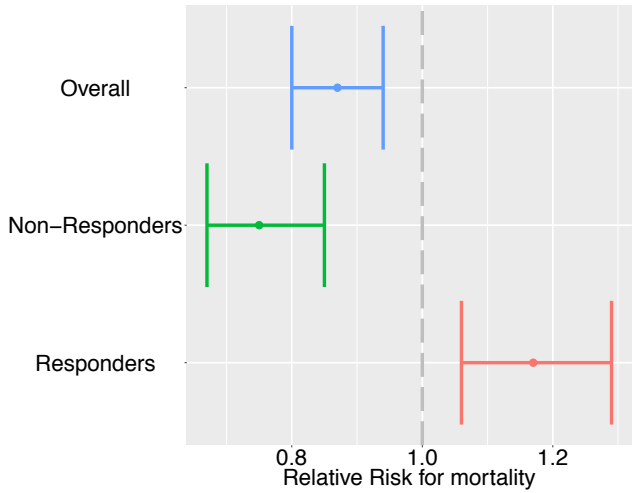

Effect Heterogeneity by Response to ACTH Stimulation

Outline

Optimal intervention allocation: “Learn as you go”

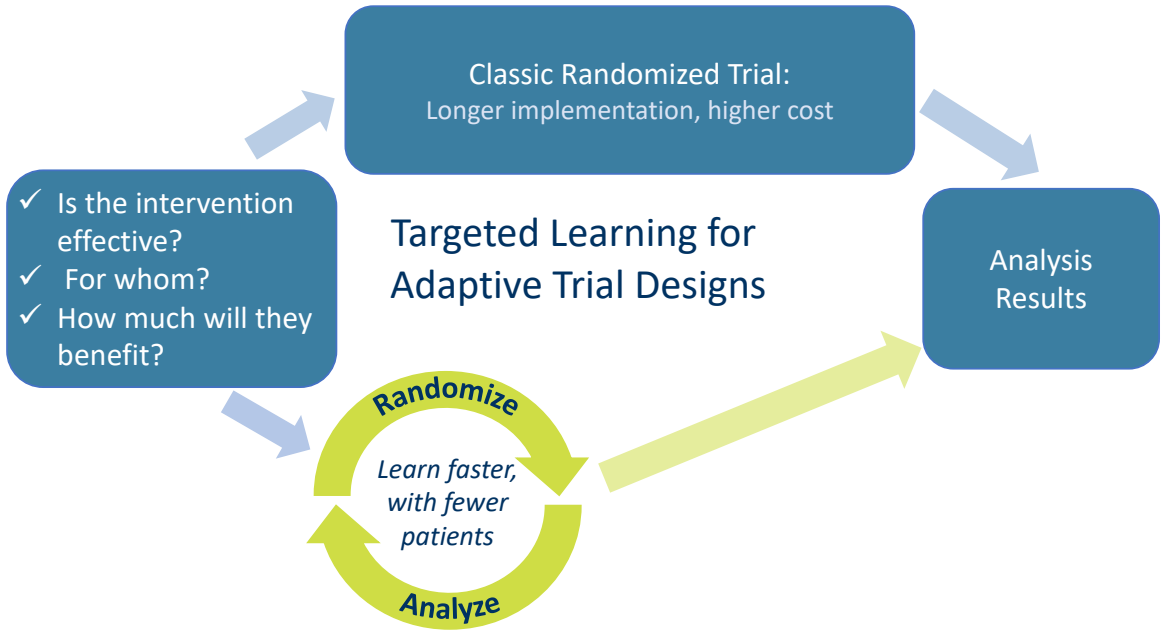

Contextual multiple-bandit problem in computer science

Consider a sequence $(W_n, Y_n(0), Y_n(1))_{n \geq 1}$ of i.i.d. random variables with common probability distribution:

- W_n, nth context (possibly high-dimensional)
- $Y_n(0)$, nth reward under action $a = 0$ (in $]0, 1[$)
- $Y_n(1)$, nth reward under action $a = 1$ (in $]0, 1[$)

We consider a design in which one sequentially,

- observe context W_n
- carry out randomized action $A_n \in \{0, 1\}$ based on past observations and W_n
- get the corresponding reward $Y_n = Y_n(A_n)$ (other one not revealed),

resulting in an ordered sequence of dependent observations $O_n = (W_n, A_n, Y_n)$.

Goal of experiment

We want to estimate

- the optimal treatment allocation/action rule d_0: $d_0(W) = \arg\max_{a=0,1} E_0\{Y(a)|W\}$, which optimizes the mean outcome EY_d over all possible rules d.
- the mean reward under this optimal rule d_0: $E_0\{Y(d_0)\}$,

and we want

- maximally narrow valid confidence intervals (primary) "Statistical...
- minimize regret (secondary) $\frac{1}{n}\sum_{i=1}^{n}(Y_i - Y_i(d_n))$ ... bandits"

This general contextual multiple bandit problem has enormous range of applications: e.g., on-line marketing, recommender systems, randomized clinical trials.

Targeted Group Sequential Adaptive Designs

- We refer to such an adaptive design as a particular targeted adaptive group-sequential design (van der Laan, 2008).
- In general, such designs aim at each stage to optimize a particular data driven criterion over possible treatment allocation probabilities/rules, and then use it in next stage.
- In this case, the criterion of interest is an estimator of reward EY_d under treatment allocation rule d based on past data, but, other examples are, for example, that the design aims to maximize the estimated information (i..e., minimize an estimator of the variance of efficient estimator) for a particular statistical target parameter.

Bibliography (non exhaustive!)

- Sequential designs
  - Thompson (1933), Robbins (1952)
  - specifically in the context of medical trials
    - Anscombe (1963), Colton (1963)
    - **response-adaptive designs**: Cornfield et al. (1969), Zelen (1969), many more since then
- Covariate-adjusted Response-Adaptive (CARA) designs
  - Rosenberger et al. (2001), Bandyopadhyay and Biswas (2001), Zhang et al. (2007), Zhang and Hu (2009), Shao et al (2010)... *typically* study
    - convergence of design ... in correctly specified parametric model
  - Chambaz and van der Laan (2013), Zheng, Chambaz and van der Laan (2015) concern
    - convergence of design, super-learning of optimal rule, *and* TMLE of optimal reward, with inference, without (e.g., parametric) assumptions.

Outline

Sequential adaptive designs adapting in continuous time

- Problem with group sequential is that one has to run a number of randomized trials sequentially, taking too much time for long term clinical outcomes.
- Suppose subjects enroll over time, possibly in groups, or one at the time.
- Each subject will go through a (say) 12-month course from entry time till final outcome: for example, one measures baseline covariates at $k = 0$, assign treatment at $k = 0$, measure surrogate outcome at time $k = 1, \ldots, k = 11$ months, and final outcome at $k = 12$-months.
- Or, one might also assign treatment at later $k > 0$ months.

Adapting the treatment decision based on observed past

- When a subject comes in at a chronological time t, $k \geq 0$ months after entry, and is subject to a treatment action, then we can take into account all the available (incomplete) data on previously or concurrently enrolled subjects.
- For example, we could use the past data to learn an optimal treatment decision at time k for maximizing the surrogate outcome at near future time-point (say) $k + 1$.
- In this manner, we can use adaptive designs for long-term clinical outcomes, adapting to optimal treatment rules w.r.t. surrogate intermediate outcomes.

Outline

Adaptive design for learning the optimal rule

- Suppose we observed a single time-series of data in which at "time" t we observe a record $O(t) = (A(t), Y(t), W(t))$, treatment $A(t)$, outcome $Y(t)$, other measurements $W(t)$, while history $\bar{O}(t-1)$ before $A(t)$ represents context, $t = 1, \ldots$.
- Suppose that the conditional distribution of $O(t)$, given past $\bar{O}(t-1)$, is parameterized by common functional parameters (stationarity).
- In a controlled setting, we can generate treatment $A(t)$ at time t from a randomization probability depending on the complete history of subject.
- These randomization probabilities can be based on learning an optimal rule for setting treatment at time t for the purpose of maximizing the next outcome $Y(t)$.
- In this manner we both learn and apply the optimal rule, while providing an estimate of its performance with inference.

Outline

Concluding Remarks

- Group sequential randomized trials can be used for short-term clinical outcomes with robust inference (as in standard RCT).
- Sequential randomized trials adapting in continuous time take into account surrogate outcomes can be used for long-term clinical outcomes, with robust inference (to be written up).
- Sequentially randomized trials within a single unit/person can be used to learn optimal rule for short term outcomes, with robust inference.
- Software in R has been developed for estimation and inference for all the first and third type of randomized trials, second is in the making.

Acknowledgements

This work has benefitted greatly from collaborations with:

- David Benkeser (Emory University), Alex Luedtke (Fred Hutchinson Cancer Research)
- Sam Lendle (Pandora), Antoine Chambaz (Paris)
- Cheng Yu, Erin LeDell (Berkeley, H20)
- Maya Petersen, Alan Hubbard (Berkeley)
- Eric Polley (NCI)
- Sherri Rose (Harvard)

This research was supported by NIH grant R01 AI074345-06.