Adaptive Designs for Epidemic Control
Test and Respond Strategies for COVID-19

Maya Petersen and Mark van der Laan

Division of Biostatistics, School of Public Health
University of California, Berkeley

August 12, 2020

Outline

- Motivation and Background: Adaptive designs for COVID-19 epidemic control
  - Adaptive Surveillance
  - Adaptive Randomized Trials
- Applications:
  - UC Berkeley
  - San Francisco County
- General Statistical Framework
- Adaptive Test and Respond Design for an Evolving Epidemic

Motivation: What is the best way to deploy finite testing resources for COVID-19 epidemic control?

- **Example Objective:** Minimize prevalence of SARS-CoV-2 infection at some future time point
  - i.e. minimize $E[Y(K)]$, where $Y(K)$ is an indicator of active infection at some future time point K
  - Many alternatives; e.g., expected hospitalizations (cumulative or max daily census) $<$ some threshold, etc.
- **Mechanisms** by which testing for SARS-CoV-2 affects future disease prevalence
  1. **Direct impact:** Isolation of infectious persons detected by **current** testing
  2. **Information:** Improved **targeting of future interventions** – amplify impact of future testing, other interventions (e.g., education, social support)

Background: Adaptive Sampling Designs

Objective: Assign measurements to optimize information

- **Design:** Sequentially sample individuals for testing
  - Sampling probability depends on individual covariates W (e.g., geospatial coordinates)
  - Sampling mechanism updated over time based on results of past testing rounds (and corresponding risk estimates)
- **Target parameter** (Examples):
  - Disease prevalence at a point $E[Y|W = w]$
  - "Hotspot" detection (i.e., $I[E(Y|W = w) > \rho]$, for some threshold prevalence $\rho \in (0, 1)$)
  - Performance of adaptive design (e.g., hotspot classification error)
- Machine Learning + Statistical Inference (e.g., Pancheco et al. 2020, Bibaut et. al., 2020)

Example: Adaptive geospatial sampling for malaria surveillance (eg, Pancheco et al. 2020,Kabaghe, et. al, 2017)

- Exploitation vs Exploration for hotspot classification with geospatial correlation

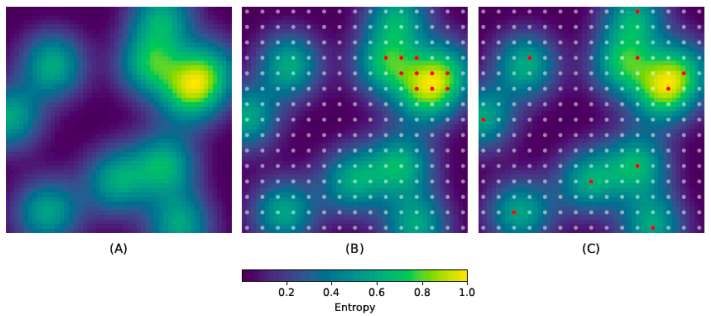

Figure 1 Exploration-exploitation trade-off. A: Spatially correlated uncertainty. B: Batch selected (red dots) by using the greedy approach of targeting the highest values of uncertainty. C: Batch of locations selected (red dots) using the acquisition function described in Eq. 7.

- Adaptive sampling improves classification accuracy

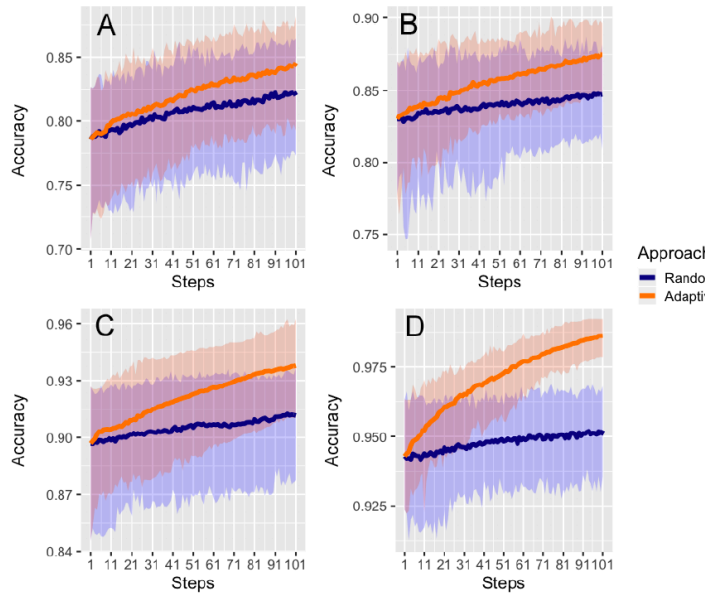

Figure 3 Out of sample accuracy (batch size = 1). The solid line represents the average v across 50 repetitions. The shaded area represents the 2.5% and 97.5% quantiles of the valu observed across all 50 repetitions at each step. A: Cote d'Ivoire ($\vartheta = 10\%$). B: Malawi ($\vartheta =$ C: Haiti ($\vartheta = 2\%$). D: Philippines ($\vartheta = 2\%$).

Background: Adaptive Trial Designs

Objective: Assign treatments to optimize impact

- **Design:** Covariate-Adjusted Response-Adaptive (CARA) designs for responding to treatment effect heterogeneity
  - Randomization probability depends on covariates W
  - Updated over time based on outcomes of persons previously randomized (and corresponding conditional effect estimates)
- **Target parameter** (examples):
  - Treatment allocation rule $d_0(W)$ that maximizes (good) outcomes over all possible rules d
  - Mean reward under this optimal rule $E(Y(d_0(W)))$ (i.e. mean counterfactual outcome had all treatments had been allocated optimally)
- Machine learning for effect heterogeneity + statistical inference (e.g., Luedtke vdL, 2016a,b,c; Chambaz et al 2011, 2017)

Example: Personalized behavioral interventions to optimize retention in HIV Care in Kenya

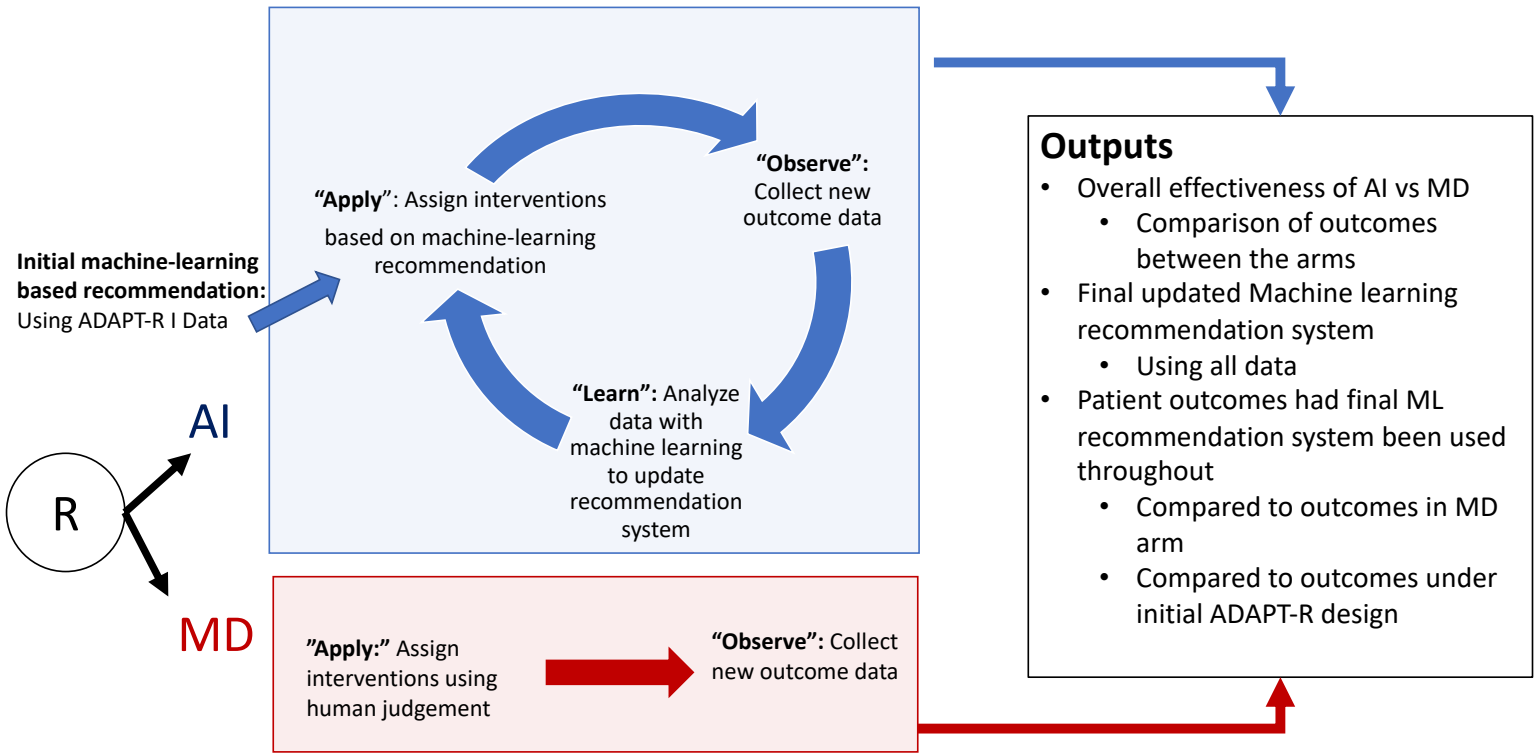

ADAPT-2: "Man vs. Machine"; NIH-sponsored R34 Pilot (Petersen, Geng)

Back to COVID-19: Adaptive Epidemic Control "in the wild"

- Testing plays a **dual role**: Information and Direct Impact
- Prevalence varies over physical and covariate space
- Impact of detecting an active infection varies depending on behaviors, networks, viral load, etc.
- Risk surface and epidemic context evolve over time
- Testing decisions are at best partially controlled (hybrid observational and randomized designs)
- Networks imply complex dependence (and are not well measured)

Application: Berkeley Safe Campus Study

- Cohort of 3000 students, workers, and researchers
- Integrated data collection, testing and response system

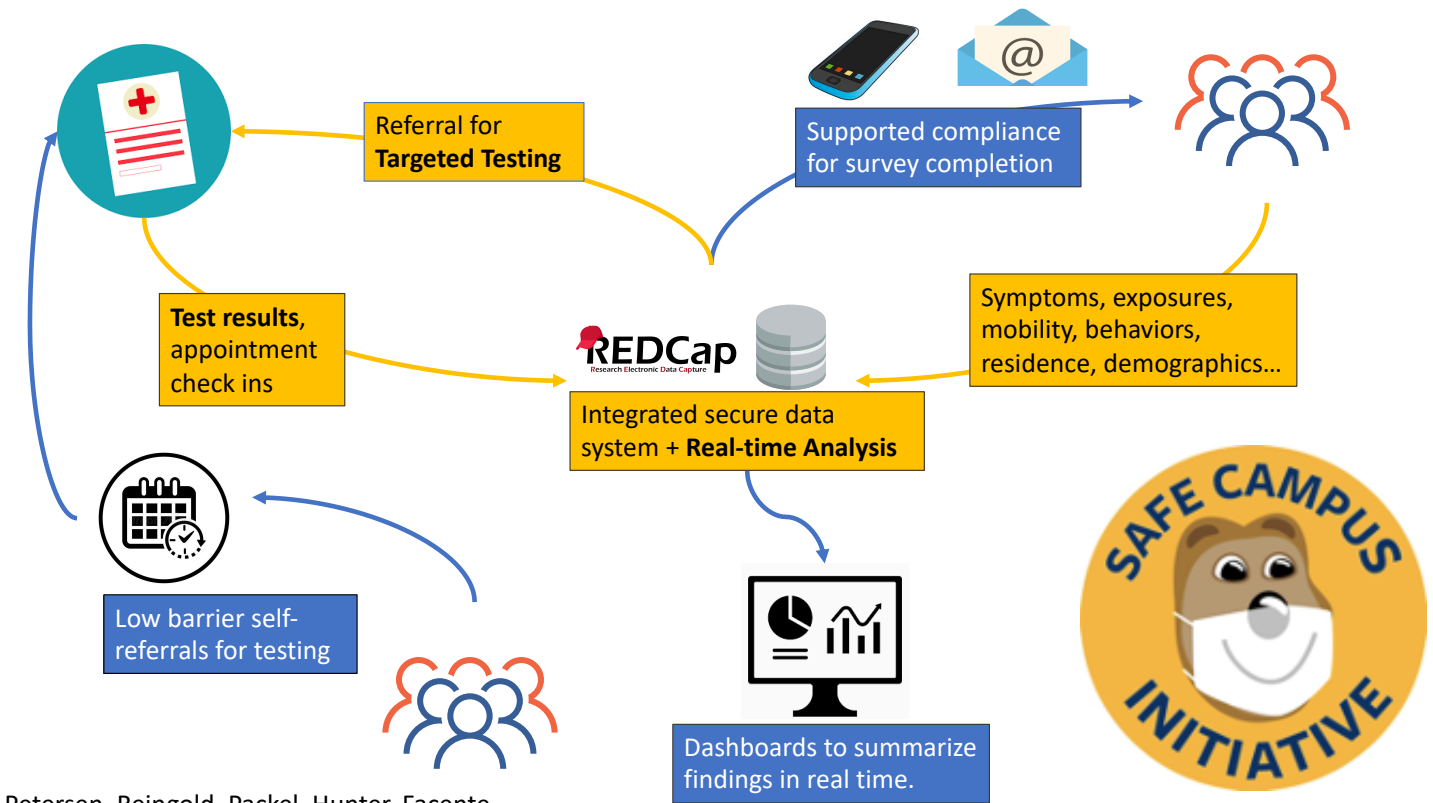

Petersen, Reingold, Packel, Hunter, Facente

Simulation-based illustration

- Agent-based mathematical model of UC Berkeley campus (N=20,000) (Malenica, Coyle)
  - Household (including group housing), classroom, and random contacts
  - Heterogeneity in risk (beyond network)
  - Infection seeding from outside of campus
- Testing strategies contrasted (200 tests/day)
  1. Random
  2. Single regular interval
  3. Static risk-stratified interval
     - 1x/wk for Dorm, remaining tests at regular intervals
  4. Based on estimated individual risk of infectiousness
     - Based on past data, accounting for testing mechanism
     - Different risk-driven testing functions $\Rightarrow$ different tradeoffs between "exploitation and exploration"

Risk-driven testing is much more efficient for epidemic control

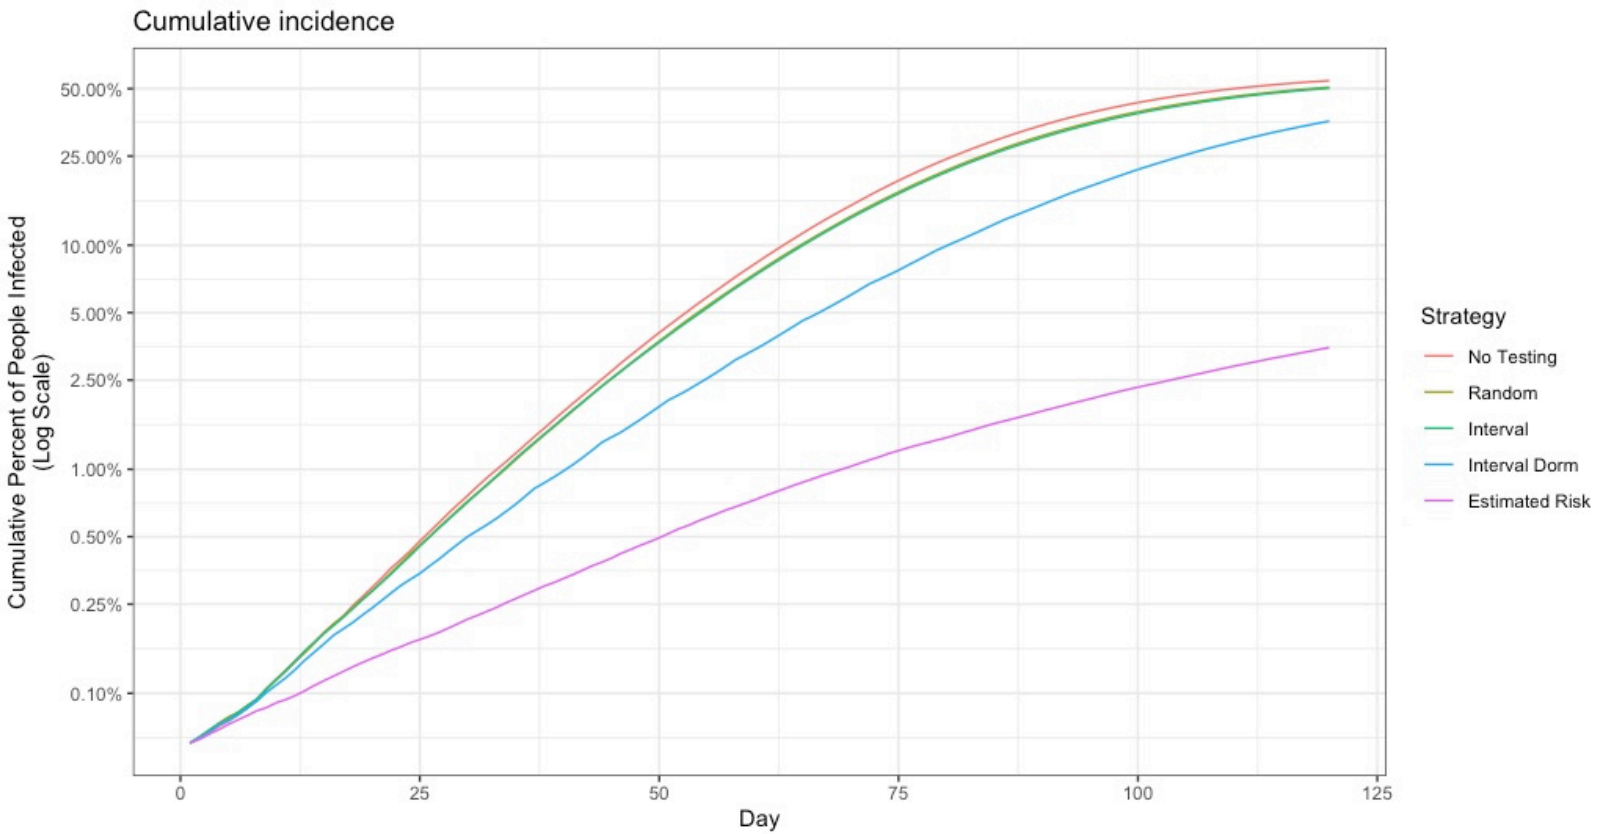

Berkeley Safe Campus Study: Preliminary Results

- **High acceptability and uptake** (98% uptake and 78% average daily completion rate)
  - Nearly 85% of students "likely" or "extremely likely" to continue to participate in the daily and weekly surveys should campus continue these beyond the study
- **Triggered testing:**
  - 58 persons with a positive test
  - 40-60% of students with a test triggered reported for testing
  - 5-10% of those tested positive
- Additional **lessons learned**: self-sreferrals, low barrier testing, social networks, de-stigmatization

Application: San Francisco County

- Same general design can (in theory) be applied to a city or region: Adapt and respond cycles to target testing and other interventions
- Different "action": Targeted outreach to neighborhoods and communities in place of individual-level testing referral
- More persons measured, but fewer covariates
  - Residential location and demographics
  - LOTS of missing values (and other data issues)

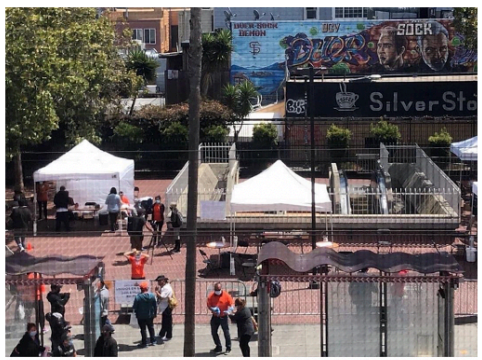

UNIDOS EN SALUD UNITED IN HEALTH
SAN FRANCISCO

Heterogeneity between and within census tracts

- Heterogeneity across census tracts
- Within census tracts with high percent positive, Latinx communities disproportionately affected

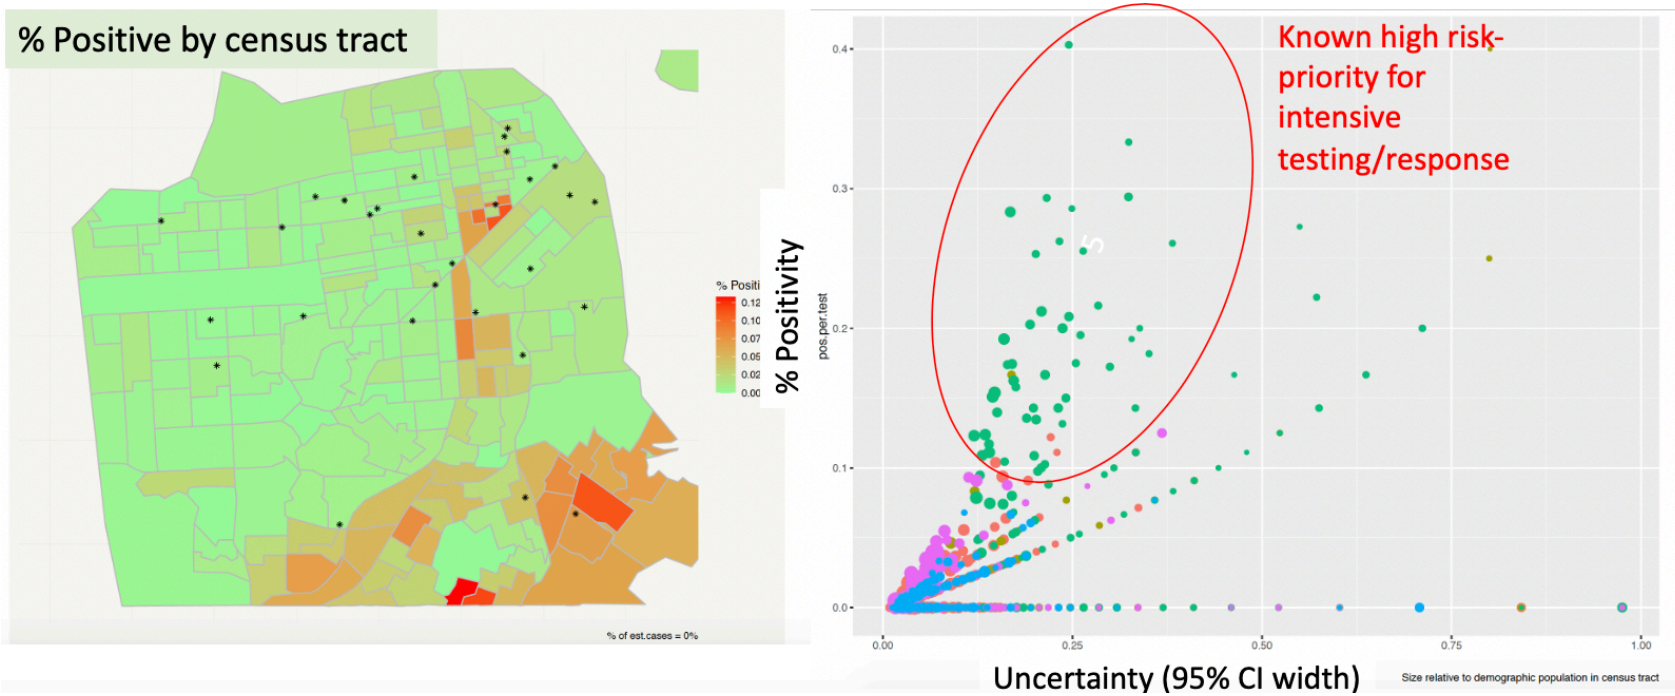

Potential for improved epidemic control through targeted testing

- Focused testing, supported isolation (e.g., Right to Recover), education and prevention
- Geographic proximity and community mobilization and partnerships

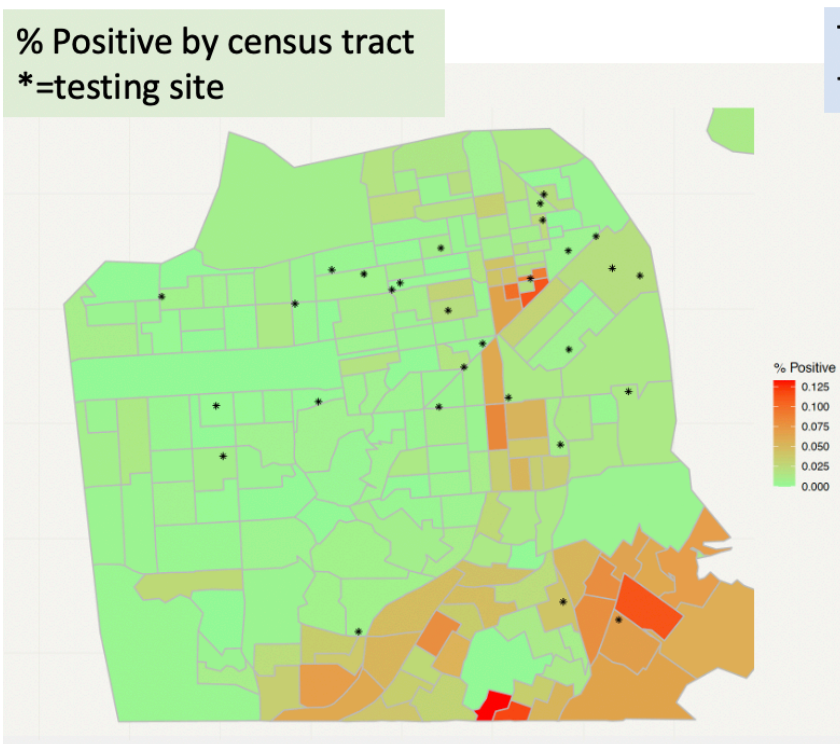

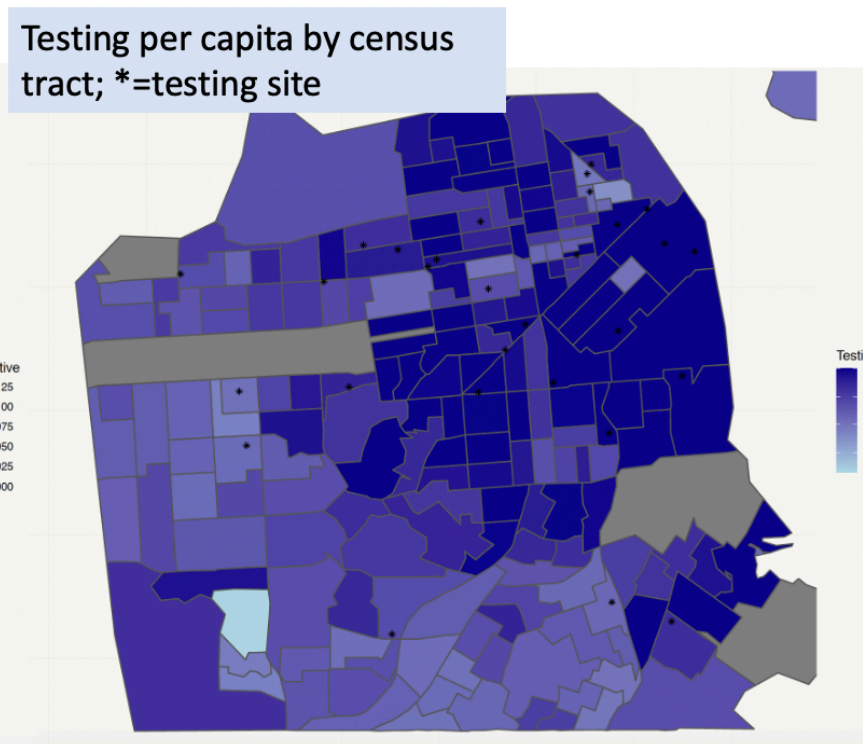

Overview: General Statistical Framework

- Contextual multiple-bandit problem in computer science
  - Need for statistical inference (both optimize a system and make inferences about it)
- Links to adaptive trial designs:
  - Optimizing and evaluating strategies for responding to treatment effect heterogeneity
- Extensions to real-world adaptive surveillance and response systems for epidemic control
  - Theoretical progress and ongoing work

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 P_0^F:

- 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 $O_n = (W_n, A_n, Y_n)$.

Contextual Multiple Bandit: General Objective

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 EY_d over all possible rules d.
- the mean reward under this optimal rule d_0: $\Psi(P_0^F) = E_0\{Y(d_0(W))\}$,

and we want

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

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

Balanced vs. adaptive sequential design

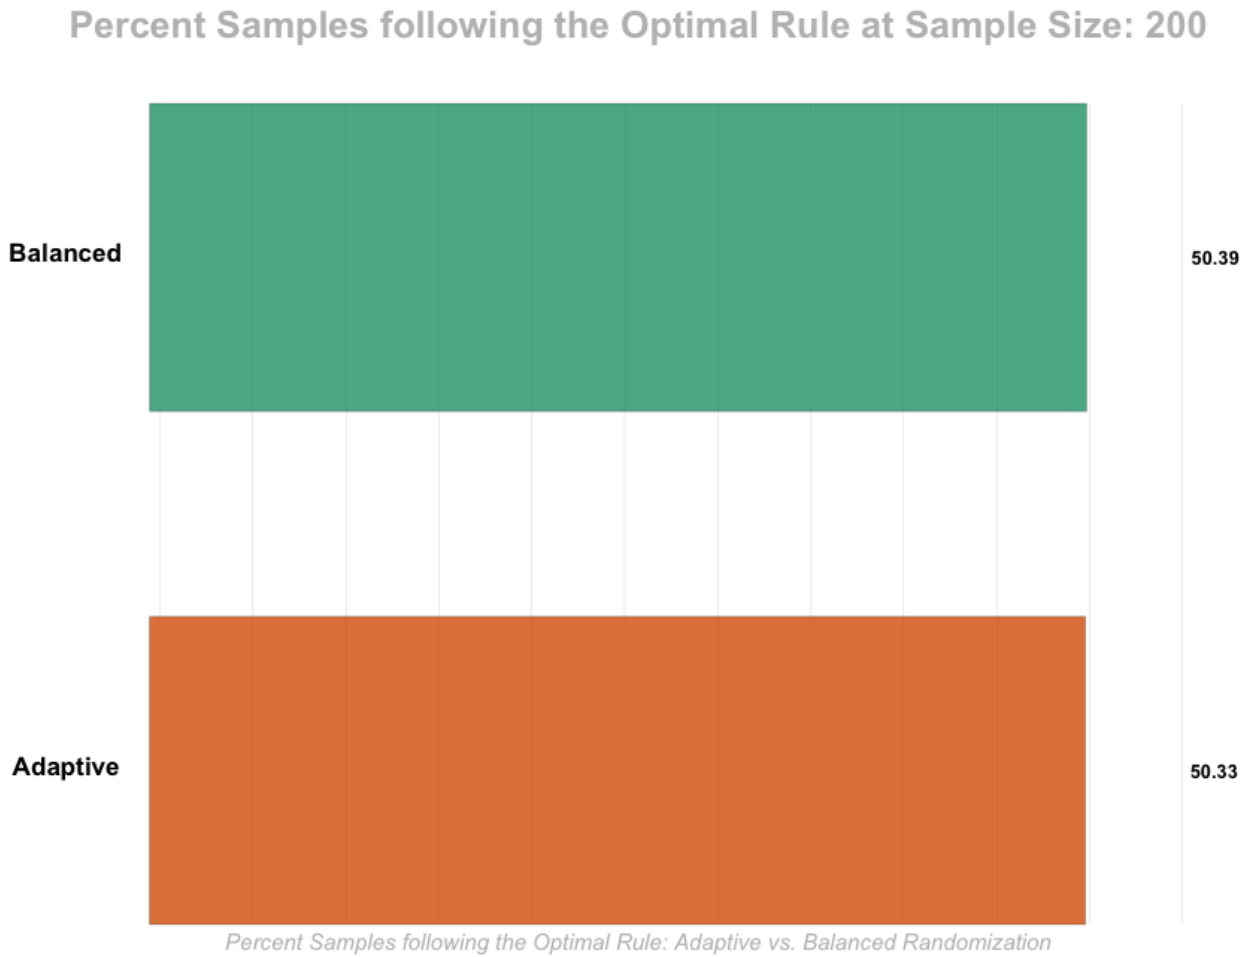

Balanced vs. adaptive sequential design

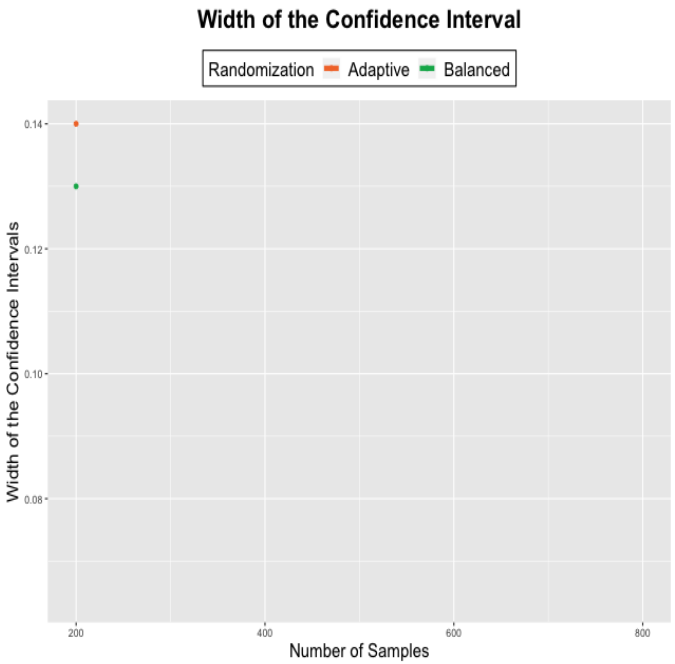

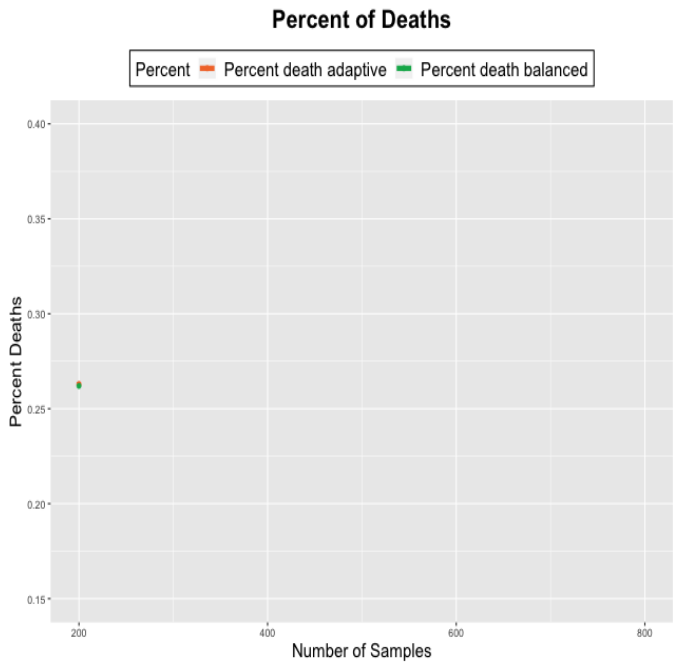

Adaptive Testing and Response for an Evolving Epidemic

Observed longitudinal data structure on population of individuals

- Consider population of N individuals.
- For each individual i, we observe a longitudinal data structure

$$O_i = (L_i(0), A_i(0), L_i(1), A_i(1), \ldots, L_i(\tau), A_i(\tau), Y_i(\tau+1)).$$

- $L_i(0)$ baseline history, including contacts of subject i, and baseline infectious status $Y_i(0)$.
- $A_i(t)$ is indicator of being sampled/tested.
- $L_i(t)$ includes $A_i(t-1)Y_i(t)$ indicator of being infected, if tested.
- $L_i(t)$ may include changes in contacts, changes in risk factors that are predictive of future infection.
- $Y_i(\tau+1) = A_i(\tau)Y_i(\tau+1)$ indicator of being infectious, if tested.
- Notation: $\mathbf{L}(t)$ and $\mathbf{A}(t)$ for N dimensional vectors. $\bar{\mathbf{L}}(t)$ history of process up till t.

Likelihood of observed data

- We can factor likelihood according to time ordering.
- Assume conditional on past at time $t-1$, independence across individuals of random $L_i(t)$:

$$P(\bar{\mathbf{O}}(\tau+1)) = \prod_t \left\{ \prod_{i=1}^{N} P(L_i(t) \mid \bar{\mathbf{L}}(t-1), \bar{\mathbf{A}}(t-1)) \right\} P(\mathbf{A}(t) \mid \bar{\mathbf{A}}(t-1), \bar{\mathbf{L}}(t-1)).$$

- Short-hand notation for density of $\mathbf{O}(\tau+1)$:

$$\prod_t \left\{ \prod_i q_{i,t}(L_i(t)) \right\} g_t(\mathbf{A}(t)),$$

where g_t denotes the sampling/testing design drawing the testing indicators $(A_1(t), \ldots, A_N(t))$ at time t.

- With the exception of symptomatic subjects, we control drawing $\mathbf{A}(t)$.

Statistical Model: Stationarity in time

- In order to be able to learn the data distribution from the data over time t, we may assume that there is a common mechanism generating $L_i(t+1)$, given the history up till time t, across time t.
- For example, we assume $q_{i,t}(l_i(t) \mid \text{Past(t)}) = q(l_i(t) \mid z_i(t))$ is modeled by common conditional density that only depends on past through fixed dimensional extraction $z_i(t)$, but we leave this function q unspecified.
- Stationarity in time assumption is often problematic, but over time, one might be able to learn and measure the factors that result in changes in the data generating process.

Stationarity across individuals

- One can also avoid making such an assumption, and rely on asymptotics in number of individuals N.
- In this case, we assume that at time t, there is a common in i conditional density $q_t(L_i(t) \mid Z_i(t))$, allowing that it changes over time t.
- In this case, one can still learn the data generating distribution, but one would only use the recent estimates of q_t to optimize the next sampling mechanisms w.r.t. the status of epidemic x-time points in the future (moving target).

Outcome of interest: Evaluation of adaptive design g

- The performance of the design could be measured by a final proportion of actively infected:

$$Y^c = \frac{1}{N} \sum_{i=1}^{N} \frac{A_i(\tau)}{g_{\tau,i}(A_i(\tau))} Y_i(\tau + 1).$$

- This inverse probability of sampling weighted proportion of actively infected has the same expectation as the actual proportion of actively infected subjects.
- A design g could therefore be evaluated by the expectation of Y^c

$$E_{g,q} Y^c.$$

- For any choice g, true q, assuming sequential randomization, this equals the mean outcome in the counterfactual world in which we would have employed adaptive design g (i.e., causal quantity, identified by G-computation formula).

Oracle adaptive design

- We could restrict to a class $\mathcal{G}$ of conditional distributions $(g_t : t = 1, \ldots)$ that are only allowed to respond to certain extractions from the past at time t.
- Among that class of possible adaptive designs, we can define an oracle design for our study with mechanism q:

$$g(q) = \arg\min_{g \in \mathcal{G}} E_{g,q} Y^c.$$

- Oracle design depends on the unknown data generating function q.

Adaptive designs that learn the oracle design

- At time t, we can use the past data, to estimate q.
- Let $q_{N,t}$ be this estimator.
- Then, we have an estimator $g_{N,t} = q(q_{N,t})$ of the oracle adaptive design.
- We can use this $g_{N,t}$ to sample the next $\mathbf{A}(t+1)$.
- As time t increases we have that $g_{N,t}$ converges to the oracle design $g(q)$.
- We can consider analogue adaptive designs for the case that we have stationarity across subjects, rely on N large, and select g to optimize

$$g \rightarrow E_{q,g} Y^c_{t+x}.$$

Theoretical advances supporting such complex adaptive surveillance and treatment systems

- Online super-learning for time-series (vdL, Malenica, 2019, Benkeser et al, 2018, vdL,Rose, 2018)
- Highly Adaptive Lasso: General nonparametric MLE machine learning (vdL 2015, vdL, Rose, 2018)
- TMLE to obtain efficient and normally distributed estimators, fully utilizing machine learning, of quantities such as EY^c, or contrasts of different counterfactual mean outcomes (vdL, Rubin, 2006, vdL, Rose , 2011, 2018)
- Weak convergence theory for processes; i.e., Martingale CLT and probability inequalities for martingales.

Conclusions and Future Work

- Building a real-world adaptive surveillance and response system poses formidable challenges
  - Statistical
  - Operational
- But also holds immense promise.
- Lots of highly relevant advances are in place,
  - But still work to do

Acknowledgements

- UC Berkeley Biostatistics
  - Ivana Malenica
  - Jeremy Coyle
  - Aurelien Bibaut
- UCSF
  - Diane Havlir
- SF DPH
- Latino Taskforce
- Unidos en Salud Team
- Berkeley Safe Campus Study
  - Art Reingold
  - Laura Packel
  - Lauren Hunter
  - Shelley Facente
  - Safe Campus Team
- ADAPT Team
  - Elvin Geng

References I

Ricardo Andrade Pacheco, Francois Rerolle, Jean Lemoine, Leda Hernandez, Aboulaye Meïté, Aurelien Bibaut, Mark van der Laan, Benjamin Arnold, and Hugh JW Sturrock. Finding hotspots: development of an adaptive spatial sampling approach, 2020. URL https://doi.org/10.1101/2020.01.10.20016964.

D. Benkeser and M.J. van der Laan. The highly adaptive lasso estimator. *2016 IEEE International Conference on Data Science and Advanced Analytics*, PMC5662030.

D. Benkeser, C. Ju, S. Lendle, and M.J. van der Laan. Online cross-validation-based ensemble learning. *Statistics in Medicine*, 37(2):249–260, 2018. ISSN 1097-0258. doi: 10.1002/sim.7320.

References II

A. Bibaut, A. Chambaz, and M.J. van der Laan. Rate-adaptive model selection over a collection of black-box contextual bandit algorithms. Technical report.

A. Bibaut, A. Chambaz, and M.J. van der Laan. Some advances in targeted sequential learning and an application to disease hotspot identification. Technical Report Technical report, 2020a.

A. Bibaut, A. Chambaz, and M.J. van der Laan. Generalized policy elimination: an efficient algorithm for nonparametric contextual bandits. *Uncertainty in Artificial Intelligence*, 2020b.

References III

Frank Blaauw and Antoine Chambaz. OnlineSuperLearner: Superlearner with online functionality for time-series analysis. https://github.com/frbl/OnlineSuperLearner, 2017. URL https://github.com/frbl/OnlineSuperLearner. R package version 0.1.0.

A. Chambaz, W. Zheng, and M. J. van der Laan. Targeted sequential design for targeted learning inference of the optimal treatment rule and its mean reward. *Ann Stat*, 45 (6):2537–2564, 2017.

References IV

Antoine Chambaz and Mark J. van der Laan. Targeting the optimal design in randomized clinical trials with binary outcomes and no covariate: Simulation study. *The International Journal of Biostatistics*, 7(1): 00001022021557467913 10, 2011. doi: https://doi.org/10.2202/1557-4679.1310. URL https://www.degruyter.com/view/journals/ijb/7/1/article-00001022021557467913 10.xml.

J. R. Coyle, N. S. Hejazi, I. Malenica, and O. Sofrygin. sl3: modern super learning with pipelines. https://github.com/jeremyrcoyle/sl3, 2018. URL https://github.com/tlverse/sl3. R package version 0.1.0.

References V

A. Luedtke and M. van der Laan. Super-learning of an optimal dynamic treatment rule. *The International Journal of Biostatistics*, 12(1):305–332, 2016a. doi: 10.1515/ijb-2015-0052.

A. Luedtke and M. van der Laan. Optimal individualized treatments in resource-limited settings. *The International Journal of Biostatistics*, 12(1):283–303, 2016b. doi: 10.1515/ijb-2015-0007.

A. R. Luedtke and M. J. van der Laan. Statistical inference for the mean outcome under a possibly non-unique optimal treatment strategy. *Ann Stat*, 44(2):713–742, Apr 2016c.

References VI

I. Malenica and M. J. van der Laan. tstmle: Data-adaptive estimation and inference for causal effects with a single time series. https://github.com/podTockom/tstmle, 2018. URL https://github.com/podTockom/tstmle. R package version 0.1.0.

Mark J. van der Laan and Ivana Malenica. Robust estimation of data-dependent causal effects based on observing a single time-series, 2019.

M.J. van der Laan. A generally efficient targeted minimum loss estimator based on the highly adaptive lasso. *Int. J. Biostat.*, 13(2). ISSN PMID 29023235.

M.J. van der Laan. The construction and analysis of adaptive group sequential designs. Technical Report Working paper 232, www.bepress.com/ucbbiostat/paper232, 3 2008.

References VII

M.J. van der Laan and S. Rose. *Targeted Learning in Data Science: Causal Inference for Complex Longitudinal Studies*. Springer Science and Business Media, 2018.

M.J. van der Laan and D. Rubin. Targeted maximum likelihood learning. *International Journal of Biostatistics*, 2(1): www.bepress.com/ijb/vol2/iss1/11, 2006.

van der Laan M.J. and Rose S. *Targeted Learning: Causal Inference for Observational and Experimental Data (Springer Series in Statistics)*. Springer, 2011.

van der Laan M.J., Chambaz A., and Lendle S. *Online Targeted Learning for Time Series*, pages 317–346. Springer International Publishing, Cham, 2018. ISBN 978-3-319-65304-4. doi: 10.1007/978-3-319-65304-4_19. URL https://doi.org/10.1007/978-3-319-65304-4_19.