Optimal spatio-temporal treatment allocations

Eric B. Laber

Department of Statistics, North Carolina State University

RDLM, June 10, 2015

So, what brings you to Edmonton?
–M. Littman

Example: white-nose syndrome

- Recent headlines
  - *Fungus that's killing millions of bats isn't going away*
    -Los Angeles Times, Nov. 1, 2013
  - *The secret bataclysm: white nose syndrome and extinction*
    -Wired, Aug. 12, 2014
  - *White-nose syndrome has almost completely wiped out some North American bat colonies*
    -Science, Feb. 15, 2015

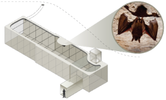

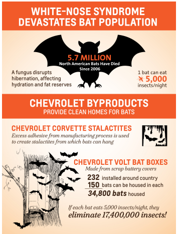

Bat images from AP photo, poster from 3blmedia.com

Example: white-nose syndrome cont'd

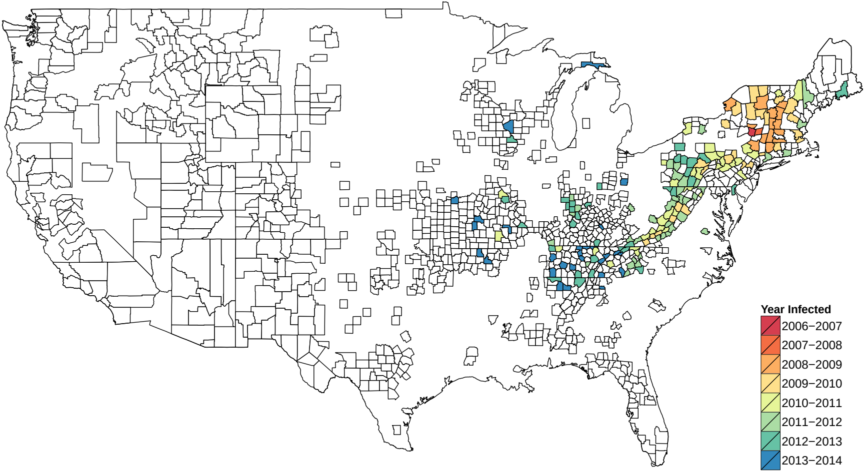

- Allocation strategy: map from current epidemic status to subset of counties marked 'high priority for treatment'

Additional motivating examples

- Other spatio-temporal allocation problems:
  - Spread of human disease (e.g., Ebola)
  - Wildlife management: allocating catch-sizes in each county
  - Adaptive conservation
  - Forest fires: where and when to do controlled burns
  - Crime prevention: assigning police officers to locations, i.e., 'cops on dots'
  - Precision agriculture: adaptively adjust nutrients over space and time

Major challenges

- Data
  - No data on intervention effectiveness at $t = 1$
  - Noisy, incomplete, and sparse state measurements
  - High-dimensional, e.g., in WNS, $L = 1,137$ locations x 10 measurements per location = 11,370 measurements per time point
- Estimation and inference
  - Online estimation means allocations are selected from a continually changing (data-dependent) strategy $\Rightarrow$ sampling distns are complex
- Computation
  - There are 2^L possible allocations at each time point
  - Enumeration is not possible

Setup and notation

- Allocation problem evolves over
  - Locations $\mathcal{L} = \{1, 2, \ldots, L\}$
  - Time points $\mathcal{T} = \{1, 2, \ldots\}$
- At each time t and location ℓ
  - Observe state $\boldsymbol{S}_\ell^t \in \mathbb{R}^p$
  - Observe outcome $Y_\ell^t \in \mathbb{R}$ (higher is better)
  - Select allocation $A_\ell^t \in \{0, 1\}$
  - Define $\boldsymbol{S}^t = \{\boldsymbol{S}_\ell^t\}_{\ell \in \mathcal{L}}$, $\boldsymbol{A}^t = \{A_\ell^t\}_{\ell \in \mathcal{L}}$, and $\boldsymbol{Y}^t = \{Y_\ell^t\}_{\ell \in \mathcal{L}}$
- Assume allocation based on $\boldsymbol{S}^t$ at time t

Stochastic allocation strategy

- Random allocations needed to learn in online setting
- Let $\mathcal{B}_L$ denote all distributions over $\{0,1\}^L$ and $\mathcal{S} = \operatorname{dom} \boldsymbol{S}^t$
- Allocation strategy π is a map

$$\pi : \mathcal{S} \to \mathcal{B}_L,$$

under π, a decision maker presented with $\boldsymbol{S}^t = \boldsymbol{s}^t$ will select $\boldsymbol{a}^t$ with probability $\pi(\boldsymbol{a}^t; \boldsymbol{s}^t)$

Class of potential strategies

- Focus on estimation within class of strategies Π
  - Enforce feasibility and cost constraints
  - Add'l considerations: parsimony, interpretability, and scalability

- Ex. probit rank model

$$\Pi^{\text{ex.}} = \left\{ \pi_{\rho,\sigma}(\boldsymbol{a}; \boldsymbol{s}) \propto \Phi \left\{ \psi(\boldsymbol{s}, \boldsymbol{a}; \rho)/\sigma \right\}, \rho \in \mathbb{R}^{p}, \sigma \in \mathbb{R}_{+} \right\},$$

  where $\psi(\boldsymbol{s}, \boldsymbol{a}; \rho)$ is a feature vector indexed by $\rho \in \mathbb{R}^{p}$, and Φ is the CDF of a standard normal random variable

Optimal strategy via potential outcomes

- Define $\mathcal{F} = \left\{ \boldsymbol{a} \in \{0,1\}^L : \boldsymbol{a} \in \operatorname{dom} \pi \text{ for some } \pi \in \Pi \right\}$

- Potential outcomes under $\overline{\boldsymbol{a}}^t = (\boldsymbol{a}^1, \dots, \boldsymbol{a}^t)$

$$\boldsymbol{W} = \left\{ \boldsymbol{Y}^{*t}(\overline{\boldsymbol{a}}^t), \boldsymbol{S}^{*(t+1)}(\overline{\boldsymbol{a}}^t) : \boldsymbol{a}^t \in \mathcal{F} \right\}_{t \in \mathcal{T}}$$

- Potential outcome under strategy π

$$\boldsymbol{Y}^{*t}(\pi) = \sum_{\overline{\boldsymbol{a}}^t} \boldsymbol{Y}^{*t}(\overline{\boldsymbol{a}}^t) \prod_{v=1}^{t} \mathcal{I}\left[\xi_\pi^v \left\{ \boldsymbol{S}^{*v}(\overline{a}^{v-1}) \right\} = \boldsymbol{a}^v \right],$$

where $\{\xi_\pi^t(\boldsymbol{s})\}_{t \in \mathcal{T}, \boldsymbol{s} \in \mathcal{S}}$ collection independent r.v.s with $P\{\xi_\pi^v(\boldsymbol{s}^v) = \boldsymbol{a}^v\} = \pi(\boldsymbol{a}^v; \boldsymbol{s}^v)$

Optimal strategy via potential outcomes cont'd

- Discounted marginal mean outcome under π is

$$V(\pi) = \mathbb{E}\left[\sum_{t\in\mathcal{T}} \gamma^{t-1} u\left\{\boldsymbol{Y}^{*t}(\pi)\right\}\right],$$

  where $\gamma \in (0,1)$ and $u(\cdot)$ a utility fn

- Optimal strategy satisfies $V(\pi^{\mathrm{opt}}) \geq V(\pi)$ for all π

Optimal strategy via generative model

- Let $\boldsymbol{H}^t$ be history at time t
- Standard assumptions:
  - (A1) Sequential ignorability: $\boldsymbol{A}^t \perp\!\!\!\perp \boldsymbol{W} \mid \boldsymbol{H}^t$
  - (A2) Consistency: $\boldsymbol{Y}^t = \boldsymbol{Y}^{*t}(\overline{\boldsymbol{A}}^t)$, $\boldsymbol{S}^t = \boldsymbol{S}^{*t}(\overline{\boldsymbol{A}}^t)$
  - (A3) Positivity: there exists $\epsilon > 0$ s.t. $P\left\{\boldsymbol{A}^t = \boldsymbol{a} \mid \boldsymbol{H}^t\right\} \geq \epsilon$ with probability one for all $\boldsymbol{a} \in \mathcal{F}$

Optimal strategy via generative model

- Let $\boldsymbol{H}^t$ be history at time t
- Standard assumptions:
  - (A1) Sequential ignorability: $\boldsymbol{A}^t \perp\!\!\!\perp \boldsymbol{W} \mid \boldsymbol{H}^t$
  - (A2) Consistency: $\boldsymbol{Y}^t = \boldsymbol{Y}^{*t}(\overline{\boldsymbol{A}}^t)$, $\boldsymbol{S}^t = \boldsymbol{S}^{*t}(\overline{\boldsymbol{A}}^t)$
  - (A3) Positivity: there exists $\epsilon > 0$ s.t. $P\left\{\boldsymbol{A}^t = \boldsymbol{a} \mid \boldsymbol{H}^t\right\} \geq \epsilon$ with probability one for all $\boldsymbol{a} \in \mathcal{F}$

Spillover effects due to spatial proximity violate SUTVA $\Rightarrow$ experimental unit is the entire collection of locations $\mathcal{L}$

Optimal strategy via generative model cont'd

- Under (A1)-(A3), $\mathbb{E}\left[u\left\{\boldsymbol{Y}^{*t}(\pi)\right\}\right]$ is

$$\int u(\boldsymbol{y}^t)\prod_{v=1}^{t}\left[f_v\left(\boldsymbol{y}^v \mid \boldsymbol{h}^v, \boldsymbol{a}^v\right)\pi(\boldsymbol{a}^v; \boldsymbol{s}^v)g_v\left(\boldsymbol{h}^v \mid \boldsymbol{h}^{v-1}\right)\right] d\lambda(\overline{\boldsymbol{y}}^t, \overline{\boldsymbol{a}}^t, \overline{\boldsymbol{h}}^t),$$

where f_v is cond'l density of $\boldsymbol{Y}^v$, g_v is cond'l density of $\boldsymbol{H}^v$

  - One observation per time point $\Rightarrow$ cannot identify f_v, g_v
  - Even if $f_v \equiv f$ and $g_v \equiv g$ nonparametric estimation is essentially impossible because $\boldsymbol{a}^t \in \{0,1\}^L$
  - Need assumptions that allow pooling over locations

Parametric estimation

- Assume Markovian, low-dimensional parametric models for systems dynamics
  - Conditional density for outcome $f_v\left(\boldsymbol{y}^v \mid \boldsymbol{h}^v, \boldsymbol{a}^v\right) = f\left(\boldsymbol{y}^v \mid \boldsymbol{s}^v, \boldsymbol{a}^v; \beta\right)$
  - Conditional density for <u>state</u> $g_v\left(\boldsymbol{s}^v \mid \boldsymbol{h}^{v-1}\right) = g(\boldsymbol{s}^v \mid \boldsymbol{s}^{v-1}, \boldsymbol{a}^{v-1}; \theta)$

- Use simulation-optimization to estimate π^{opt}

(S1) Constructed estimators $\widehat{\beta}$, $\widehat{\theta}$ of β, θ

(S2) Draw $\widetilde{\beta}, \widetilde{\theta}$ from sampling distn of $\widehat{\beta}, \widehat{\theta}$

(S3) Simulate process under $\widetilde{\beta}, \widetilde{\theta}$ to form estimator $\widehat{V}(\pi)$ of $V(\pi)$

(S4) Stochastic programming to compute $\widehat{\pi} = \arg\max_{\pi \in \Pi} \widehat{V}(\pi)$

Online parametric estimation

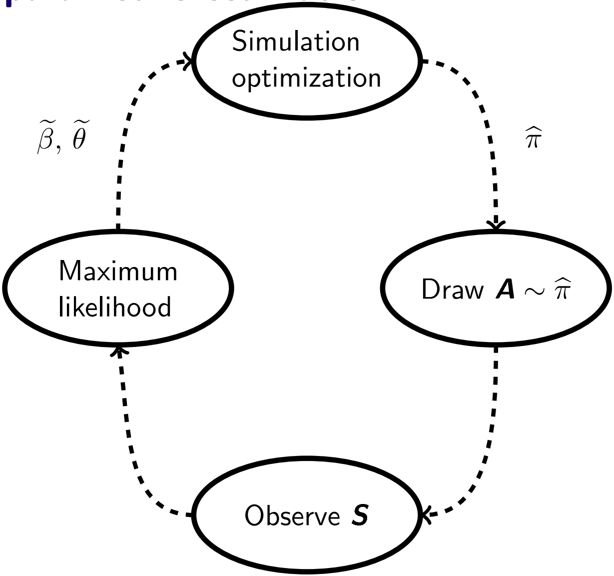

Note: this is Thompson sampling.

Parametric estimation discussion

- Positives
  - Intuitive
  - Leverage existing scientific theory
  - Low variance, applies in data impoverished setting
  - Nested parametric models can be used for hypothesis testing, e.g., is a disease spread diffusive
  - Extensions to non-stationary setting via state-space models
- Drawbacks: the devil is in the details
  - Potentially high bias
  - Computation is non-trivial, requires climbing up-hill in the strategy space Π using stochastic programming
  - Choosing expressive but scalable Π difficult

One-step solution improvement

- Define $\nu(\boldsymbol{s}, \pi) \triangleq \mathbb{E}\left[\sum_{t\geq 1} \gamma^{t-1} u\left\{\boldsymbol{Y}^{*t}(\pi)\right\} \big| \boldsymbol{S}^1 = \boldsymbol{s}\right]$

- Optimal regime, π^{opt}, satisfies

$$\pi^{\mathrm{opt}}(\boldsymbol{s}) = \arg \max_{\boldsymbol{a}^t \in \mathcal{A}(\boldsymbol{s})} \mathbb{E}\left\{u(\boldsymbol{Y}^t) + \gamma\nu(\boldsymbol{S}^{t+1}, \pi^{\mathrm{opt}}) \big| \boldsymbol{S}^t = \boldsymbol{s}, \boldsymbol{A}^t = \boldsymbol{a}^t\right\}$$

- One-step estimator

$$\widetilde{\pi}^t(\boldsymbol{s}) = \arg \max_{\boldsymbol{a}^t \in \mathcal{A}_{\widehat{\pi}^t}(\boldsymbol{s})} \widehat{\mathbb{E}}\left\{u(\boldsymbol{Y}^t) + \gamma\widehat{\nu}(\boldsymbol{S}^{t+1}, \widehat{\pi}) \big| \boldsymbol{S}^t = \boldsymbol{s}, \boldsymbol{A}^t = \boldsymbol{a}^t\right\},$$

where $\widehat{\mathbb{E}}$ and $\widehat{\nu}$ estimated via Monte Carlo

Simulation experiments: overview

- Simulate spread of disease across nodes in a network
  - Randomly select 1% of nodes to infect at baseline
  - No interventions for 6 time steps
  - Select 6% of nodes for treatment at $t = 7, \ldots, 15$
- Generative model built from white-nose syndrome data
  - Variant of gravity model (Maher et al., 2012), probability that node i infects node j is linear on logit scale
  - Coefficients estimated using white-nose data then scaled s.t. spread to 70% of the network at $t = 15$ under no treatment
  - Additional initial state values at each node generated from Normal$(\mathbf{0}_{10}, \mathbf{I}_{10})$
- Goal: minimize number infected nodes at $t = 15$

Simulation experiments: setup

- Competing allocation strategies
  - NoTxt: no treatment
  - Myopic: allocate to nodes with highest predicted probability of infection at next time point
  - Proximal: ad hoc strategy proposed by USFWS that treats locations on the 'border' of a spreading infection
  - SimOpt: simulation optimization with perturbed linear ranks
  - OneStep: one-step updated estimator
- Consider correct and incorrect dynamics models
  - Correct: time-dependent SI model (Maher et al., 2012)
  - Incorrect: estimated dynamics only depend on distance
- Use probit rank class of strategies with linear features constructed from postulated dynamics model

Simulation: scale-free example

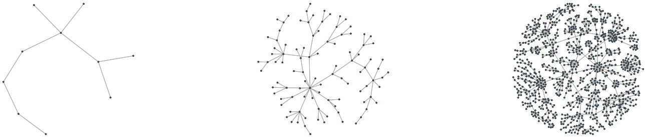

Correctly specified dynamics model

Nodes	NoTxt	Myopic	Proximal	SimOpt	OneStep
100	0.70	0.64	0.56	0.51	0.50
500	0.70	0.61	0.62	0.47	0.47
1000	0.70	0.63	0.64	0.52	0.52

Incorrectly specified dynamics model

Nodes	NoTxt	Myopic	Proximal	SimOpt	OneStep
100	0.70	0.65	0.56	0.52	0.50
500	0.70	0.63	0.62	0.53	0.52
1000	0.70	0.66	0.64	0.55	0.55

Simulation: random 3-nearest neighbors example

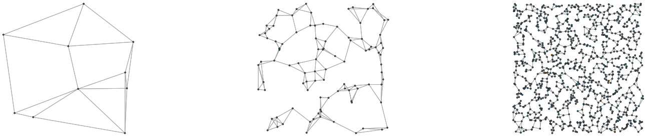

Correctly specified dynamics model

Nodes	NoTxt	Myopic	Proximal	SimOpt	OneStep
100	0.70	0.36	0.31	0.27	0.27
500	0.70	0.45	0.37	0.31	0.30
1000	0.70	0.53	0.48	0.46	0.46

Incorrectly specified dynamics model

Nodes	NoTxt	Myopic	Proximal	SimOpt	OneStep
100	0.70	0.45	0.31	0.41	0.29
500	0.70	0.55	0.37	0.45	0.32
1000	0.70	0.58	0.48	0.54	0.46

Simulation: white-nose syndrome

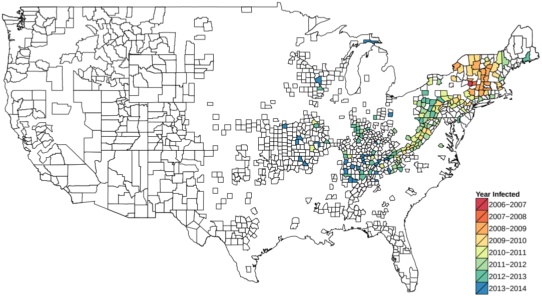

Correctly specified dynamics model

Nodes	NoTxt	Myopic	Proximal	SimOpt	OneStep
1137	0.60	0.40	0.45	0.30	0.33

Incorrectly specified dynamics model

Nodes	NoTxt	Myopic	Proximal	SimOpt	OneStep
1137	0.60	0.54	0.45	0.49	0.36

Discussion

- Management of epidemics can be posed as spatio-temporal allocation problems
- An effective allocation strategy:
  - Accounts for spillover effects (no SUTVA)
  - Implemented online
  - Computationally feasible
  - Accommodates evolving logistical constraints
- We proposed a parametric, model-based estimator

Discussion cont'd

- Many open and exciting problems
  - Imperfect measurement/detection
  - Scaling to networks with millions or billions of nodes (we currently have heuristics that apply to $\sim$10 million nodes)
  - Semi-parametric and non-parametric estimation
  - Choose amount of stochasticity in allocation strategy to optimize exploration/exploitation trade-off

Acknowledgments

Thanks to

- RLDM Organizing Committee
- NCSU/UNC AI and Personalized Medicine Lab
- PR Department of Natural Resources
- Patuxent Wildlife Research Center

Joint work with

- Jaime Collazo
- John Drake
- Nick Meyer
- Krishna Pacifici
- Brian Reich

Thank you. Comments, suggestions, collaborations?

laber@stat.ncsu.edu

Simulation: white-nose syndrome

- Build systems dynamics model using data from 2006-2012
- Use linear strategies $\pi(\boldsymbol{s};\gamma) = \arg\max_{\boldsymbol{a}\in\mathcal{A}^t} \phi(\boldsymbol{s},\boldsymbol{a})^{\intercal}(\gamma+\xi)$
- Spatio-temporal gravity model, Y_ℓ^t, infectious status

$$P\left(Y_\ell^t = 0 \middle| \boldsymbol{H}^t, \boldsymbol{A}^t\right) = \begin{cases} 0 & Y_\ell^{t-1} = 1 \\ \prod_{j\in\mathcal{I}^{t-1}}(1-p_{j\ell}^t) & Y_\ell^{t-1} = 0, \end{cases}$$

where $\mathcal{I}^t$ is the set of infected counties at time t, $p_{j\ell}^t$ is prob. of spread from j to ℓ at time t

Simulation: white-nose syndrome cont'd

- Probability of spread from j to ℓ at time t

$$\mathrm{logit}\left(p_{j\ell}^{t}\right) = \beta^{\intercal}\boldsymbol{X}_{j}^{t} + \alpha A_{\ell}^{t} - \frac{\rho d_{j\ell}}{(m_j m_\ell)^{\nu}},$$

  where:
  - $\boldsymbol{X}_{j}^{\intercal}$ are covariates and β unknown coeffs
  - A_{j}^{t} indicates treatment, α treatment effect
  - Final term controls spatial dependence
    - $d_{j\ell}$ distance between j and ℓ
    - m_j nbr caves in j
    - ρ and ν unknown parameters
- See Maher et al. (2012) for details

Simulation: white-nose syndrome cont'd

- Simulation procedure:
  1. Simulate data from estimated model, no treatments in 2006-2012
  2. Use ad hoc proximity-based strategy to choose locations in 2012
  3. Start estimating optimal strategy in 2013
  4. Run through 2021
- Restrict number of treated counties to 60 at each time point
- Compare with proximity-based strategy and myopic prediction-based strategy

Simulation: white-nose syndrome results

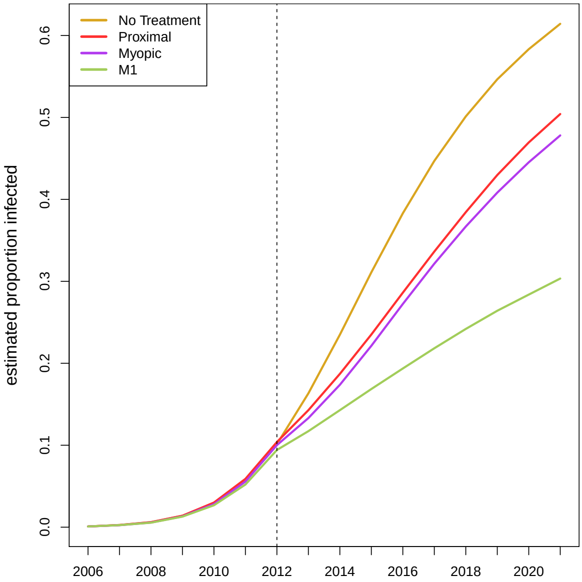

Simulation: white-nose syndrome discussion

- Online anchored estimation performs favorably with competing strategies
- Computational burden is significant information in plot took approximately 20 hours on 64 cores to construct
- Class of policies has major impact on complexity/performance, we are currently working on this

Example two: ebola virus disease

- Recent headlines
  - *In West Africa, disease just as devastating as war*
    –Globe and Mail, Nov. 26, 2014
  - *Despite Aid Push, Ebola Is Raging in Sierra Leone*
    –New York Times, Nov. 27, 2014
  - *Global Ebola Death Toll Exceeds 5,600*
    –Wall Street Journal, Nov. 26, 2014
  - *Ebola Fight Far From Over*
    –Global Times, Nov. 27, 2014

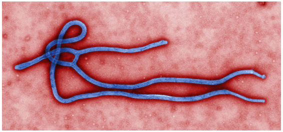
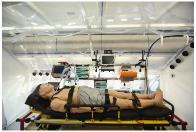


Ebola virion image from CDC PHIL, plane image and burial worker images from AP photo

Example two: ebola virus disease cont'd

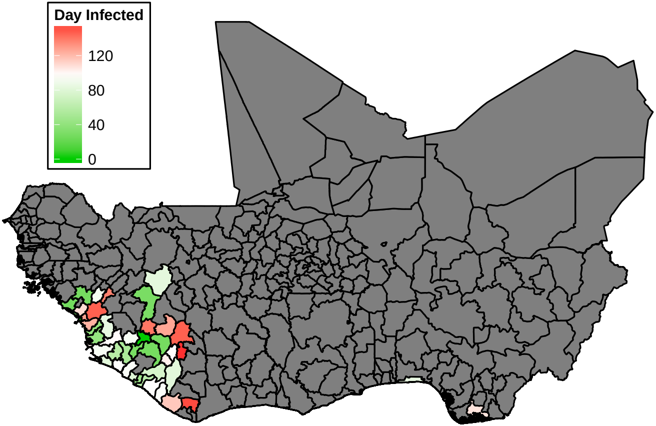

Example two: ebola virus disease cont'd

- Goal: minimize number infected individuals
- Impact
  - Thousands of human lives
  - World bank estimated cost of management at 33 billion USD
  - Political and economic stability
- Potential interventions
  - Quarantine protocols / movement limitations / border closings
  - Vaccination*
- Allocation strategy: map from current epidemic status to subset of locations for mobile treatment units

Simulation: Ebola

- Build systems dynamics model using data from last 157 days
- Model currently under development, but similar to spatio-temporal gravity model
  - Population density
  - Population movements (historical)
- Simulate 14 weeks of treatment
  - Apply treatments weekly
  - Restrict number of treated locations
  - Compare with proximity-based and myopic strategies
  - Use linear strategies

Simulation: Ebola results

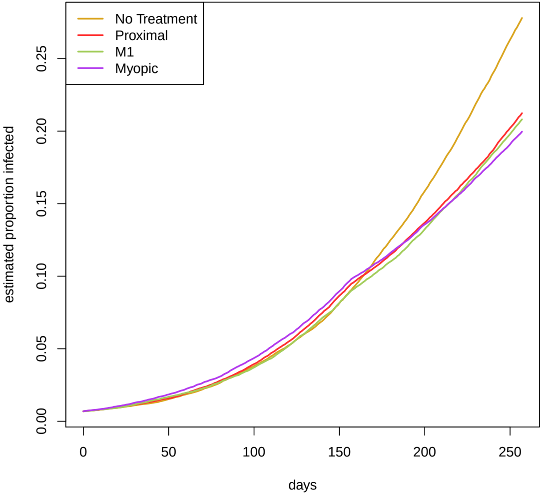

Simulation: Ebola discussion

- No difference between ad hoc and anchored strategies
- Spread model very coarse, more data are coming
- Small number of locations, more expressive strategies are possible