Analytics for community first response

Dutch Optimization Seminar
23-04-2024

Caroline Jagtenberg
Dept. of Operations Analytics, Vrije Universiteit

joint work with:

Shane Henderson Pieter van den Berg Maggie Li

Bridget Dicker Océane Fourmentraux

Community First Responder Systems

OHCA

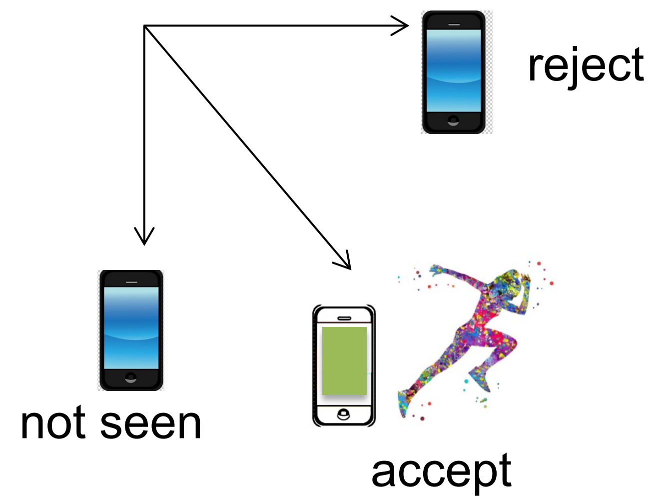

Goal:

- shorten time to CPR
- improve survival rates

VU VRIJE UNIVERSITEIT AMSTERDAM
School of Business and Economics

Community First Responder Systems

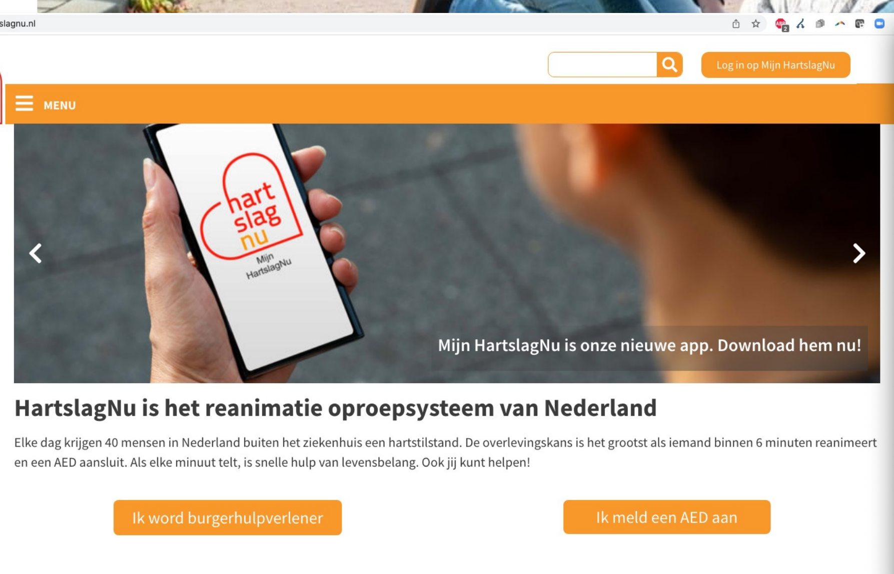

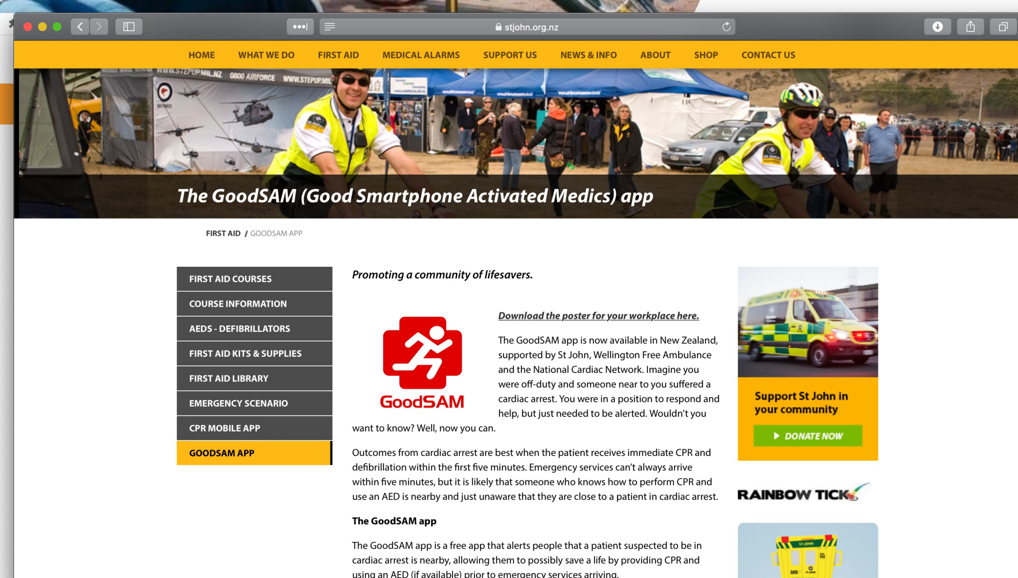

Previous work

Common are **retrospective** studies.

Uncommon are studies **proactively** investigating a system.
What if:

- we recruit more volunteers?
- we optimizing system's alerting settings?

Volunteer recruitment

How to quantify the impact of *n* volunteers on patient survival?

- What is the **arrival-time distribution** of the first-arriving responder?

OHCA survival function

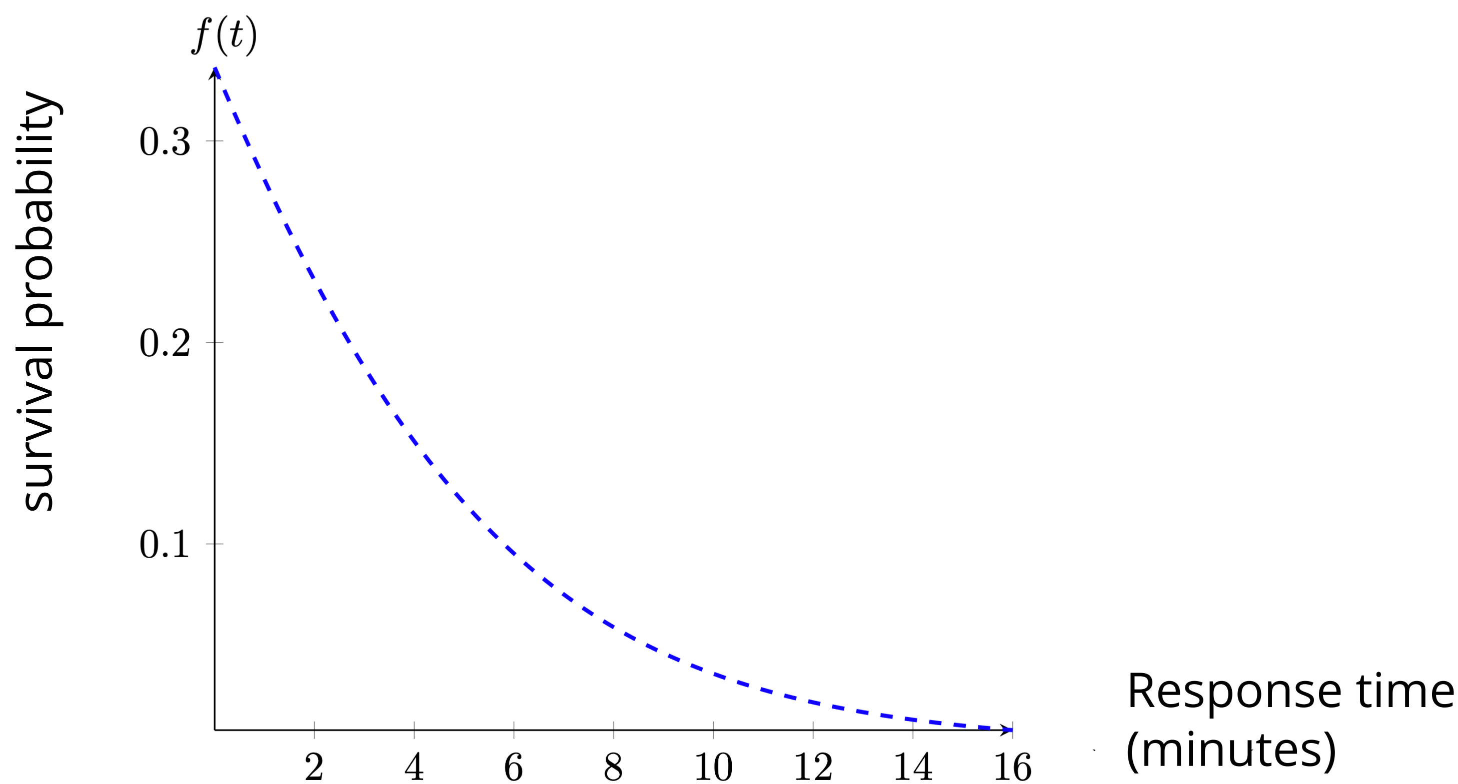

Approach

We'll model the locations of volunteers as a Poisson Point Process.

It may have a different μ in different parts of the city, though.

And recruiting would influence the μ.

Tolhuistuin
IJPLEIN
STAATSLIEDENBUURT
Kattensloot
Sexmuseum Amsterdam
FREDERIK HENDRIKBUURT
Marnixstraat
Anne Frank Huis
JORDAAN
De Oude Kerk
NEMO Science Museum
NIEUWMARKT EN LASTAGE
Dijksgracht
Rozengracht
Singel
DE WALLEN
X
Het Scheepvaartmuseum
De 9 Straatjes
Begijnhof
Museum Rembrandthuis
KADIJKEN
DA COSTABUURT
Bilderdijkgracht
Kloveniersburgwal
ARTIS - Micropia
Hortus Botanicus Amsterdam
ARTIS
Plantage Middenlaan
OUD-WEST
GRACHTENGORDEL
H'ART Museum
Foam
Holland Casino Amsterdam
citizenM Amstel Amsterdam
Best beoordeeld
Wereldmuseum Amsterdam
Spiegelgracht
VONDELPARKBUURT
Rijksmuseum
Lijnbaansgracht
Oosterpark
OOSTERPARKBUURT
S100
S116
S105
S112

Tolhuistuin
STAATSLIEDENBUURT
S100
Kattensloot
s116
IJPLEIN
S100
Sexmuseum Amsterdam
FREDERIK HENDRIKBUURT
Marnixstraat
s105
S100
Anne Frank Huis
JORDAAN
De Oude Kerk
NEMO Science Museum
Dijksgracht
S100
NIEUWMARKT EN LASTAGE
s116
Het Scheepvaartmuseum
Rozengracht
Singel
De 9 Straatjes
Begijnhof
Kloveniersburgwal
museum Rembrandthuis
KADIJKEN
DA COSTABUURT
Bilderdijkgracht
ARTIS - Micropia
Hortus Botanicus Amsterdam
ARTIS
Plantage Middenlaan
OUD-WEST
GRACHTENGORDEL
H'ART Museum
Foam
Holland Casino Amsterdam
citizenM Amstel Amsterdam
Best beoordeeld
Wereldmuseum Amsterdam
Spiegelgracht
S100
VONDELPARKBUURT
Rijksmuseum
Lijnbaansgracht
S112
Oosterpark
OOSTERPARKBUURT

Response-time distribution

Consider a patient at location *l*. Let *R* be the (random) response time of the <u>closest</u> available volunteer.

$P(R > \text{t minutes}) =$

$P(\text{0 volunteers within distance } d_t)$

$= e^{-\mu(B(l,t))}$

$= \exp(-\text{density}_l\, \boldsymbol{\pi}\, d_t^2\ )$

density of *accepting* volunteers (so after thinning) around location l

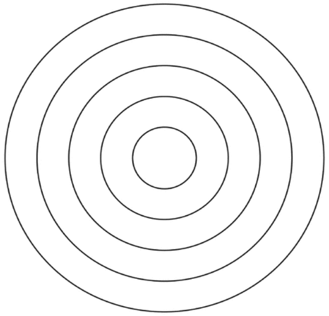

Response time distribution

Assuming volunteers walk at 6 km/h, we obtain an exact expression for the on-foot response time of closest volunteer:

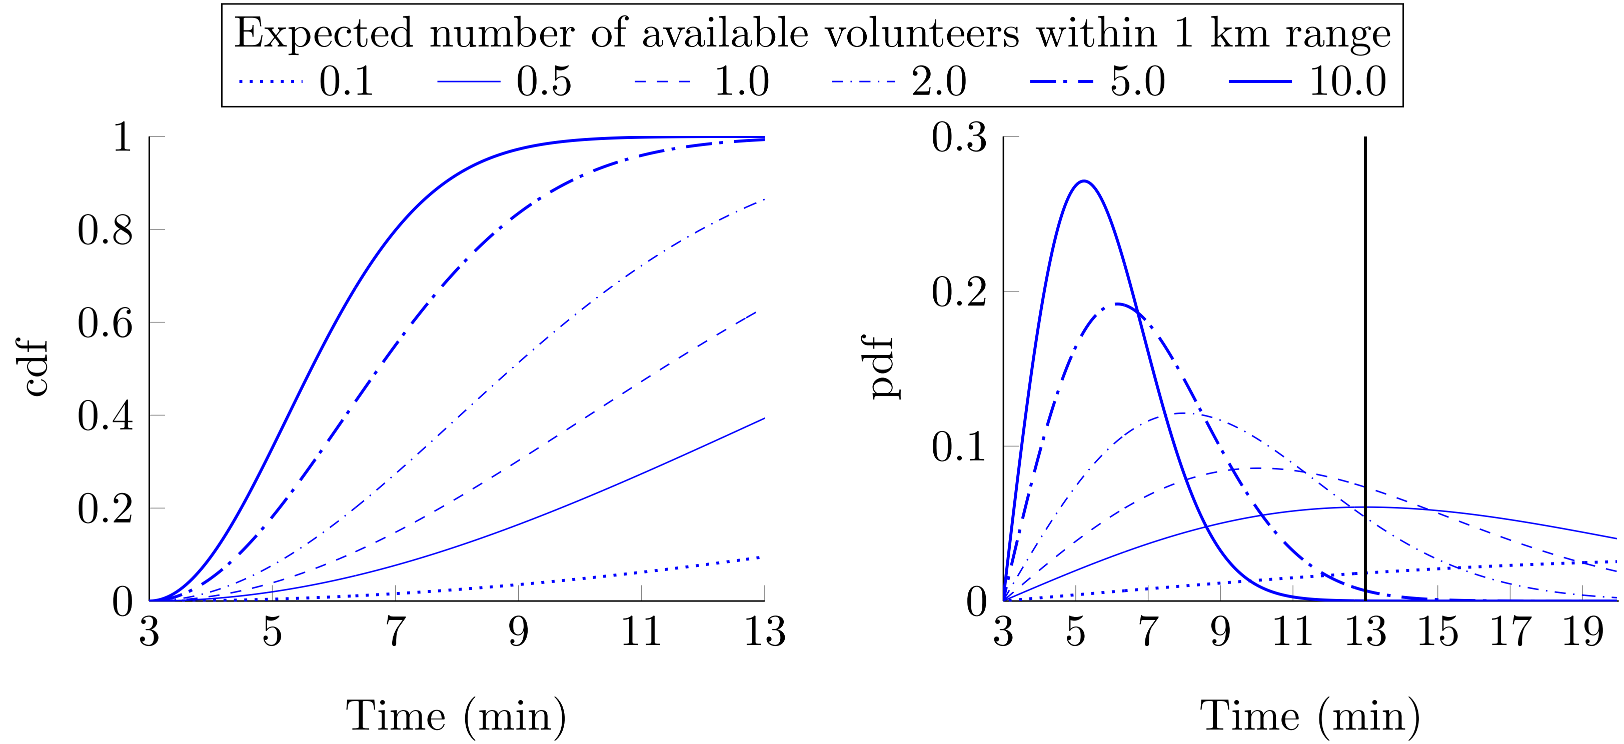

First result

Required density of available volunteers
(per km^2) to meet targets

		Response-time target (minutes)									
		3	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**
Cov.	**0.5**	22.06	5.52	2.45	1.38	0.88	0.61	0.45	0.34	0.27	0.22
	0.7	38.32	9.58	4.26	2.40	1.53	1.06	0.78	0.60	0.47	0.38
	0.9	73.29	18.32	8.14	4.58	2.93	2.04	1.50	1.15	0.90	0.73

But remember: we have more than just response-time goals.
We have survival goals.

So we'll have to integrate the survival function against our obtained response-time PDF, to get probability of survival.

Extend to a heterogeneous area (e.g. city)

Partition the city into regions indexed by $l \in \{1, \ldots, L\}$.

- Let λ_l be the OHCA rate of region *l* (input).
- Let ν_l be the probability of finding a volunteer in region *l*. (Unknown, but at least $\sum_{l=1}^{L} \nu_l = 1$.)
- Assume volunteer density is uniform within a region.

Case study

We consider an urban area of **Auckland, New Zealand** that is discretized into 287 so-called area units.

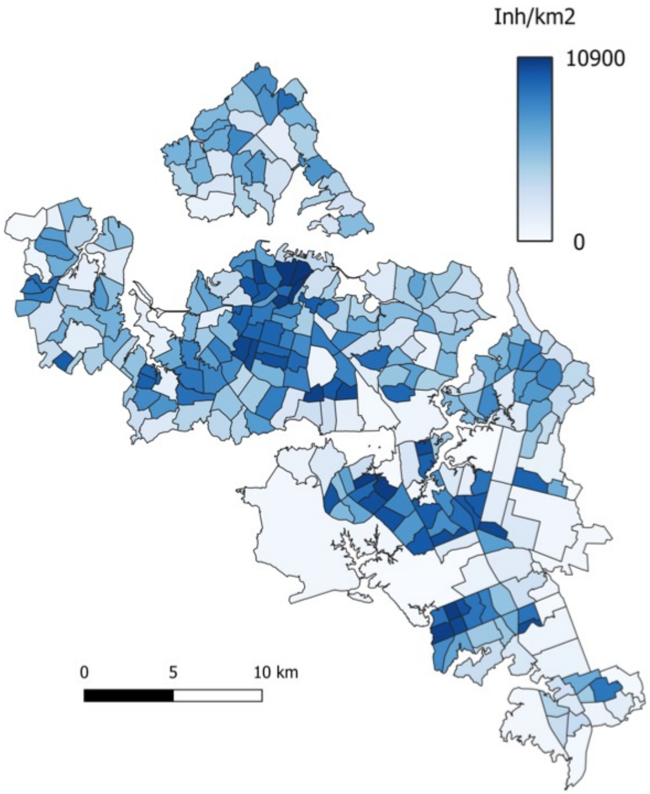

Extending to a heterogeneous area (e.g. city)

$$\mathbb{P}(T > \tau) = \sum_{l} \lambda_l \mathbb{P}(T(l) > \tau) = \sum_{l \in \mathcal{L}} \lambda_l \exp\left(-\pi d_\tau^2 \overbrace{n\alpha\nu_l/a_l}^{\text{density}}\right)$$

We have good estimates for λ_l, but not for ν_l

We can make some assumptions on ν_l , for example:
- proportional to inhabitants of location l

and evaluate this function above. Also transforming this to survival probabilities is no problem.

Auckland, ν_l proportional to inhabitants

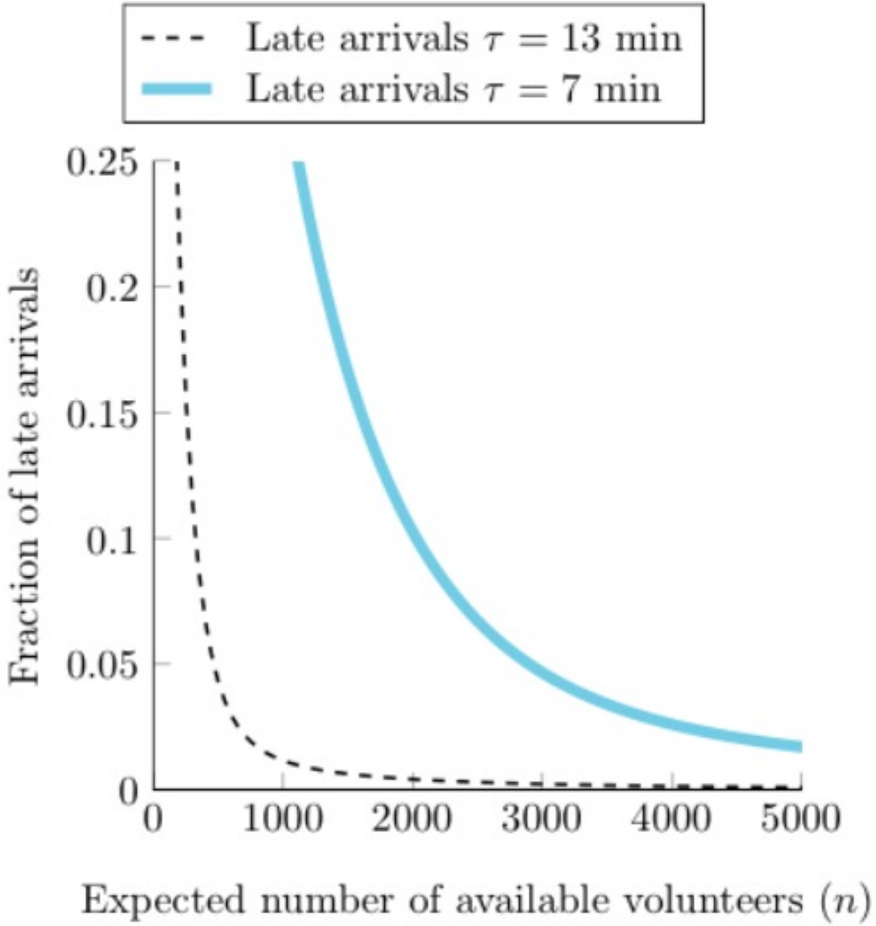

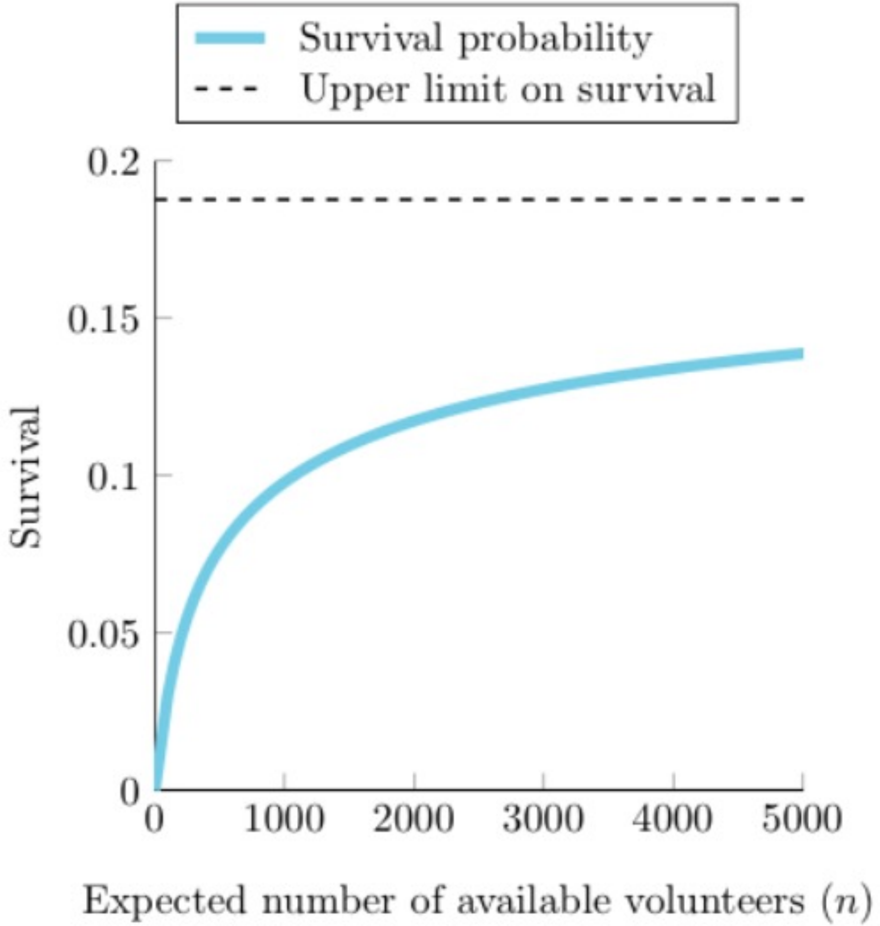

How to choose ν_l

Let's turn this into an optimization question:

What location measure υ gives the best survival over the whole city?

- This provides a bounds on what can be achieved with n volunteers.
- Can also guide recruitment efforts.

Optimizing where volunteers are in the city

$$\mathbb{P}(T > \tau) = \sum_{l} \lambda_l \mathbb{P}(T(l) > \tau) = \sum_{l \in \mathcal{L}} \overset{\text{input}}{\lambda_l} \exp\left(-\pi d_\tau^2 \overbrace{n\alpha \underset{\text{variables}}{\nu_l} / a_l}^{\text{density}}\right)$$

Proposition: This function is convex in the probabilities ν_l

→ Can use convex optimization methods to

minimize P(T > τ) ← can even do this exact

maximize survival ← exact up to step size ε

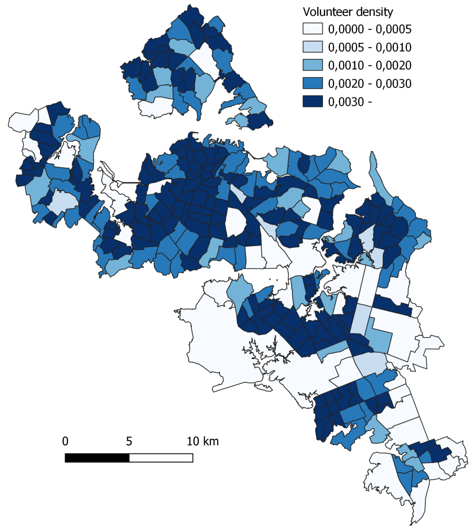

(a) Late arrival, $n = 500$

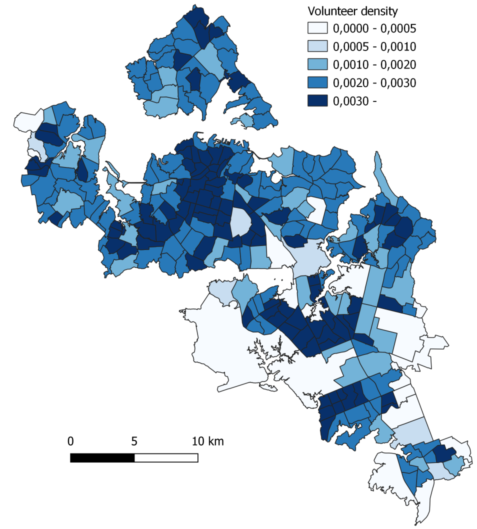

(b) Survival, $n = 500$

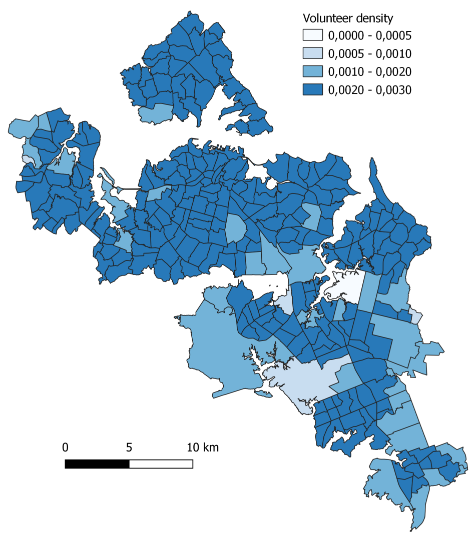

(d) Late arrival, $n = 5000$

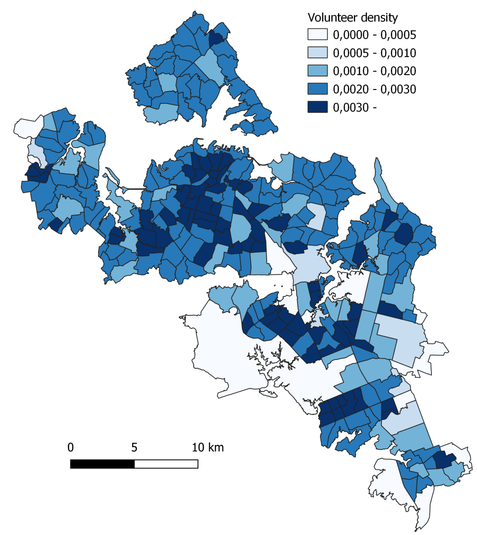

(e) Survival, $n = 5000$

Ambulances + volunteers

$$s(t_{\mathrm{CPR}}, t_{\mathrm{EMS}}) = \left(1 + e^{0.04+0.3t_{\mathrm{CPR}}+0.14(t_{\mathrm{EMS}}-t_{\mathrm{CPR}})}\right)^{-1}$$

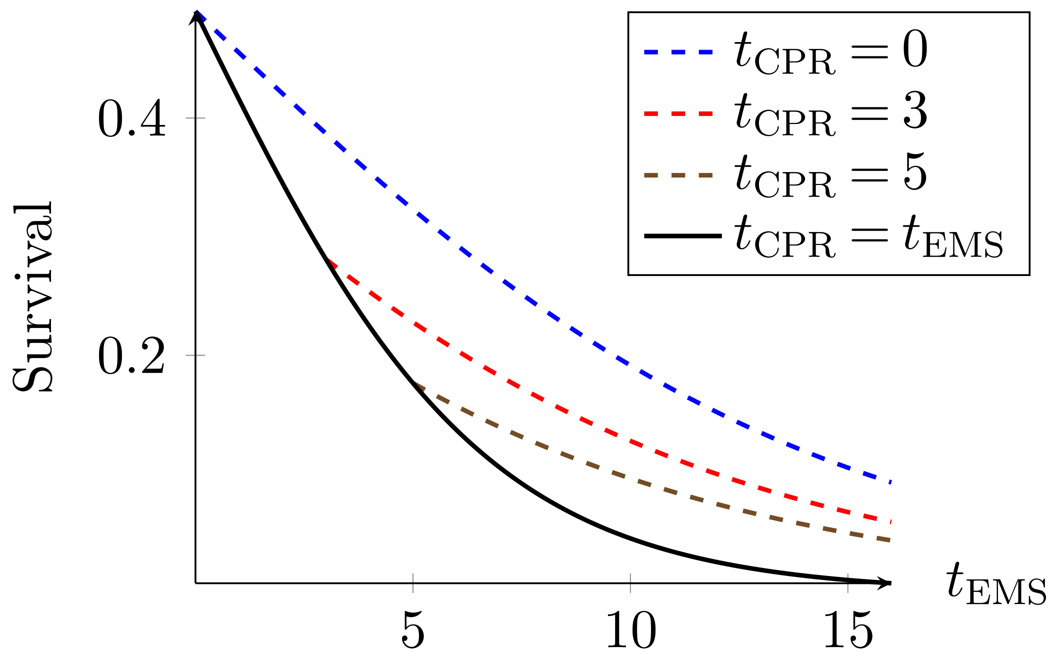

Ambulances + volunteers

Input:

- Ambulance response-time distribution per area unit
- Total number of volunteers in the city (n)

Variables: how the n volunteers are distributed over the area units.

Objective: maximize survival.

→ Still convex!

Auckland: 25 ambulances

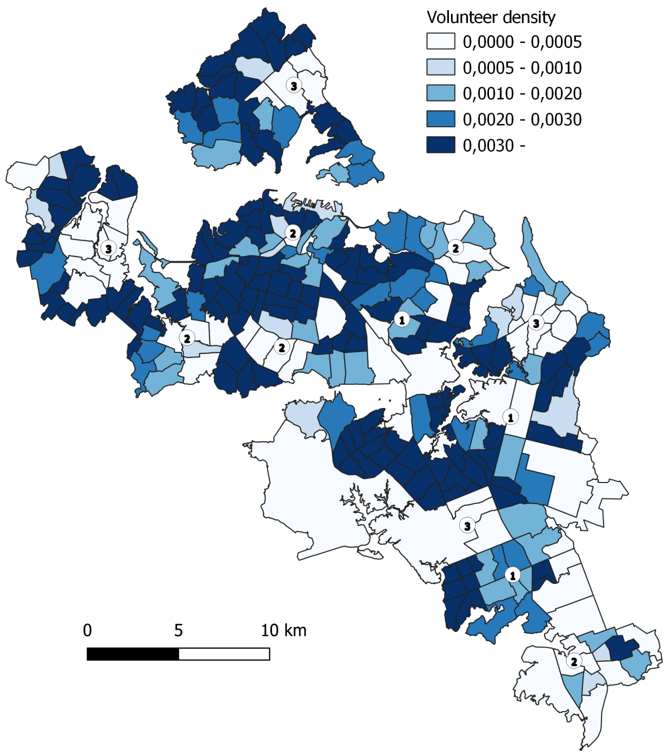

Survival, $n = 500$

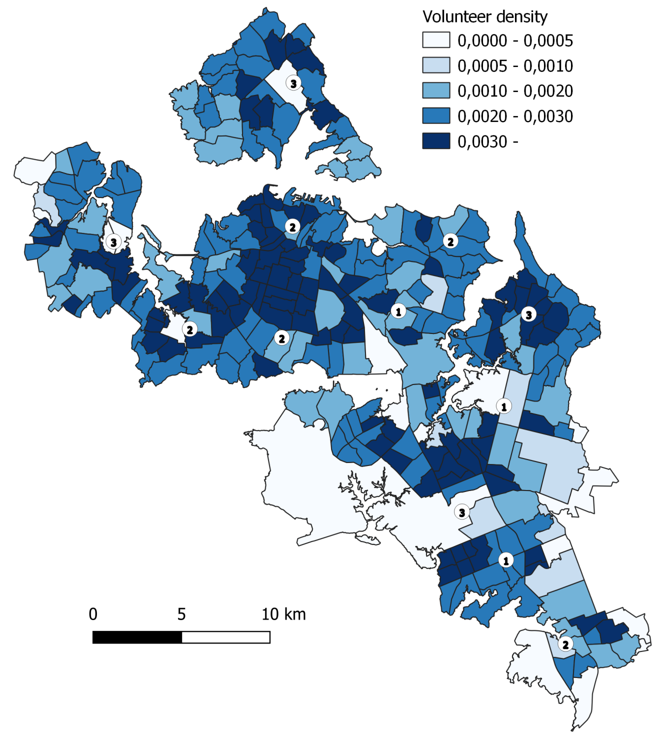

Survival, $n = 5000$

Operational planning

Until now we have discussed *strategic* planning.

Let's talk about what we can optimize
for a CFR system in *real time*.

Phased alerting

<u>Problem definition</u>
Which volunteers should be alerted (and when)?

Given a single patient, observe:

- Response time of the ambulance
- Locations of nearby volunteers

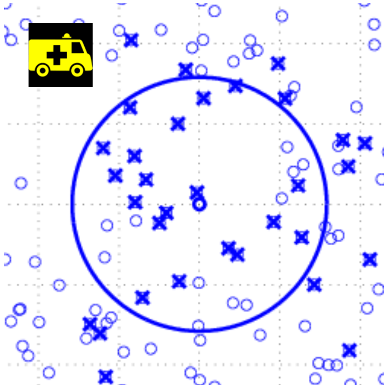

Goal: maximize survival

Avoid multiple volunteers arriving on site

GoodSAM NZ's current dispatch policy

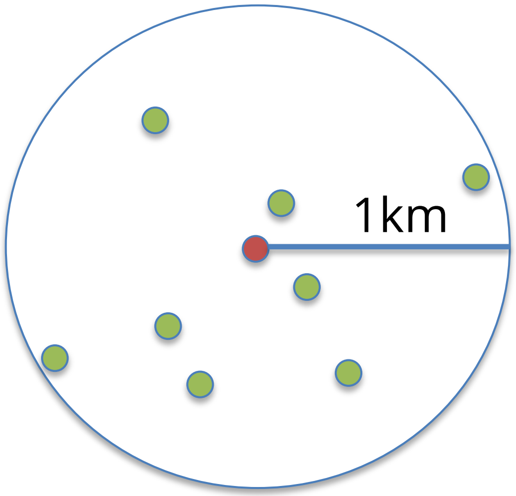

Alert in batches of 3
- ✓ with 30-second time lags
- ✓ until someone has accepted

Never retract alerts

Phased alerting

Response time = min (alert time$_v$ + acceptance delay$_v$ + travel time$_v$)

v in volts

decide stochastic deterministic

Phased alerting

Response time = min (alert time$_v$ + acceptance delay$_v$ + travel time$_v$)

v in accepting volts

Strategy may be adaptive: depending on real-time information (accepts/rejects).
Example: New Zealand’s current policy is adaptive.
But we can imagine even more situation-specific policies.

Volunteer reactions from empirical data

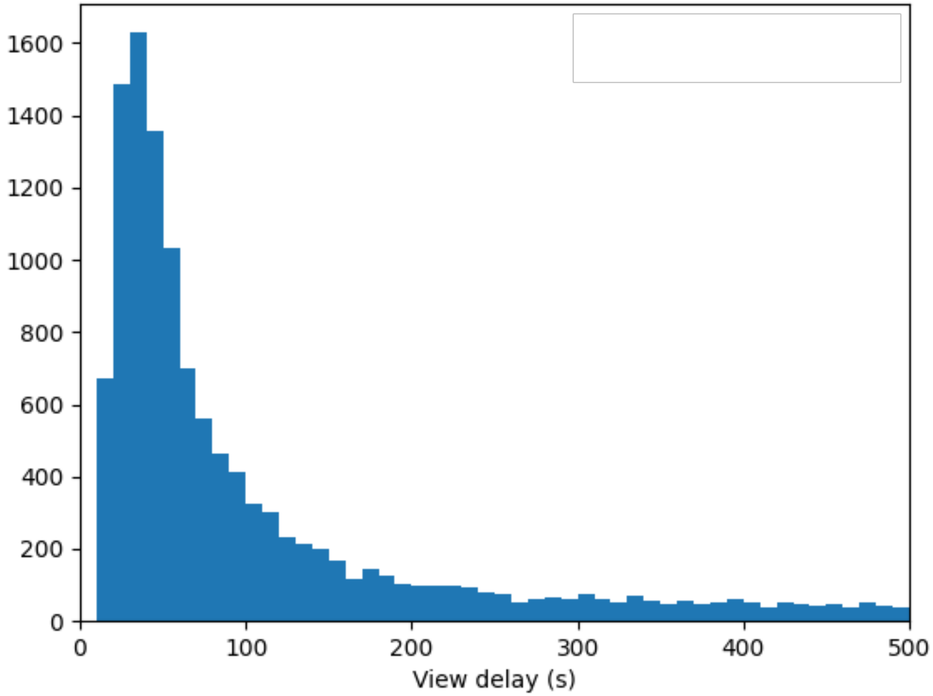

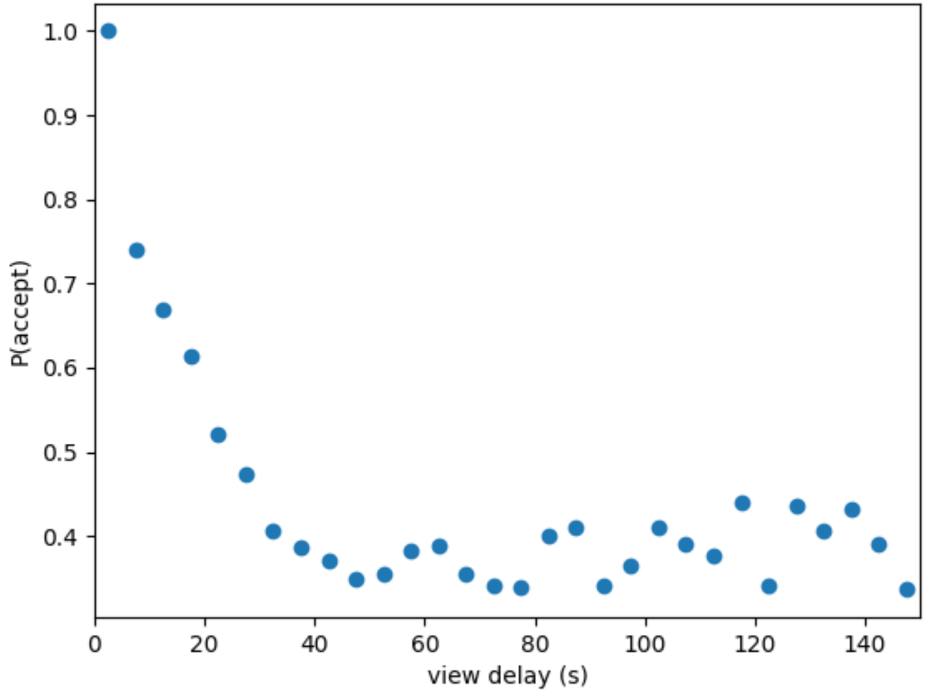

Time between the alert and the reaction (accept/reject).
Based on 12,591 observations from GoodSam NZ.

Monte Carlo Simulation

Compare a number of policies:
- Send all at time 0
- Send n at time 0
- Keep-n-active
- NZ current policy

Generate distances by drawing volunteer locations uniformly at random in a 1-km circle around the patient.

Sample from GoodSAM data: view delays and accept/rejects, distance-dependent travel speed.

Fix EMS time at 12 minutes.

Simulate (often enough to reduce confidence intervals to almost zero).

VU VRIJE UNIVERSITEIT AMSTERDAM
School of Business and Economics

Monte Carlo Simulation

Shows trade-offs between three metrics (for 10 volunteers in the circle):

Policy	**Survivors per year**	**Redundant arrivals**
Send all at time 0.	191	0.918
Keep 6 alerts active.	178	0.453
Keep 7 alerts active.	184	0.587
Keep 8 alerts active.	188	0.718
Keep 9 alerts active.	190	0.834
Keep 10 alerts active.	191	0.918
Send 7 alerts at time 0.	179	0.498
Send 8 alerts at time 0.	184	0.630
Send 9 alerts at time 0.	188	0.771
Send 10 alerts at time 0.	191	0.918
NZ current strategy.	181	0.574

Monte Carlo Simulation

Shows trade-offs between three metrics (for 100 volunteers in the circle):

Policy	**Survivors per year**	**Redundant arrivals**
Send all at time 0.	259	16.768
Keep 7 alerts active.	222	0.595
Keep 8 alerts active.	228	0.743
Keep 9 alerts active.	233	0.897
Keep 10 alerts active.	237	1.056
Send 1 alerts at time 0.	129	0.000
Send 5 alerts at time 0.	196	0.265
Send 10 alerts at time 0.	231	0.918
Send 15 alerts at time 0.	246	1.718
NZ current strategy.	229	1.021

Insight

You can choose your favorite policy from this list and always do that.

But, you may also do something more clever: decide your policy after having observed where the volunteers are.

Machine Learning

Pre-define a list of 40 dispatch strategies, e.g.:

- *Send all*
- *Send n*
- *Keep n active*

For various *n*.

Objective: patient survival – *w* * redundant arrivals

Generate lots of scenarios (e.g. [55, 129, 300, 499, 540 ,588]), simulate all strategies. Store the best strategy.

Example: [55, 129, 300, 499, 540, 588] , #37
[71, 120, 136, 377, 520, 578] , #21
[32, 129, 300, 499, 540, 588] , #4
[85, 190, 298, 360, 387, 440] , #1
[55, 182, 209, 361, 405, 540] , #37

Build a tree that predicts which of the 40 strategies is best, depending on scenario. (Multiclass classification)

VU VRIJE UNIVERSITEIT AMSTERDAM
School of Business and Economics

Small case study

Let's keep it simple: study the process when we have exactly 6 volunteers in the circle.

Step	Runtime
Generate 500 random scenarios	microseconds
For each scenario, evaluate each alerting strategy by simulating it 10.000 times	20 mins
Store the strategy that performed best.	-
Build tree (Python sklearn)	seconds

Small case study

Result:

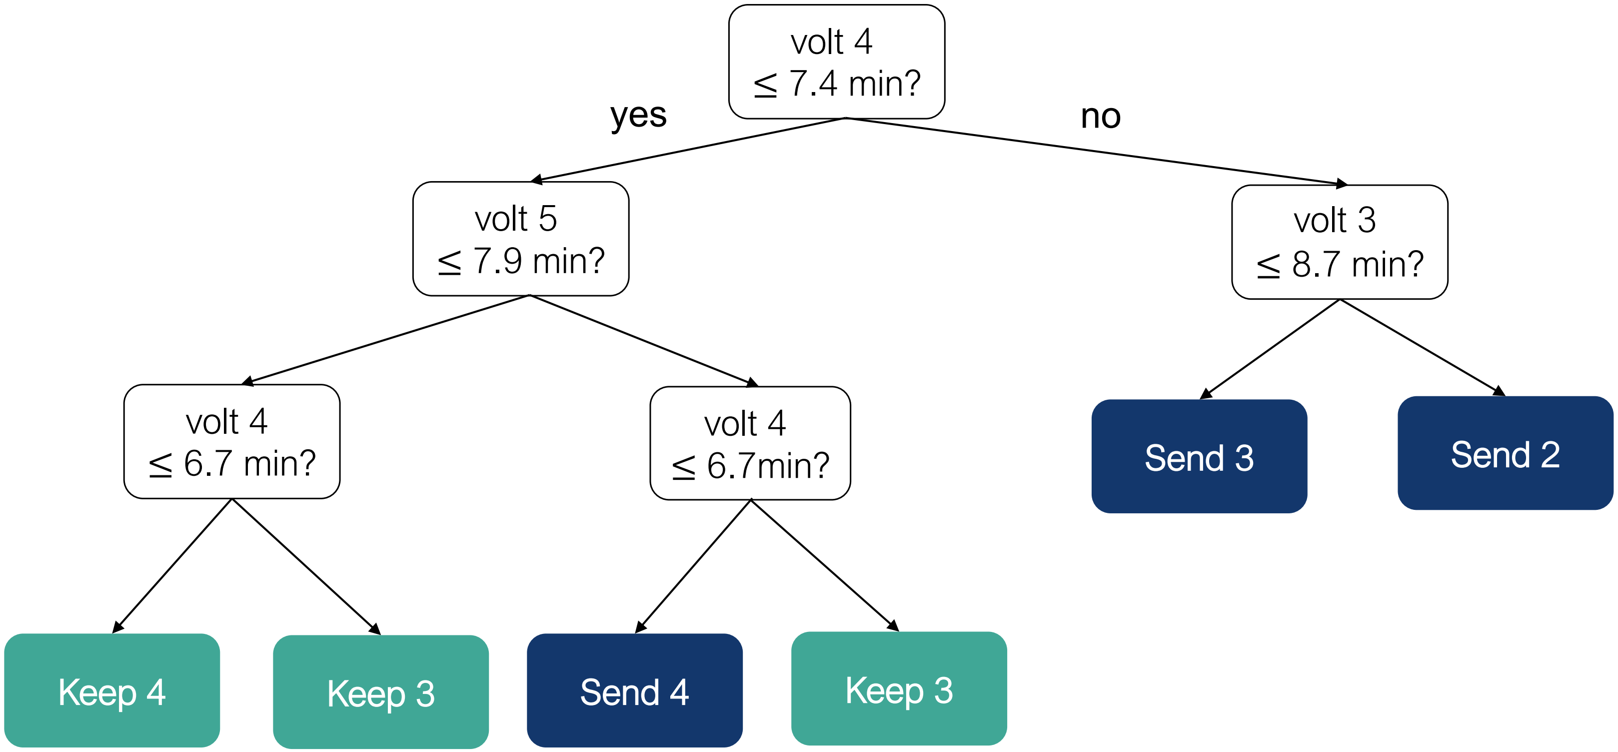

THREE WAYS FORWARD

More realistic input scenarios

Is our finite list of policies covering enough options?

Better trees

Compare against Dynamic Programming

How many volunteers to alert when, is a (stochastic) optimization problem.

Formulate it as a finite-horizon Markov Decision Problem.

5-second time epochs.

To allow a smaller state space, pretend the duration until a volunteer replies yes/no is memoryless.

Actions are how many volunteers to alert (assumed always choose the next-closest one).

Solve by dynamic programming.

Slow: can handle just 5 volunteers

DP versus best-in-the-list

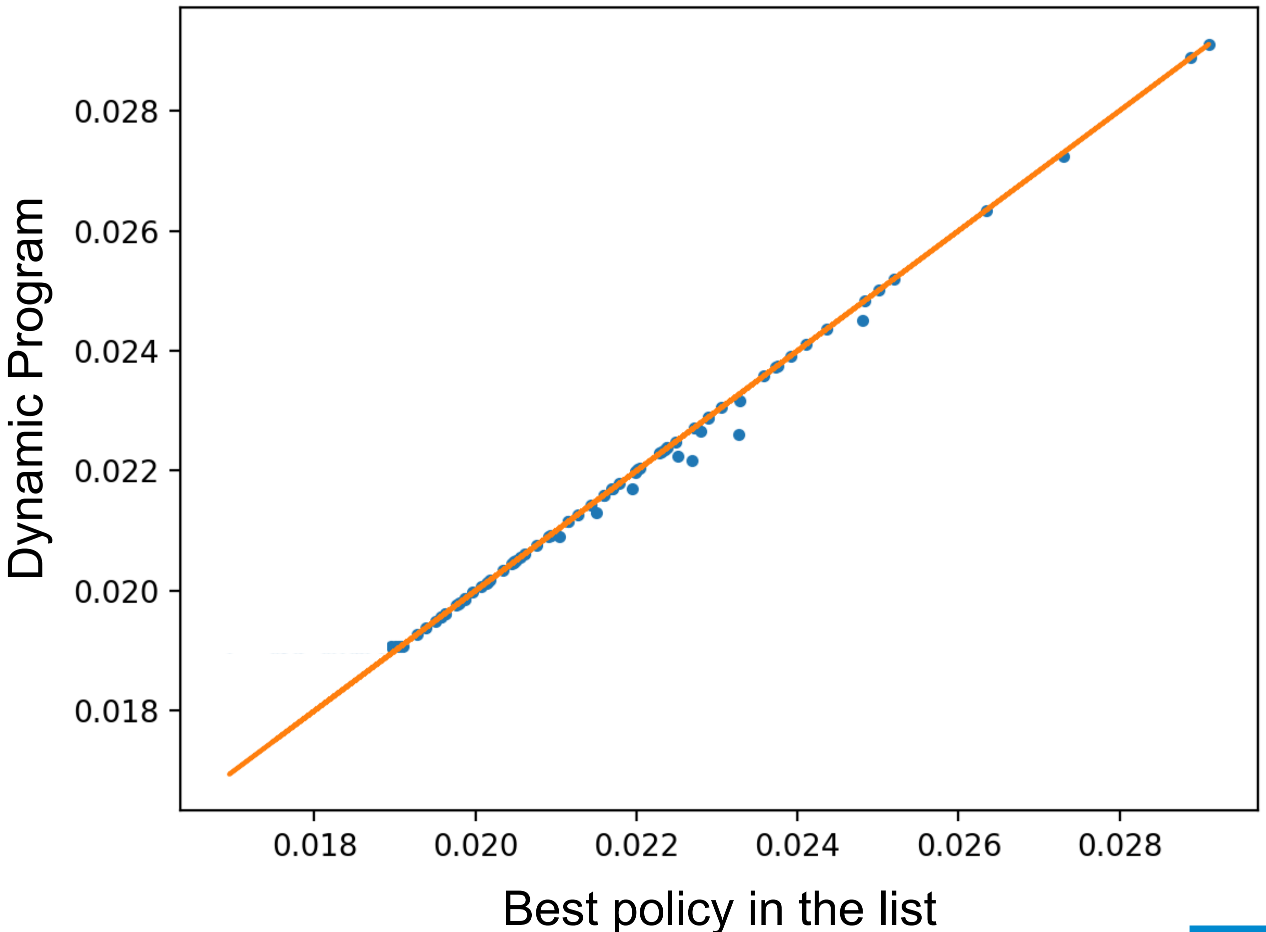

THREE WAYS FORWARD

- More realistic input scenarios
- Is our finite list of policies covering enough options?
- Better trees

More realistic input scenarios

We don’t always see exactly 6 people in the 1-km ball.

How to get a realistic estimate?

More realistic input scenarios

Auckland with a 0.1% signup rate.

Taking into account where the people live.

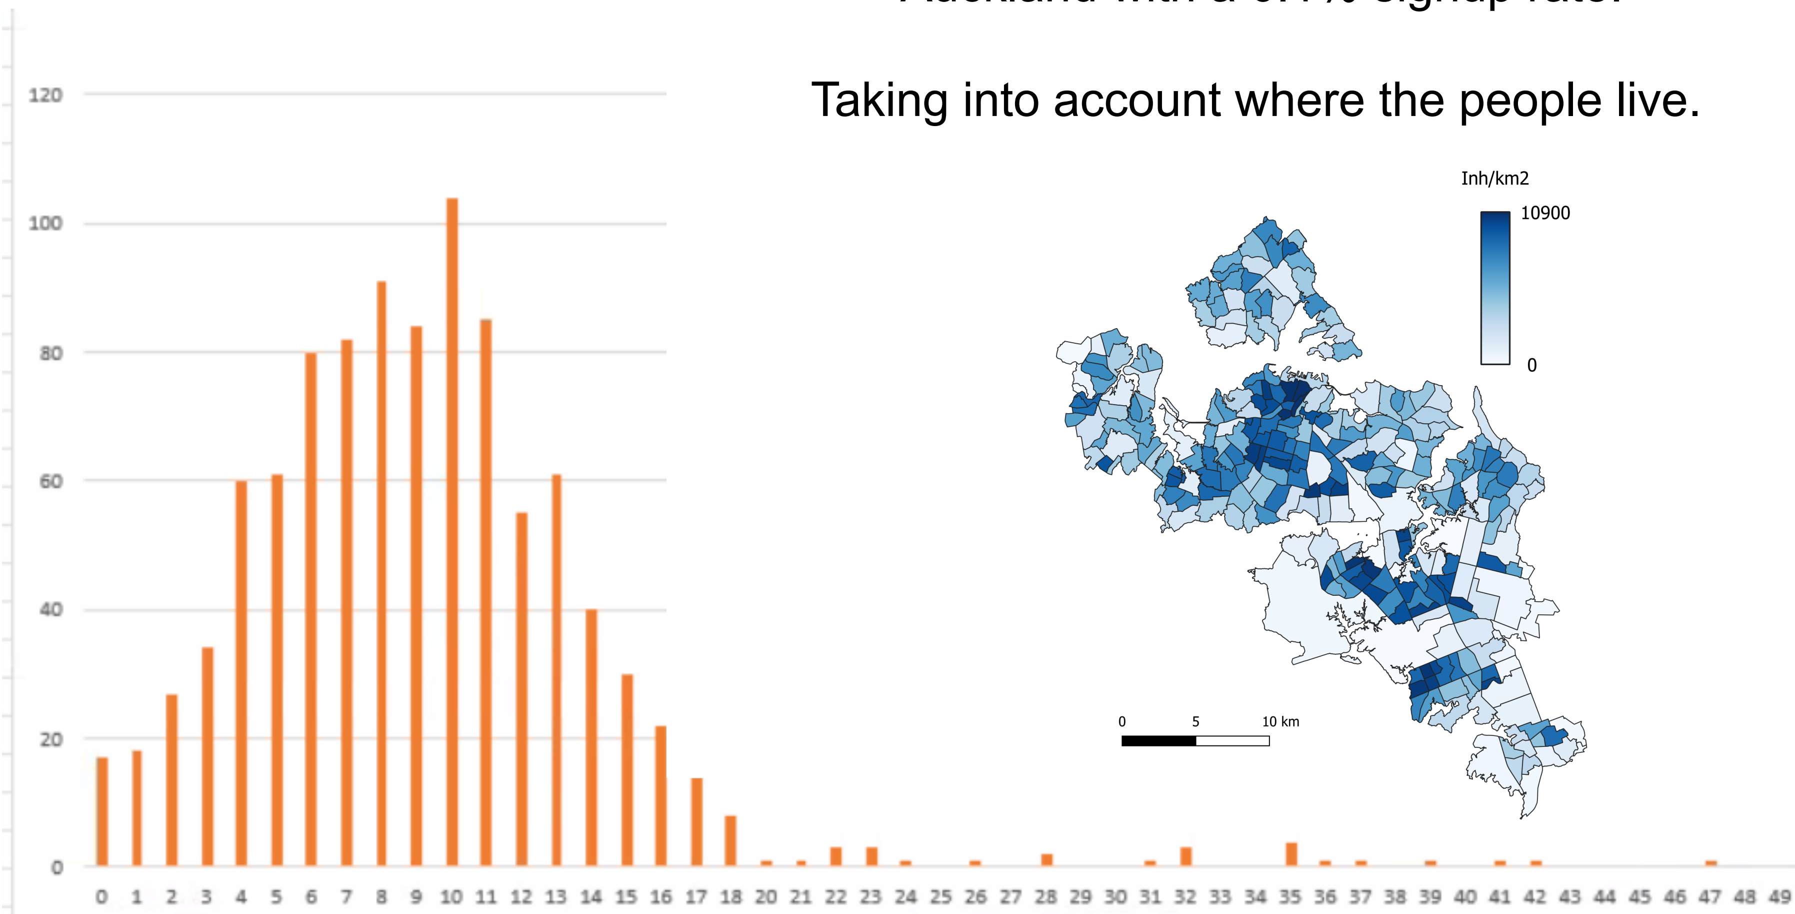

Nr of volunteers in the 1-km ball

Input now looks like this

Example: [55, 129, 300, 499, 540] , #37
[55, 102, 225, 369, 499, 540, 588] , #4
[75, 190] , #1
[55, 187, 300, 477, 545] , #37

Ambulance response times

Obviously, it's not always 12 minutes.

Use realistic estimates that vary across the region.

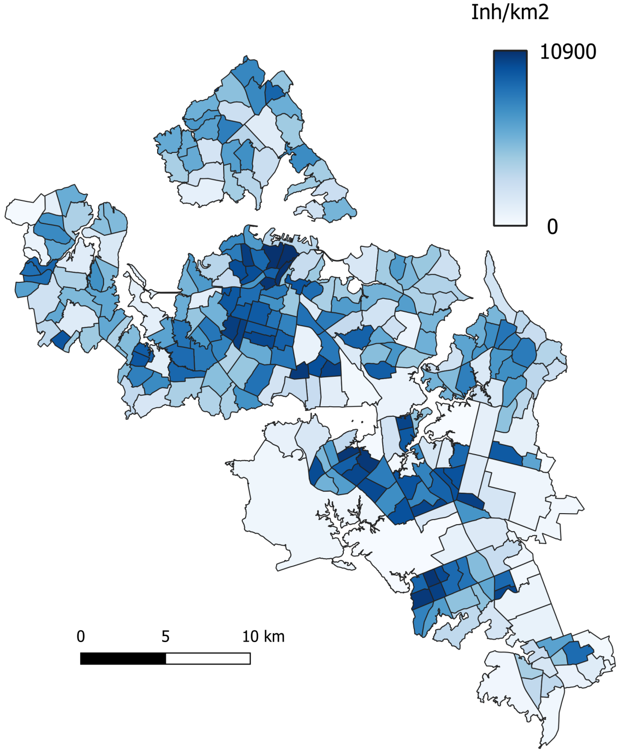

Input now looks like this

Example: [55, 129, 300, 499, 540] , 672, #30
[55, 102, 225, 369, 499, 540, 588] , 590, #4
[75, 190] , 274, #2
[55, 187, 300, 477, 545] , 588, #37

Volunteers EMS time best policy

Triage time affects survival

By the time you send the alert, you know how long you've triaged!

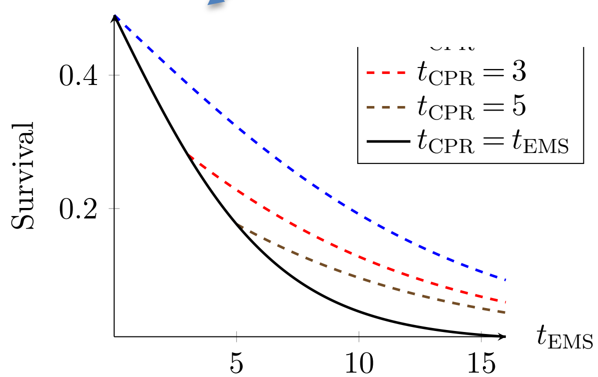

Input now looks like this

Example: [55, 129, 300, 499, 540] , 601, 45, #19
[55, 102, 225, 369, 499, 540, 588] , 590, 77, #24
[75, 190] , 274, 126, #3
[55, 187, 300, 477, 545] , 588, 189, #25

Volunteers EMS time triage time best policy

Result (accuracy 0.83)

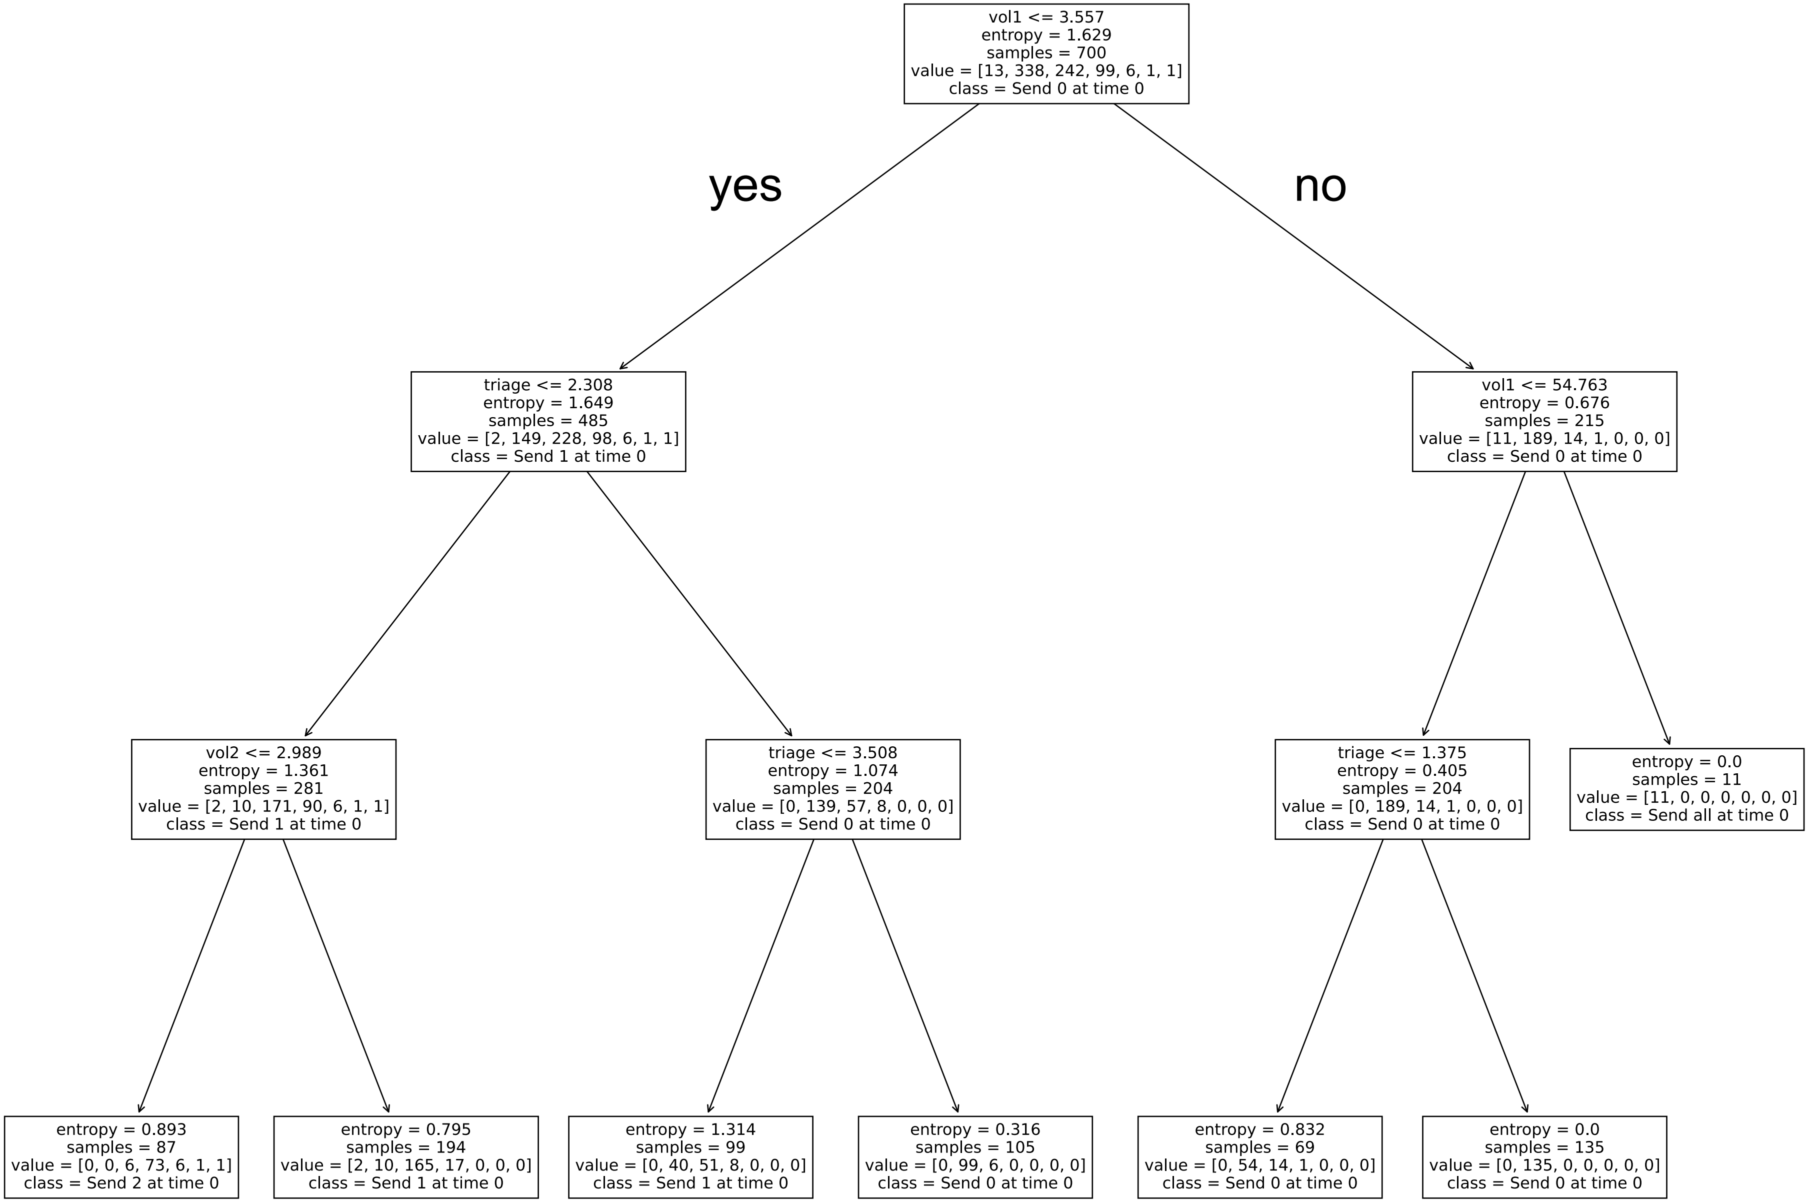

THREE WAYS FORWARD

More realistic input scenarios

Is our finite list of policies covering enough options?

Better trees

Optimal Trees

Who is familiar with *Optimal trees* by Dimitris Bertsimas & Jack Dunn?

"Classification and Regression Trees (CART) build the decision tree using a recursive approach based on a greedy heuristic. We study the benefits of an optimal decision tree approach, which creates the entire decision tree at once using Mixed Integer Optimization".

Benefits:

- A better performing tree (at least true in-sample, hopefully also out-of-sample)
- Allows complex error functions (more than just the % of misclassifications)
- Allows hyperplane splits $4x + 7y - 8z < 17.5$

Optimal trees

- Cool? Yes
- Easy? No

- Not quick to solve, even using commercial MIP solvers
- Tips:
  - Generate a bunch of potentially decent trees using conventional ML packages
  - Use these as warm-starts for the MIP solver
  - Terminate while there is still an optimality gap
  - End with a local search around the best-found solution

VRIJE UNIVERSITEIT AMSTERDAM
School of Business and Economics

Four papers

recruitment

Volunteer recruitment
- Modeling the Impact of Community First Responders

Management Science 2024

alerting

The alerting question
- Alerting in batches with time lags in between?

Queueing Systems 2022

Simulation
- Phased alerting of community first responders for cardiac arrest.

Submitted to Annals of Emergency Medicine

Optimization
- DP & ML

Work in progress

Future work

- Modeling AED pickups
- Or….

- CFR beyond the scope of cardiac arrest

Thank you!

Questions ?