Evaluating influenza mitigation strategies using preventive bandits

Pieter Libin
AI Lab, Vrije Universiteit Brussel

ACAI RL Intro
Nieuwpoort, 6 September 2017

Pandemic influenza

- Influenza
- Seasonal vs pandemic
- High morbidity and mortality
- Problematic mitigation

Epidemiological models

- Model
  - Spread of pathogen(s) through a population
  - Preventive strategies

- Compartment vs individual-based models

Compartment model

- Discrete homogeneous state + transitions

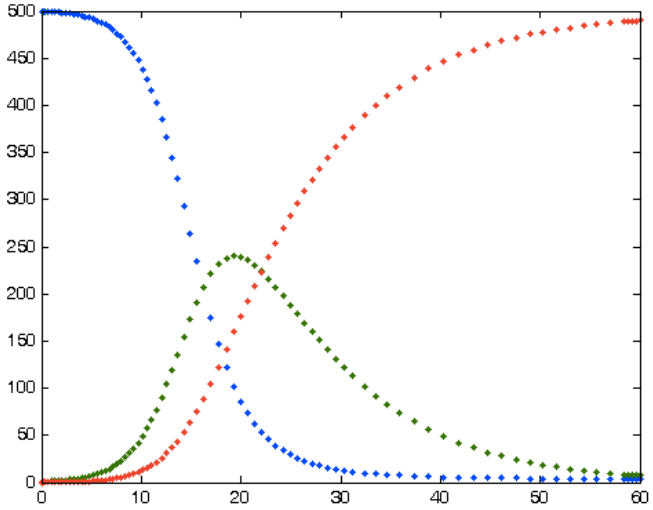

Wikipedia, 2017

Individual-based model

- Individuals + social network

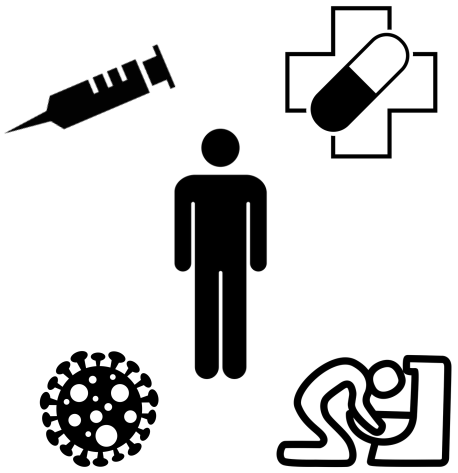

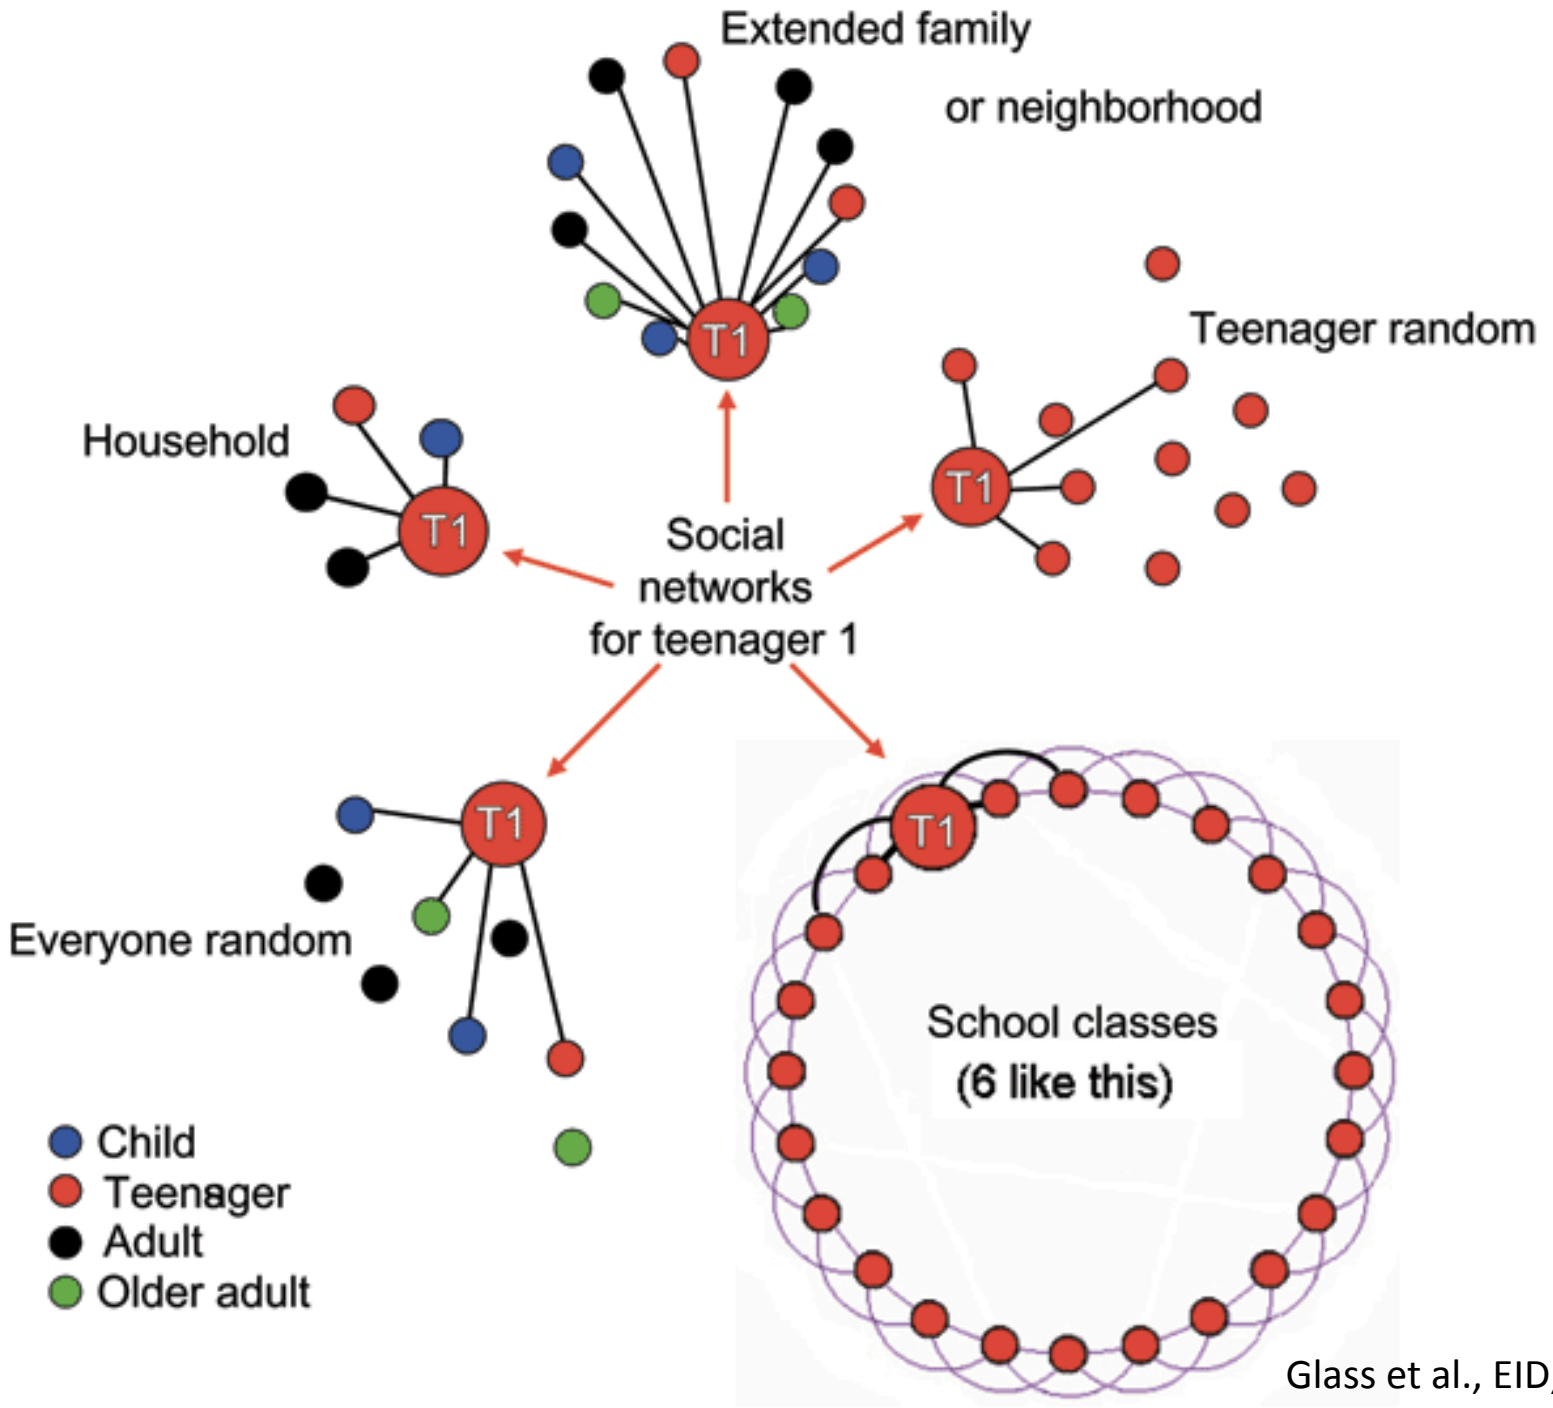

Glass et al., EID, 2006

Modelling influenza

- FluTE individual-based simulator
  - Social network
    - Hierarchy of social mixing groups
    - Worker's commute vs local community
  - Disease progression
  - Interventions

Objective

- Evaluate a set of prevention strategies
  - In an epidemiological model
  - As little simulations as possible

Multi-armed bandit

- Slot machine with *k* arms
- Arm associated with reward

- Play sequence of arms to maximize reward
- Exploit/Explore

Preventive bandit

- Epidemiological model:
  - outcome ~ E(c,p)
- Preventive bandit
  - Set of preventive strategies $P=\{p_1, \ldots, p_n\}$
  - Bandit with $|P|$ arms
  - Playing arm i ~ Evaluating $E(c,p_i)$
- Identify the optimal prevention strategy:
  - ε-greedy/UCB1

Experiment

- Context of pandemic influenza

- Limited vaccine available (Medlock, 2008)
  - Prioritize vaccine distribution

0-4	5-18	19-29	30-64	>65

Experiment

- To find best distribution policy:
  - Consider 2^5 distribution strategies
  - Evaluate using ε-greedy/UCB1

Experiment: FluTE configuration

- Base configuration
  - 2000 individuals
  - 10 seeds
  - 180 days
  - no pre-existing immunity
  - 100 doses of vaccine
  - 1 type of vaccine
- Experiment-specific configuration
  - **R_0 = {1.3, 1.4}**

FluTE preventive bandit

- Play arm:
  - Invoke FluTE
  - Extract the proportion of symptomatic infections
  - Return the reward based on this proportion

Outcome distribution

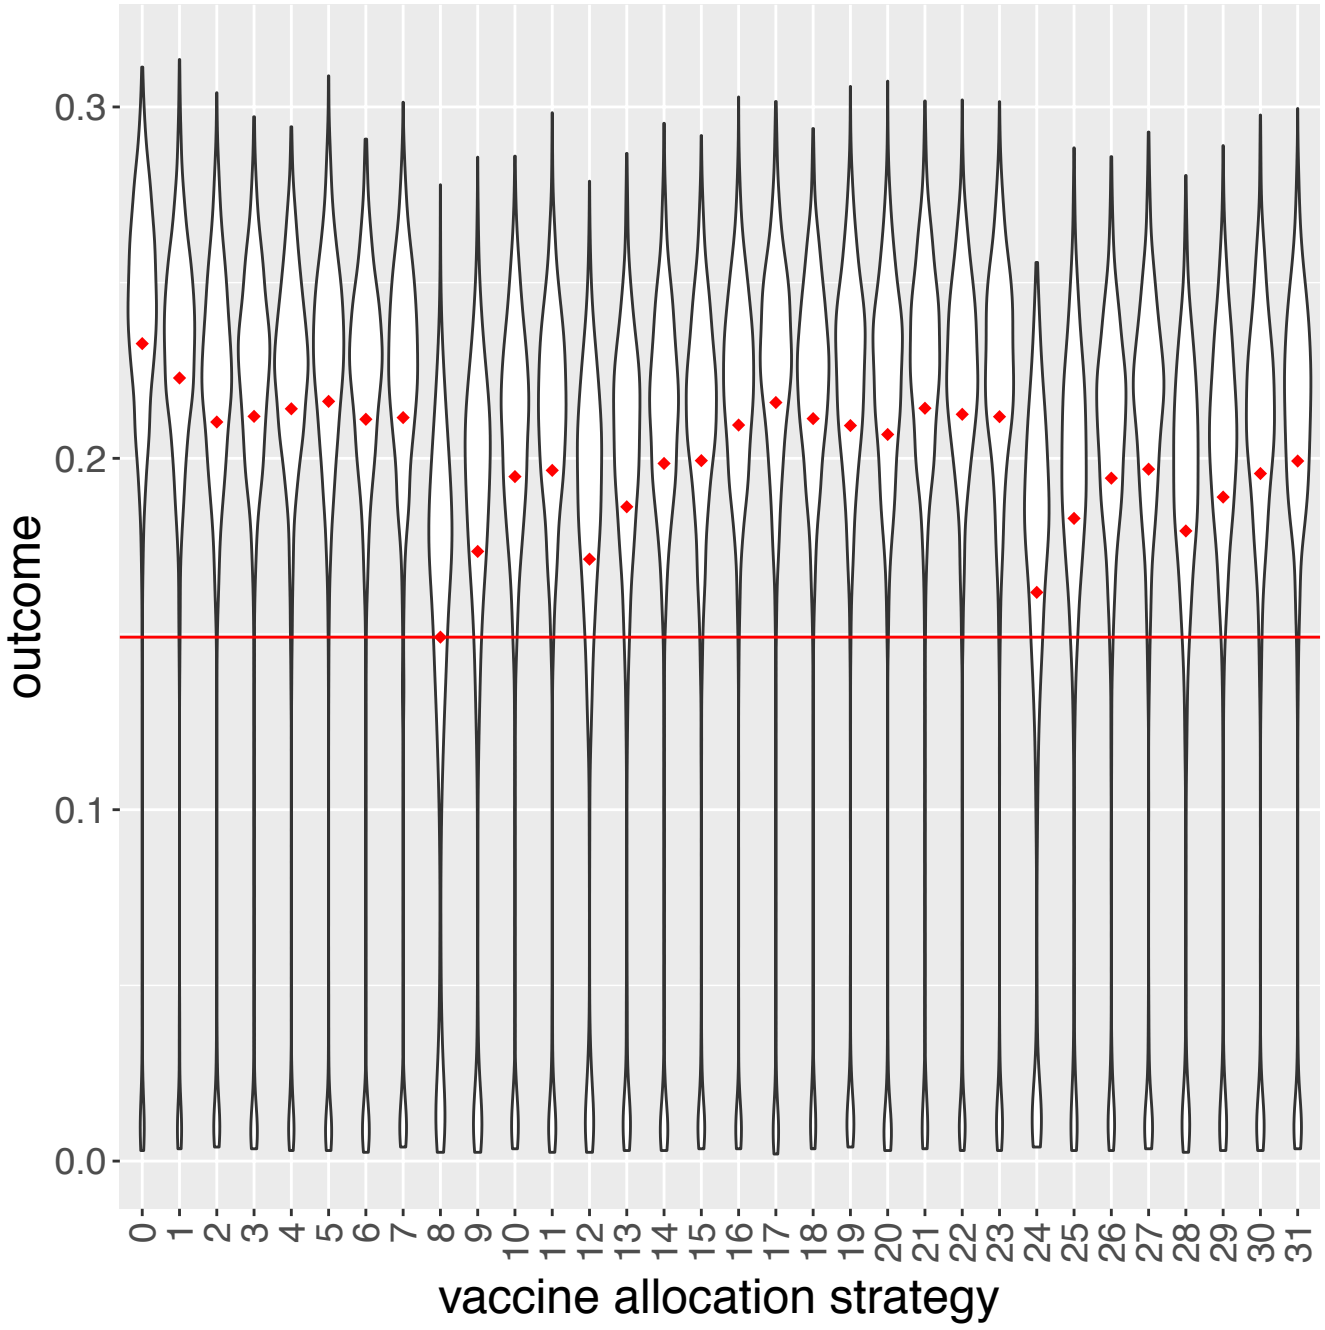

$R_0 = 1.3$

Results

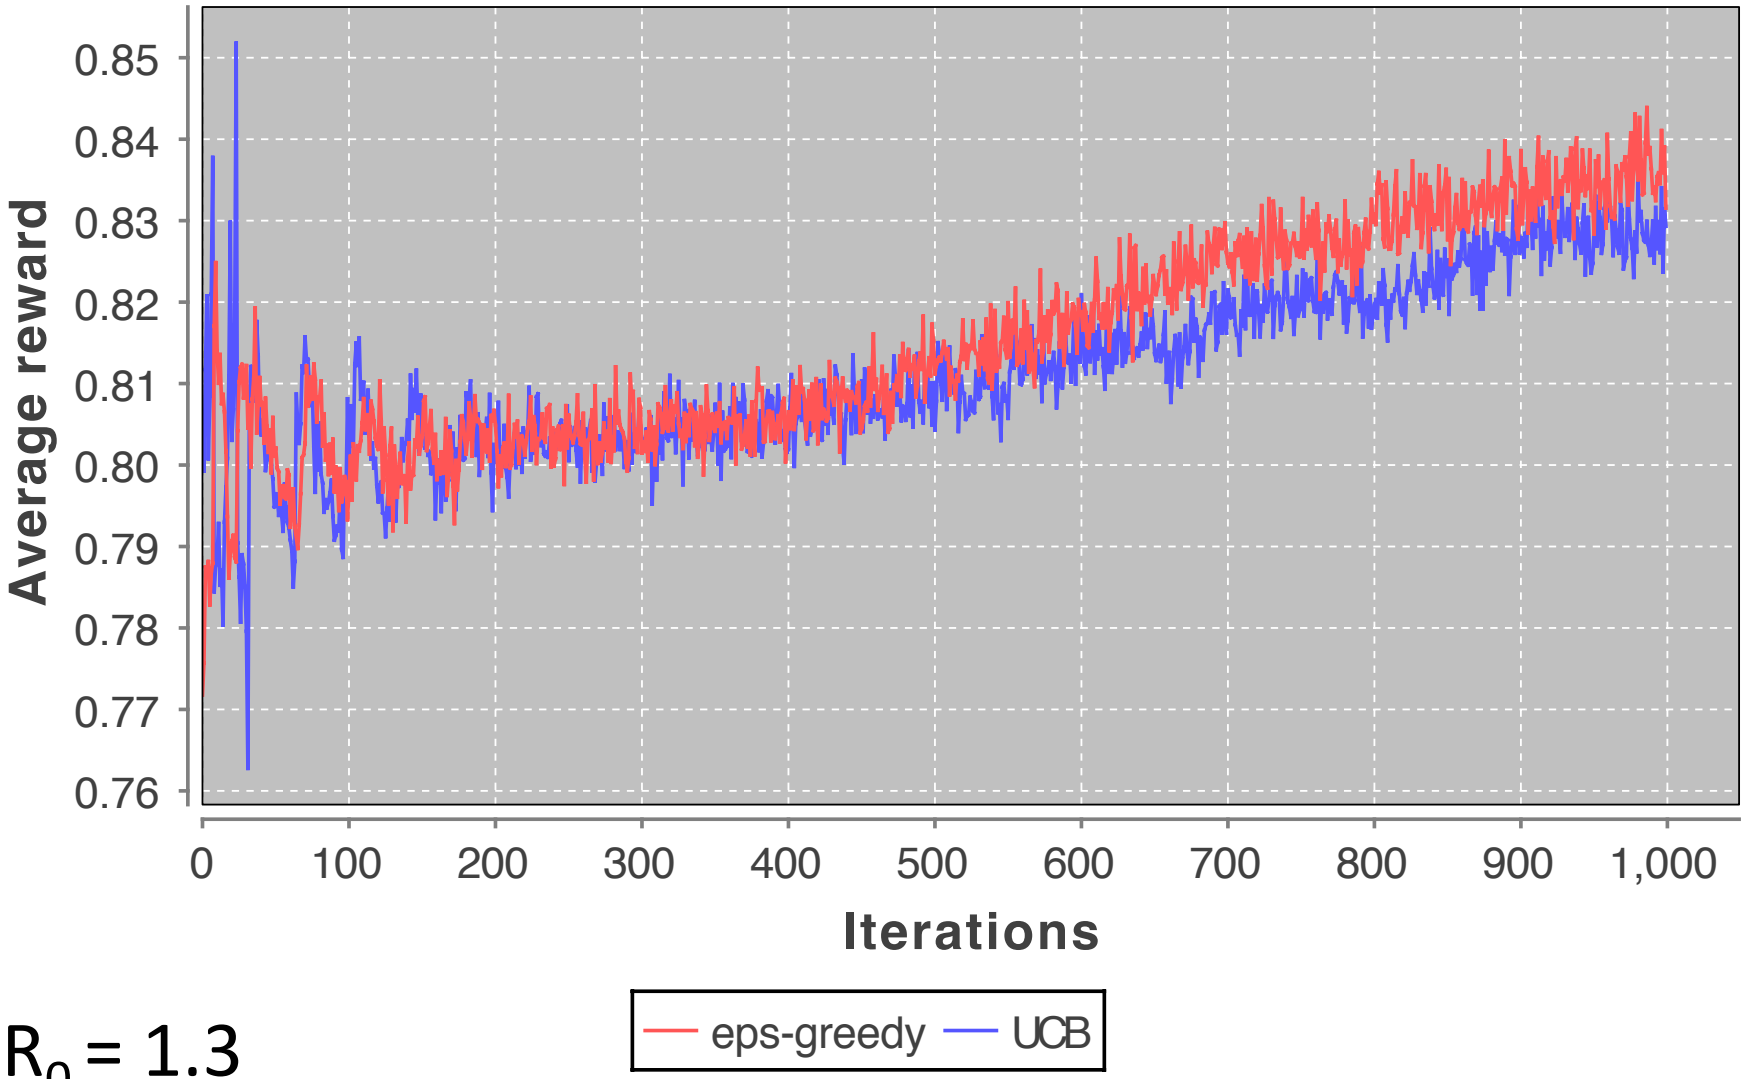

$R_0 = 1.3$

Results

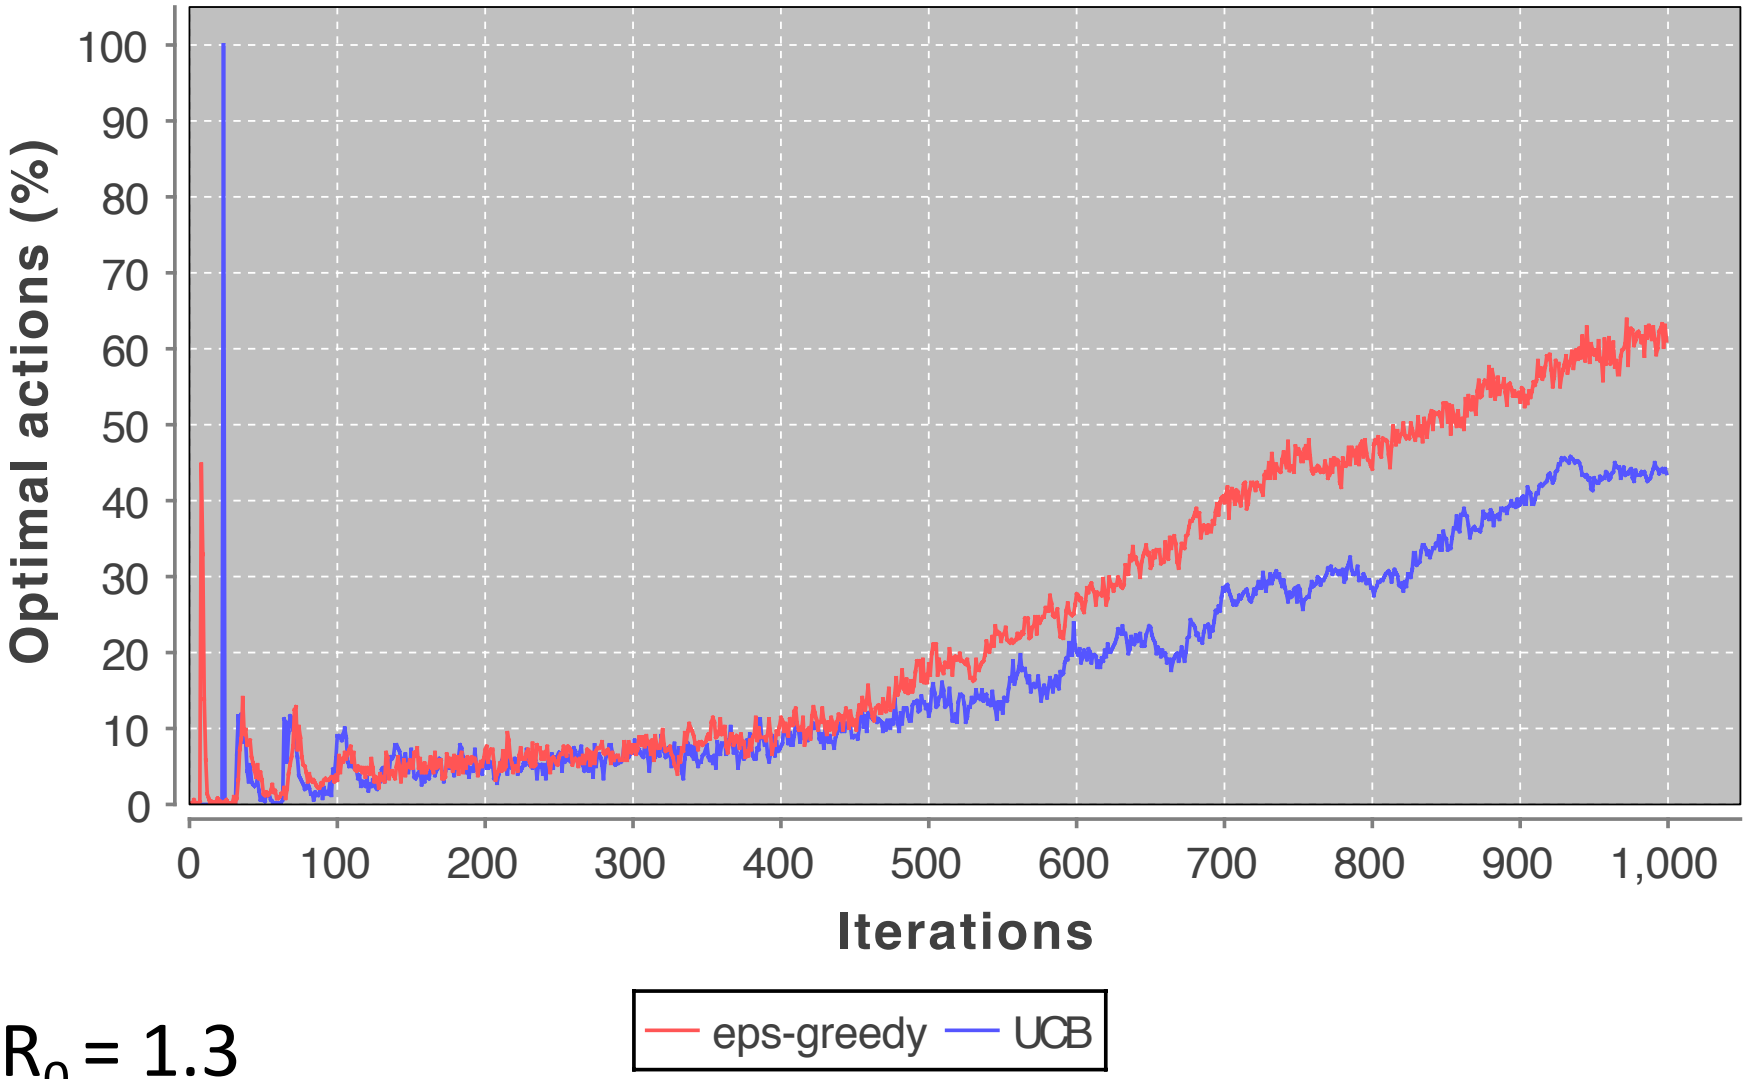

$R_0 = 1.3$

Discussion

- Formalisms and methods are generic
  - https://github.com/vub-ai-lab/FluTE-bandits
- Multi-objective bandits
- Stateful reinforcement learning
- Finds optimal strategy

In collaboration with

Timothy Verstraeten

Diederik Roijers

Peter Vrancx

Ann Nowé

Kristof Theys

Philippe Lemey

Credits

- thenounproject.com

  Adriano Emerick 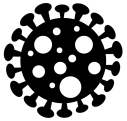

  Wilson Joseph 

  Viktor Fedyuk 

  Edward Boatman 

- References

Medlock, Jan, and Alison P. Galvani. "Optimizing influenza vaccine distribution." Science 325.5948 (2009): 1705-1708.

Glass, Robert J., et al. "Targeted social distancing design for pandemic influenza." Emerg Infect Dis 12.11 (2006): 1671-1681.