AI for Social Good: Key Techniques, Applications, and Results

MILIND TAMBE

Founding Co-director, Center for Artificial Intelligence in Society (CAIS)

University of Southern California

tambe@usc.edu

Co-Founder, Avata Intelligence

USC Center for Artificial Intelligence in Society (CAIS.USC.EDU)

Mission Statement: Advancing AI research driven by…

American Academy of Social Work & Social Welfare

Grand Challenges of Social Work

SUSTAINABLE DEVELOPMENT GOALS

17 GOALS TO TRANSFORM OUR WORLD

- Ensure healthy development for all youth
- Close the health gap
- Stop family violence
- Advance long and productive lives
- Eradicate social isolation
- End homelessness

…

Overview of CAIS Project Areas

AI for Assisting Low Resource Communities

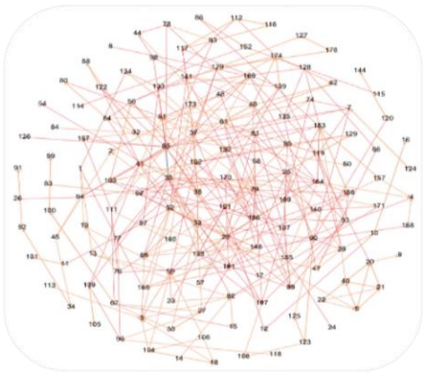

- Social networks: Spread HIV information
- Maximize influence under uncertainty
- Real-world pilot tests: Big improvements

AI for Protecting Endangered Wildlife

- Machine learning/planning: Anti poaching
- Scale, boundedly rational poachers,…
- Real-world: Uganda, South Asia…

Overview of CAIS Project Areas

AI for Public Safety and Security

- Game theory: security optimization
- Solve massive “security games”
- Real-world: US Coast Guard, FAMS…

Gangs, Substance abuse, Veterans mental health

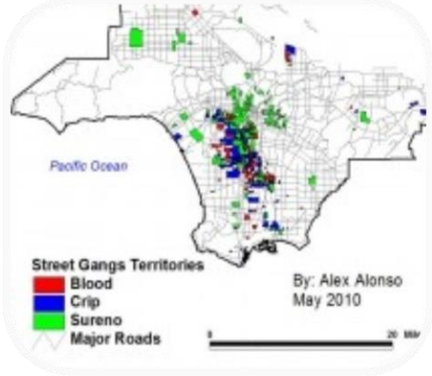

- Social networks, robust optimization,…
- Behavioral models…
- Real-world: Research in progress

Outline

- Introduction
- HIV Information among homeless youth
- Wildlife Conservation

AI Program: HEALER

Outline: HIV Information & Homeless Youth

- Domain of homeless youth and HIV information dissemination
- Real World Challenges in Influence Maximization
- POMDP Model and algorithms
- Pilot Study

Influence Maximization Background

- Input:
  - Graph G
  - Influence Model I
  - Choose K nodes per time step
  - Number of time steps for influence spread T

- Output:
  - K nodes per time step maximizing expected # influenced nodes

Independent Cascade Model

$$G = (V, E)$$

- Propagation Probability (for each edge)

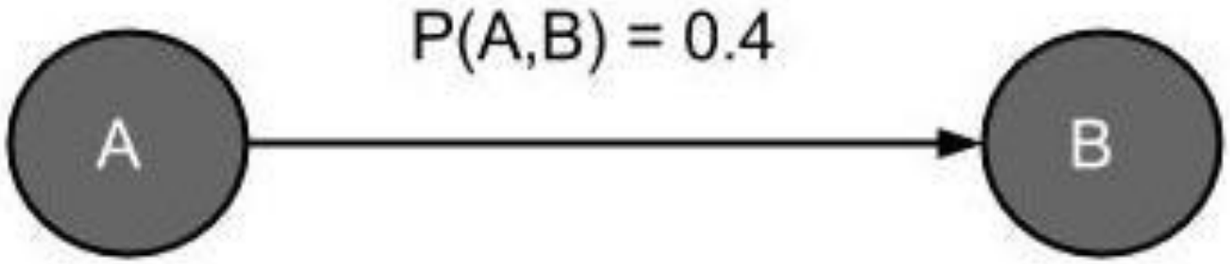

Real World Challenges

- Uncertain network state
- Uncertainty in network structure
- Adaptive selection

Challenge 1: Uncertain Network State

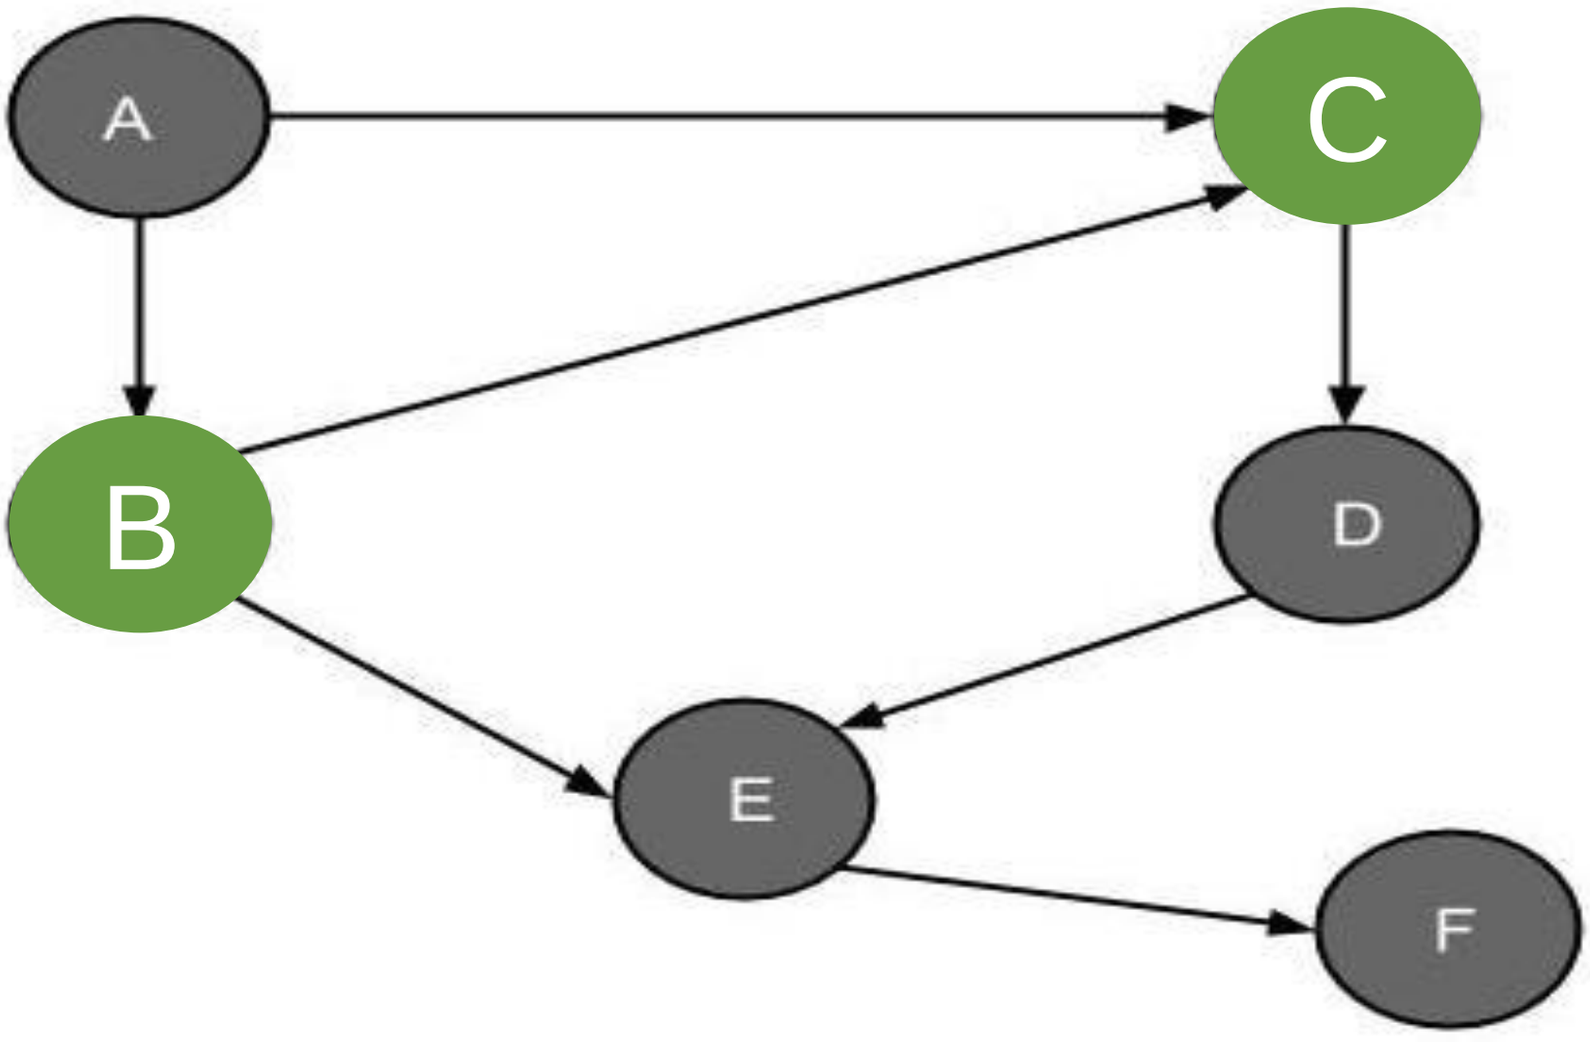

Challenge 2: Uncertain Network Structure

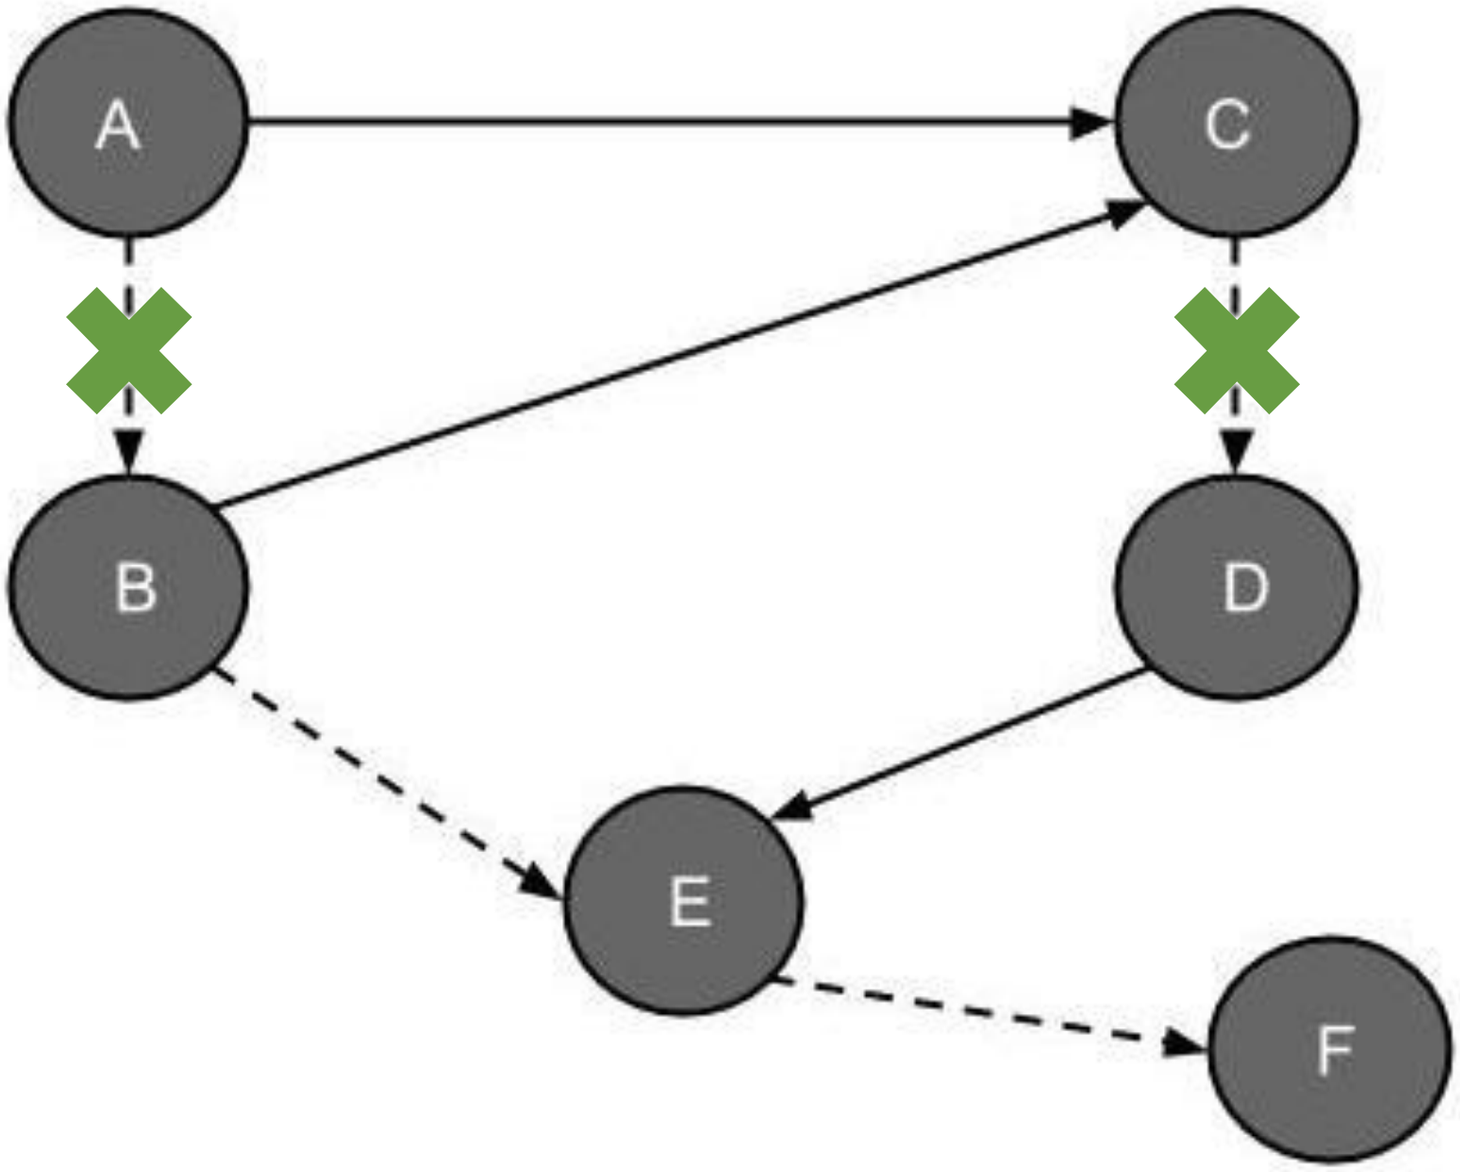

Independent Cascade Model

$$G = (V, E) \qquad E = E_{cert} \cup E_{uncert}$$

- Propagation Probability (for each edge)

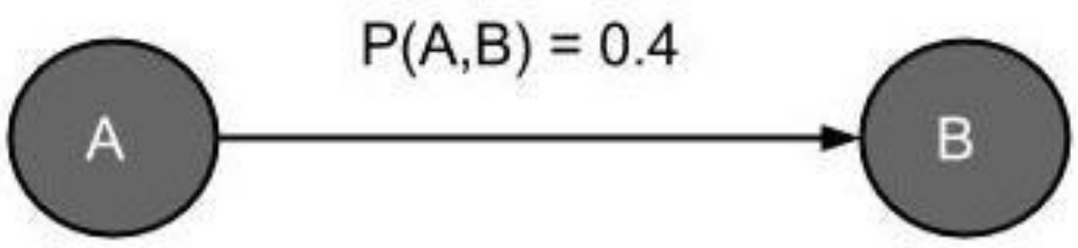

- Existence Probability (for uncertain edges only)

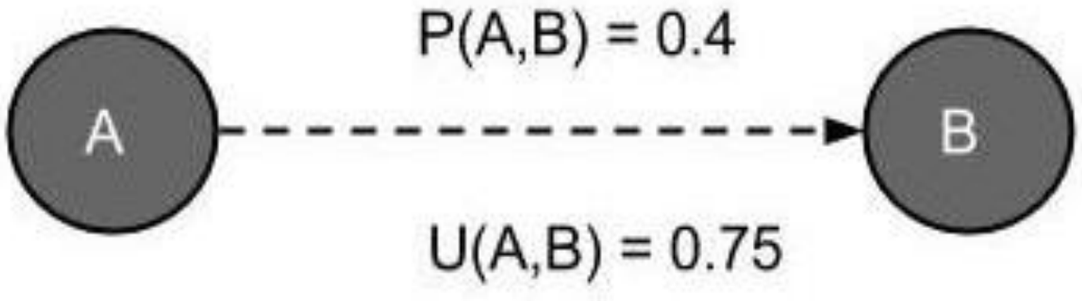

HIV Prevention Programs:
Using Social Networks to Spread HIV Information

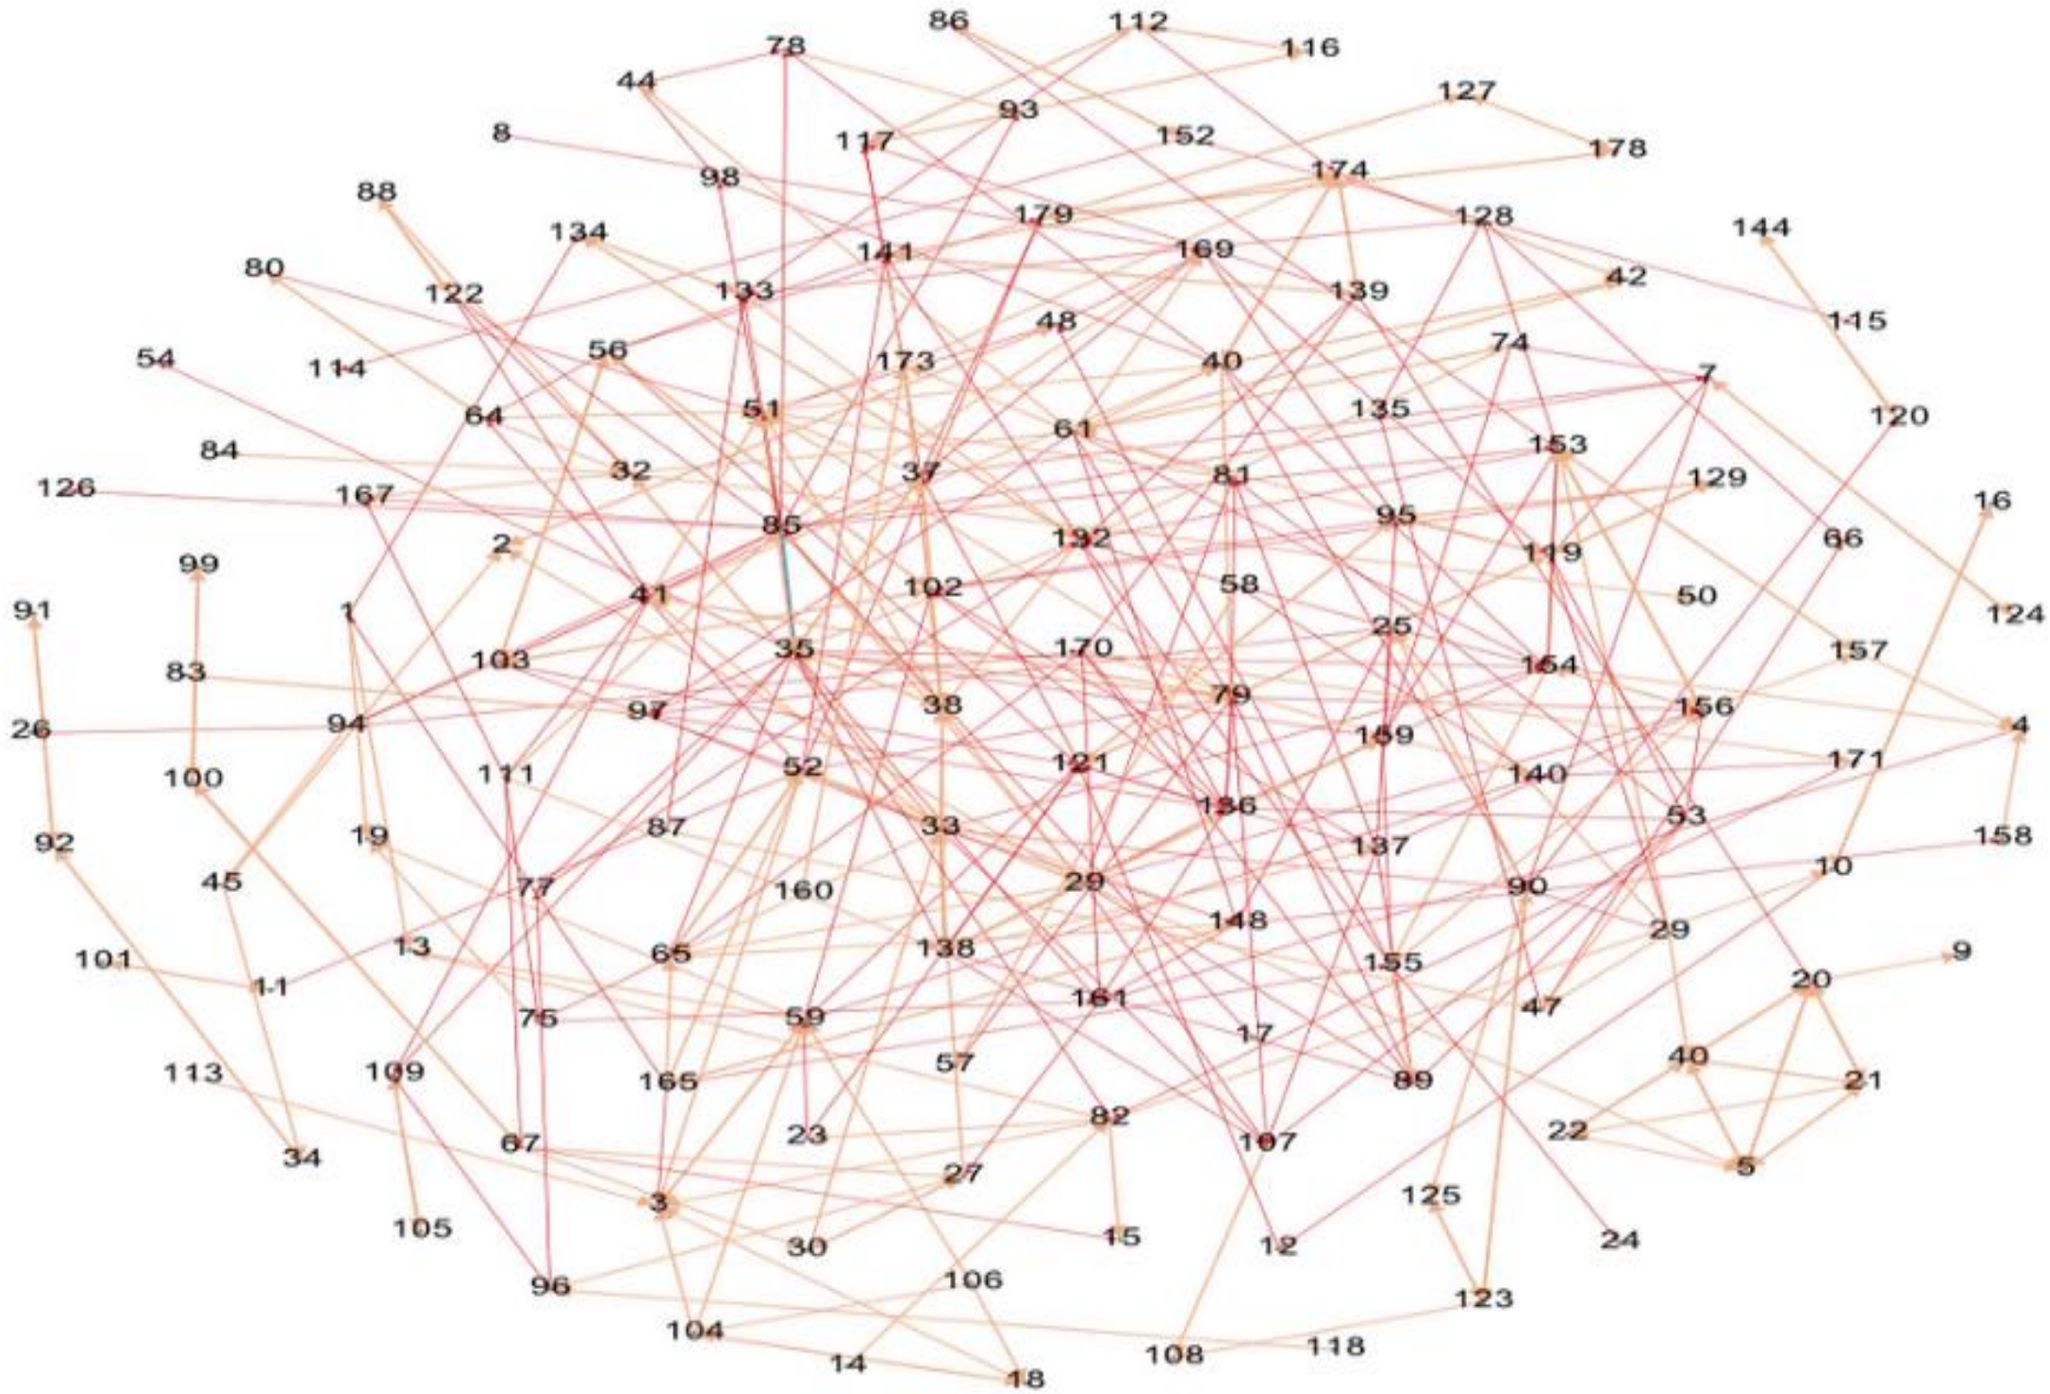

Challenge: Adaptive selection in Uncertain Network

K = 5
1st time step

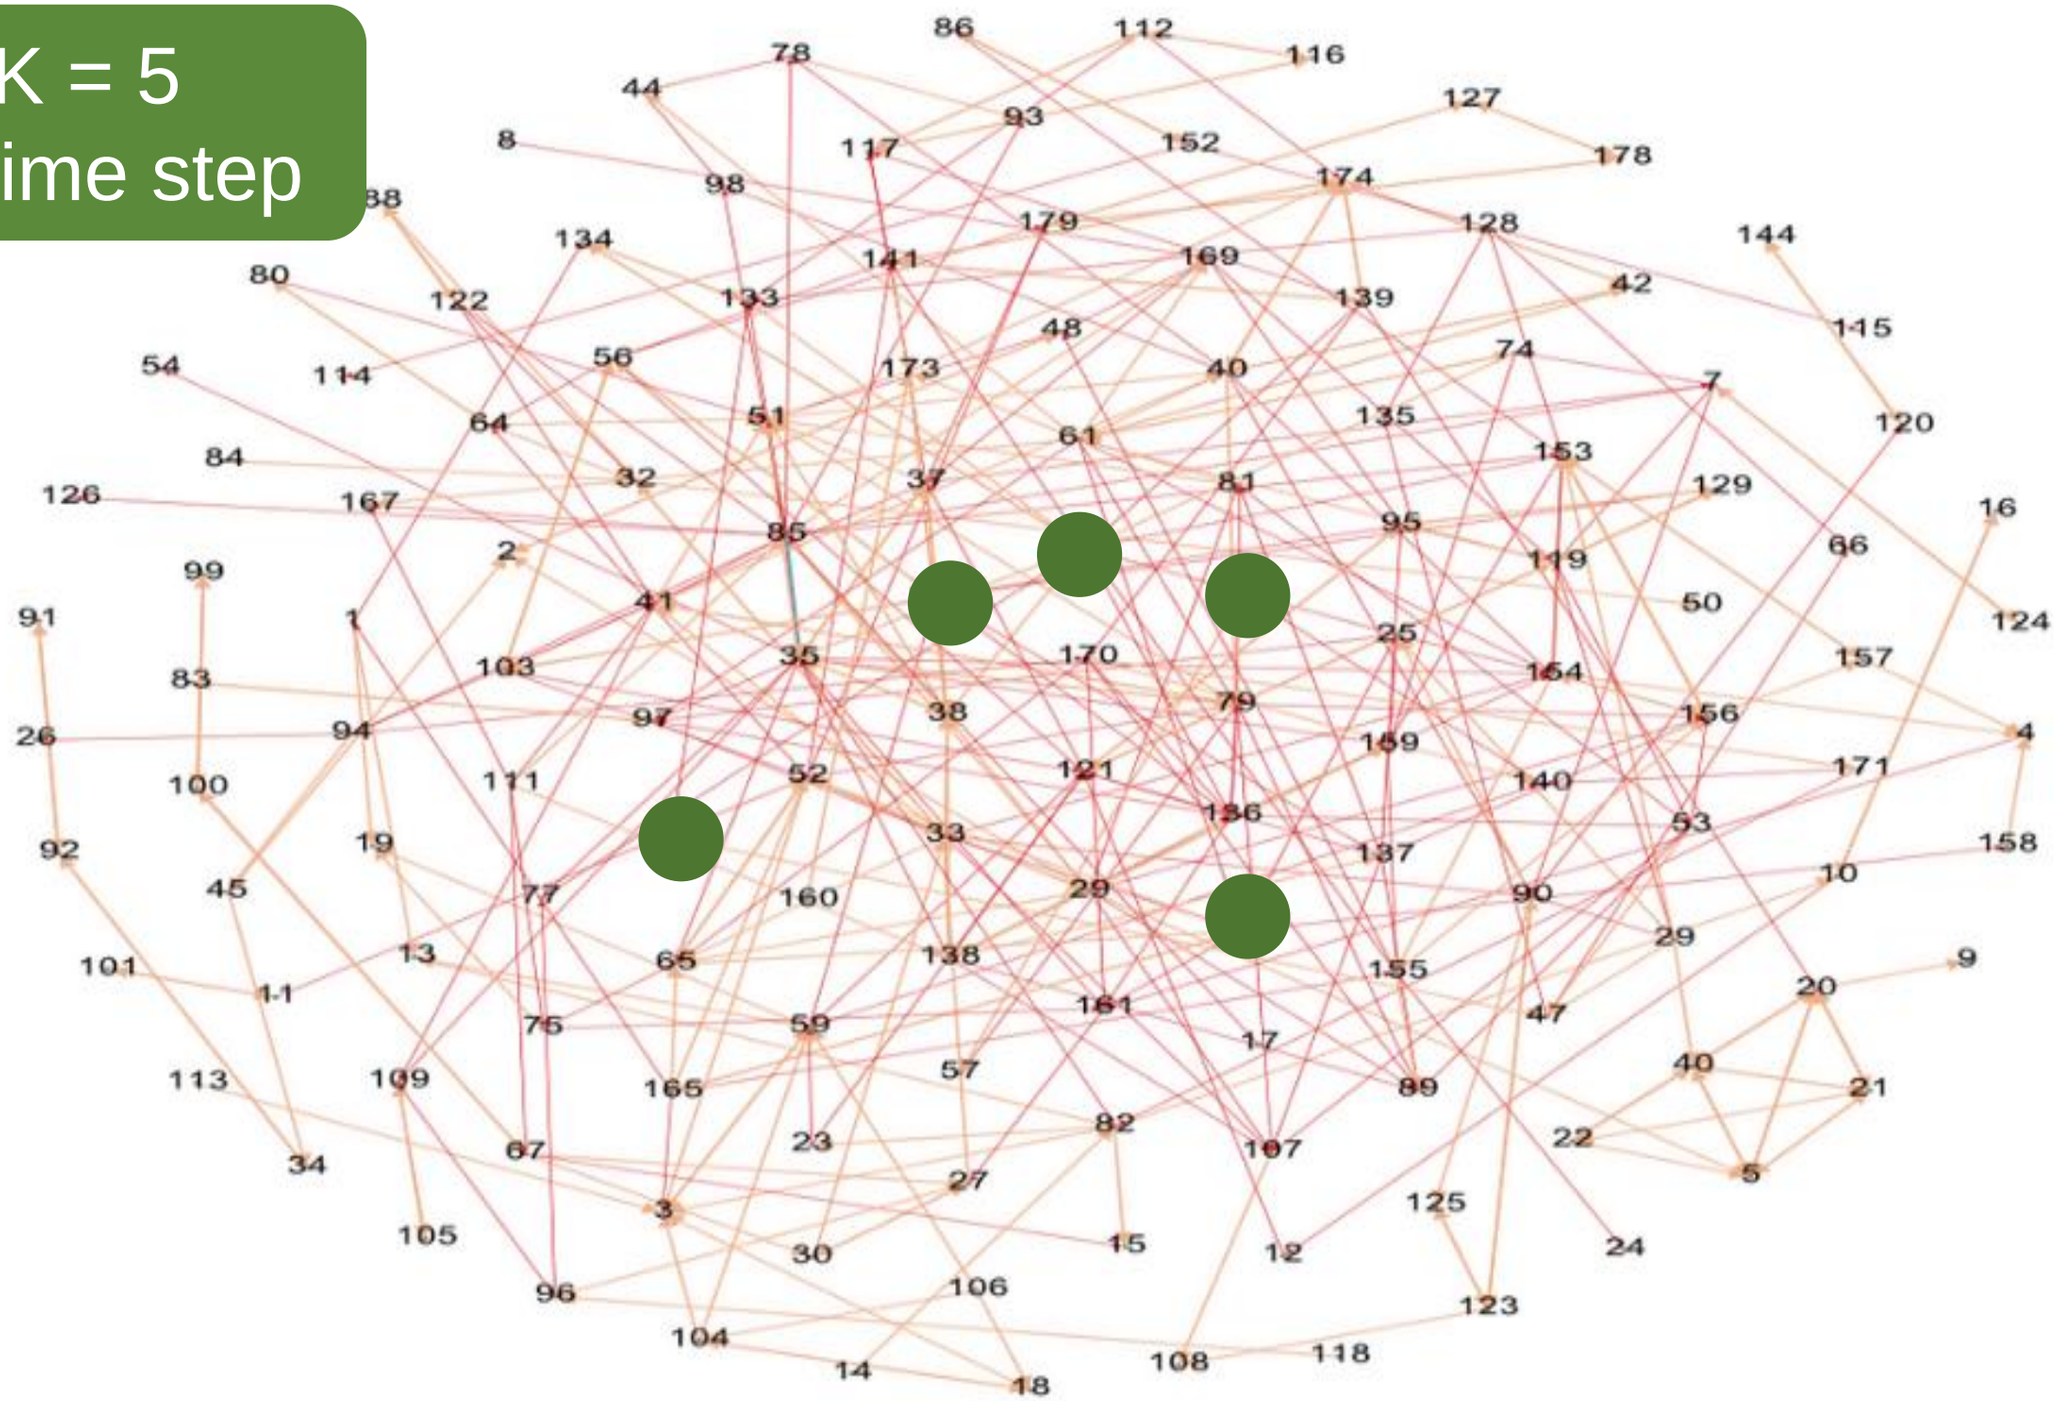

Challenge: Adaptive selection in Uncertain Network

K = 5
2nd time step

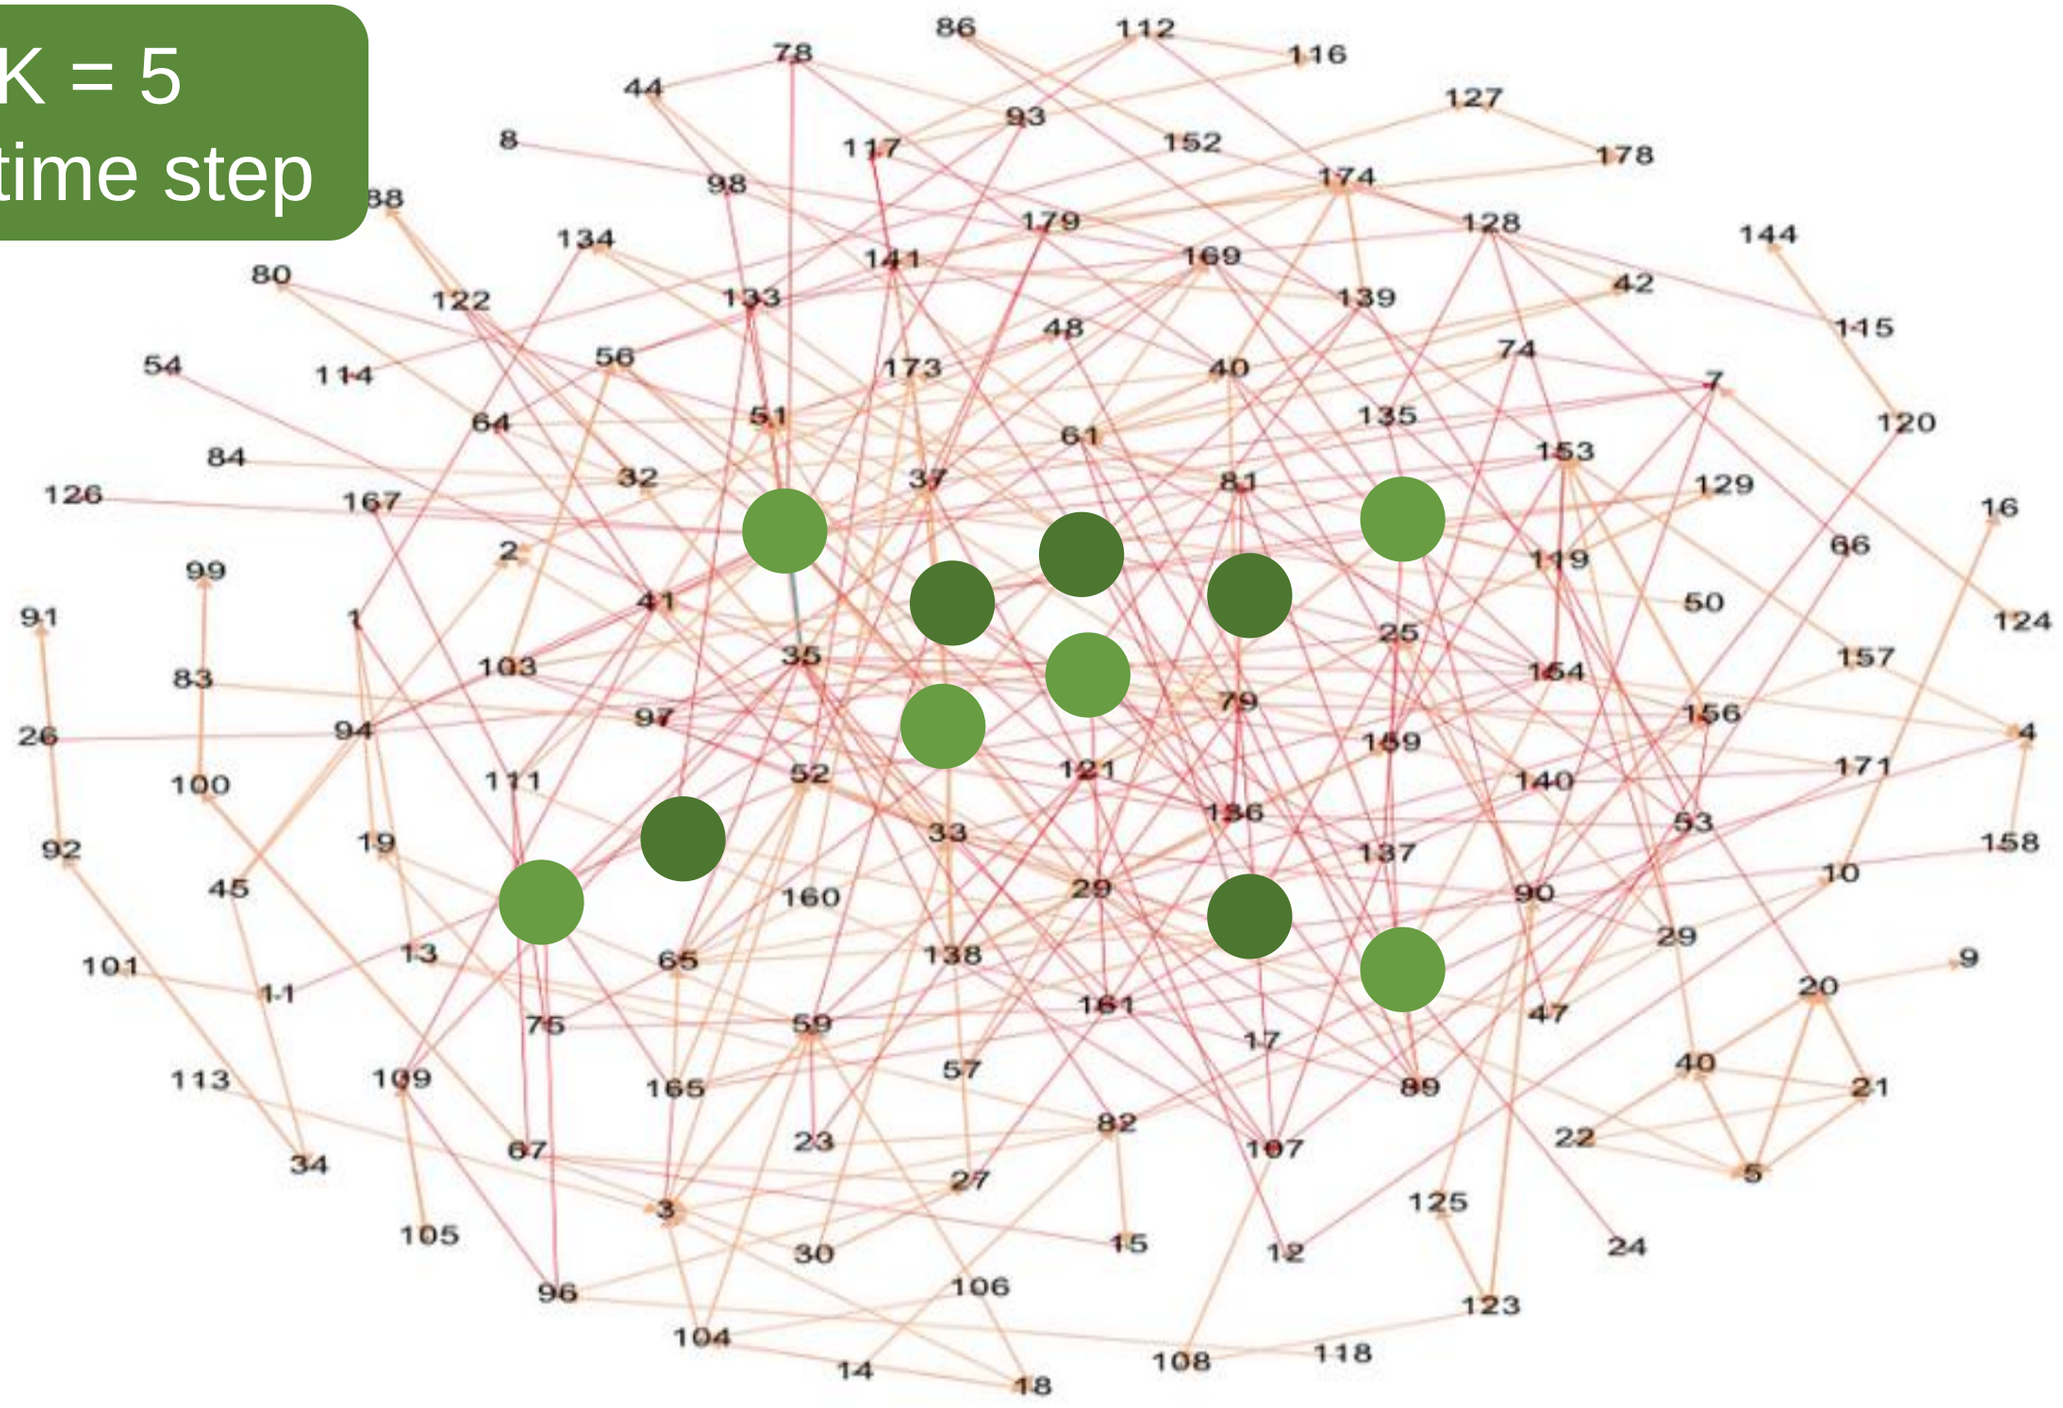

Challenge 3 : Adaptive selection

K = 5
3rd time step

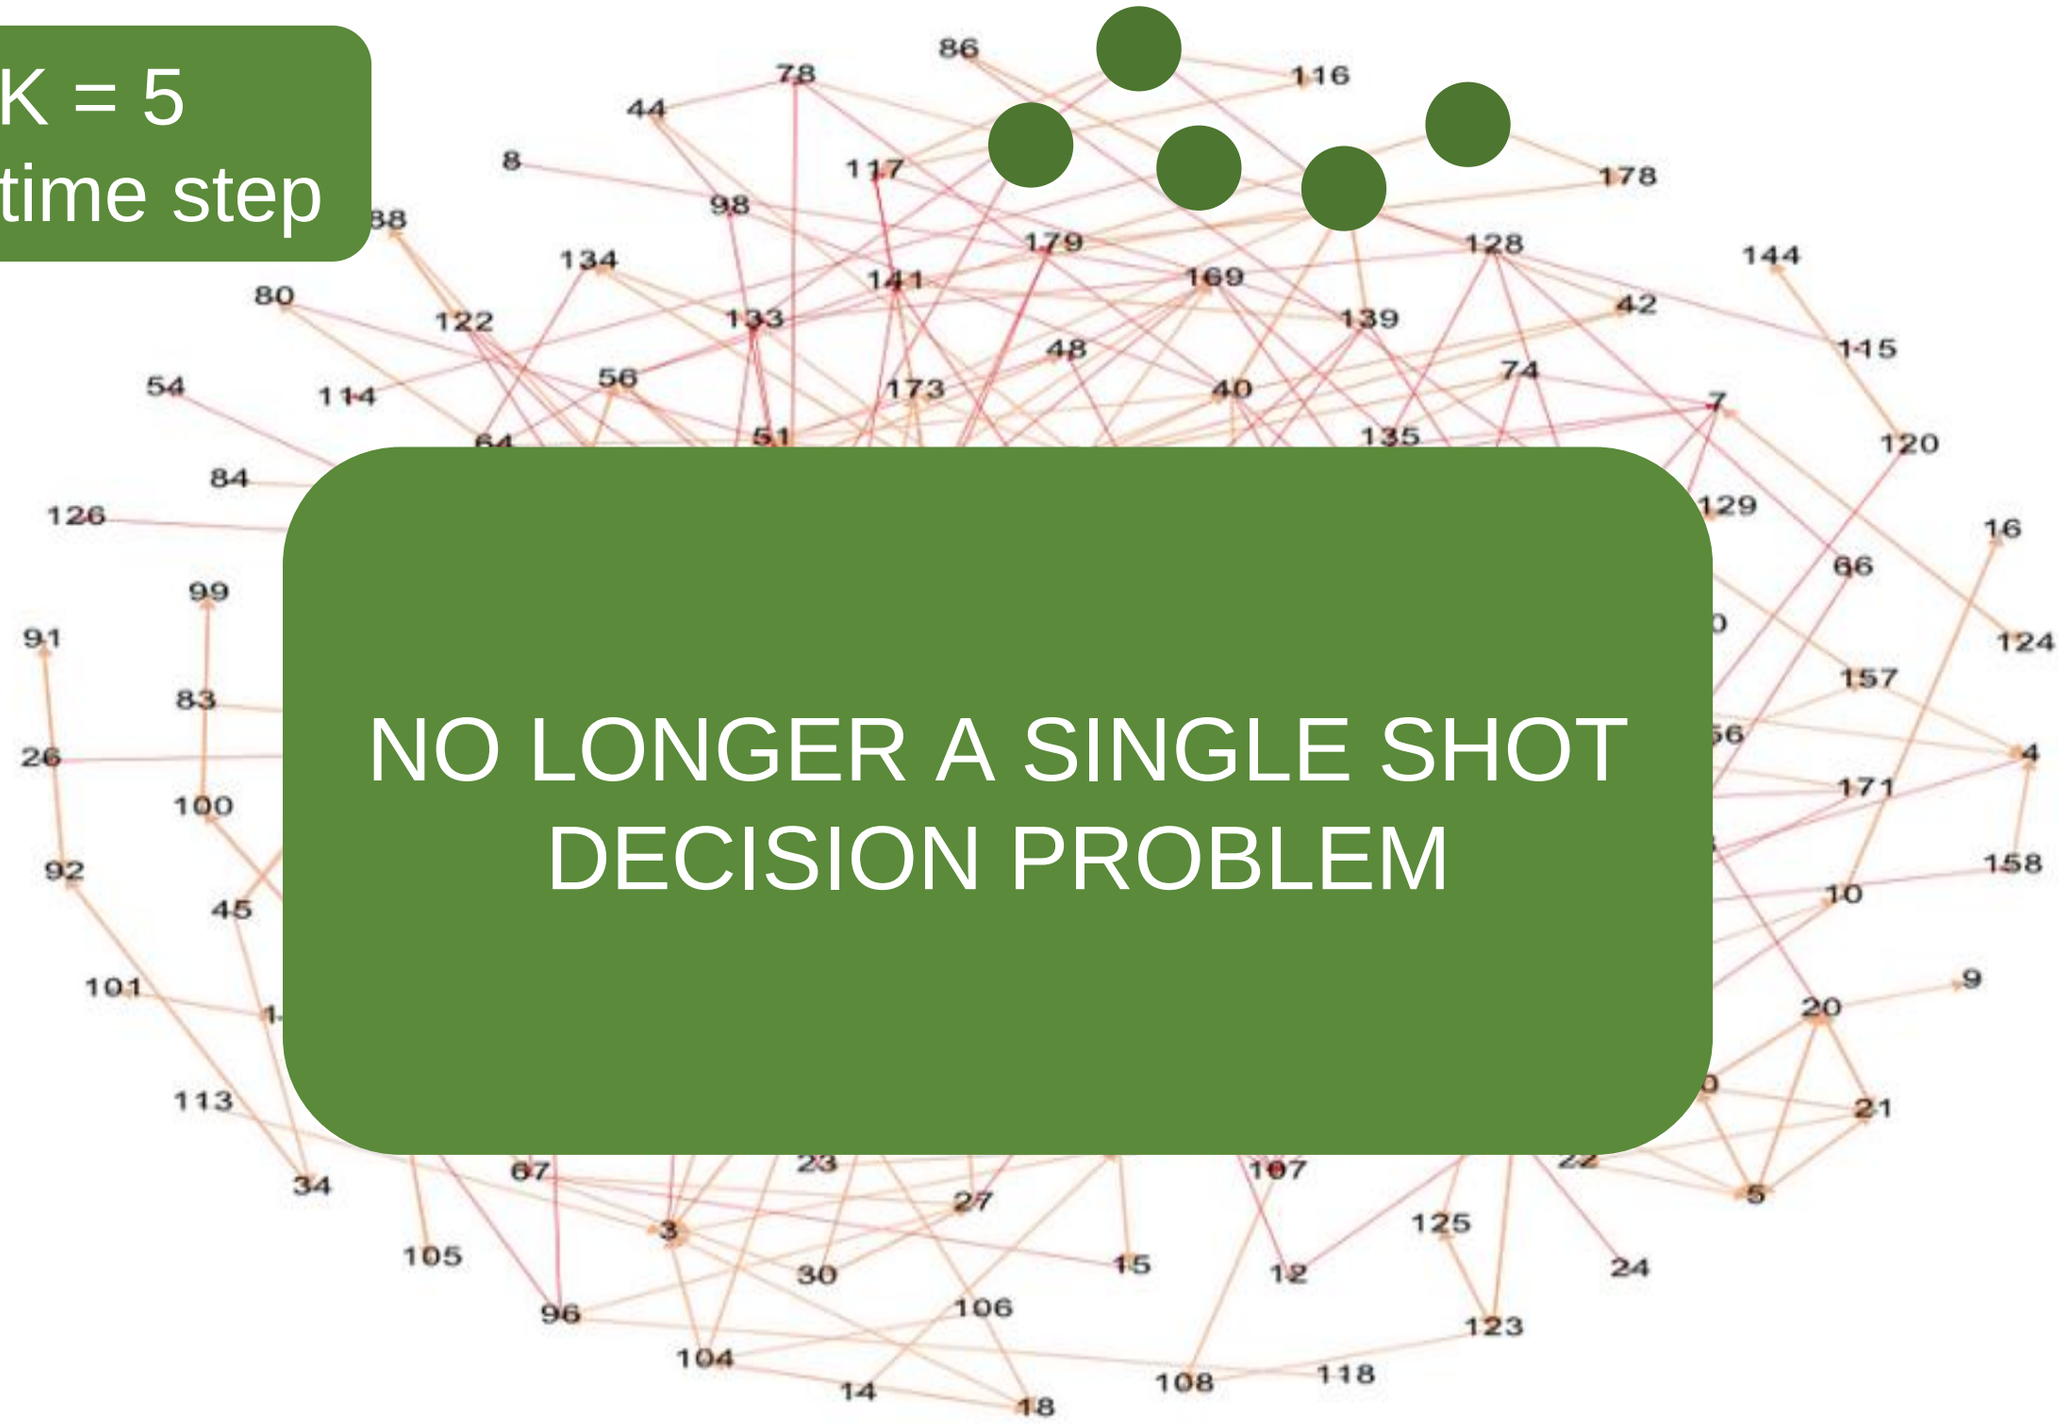

NO LONGER A SINGLE SHOT DECISION PROBLEM

Outline: HIV Information & Homeless Youth

- Domain of homeless youth and HIV information dissemination
- Real World Challenges in Influence Maximization
- POMDP Model and algorithms
- Pilot Study

POMDP Model

- Sequential decision making under uncertainty
  - Homeless shelters – sequentially select nodes
  - Homeless shelters – network state not known

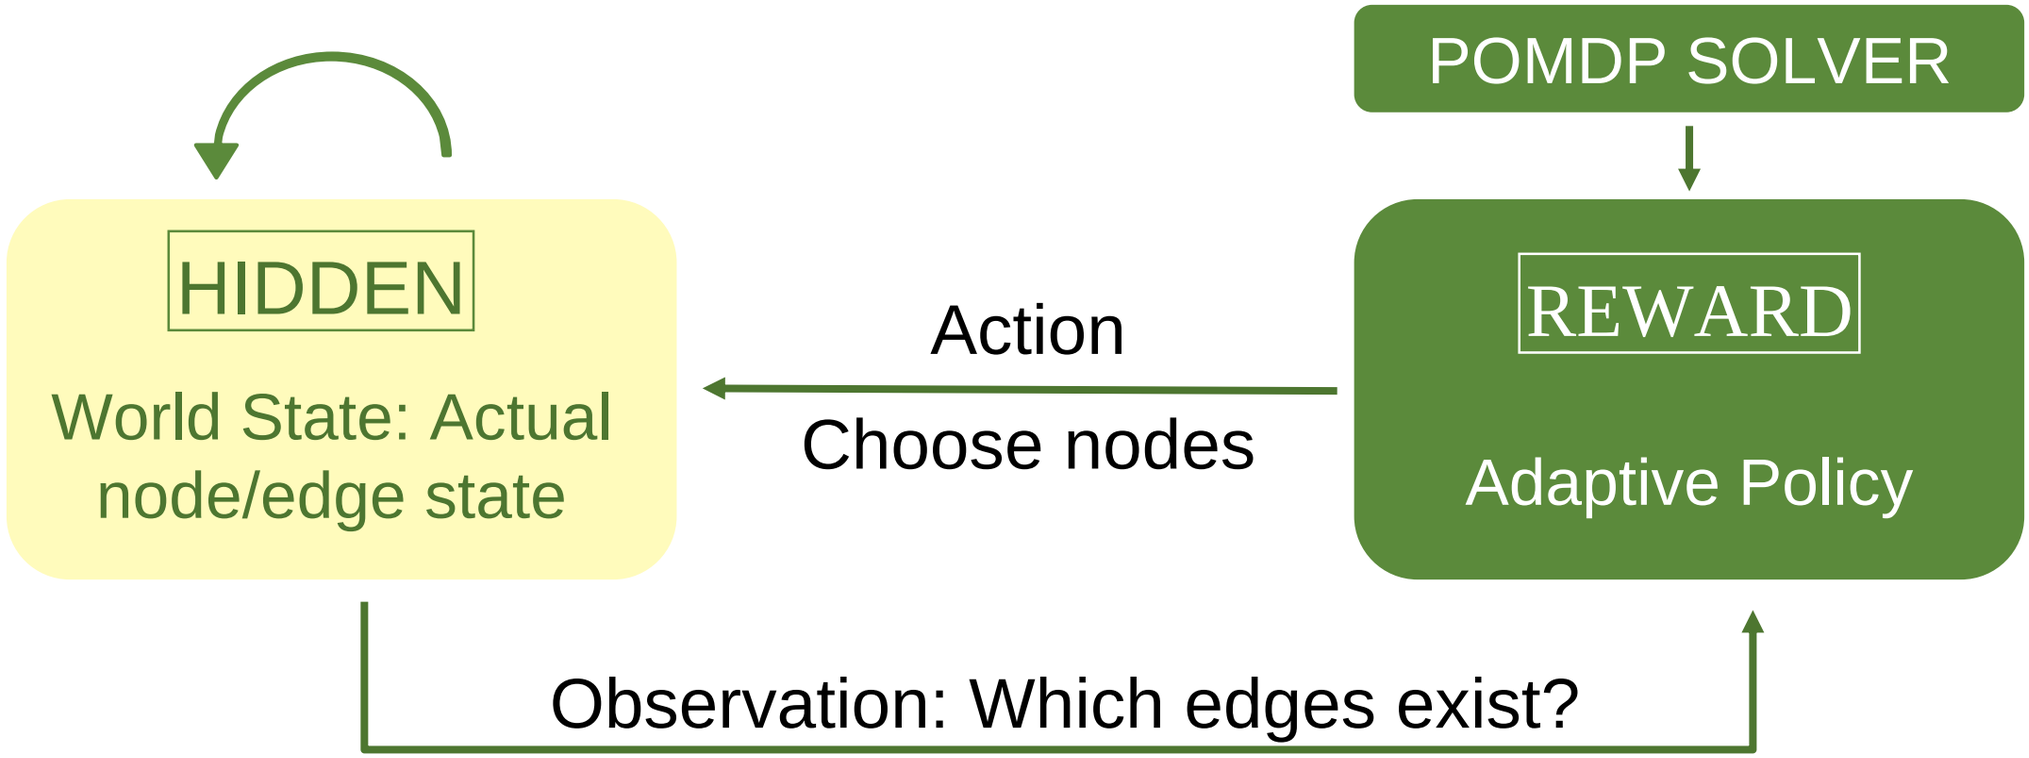

Real world scale: Why is it hard to solve?

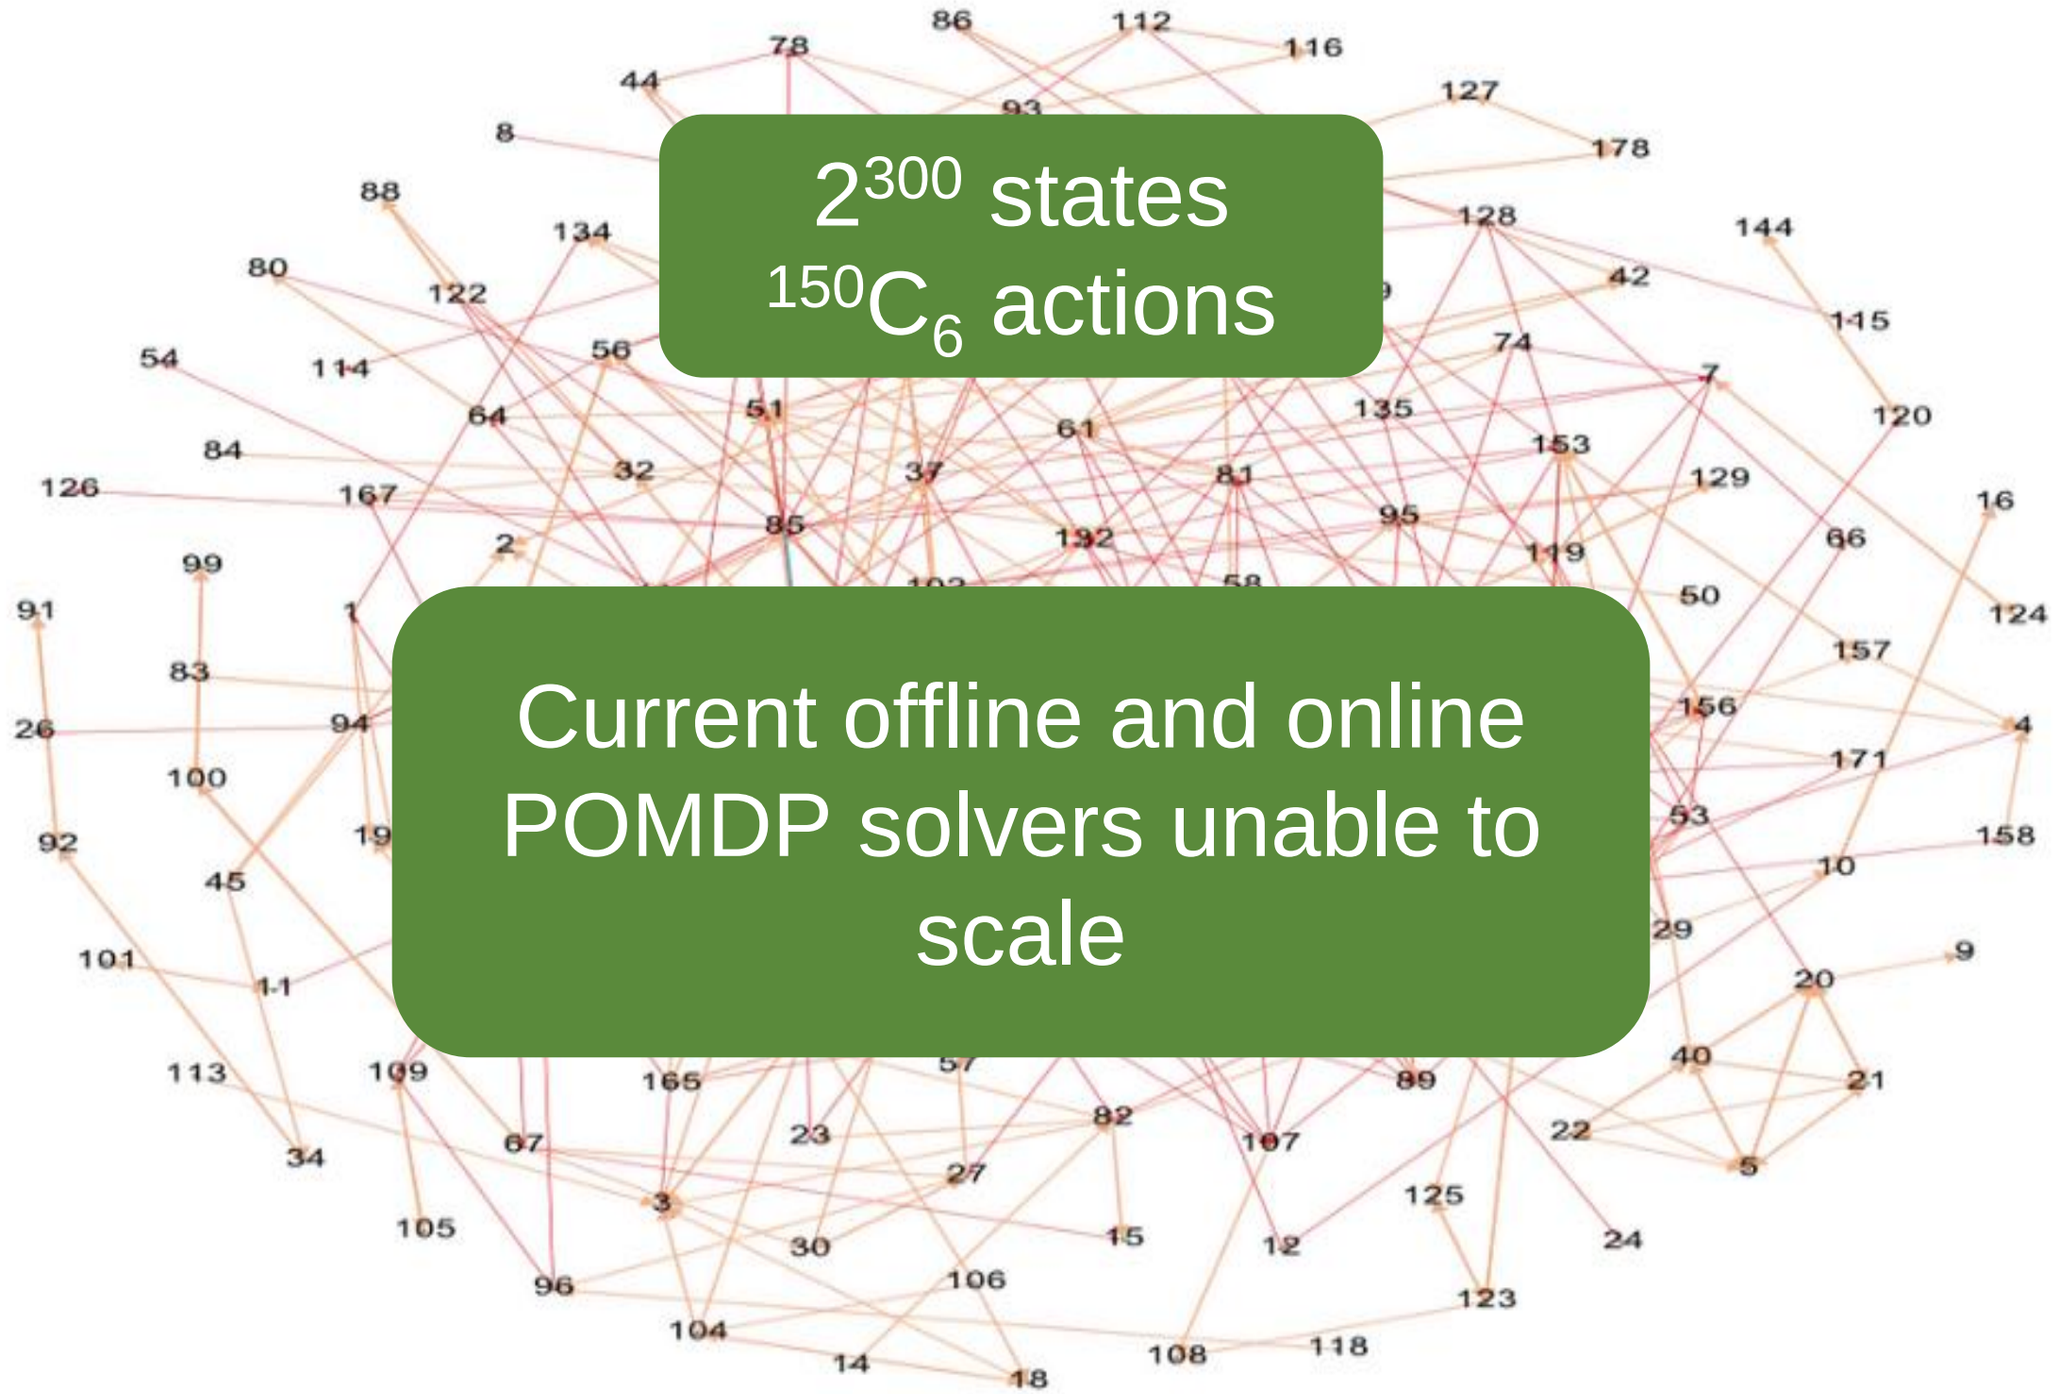

Real world scale: Why is it hard to solve?

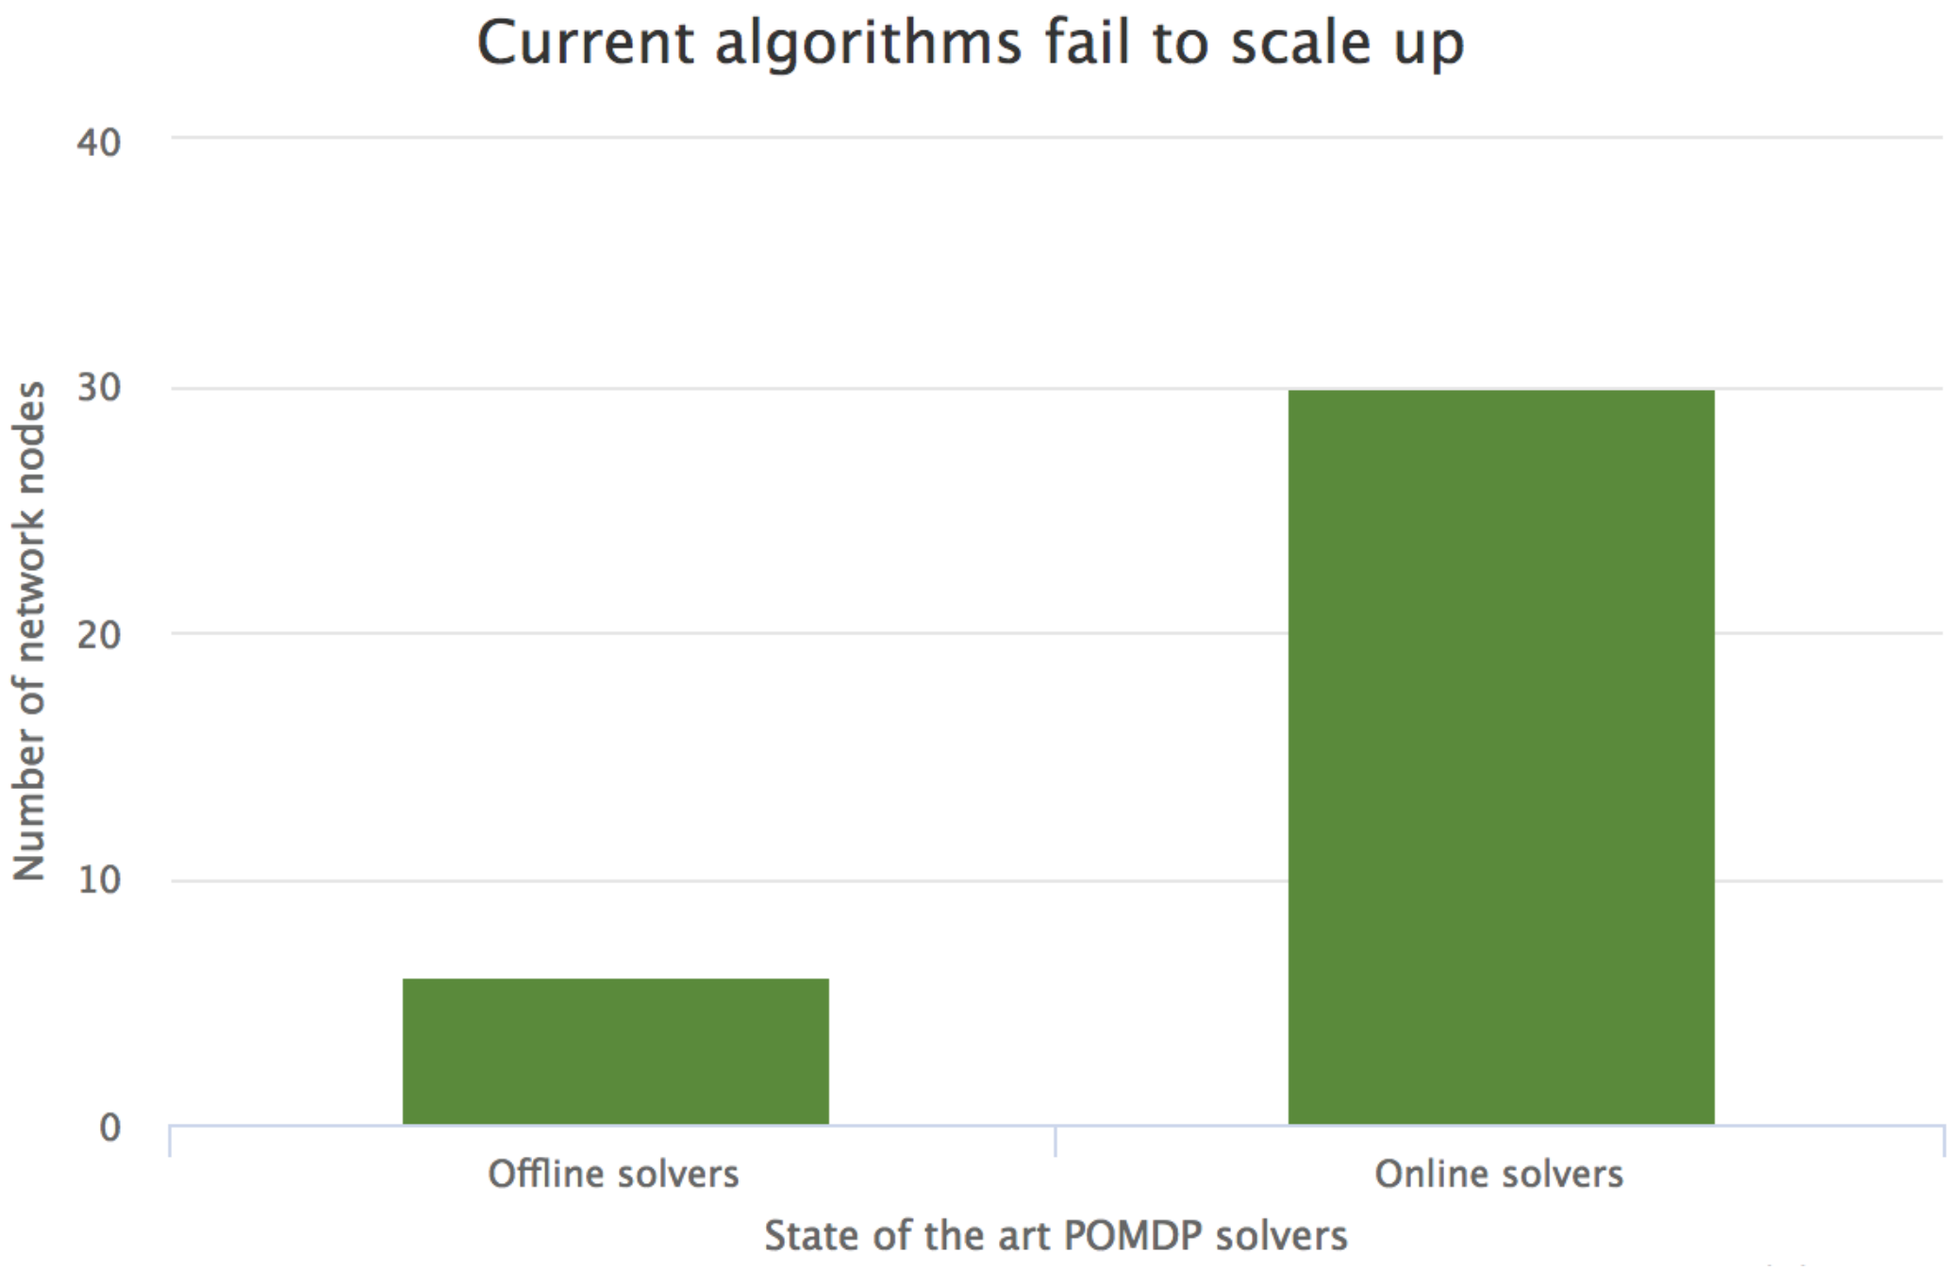

Real world networks have community structure

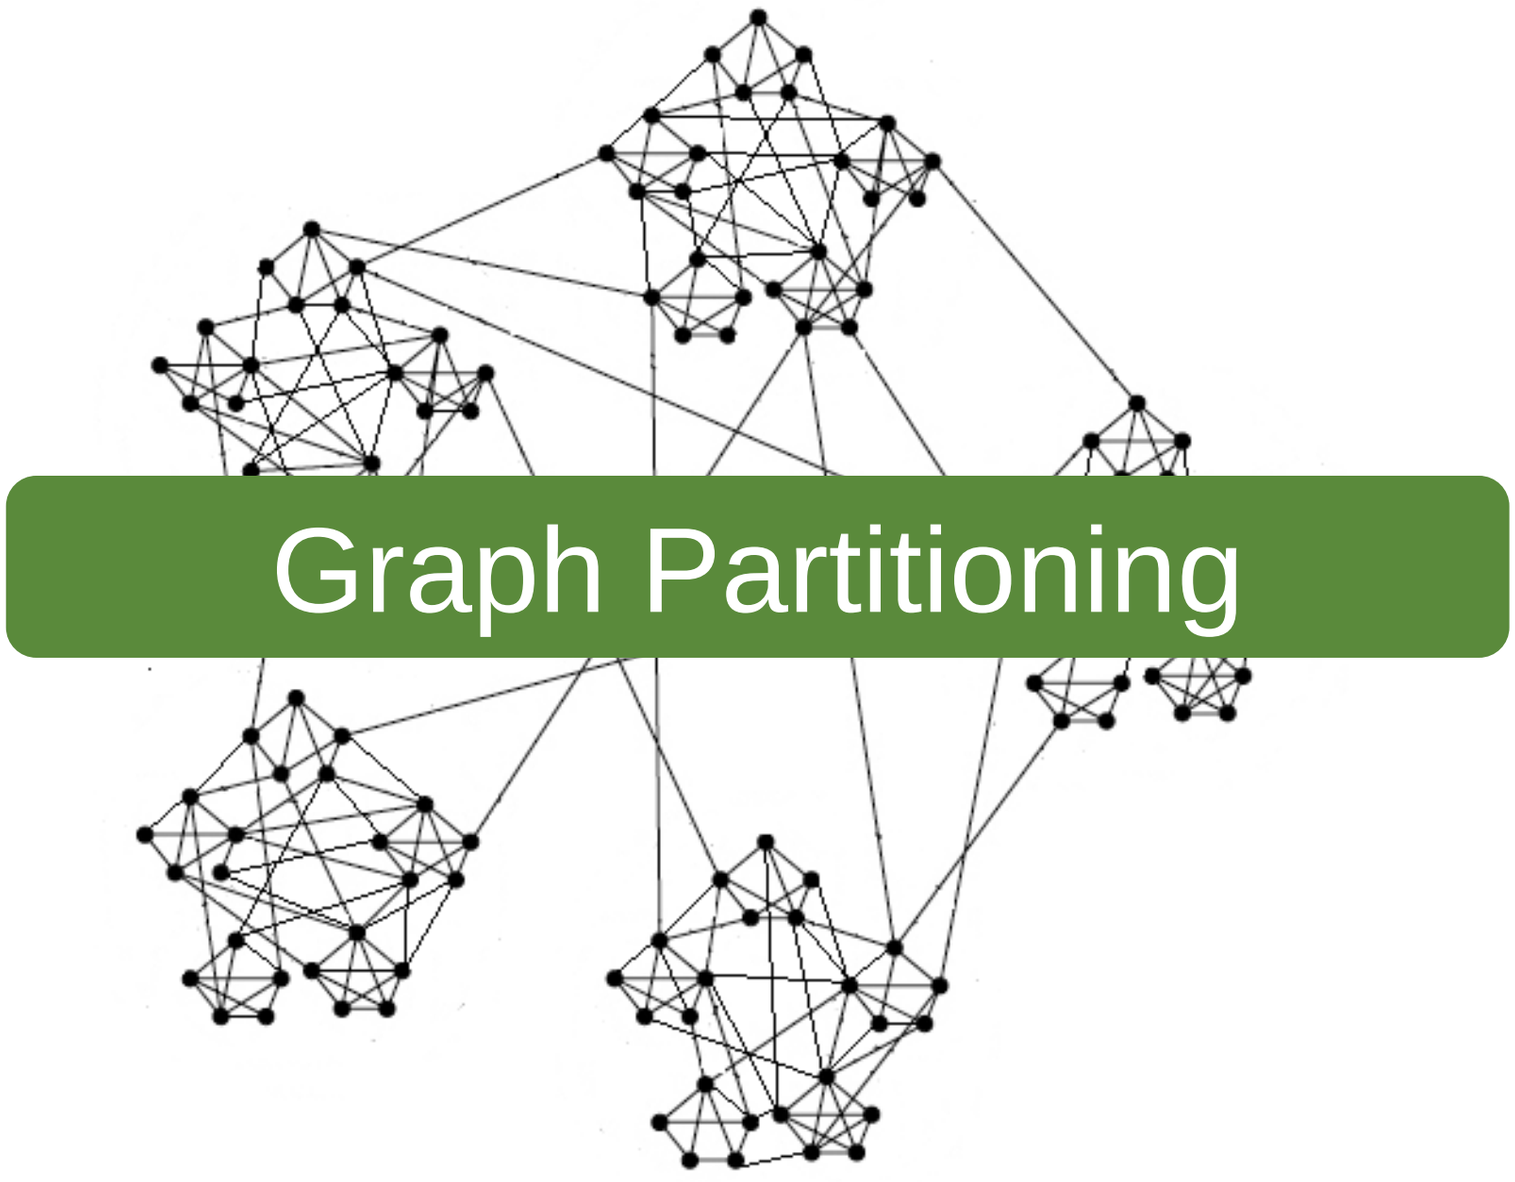

Graph Partitioning

Graph Partitioning

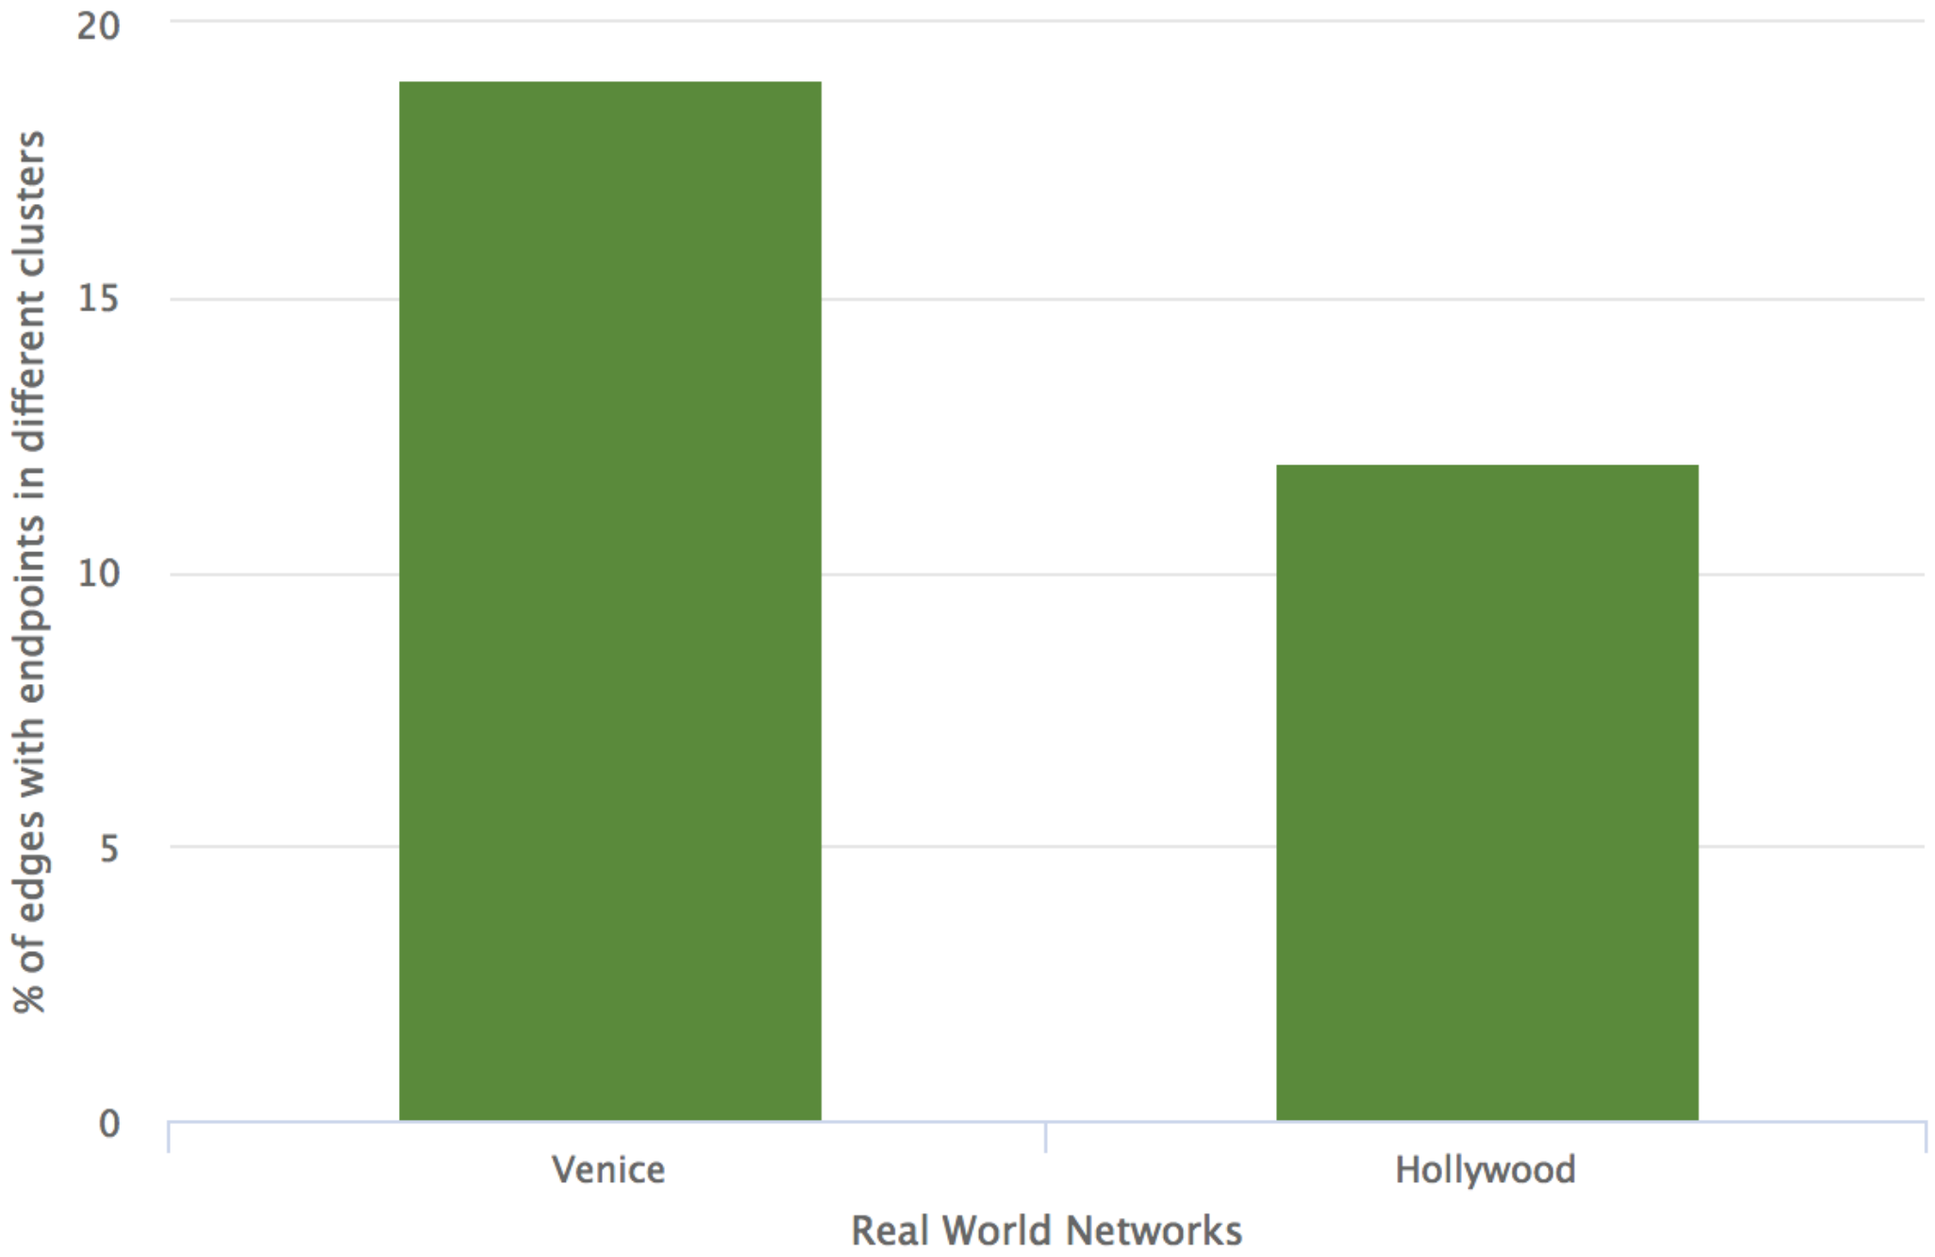

HEALER : Hierarchical Ensembling

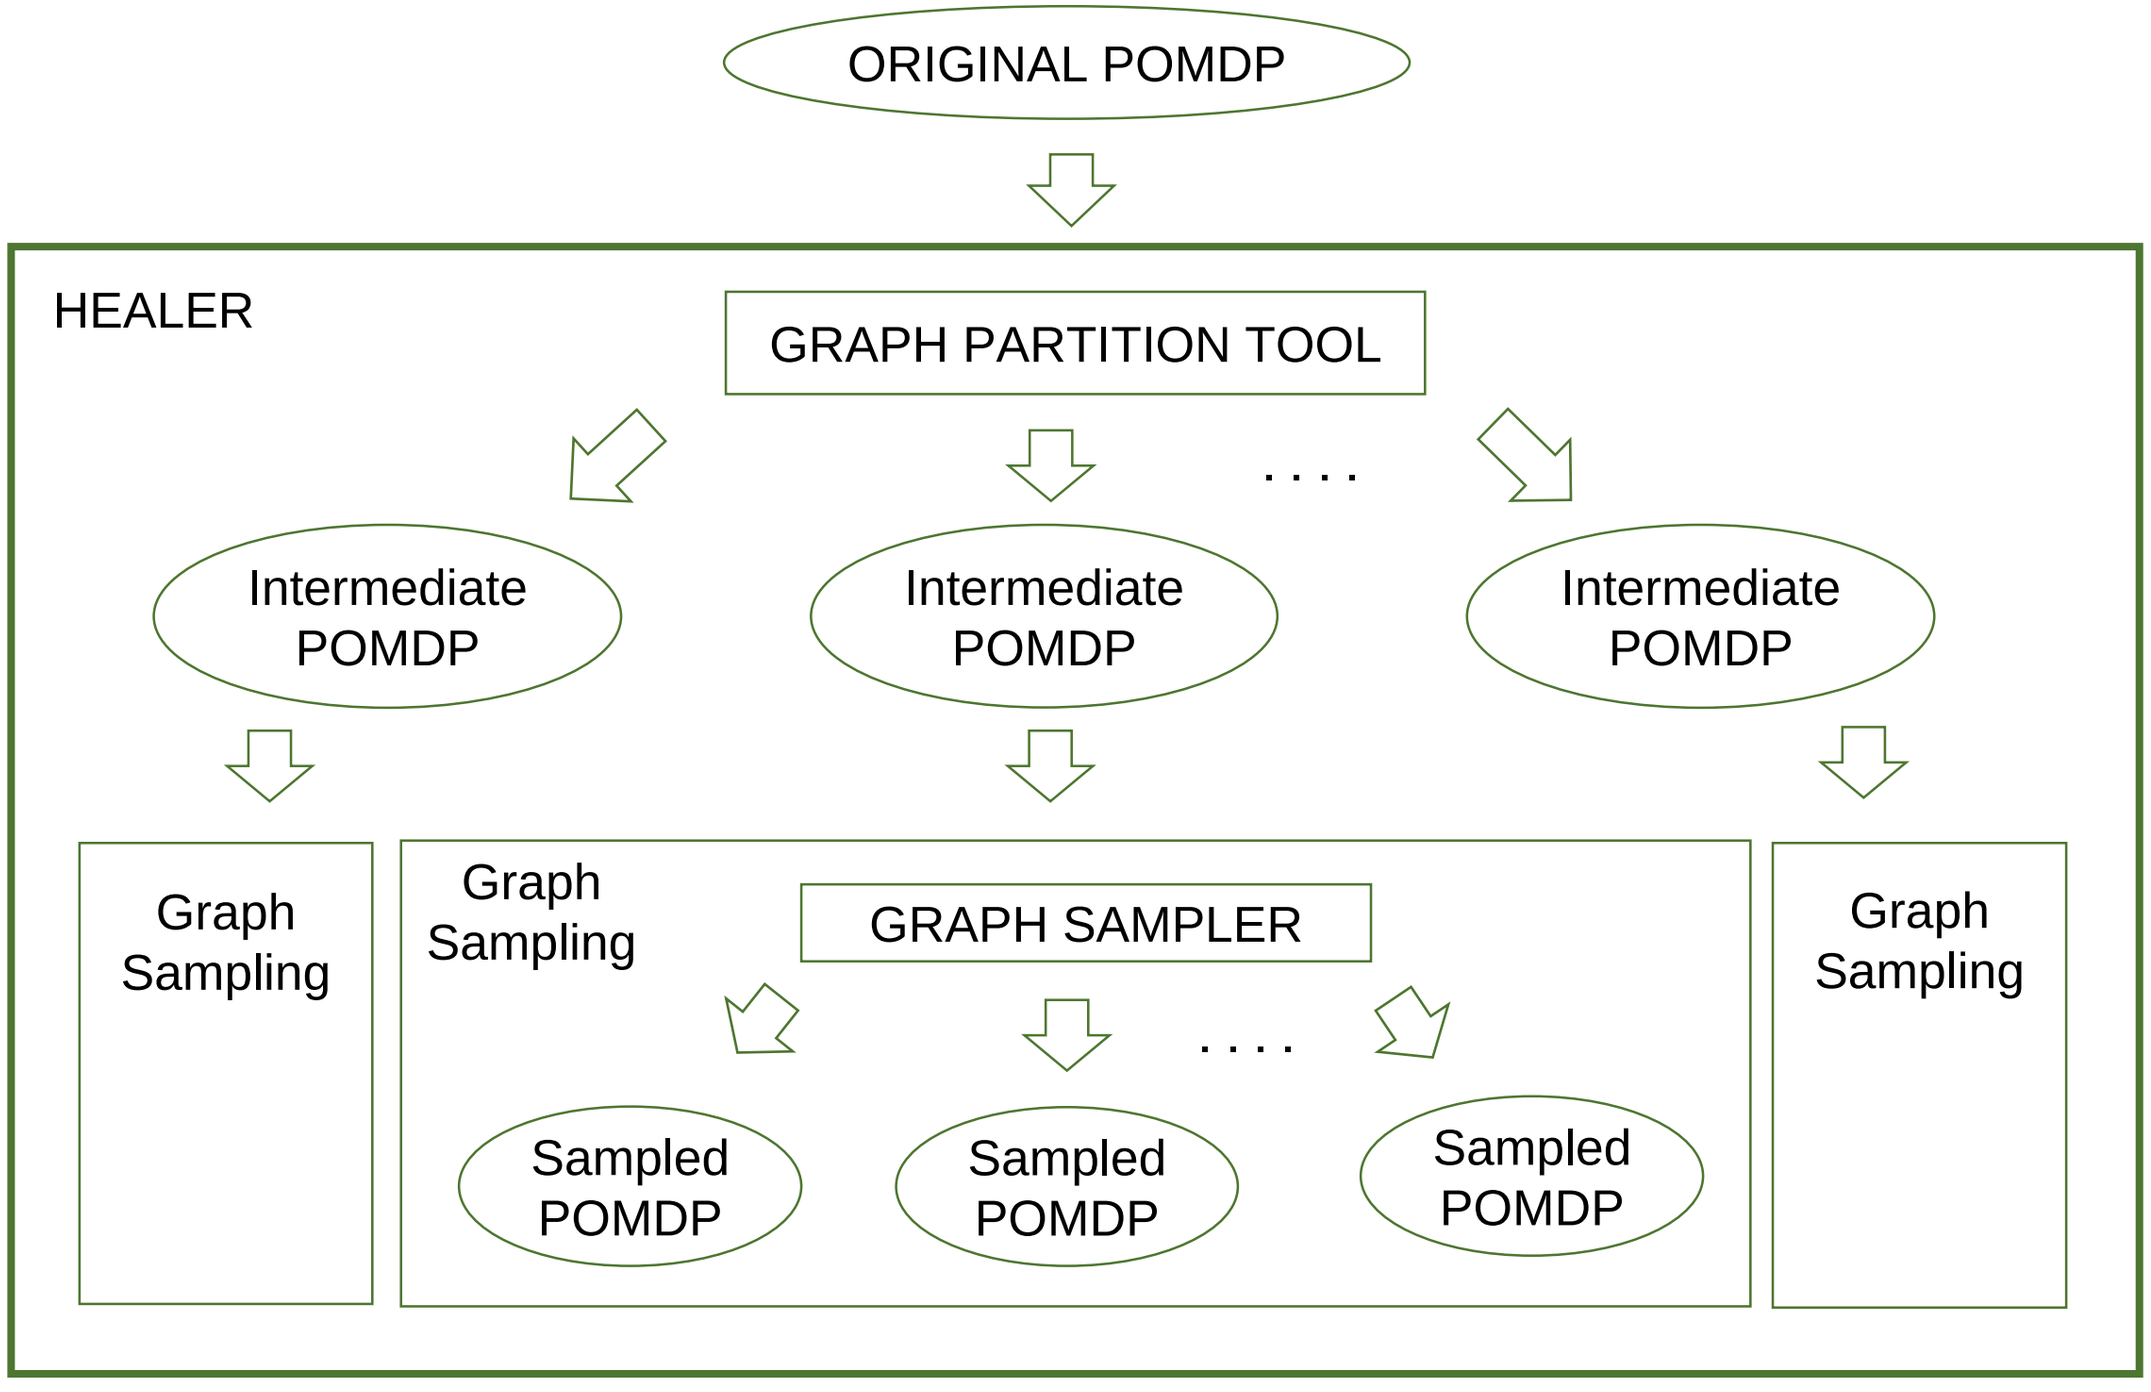

Graph Partitioning: Why partition?

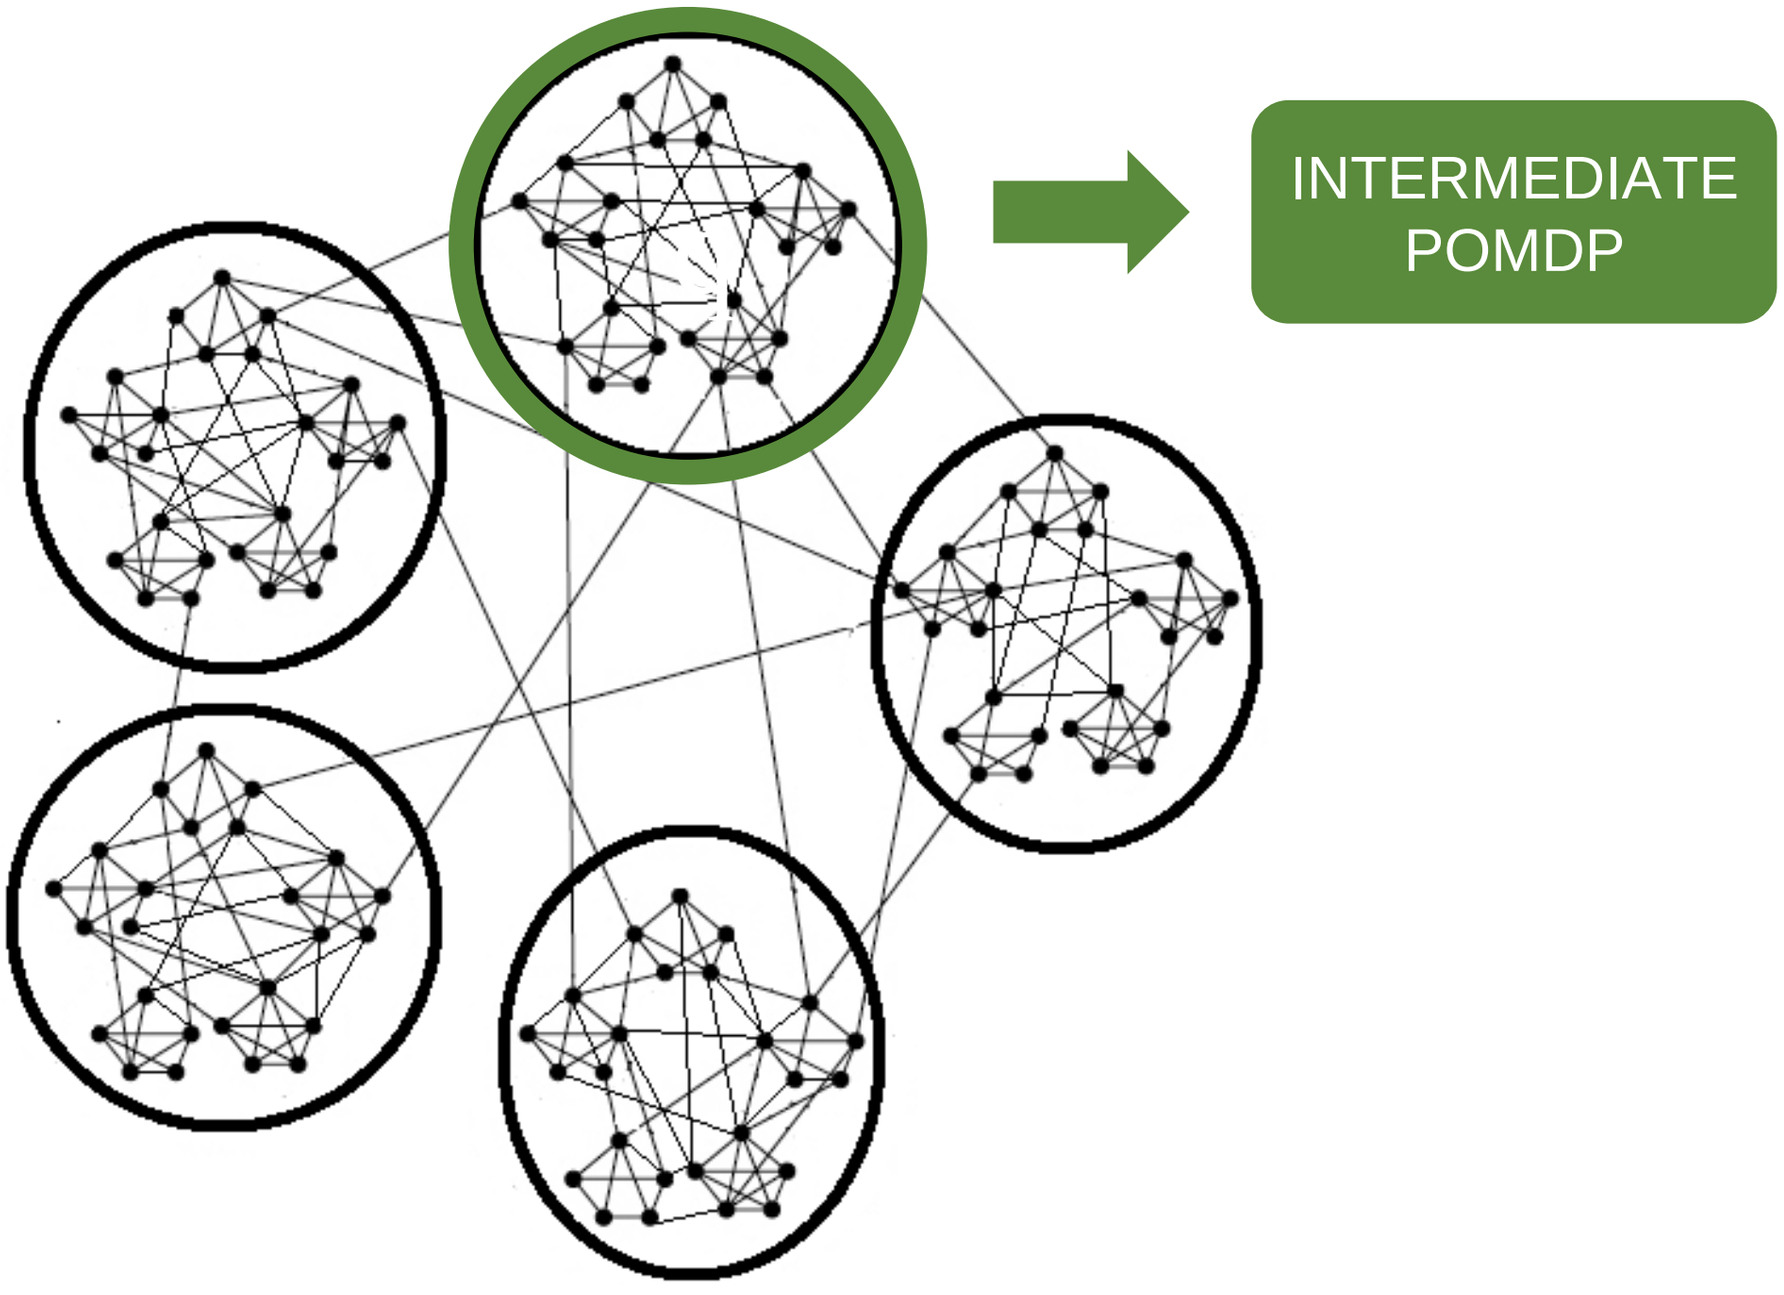

Real Networks - Solution Quality

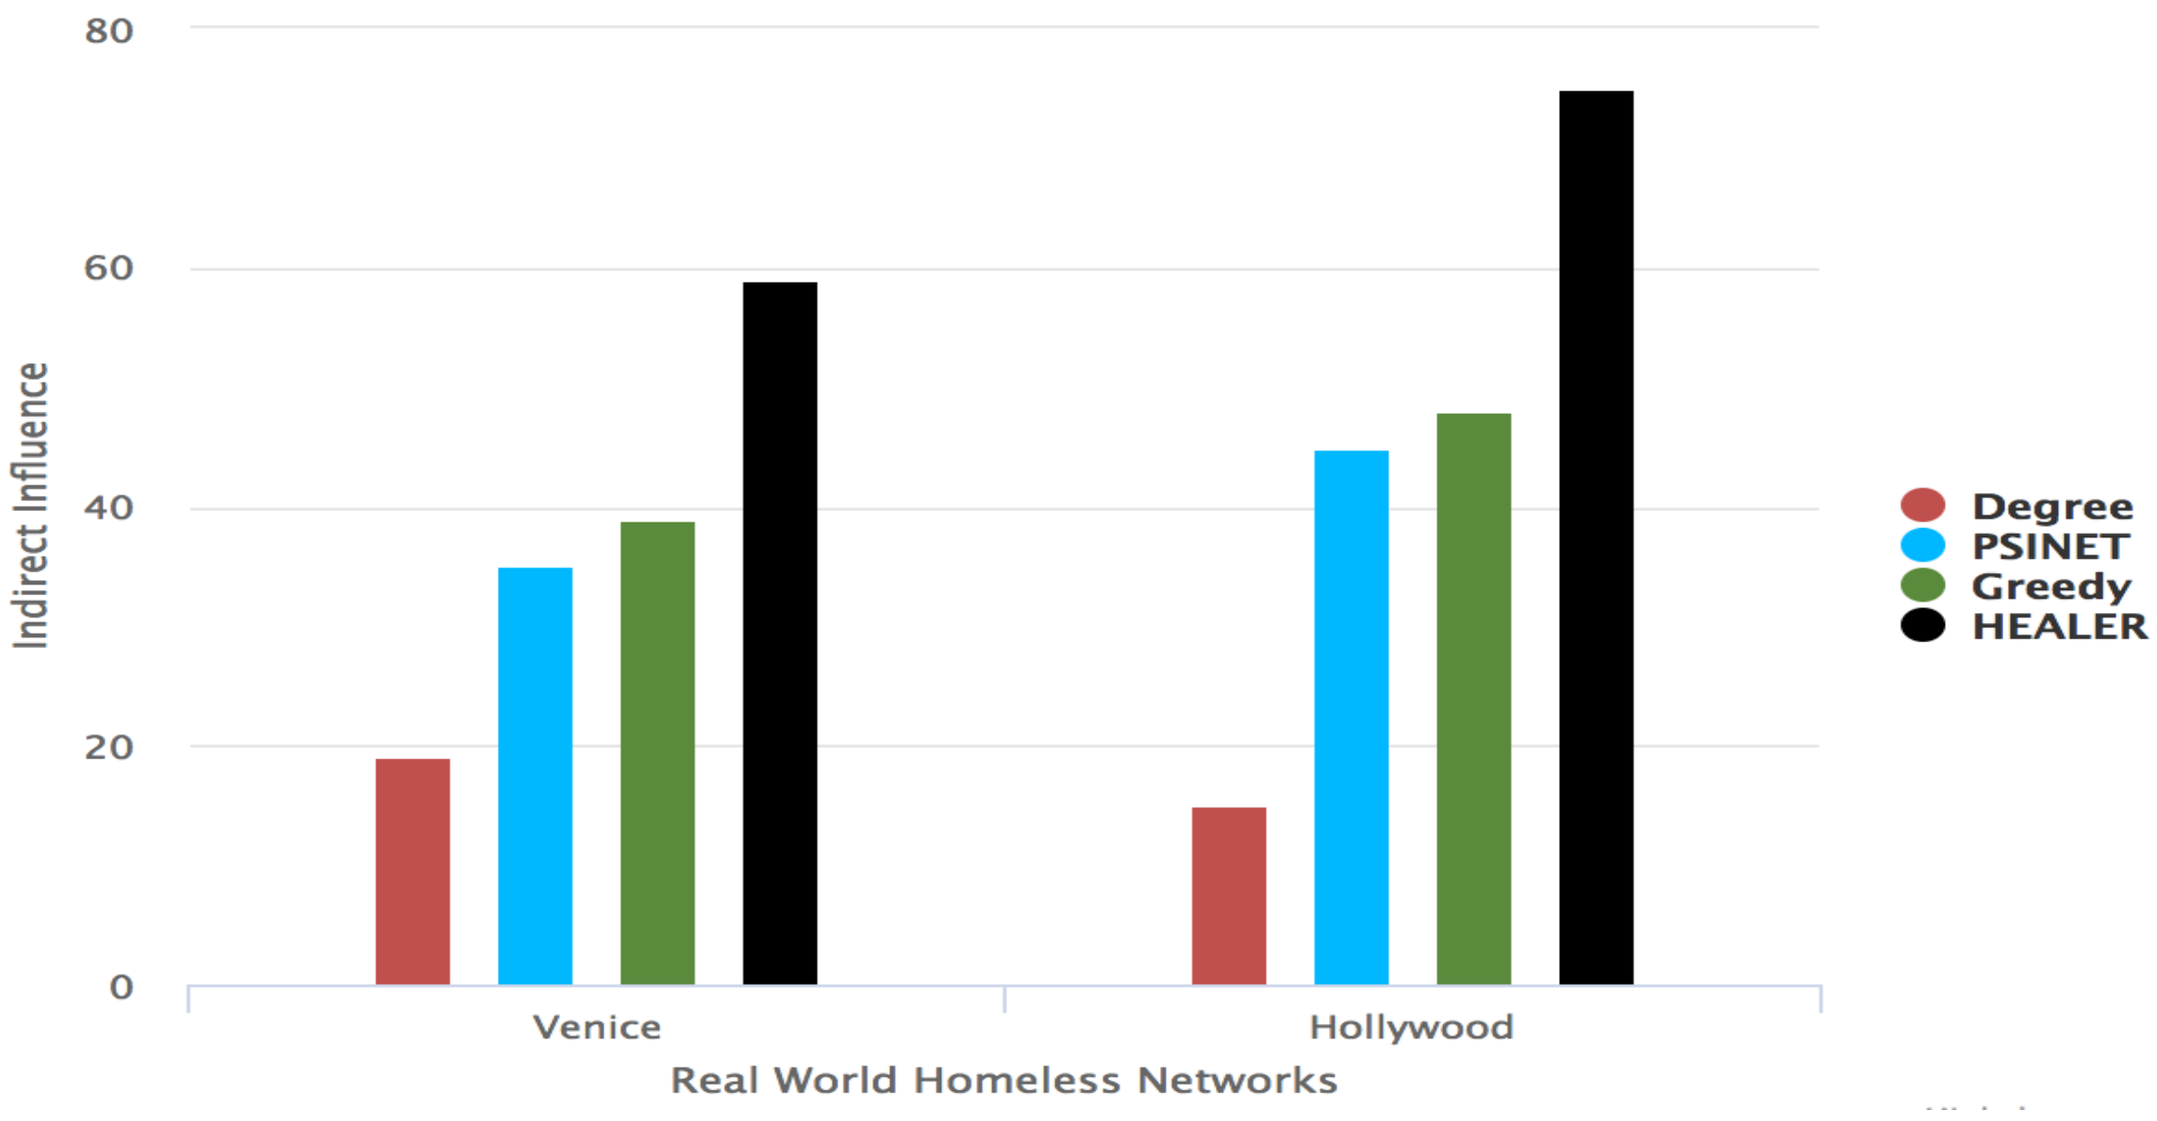

Robustness & Parameter Uncertainty

- HEALER: fixed propagation and existence probability

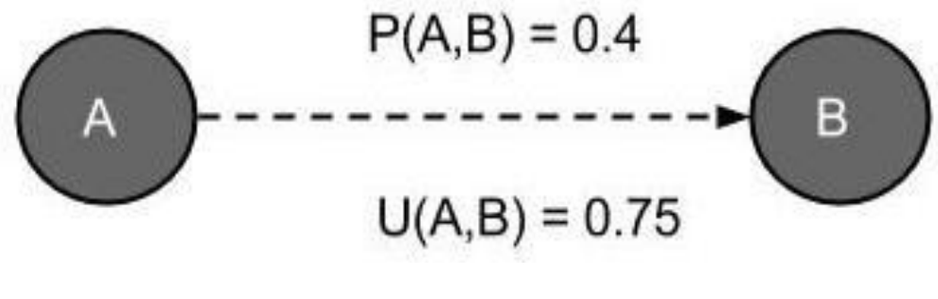

- Want policies robust to different possible values of P(A,B) and U(A,B)
  - Express as ranges of values, e.g., U(A,B) is in [0.4, 0.8]

HEALER++
Robustness & Parameter Uncertainty

- Worst case parameters: a zero-sum game against nature

Nature	vs	**Algorithm**
Chooses parameters P(A,B) and U(A,B)		Chooses policy in POMDP state-action space

- Payoffs: (performance of algorithm)/OPT

Outline: HIV Information & Homeless Youth

- Domain of homeless youth and HIV information dissemination
- Real World Challenges in Influence Maximization
- POMDP Model and algorithms
- Pilot Study

Pilot Tests with 170 Youth in LA Area

Recruited youths:

HEALER	HEALER++	DEGREE CENTRALITY
62	56	55

Preliminary network —> HEALER

Bring 4 youth for training, get edge data —> HEALER

Bring 4 youth for training, get edge data —> HEALER

Bring 4 youth for training

Results: Pilot Studies

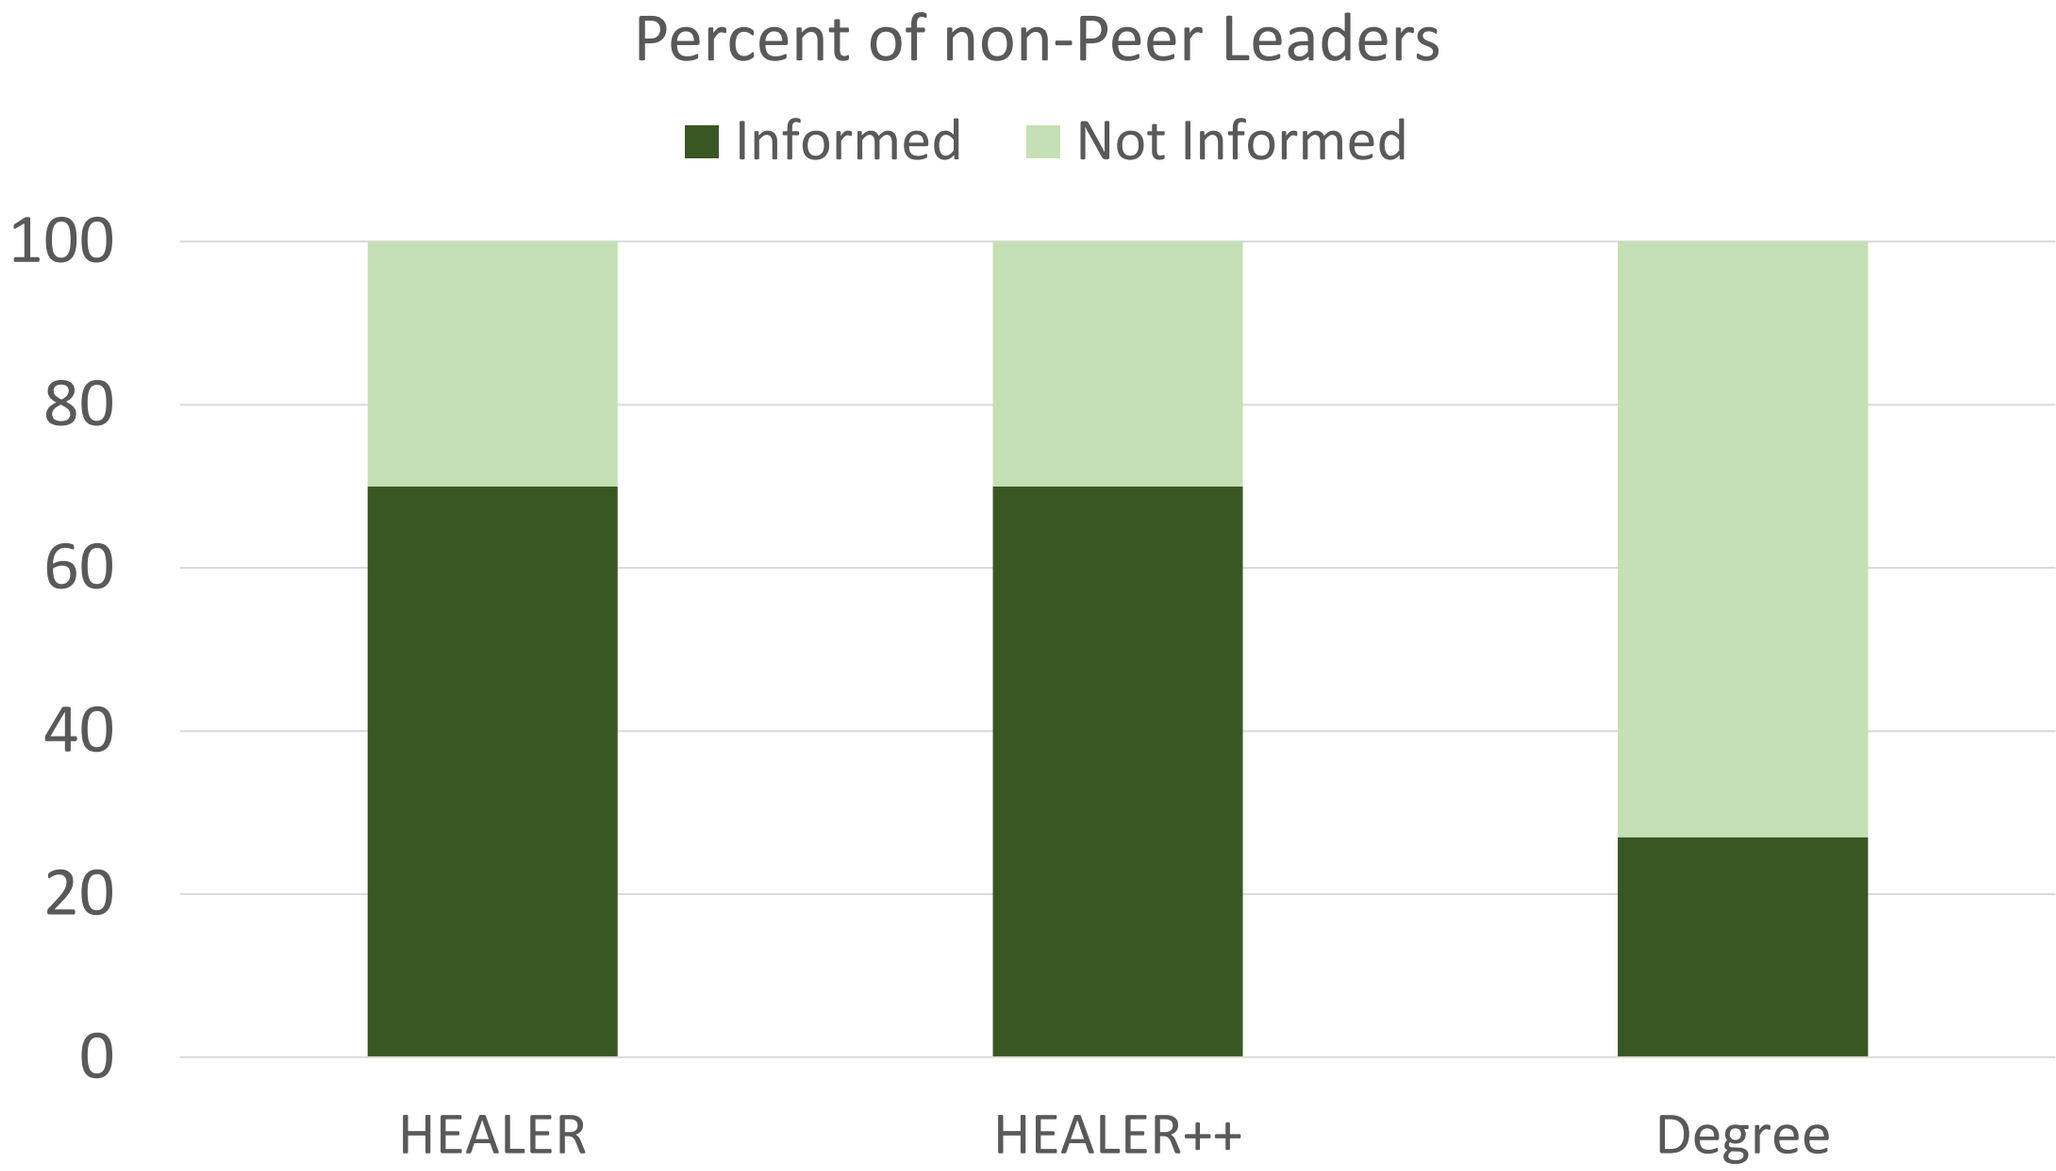

Results: Pilot Studies

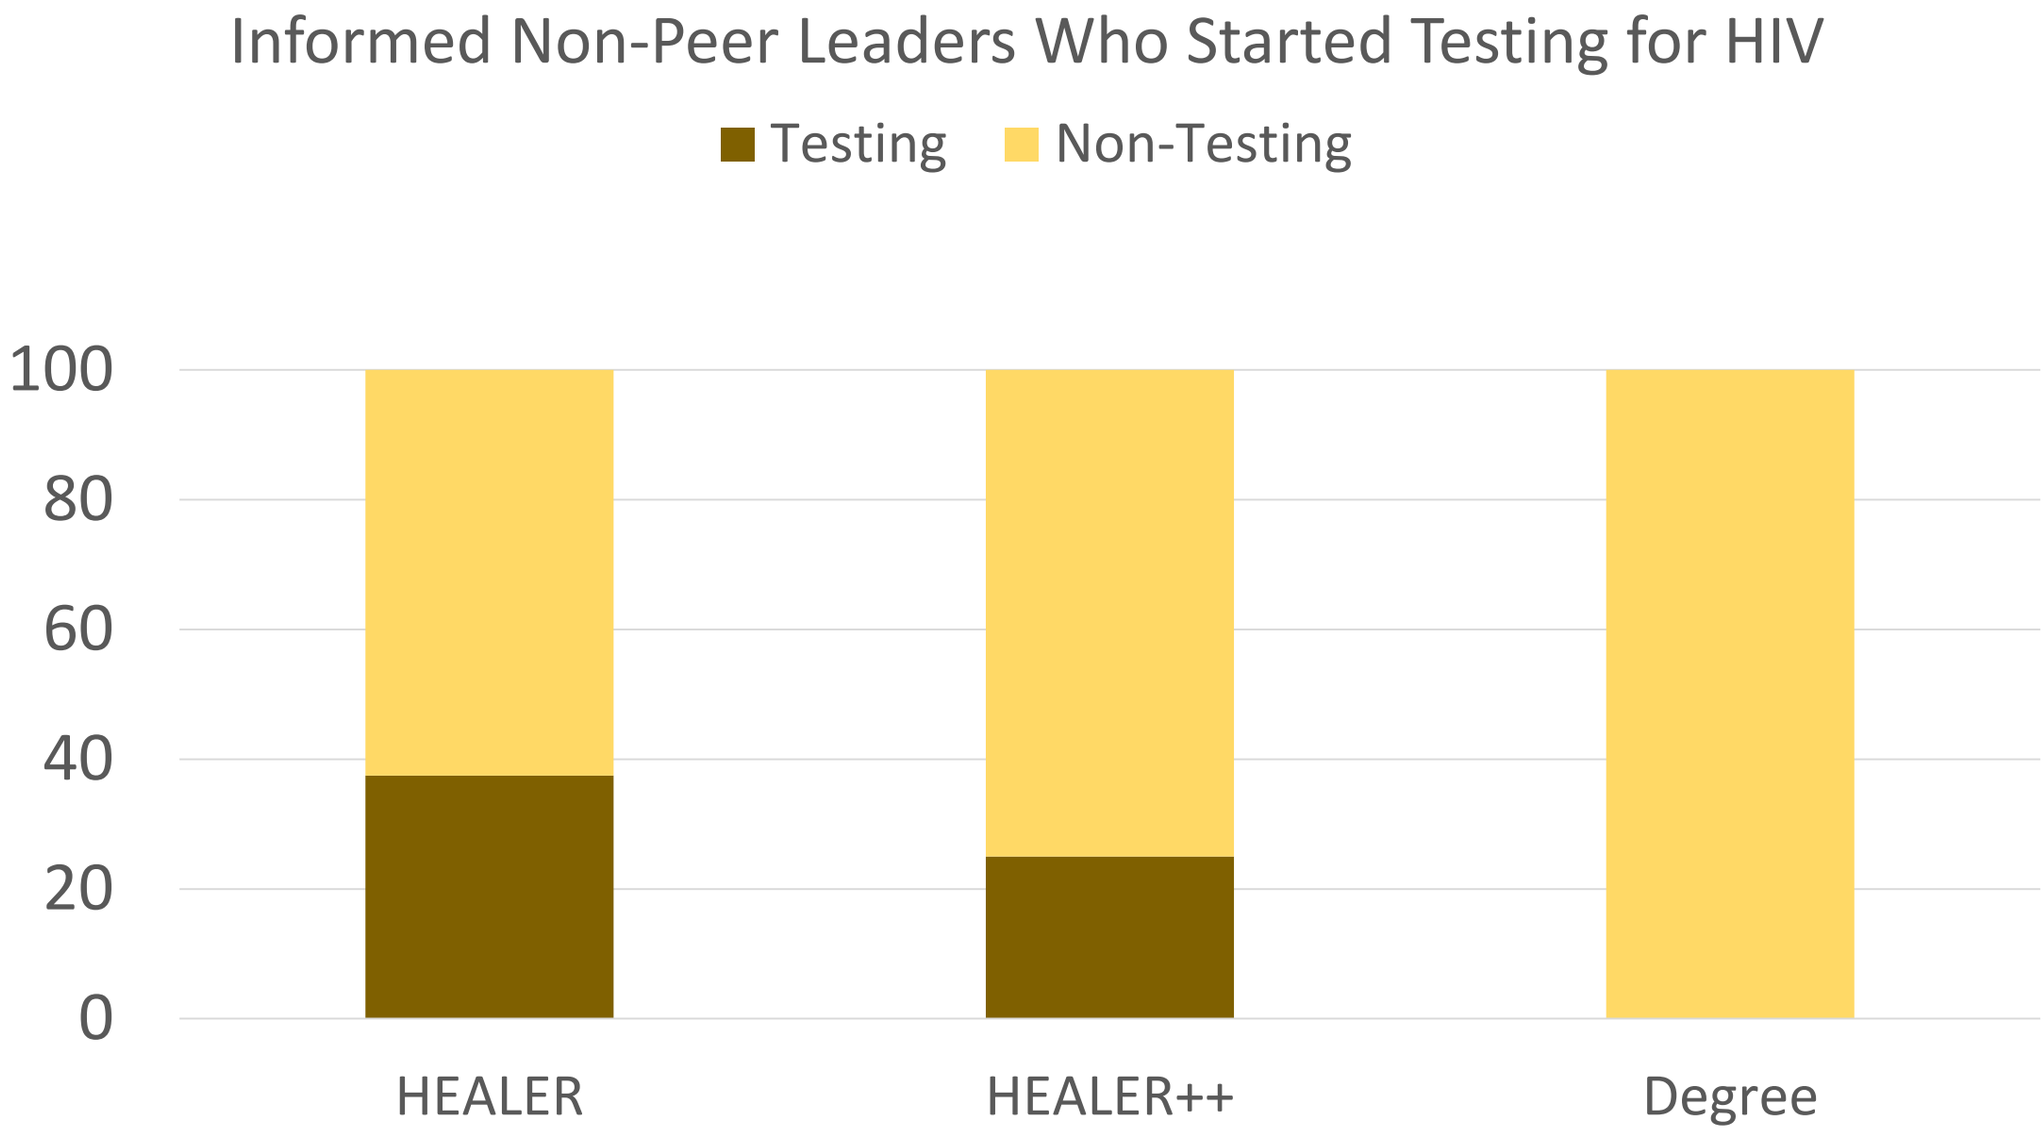

AI Program: HEALER

Next Steps

- 900 youth study begun at three locations in Los Angeles
  - 300 enrolled in HEALER/HEALER++
  - 300 enrolled in no condition
  - 300 in Degree centrality
- IRB approvals
- Presenting video updates every few months

Outline

- Introduction
- HIV Information among homeless youth
- Wildlife Conservation

Protecting Wildlife in Uganda

Uganda Wildlife Authority & Wildlife Conservation Society

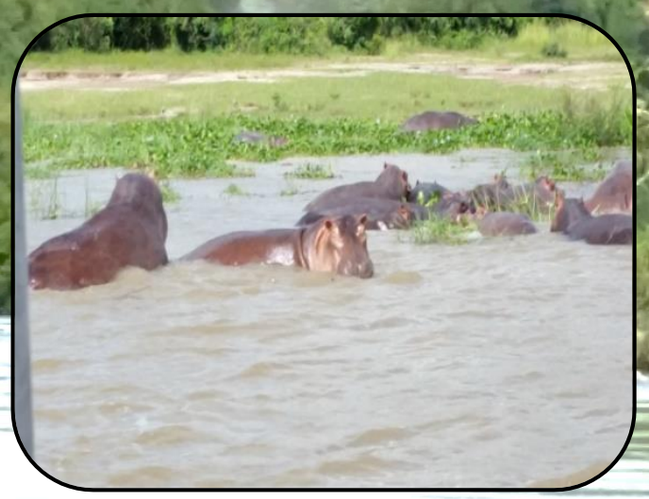

PAWS: Applying AI for protecting wildlife

Poacher Behavior Prediction

Predicting Poaching from Past Crime Data

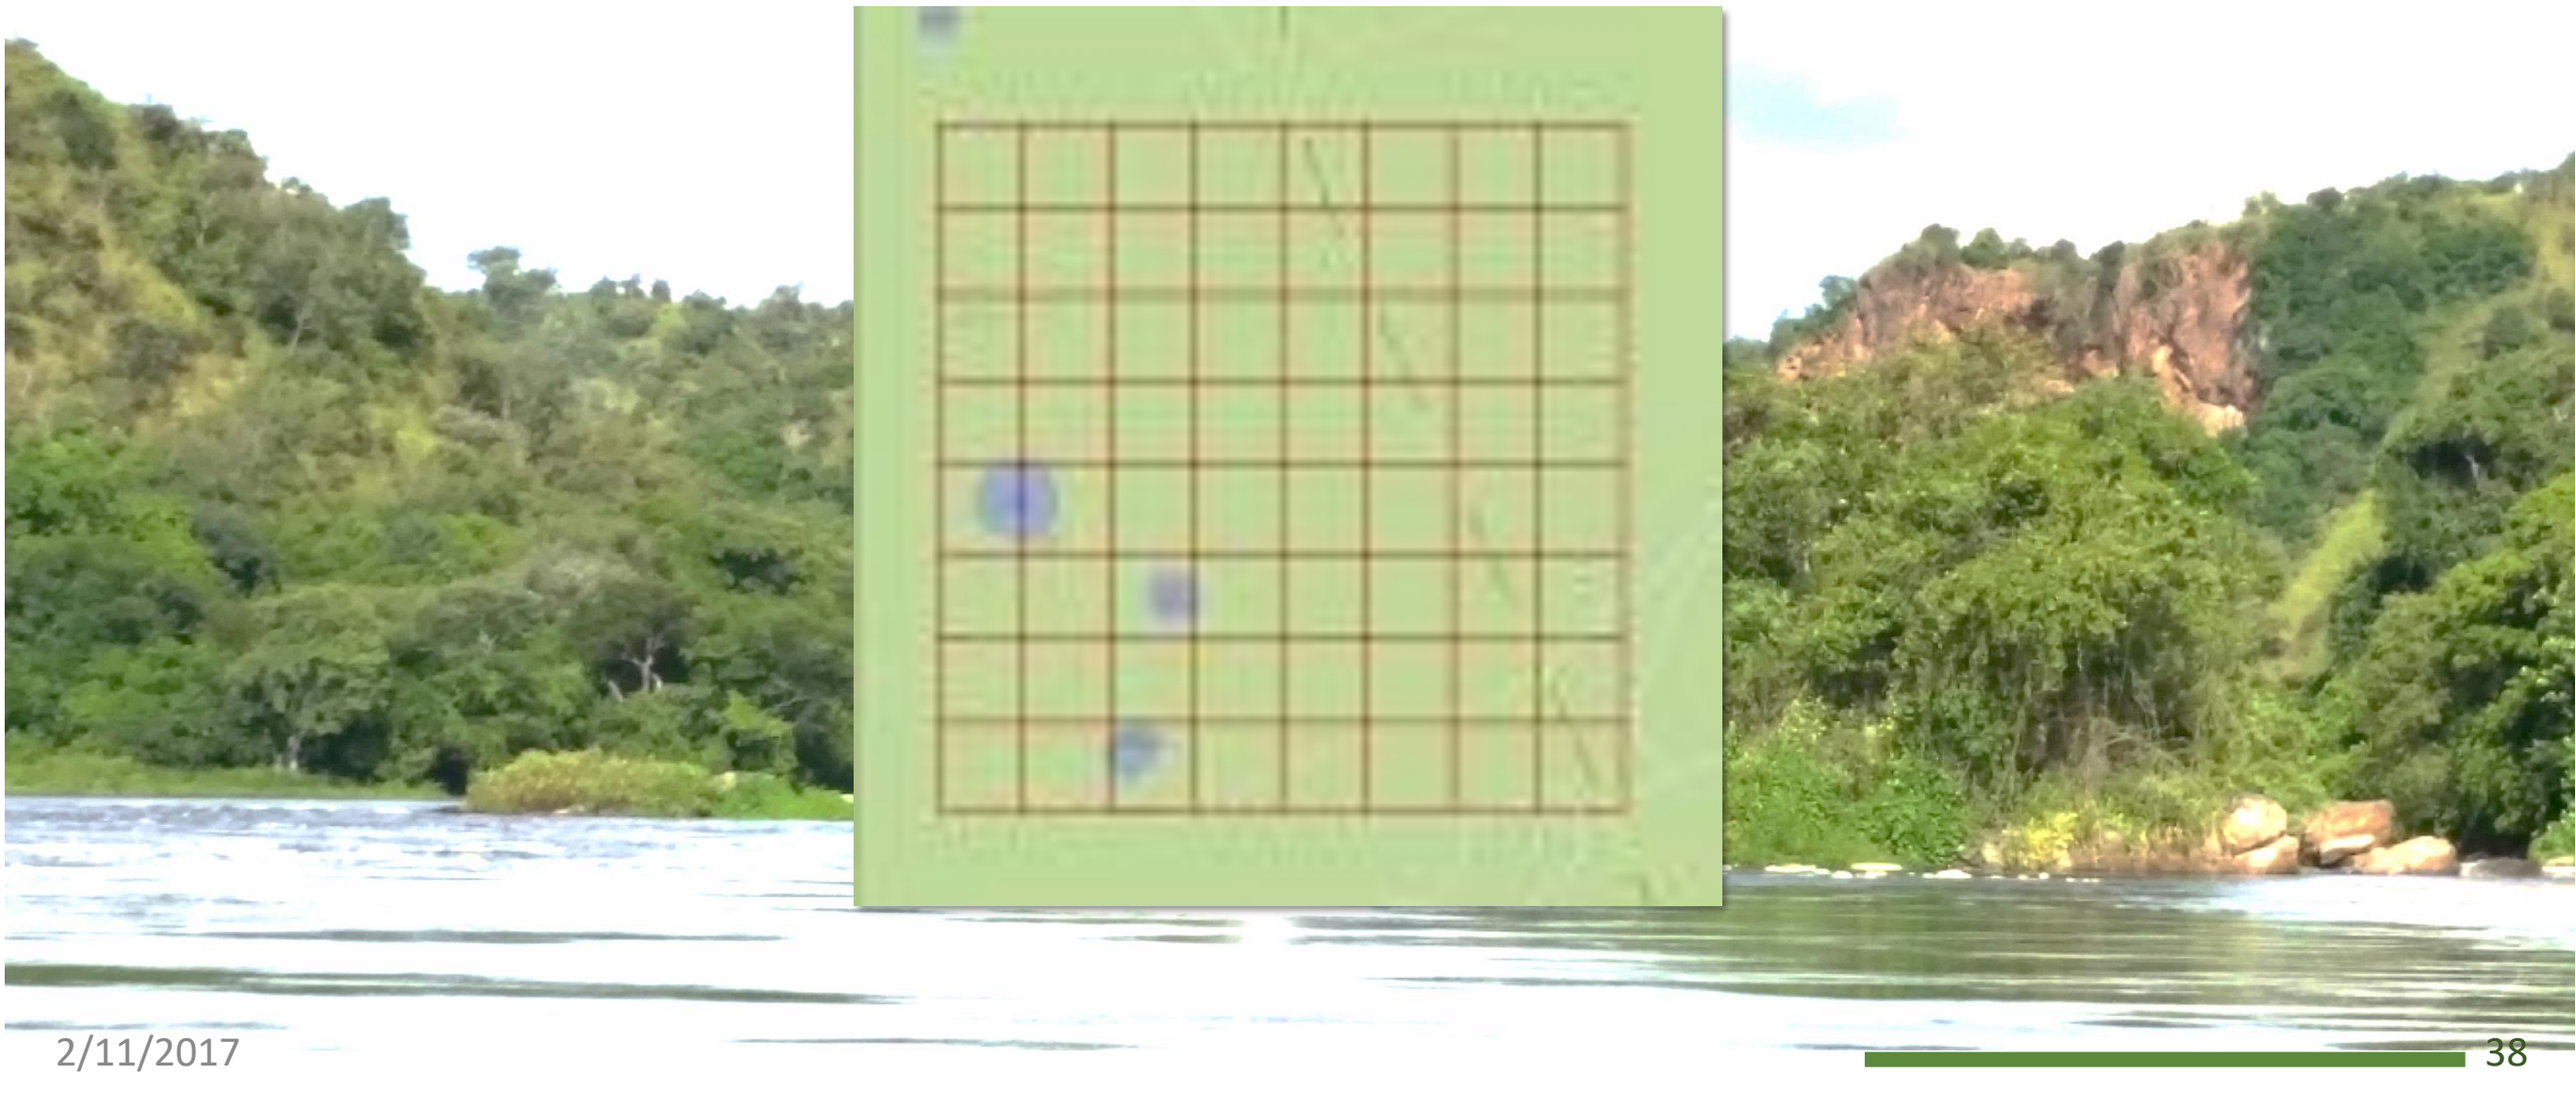

Poacher behavior prediction

Queen Elizabeth National Park, Uganda

1300 targets, 12 years of patrol data

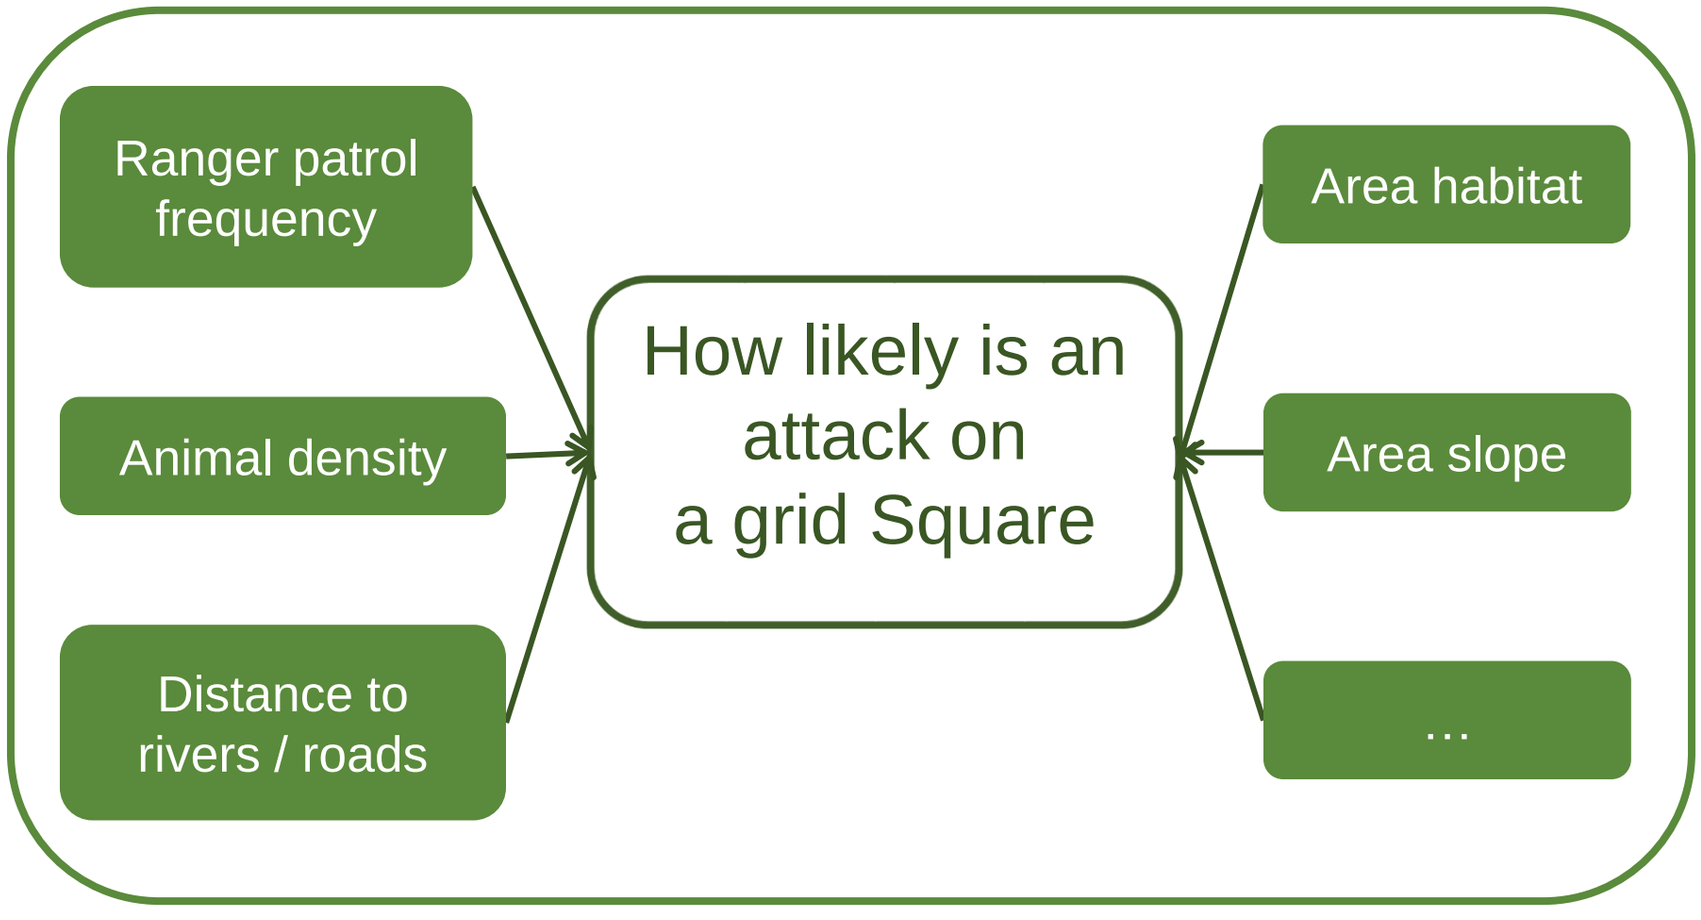

Initial Attempt Using Dynamic Bayes Net: Time Dependency & Imperfect Observation

$$\text{Adversary's probability of choosing target j} = \frac{e^{SEU^{adversary}(x, j)}}{\sum_{j'=1}^{M} e^{SEU^{adversary}(x, j')}}$$

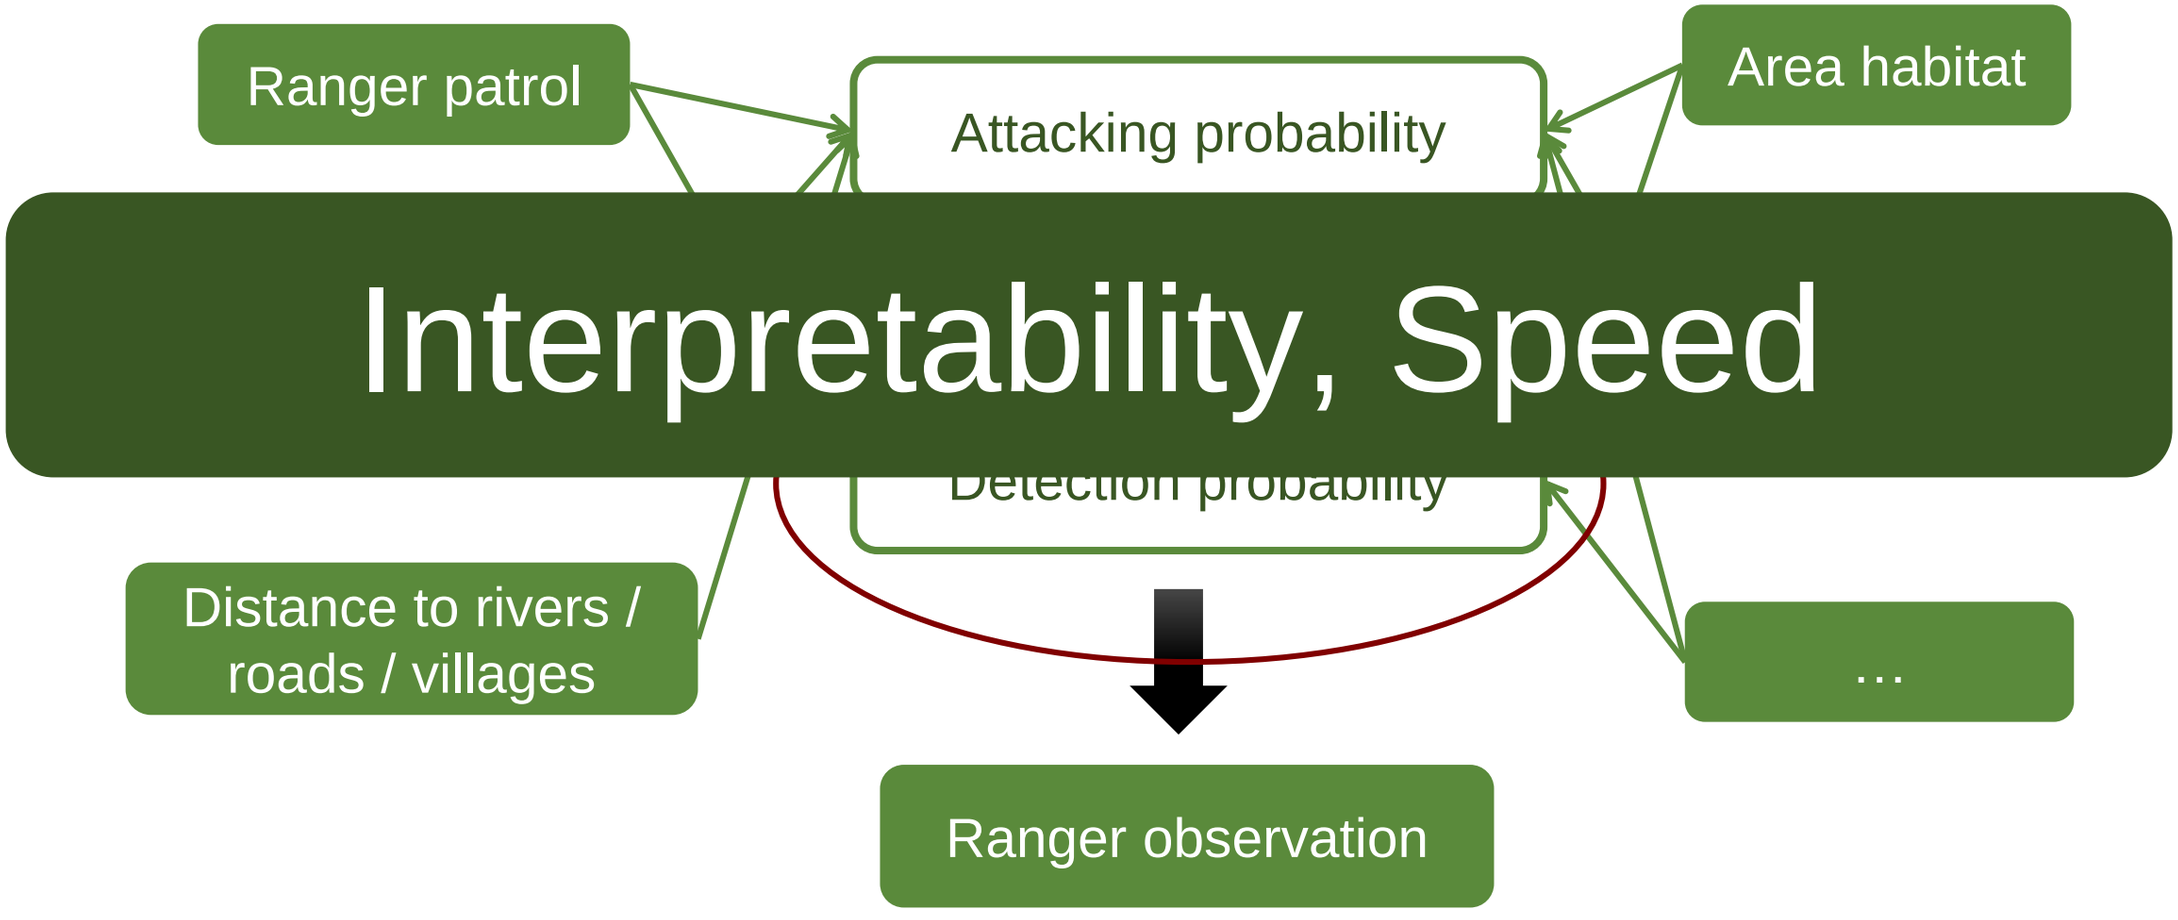

Poacher behavior prediction

Poacher Behavior Prediction

Ensemble of Decision Trees

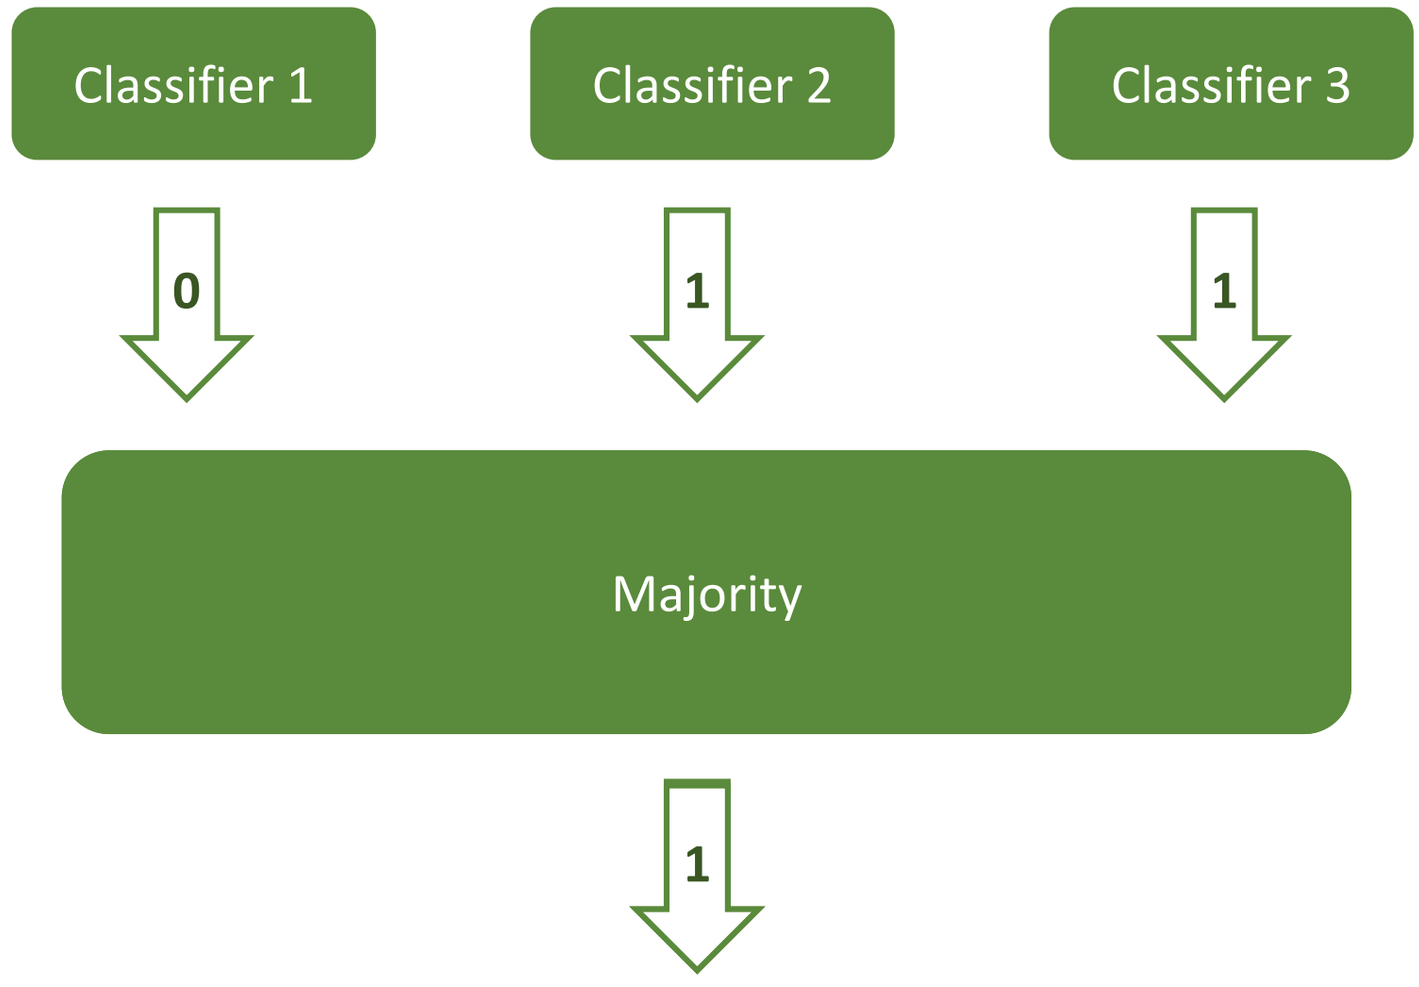

Poacher attack prediction

Poacher Behavior Prediction

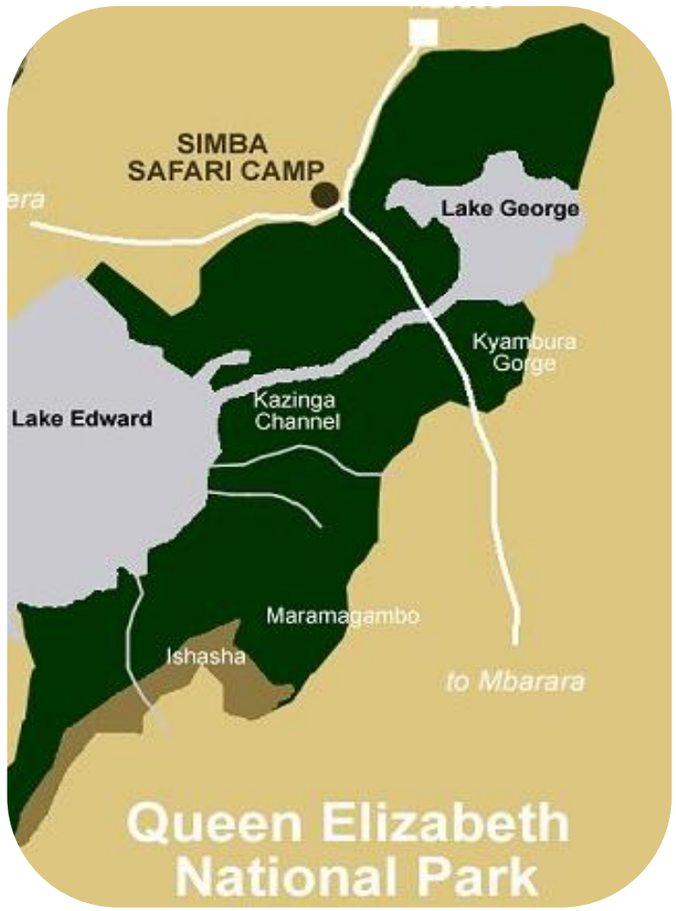

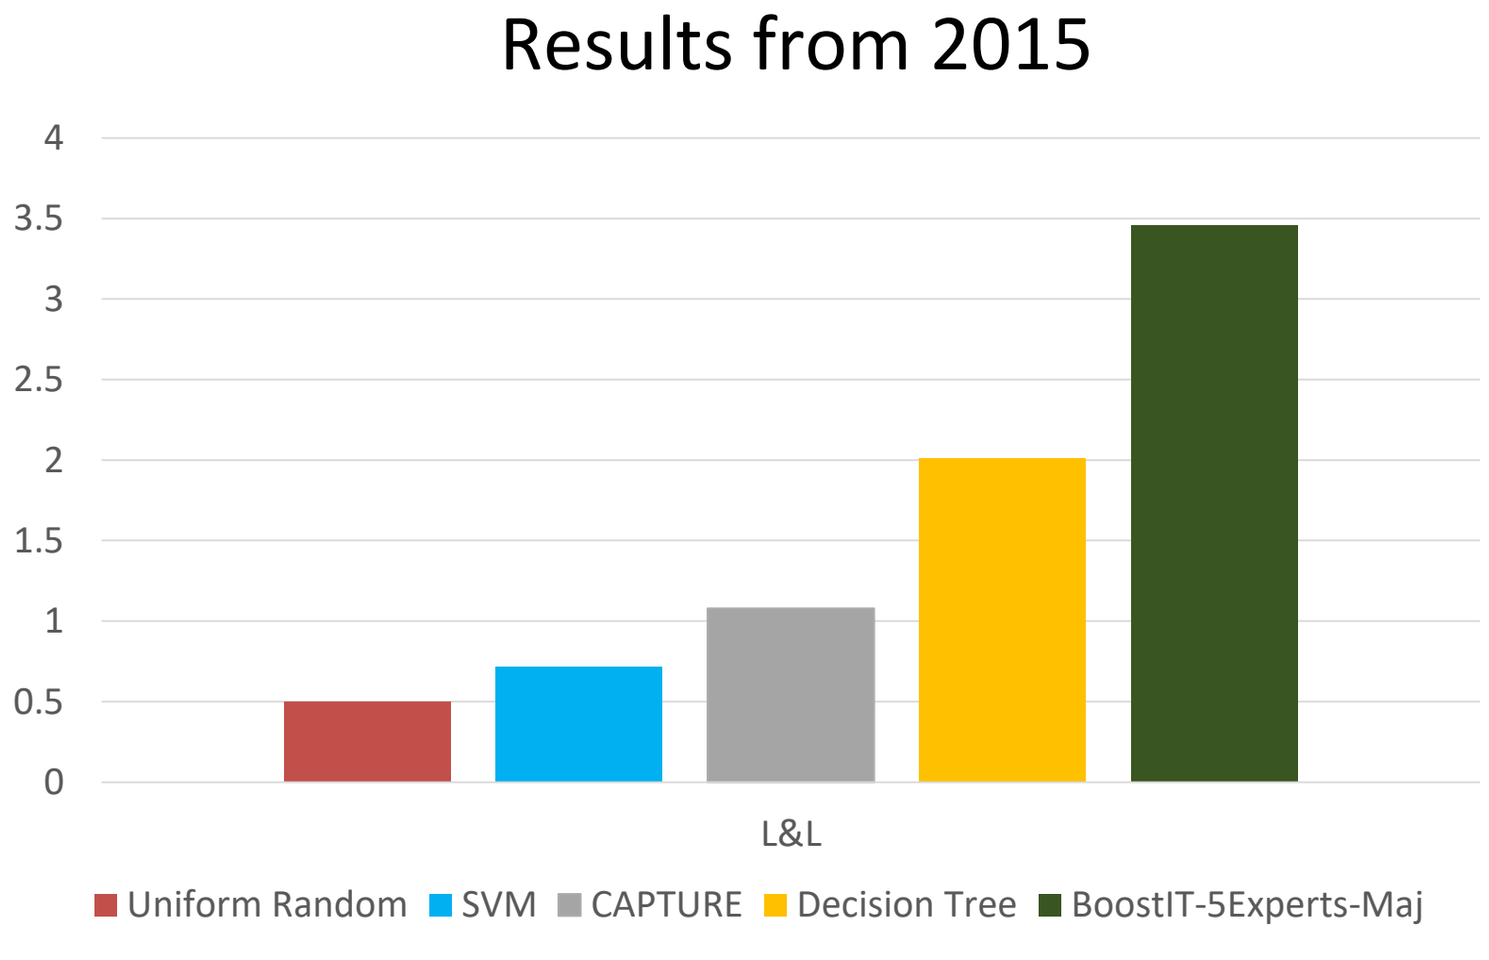

Real-world Deployment: Results

- Two 3 sq km patrol areas: Predicted hot spots with infrequent patrols

- Trespassing: 19 signs of litter etc
- Snaring: 1 active snare
- Poached Animals: Poached elephant
- Snaring: 1 elephant snare roll
- Snaring: 10 Antelope snares

- Hit rates (per month)
  - ➢ Ours outperforms 91% of months

Historical Base Hit Rate	Our Hit Rate
Average: 0.73	3

On-Going Experiments:
Queen Elizabeth National Park

- Red: Group 1 (highest attack prediction rate)
- Yellow: Group 2
- Green: Group 3

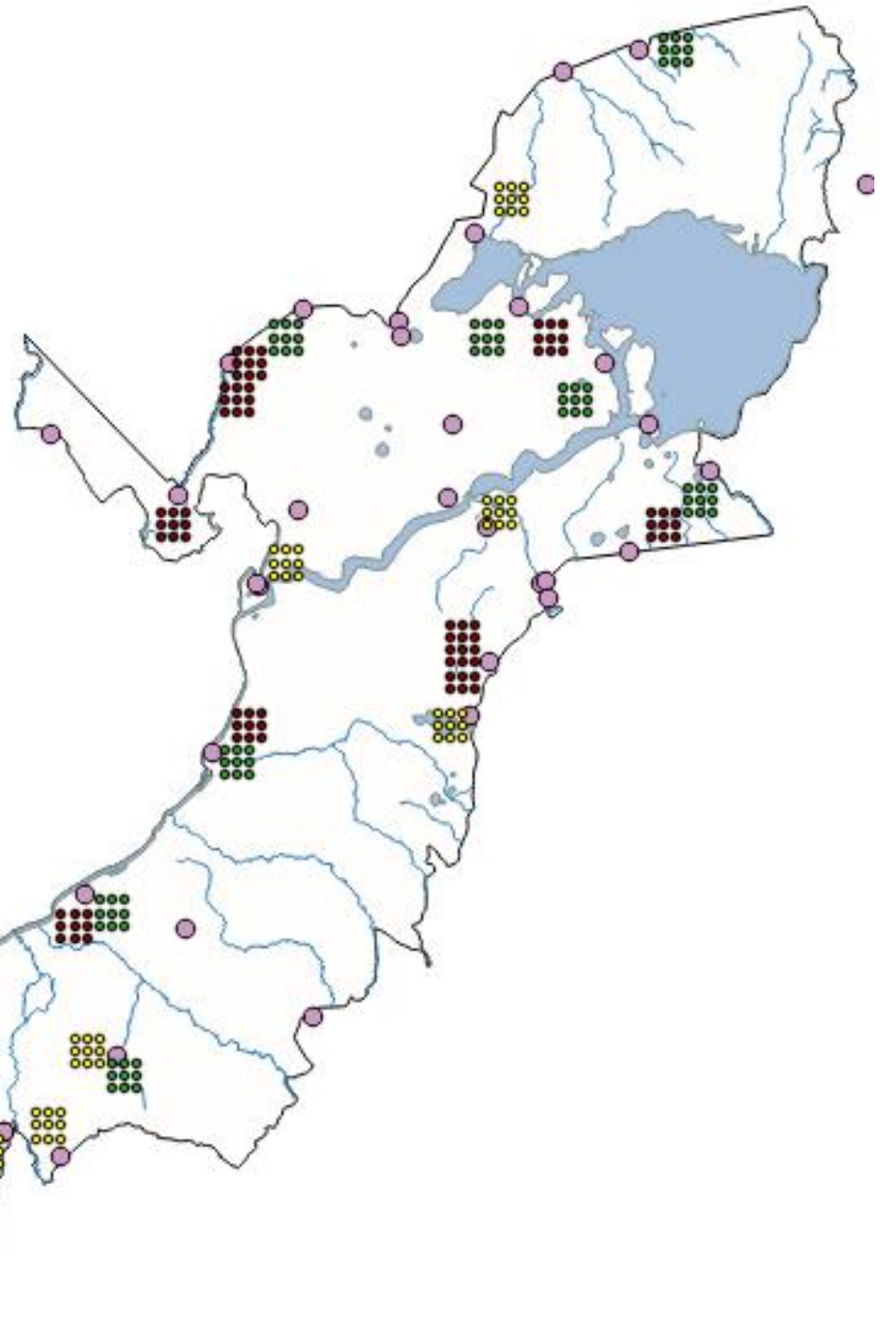

On-Going Experiments: Queen Elizabeth National Park

- Rangers followed poachers' trail; ambushed camp
  - *Arrested one (of 7) poachers*
  - *Confiscated 10 wire snares, cooking pot, hippo meat, timber harvesting tools.*

- Indirect poaching signs; pursuit of poachers
- Signs of road building, fires, illegal fishing

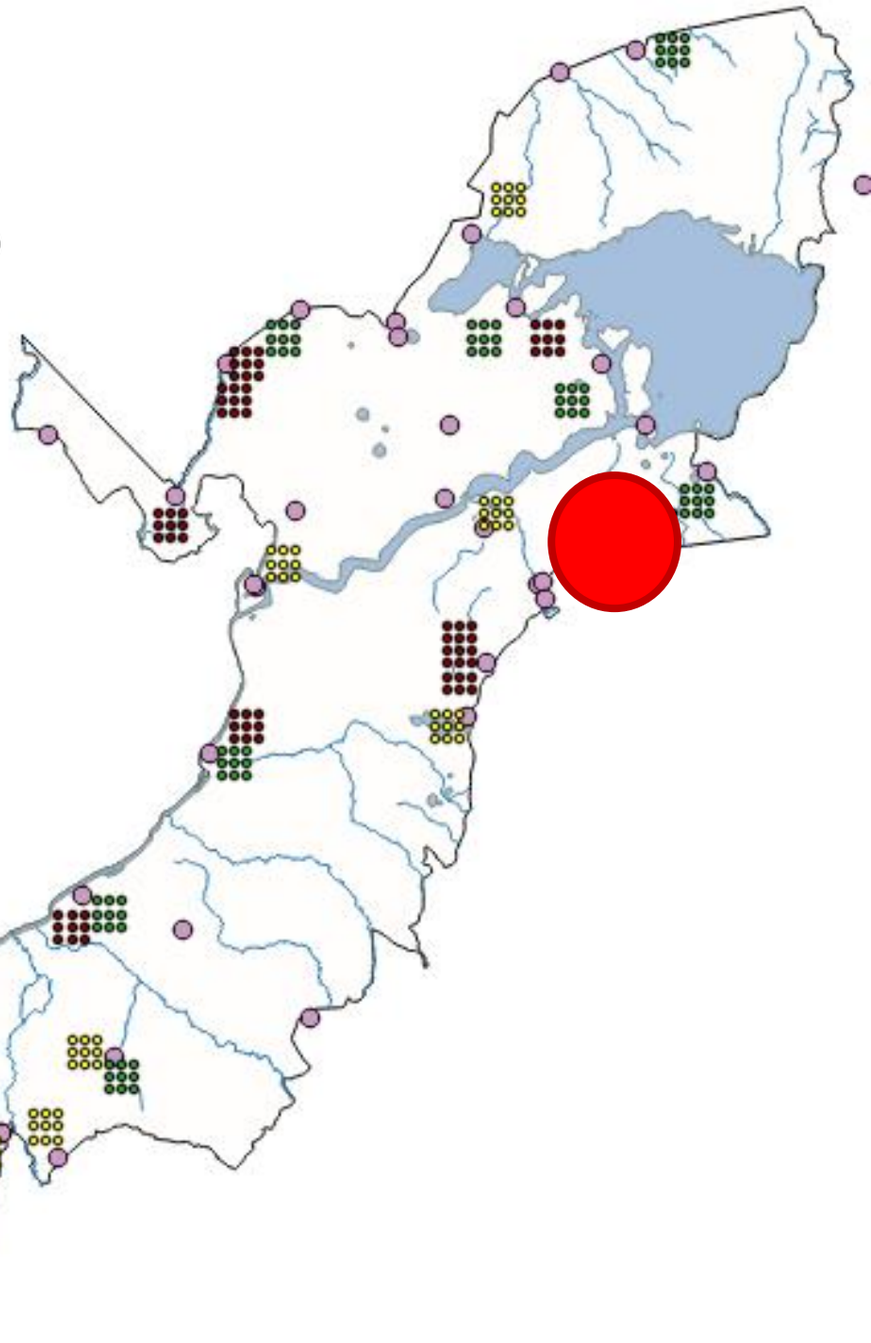

PAWS: Applying AI for protecting wildlife

Game Theory + Poacher Behavior Prediction

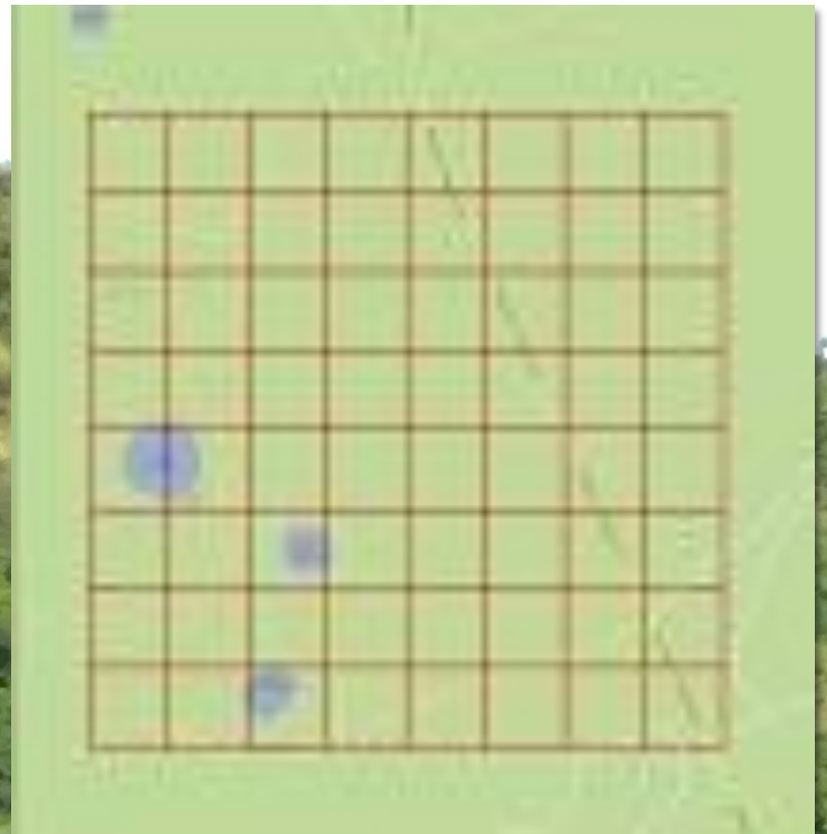

Predicting Poaching from Past Crime Data

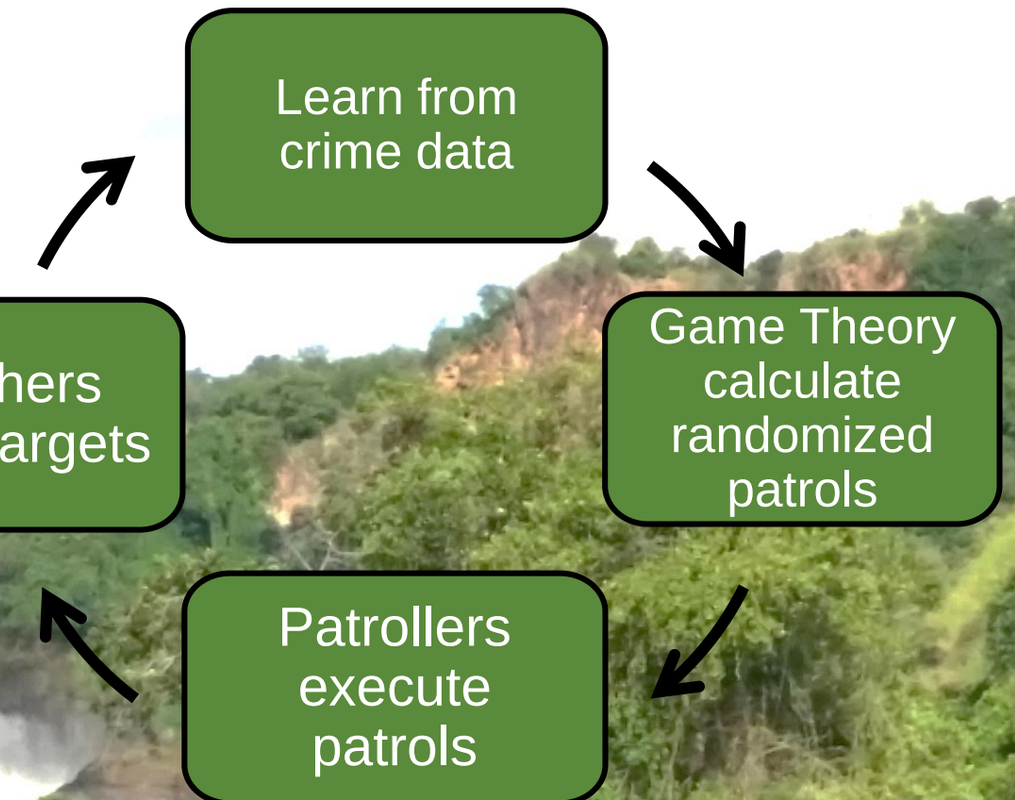

Towards the Future

- Significant potential: AI for low resource communities, emerging markets
  - Direction of AI research in our hands

- Not just applications; novel research challenges:
  - Fundamental computational challenges from use-inspired research
  - Designing AI systems in society:
    - Interpretability
    - Maintaining human autonomy

- Methodological challenges:
  - Encourage interdisciplinary research: measures impact in society

AI for Social Good

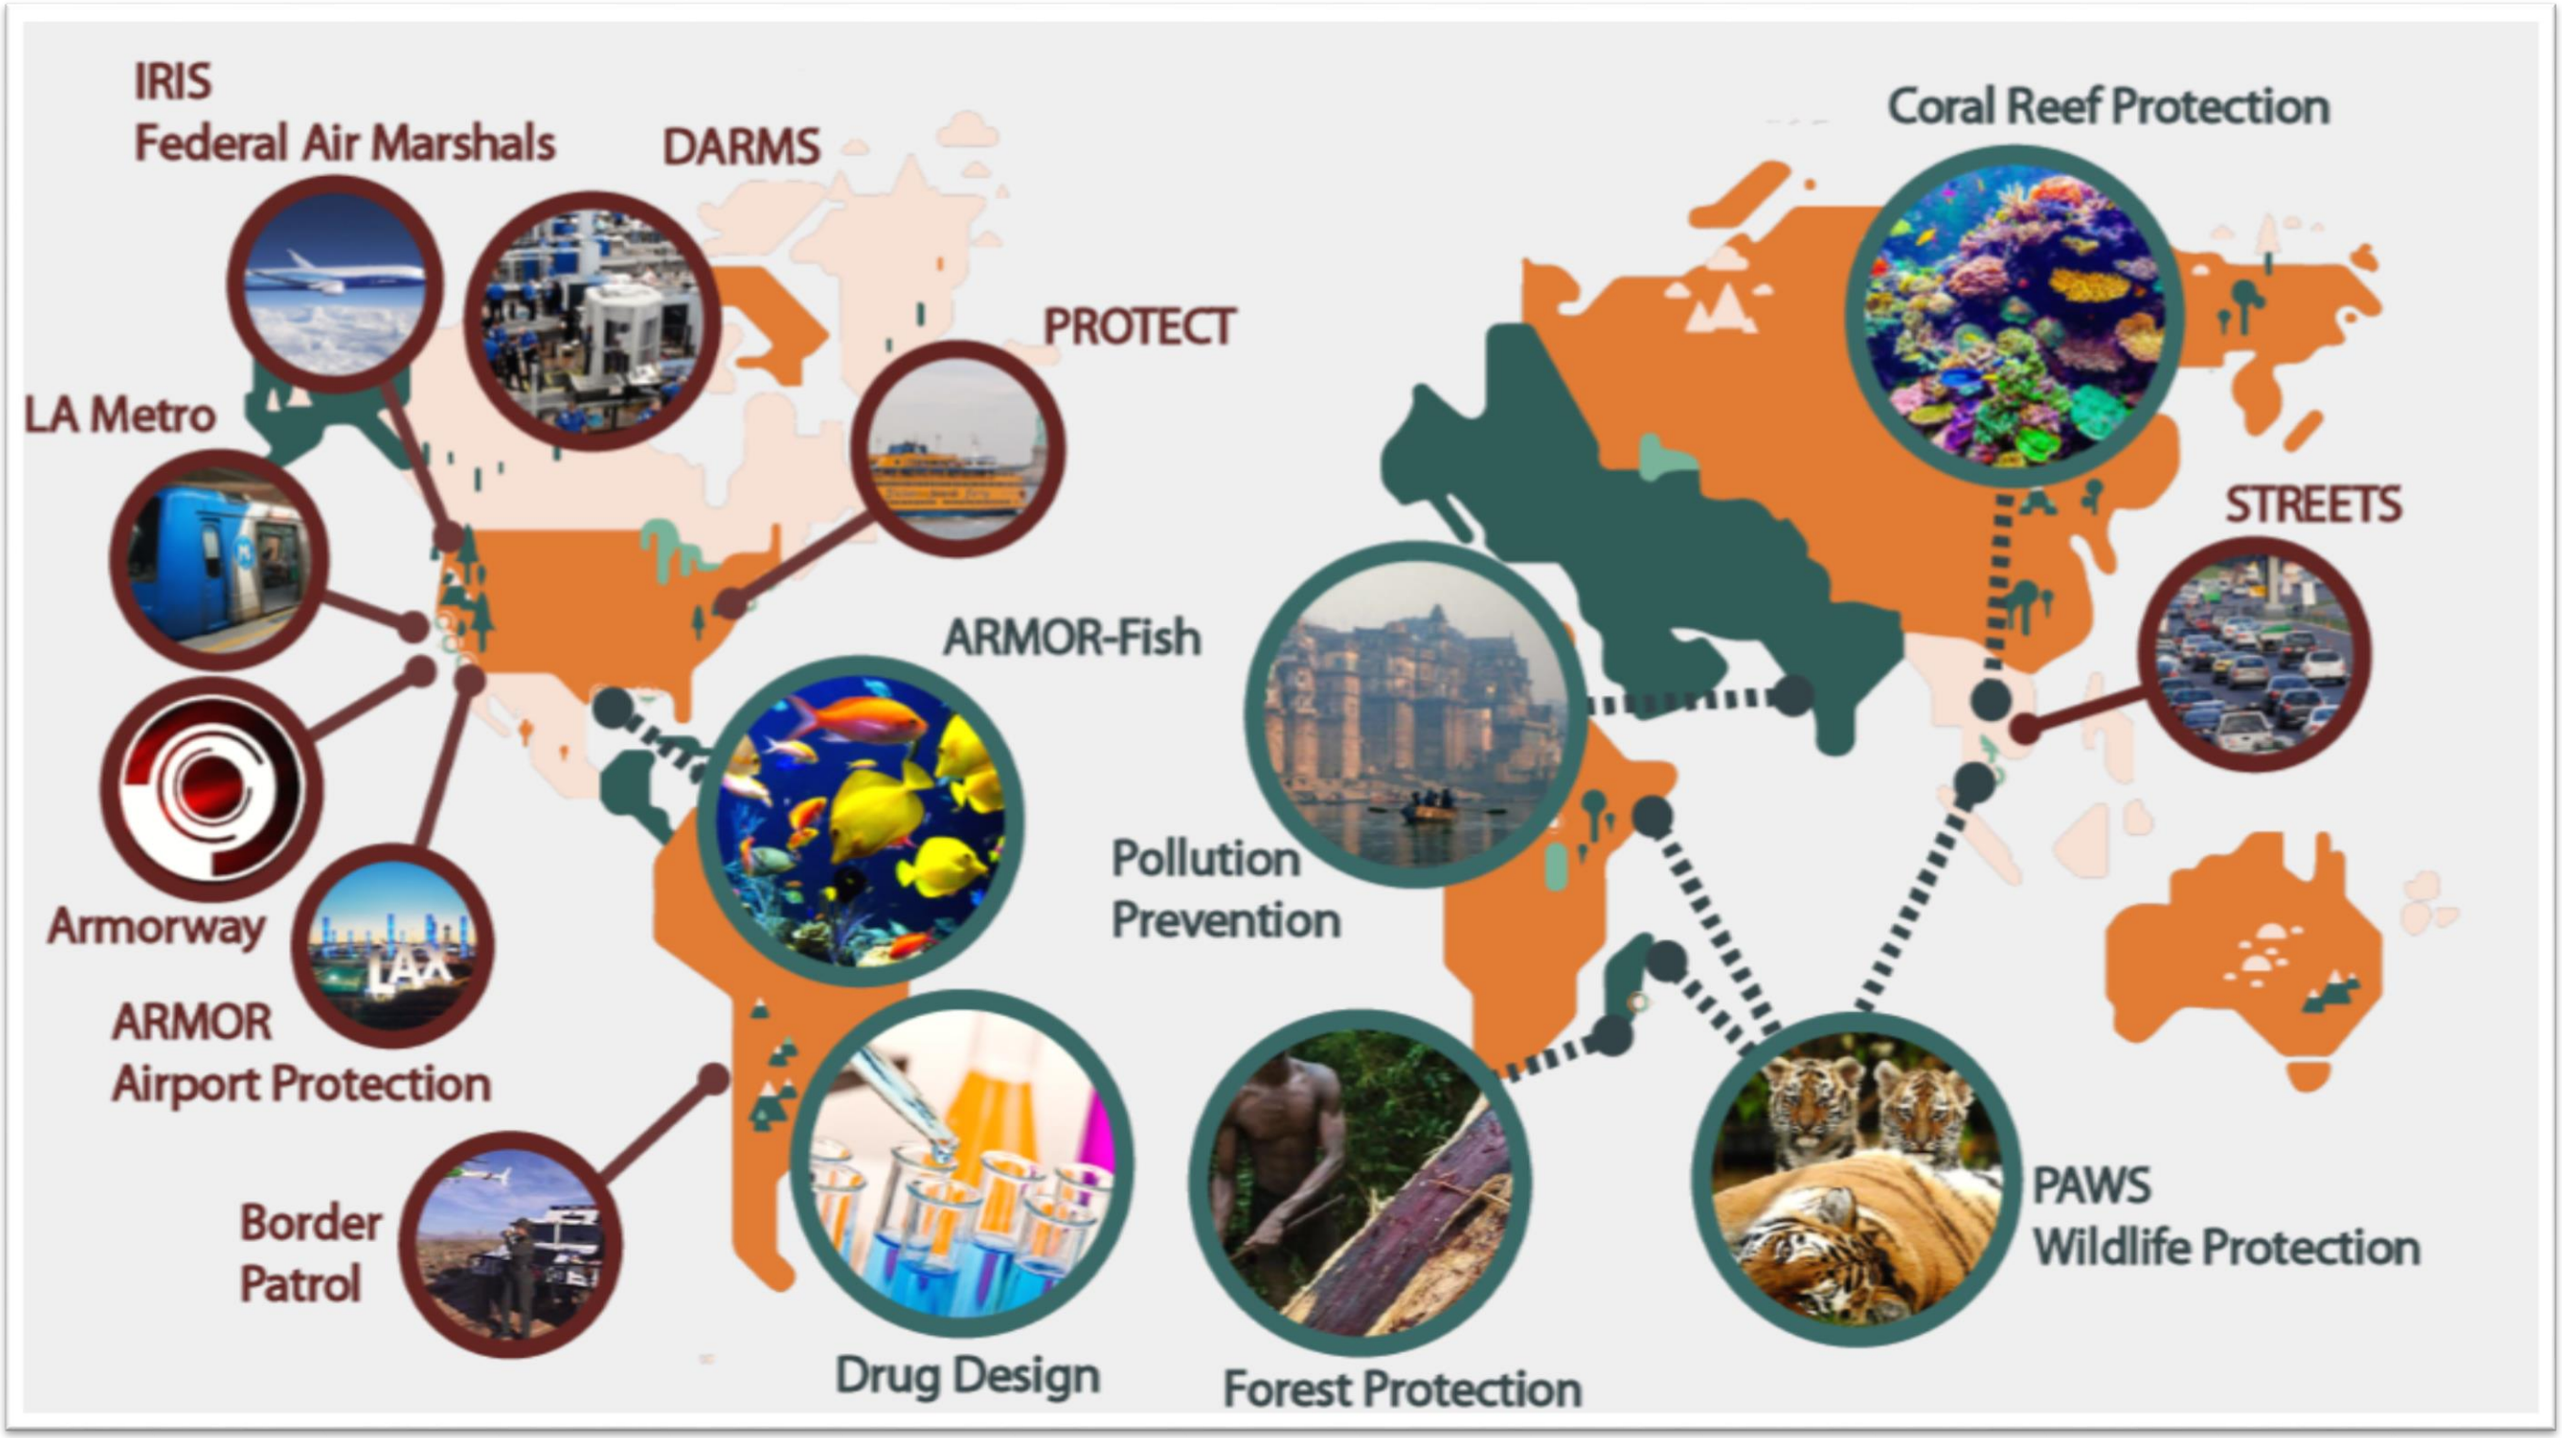

THANK YOU

tambe@usc.edu

http://teamcore.usc.edu/security

BACKUP

tambe@usc.edu

http://teamcore.usc.edu/security

POMDP Model: Challenges

- Number of states (node & edge uncertainty): ~ 2^{300}
- Number of actions (N-choose-K): > 1,000,000,000
- Number of observations (Edges exist or not): Exponential

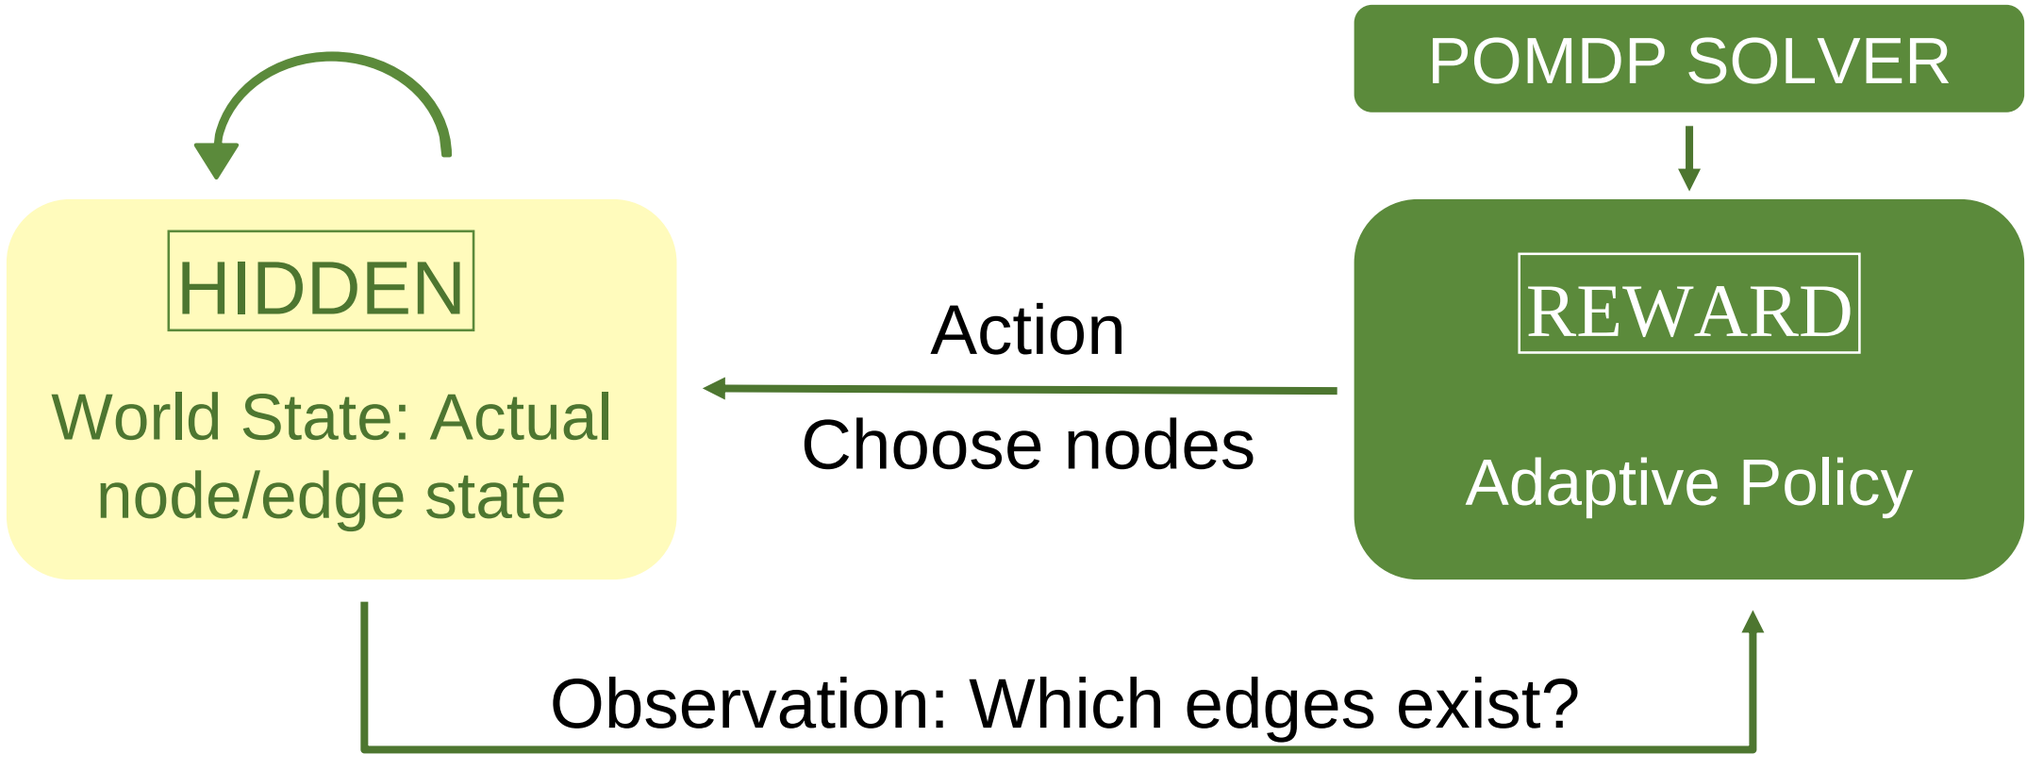

Mission Statement: Advancing AI research driven by…

Divide and Conquer

- Main idea:
  - Divide POMDP into smaller POMDPs
  - Combine solutions of smaller POMDPs

- Two different ways of dividing POMDPs
  - Using uncertain edges: PSINET
  - Using graph partitioning and sampling: HEALER

…And the Past

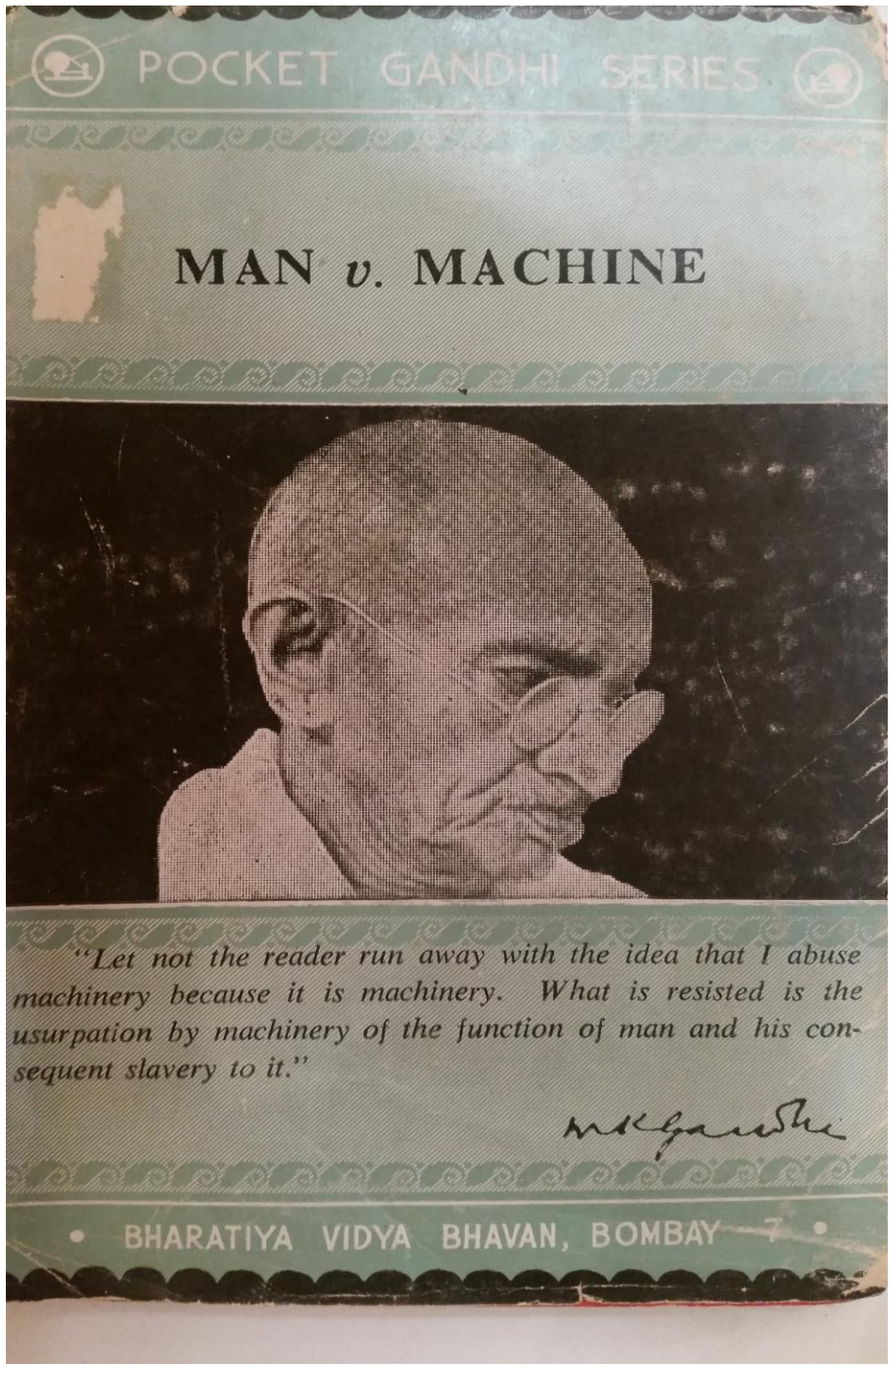

"…prize every invention of science made for the benefit of all"

Intermediate POMDP

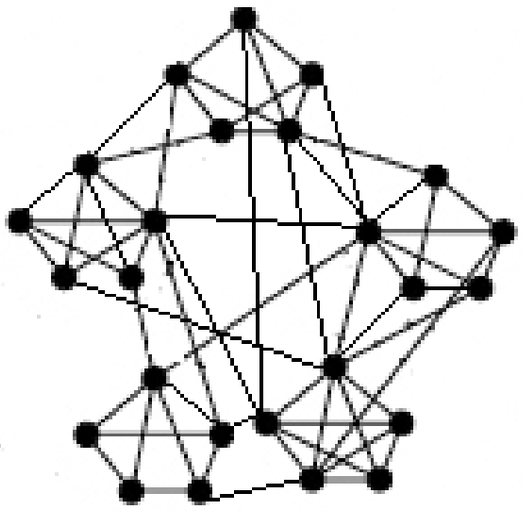

Sampled POMDP

Sampled POMDP

.
.
.

Sampled POMDP

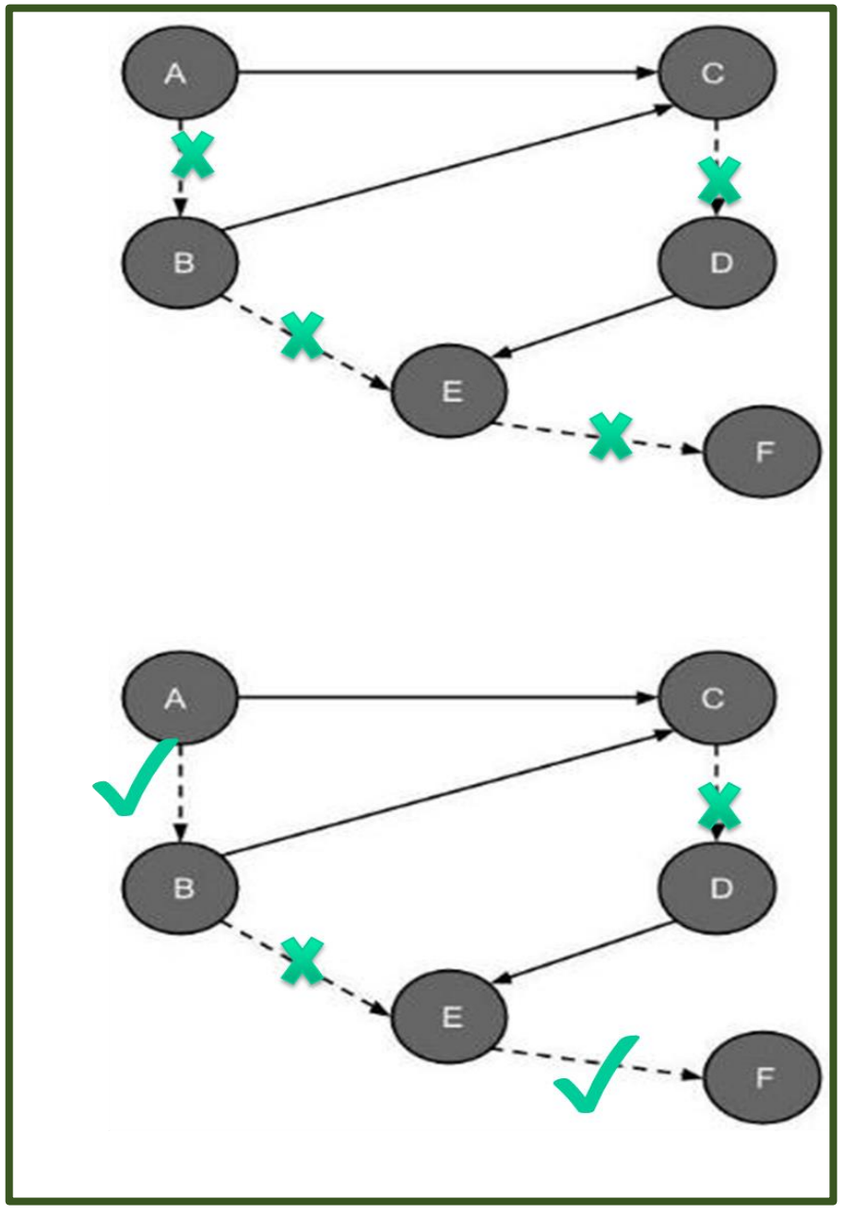

Intermediate POMDP

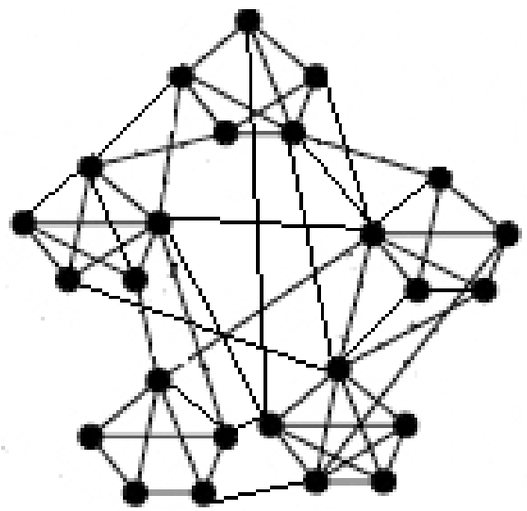

Sampled POMDP

Sampled POMDP

.
.

Sampled POMDP

Value	4	7	3.4	10	9
	1	2	3	4	5

Value	6	2	1	3.1	4.6
	1	2	3	4	5

Value	4	7	3.4	10	9
	1	2	3	4	5

Expected Value	**4.2**	**5**	**2.6**	**7**	**7.1**
	1	2	3	4	5

HEALER++ Algorithm

- Computes an equilibrium strategy for this game
- Exponentially large strategy space: incremental generation, double oracle
- Under some conditions, provably converges to approx. equilibrium

	Params #1	Params #2	Params #3
Policy #1	0.8, -0.8	0.3, -0.3	0.4, -0.4
Policy #2	0.7, -0.7	0.5, -0.5	0.6, -0.6
Policy #3	0.6, -0.6	0.4, -0.4	0.7, -0.7

Nature's oracle

	Params #1	Params #2
Policy #1	0.8, -0.8	0.3, -0.3
Policy #2	0.7, -0.7	0.5, -0.5

Influencer's oracle

	Params #1	Params #2	Params #3
Policy #1	0.8, -0.8	0.3, -0.3	0.4, -0.4
Policy #2	0.7, -0.7	0.5, -0.5	0.6, -0.6

Real networks - robustness