




Comp Sci 288: AI for Social Impact


Lecture 3: Case Studies from past 15 years of work In AI for Public Health and Conservation

MILIND TAMBE

Any questions from our last lecture?



Outline



Public Health: Influence maximization & social networks

Conservation/Wildlife Protection: Green Security Games

Information dissemination & behavior change

Optimizing Limited Intervention (Social Worker) Resources


Preventing HIV in homeless youth: Rates of HIV 10 times housed population

- **Shelters:** Limited number of peer leaders to spread HIV information in social networks
- “Real” face-to-face interactions; not Facebook etc




Influence Maximization in Social Networks


- Given:
 - Social network Graph G
 - Choose K “peer leader” nodes
 - Assume: Independent cascade model of information spread
- Objective:
 - Maximize expected number of influenced nodes



Independent Cascade Model




- Propagation Probability (for each edge)



Influence Maximization (Budget = 5 nodes)

Algorithm to select best 5 nodes?



-
- Any clarifications?


Influence Maximization in Social Networks

Three Key Challenges Combined Together

- **Uncertainty in propagation probability over edges**
- **Multi-step dynamic policies to handle peer leader “no shows”**
- **Unknown social network, limited query budget to SAMPLE network**


New research problems uncovered

Challenge 1: Uncertainty in Real-world Physical Social Networks



Robust Influence Maximization

(AAMAS 2017)



Wilder

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

$$\max_{x \in \Delta^{|P|}} \min_{\mu, \sigma} \sum x_p \frac{(\text{Outcome}(p))}{OPT(\mu, \sigma)}$$

Algorithm

Choose Peer Leaders $p \in P$
generating mixed strategy

“ $x \in \Delta^{|P|}$ ”

vs

Nature

Chooses parameters

μ, σ


Challenge 2: POMDPs for Multi-step Policy

(AAMAS 2018a)




Yadav

$K = 4$
1st time step



$K = 4$
2nd time step




HIDDEN STATE

Action
Choose nodes

POMDP
Policy

Observation: Node presence




Partition POMDPs:
Exploit community structure

Challenge 3: Sampling to avoid Data Collection Bottleneck (AAAI 2018)




Wilder

Data collection costly



Sample 20%



Sampling Algorithm

Sample node randomly
& estimate size of its
community;
Choose seeds from
largest K communities

- Input: total number of nodes, n
- Query upto query budget
- Output K peer leader nodes to spread influence
- Perform similar to OPT , best influence spread with full network


“Sampling-HEALER”



Yadav



Wilder



How do we measure impact?

“Sampling-HEALER”

Pilot tests with 230 Homeless Youth


(IJCAI 2018)



Yadav



Wilder



12 peer leaders

Sampling HEALER (Sampled Network)	HEALER (Full Network)	HEALER+ (Full Network)	DEGREE CENTRALITY (Full Network)
60 youth	62 youth	56 youth	55 youth

Results: Pilot Studies


(Journal of Society of Social Work & Research 2018)



Yadav



Wilder



Results of 750 Youth Study [with Prof. Eric Rice]


Actual Change in Behavior?

(AAAI 2021)


First large-scale application of influence maximization for public health



Reduction in condomless anal sex (1 month)




Reduction in condomless anal sex (3 months)




Results of 750 Youth Study [with Prof. Eric Rice]



LOS
ANGELES
LGBT
CENTER



What our collaborators are saying:



Smiling Youth



*“Picking youth as peer leaders was changing their **self esteem and the sense of confidence**”*

Eric Rice

Breakout Group

What might be some concerns from:

-- Fairness

“AI algorithms and datasets can reflect, reinforce, or reduce unfair biases”

“...avoid unjust impacts on people, particularly those related to sensitive characteristics such as race, ethnicity, gender, nationality, income, sexual orientation, ability, and political or religious belief.” [From Google’s AI principles]


-- Accountability, Transparency

“...design AI systems that provide appropriate opportunities for feedback, relevant explanations, and appeal. Our AI technologies will be subject to appropriate human direction and control.” [From Google’s AI principles]

[Could look up Microsoft responsible AI principles etc if you wish]

Fairness

- Typically, algorithms optimize **total welfare**
 - Expected number reached
- But how is influence distributed?
- Neglect small/less connected groups?



Algorithms for Fair Influence Maximization

Do standard influence maximization algorithms neglect minority groups?

How can disparities be avoided?

What else might be missing?

Interpretability?

Simplicity of interaction

Next steps

- Trying to get homeless drop-in centers to adopt these techniques
- On-going discussions

Lessons

Slower timelines

- First meeting with social work professor: late 2013
- Many many discussions to define a problem statements
- Actual work started 2014 summer
- Final results Summer of 2020

Bridging different academic cultures, different languages:

- What is an objective?
- Very interesting to walk from CS department to School of Social work

Health Program Adherence Maternal & Child Care in India

(Under submission)

India: Woman dies in childbirth every 15 min; 4 of 10 children too thin/short

- *ARMMAN: 18 Million women enrolled, 160000 health workers...*
- *mMitra: Weekly call to new/expectant moms; friendly 3 minute messages about health*
- *mMitra: Significant benefits shown; 2.2 million women enrolled*
- *Unfortunately, significant fraction low-listeners or drop-outs*





Passive Adherence Monitoring Maternal & Child Care in India

(with B Ravindran IIT Madras)

Predict beneficiaries likely to drop out: Allows ARMMAN to focus intervention

- Neural networks for prediction: Extensive tests with past data
- Results of pilot with 18000 beneficiaries: High precision, recall, accuracy
- **Prediction software deployed: helps 300,000 beneficiaries in mMitra**



ARMMAN Pilot

- 18000 Beneficiaries
- Nov & Dec 2019
- Test: Jan-April 2020


Passive Adherence Monitoring Preventing Tuberculosis in India

(KDD 2019)




Tuberculosis (TB): ~500,000 deaths/year, ~3M infected in India

- *Non-adherence to TB Treatment: Digital adherence tracking via daily phone calls*
- *Intervene before patients miss dose*
- *Predict adherence risk from phone call patterns?*
- *Results from Mumbai, India: 15,000 patients, 1.5 million phone calls*



Outline

Public Health: Influence maximization & social networks



Conservation/Wildlife Protection: Green Security Games


Poaching of Wildlife in Uganda

Limited Intervention (Ranger) Resources to Protect Forests

Snare or Trap




Wire snares



Focus on Threat prediction


Diving up the area into 1 Km x 1 Km Grid cells



Learning Poacher Model: Uncertainty in Observations



Nguyen Gholami



What uncertainty might you expect as rangers walk to recover snares and report that data?


Learning Poacher Model: Uncertainty in Observations



Nguyen



Gholami



Positive and unlabeled datasets




Learning Adversary Response Model: Uncertainty in Observations






Nguyen




Gholami



Training: Filtered Datasets



Predict: Ensemble of Classifiers



PAWS: First Pilot in the Field

(AAMAS 2017)



Ford






Gholami

- Two 9-sq.km areas, infrequent patrols



- Poached elephant
- 1 elephant snare roll
- 10 Antelope snares




PAWS Predicted High vs Low Risk Areas: 3 National Parks, 24 areas each, 6 months (ECML PKDD 2017, ICDE 2020)




Xu




Gholami



Queen Elizabeth National Park




Murchison Falls National Park




Srepok Wildlife Sanctuary


Snares per patrolled sq. KM



Snares per patrolled sq. KM



Snares per patrolled sq. KM



PAWS Real-world Deployment

Cambodia: Srepok Wildlife Sanctuary

(ICDE 2020)



Xu



- *521 snares/month our tests*
- VS
- *101 snares/month 2018*

PAWS GOES GLOBAL with SMART platform!!



**Protect Wildlife
800
National Parks
Around the Globe**

Also: Protect Forests, Fisheries...

Questions

- What assumptions do we need to make PAWS predictions work in general?
- Why not just ask rangers to go to hotspots?
- Is measuring by snare captures sufficient?

Questions

- What assumptions do we need to make PAWS predictions work in general?

Availability of sufficient, diverse data

- Why not just ask rangers to go to hotspots?


Game theory

- Is measuring by snare captures sufficient?

How do we measure overall impact?

Green Security Games

(IJCAI 2015)




- *Stackelberg security games (SSG)*
- *With boundedly rational poachers*
- *Learn adversary response model at targets “i”*




	Area1	Area2
Area1	4, -3	-1, 1
Area2	-5, 5	2, -1

Overall model




Sometimes
AI is not
the right
solution



Visiting Malaysia was important!

Feedback: Paths Difficult to Follow




Understand Path Following Difficulty: One-day Patrol In the Forest

- 8-hour patrol in April 2015
- 8 hours = **2** miles
- A typical day of patrollers: **6** miles





Our Patrol in Malaysia



PAWS: PROTECTION ASSISTANT FOR WILDLIFE SECURITY [2016]

Game Theory + Poacher Behavior Prediction + Forest Street Map






Indonesia: Protected Area disappeared




Lack of Data



Xu



Srepok, Cambodia
43,269 patrol observations
2013 – 2018



Royal Belum, Malaysia
824 patrol observations
June – August 2018


1000km
ees

Bandit Algorithms



Xu

Data-rich parks: build predictive models to plan patrols




[Gholami et al., AAMAS-18,
Xu et al., ICDE-20]

Data-scarce parks: conduct patrols to detect illegal activity and improve the predictive model

exploitation

exploration



LIZARD: Multiarmed Bandit

Lipschitz Arms with Reward Decomposability (AAAI 2021)



Xu

- Input: N Targets with features, T Time Horizon
- Stochastic adversary, who places snares at targets
- Patrolling algorithm: Specify patrol effort in each target up to budget B
- Reduce regret wrt OPT , optimal patrol effort, for capturing snares




Lizard exploits decomposability, smoothness, monotonicity





Green Security Games: Integrating Real-Time “SPOT” Information (IAAI 2018)



Bondi



Goal: automatically find poachers



Discussion

Community involvement