Reminder

- PRA6 due 4/16
- HW6 due 4/25
- Course project presentation 4/23 and 4/25
  - Schedule is posted
  - Come to OH for discussions!
- Course project final report due 5/2

Artificial Intelligence Methods for Social Good

Lecture 24: Common Challenges in AI for Social Good Projects

17-537 (9-unit) and 17-737 (12-unit)

Fei Fang

feifang@cmu.edu

Outline

- Typical Pipeline for AI for Social Good Projects
- Common Challenges and Practical Guide by Stage
  - Problem Formulation
  - Method Development
  - Evaluation & Deployment
- Q&A + Discussion

Typical Pipeline for AI for Social Good Projects

- What are the steps need to be taken to work on an AI project aimed for social impact?

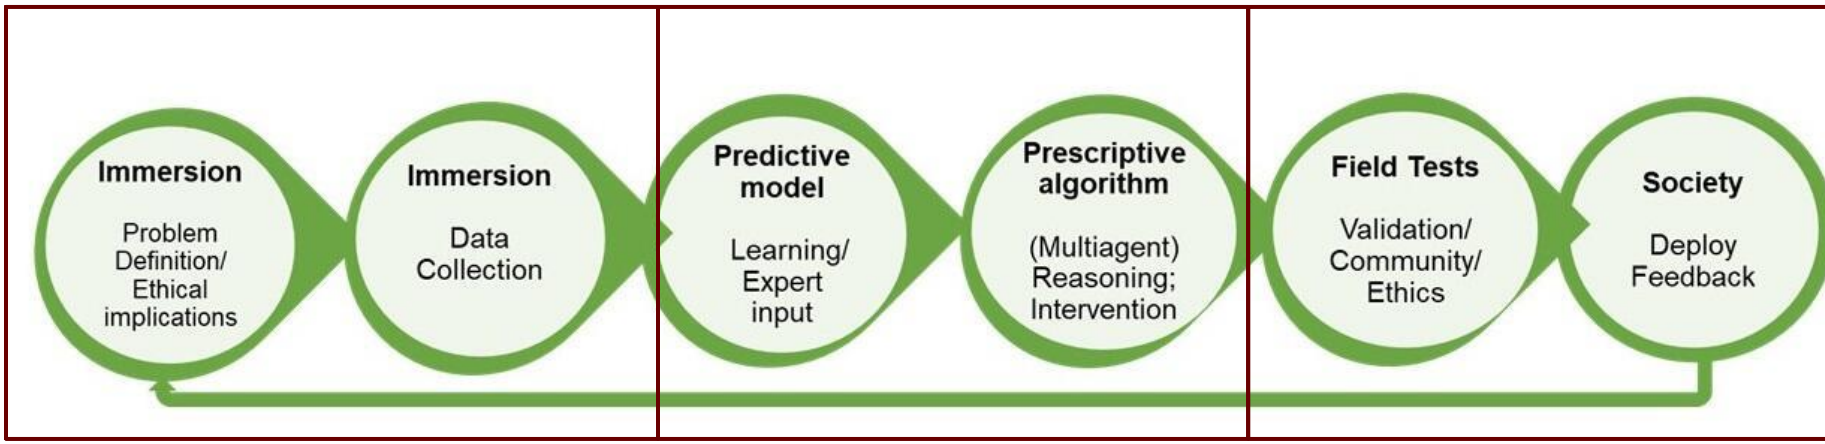

Problem Formulation | Method Development | Evaluation and Deployment

Typical Pipeline for AI for Social Good Projects

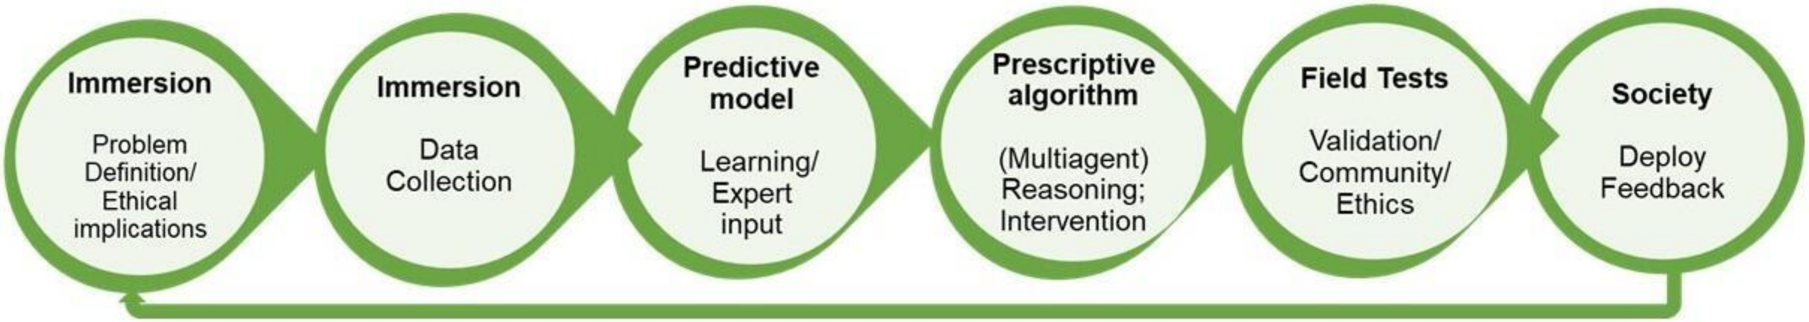

- Immersion in the domains (Problem Definition)
  - Crucial to get a critical understanding of the problems, constraints and datasets
  - Goal: come up with a clearly-defined problem

Typical Pipeline for AI for Social Good Projects

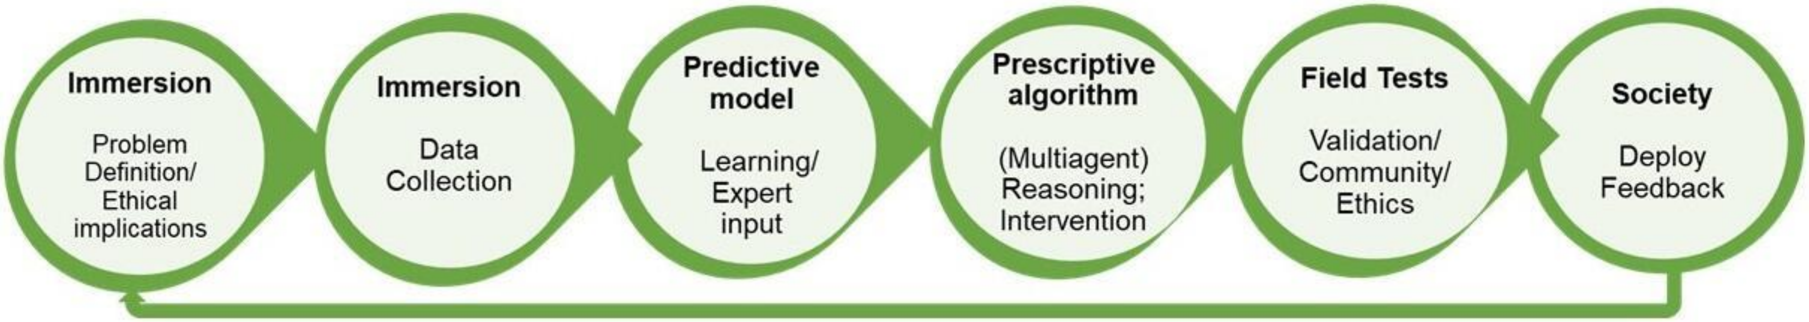

- Immersion in the domains (Data Collection)
  - Collect and clean the data needed

Typical Pipeline for AI for Social Good Projects

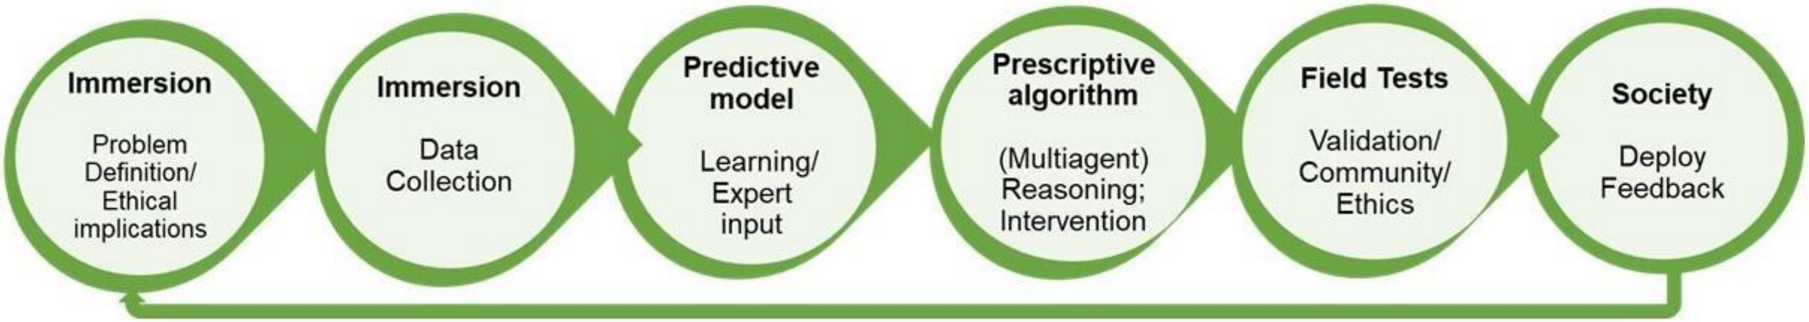

- Build a predictive model
  - Using machine learning or domain expert input

Typical Pipeline for AI for Social Good Projects

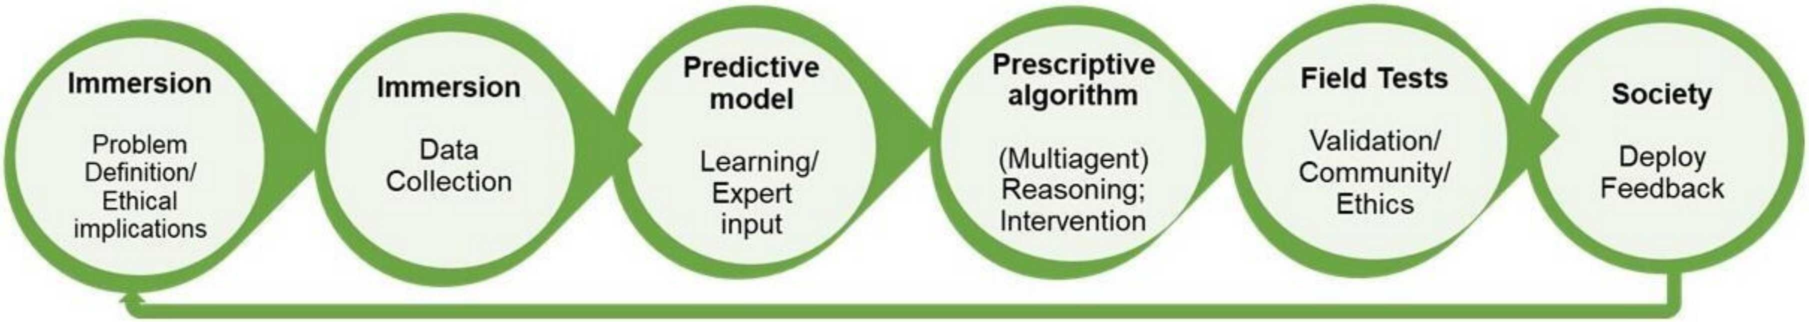

- Develop a prescriptive algorithm
  - Assist decision making
  - Suggest actions to take

Typical Pipeline for AI for Social Good Projects

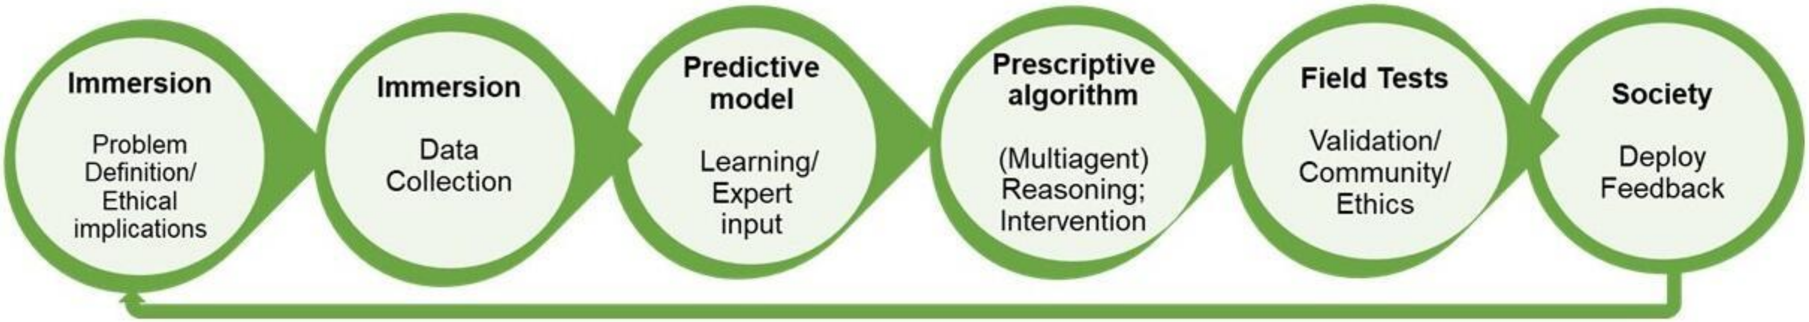

- Evaluate our models and algorithms in the field
- Learn key limitations of our models and algorithms
- Improve the models and algorithms
- Prepare for larger-scale deployment

Typical Pipeline for AI for Social Good Projects

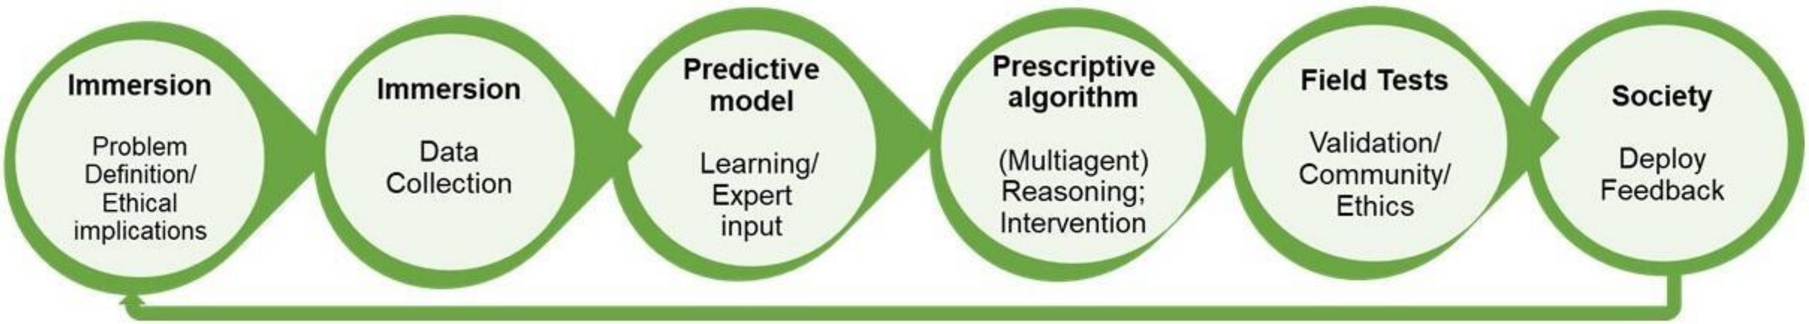

- Sustainably deploy the AI-based system
- Collect feedback and measure long-term impact on society

Typical Pipeline for AI for Social Good Projects

- What are challenges in these stages and how to tackle them?

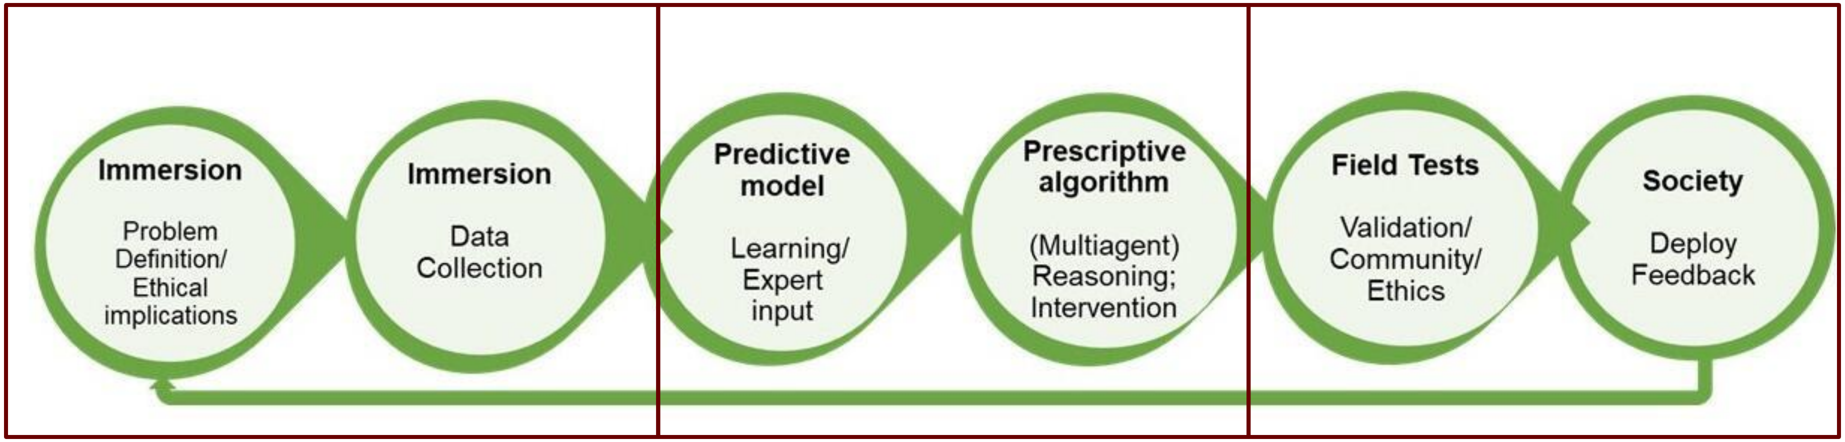

Problem Formulation | Method Development | Evaluation and Deployment

Outline

- Typical Pipeline for AI for Social Good Projects
- Common Challenges and Practical Guide by Stage
  - Problem Formulation
  - Method Development
  - Evaluation & Deployment
- Q&A + Discussion

Common Challenges in Problem Formulation Stage

- How to find problems to work on?
- How to formulate them as problems that AI can solve?
- How to get the data needed?

How to Find Problems to Work on?

- Choose a domain that one is familiar with and some existing efforts that can have a significant social impact
- Identify pain points in the current practice: steps that currently
  - (1) heavily rely on human experience
  - (2) requires a huge amount of human efforts
  - or (3) is done in an ad-hoc way but is crucial to the outcomes

How to Find Problems to Work on?

- Identify pain points that can be tackled with AI
  - Understand what AI is capable of doing in general:
    - Prediction and estimation
    - Clustering
    - Suggest actions or facilitate decision-making
    - Generate content (text, image) given instructions
    - …
  - Review existing literature to understand how AI has been used to tackle pain points

How to Find Problems to Work on?

- Envision the benefit and impact an AI-based tool can bring
  - Save human hours, less cognitive burden, better understanding, better outcomes

- Think about the cost and ethical implications of using an AI-based tool
  - Cost of training humans, negative outcomes due to misuse

How to Find Problems to Work on?

- For AI researchers: Immersion in the domains is important
  - Crucial to get a critical understanding of the problems, constraints, and datasets
  - Build interdisciplinary partnerships
  - Understand the challenges from the perspective of domain experts
  - Reach out to stakeholders actively (via emails, phone calls, and remote meetings…)
  - Discuss with stakeholders, including the impacted community
  - May cite domain experts' words with their consent in the publication

Example: Food Rescue

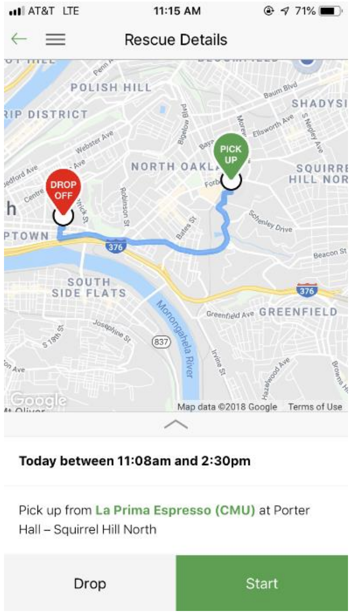

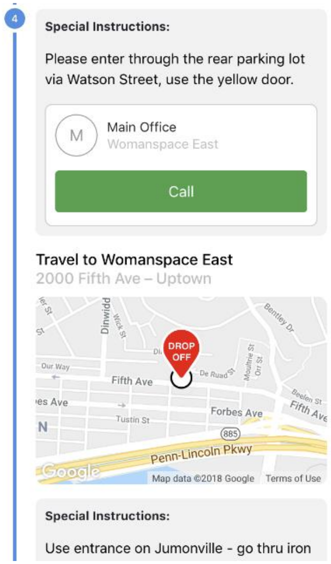

How to Formulate Them as Problems that AI can Solve?

- Determine what type of problem it is
  - Prediction and estimation
  - Clustering
  - Suggest actions or facilitate decision-making
  - Generate content (text, image) given instructions
- Choose the corresponding formulation

How to Get Data Needed?

- Determine what data is readily available
  - Data shared by collaborators
- Determine what kind of data is needed and how much is needed
- Investigate what data can be collected?
  - Collect data through e.g., crowdsourcing, human subject experiments
  - Get data from publicly available sources

How to Get Data Needed?

- Some publicly available data sources
  - Landsat-9 (USGS)
  - Google Earth Engine (Google)
  - Earthdata (NASA)
  - Google Public Data Explorer (Google)
  - AWS Open Data Registry (Amazon)
  - Global Health Observatory data (WHO)
  - World Bank Open Data (World Bank)

How to get the data needed?

- Preprocess the data
  - Understand data limitations and check if necessary to preprocess the data
  - Typical limitations that can be mitigated through preprocessing:
    - Missing entries
      - Discuss with domain experts to see if there is a way to interpolate the missing data
    - Noise in data
      - Discuss with domain experts to understand what kind of noise exist and whether it is possible to denoise

Example: Wildlife Corridor Design

- Acquisition cost
  - Tax records
  - Information on conserved lands
  - Other information: water body, urban parcel, etc

- Resistance
  - Geographical information and other landscape features
    - Grizzly bears: vegetation, human development, road density
    - Wolverines: snow cover, housing development, forest edge

Common Challenges in Problem Formulation Stage

- How to find problems to work on?
- How to formulate them as problems that AI can solve?
- How to get the data needed?
- Discussion: Any other challenges (in problem formulation stage) you faced or would like to learn more about?

Outline

- Typical Pipeline for AI for Social Good Projects
- Common Challenges and Practical Guide by Stage
  - Problem Formulation
  - Method Development
  - Evaluation & Deployment
- Q&A + Discussion

Common Challenges in Method Development Stage

- How to choose or develop the right AI method for the problem, while accounting for domain-specific considerations (e.g., practical constraints on computing resources or runtime, uncertainties and noise)?

- Decompose the problem into smaller tasks, choose the right AI method for each task based on the type of the task
  - Prediction task: Predict or estimate certain values
  - Prescription task: Suggest actions or facilitate decision-making

How to choose or develop the right AI method for the problem?

- Prediction task
  - Select features based on available data, intuition and discussion with domain experts
  - Start with existing Machine Learning methods based on data type
    - Basic discrete or continuous-valued data: Linear Regression, Random Forest, Gaussian Process, Neural Network, XGBoost
    - Image data: ResNet+refinement
    - Text data: BERT/RoBERTa/GPT3+refinement, leverage ChatGPT API
    - Graph data: Graph neural networks
  - Train and evaluate the model
    - Metrics for regression or classification
    - Domain-specific metrics

Example: Food Rescue

- Model: neural network
- Evaluate:
  - AUC
  - Domain specific metric: hit ratio

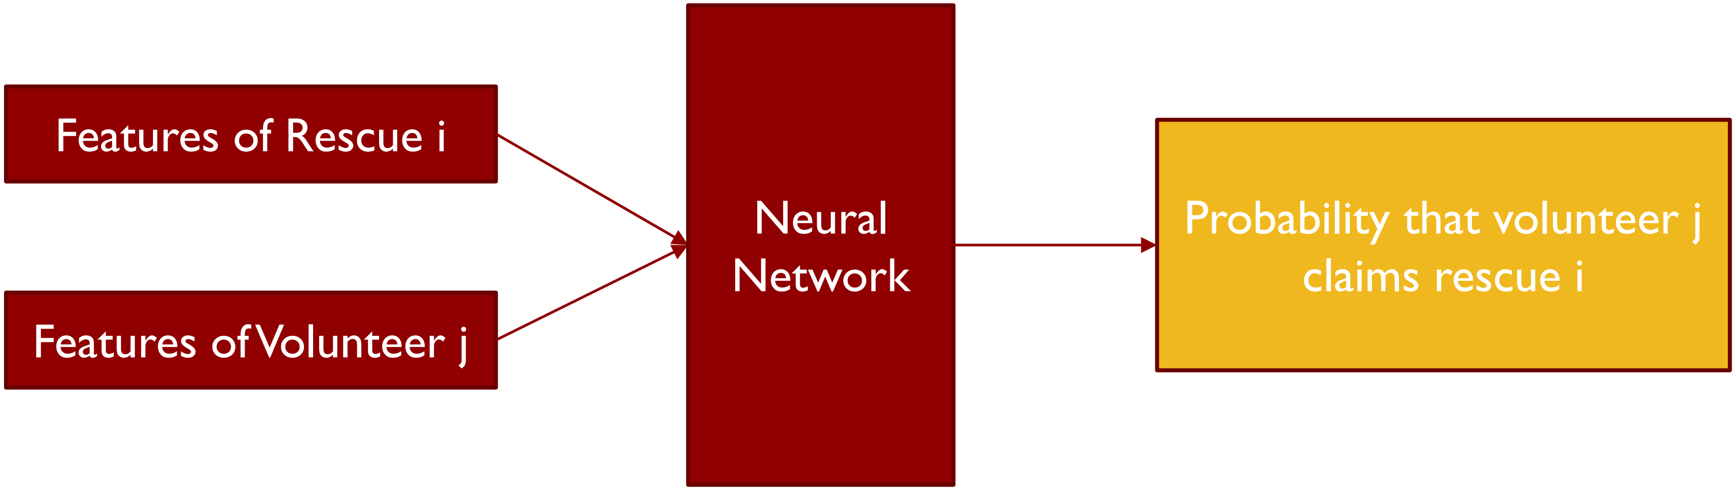

How to choose or develop the right AI method for the problem?

- Prediction task
  - Identify domain-specific challenges or limitations of existing methods
  - Develop methods to tackle those challenges

Learning from Limited Data

- Collaborators often do not have enough data required by modern AI techniques

- Training data is too small → cannot generalize well to unseen data

- How to deal with limited data?
  - Try to collect more data: Active learning to get labels for selected data points
  - Revisit the feasibility of the problem
  - Come up with methods that can learn from limited data

Learning from Limited Data

- Use less data-greedy approaches
  - Ensemble methods based on decision trees
  - Build models based on domain knowledge with a very small number of parameters to be learned from data
- Transfer learning
  - Learn from a relevant domain with rich data, apply (part of) the trained model in the target domain [Jean et al., 2016, Shen et al. 2018]
- Semi-supervised learning
  - Leverage abundant unlabeled data [Ma et al. 2018, Fan et al. 2018]

Learning from Limited Data

- Dimension reduction
  - Use some low-dimensional statistics such as the count of different pixel types as features instead of images [You et al., 2017]
  - Cut a long sequence into several shorter sequences [Zhou et al., 2019]

- Deal with missing features of some data points instead of dropping those data points
  - Fill the entries with imputation, e.g., deductive imputation, mean/median imputation
  - Model correlations between data points [Yan et al. 2013]

Tackling Biased Data

- Noisy labels
  - Sometimes one-sided: e.g., positive labels are indeed positive, but negative labels could be positive

- Learning from data with noisy labels
  - Only use data points with high confidence to train the model [Zhou et al., 2019]
  - Denoise data based on domain knowledge [Shankar et al, 2019]
  - Use learning algorithms designed for noisy labels, e.g., [Natarajan et al., 2013, Cheng et al. 2020]
  - Use noise correction algorithm, e.g., CORES^2 loss (Cheng et al. 2021) or peer loss (Liu and Guo 2020)-based noise correction algorithm

Tackling Biased Data

- Distributional shift
  - Machine learning model is developed using dataset D sampled from some distribution $p(x)$ but will be evaluated on data D' following some other distribution $q(x)$
    - E.g., Rich data in some geographical regions only in citizen science

- Deal with distributional shift
  - Factor the distribution shift into the model construction phase [Chen and Gomes]
  - Change loss function to $\mathbb{E}_{(x,y)\sim p}[L\frac{q(x)}{p(x)}]$

Tackling Biased Data

- Label imbalance
  - E.g., A lot more positive labels than negative labels
- Deal with label imbalance
  - Over-sampling/down-sampling in training data to get a balanced data set
  - Sample additional points based on domain knowledge
  - Collect soft labels from domain experts [Gurumurthy et al., 2018]

How to choose or develop the right AI method for the problem?

- Prescription task
  - Determine what actions are available, whether it is sequential decision making, how many decision-makers
  - Candidate AI methods for prescription
    - Mathematical programming
    - Game theoretic modeling
    - Multi-armed bandit (MAB) or restless MAB
    - Monte Carlo Tree Search
    - Reinforcement Learning
    - Imitation Learning
  - Address computational issues

Example: Ferry Protection

Probability flow over each edge (N^2T variables)

$$\min_{\boxed{f}\,v} v$$

$$\boxed{\text{s.t. } v \geq AttEU(i,t),}$$

$$\boxed{\begin{aligned} \sum_{e\in(i,t)\rightarrow} f(e) &= \sum_{e\in\rightarrow(i,t)} f(e) \\ \sum\nolimits_{e\in(*,0)\rightarrow} f(e) &= 1 \end{aligned}}$$

Best response

f is a unit flow

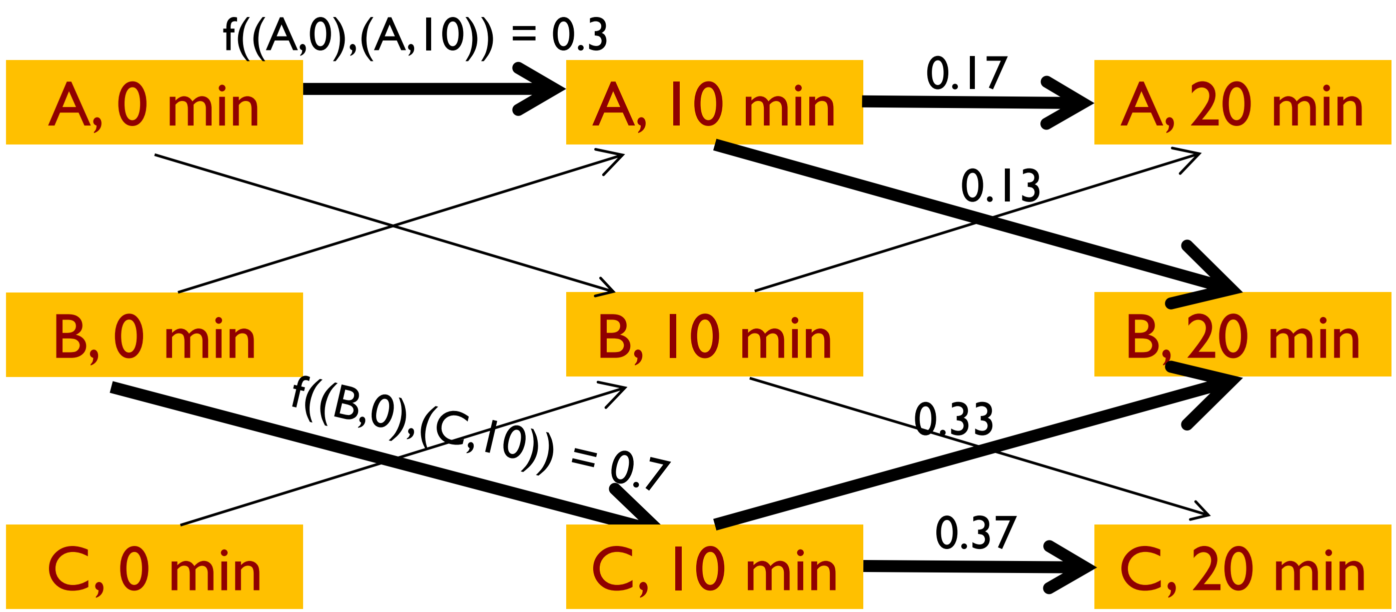

Common Challenges in Method Development Stage

- How to choose or develop the right AI method for the problem, while accounting for domain-specific considerations (e.g., practical constraints on computing resources or runtime, uncertainties and noise)?

- Discussion: Any other challenges (in method development stage) you faced or would like to learn more about?

Outline

- Typical Pipeline for AI for Social Good Projects
- Common Challenges and Practical Guide by Stage
  - Problem Formulation
  - Method Development
  - Evaluation & Deployment
- Q&A + Discussion

Common Challenges in Evaluation and Deployment Stage

- How to run a field experiment?
- How to evaluate the impact?
- How to get the solution deployed in a scalable and sustainable fashion?

How to run a field experiment?

- Need to convince the stakeholders first
  - Show to stakeholders the potential impact of AI-based solution through in-lab simulations with real-world data (dry-run)
  - Explain the AI-based solution -- ideally make the AI solution interpretable such that all the stakeholders can understand and trust the AI-based solution
  - Small-scale pilot test, learn key limitations, get feedback from stakeholders, and improve solution

- Randomized control trial (a.k.a. A/B testing)

How to evaluate the impact?

- Determine the stakeholders who might be impacted
  - Not just the direct users, but also other stakeholders

- For each class of stakeholder, design evaluation metrics and ways to collect data to evaluate
  - Survey, interviews
  - Quantitative measures

How to get the solution deployed in a scalable and sustainable fashion?

- Cloud service
  - Example: food rescue (AWS)

- Integrate into the software or devices stakeholders are already using (through API, file format that can be directly imported into their software or devices)
  - Example: PAWS (Microsoft, SMART)

Example: Ferry Protection

- Deployed since 2013

- US Coast Guard evaluation
  - Point defense to zone defense
  - Increased randomness
  - Mock attacker

- Patrollers feedback
  - More dynamic (speed change, U-turn)

- Professional mariners' observation
  - Apparent increase in Coast Guard patrols

Outline

- Typical Pipeline for AI for Social Good Projects
- Common Challenges and Practical Guide by Stage
  - Problem Formulation
  - Method Development
  - Evaluation & Deployment
- Q&A + Discussion

Q&A

- Questions about course project?

Discussion

- Share your ideas that you think others might consider for course project
  - From your discussion sections of PRAs

References

- Di Chen and Carla P Gomes. Bias reduction via end-to-end shift learning: Application to citizen science. AAAI-19
- Cheng, Hao, Zhaowei Zhu, Xingyu Li, Yifei Gong, Xing Sun, and Yang Liu. "Learning with instance-dependent label noise: A sample sieve approach." *arXiv preprint arXiv:2010.02347* (2020).
- Natarajan, N., Dhillon, I. S., Ravikumar, P., & Tewari, A. (2013, December). Learning with noisy labels. In *NIPS* (Vol. 26, pp. 1196-1204).
- Shiv Shankar, Daniel Sheldon, Tao Sun, John Pickering, and Thomas G. Dietterich. Three-quarter sibling regression for denoising observational data, IJCAI-19
- Swaminathan Gurumurthy, Lantao Yu, Chenyan Zhang, Yongchao Jin, Weiping Li, Xiaodong Zhang, Fei Fang. Exploiting Data and Human Knowledge for Predicting Wildlife Poaching. In COMPASS-18
- Shi, Zheyuan Ryan, Claire Wang, and Fei Fang. Artificial intelligence for social good: A survey. arXiv preprint arXiv:2001.01818 (2020).

References

- Yujie Fan, Yiming Zhang, Yanfang Ye, and Xin Li. Automatic opioid user detection from twitter: Transductive ensemble built on different meta-graph based similarities over heterogeneous information network. In IJCAI-18
- Neal Jean, Marshall Burke, Michael Xie, W. Matthew Davis, David B. Lobell, and Stefano Ermon. "Combining satellite imagery and machine learning to predict poverty." *Science* 353, no. 6301 (2016): 790-794.
- Tengfei Ma, Cao Xiao, Jiayu Zhou, and Fei Wang. Drug similarity integration through attentive multi-view graph auto-encoders. In AAAI-18
- Tiancheng Shen, Jia Jia, Guangyao Shen, Fuli Feng, Xiangnan He, Huanbo Luan, Jie Tang, Thanassis Tiropanis, Tat-Seng Chua, and Wendy Hall. Cross-domain depression detection via harvesting social media. In AAAI-18
- Junchi Yan, Yu Wang, Ke Zhou, Jin Huang, Chunhua Tian, Hongyuan Zha, and Weishan Dong. Towards effective prioritizing water pipe replacement and rehabilitation, IJCAI-13
- Jiaxuan You, Xiaocheng Li, Melvin Low, David Lobell, and Stefano Ermon. Deep gaussian process for crop yield prediction based on remote sensing data. In AAAI-17
- Yuxi Zhou, Shenda Hong, Junyuan Shang, Meng Wu, Qingyun Wang, Hongyan Li, and Junqing Xie. K-margin-based residual-convolution-recurrent neural network for atrial fibrillation detection. In IJCAI-19

References

- Cheng, H.; Zhu, Z.; Li, X.; Gong, Y.; Sun, X.; and Liu, Y. 2021. Learning with Instance-Dependent Label Noise: A Sample Sieve Approach. In ICLR.
- Liu, Y.; and Guo, H. 2020. Peer Loss Functions: Learning from Noisy Labels without Knowing Noise Rates. In ICML'20

Backup Slides

How to start an AI for Social Good project

- What problem to work on?
  - 1) Start from interest in a domain/real-world problem
    - E.g., Found something unsatisfactory from your volunteer experience? Found an interesting real-world problem from news articles?
    - Brainstorm: How that AI can help

How to start an AI for Social Good project

- What problem to work on?
  - 2) Start from interest in a certain technique
    - E.g., Read/implemented state-of-the-art methods in computer vision? Had research experience in natural language processing?
    - Brainstorm: What real-world problems can the technique potentially be used for?

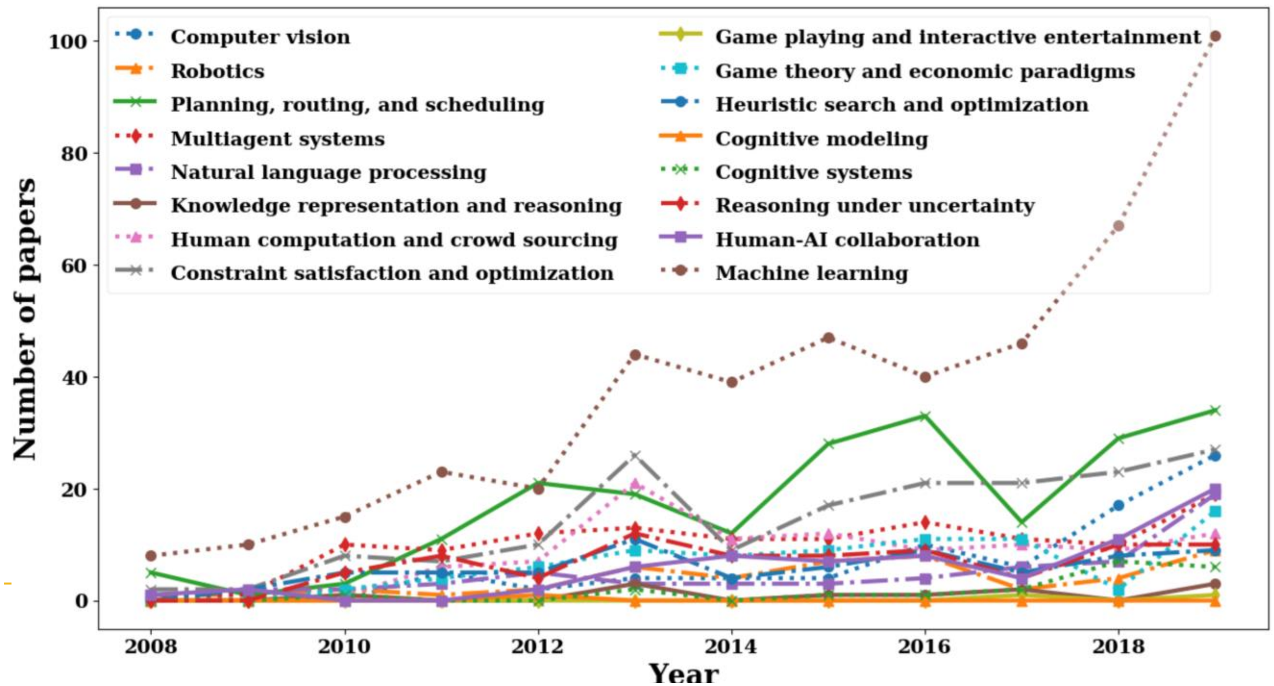

How to start an AI for Social Good project

- What problem to work on?
  - 3) Start from a domain/technique combination

Technique	Agriculture	Education	Environmental sustainability	Healthcare	Combating information manipulation	Public safety	Social care and urban planning	Transportation	Total
Cognitive modeling	0	0	0	0	0	1	0	0	1
Constraint satisfaction and optimization	2	5	31	48	20	26	9	59	173
Cognitive systems	1	2	2	7	2	3	1	5	20
Computer vision	3	8	12	20	6	12	7	19	79
Game playing and interactive entertainment	0	1	0	1	0	0	0	0	2
Game theory and economic paradigms	3	5	30	6	11	31	1	16	78
Human-AI collaboration	1	8	11	23	9	6	6	17	69
Human computation and crowd sourcing	1	5	6	20	45	12	11	15	98
Heuristic search and optimization	1	3	11	14	8	8	6	26	69
Knowledge representation and reasoning	0	0	0	5	3	2	0	1	11
Multiagent systems	2	7	47	19	16	22	8	31	122
Machine learning	12	27	65	174	53	65	36	92	460
Natural language processing	4	12	6	18	10	10	5	3	58
Planning, routing, and scheduling	9	4	48	43	14	28	31	84	210
Robotics	3	4	12	10	4	5	4	10	47
Reasoning under uncertainty	4	3	30	23	8	6	6	13	78
Total	40	78	225	344	155	177	90	253	1176

How to start an AI for Social Good project

- What problem to work on?
  - 4) Start from an existing work or established challenges
    - E.g., Go through existing work that applies AI technique to tackle societal challenges
      - IJCAI/AAAI special tracks, AI for Social Good workshops/symposiums
      - COMPASS, AIES, etc
      - Lists of previous course projects (Lec 1)
    - E.g., Check Kaggle competitions
    - Brainstorm: Is there room for improvement?
      - Existing model missing some critical practical aspects?
      - A new algorithm can lead to better performance?
    - Brainstorm: Is there a similar problem that can use similar framework?

Common Challenges in AI for Social Good Problems

- Learning from Limited Data
- Tackling Biased Data
- Stackelberg Leadership Models
- Privacy-preserving ML
- Human in the Loop

Typical Frameworks

- First identify a concrete social good problem that AI methods can potentially help

- Option 1: Data-centric
  - Look for real-world data and clean/Preprocess data
  - Identify or Propose AI algorithm that can be applied to the data
  - Evaluate algorithm, summarize/visualize result
  - Discuss insights and lessons learned
  - Example: "Detecting Mining Sites from Satellite Imagery Using Faster R-CNN"

- Option 2: Model/algorithm-centric
  - Mathematically model the challenge
  - Propose AI-based solution
  - Theoretically analyze of the model/algorithm
  - Implement the algorithm and test on simulated or real-world instance
  - Example: "Optimizing Inspection Strategy to Reduce Air Pollution"
  - For Ph.D. students: recommended to talk to your Ph.D. advisor and choose a project