Carnegie Mellon University

Artificial Intelligence Methods for Social Good

Meeting Days, Times, Location: Tue/Thu 10:30am-11:50am, GHC 5222

Semester: Spring**, Year:** 2018

Units: 9/12**, Section(s):** 08-537/08-737

Instructor information

Name	Dr. Fei Fang
Contact Info	Email: feifang@cmu.edu
Office location	Wean Hall 4126
Office hours	Tue/Thu 1pm-2pm

TA Information

TA name	Chun Kai Ling
TA Contact Info	Email: chunkail@andrew.cmu.edu
Office location	GHC 6507
Office hours	

Course Description

The rapid advance in artificial intelligence (AI) has opened up new possibilities of using AI to tackle the most challenging societal problems today. This course brings together a set of advanced AI methods that allow us to address such challenges and promote social good:

1) Optimization: mathematical programming, robust optimization, influence maximization
2) Game Theory and Mechanism Design: security games, human behavior modeling, auction and market equilibrium, citizen science
3) Machine Learning: classification, clustering, probabilistic graphical models, deep learning
4) Sequential Decision Making: Markov Decision Processes (MDPs), partially observable MDPs, online planning, reinforcement learning

In addition to providing a deep understanding of these methods, the course will introduce which societal challenges they can tackle and how, in the areas of (i) healthcare, (ii) social welfare, (iii) security and privacy, (iv) environmental sustainability. The course will also cover special topics such as AI and Ethics and AI and Humans. Example research projects and social good outcomes can be found at http://aiandsocialgood.org

The course content is designed to not have too much overlap with other AI courses offered at CMU. Although the course is listed within SCS, it should be of interest to students in several other departments, including ECE, EPP and SDS.

(9 Unit) The students in this 9-unit course are expected to have taken at least three mathematics courses covering linear algebra, calculus, and probability. The students will work in groups on a systematic literature review or a project exploring the possibility of applying existing AI tools to a societal problem, with a survey paper or technical report and presentation delivered at the end of the semester.

(12 Unit) This 12-unit course is only open to graduate students (master and Ph.D. students) with previous programming experience and background knowledge in artificial intelligence. The students will work in groups on a research project with a research-style paper and an oral presentation delivered at the end of the semester. Please see the instructor if you are unsure whether your background is suitable for the course.

Learning Objectives

At the end of the course, the students should be able to

1) Identify societal challenges that can potentially be tackled by AI methods, and determine which AI methods can be applied
2) Describe the AI methods covered in the course, including the basic concepts, the key algorithms, and the commonly-used implementation of the methods
3) Model the societal challenges as mathematical problems that AI techniques can be applied and propose how to adjust and modify the AI techniques to fit the problems
4) Describe evaluation criteria and methodologies of applying AI methods for social good
5) Deliver written and oral presentation on research projects or research survey

Learning Resources

No formal textbook. References and additional resources will be provided in slides and on Canvas.

Assessments

The final course grade will be calculated using the following categories:

Assessment	**Percentage of Final Grade**
Class participation	10 points
Paper Summaries	20 points
Written Answers Assignment	20 points
Final Project	50 points

- Class participation. The grading of the class participation will be mostly based on attendance, checked by in-class quizzes and asking and answering questions in class. Other factors include asking and answering questions on Canvas.
- Paper reading assignment. The course will require all students to complete weekly paper reading assignments individually. In each assignment, the students are required to provide a summary of the paper/article, a list of questions, and a few brainstorming ideas. The assignments will be submitted through Canvas and will be peer reviewed, but the final score will be provided by the instructor and the TA. For each student, the lowest scored assignment will be dropped.
- Written answers assignment. The course will require all students to complete biweekly written answers

assignment individually. Each written answers assignment will involve checking the understanding of basic concepts and working through the algorithms presented in class on example problems. Most questions are multiple-choice questions or numerical answer questions. In addition to providing the answers, the students are required to provide explanations to the answers. The answers will be submitted and auto-graded through Canvas. The explanations will be checked through peer review. For each student, the lowest scored assignment will be dropped.

- Final project. The students will work in small groups (1-3 students in each group). The students are expected to focus on one or more societal challenges, summarize or propose models and AI-based solutions to tackle the challenges, and evaluate the solutions. The students are required to submit project report through Canvas and deliver oral presentation. The instructor will provide suggested project topics. The students can also propose their own projects topics related to AI and Social Good but they will need consent from the instructor. The progress of projects will be checked through Project Proposal, Project Progress Report, Project Presentation, and Final Project Report. The proposal and progress report will be peer reviewed. The presentation and the final report will be evaluated by instructor and TA directly.
    - The students who choose the 12-unit section will work on a research project, with a research-style paper and an oral presentation delivered at the end of the semester. Ph.D. students may choose a paper format with the guidance of their advisors. For master students, a sample research-style paper can be found here: https://feifanginfo.files.wordpress.com/2016/11/2015_aaai_csworkshop_repeatedssg.pdf
    - The students who choose the 9-unit section will work on a systematic literature review or a project exploring the possibility of applying existing AI tools to a social good problem, with a survey paper or technical report and presentation delivered at the end of the semester.
        - Sample survey paper: http://www.sigecom.org/exchanges/volume_15/1/FANG.pdf
        - Sample technical report: https://www.aaai.org/ocs/index.php/SSS/SSS12/paper/view/4310/4644

Students will be assigned the following final letter grades, based on calculations coming from the course assessment section.

Grade	Percentage Interval
A	[90,100] points
B	[80,90) points
C	[70,80) points
D	[60,70) points
R (F)	[0,59) points

Grading Policies

- **Late-work policy and Make-up work policy**: No late days or make-up work is allowed. However, for both paper reading assignments and written answer assignments, the lowest scored assignment in each category will be dropped.
- **Re-grade policy**: To request a re-grade, the student need to write an email to the instructor titled “Re-grade request from [Student’s Full Name]” within one week of receiving the graded assignment.
- **Attendance and participation policy**: Attendance and participation will be a graded component of the course. The grading of the class participation will be mostly based on attendance, checked by in-class

quizzes and asking and answering questions in class. Other factors include asking and answering questions on Canvas.

Course Policies

- **Academic Integrity & Collaboration**: For both paper reading assignments and written answer assignments, a student can discuss with other students, but he need to specify the names of the students he discussed with in the submission, and complete the calculations and writing of explanations, summary, and questions on his own. For the final project, the students can discuss and collaborate with others (including students, faculty members, and domain experts), but the students need to give proper credits to whoever involved, and report the contributions of each group member in the final report and presentations, which will be considered in the grading. For assignments and final project, it is allowed to use publicly available code packages but the source of code package needs to be specified in the submission. Plagiarism is not allowed. The policy is motivated by CMU policy on academic integrity https://www.cmu.edu/student-affairs/ocsi/academic-integrity/index.html
- **Mobile Devices**: Mobile devices are allowed in class. Cellphones should be in silent mode. Students who use tablet in upright position and laptops will be asked to sit in the back rows of the classroom.
- **Accommodations for students with disabilities**: If you have a disability and require accommodations, please contact Catherine Getchell, Director of Disability Resources, 412-268-6121, getchell@cmu.edu. If you have an accommodations letter from the Disability Resources office, I encourage you to discuss your accommodations and needs with me as early in the semester as possible. I will work with you to ensure that accommodations are provided as appropriate.
- **Statement on student wellness**: As a student, you may experience a range of challenges that can interfere with learning, such as strained relationships, increased anxiety, substance use, feeling down, difficulty concentrating and/or lack of motivation. These mental health concerns or stressful events may diminish your academic performance and/or reduce your ability to participate in daily activities. CMU services are available, and treatment does work. You can learn more about confidential mental health services available on campus at: http://www.cmu.edu/counseling/. Support is always available (24/7) from Counseling and Psychological Services: 412-268-2922.

Course Schedule (Subject to Change)

Date	Theme/Topic	Assignment Due	Reference
1/16	M0: Introduction, Logistics, Course Project		
1/18	M1-1 [Optimization]: Optimization Problems Cover: Convex optimization, Linear Programming (LP) and Mixed Integer Linear Programming (MILP)		*Convex Optimization, Chapters 1-4* Stephen Boyd and Lieven Vandenberghe (Cambridge University Press) *Applied Mathematical Programming, Chapters 2, 9* Bradley, Hax, and Magnanti (Addison-Wesley, 1977)
1/23	M1-2 [Optimization]: Conservation Planning Cover: Wildlife corridor design	PRA1 due	(PRA1) *Solving Connected Subgraph Problems in Wildlife Conservation* Bistra Dilkina & Carla P. Gomes *Trade-offs and efficiencies in optimal budget-constrained multispecies corridor networks* Bistra Dilkina, Rachel Houtman, Carla P. Gomes, Claire A. Montgomery, Kevin S. McKelvey, Katherine Kendall, Tabitha A. Graves, Richard Bernstein,

			Michael K. Schwartz *Robust Network Design for Multispecies Conservation* Ronan Le Bras, Bistra Dilkina, Yexiang Xue, Carla P. Gomes, Kevin S. McKelvey, Michael K. Schwartz, Claire A. Montgomery
1/25	M2-1 [Game Theory]: Basics of Game Theory Cover: Minimax theory, Nash Equilibrium, Stackelberg Equilibrium	HW0 due	*Algorithmic Game Theory, Chapters 1-3* Editors: Noam Nisan, Tim Roughgarden, Eva Tardos, Vijay V. Vazirani (Cambridge University Press)
1/30	M2-2 [Game Theory]: Security Games Cover: Ferry protection, ranger patrol planning	PRA2 due	(PRA2) *Deployed ARMOR Protection: The Application of a Game Theoretic Model for Security at the Los Angeles International Airport* James Pita, Manish Jain, Janusz Marecki, Fernando Ordóñez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri, Sarit Kraus *Optimal Patrol Strategy for Protecting Moving Targets with Multiple Mobile Resources* Fei Fang, Albert Xin Jiang, Milind Tambe *Deploying PAWS: Field Optimization of the Protection Assistant for Wildlife Security* Fei Fang, Thanh H. Nguyen, Rob Pickles, Wai Y. Lam, Gopalasamy R. Clements, Bo An, Amandeep Singh, Milind Tambe, Andrew Lemieux
2/1	M2-3 [Game Theory]: Human Behavior Modeling Cover: Prospect theory, quantal response, cognitive hierarchy	HW1 due Submit Final Project Group Member List	(PRA4) *Comparing Human Behavior Models in Repeated Stackelberg Security Games: An Extended Study* Debarun Kar, Fei Fang, Francesco M. Delle Fave, Nicole Sintov, Milind Tambe, Arnaud Lyet *Improving Resource Allocation Strategy Against Human Adversaries in Security Games* Rong Yang, Christopher Kiekintveld, Fernando Ordonez, Milind Tambe, Richard John *Predicting human behavior in unrepeated, simultaneous-move games* James R. Wright, Kevin Leyton-Brown
2/6	Guest Lecture by Prof. Illah Nourbakhsh (Carnegie Mellon University) AI and Ethics (may move to other dates in Feb)	PRA3 due	(PRA3) *The Rhetoric of Robotics* Illah Reza Nourbakhsh
2/8	M3-1 [Machine Learning]: Classification and Clustering Cover: Decision trees, k-means, Gaussian Mixture Models (GMMs)	PRA4 due	*Pattern Recognition and Machine Learning, Chapters 4, 9* Christopher Bishop

2/13	M3-2 [Machine Learning]: Probabilistic Graphical Models Cover: Dynamic Bayesian Networks (DBNs), Markov Random Fields (MRFs)	HW2 due	*Pattern Recognition and Machine Learning, Chapter 8* Christopher Bishop
2/15	M3-3 [Machine Learning]: Predicting Illegal Activities Cover: Predict poaching threat, predict urban crime	Project Proposal due	(PRA5) *Keeping Pace with Criminals: An Extended Study of Designing Patrol Allocation against Adaptive Opportunistic Criminals* Chao Zhang, Shahrzad Gholami, Debarun Kar, Arunesh Sinha, Manish Jain, Ripple Goyal, Milind Tambe *Taking it for a Test Drive: A Hybrid Spatio-temporal Model for Wildlife Poaching Prediction Evaluated through a Controlled Field Test* Shahrzad Gholami, Benjamin Ford, Fei Fang, Andrew Plumptre, Milind Tambe, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba, Mustapha Nsubaga, Joshua Mabonga *Cloudy with a Chance of Poaching: Adversary Behavior Modeling and Forecasting with Real-World Poaching Data* Debarun Kar, Benjamin Ford, Shahrzad Gholami, Fei Fang, Andrew Plumptre, Milind Tambe, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba
2/20	M3-4 [Machine Learning]: Regression Cover: Linear Regression, Regularization, Support Vector Machines (SVMs)	PRA5 due	*Pattern Recognition and Machine Learning, Chapters 3, 6* Christopher Bishop
2/22	M4-1 [Sequential Decision Making] Markov Decision Process (MDP) Cover: Value iteration	HW3 due	*Reinforcement Learning: An Introduction, Chapter 3* Richard S. Sutton and Andrew G. Barto
2/27	M1-3 [Optimization] Combinatorial Optimization and Robust Optimization Cover: Duality, branch and price, maximin model, minimax regret, prevent illegal fishing	PRA6 due	*Combinatorial Optimization: Algorithms and Complexity, Chapters 3* Christos H. Papadimitriou, Kenneth Steiglitz *Branch-and-price: Column generation for solving huge integer programs* Cynthia Barnhart, Ellis L. Johnson, George L. Nemhauser, Martin W. P. Savelsbergh, Pamela H. Vance *Robust protection of fisheries with COmPASS* William Haskell, Debarun Kar, Fei Fang, Milind Tambe, Sam Cheung, Elizabeth Denicola
3/1	Guest Lecture by Prof. Tuomas Sandholm (Carnegie Mellon University)	HW4 due	(PRA6) FutureMatch: Combining Human Value Judgments and Machine Learning to Match in Dynamic Environments.

	Cover: Kidney exchange		Dickerson, J. and Sandholm, T. In Proceedings of the *AAAI Conference on Artificial Intelligence*. Extended version with appendix. Position-Indexed Formulations for Kidney Exchange. Dickerson, J., Manlove, D., Plaut, B., Sandholm, T., and Trimble J. In Proceedings of the *ACM Conference on Economics and Computation (EC)*. Extended version.
3/6	M4-2 [Sequential Decision Making]: Partially Observable MDPs Cover: Online Planning, Monte-Carlo Tree Search (MCTS)	PRA7 due	*Monte-Carlo Planning in Large POMDPs* David Silver, Joel Veness *Bandit based Monte-Carlo Planning* Levente Kocsis and Csaba Szepesvari
3/8	Guest Lecture by Prof. Norman Sadeh (Carnegie Mellon University) Cover: Security and privacy, Livehoods: understand city using social media		(PRA7) *The Livehoods Project: Utilizing Social Media to Understand the Dynamics of a City* Justin Cranshaw Raz Schwartz Jason I. Hong Norman Sadeh
3/13	Spring Break, no class		
3/15	Spring Break, no class		
3/20	Guest Lecture by Prof. Phebe Vayanos (University of Southern California) Cover: Kidneys for translation, housing for homeless youth	PRA8 due Project Progress Report due	*Robust Multiclass Queuing Theory for Wait Time Estimation in Resource Allocation Systems* Chaithanya Bandi, Nikolaos Trichakis, Phebe Vayanos
3/22	Guest Lecture by Prof. Stephen Smith (Carnegie Mellon University) Cover: Smart traffic light control	HW5 due	(PRA8) *Accomodating High Value-of-Time Drives in Market-Driven Traffic Signal Control* Isaac Isukapati and Stephen Smith
3/27	M4-3 [Sequential Decision Making]: Invasive Species Management Cover: Reinforcement Learning, Invasive species management	PRA9 due	(PRA9) *PAC Optimal MDP Planning with Application to Invasive Species Management* Majid Alkaee Taleghan, Thomas G. Dietterich, Mark Crowley, Kim Hall, H. Jo Albers *PAC Optimal Planning for Invasive Species Management: Improved Exploration for Reinforcement Learning from Simulator-Defined MDPs* Thomas G. Dietterich, Majid Alkaee Taleghan, Mark Crowley
3/29	M1-4 [Optimization]: Influence Maximization Cover: Influence propagation models, submodular function optimization		*Maximizing the spread of influence through a social network* David Kempe, Jon Kleinberg, Éva Tardos *Submodular Functions: Extensions, Distributions, and Algorithms. A Survey* Shaddin Dughmi

			Information and Influence Propagation in Social Networks Wei Chen, Laks V.S. Lakshmanan, Carlos Castillo
4/3	Guest Lecture by Prof. Milind Tambe (University of Southern California) Cover: AI and Social Work; HIV prevention among homeless youth	PRA10 due	(PRA10) *Using Social Networks to Aid Homeless Shelters: Dynamic Influence Maximization under Uncertainty* Amulya Yadav, Hau Chan, Albert Jiang, Haifeng Xu, Eric Rice, Milind Tambe *Influence Maximization in the Field: The Arduous Journey From Emerging to Deployed Application* Amulya Yadav, Bryan Wilder, Eric Rice, Robin Petering, Jaih Craddock, Amanda Yoshioka-Maxwell, Mary Hemler, Laura Onasch-Vera, Milind Tambe, Darlene Woo *Uncharted but not Uninfluenced: Influence Maximization with an Uncertain Network* Bryan Wilder, Amulya Yadav, Nicole Immorlica, Eric Rice, Milind Tambe
4/5	M3-5 [Machine Learning]: Deep Learning Cover: Neural Networks (NNs), Convolutional NN, Faster RCNN, Detecting human and wildlife from UAV videos	HW6 due	*Deep Learning, Chapter 6, 9* Ian Goodfellow and Yoshua Bengio and Aaron Courville *Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks* Shaoqing Ren, Kaiming He, Ross Girshick, Jian Sun *SPOT Poachers in Action: Augmenting Conservation Drones with Automatic Detection in Near Real Time* Elizabeth Bondi, Fei Fang, Mark Hamilton, Debarun Kar, Donnabell Dmello, Jongmoo Choi, Robert Hannaford, Arvind Iyer, Lucas Joppa, Milind Tambe, Ram Nevatia
4/10	Guest Lecture by Prof. Stefano Ermon (Stanford University) Cover: Deep learning for developing countries	PRA11 due	(PRA11) *Combining satellite imagery and machine learning to predict poverty* Neal Jean, Marshall Burke, Michael Xie, Matthew Davis, David B. Lobell, Stefano Ermon *Deep Gaussian Process for Crop Yield Prediction Based on Remote Sensing Data* Jiaxuan You, Xiaocheng Li, Melvin Low, David Lobell, Stefano Ermon
4/12	M2-4 [Game Theory]: Mechanism Design with Money Cover: Auction, Truthfulness, Price-of-Anarchy		*Algorithmic Game Theory, Chapters 9, 11* Editors: Noam Nisan, Tim Roughgarden, Eva Tardos, Vijay V. Vazirani (Cambridge University Press)
4/17	M2-5 [Game Theory]: Spatio-Temporal Pricing in Ridesharing Platforms Cover: scheduling and pricing, market equilibrium	HW7 due	(PRA12) *Spatio-Temporal Pricing for Ridesharing Platforms* Hongyao Ma, Fei Fang, David C. Parkes

4/19	CMU's Carnival, No class	PRA12 due	
4/24	Guest Lecture by Prof. David Danks (Carnegie Mellon University) Cover: AI and Humans	PRA13 due	(PRA13-a) *Regulating Autonomous Systems: Beyond Standards* David Danks and Alex John London (PRA13-b) *Finding trust and understanding in autonomous technologies* David Danks
4/26	M2-6 [Game Theory]: Citizen Science Cover: Citizen science		(PRA14) *Avicaching: A Two Stage Game for Bias Reduction in Citizen Science* Yexiang Xue, Ian Davies, Daniel Fink, Christopher Wood, Carla P. Gomes *Behavior Identification in Two-stage Games for Incentivizing Citizen Science Exploration* Yexiang Xue, Ian Davies, Daniel Fink, Christopher Wood, Carla P. Gomes *Improving Your Chances: Boosting Citizen Science Discovery* Yexiang Xue, Bistra Dilkina, Theodoros Damoulas, Daniel Fink, Carla P. Gomes and Steve Kelling
5/1	Course Project Presentation 1	PRA14 due	
5/3	Course Project Presentation 2	HW8 due	
5/10		Final Project Report due	