A Summary of IST 402 (Week 14)

Emma Dodoo, Grace Morgan, Sarah Dragon, Andrew Witherite

AI and Multiagent Systems for Social Good

1. Public Safety and Security
2. Conservation
3. Public Health

KEY CHALLENGE

How can we optimize our limited intervention resources when we are interacting with other agents?

Optimizing Limited Intervention Resources

Public Safety and Security

- Game Theory for Security Resources Optimization
- Real-world: US Coast Guard, US Federal Air Marshals Service
    - i.e. African Countries: Ghana, Uganda

Public Safety and Conservation

- Security games and adversary
- Green security games
    - Intelligent patrol patterns
- Real-world: Parks
    - i.e. Zambia, Mozambique
    - PAWS
        - Increased hit rate on poachers and related items to 3 from .73

Solving These Problems: Overall Research Process

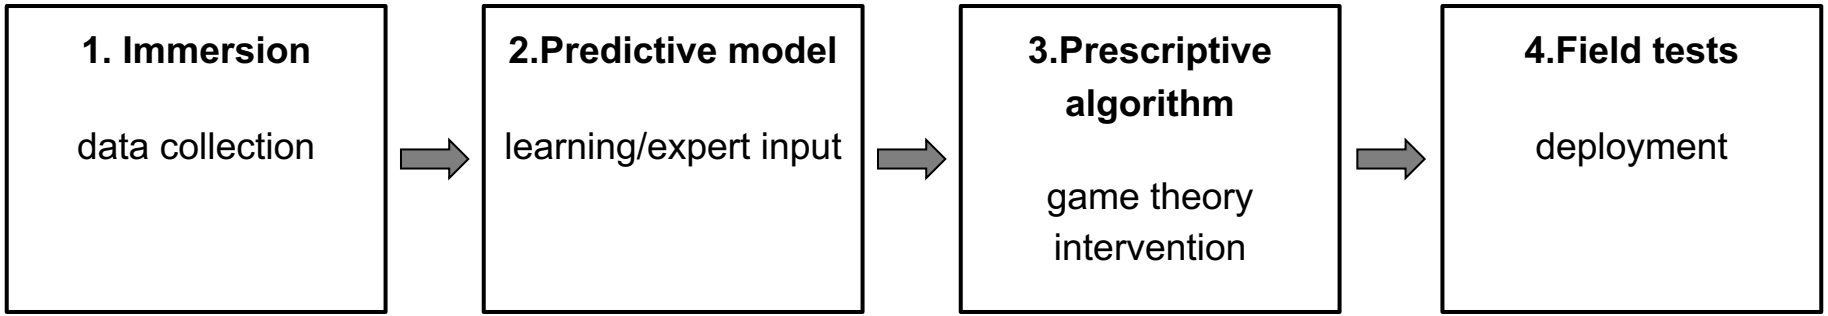

Game Theory: Security Resource Optimization

- Based on Stackelberg Security Game
  - Stackelberg Example: Defender commits to randomized strategy, adversary responds
    - shows a strategy and game can be as complex as you want
    - For simplicity, numbers are determined by surveillance, etc.
  - Optimization: Not 100% security (i.e defender would always win), increase cost/uncertainty to attackers
    - Can't guarantee 100% security due to limited resources
- Real-World: ARMOR at LAX [2007], IRIS for Federal Air Marshals Service [2009]

		Adversary	
		Terminal #1	Terminal #2
Defender	Terminal #1	4, -3	-1,1
	Terminal #2	-5, 5	2, -1

Why don’t Countries use Machine Learning?

- We simply do not have enough data
  - A lot of countries simply do not have frequent breaches, to gather datasets of information

Class Recap

Key Takeaways

- AI has blossomed into one of the biggest fields in computer science
  - Led by AI and ML
  - More data - big data
  - Better computing power

AI Trajectory: Valid Concerns

- People: Elon Musk, Bill Gates, Stephen Hawking
  - Fake news generator (recently been released)
- AI Chatbot that was discovered to be racist
- Facebook tagged two people as gorillas among AI discrimination issues
- AI Discrimination amongst employee hires

But...

- Despite these issues, AI still has enormous potential to help underserved people
- Course has been exercise in making you believe in vision
- Numerous applications of how humanity can be helped through the use of AI

What cool techniques did we learn?

- Train, validation, test (supervised learning)
  - Training set is a set of examples used for learning a model
  - Validation set is a set of examples that cannot be used for learning the model but can help tune model parameters
    - Need a validation set to figure out what parameters should be used in your model
    - Helps control overfitting
  - Test set is used to assess the performance of the final model and provide an estimation of the test error
    - Results on the testing set are the only results that matter for final product
  - NOTE: Never use the test set in any way to further tune the parameter or revise the model
- Decision Trees

Training and Testing

- Underfitting - High bias
- Overfitting - High variance
- Bias Variance Tradeoff

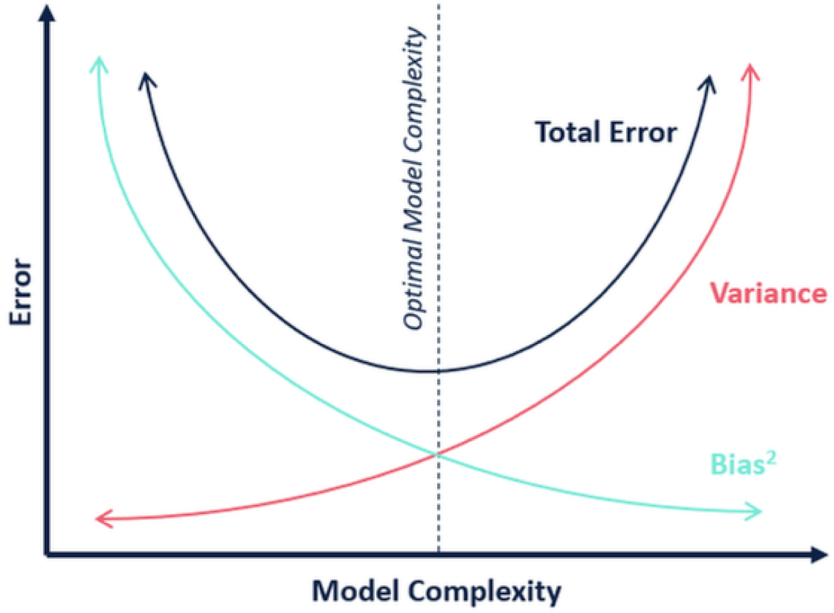

Evaluation Metrics

- Predicted labels
  - 2 options: Positive or Negative
  - The predictive and the actual labels can be either positive or negative

		Actual	
		Positive	Negative
Predicted	Positive	**True Positive**	**False Positive**
	Negative	**False Negative**	**True Negative**

Bagging

- More evolved version of a decision tree
  - Essentially create multiple decision trees
- After making all the decision trees, the decision that you will make will be the average of all the decision trees
  - Use an average model as the final model

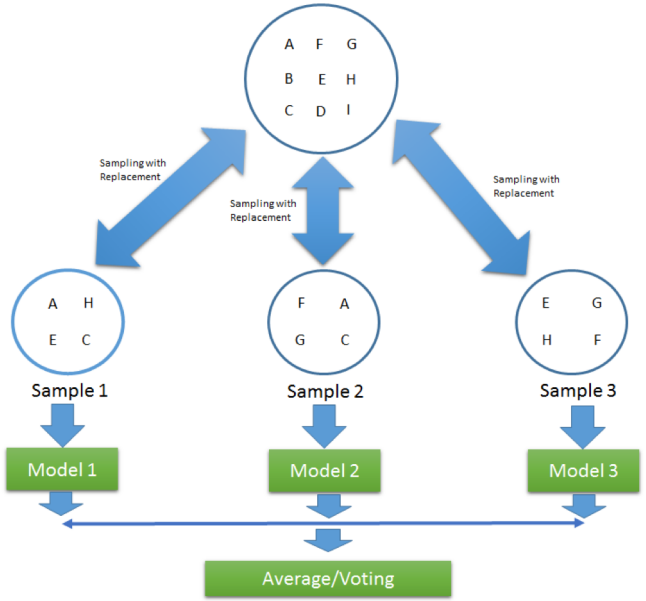

Prisoner’s Dilemma

- Rational way of solving: it is always optimal to defect
- In rational play, both players will choose to defect

	Confess Lie	
Confess	1 month, 1 month	3 years, Free
Lie	Free, 3 years	1 year, 1 year

Nash Equilibrium

- It is suboptimal for each player to deviate from their strategies, if the other play continues to play the same strategy
  - Given they are at equilibrium
  - Every finite game has a nash equilibrium

Stackelberg Game: Non-simultaneous moves

- One player is a defender - They move first
  - Must commit to random strategy
- Other player is the adversary - They move after the leader
  - Responds to random strategy
- Not 100% security - optimizing chances based on limited resources

Deep Learning

- Neural Networks
- Backpropagation algorithm
- Different kinds of Neural Networks
  - Vanilla deep nets - numeric data
  - Convolutional neural networks (CNN) - image data
  - Generative adversarial nets (GAN) - create data
  - Recurrent neural nets (RNN) - sequentially structured data (text, videos,etc.)

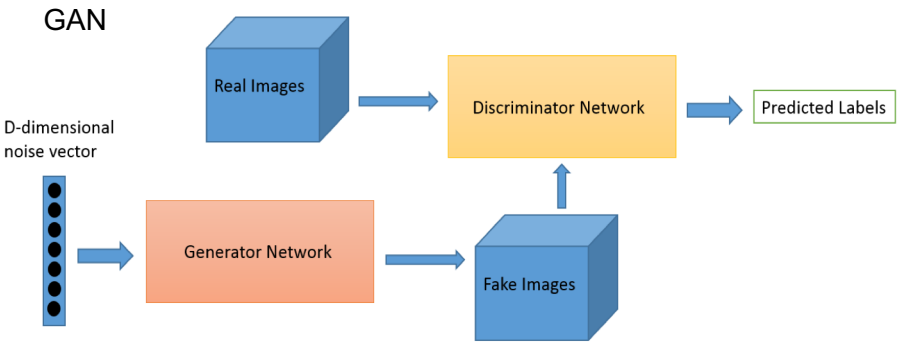

More than Accuracy...

- Interpretability - Medical community
- Fairness - COMPAS
- Ethics - Self-driving cars
- Machine application for social good - predicting childhood influenza