Computational Applications to Policy and Strategy (CAPS)

Johns Hopkins School of Advanced International Studies, Fall 2019

Leo Klenner, Henry Fung, Cory Combs

Course website: https://capsseminar.github.io/

Last updated: October 7, 2019

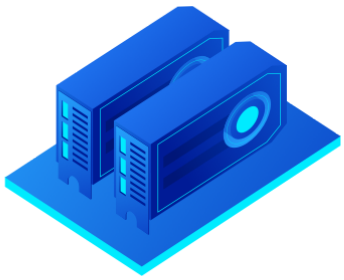

Course Overview

Summary

CAPS prepares international relations students to critically evaluate AI algorithms, specifically for applications that challenge conventional modes of policy and strategic decision-making. The course focuses on providing students with a deployable toolkit for assessing the trade-offs of common AI algorithms in the fields of supervised learning, unsupervised learning and reinforcement learning. CAPS places strong emphasis on understanding how data and the environments in which an algorithm is applied influence the algorithm's performance and the unexpected failure modes this can create. The course addresses how domain knowledge, from the regional or strategic fields taught at SAIS, can make critical contributions towards minimizing such failure modes and ensuring an optimal application of AI to complex problems. Upon completion of the course, students will have acquired a rigorous conceptual understanding of AI algorithms. Students will have gained the foundational skills needed to evaluate and make recommendations on the application of AI algorithms to the manifold challenges of today's complex world.

Why take this course?

AI is reshaping the problems that policymakers need to resolve – from strategic engagements to situations in the broader economy – and it is reshaping the means available to policymakers to resolve these problems – from deploying algorithms that augment human operators to learning new models of the world from data. In the public and private sectors, demand is rising fast for professionals who can connect engineers and policymakers around AI and continuously optimize how these two sides interact. CAPS teaches students core skills at the intersection of algorithms and policymaking that are needed to meet this demand and start a career in the emerging fields of AI policy and strategy. The course is highly interactive and transfers technical knowledge through a series of hands-on case studies focused on a fictional career path, the Forward Deployed Policy Engineer. Previous versions of the course have been developed and taught together with senior leadership from the public and private sectors.

Course Structure

Lecture Overview

Lecture	Date	Title	Description
1	10/29	Learning and Decision-Making	Types and measures of learning, approximation, complexity, efficiency, rule-based algorithms
2	11/5	Supervised Learning I	Learning from data, logistic regression, standard regression, tree-based models, k-nearest neighbor (KNN)
3	11/12	Supervised Learning II and Unsupervised Learning	Learning about data, principle component analysis (PCA), K-means clustering, ensemble methods, differences between human and machine learning
4	11/19	Reinforcement Learning	Dynamic data acquisition, goals and rewards, exploration vs exploitation, online vs offline learning, model-based vs model-free learning, human-machine teaming
5	11/26	Neural Networks I	Understanding the design of the Perceptron
6	12/03	Neural Networks II and Operational Aspects of AI	Feature selection, hyperparameter tuning, bias-variance trade-off, transparency, interoperability, safety

Topics Covered

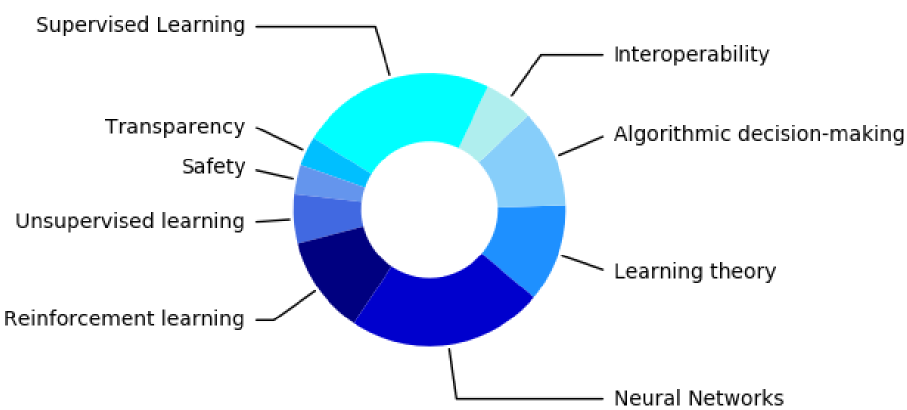

Syllabus

1. Learning and Decision-Making

This opening session covers the principles of human and algorithmic learning and decision-making, the similarities of which have propelled modern machine learning advancements and the differences between which remain a central barrier. We introduce the central orientation of the course around tradeoffs, both technical and ethical, and then the fundamentals of algorithms as needed to engage with contemporary machine learning systems, all geared toward engaging four central questions:

- How can policymakers and engineers jointly establish the most reliable, consistent paths of decision-making?
- How can we assess whether these are “good” paths? What are the relevant dimensions of “good” in a given strategic setting?
- What are the fundamental limitations of decision-making, from both human and algorithmic perspectives? What aspects can be improved?
- What problems do we face in instituting decision-making practices, specifically given differences in human and machine decision-making?

```
Case: Analyzing patterns of learning in a counterinsurgency team
```

2. Supervised Learning I: Bridging Policymaking and Engineering

How to algorithms learn *from* data? In this session, students engage the essential techniques of classification and regression data and explore their application to strategy and policy. Through live demonstrations and hands-on case studies, students explore the benefits and challenges of four essential techniques: logistic regression, standard regression, tree-based models, and the KNN algorithm. Anchoring discussion are the key questions:

- What does “learning” mean in supervised learning? How can we measure learning?
- What are the key data risks? How do we manage them, both through measurement and alternative techniques?
- What knowledge and language must policymakers and engineers share for productive collaboration?

```
Case: Training human data labelers in a complex organizational ecosystem
```

3. Supervised Learning II and Unsupervised Learning: Bridging Learning Types

How do algorithms learn *about* data? Continuing the lessons of labeled data from the first session on supervised learning, we move into the realm of unlabeled data using the techniques of unsupervised learning. We focus on current industry-standard techniques of principle component analysis (PCA), K-means clustering and hierarchical clustering, and broach emerging uses of ensemble methods and human-machine teaming. Continuing the hands-on cases and demonstrations, participants will discuss:

- Based on our findings, what can we and can’t we learn from unsupervised learning?
- How can we connect supervised and unsupervised machine learning to optimize our own understanding of policy or strategy issues?
- Looking ahead to our next section: how does machine learning compare to human learning?

```
Case: Deploying an autonomous supply agent to assist a human team
```

4. Reinforcement Learning: The Future of Artificial Systems?

Agent-environment interaction: the basis of human learning and the central mode of reinforcement learning. Present at both the origins and the current cutting edge of machine learning, reinforcement learning remains one of the most exciting areas of machine learning research and one of the most difficult to implement in real-world situations. Presenting its immense challenges and immense potential, we explore the key tenets of reinforcement learning, including dynamic data acquisition, goals and rewards, key tradeoffs such as between exploration and exploitation, online and offline learning, and model-based and model-free learning, and the potential for human-machine teaming.

- How can machines learn from interacting with the environment? With what limitations and risks?
- What are the key tradeoffs involved in reinforcement learning? How do they parallel and diverge from human learning?
- How can human-machine teaming be leveraged – and to what ends?

```
Case: Reverse engineering the agency of an unknown aerial vehicle
```

5. Neural Networks I: Lifting the Veil

Perhaps the most captivating and yet most opaque areas of machine learning, neural networks offer what is ostensibly the closest approximation to human learning yet created. This session focuses on demystifying their mechanics and capabilities through hands-on engagement. Returning to the course's earlier themes, we step through the mechanics of neural networks, starting with the perceptron, to further explore the question: what are the limits of algorithmic learning? In contrast to the earlier topics, which focused on individual types of learning, this session explores how neural networks can be designed to engage in myriad forms of learning, with examples applications to strategy and policy.

- How do neural networks operate?
- How can they be designed to implement a range of learning types?
- What are the requirements for neural networks to perform well, and how does this enable or constrain them in tackling strategic and policy issues?

```
Case: Training a simple neural net in Python and Google Collab
```

6. Neural Networks II and Operational Aspects: Deploying Machine Learning Systems

With the foundations of neural networks in place, this final session delves into practical examples, grounded in the more technical content essential for governance and appropriate implementation. Students will undertake scenarios to explore essential aspects such as feature selection and hyperparameter tuning with an eye to managing issues such as overfitting and various forms of bias. These issues will be related back to the core learning types discussed earlier in the course, reinforcing the full suite of concepts learned. Finally, students will have the opportunity to conceptualize their own machine learning approaches to an archetypal strategic challenge.

- What key data risks do we face when implementing machine learning? What algorithmic risks?
- As a policy or strategy professional, what tradeoffs may be most essential for your assessment of deployed neural network systems?
- In what scenarios might you commission a neural network versus another machine learning system?

- Using the foundations learned in this course, what next steps should you take as a policy or strategy professional interested in the future of machine learning?

```
Case: Devising levels of stakeholder access to an advanced AI-based model
```

FAQ

Who are the instructors?

CAPS is taught by recent SAIS graduates Leo Klenner, Henry Fung, and Cory Combs.

Leo and Cory work as Research Fellows at SAIS and Henry works as a Machine Learning Engineer for a Virginia-based start-up.

We started CAPS as SAIS students in the Spring 2018 semester and this is the fourth time that we're teaching the course. We always adapt CAPS from semester to semester, so this is the first time we're teaching the above syllabus.

You can contact us with any questions at capsseminar@gmail.com

Does the course have prerequisites?

There are no prerequisites for taking this course. We will make sure that everyone is on board with the technical aspects of algorithms or working with Python code. It helps if you're curious about acquiring a foundational but technically-grounded understanding of AI.

Do I need to install software for the course?

No. We will run all code through in-browser solutions like Google Collab, so you don't need to install Python or other software. However, if you want to install and run Python locally on your computer, that's fine, too.

Does the course have lecture notes?

Yes. We will make slides, lecture notes, and case studies available at the start of the course. At the start of the course, you will be able to find these documents on the course website https://capsseminar.github.io/.

What is a Forward Deployed Policy Engineer and why does it matter?

To help you navigate the complex landscape of AI and how AI interacts with policy and strategy, this course is structured around a fictional career path, the Forward Deployed Policy Engineer (FDPE).

The core responsibilities of a FDPE are:

- provide analytic guidance in domains that have an opaque problem space, defy precise solutions and feature decision-making involving a combination of human actors and autonomous agents
- ensure that each stakeholder of a decision-making process has the best possible shared understanding of the data that drives a decision and how the decision will alter the environment in which it is carried out
- contribute to stable ecosystems of human-machine decision-making through technical documentation, policy analysis, and model testing, especially in an ad-hoc manner when relevant information is sparse

- build robust connections between high-level goals from policy and low-level operational aspects from technology to ensure mission success in complex environments

Why does the FDPE matter?

There are two main reasons why the FDPE matters for this course: integration of distinct skill sets and adaptation in a shifting labor market.

Integration of distinct skill sets: In this course, we'll cover three distinct skill sets, the analysis of technical systems, the analysis of policy and strategy directives, and engineering. In today's economy, each of these skill sets is associated with a separate career. We use the FDPE career path to present these distinct skill sets, and careers, in an integrated and immersive way. In this way, the FDPE career path provides a coherent learning experience around a diverse portfolio of skills.

Adaptation to a shifting labor market: Although you might never work as a FDPE, you might transition between two careers whose skill set the FDPE encompasses, such as from a career in engineering to a career in policy. In this case, the FDPE career path provides you with a means to show that you can transfer skills from your present career in engineering to your next career in policy, facilitating your career transition. You can think of the FDPE like a meta-career that links your current career with careers that you might hold in the future, thereby providing you with greater adaptability in today's shifting labor market.

If you have questions about the FDPE career path, reach out to us at capsseminar@gmail.com