

# THESIS DEFENSE

## Forecasting and Controlling Behavior by Learning from Visual Data

Tuesday, August 27, 2019  
4305 Newell Simon Hall  
1:00p.m.

### Thesis Committee:



**Kris M.  
Kitani**  
Chair



**Martial  
Hebert**



**Ruslan  
Salakhutdinov**



**Sergey  
Levine**  
University of  
California, Berkeley



**Paul  
Vernaza**  
Aurora



## Nicholas Rhinehart

### Abstract

Achieving a precise predictive understanding of the future is difficult, yet widely studied in the natural sciences. Significant research activity has been dedicated to building testable models of cause and effect. From a certain view, a perfect predictive model of the universe is the “holy grail”; the ultimate goal of science. If we had it, we could anticipate, and therefore (at least implicitly) understand all observable phenomena. We approach the difficulty of modeling the future by deferring as much of the modeling as possible to be computationally learned. In this work, we seek to computationally model the

future in the presence of agent behavior given rich observations of the environment. The brunt of our focus is to reason about what agents will do, instead of other dynamic aspects of the environment. Whereas many natural science theories offer human-crafted predictive models of physical phenomena, we instead offer a paradigm of learned correlation-based predictive models of behavior-based phenomena.

The human capability to forecast offers complementary motivation. Humans use rich environment observations to inform their understanding, and ultimately, their future behavior. Critical to our intelligence is our ability to plan behaviors by considering how our actions are likely to result in future payoff, especially in the presence of other collaborative and competitive agents. We argue that a system cannot be intelligent if it cannot explicitly reason about the future of itself and other meaningful entities. By this logic, explicitly reasoning about the future is a necessary component of intelligence. Therefore, as scientists, we must design systems to explicitly forecast in order to have any hope of building intelligent systems.

The field of Computer Vision (CV) is focused on designing algorithms to automatically understand images, videos, and other perceptual data. However, the field’s effort to-date focuses on non-interactive, present-focused tasks, like object detection, scene classification, geometric understanding, and activity classification. Most CV contributions are algorithms to answer questions like “what is that”, and “what happened”, rather than “what will happen”, “what could happen”, or “how could I achieve X”. Computer Vision has under-explored reasoning about the interactive and decision-based nature of the world.

In contrast, Reinforcement Learning (RL) prioritizes modeling interactions and decisions by focusing on how to design algorithms to evoke behavior that maximizes a scalar reward signal. The resulting learning agents, in order to perform well, must have an understanding of how their current behaviors will affect their prospects of future reward. However, in the dominant paradigm of model-free RL, agents reason implicitly about the future. In contrast, model-based RL learns one-step dynamics as  $P(s'|s,a;\theta)$  or  $s' = f(s,a;\theta)$ . One-step dynamics provide an explicit estimate of “what could happen in the near future”. Combined with knowledge of how the agent will react to any given situation as a policy  $\pi(s'|s,a)$  or  $s' = \pi(s,a)$ , these objects enable us to forecast the distribution of future outcomes at arbitrarily-long time horizons. Unfortunately, in multi-agent systems, these objects are insufficient to forecast the future. We must also estimate how all agents will behave, in combination with the one-step world dynamics, in order to achieve a distribution of future outcomes over multiple time-steps.

In this thesis, we consider the problem of designing algorithms to enable computational systems to reason about the future behavior of intelligent agents given rich observations of their environments, as well as to use this reasoning for control. We primarily employ the frameworks of Imitation Learning and Reinforcement Learning to formulate and situate our work. We contribute forecasting approaches to excel in diverse, realistic, single and multi-agent domains. These include (1) sparse models to generalize from few demonstrations of human daily activity, (2) adaptive models to continuously learn from demonstrations of human daily activity, (3) high-dimensional generative models learned from demonstrations of human driving behavior, and (4) factorized multi-agent models to reason about the future interactions of agents. We also contribute control approaches to excel in diverse domains, including (5) incentivized forecasting, which encourages an artificial agent that only has access to partial observations of state to learn predictive state representations in order to perform a task better, and an approach to (6) leverage a probabilistic forecasting model as the backbone for flexible, controlled behavior.