Applying Markov Decision Processes to Optimize Aggregate Workforce Planning with Uncertain Funding Conditions

UNIVERSITY OF ARKANSAS
College of Engineering
Industrial Engineering

Blake Sooter (Team Leader), Katie Augsburger, James Fite, Nathanael Harris, Deniz Vural
Stephen Sieck (Division Lead | Transformational Technologies), David Hillstrom (Simulation Scientist)
May 1st, 2024

i∞ LABS

Infinity Labs LLC

Infinity Labs is a defense firm, founded in 2020 with 7 employees, that seeks to solve the world's hardest problems in modeling, simulation, and cybersecurity by developing innovative technology solutions. Today, Infinity Labs has ~70 employees across 14 locations in the United States. Our project improves their hiring decision process given their rapid growth and grant-driven revenue model.

Current Hiring Process under Uncertainty

Estimate Project Level Funding Risk → Estimate FTE Need by Project/Month → Calculate FTE Gap

The Current Process: To generate an expected value of each contract, Infinity considers their funding success rate, proposal quality, relationship with the funding entity, and the percent of a contract that will be fulfilled in house to generate a 0-100 project confidence estimate. These expected values for each grant are converted to a Full Time Equivalent (FTE)/project/month using start and end dates, the hourly cost of labor, and the number of working hours in a month. This value is compared to the existing FTE capacity to determine the FTE gap. This process is overseen by program managers and is largely ad-hoc.

Characterizing Uncertainty: The existing process is deterministic. To acknowledge uncertainty in project information, a triangular distribution is assumed around the project confidence estimate. An empirical distribution for lateness is applied to project start dates. Random variates are drawn from these distributions, and the Monte Carlo technique is applied to generate improved expected values for FTE need and a distribution of possible outcomes.

Sequential Hiring Decisions as an MDP

S = <-20 Gap, 20 Gap>
for all decision epochs (t)
A = <Hire 0, Hire 20>
P = <$P_{-20\ gap}, P_{20\ gap}$ >
for all decision epochs (t)
$R = \max(gap) - |gap_t|$

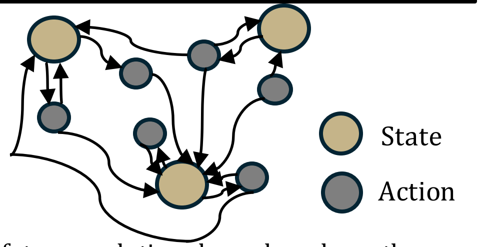

Markov Property: A system's future evolution depends only on the present state.

discount factor	0.9	reward wt	[1,1]
max hiring action	20	triangular	[c-0.2,c+0.1]
projects/staff	3	lateness	[0.2,0.7,0.9,0.9,1]
horizon	9 mos	iterations	10,000

Transition Matrix: The Monte Carlo estimate and distribution of likely outcomes is used to determine the probability that a particular action will lead to a particular outcome. Each iteration of the simulation is binned into a group of size of 1 FTE to estimate the probability density function.

$$\overset{57}{P(FTE\ Achieved_{t+1})} \overset{=}{=} \overset{58}{P(FTE\ Est_t} + \overset{-4}{State_t} + \overset{5}{Action_t} + \overset{-2}{State_{t+1})}$$

Apr	May	June	Jul	Aug	Sep	Oct	Nov
55	58	63	61	60	58	55	53

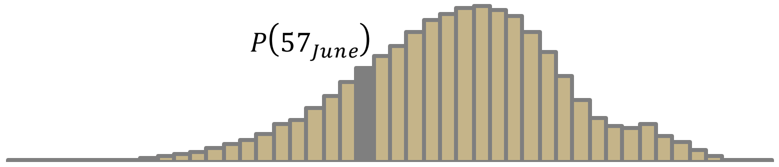

Use of Model and Optimal Policy

We implemented the Monte Carlo simulation using the numpy random number generator and our own implementations of probability distributions. The Markov Decision Process is solved using an implementation of value iteration from mdptoolbox. If project funding information changes substantially month to month, the policy can be reoptimized. Our script allows for what-if analysis on the previously described set of key assumptions, with both heatmap visualization and table-based policy output.

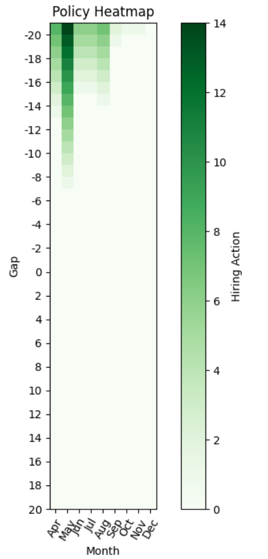

Example Use: If the decision maker observes that they are 15 employees short of the estimated need in May, they should hire 2 employees between April 1 and May 1.

Results and Future Recommendations

We compared the impact of the sequential decision model on rewards compared to a simple rule subject to the same underlying transition probabilities. By requiring that an action be taken as prescribed by a simple rule, we reduce a Markov Decision Process to a Markov Chain. Each simple rule performed at least 14.48% worse than the equivalent MDP policy.

	Performance vs. MDP (Linear Rewards)
Hire 50% of any negative gap	−16.41%
Hire 100% of any negative gap	−14.48%

As best practices around project data collection evolve at Infinity, we recommend that project managers estimate 10th and 90th percentile confidence estimates to improve the triangular distribution and eventually add additional factors to the input FTE simulation.