Optimize the Next Best Action Engine

with Reinforcement Learning

COMCAST | MIT MANAGEMENT BUSINESS ANALYTICS

Project Sponsors: Michael Nipple, Sandeep Balasundaram, Natalya Shcherban **Faculty Advisor**: Prof. Retsef Levi **MBAn Students**: Katherine Wang & Sylvia Zhang

Problem Statement

Problem

What is Recommendation Engine?
Comcast delivers approximately 400 million product recommendations annually to its customers through the Xfinity Mobile App channel, with the product recommendation engine called Nexus.

What is the current method?
The existing approach utilized by Nexus system relies on micro-segmentation, a clustering-based technique, which primarily depends on static data.

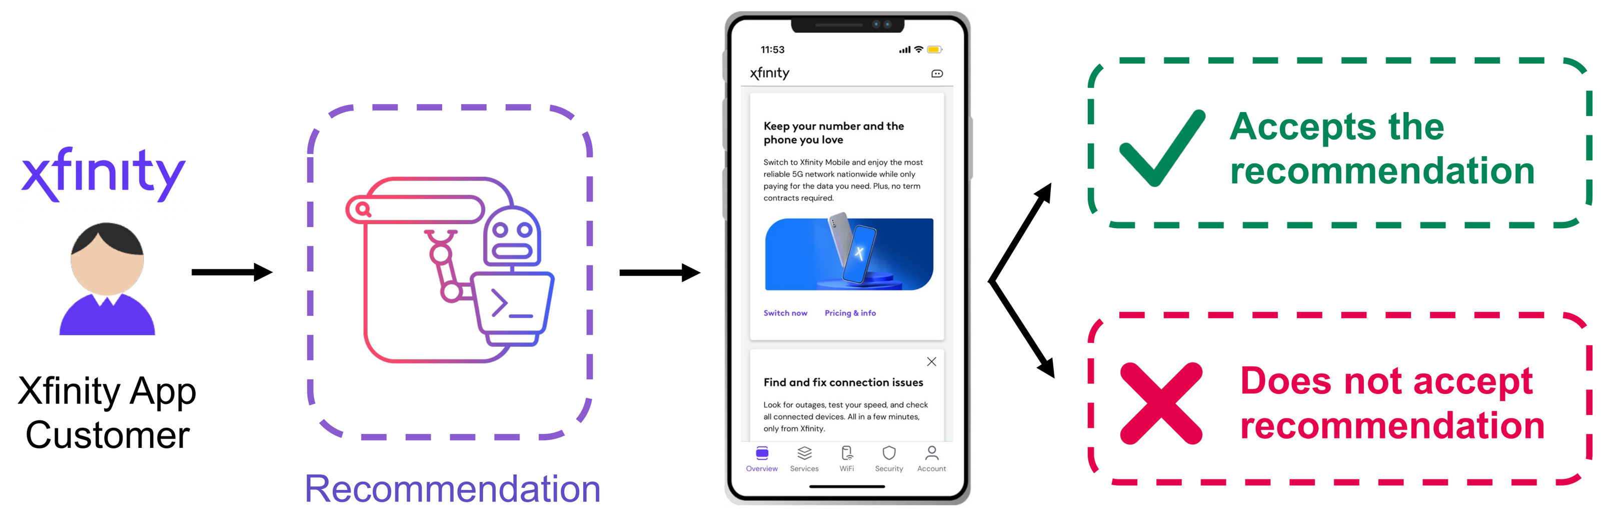

Data

Micro-segmentation Output Features:
Includes segments, probabilities, delivered date and *success flags (target variable)*

Customer Features:
Monthly recurrent revenue/charges, tenure, CLV and contact rates

Product Features:
Product name, New product flags, placement (engagement or sales) and product category

In 2023 Feb, on Xfinity App, there are:

37M recommendations
4M customers
39 unique products

Micro-segmentation Success rate is only **2.3%**

Objective

Goals: The project aims to improve the Nexus recommendation engine with the following goals:

- Improve **success rate** by developing a personalized and interactive approach.
- Use **Reinforcement Learning** to drive this improvement.

Modelling Approach

Reinforcement Learning Framework

We use ***contextual multi-armed bandits model*** which

- uses the customer and product information to make recommendations
- compares its recommendations with historical data
- learns and adjusts its future choices and continually evolves

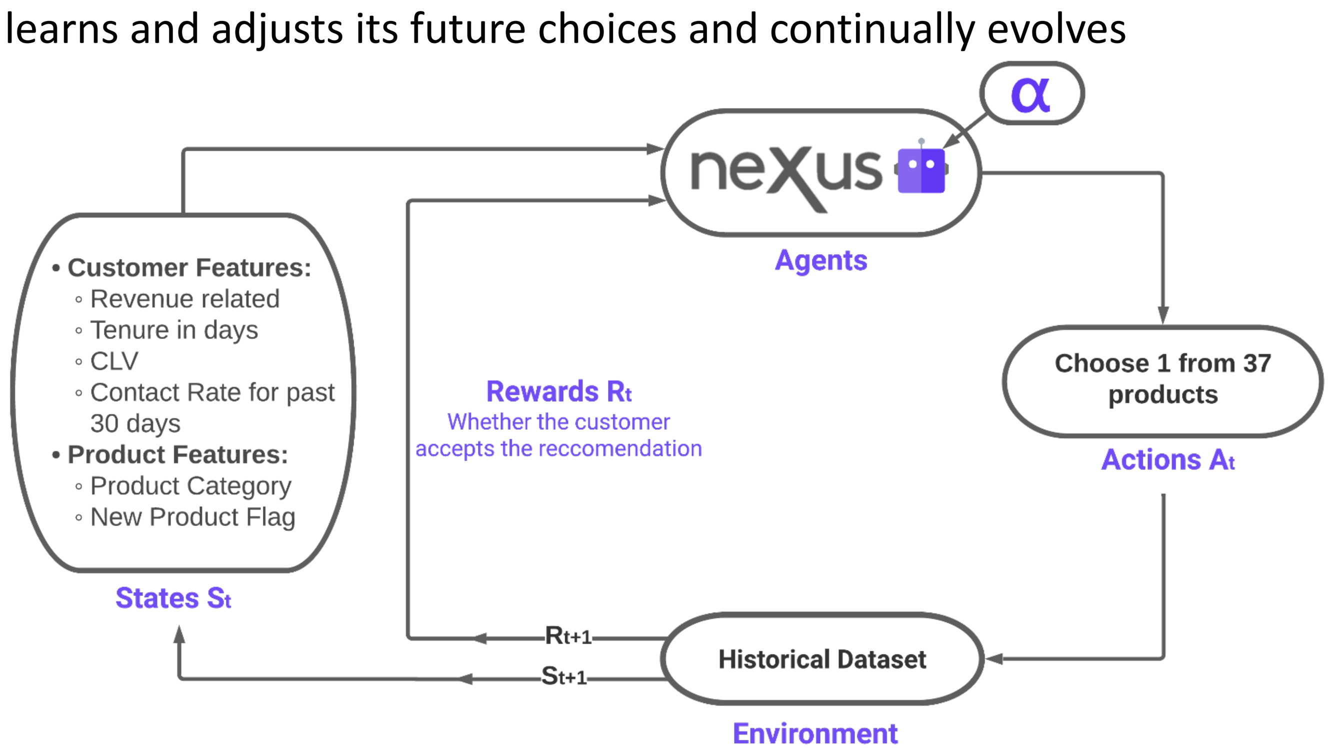

To incorporate contextual information, we use ***LinUCB (Linear Upper Confidence Bound)***, which also balances the ***trade-off between exploration and exploitation*** (through parameter ***α***) while making recommendations based on estimated confidence levels.

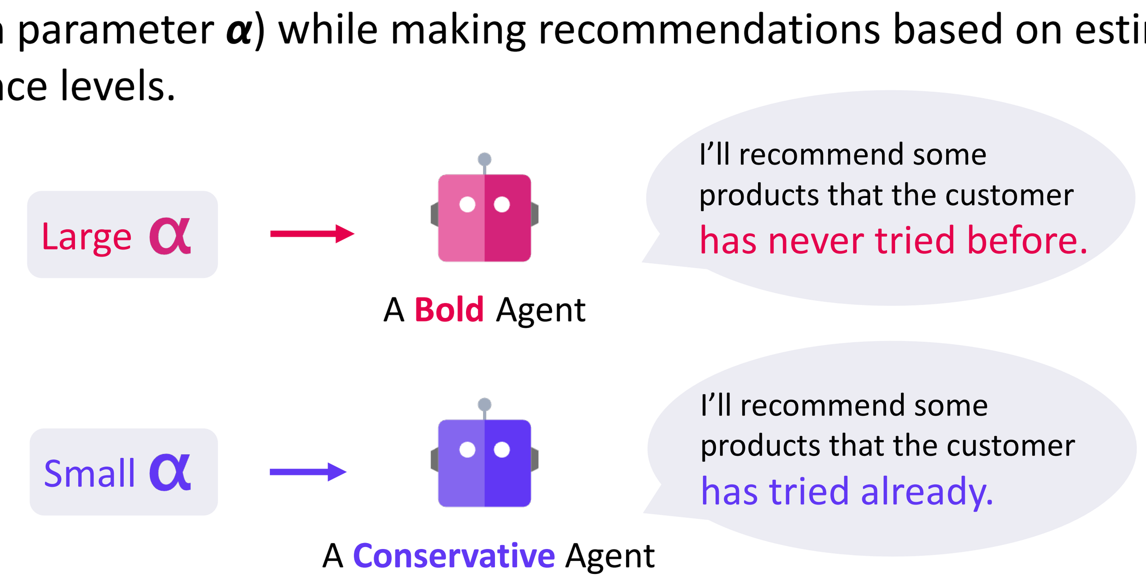

Counterfactual Reward Estimation

Understanding feedback on recommended products is straightforward, but how about unseen ones? We group customers into segments, estimate product preferences for each group, and use these estimates to predict satisfaction. By using these predictions, we create a 'what-if' scenario to understand how well different, untested recommendations might have performed.

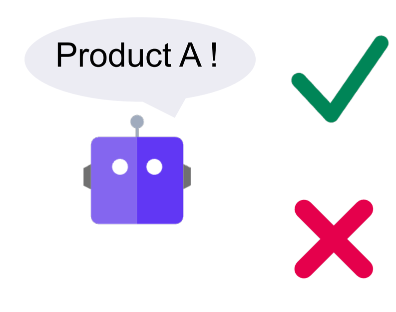

Matches with historical recommendation
- Record customer response (purchased or not)

Does NOT match with historical recommendation
- Estimate customer response

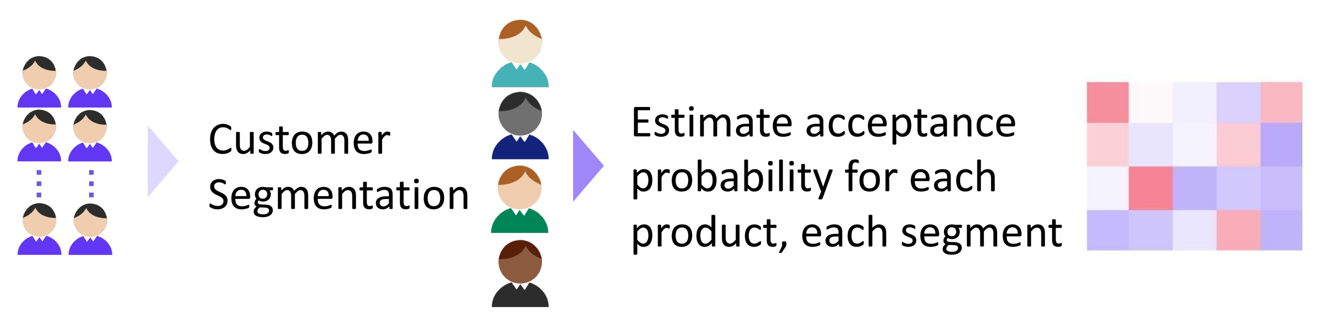

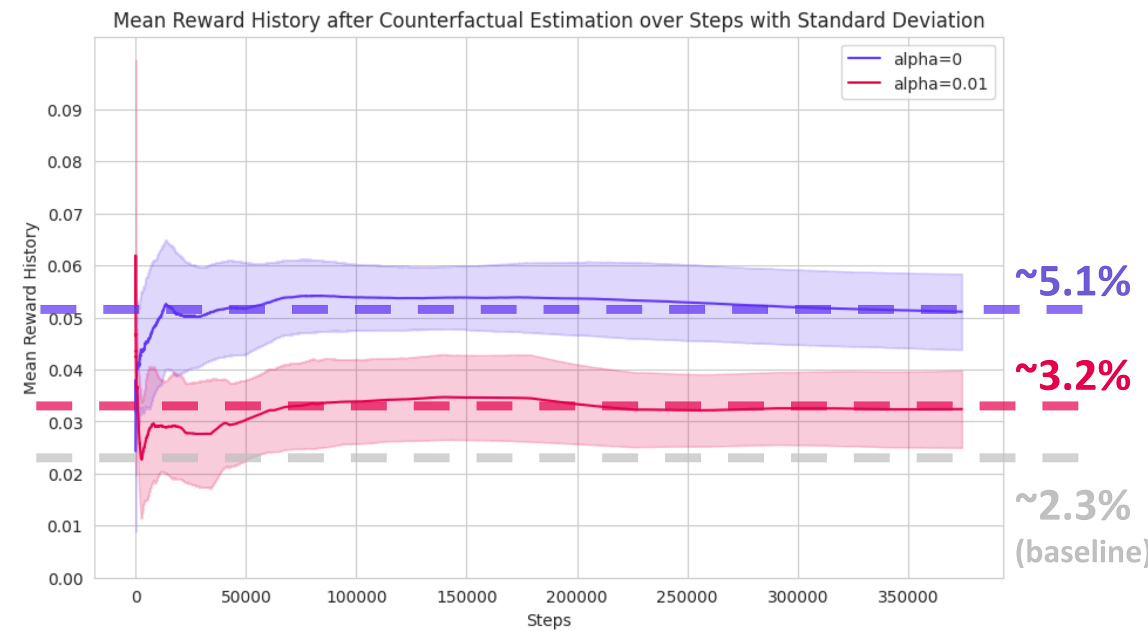

Key Takeaway:

Both models are able to outperform the current baseline, and we will adopt the model with alpha=0.01 because it encourages more exploration.

Results and Impact

Impact

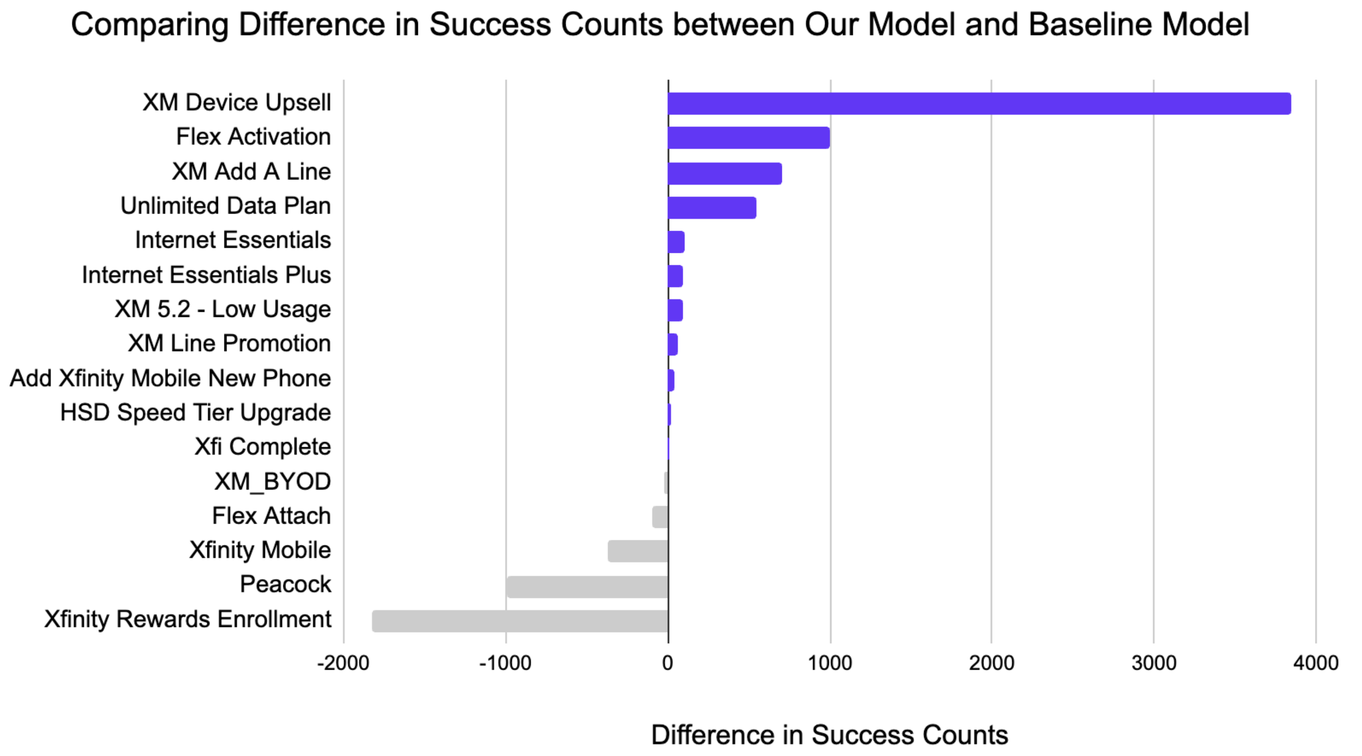

- Provide nearly **400 Million** recommendations per year
- Build **interactive** recommendation model based on feedback
- Adapts to **individual** customer preferences

+39% Success rate

+30M/yr Revenue increase

Future Work

- Conduct A/B test to compare the performance with old model
- Put model in production to get customer's real-time feedbacks
- More advanced model for estimating counterfactual rewards, e.g., Tensor Completion