

The Business Use Case of Reinforcement Learning in Product, Regtech, Marketing

Oneconnect AI Team
Gamma Lab

Reported by

Wenyi Tao

Contents

1 Introduction


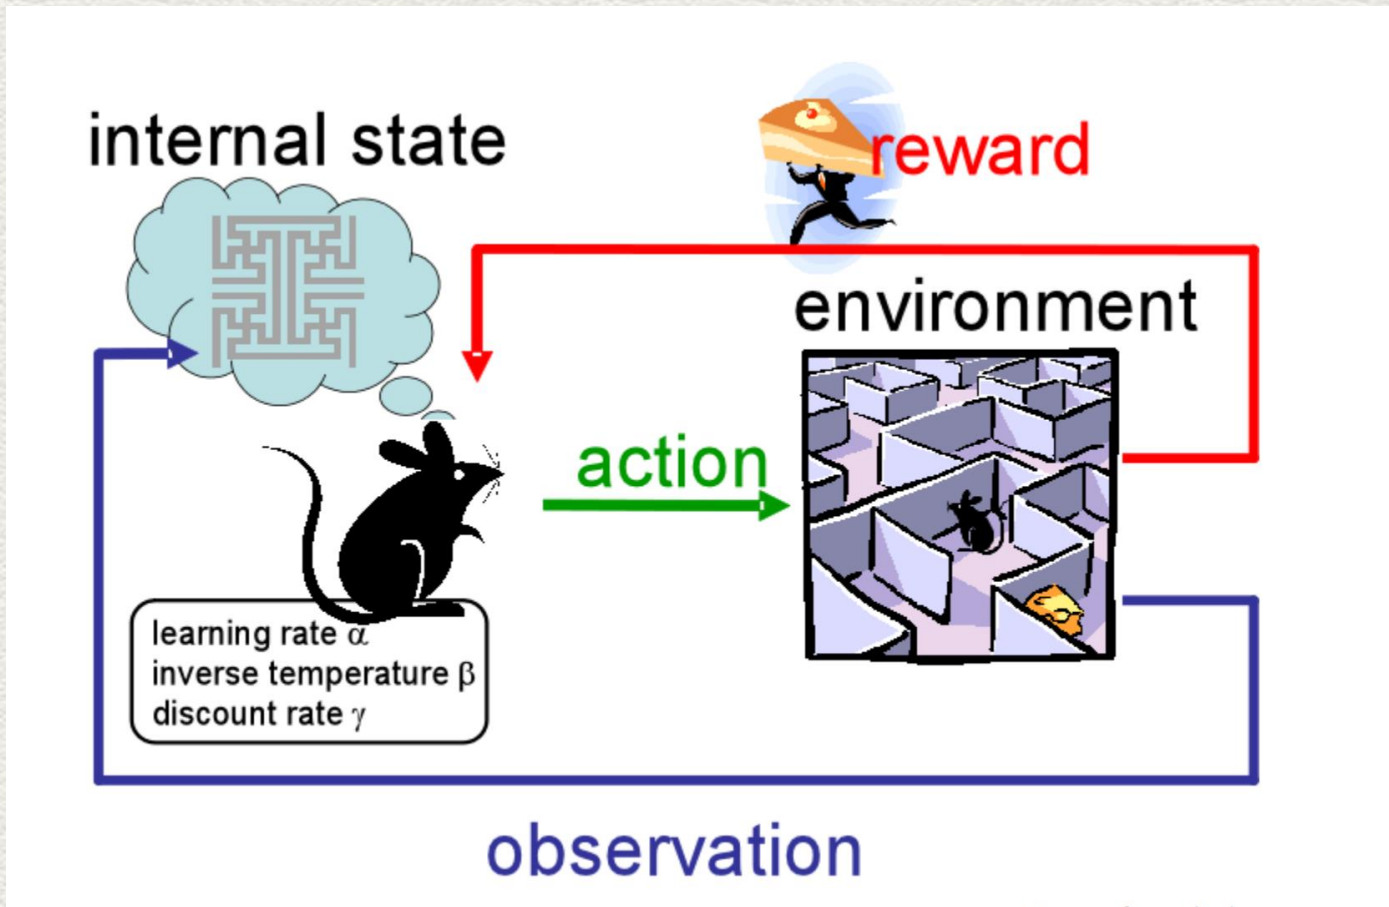
2 Product improvement & Horizon


3 Marketing

4 Regtech

5 Discussion

1 General Introduction





AGENT

current state
action



new state
reward



ENVIRONMENT

$$Q(s_t, a_t) \leftarrow (1 - \alpha) \cdot \underbrace{Q(s_t, a_t)}_{\text{old value}} + \underbrace{\alpha}_{\text{learning rate}} \cdot \left(\underbrace{r_t}_{\text{reward}} + \underbrace{\gamma}_{\text{discount factor}} \cdot \underbrace{\max_a Q(s_{t+1}, a)}_{\text{estimate of optimal future value}} \right)$$

learned value

Contents

1 Introduction

2 Product improvement & Horizon

3 Marketing and Related Research

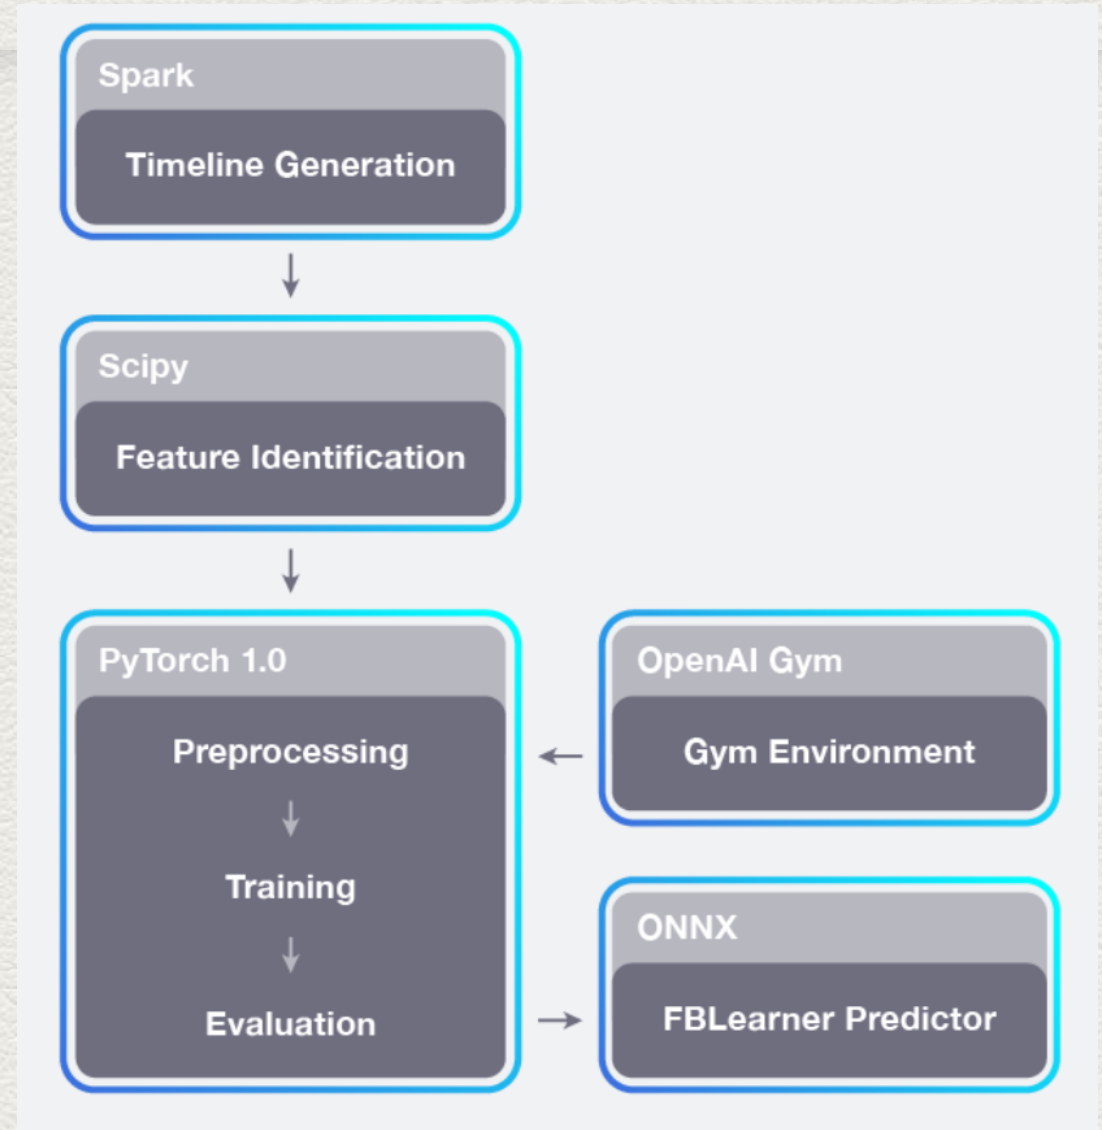
4 Regtech

5 Discussion

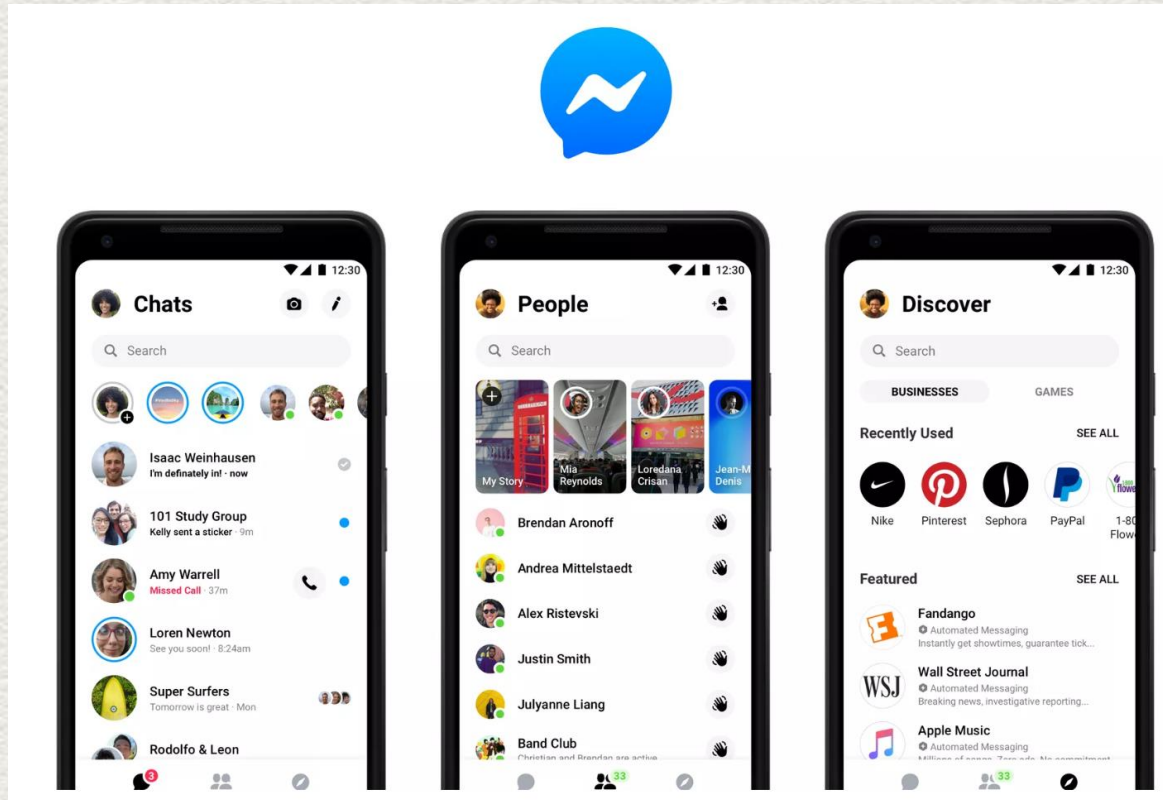
2.1 Horizon and its use case

The first open source reinforcement learning platform for large-scale products and services

— JASON GAUCI, Facebook AI



2.2 Messenger

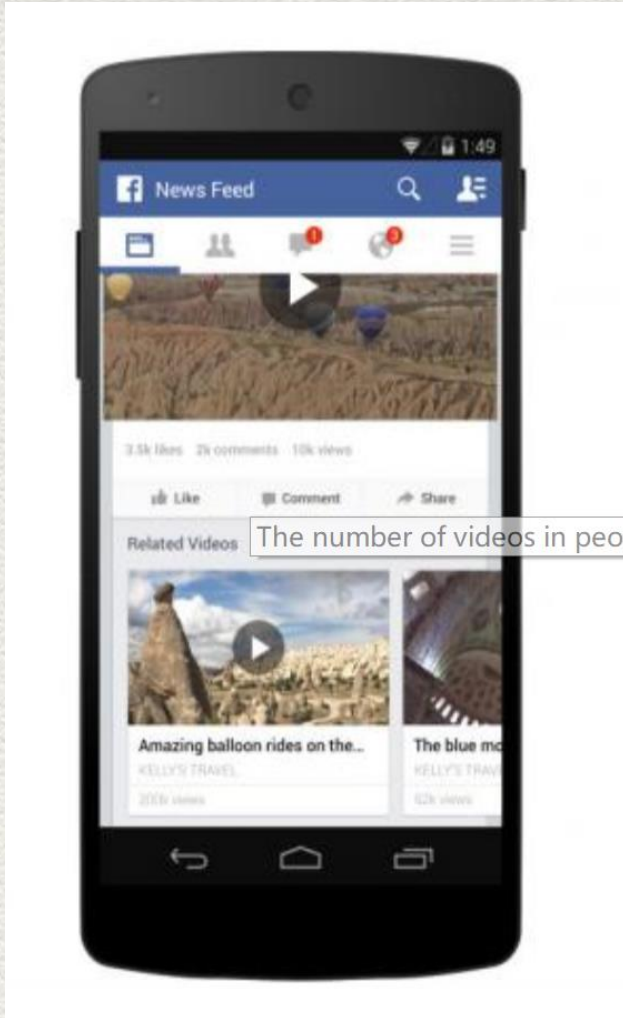


Horizon – Push Notification based on real-time user feedback

Adaptive Personalization:

- If people interact with one suggestion with more regularity, Messenger might surface it more.
- Frequent visitors may not need a notification for a post they would have seen anyway, whereas a less active person could benefit from notifications to make sure they do not miss noteworthy posts from family and friends.

2.3 Video Buffer and Quality



Intelligent video buffering and pre-fetching

Real-Time Quality Adjustment

- Automated decision on whether to send people low or high quality videos depending on variables, like people's cellular connections, whether they are on a subway, or if they have just exited a tunnel.

Contents

1 Introduction

2 Product improvement & Horizon

3 Marketing

4 Regtech

5 Discussion

3 Marketing: Goal of Reinforcement Learning Algorithm



minimizing
budget



increasing
customer
acquisition

3 Marketing

Case 1 Digital Campaign and Personalized Ads

- Determine the best offer to pitch to prospects, an Alternative for AB testing
- Decision based on the reward/return (click through rate(CTR)) in test campaign, the winning ad will be shown to the home page, social media and other partner sites.
- Two separate goals:
 - 1. optimize the total immediate clicks
 - 2. optimize for lifetime value(LTV)

Case 2 Ads Budgets Spending Optimized for Realtime

Real Time bidding:

- MARL (multi-agent reinforcement learning) algorithm to optimize bidding on the largest e-commerce platform

Self-Improved based on Feedback:

- RL algorithms can self-improve their performance even further by aggregating more historic auction data, user feedback, and being challenged with more budget constraints.

3.1 Realtime Digital Campaign

(1) An alternative for AB-testing



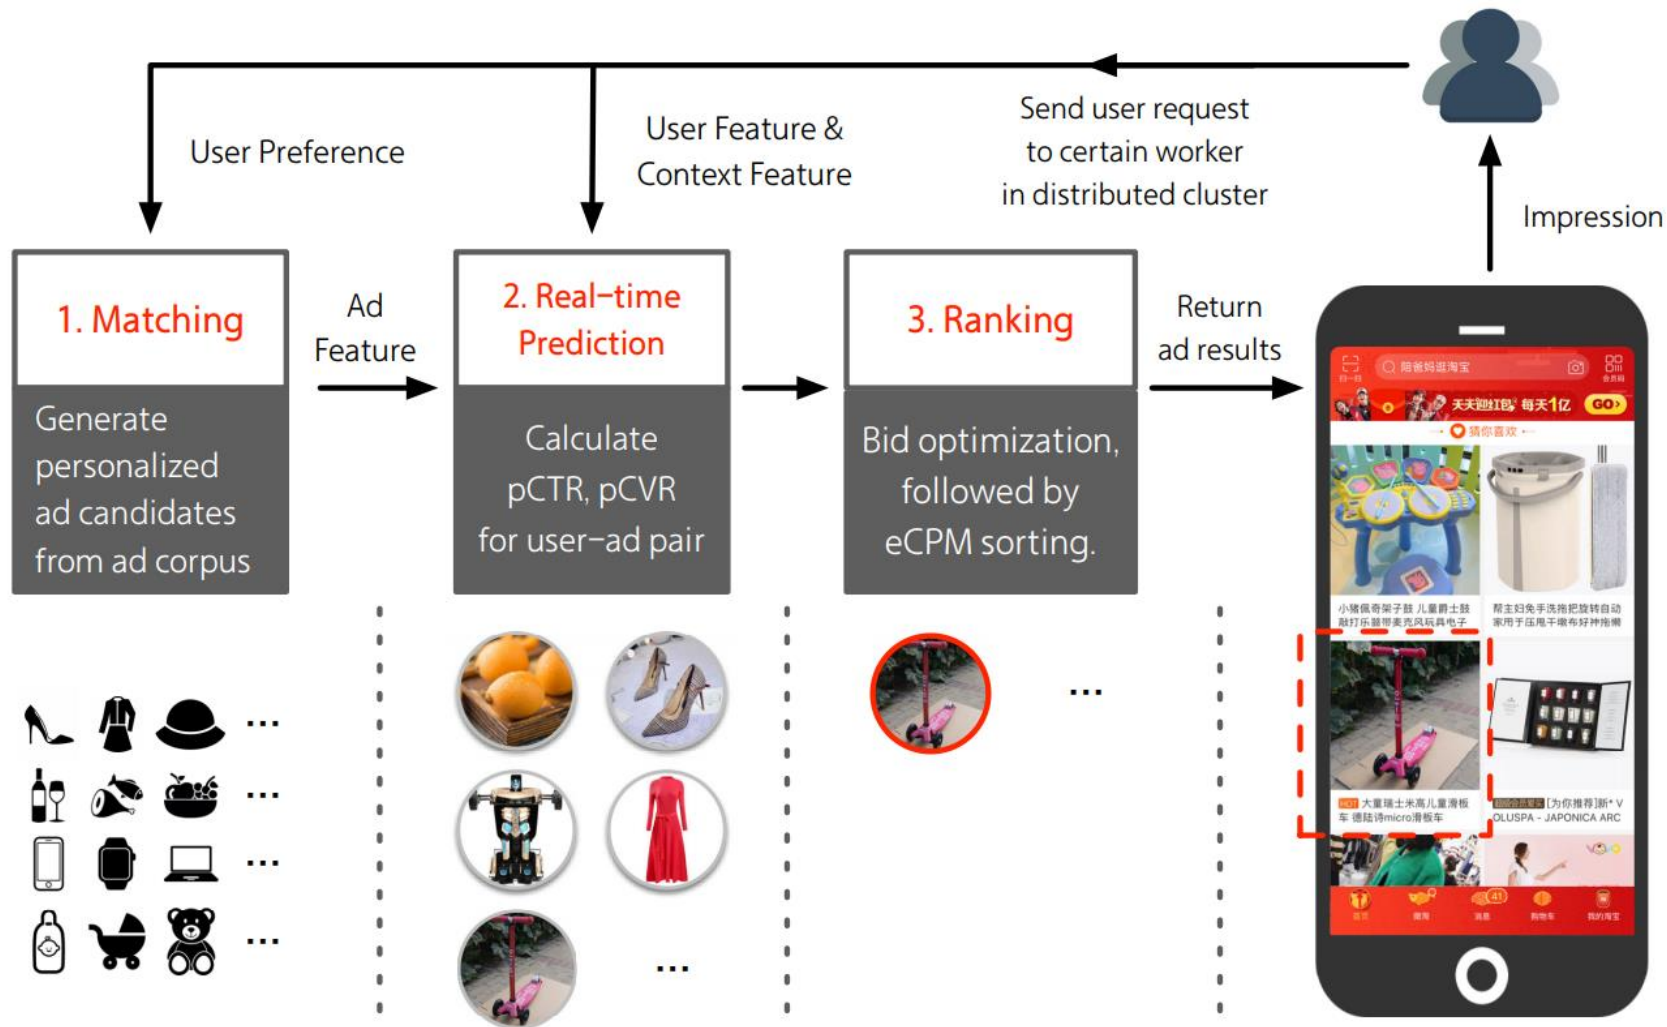
**_Images are taken from

https://www.google.co.in/search?q=creative+credit+card+campaign&source=images&tbm=isch&sa=X&ved=0ahUKEwiml_uilL3XAhWDul8KHd_NDq_Q_AUICigB&biw=1242&bih=614

① Which Version is better?

- Considering Time Cost:
Exploration vs exploitation
Multi-armed bandit problem

3.2 Realtime Ads Bidding with multi-merchants



Contents

1 Introduction

2 Product improvement & Horizon

3 Marketing

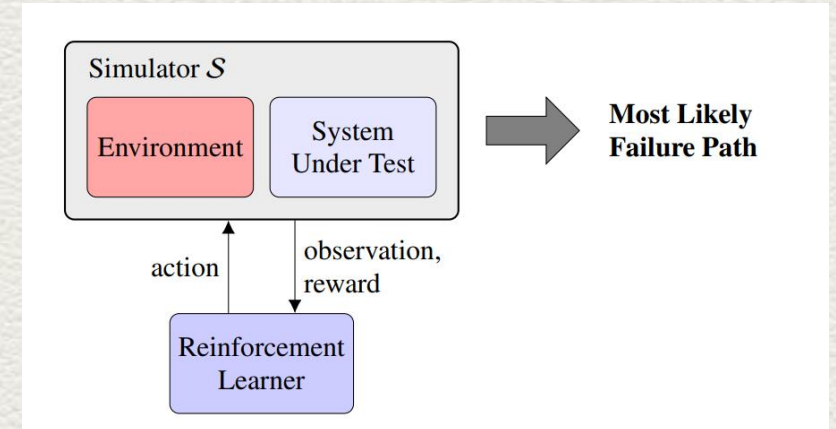
4 Regtech

5 Discussion

4 Regtech

Propose: 1 Adaptive Stress Testing: Find Failure Events With RL

- Sequence of Decision Making with huge Uncertainty
- Searching a simulator for the most likely path to a failure event
- Monte-Carlo Tree Search



Propose: 2 Fraud Detection and Anomaly Detection

- Insurance fraud, credit card fraud and accounting fraud
- detecting insider trading in transaction tables

Thanks

Reported by

Wenyi Tao

5 Reference

1. HORIZON: FACEBOOK' S OPEN SOURCE APPLIED REINFORCEMENT LEARNING PLATFORM

<https://research.fb.com/wp-content/uploads/2018/10/Horizon-Facebooks-Open-Source-Applied-Reinforcement-Learning-Platform.pdf>

2. Real-Time Bidding with Multi-Agent Reinforcement Learning in Display Advertising

<https://arxiv.org/pdf/1802.09756.pdf>

3. Adaptive Stress Testing: Finding Failure Events with Reinforcement Learning

<https://arxiv.org/pdf/1811.02188.pdf>

4. Fraud detection in banking using deep reinforcement learning

<https://ieeexplore.ieee.org/document/8102446>

5. Artificial Intelligence Modeling of Financial Profit and Fraud

http://www.iaeng.org/publication/WCE2011/WCE2011_pp381-383.pdf

6. Confronting Deep Uncertainties in Risk Analysis

<https://onlinelibrary.wiley.com/doi/full/10.1111/j.1539-6924.2012.01792.x>