

Grid Dynamics Portfolio

PERSONALIZING IN-GAME EXPERIENCE USING REINFORCEMENT LEARNING

CASE STUDY OF THE VIDEO GAME PUBLISHER

www.gridynamics.com



JANUARY 2021

Personalizing In-game Experience Using Reinforcement Learning

Business Problem

The client is one of the world's leading video game publishers. Superior user experience is one of the key success factors for them, and the company continuously innovates on game design, in-game personalization, and user engagement strategy and techniques.

The personalization of in-game promotions and special offers is one of the important use cases directly impacting customer experience and the efficiency of in-game monetization. This problem can be approached using a variety of traditional personalization methods, but the publisher was looking to build an innovative solution that overcomes the limitation of standard methods and delivers both superior user experience and marketing efficiency.

The company partnered with Grid Dynamics to design and implement an innovative personalization platform that leverages the latest advancements in reinforcement learning to achieve the above goals.

Business Solution

The project was focused on delivering the following business capabilities:

- Personalize in-game experience to ensure the relevancy of offers and other features to every single user
- Optimize user experience in a strategic context to ensure long-term engagement and satisfaction with the game rather than optimizing immediate outcomes, such as click-through rates
- Provide the flexibility to set various business objectives, such as the improvement of customer lifetime value and offer redemption rates
- Establish a highly automated continuous optimization process that observes ongoing feedback data, updates the personalization model, and evaluates the model using historical data and then in production
- Establish a platform that allows for the rapid onboarding of new marketing and personalization use cases. The platform should provide generic optimization components that achieve the previous three goals for any use case with minimal engineering effort.

These capabilities collectively enable the comprehensive and continuous improvement of user experience; the flexible optimization of in-game revenue, engagement, and retention metrics, and the reduction of engineering costs.

Technical Solution

The business capabilities specified in the previous section could be partly implemented using traditional personalization methods, such as look-alike modeling or collaborative filtering. However, these methods generally require putting a significant engineering effort into each new use case (including feature engineering, model development, and recurring model re-training) and lack strategic optimization capabilities. To overcome these challenges, the team turned to reinforcement learning technology.

Reinforcement learning has gained much attention in recent years mainly because of impressive advancements in the development of self-learning algorithms for board and video games that defeated top professional players in Go, StarCraft, and Dota. In the industry, however, reinforcement learning has a relatively limited use in robotics and some other applications that require exploring the environment and learning from experience. The potential of reinforcement learning for marketing, supply chain, and other enterprise operations was recognized just recently in connection with the following capabilities:

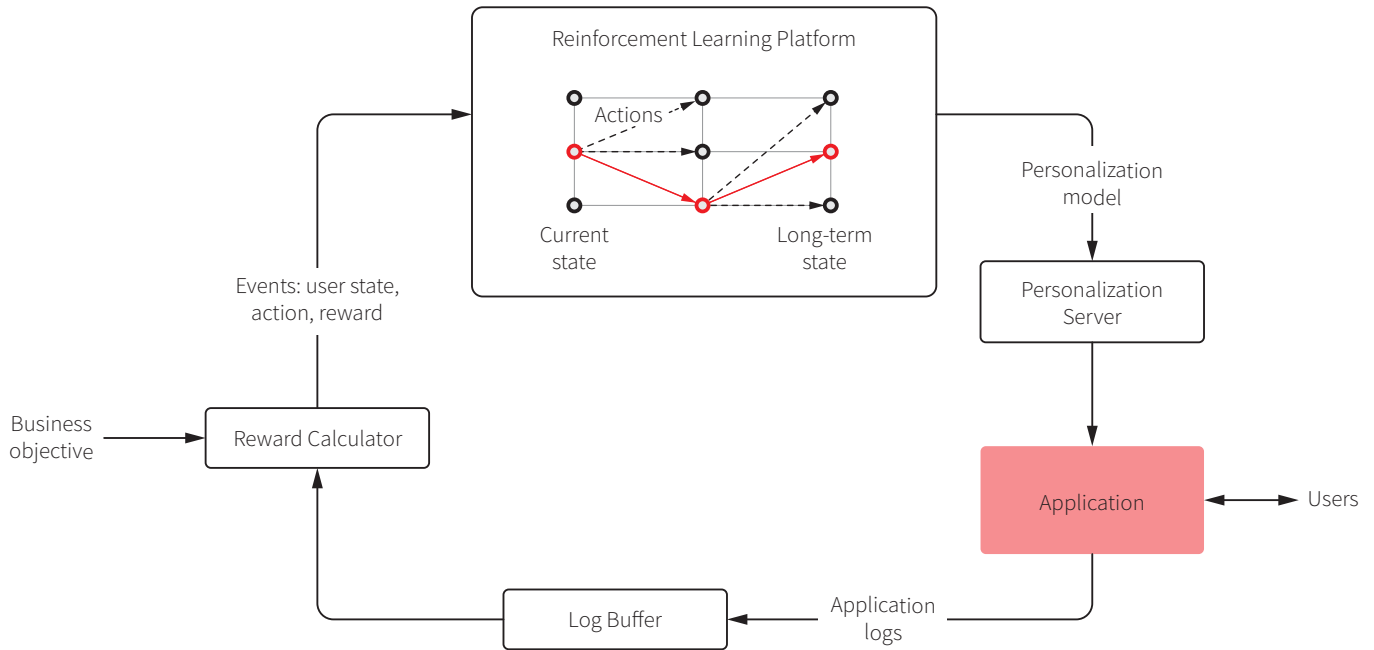
- Reinforcement learning provides a solid framework for strategic optimization, that is, optimization for multiple steps ahead. In marketing and personalization, this helps to determine not just the offer or recommendation with the maximum click-through rate but the optimal sequence of actions maximizing long-term engagement metrics.
- Reinforcement learning algorithms are designed to continuously try actions, observe feedback, and update the decision-making policy in near real time. This allows the system to rapidly accommodate in dynamic environments and initialize itself even if historical data are not available.
- Finally, reinforcement learning methods provide much flexibility in specifying objective functions and balancing long- and short-term goals.

At the same time, many traditional personalization models can be expressed in reinforcement learning terms as particular cases. Consequently, the solution was designed as a multipurpose reinforcement learning platform with the following lifecycle:

- User activity logs are collected from the applications and buffered. The buffer is needed because a user can react to the system's actions with a significant delay.
- Business users define objectives, such as the optimization of in-game revenue or engagement metrics, and user activity records are attributes with the corresponding rewards.
- The reinforcement learning platform replays the batch of events and updates the policy (personalization model) that determines the optimal action based on the user state.
- The new version of the policy is deployed into production, and the personalization server uses it to make personalization decisions. The process is then repeated.

This conceptual design is shown in Figure 1. Although conceptually simple, it supports a wide range of marketing use cases and can outperform most traditional personalization methods.

Figure 1. A high-level architecture of the reinforcement learning-based personalization system.



Business Outcome

The first version of the platform and the pilot use case were implemented in just about eight weeks. Production trials demonstrated the high efficiency of the new solution, delivering up to 25% improvement compared with the baselines. The new platform created a solid foundation for experimenting with various objective functions, balancing strategic and myopic targets, and quickly onboarding new use cases to create superior gaming experiences.

About the Client

The client is a leading video game company that develops and publishes interactive entertainment for PC, console, and mobile platforms.

The company entered the video game industry in the early 1990s and currently employs more than 5,000 people. The studio’s portfolio includes multiple critically acclaimed franchises. The company has multiple offices worldwide and holds strong market positions in the US, Europe, and Asia.

Solution Highlights



Use Cases

In-game experience personalization, in-game revenue maximization



Technical Features

Strategic optimization, online optimization, customizable business objectives, easy use case onboarding, low maintenance effort



Business Outcome

Up to 25% dollar-per-user improvement compared with the baselines



Required Data Sources

In-game events and purchases, promotion metadata



Implementation

Open source libs and frameworks



Infrastructure

Public cloud



Pilot Project Timeline

8 weeks

Grid Dynamics Inc.

5000 Executive Parkway, Suite 520
San Ramon, CA

650-523-5000
info@griddynamics.com
www.griddynamics.com

The information in this document is provided as is and does not warrant the accuracy, completeness and fitness for a particular purpose. In no event shall Grid Dynamics International be liable for any damages whatsoever arising from your reliance on any information contained in this document. Except as permitted under the Copyright Act no part of the document may be reproduced, transmitted, stored in a retrieval system, or translated into any human or computer language in any form by any means, electronic, mechanical, magnetic, optical, chemical, manual or otherwise without the expressed permission of Grid Dynamics International.