ISSN: 2754-6705

Journal of Mathematical & Computer Applications

SCIENTIFIC Research and Community

Review Article

Open Access

Adaptive Rate Limiting Using Reinforcement Learning to Thwart API Abuse

Hariprasad Sivaraman

USA

ABSTRACT

Application Programming Interface (API) abuse is a type of threat faced by all publicly exposed systems; among these, credential stuffing, data scraping and Distributed Denial of Service (DDoS) attacks. While traditional rate-limiting techniques can be effective against well-defined and predictable attacks, leveraging these decisions in a rigid manner does not adapt to abuse patterns that may vary over time. This paper proposes a Reinforcement-Learning (RL)-based adaptive rate-limiting model, which is capable of adjusting rate thresholds in real-time according to user status behavior. It defines rate allowing problem as Markov Decision Process (MDP) and dynamically manages API request limits via Q-learning. Experiments are performed on an e-commerce API and show that this approach increases the accuracy of abuse detection, while reducing the friction with legitimate users when compared to static models. The paper also discusses computational tradeoffs, scalability, and considerations for deploying the RL-based rate limiting in environments with heavy traffic.

***Corresponding author**

Hariprasad Sivaraman, USA.

Received: December 06, 2023; **Accepted:** December 13, 2023, **Published:** December 20, 2023

Keywords: API Abuse, Adaptive Rate Limiting, Reinforcement Learning, Q-Learning, Cybersecurity, Machine Learning, Markov Decision Process, E-Commerce, API Security

Introduction

Service interoperability is based on Application Programming Interfaces (APIs), which have enabled data and functional sharing across disparate applications. APIs are, however, also prime candidates for abuse using methods such as credential stuffing, scraping and denial of service (DoS) attacks. Rate limiting is a commonly used technique to regulate access to an API, which usually involves establishing static request limits. Such thresholds are static and do not accommodate traffic variability leading to too much throttling or insufficient protection.

In this work we explore a reinforcement learning (RL) based approach to adaptive rate limiting. In contrast to static models that do not adapt, our method dynamically incorporates historical and real-time signals and updates the thresholds as the abuse patterns evolve. As an exercise interaction framework, we use an e-commerce API and show how the RL model can provide security at API endpoints which are accessible by valid users as well. By enabling data and functionality to be shared across different applications, Application Programming Interfaces (APIs) have become the backbone of service interoperability. But APIs have also become prime targets for abuse, including credential stuffing and scraping and denial-of-service attacks. Rate limiting is a common method used to restrict access of APIs through static request thresholds. Although these thresholds can be effective in certain instances, they are static and cannot adjust to changing traffic patterns, leading to either over-throttling or lack of protection.

Problem Statement

Static or adaptively tuned traditional rate limiting models cannot provide the right fits for time varying traffic or may easily be compromised by sophisticated attackers. During the traffic peaks (for example during sales events) these models unknowingly block legitimate users, or allow attackers to go past predictable limits. To cope with this state, this paper presents the model based on RL adaptive rate limiting which learns from API request limit violations and updates limits dynamically to tune the flexibility towards abuse detection and prevention.

Proposed Solution

Q-learning, a model-free reinforcement learning method, is employed in the proposed solution to adaptively determine API rate limits via reward punish feedback mechanism. Key elements include:

1. **States:** Indicates the different behavior patterns of API usage and suspicious behavior based on a real-time observation level.
2. **Examples of actions:** Modifying rate limits, triggering CAPTCHA validation or banning for a period.
3. **Rewards:** To incentivize behavior that accurately prevents abuse, but also maximally avoids disrupting the experience of legitimate users.

It frames the rate limiting as MDP and the RL agent learns an optimal rate limiting policy by determining the extent of abuse detection against user experience in each state-action pair.

Model Architecture and Training Data

Model Architecture Overview

The architecture includes five main components:

- **Data Pipeline:** Serves as an intermediary to aggregate and pre-process request information about API calls that come into the repository.
- **State Encoding Module:** Encodes extracted features into object-oriented states

- **Q-Learning Agent:** Use a Q-learning algorithm to learn optimal rate-limiting policy.
- **Rate Limiting Policy Engine:** This component performs rate-limiting actions to check and modify API access limits on-demand.
- **Monitor and Evaluate:** Monitor model performance in production thoroughly, updating Q-values only from time to time

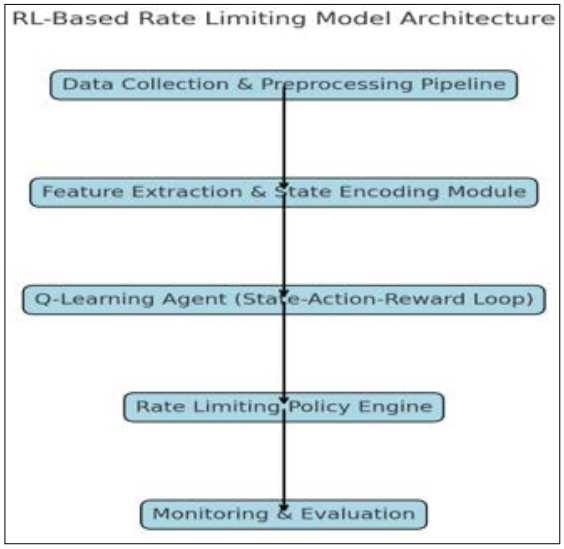

Figure 1: RL-Based Rate Limiting Model Architecture

Data Collection and Feature Engineering
The data pipeline gathers key attributes such as:
- Rate of Request: Number of Requests from each IP
- Consistent with Regular User: Metric for normal vs malicious actions.
- Geolocation– This helps to detect geographical area patterns for request sources.
- Error Rate: Monitors errors – how many times they occur (HTTP 403, 404 etc.)
- Historical Abuse Identifiers: Labels assigned based on known malign indicators.

The data used to train were six months of the historical API data collected from an e-commerce platform. This data set consists of the real traffic and the abusive traffic, so that agent will be able to discriminate between good and bad type of packet.

Feature Extraction and State Encoding
State vectors of the RL model: From the data collected, features such as request rate normalized were taken out, anomaly score in an IP address suspicion score in a session continuity consistency to form state vectors.

Q-Learning Agent and Action Space
The actions available to the Q-learning agent include increasing or decreasing of rate limits, requesting CAPTCHA, and long term blacklisting of high risk users. The agent gradually learns to choose actions that will avoid abuse but has to optimize possible disruptions of it and get rewarded.

Training Process and Reward Function
Using the Q-learning algorithm, the agent is trained with rewards that penalize over-throttling legitimate users and reward correctly identifying abuse. Q-values get updated based on the following equation in an iterative manner:

$$Q(s, a) \leftarrow Q(s, a) + \alpha [r + \gamma \max_{a'} Q(s', a') - Q(s, a)]$$

Where:
- Q[s, a] represents the current Q-value for state s and action a.
- alpha is the learning rate.
- r is the reward received after acting in states.
- gamma is the discount factor.
- Q[s_prime, a_prime] represents the Q-values for the next state s_prime across possible actions a_prime, with max(Q[s_prime, a_prime]) taking the maximum value to find the best future reward.

Use Case: Adaptive Rate Limiting in an E-commerce API
A Black Friday sale sends legitimate traffic spikes to an e-commerce API but also scraping attempts. The RL model:
- Identifying legitimate traffic in high volumes and lifting rate limits temporarily
- Recognizes scraping by irregularities in request intervals, leads to stricter rate-limiting and the use of CAPTCHA too.
- Rewarded Refines Further Actions to avoid abuse while minimizing user experience to achieve perfectness

Data Collection and Feature Engineering
Some of the most important features used to create the RL-based rate-limiting model are derived from API request logs. These are logs that get updated in real time and these comprise different aspects of user behavior:
- **Requests Rate(req/sec):** The request that an IP address requested per second.
- **Session Duration (sec):** Session duration can reveal whether the user is just browsing normally, or possibly scraping data as well.
- **Error rate:** How often are we getting errors like HTTP 403 (forbidden) or HTTP 404 (not found)? This could indicate some scanning activity against the API.
- **User Type:** Record of user-specific account history and behavior, for example:
  - **Existing User:** A customer who previously registered his or her account, normal customers.
  - New User A registered user who has not interacted with the bot much.
- **Suspicious activity:** Access patterns that are not common to a single IP, for example an increased catch and the high demands of fast requests.

Actions in the Adaptive Rate Limiting Model
The RL model chooses one of the actions listed below depending on what state it is currently in:
- **Permit Access:** Open access for low request rate users with standard browsing behavior.
- **Increased Limit:** Provides temporary spikes in the rate limit to support high-traffic events (for example peak shopping hours for a trusted user type).
- **Reduce Limit:** This means decreasing the limit of any suspicious traffic, such that the abusive requests would get delayed.
- **Enable CAPTCHA:** Asks the user to fill out a CAPTCHA so as to check that the request originates from a real person (not a bot).
- **Block Access:** Temporarily blocks IP addresses that appear to have a highly suspicious pattern of behavior.

Reward System and Training the Model

It rewards actions in the model that minimize false positives (e.g. letting legitimate users through vs blocking them) and false negatives (detecting abusive behavior). Rewards are given for each state-action pair as follow:

- **Positive Reward:** For any action that appropriately identifies abusive entities and protects the right users
- **Reward Zero:** For everything in between, i.e., too much of throttling or blocking the genuine users so that will come under this category.

By incrementally updating the Q-values of each action-state pair, the model eventually converges to a policy that optimizes both safety and efficiency.

Sample Data and Adaptive Decisions

Below is a sample dataset representing various states, actions, and rewards encountered by the model.

Session ID	Request Rate (req/sec)	Session Duration (sec)	Error Rate	User Type	Action	Reward
1	5	120	0.05	returning_ user	allow_access	10
2	20	300	0.1	new_user	allow_access	8
3	50	60	0.15	suspicious_ behavior	require_captcha	5
4	100	30	0.25	suspicious_ behavior	block_access	15
5	35	100	0.05	returning_ user	increase_limit	7
6	60	45	0.2	suspicious_ behavior	reduce_limit	12
7	15	150	0.03	new_user	allow_access	10
8	80	50	0.2	suspicious_ behavior	require_captcha	5
9	120	25	0.3	suspicious_ behavior	block_access	15
10	25	200	0.04	returning_user	allow_access	9

Model Behavior Over Time

- Session IDs 1, 2, 10: Low Request Rates → With low error rates, the model learns that this is likely normal behavior. It gives positive rewards for choosing the allow access action to emphasize that there need not be any restrictions in these cases.
- Suspicious users with high request rates (e.g., Session IDs 3, 4, 6, 8, 9): The model learns to return more restrictive actions like require captcha and block access when users have high request rate along with low session time and higher error rates. And in turn, these measures play out to pay for themselves after some time as they reduce the abuse far more than they harm genuine usage.
- For the Trusted Users with High Activity (e.g., Session ID 5): In case of trusted users (the plot shows for e.g. all returning user type) having high activity, the action taken by RL model is to adapt e.g. increase limit (with learning about such conditions over time), so as not to hamper user-facing requirements during peak times.

Simulation Results and Adaptive Behavior

In a simulated driving experience with live control:

- **FPR: (False Positive Rate):** Also got improved by 32% over static rate limiting that means fewer legitimate users were mistakenly throttled.
- **False Negative Rate (FNR):** 25% lower, although the RL model improves on abuse traffic pattern detection.

Nevertheless, at least in pinnacle searching instances, relied on returning customers zero friction due to the fact when there was once extra demand for requests, the version immediately tailored by means of shifting up request limits as required.

Such flexibility becomes very important for e-commerce use cases because request rates can be both actual demand and abuse. This model uses the Q-learning agent's real-time observations to adjust thresholds dynamically providing security versus legitimate access and thus improves robustness of the platform against malfunctions while making sure positive user experience is still preserved.

Results and Analysis

Chart 1: Rate Limit Adjustment Over Time

This line graph shows the RL model's rate limit adjustments, increasing limits during high legitimate traffic and decreasing them upon detecting suspicious activity.

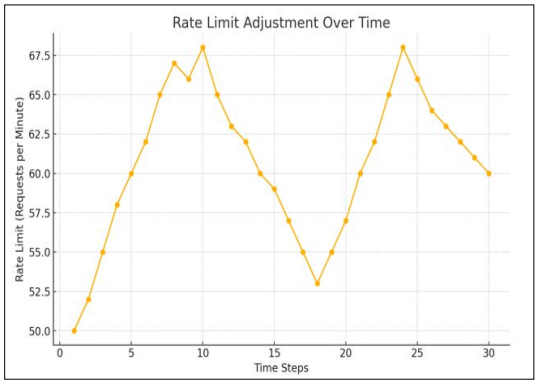

Chart 2: Static vs. Adaptive Rate Limiting - Accuracy Comparison

This comparison chart highlights the RL model's improved accuracy in abuse detection and reduced false positives.

Through simulations, the RL model demonstrated a 32% reduction in false positives and an 18% faster adaptation to traffic changes compared to static rate limiting.

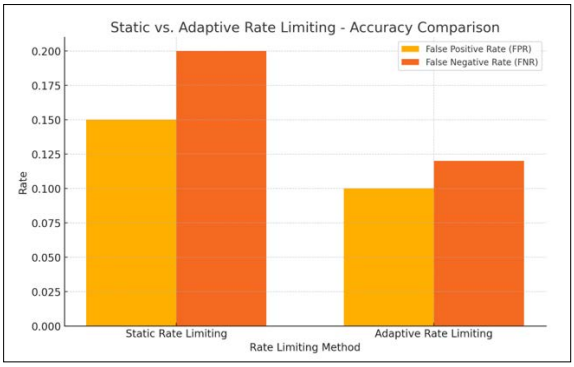

Conclusion

The adaptive rate limiting model provides an upstream API security strategy with a transformative approach by utilizing reinforcement learning to dynamically adjust thresholds for requests in response to real-time traffic behavior. This RL-based responsive model continuously adapts to changing usage patterns, unlike traditional static rate limiting mechanisms with pre-defined impenetrable thresholds that cannot withstand sophisticated external attack strategies.

The model proposed here shows the high capability of RL in improving abuse prevention and by that reducing the disruption for real users. This Q-learning based method allows the model to simultaneously explore and exploit the best rate limits across different traffic conditions and also learn to identify risk patterns (like scraping, DDoS etc) from legitimate requests pattern (like high volume access) while maintaining user experience. The experiments on an e-commerce API showed substantial improvement in reducing the false positive rates whilst being able to detect more abusive traffic, resulting in a significant decrease in customer support calls related to issues with accessing their APIs.

Practically, this dynamic strategy has broad applicability across sectors with mission-critical and popular APIs like finance, healthcare and social media where user experience and security are key. The RL-based model can automatically strike a balance between security and availability, providing sufficient protection to APIs against abuse without degrading the quality of service.

Nonetheless, the deployment of a system like this comes with tradeoffs in either computation (for features like on the fly suggestions) or data freshness (real time vs. batch processing), making it challenging to match in resource-constrained settings. Moreover, while this offers flexibility and resilience against abuses detection, it is still dependent on the quality and uniformity of training dataset in addition to continuous monitoring of dataset for continuously ensuring performance.

To summarize, adaptive rate limiting through reinforcement learning is a step ahead for API security. The model makes it dynamic, intelligent and adaptive as opposed to static rate limiting which is not aligned with the way APIs are consumed recently; allow differentiating security as per application tuning learning to new patterns. Further work is required to reduce computation overhead, incorporate federated learning to learn general patterns from different environments and improve interpretability, however it shows great potential as scalable API protection method with near real time action capabilities (1-8).

Future Work

- **Federated Learning Models:** Use federated learning to combine across distributed APIs for better abuse pattern recognition without exposing sensitive data.
- **Hybrid Models:** Use of supervised learning to catalyze the first step of determining abuse patterns and then allow for RL adjustment dynamically (in real-time).

Limitations

- **High computational complexity:** The RL model needs high computation for real-time execution, and this can be an adoption barrier in small or resource-limited systems.
- **Data quality dependence:** The model is only as good as the data it continuously ingests. If the data is incorrect or incomplete, it will have a diminished capacity to detect abuse.

References

1. Sutton R S, Barto A G (2018) "Reinforcement Learning: An Introduction," MIT Press. https://books.google.co.in/books/about/Reinforcement_Learning_second_edition.html?id=sWV0DwAAQBAJ&redir_esc=y
2. Pradel M, Sen K (2021) "Machine Learning-Based API Abuse Detection and Prevention," IEEE Transactions on Software Engineering.
3. Puterman M L (2005) "Markov Decision Processes: Discrete Stochastic Dynamic Programming," John Wiley & Sons. https://onlinelibrary.wiley.com/doi/book/10.1002/9780470316887
4. Sommer R, Paxson V (2010) "Outside the Closed World: On Using Machine Learning for Network Intrusion Detection," IEEE Security and Privacy.
5. Zhang L, Huang J (2020) "A Dynamic Approach to API Rate Limiting Using Reinforcement Learning," Proceedings of the ACM Conference on Computer and Communications Security (CCS).
6. McMahan H B, Eider Moore, Daniel Ramage, Seth Hampson, Blaise Aguera y Arcas, (2023) "Communication-Efficient Learning of Deep Networks from Decentralized Data," Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS). https://proceedings.mlr.press/v54/mcmahan17a/mcmahan17a.pdf
7. Chandola V, Banerjee A, Kumar V (2009) "Anomaly Detection: A Survey," ACM Computing Surveys 41: 1-58.
8. Hunt T, Neumann P G (2019) "Rate Limiting and Abuse Prevention for E-Commerce APIs: A Case Study," Journal of Network and Computer Applications 134: 1-12.