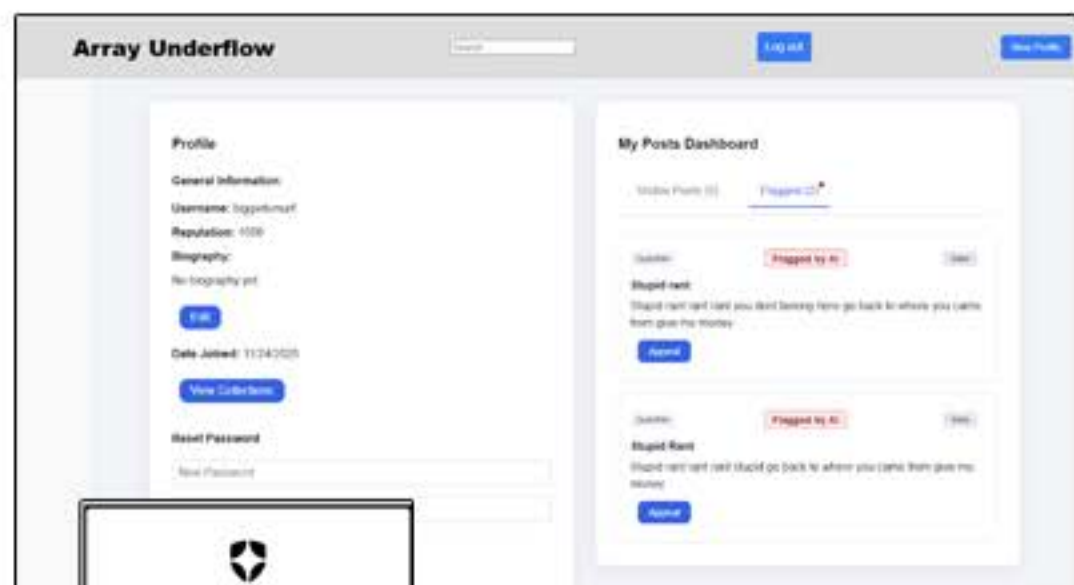


# CS4300 Final Project: "Array Underflow"

## Group 202: Aarush, Manish, Reagan, Tarun

### Project Description

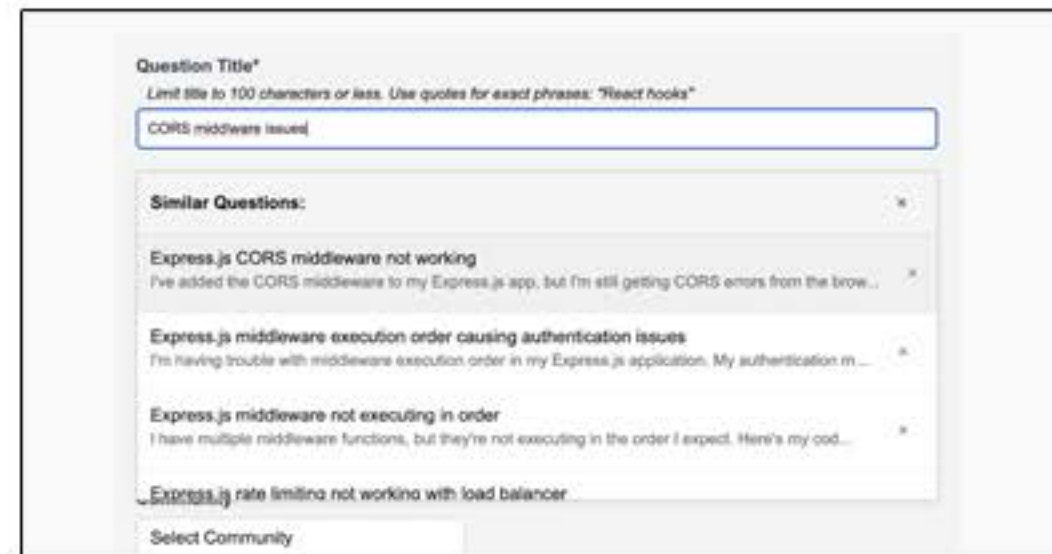
This project delivers three integrated platform features that enhance safety, usability, and community engagement. An AI-powered content moderation system automatically analyzes all user-generated content for policy violations while supporting secure authentication, transparent appeals, and human review. An intelligent autocomplete system improves the question-posting experience by suggesting relevant previously asked questions from communities the user follows, updating results every 300 ms, supporting quoted exact-match queries, and handling edge cases like empty inputs and extra spaces. A dynamic bounty system further boosts content quality by allowing users to attach non-reversible, time-limited reputation rewards to their questions, with automated reputation management and engaging countdown and celebration UI interactions.



Dashboard of flagged posts



Auth0 Login



Similar Questions

### Tech Overview and Design

We implemented three major features in the Array Underflow platform. First, we built an AI-driven moderation workflow using a queue-and-batch pipeline where new posts enter a Moderation Queue and are processed every minute by Google's Gemini 2.5 Flash Lite, which creates Flag objects and automatically hides harmful content. This system includes new Flag and Appeal models, Auth0-secured moderator routes, real-time UI updates, and a moderator dashboard.

Second, we added an embedding-based autocomplete feature that triggers about 300 ms after the user stops typing. The backend retrieves the user's communities, handles quoted phrases, and uses Voyage-3-Large embeddings with MongoDB Vector Search to return the top five similar questions.

Third, we introduced answer voting and a bounty system supported by socket events for live updates. Bounties have their own schema, adjust user reputation immediately, and award the answerer when redeemed.

Across all features, we chose simple and maintainable designs to support clarity, correctness, and rapid iteration during development.

### Future Work

Several extensions could further strengthen and scale our platform. The moderation system could be expanded with ban/unban/timeout actions, a structured moderator application process, an admin role with elevated privileges, cleaner dashboard UX, and implementation of our threat-assessment mitigation plans. The autocomplete system could be improved by replacing the current naïve search with more efficient filtering, caching, and benchmarking alternative embedding models to improve result quality at scale. For bounties and reputation, future additions could include AI-driven difficulty classification to auto-set bounty amounts, a sitewide reputation leaderboard, and a badge or ranking system similar to Duolingo to reward user progress and encourage engagement.



Bounty placed on a user's question