Cattlelog Summer Engineering Brief

September 25, 2025 - Written by Jake Roggenbuck

This document is titled “brief”, in that it’s meant to only scratch the surface of the engineering done this summer for Cattlelog. This write-up cannot be all-inclusive of the work done, but it tries to give a good overview.

The audience of this document is rather general. The main purpose is to give context with examples to our new PM, engineer, and product designer. It’s written generally enough that anyone else who is curious will likely be able to follow along. Feel free to ask questions if they do arise.

Thank you to the entire team who worked during the summer, ideating, designing, marketing, and engineering these features. I appreciate all of your contributions.

Metrics since June

Since we took a break for the summer at the start of June (finals week), we’ve gained 8,978 unique users. We saw massive spikes during different key dates e.g. when pass times were released, pass 1, pass 2, open registration, and anytime we pushed out a social media post.

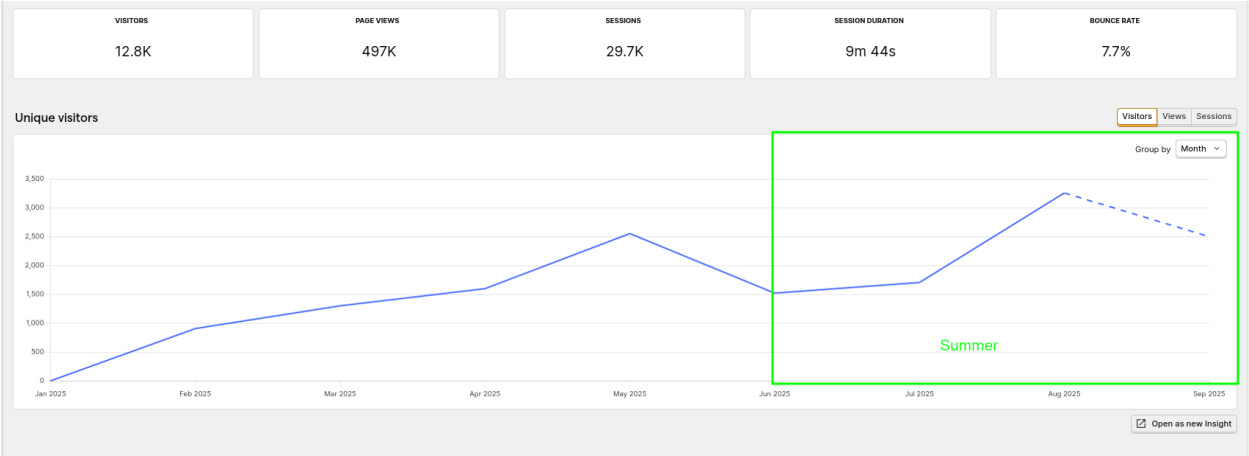

Projects

The few largest projects we worked on this summer included the new landing page, the grade distribution page, find a tutor button, professor search (beta), our gzip data cache called Neptune, two browser extensions, a new blog page, a professor timeline feature, an AI feature coming soon, and dozens of optimizations + UI improvements.

1. New Landing Page

Previously, our landing page included a search on the left and buttons on the right:

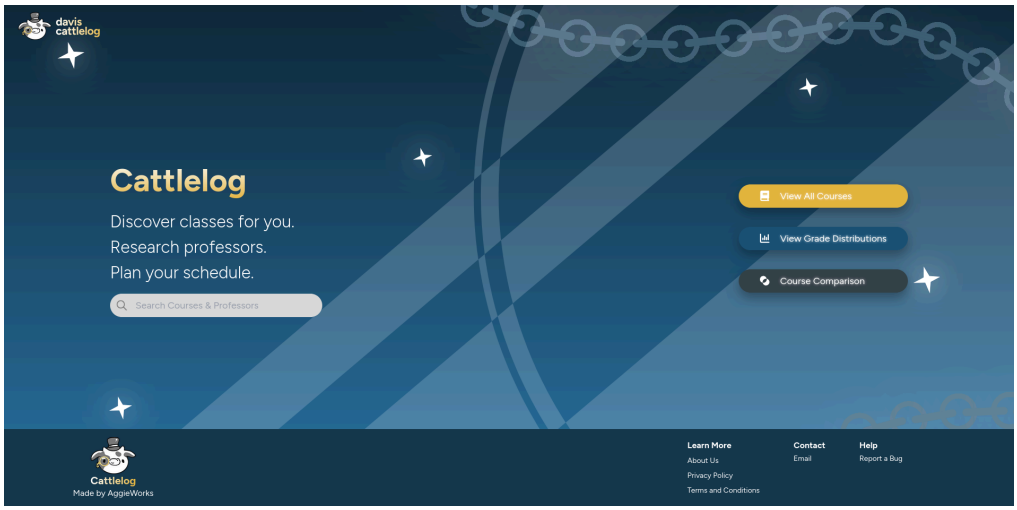

With collaboration with Jason, Eugene created the new landing page:

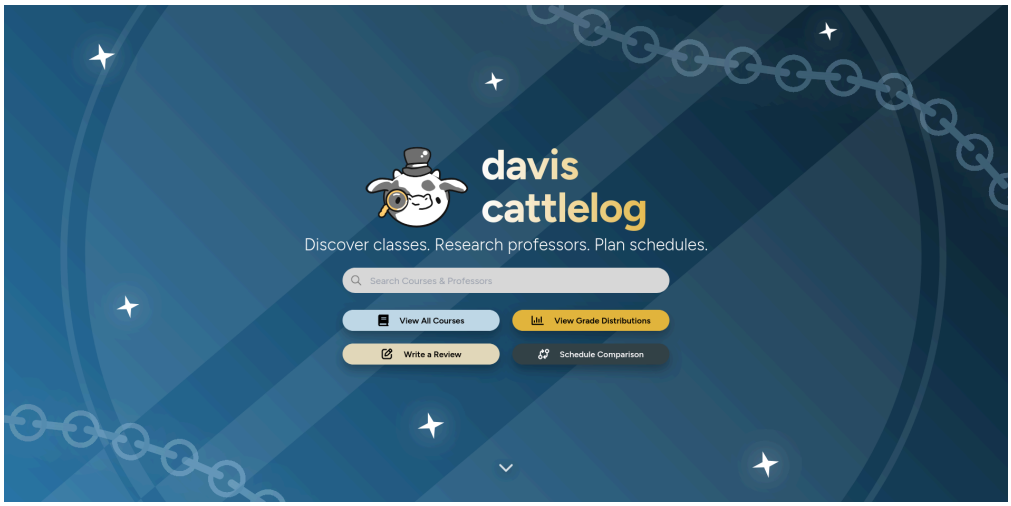

Many things improve the look and feel of this new landing page. The search bar aligns with the buttons. The logo and icon are more prominent on the page. Above that, it's also more functional with the added "About" section that has been included on the home page.

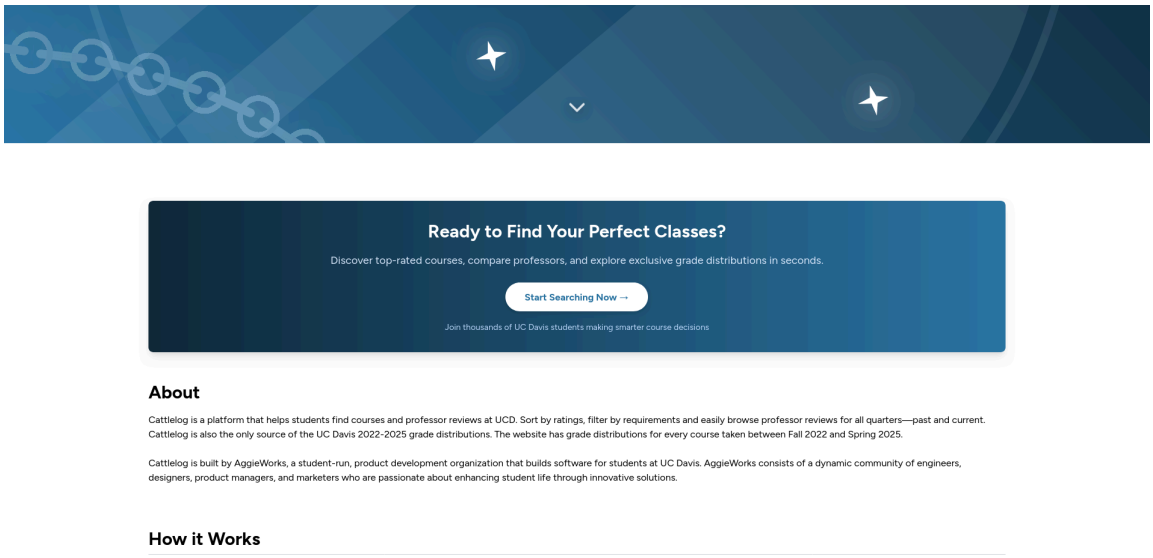

This "About" section includes details of the product that improve its search engine optimization (SEO).

2. The Grade Distribution Page

Previously, the only way to view a grade distribution was to search for the class on the home page, and then click the grade distribution button. We realized that some may not be aware that grade distributions are available because of all of the other information presented when searching a class.

To address this issue, we made a grade distribution page with a button on the landing page that directly navigates to it. You're able to search for any class and get an entire page dedicated to the chart of grade data. Side note: the normal grade distribution page has also been updated.

This new page includes the ability to view a grade "trend" across time. You can find which quarter has students ending up with the highest score. During this analysis, we did find that summer classes generally have higher GPAs.

Kudos to Eugene for creating an excellent feature here.

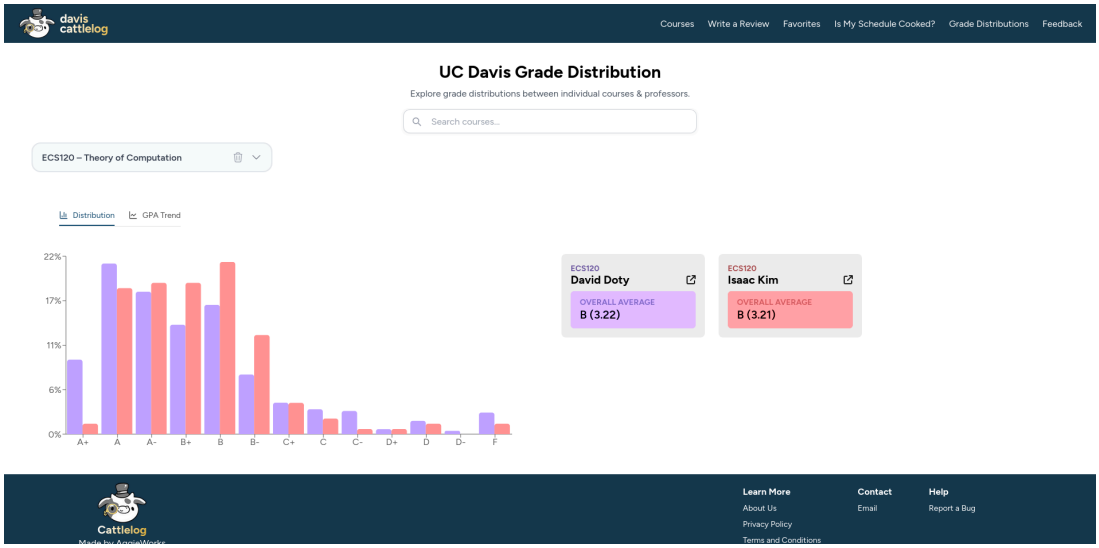

3. “Find Tutor” button

When ideating a way to improve impact for students, we realized that when students look at the grade data for a class, they might also be interested in preparing for their success in that class. We added a button to help students find tutors. This is still in an early state and is being piloted with ECS classes.

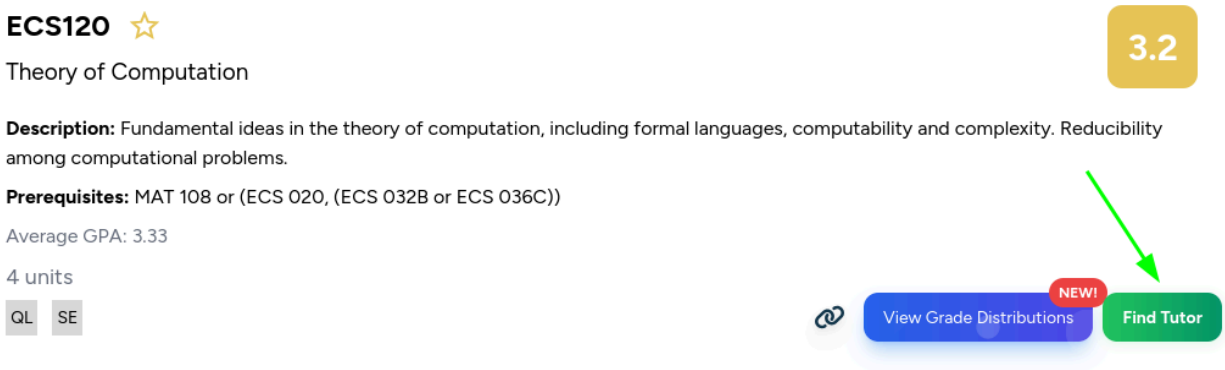

When you click the “Find Tutor” button, you are given a pop-up with available tutoring options. CS has a great organization called CS Tutoring who offer free tutoring for students.

We will continue to pilot this feature during fall and possibly expand it to classes from other departments.

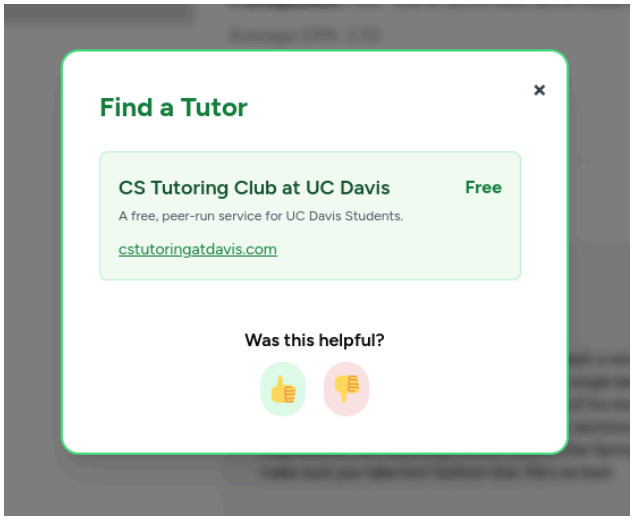

Here is what the pop-up includes. We also took the opportunity to add a thumbs up and thumbs down button to allow for feedback.

3. Prof Search - Beta

We've spoken to users, one of which said "The best feature you don't have is a professor search". Although we have a search on our home page where you can jump to professors, this user is right. We don't have a search similar to the course search for professors. Often, students are just looking for good professors and don't care exactly what classes they take. Maybe they want a CS upper-div credit with a professor who has a history of being very interesting and fair with grading.

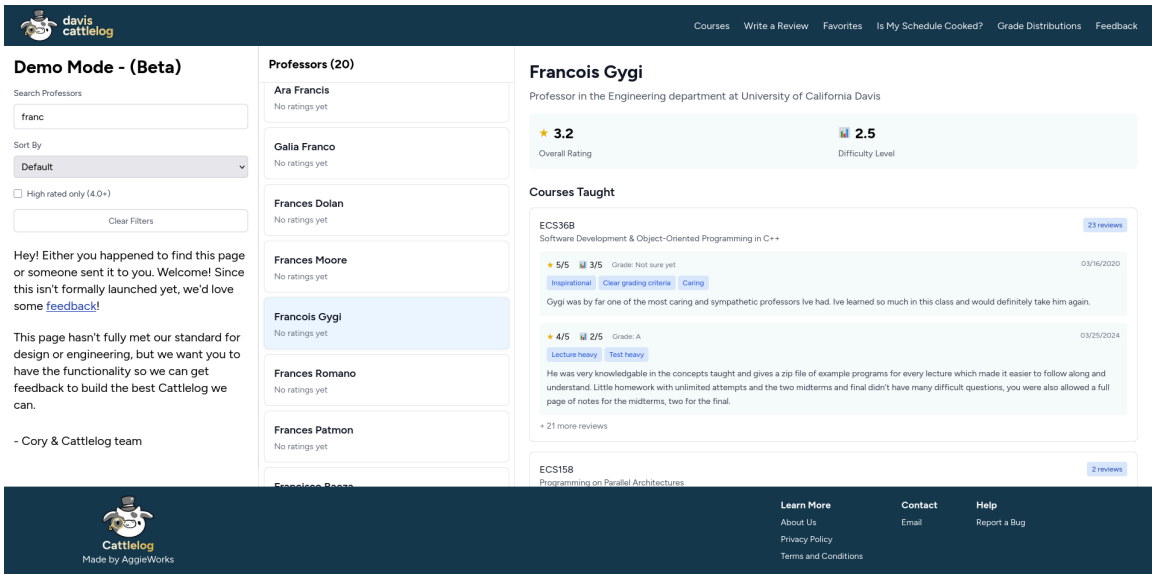

We created this professor search as a demo feature, hence “Beta”. The webpage reads “This page hasn't fully met our standard for design or engineering, but we want you to have the functionality so we can get feedback to build the best Cattlelog we can.” and is intended to get feedback. We don’t link to this page directly, so only users who have been sent a link or guess the URL can find it. We don’t intend to hide it from users, we just don’t want to funnel a bunch of traffic to it just yet, while we improve it to a standard that we adhere to.

4. Neptune - gzip data cache

Before the summer, we were using an r2 bucket to store our main json data. We found this to be **6.37** faster than serving the JSON directly from our backend.

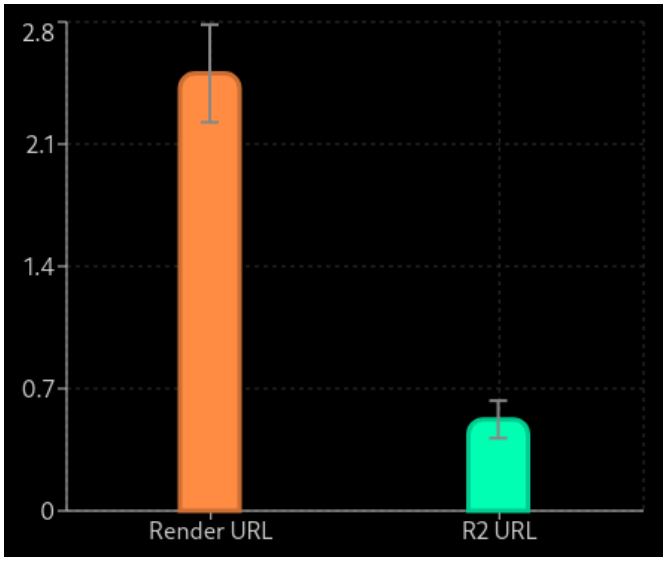

This change to r2 happened in Spring, but since then, the data has gotten bigger. Bigger data has resulted in slower time. This new data made the API speed half, from ~2.6 seconds to around 6.3 seconds over the summer. A ~2.6 second load time is okay, but after it hit 6 seconds, we needed to fix something.

We started to work on a system that would cache this data in a zipped format to improve load times. A technical detailing of this can be found on my blog in an article called Optimizing Course Search by over 8x.

During the week-long construction of this new backend cache, everything broke! Our data increased in size once more, and it went over the threshold that r2 will cache in their CDN. This made the time go from 6.3 seconds to over 30 seconds!! That was

certainly not possible to have as part of our product. We had to quickly finish our new cache called Neptune and deploy it. We had to scramble to get a new server, so it was hosted temporarily for a month on the servers for algoboard.org.

After optimizing this, we ran benchmarks to compare the original 6.3 seconds to the new cache we built in Rust.

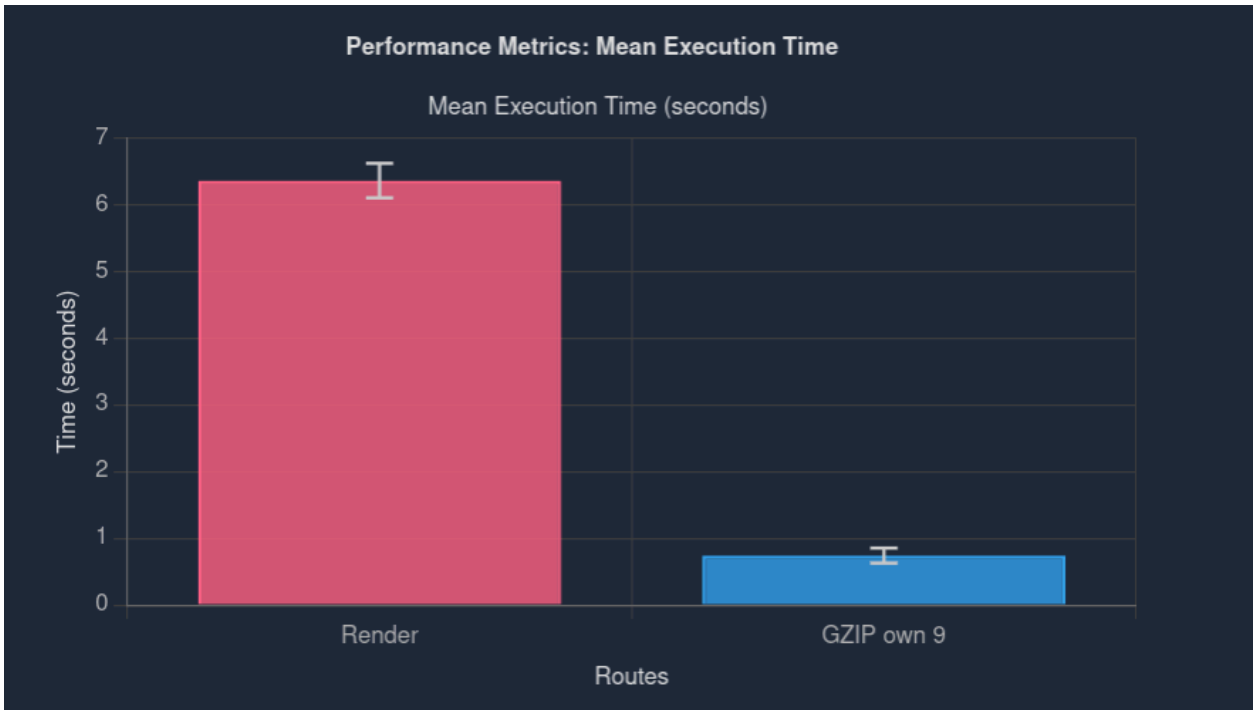

This new cache (as spoiled by the title of the article) increased the speed of the data by 8x. Along with this speed change, we also found a method to pre-load the data for Cattlelog, so when you navigate to the search page (home), you find that the data has already loaded in about half a second. It's unlikely you will ever see the loading spinner again thanks to these two changes.

Having this data cached has taken significant strain off our backend, which has handled over 1.3 million fetches so far. Users typically visit many pages and conduct many different searches. Each needs its own request and our backend was simply getting too many requests per minute, causing it to run out of memory. The addition of the cache has fixed that.

It's also worth noting that we added redis to the backend as well as some other optimizations and speed testing.

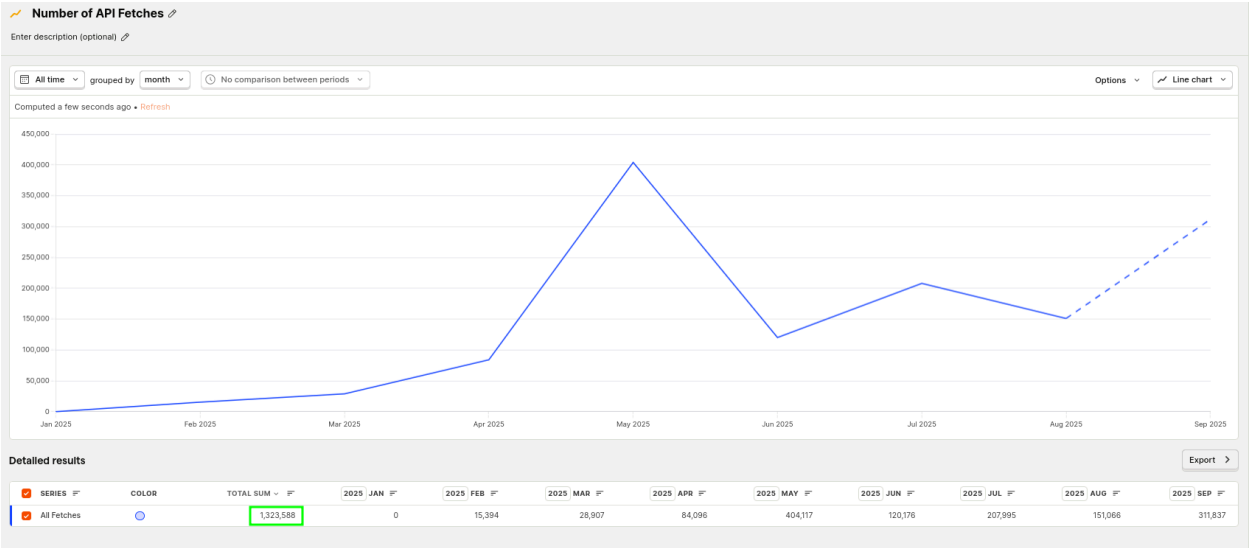

5. Two browser extensions

We created two browser extensions this summer. The first was designed after we had found a new library for making extensions, so we built one to test with. We mostly set this aside to launch in fall or sometime later.

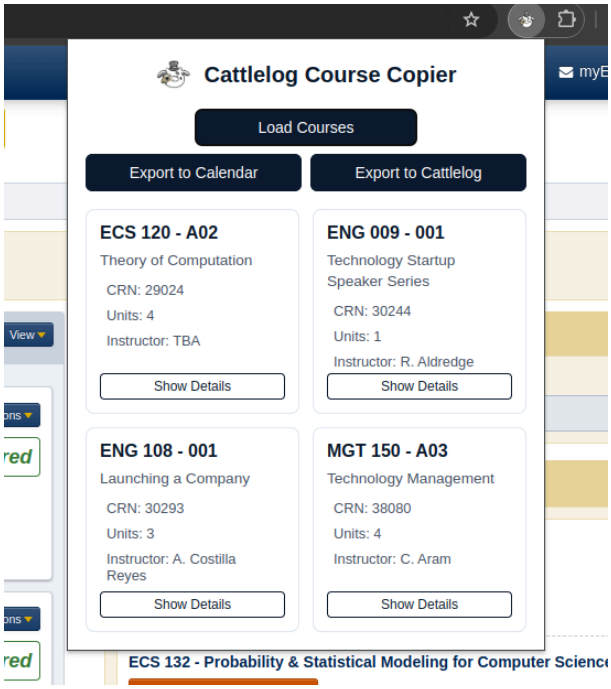

A few weeks after this, Sohan also created an extension and after he was finished making it, he shared it with us. We have now integrated this new extension into our codebase and will work towards an eventual release of this extension.

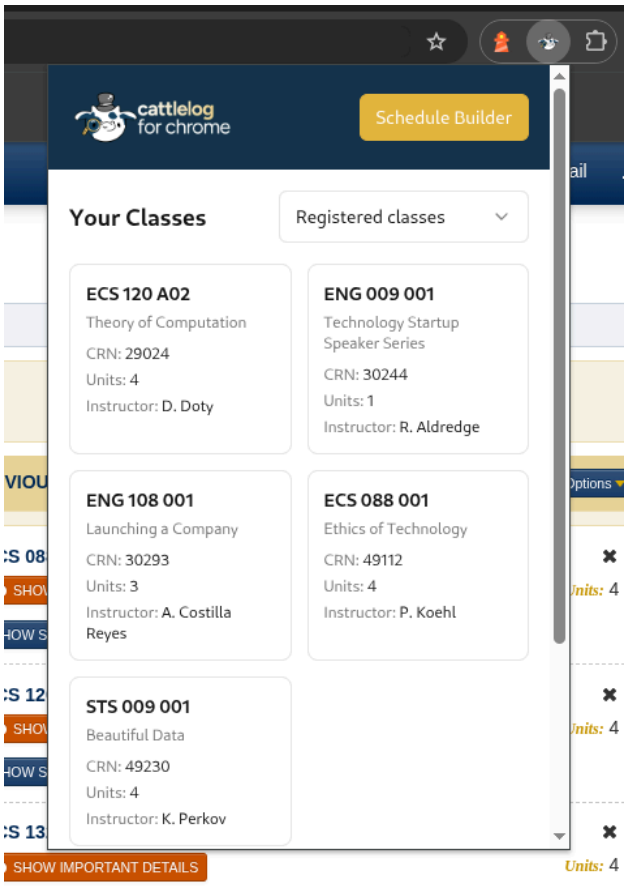

6. New Blog Page

After viewing the analytics from the original blog post, we noticed that the blog has gained us over 4,100 clicks to our product. We realized how great this is as a way to direct searches and ChatGPT responses to our website.

Content like our blog has helped us rank highly on Google and other search engines. We now can show up in the AI generated response when you ask grade questions to Google.

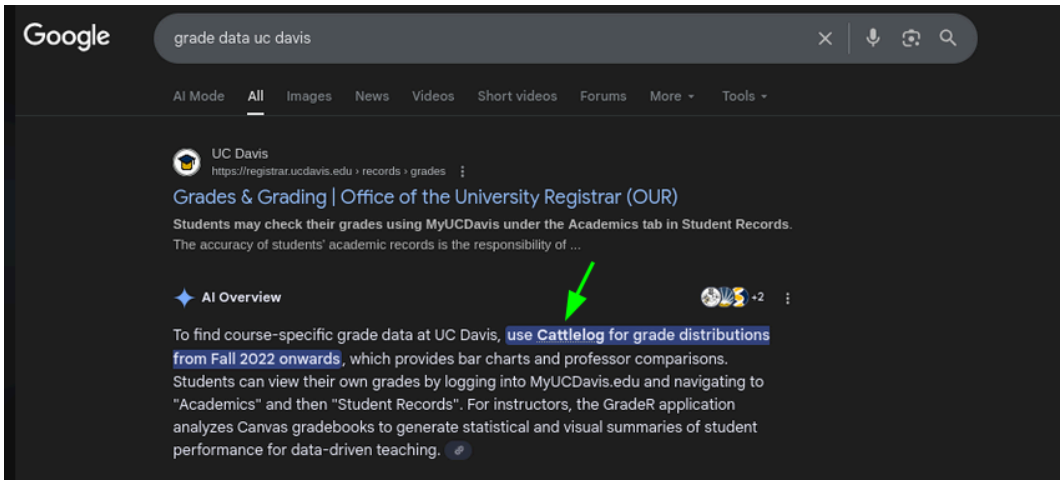

This improvement to SEO took many pull requests and changes to different pages. Ultimately it's the blog page that shows up on Google.

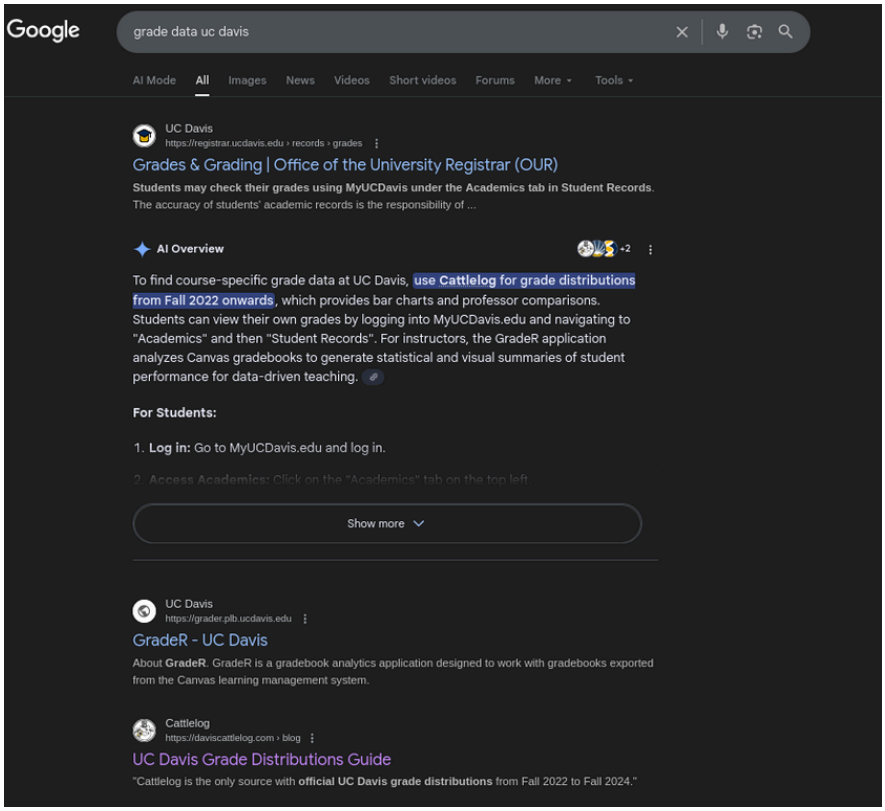

7. Prof Timeline Feature

We often see people asking on online discussions about “Who taught X in Fall quarter”, or “Who usually teaches Y”. It’s a very common question and we realized that Cattlelog is the only website on the internet that has this data for the last 5 years. We have it as a byproduct of having the grade data of all of these professors. It’s not set up to find this data without clicking through all of the grade data.

For this reason, we made the prof timeline feature. This feature isn’t currently available to everyone and is turned off with a feature flag. See more about [feature flags on our docs] or this other write-up about [A/B testing].

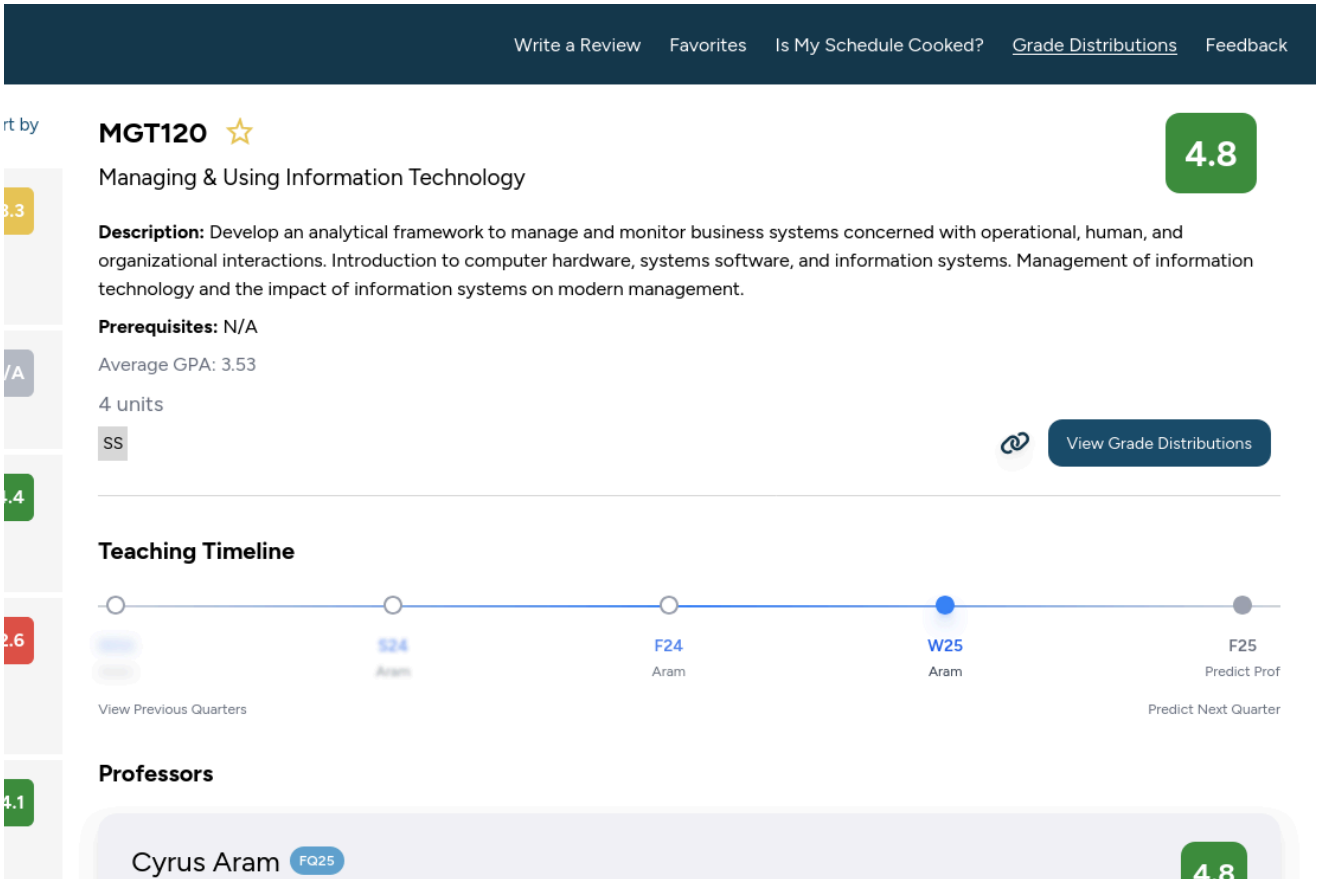

We plan to predict future professors as well. This predication has been heavily requested recently. Especially when CS professors are not listed until right before classes start.

8. Chat Mode

We thought of making a chat mode to return grade, course, and professors info to you via a chat. This feature has long been talked about, but it may be time for us to build it. We've made a lot of the other easier requested features and this would be great.

Here, you're given an option to change to "Chat" mode, where you can ask questions in plain english and get a response with retrieved information. Look out for this feature sometime in the future.

9. Misc Improvements

There are too many improvements to count, but over the summer we created **514 commits** in **84 pull requests**. Not that it's a good metric, but that ended up being **71,187 lines of code**.

This document ended up being less detailed than expected but also longer than expected.

A full list of PRs can be found below.

Pull Requests

calculating GPA sql scripts and endpoint naming convention #499
Chat mode - preview #502
Posthog for rec classes #501
Update on New data + some formatting #496
Update README.md #498
Adjusting Scroll #472
UI/UX updates #489
Add back full data - used in course route #494
Remove benchmark code for neptune #492
Add new data #491
Prof search posthog #488
Prof search #463
increased scaled width of buttons on landing page #480
Fix easy url #487
Fixed calendar export issue #486
Update docs for sub projects #485
Sohan's Chrome Ext Build #483
Updating small frontend fixes and load_to_db #478
Fix test to exclude zero #482
Add promotion for startup classes #475
More tests #481
Bug - Fixing rating issue #477
Frontend adjustments #468
Fix button text wrap in about section #467
new landing page, other random stuff #466
Remove timeline #465
Add more analytics #464
Fix color on timeline #461
Add prof timeline #457
Fix link for class #456
Easy classes blog #455
Small fixes #452
Search option #451
About on landing page #450
Frontend adjustments and cleanup #448
readme update and sql script moving #446
Update data #447

Trigger Build #445
Revert "Fix url for grades" #444
Fix url for grades #443
Add courses api #442
11.9x speed increase with Neptune API gzip cache layer #441
Format check for frontend #440
bunch of refactoring and renaming for clarity #438
Allowing grade dist professors to show up in professors #437
Add posthog capture for API fails #436
Reorganization and Refactoring Home #435
Fix prefetch to happen in background #434
Fix prefetch to happen in background #433
Move prefetch to work in background #432
Add prefetch for course data to reduce wait from 4 seconds to 0 #431
Fix URL #429
Frontend adjustments #428
Change email to link #427
Self host our fonts to save 2 entire round trips for the font data #424
Image perf - reduce image size by >80% #422
Add asterisk with span for email #423
Add capture to about page button #420
Frontend adjustments #418
Landing Grade Dist Page Trends Section #414
Redis cache implementation #413
Landing page update #412
Revert "initial setup" #411
landing page grade dist page #410
Frontend adjustments #405
moving files out of staging #406
File structure reorg #404
Add more gradient docs #402
Add basic gradient docs #401
search be speedy #400
Reduce docker image size by over 75% #399
api endpoint optimization, intro to database staging, adding columns, fixing avg gpa issue #395
course backend endpoint #391
updating pool #398

Change URL to one that works #394
Revert "Change url to faster URL" #393
Change url to faster URL #392
updating pool_size #390
updating pool size #388
Slug integration ; Refactored Course Endpoint ; Support for ALL courses (ie courses w/ no grade dist, no reviews, or anything in between) #387
Backend logic refactor ; Database work #386
Favorites Page rewrite and reorganization, minor UI fixes #383
Find tutor #380
Speed test for query async #379
Add optimization check #363
New Grade Dist Page #371
File Reorganization #370
Added Documentation for Loading and Scraping #369
Add extension docs #368
Add cattlelog course copy chrome extension #367
Create speed test and fix issue with embed #366
Add tt commit #364
comment out content locker #362
Make better default #361
fixing rando bugs #360
Make capture for posthog unique #359
Remove Write Review button and change color of View Grade Dist #358