

Using AI to Analyze 600+ Tariff Comment Letters

May 19, 2026 – Dorsey Law & Policy Notebook

by Troy M. Keller



Like many, I've found myself experimenting with AI to see if it would enable me to take on projects that are otherwise out of reach. I've had success in some narrow cases, but other times I've ended up in a rabbit hole to nowhere. My most recent project has landed

somewhere in the middle.

When USTR opened two Section 301 tariff comment dockets in March, I had aspirations of creating an AI-powered process that would allow me to effectively harvest insights and trends from the large amount of commentary that was sure to follow. I imagined how useful it would be to create a database using automated tools and then slice and dice the information to provide new perspectives. Now, after a few weeks of iterating, I've got a tool that provides much of that functionality though with some drawbacks. You can see it for yourself [here](#).

Background on 301 Tariff Dockets

As we discussed in [an earlier blog post](#), USTR initiated a suite of 301 investigations on March 11, 2026 and then again on March 12 into the trade practices of sixty foreign economies. The investigations were a necessary procedural step for the USTR prior to determining whether new schedules of 301 tariffs can be applied in those jurisdictions. While the prospect of renewed or even potentially increased tariffs was an unwelcome development for many, the ability to comment in the investigations was an opportunity many organizations seized. In the end, hundreds of comment letters were filed by the deadline of mid-April. After hearings in early May, USTR extended the comment period to allow for post-hearing submissions; that window closes later this week.

My Process

For this exercise, I downloaded 633 comment letter pdfs that had been filed at the time the dockets originally closed in mid-April. I have not included letters filed since they re-opened, nor have I included data from the transcripts of hearings earlier this month.

Related People

- Troy M. Keller

Related Capabilities

- Mergers & Acquisitions
- Securities & Financial Services Litigation & Enforcement
- Government Enforcement & Corporate Investigations

Those may come in a phase 2. To keep the dataset manageable and to focus on the more substantive comments, I also did not include comments that were not accompanied with a filed letter. This information was then aggregated in a datasheet and displayed on the [dynamic dashboard](#) linked to above. The dashboard includes summary information that can be filtered by industry, type of commentator, type of relief requested, targeted country and other factors. Each comment summary links to a stored copy of the commentator's pdf letter.

AI Tools I Used

The first step was to download 633 letters into a folder. I relied on an openclaw bot to manually download them one at a time (as they weren't available for bulk download on the Federal Register). As usual, getting the bot permissioned was most of the battle. A couple of times, I thought the process was moving along only to discover later that it had basically hallucinated most of the letters (which was impressive in itself, as these were real-looking pdfs from known companies--they just hadn't actually filed letters). I also relied on the openclaw bot to make updates to a database in Google Sheets where I stored extracted information from the letters. The bot's ability to write to the database so that I didn't have to continually cut and paste corrections from various chats was where it shone.

While the bot was useful for moving files around and editing the database, I used Gemini and Claude for the thinking and analysis. They were more accurate and less likely to generalize or hallucinate. Still, because of the number of letters, these tools had to be prompted in small batches to read and pull summary data from them. More than 20 letters at a time would result in them hitting limits and either stopping mid-analysis or returning extrapolations that were useless. Running the prompts over and over was one thing, but re-explaining the whole project every time they lost context due to compaction was painful. I also used Claude to develop the dynamic dashboard to display the results and allow users to filter and sort them. The dashboard was published using Google Sites.

Limitations and Caveats

Every step involved accuracy challenges and more effort than I would want to repeat to catch them. At this moment, bugs still remain, mostly in the categorizations. As a result, the data is good enough for spotting trends, but probably not good enough for citation. Other limitations are by design. The categories we ended up with for industry sectors we more or less evolved into while trying to organize the letters into manageable groupings. In hindsight, a better route probably would have been to apply NAICS codes from the outset--though that might have led to dozens of single-entry categories, which I was hoping to avoid.

Some Findings

The main value from this exercise is the ability to go to the dashboard and filter results in a variety of ways to see what types of groups are commenting and what they are asking for. To give you a flavor, here are a few observations for each filter.

Comment Letter Asks. No surprise, a significant number of the letters (261) sought tariff exclusions as their primary request to enable them as importers to better compete. However, there were more letters than I would have guessed going the other direction—114 in total. Generally, these were organizations pushing for tariffs that would be beneficial to their U.S. operations. A number of commentators used the opportunity to highlight non-tariff trade barriers or other trade affecting dynamics. We categorized most of these asks as "nuanced" given they covered concerns that did not fit cleanly into the for-or-against tariff framing.

Who Filed. Individual corporations filed the bulk (305) of the submissions, followed by trade associations (200). NGOs and think tanks accounted for 60. Submissions also came from foreign governmental entities (11) and labor unions (7), with 50 falling into an "other" bucket.

Targeted Countries. In terms of countries that were the subject of the comment letters, China came up the most (in 235 letters). Outside of China, the EU and India were the only economies to feature in over 100 letters. Mexico (80), Vietnam (73), Japan (48) and Taiwan (45) round out the rest of most-mentioned jurisdictions.

Most Represented Industries. As expected, heavy manufacturing interests (industrial equipment, automotive, chemicals) represent the bulk of the comment letters. However, other sectors were also quite active. Our transportation/logistics category included 35 letters, and semiconductors and electronics category involved 24 letters. We counted 23 medical and pharmaceuticals related letters and 21 for textiles & apparel. Interestingly, 13 letters landed in our arts & antiques category, 15 in dairy & cheese and 12 in bicycles, motorcycles & e-bikes. These are high numbers for niche areas and demonstrate more sophisticated supply chain dynamics than might be expected.

Parting Thoughts

Overall, I'm pleased with the results and the capability of the dynamic dashboard. The lack of full confidence in the accuracy is not a small drawback. However, it's not so large a database that flaws are hidden, which means there is still good utility for those who are a little forgiving when they notice some mis-categorizations.