

FDA Automating Active Surveillance of Online Health Fraud using Artificial Intelligence

Marla Hendriksson¹, Sunita Joshi-Taylor¹, Ben Ballintyn², Tyree Giles², Alex Hanson², Anna Hoffman², Caroline Trier²

¹US Food and Drug Administration, ORAHQ, OC ²Booz Allen Hamilton

Project Overview

Problem Statement

The Food and Drug Administration (FDA) faces challenges in the rapid and dramatic increase in FDA-regulated products sold online. For example, revenue from e-commerce grew by over 300% from \$340 billion in 2019 to \$1.03 trillion in 2022. The Online Marketplace Research & Surveillance Tool (OMRST) automates the monitoring of online marketplaces. Designed with extensibility in mind, OMRST's first use case is assisting ORA's Health Fraud Branch (HFB) protect public health by facilitating the identification of fraudulent products sold through e-commerce. In 2020 alone, HFB received approximately 5,000 consumer complaints and investigated over 1,270 products – resulting in a ~250% increase in federal actions from the previous year. OMRST identifies and reports an ordered list of the most impactful potentially violative products.

Goals:

1. Active surveillance of the potentially violative products in the online health supplement market
2. Ordered results to increase efficiency of FDA review and action

Cloud Architecture

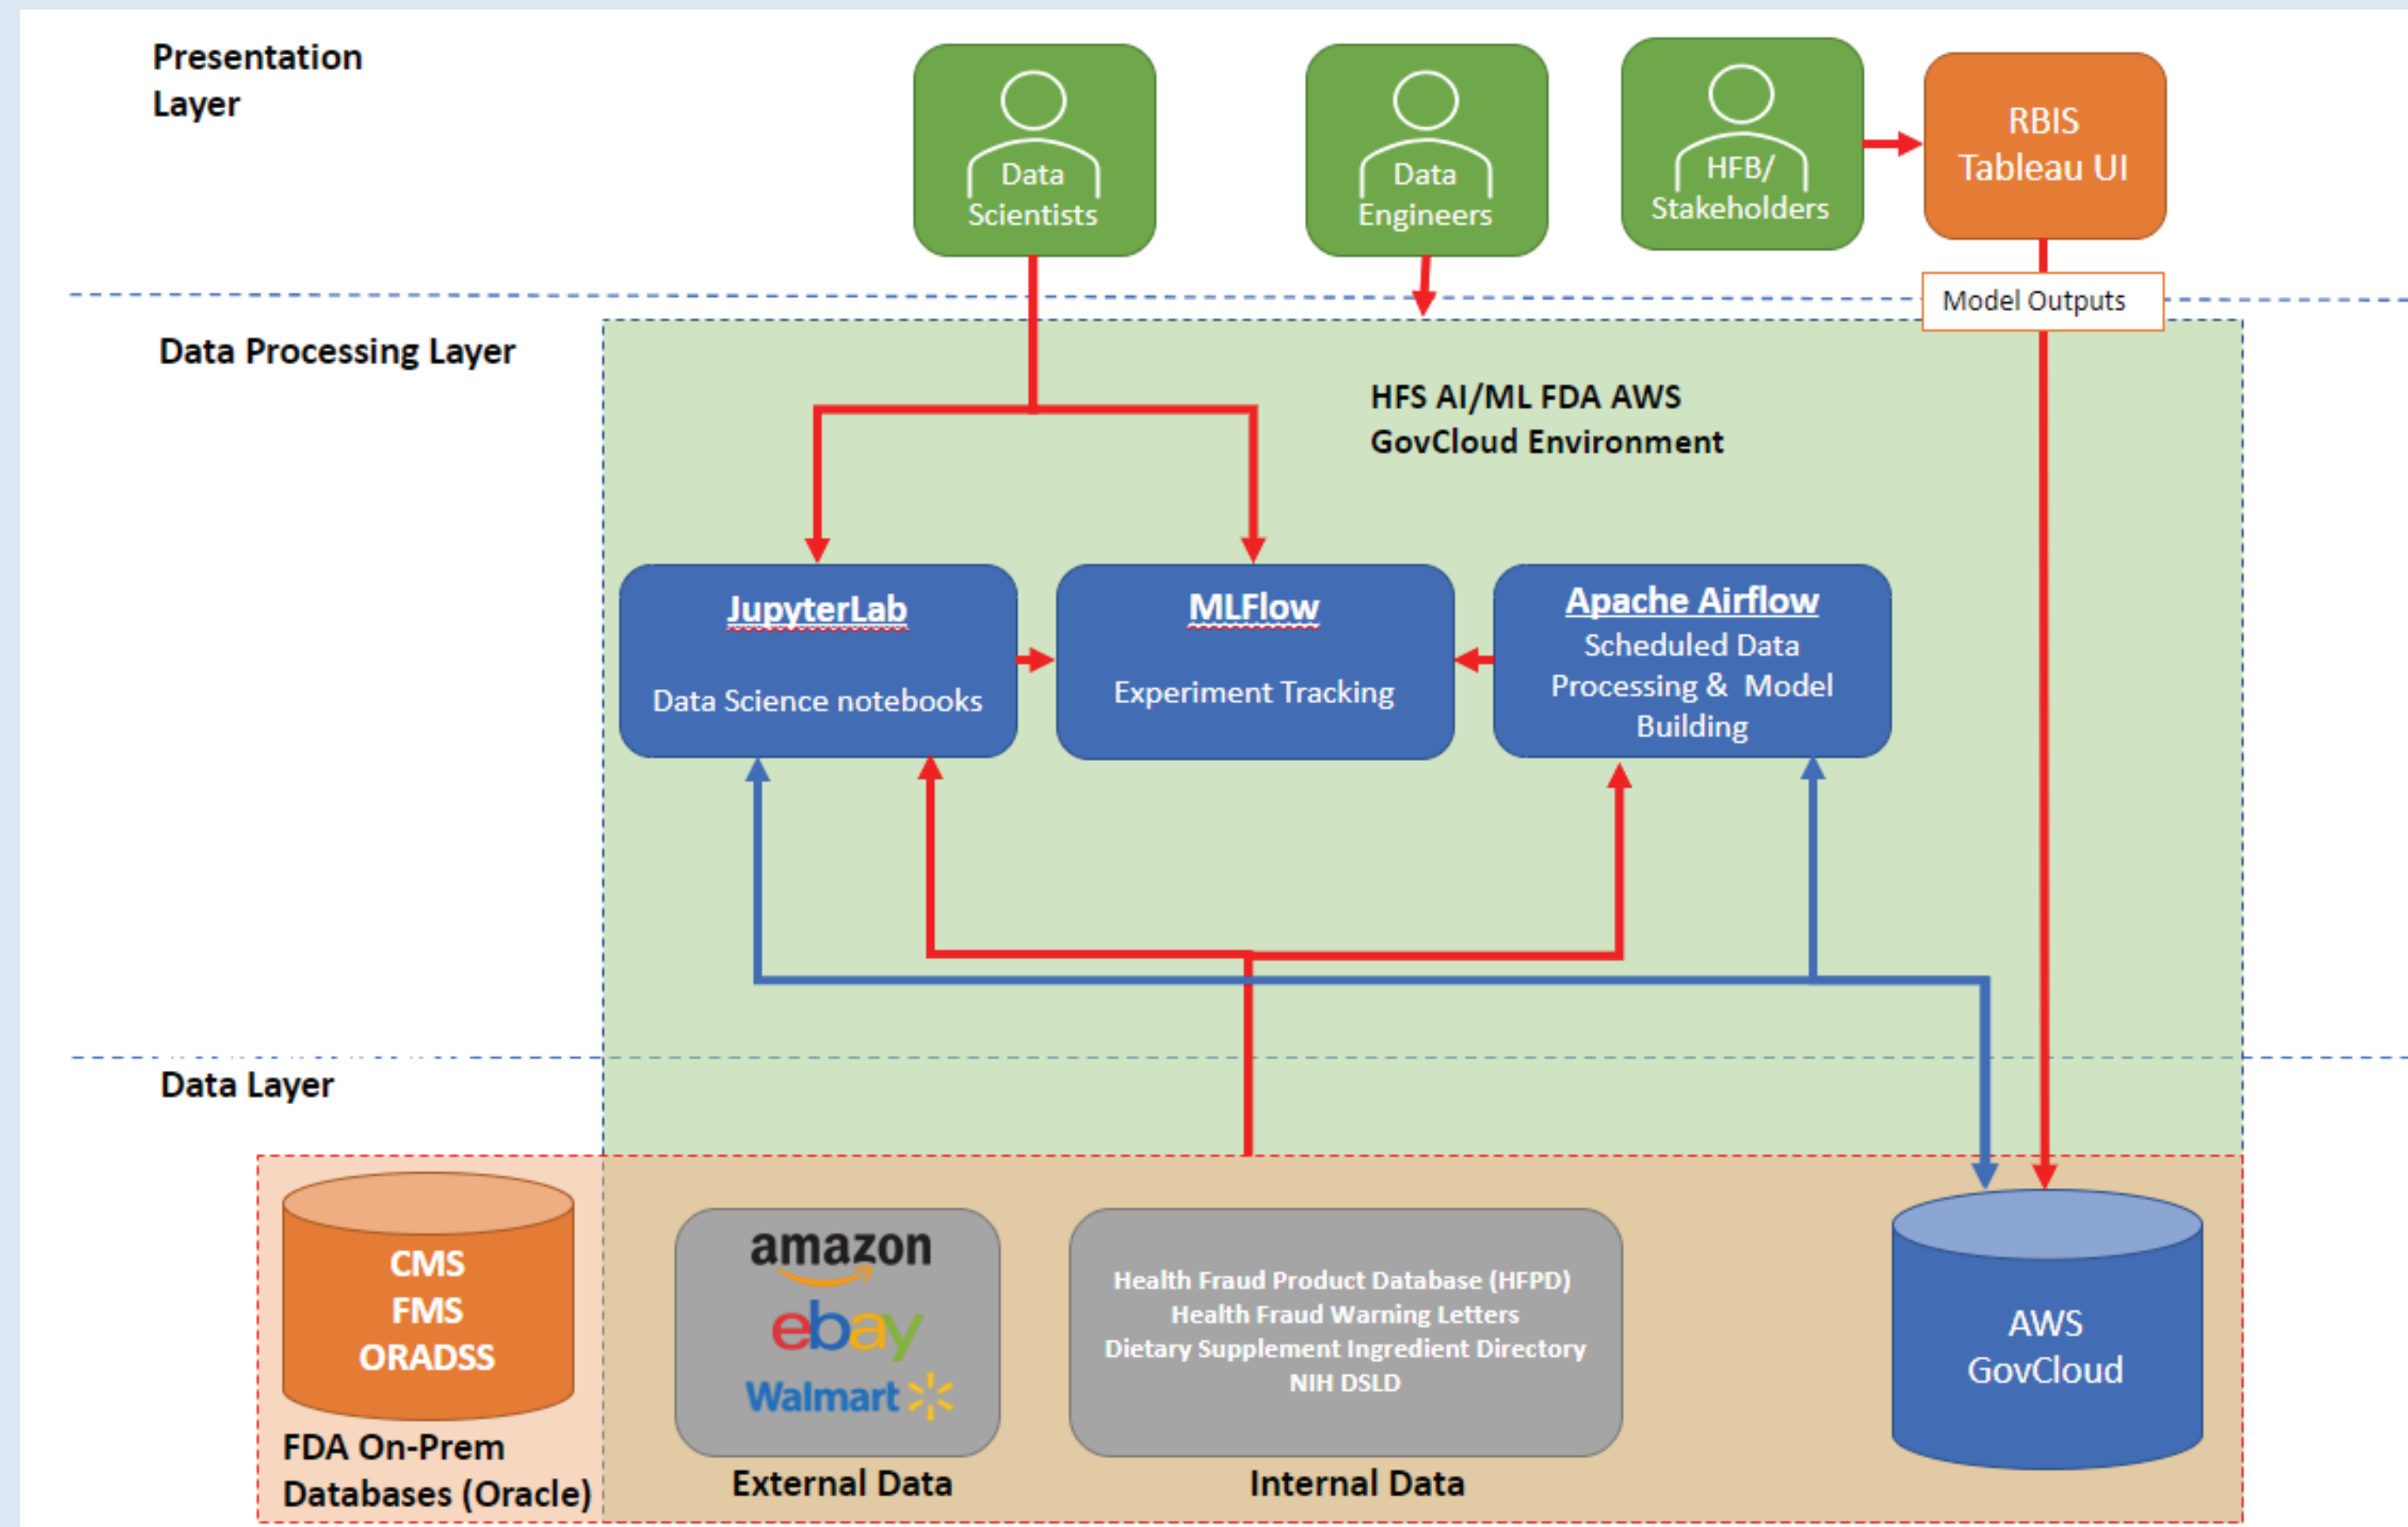


Figure 1. Cloud Architecture of OMRST

As shown in Figure 1, OMRST consists of a data layer with live connections to FDA and external marketplace data. Product information from Amazon, eBay, and Walmart APIs is collected daily and processed using modern NLP techniques to extract signals that indicate potential fraud and its severity in the data processing layer. These signals currently include illegal ingredients, and problematic claims. We aim to incorporate adverse events, product popularity, and firm history in future iterations. The resulting model output of potentially violative listings is ordered and visualized as leads in the presentation layer's Tableau dashboards for review by HFB.

Marketplace Surveillance Coverage



Figure 2. Unique products returned from our marketplace APIs (orange) and then filtered to only health supplements (blue). The gray line depicts the estimated total number of health supplements available per marketplace based on the search "health supplements".

Of the 1 trillion dollars spent in the US on e-commerce, 52% of the market share is covered by Amazon, Walmart and eBay. Using these marketplaces' APIs and hundreds of supplement searches, we queried hundreds of thousands of unique products for analysis with a limited number of API calls. Figure 2 shows the number of unique products and health supplements returned through our searches, as well as the estimated available health supplements. On average, ~75% of the available supplements were returned, which corresponds to 39% of total estimated e-commerce listed health supplements.

Online Marketplace Surveillance Tool Models

Claims Model: Automatic identification of problematic claims made by products

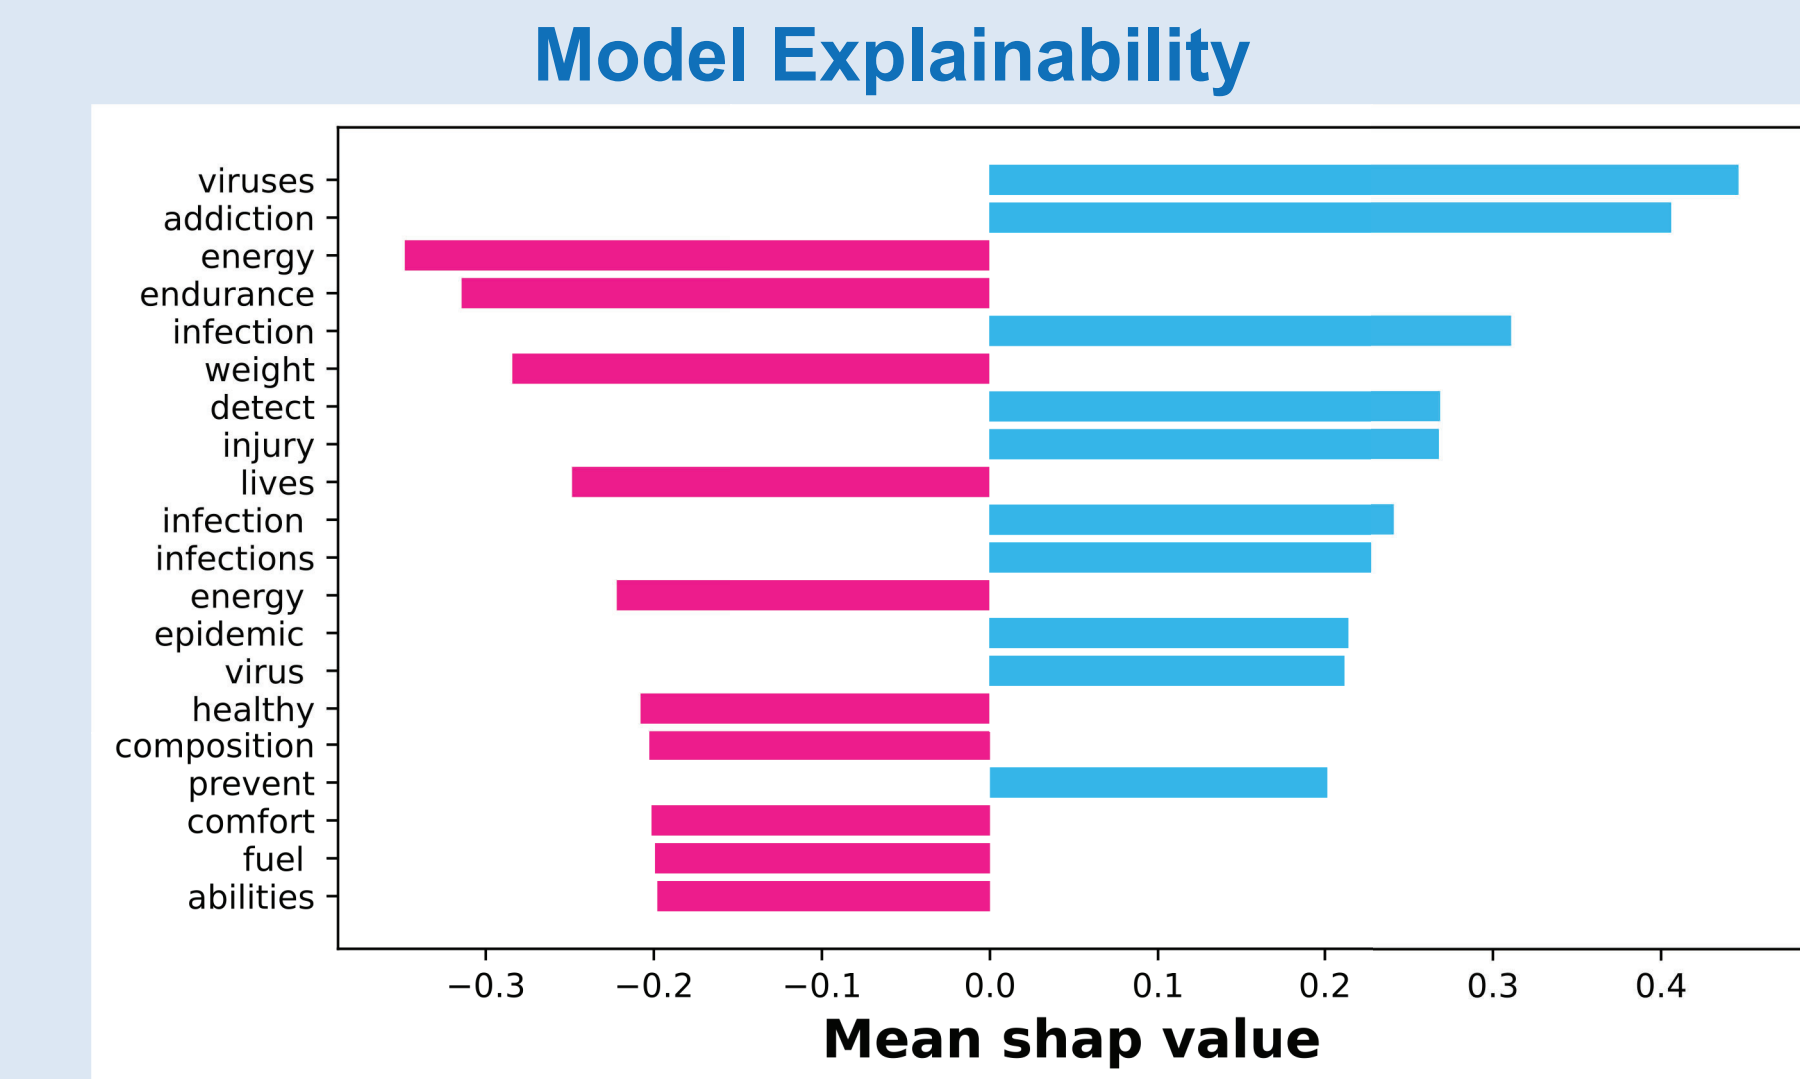
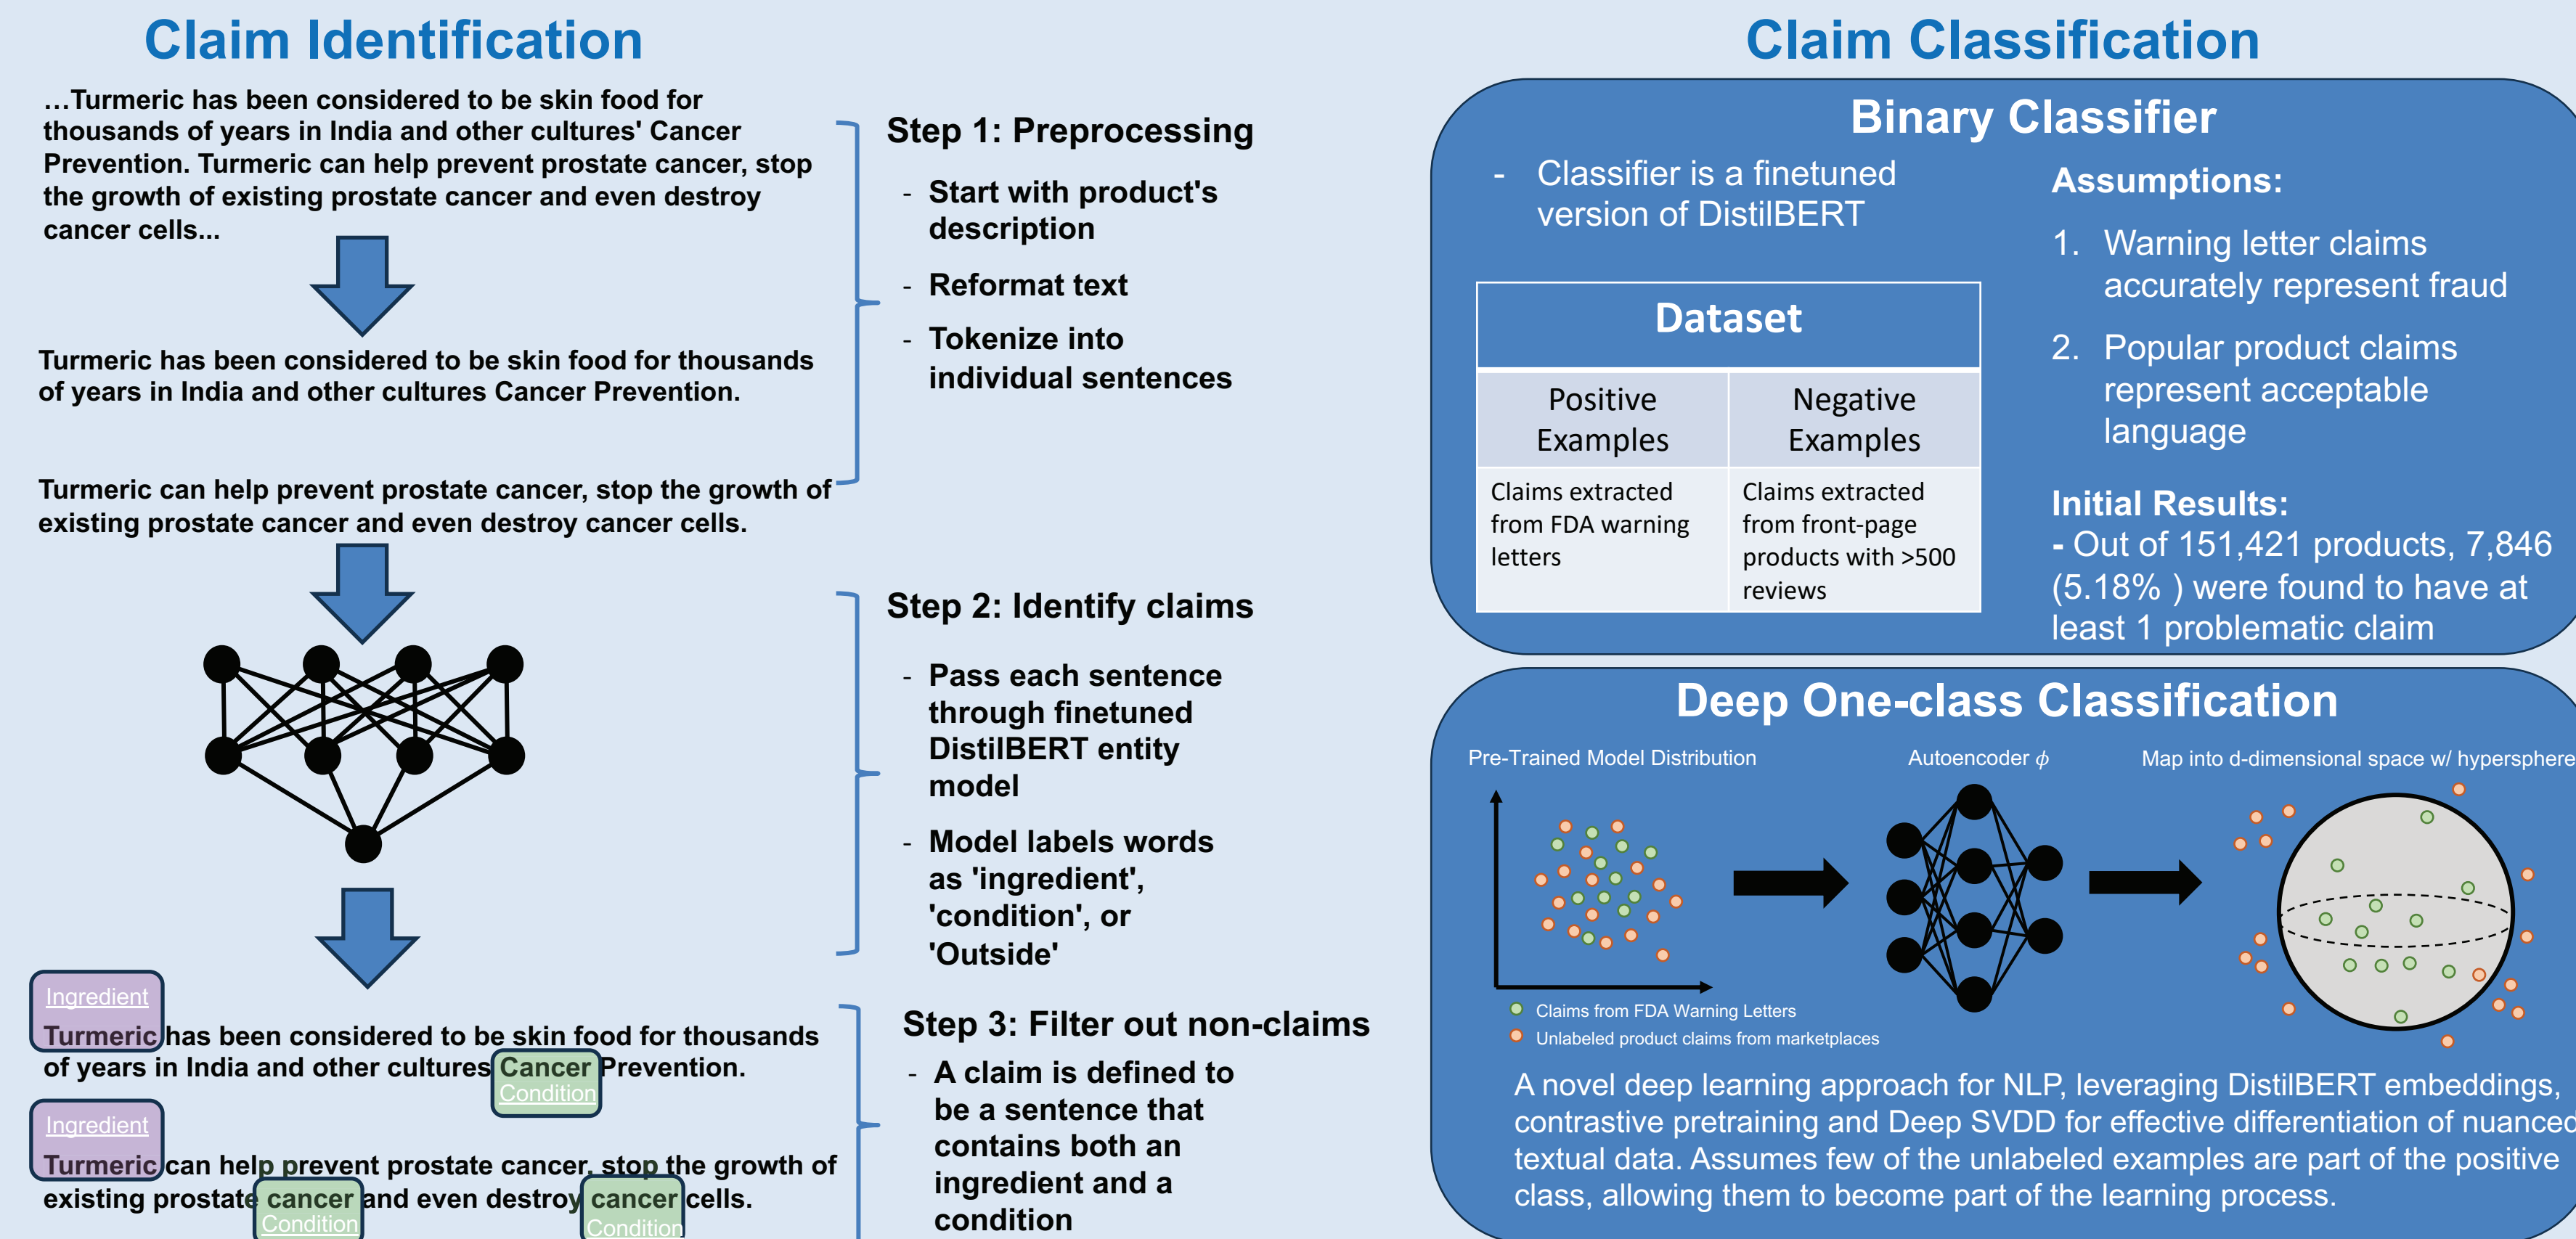


Figure 3. Top 20 most impactful tokens as defined by mean shap value across 5000 problematic claims run through the binary classifier. Blue (+) / Red (-) values increase/decrease output score (where 1 = problematic).

Ingredient Model

One of our objectives is to identify violative products based on the presence of an ingredient from multiple sets. Our solution uses tries, which simultaneously compare an input against all set words at once. This deterministic approach provides a fast and efficient solution compared to other algorithms.

Using a trie structure minimizes the computation cost of scaling up the number of sets and number of ingredients in each set.

In benchmarks of 130,000 products with over a million ingredients, the trie finished the task ~100x more quickly than traditional match methods, with no difference in the output. Our model identified 87 products on the marketplaces with banned ingredients.

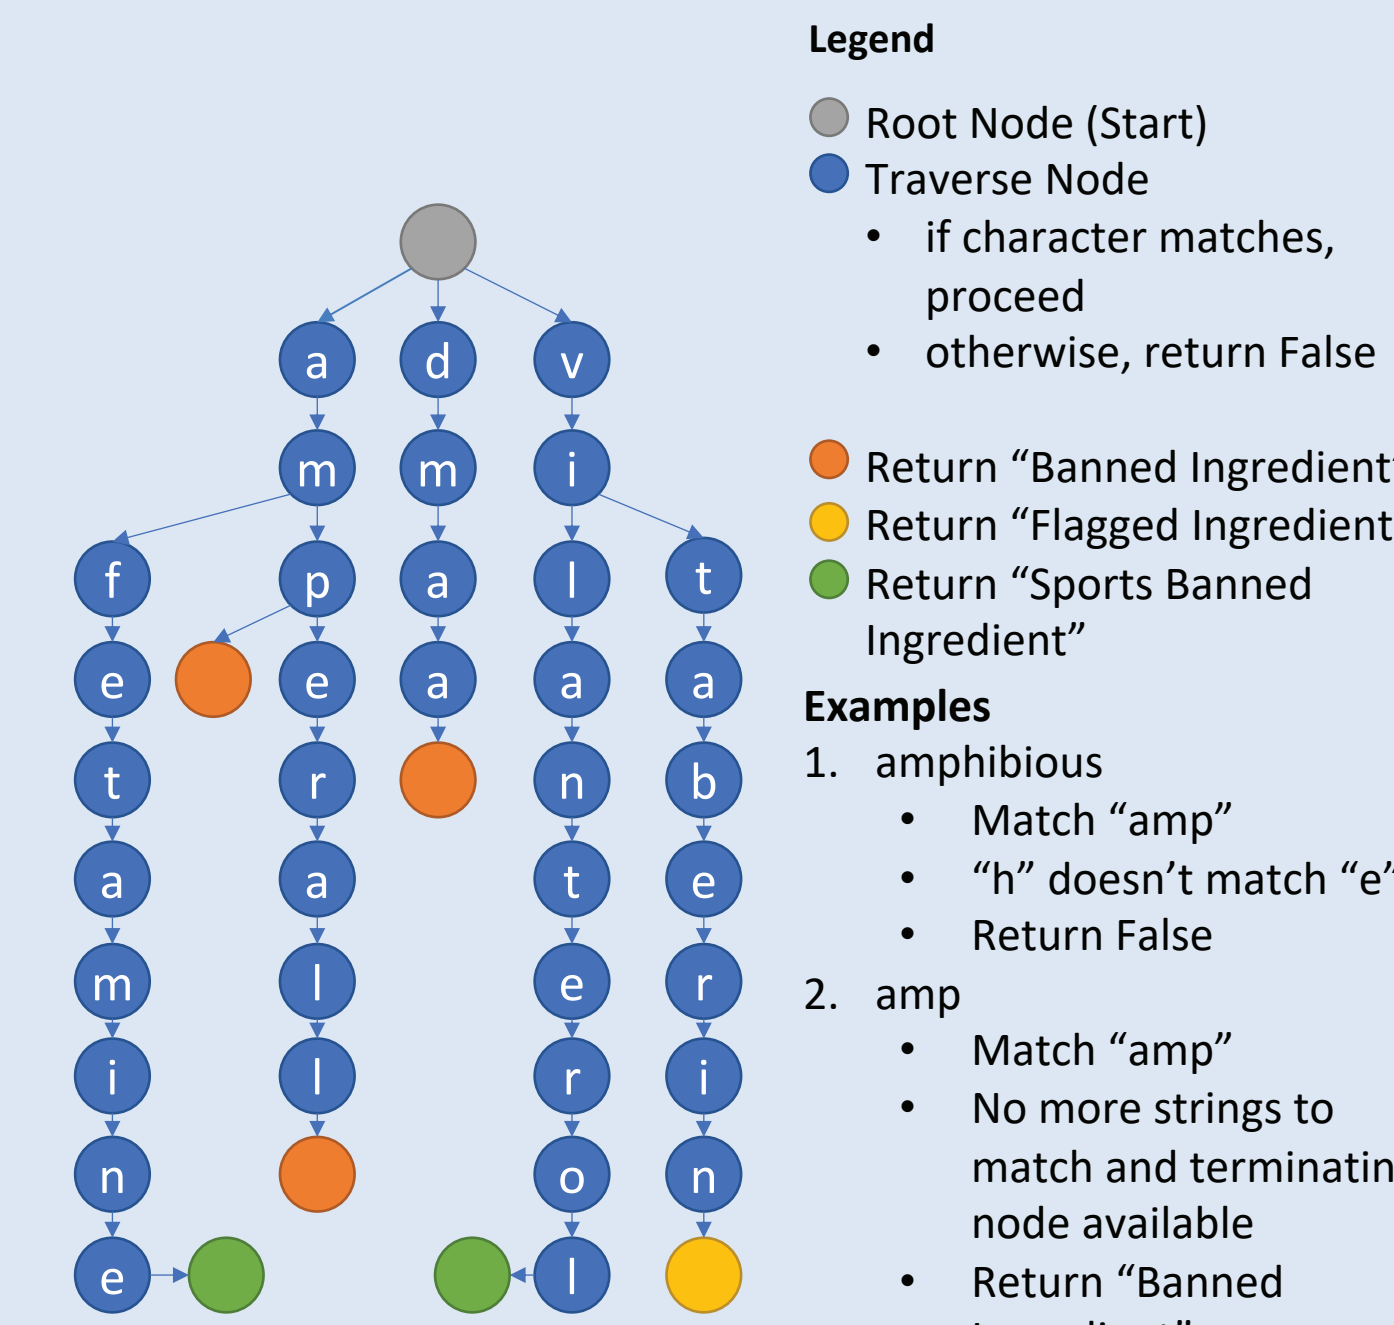


Figure 4. A visualization of a trie with the ingredients amfetamine, amp, amperall, dmaa, vilanterol, and vitaberin

Future Directions

Modeling:

1. Unsupervised pretraining
 - Obtain domain-specific embeddings to use as base model for NER and claim classification
2. Refinement of claims training dataset through human labeling
3. Detection of user-reported adverse events in product reviews
4. Optical character recognition (OCR) to transcribe label information from images

Result ordering:

1. Prioritization for claims mentioning serious diseases/conditions
2. Inclusion of firm compliance history in prioritization

System:

1. Customizable search terms for product queries
2. Automatic verification of product removal from marketplaces by firms

Ordering Results

To assist review, we order products for HFB analysis by the following algorithm: $\max(0, 1) \text{ banned ingredients}, (0-1) \max(\text{claims confidence score})$ and display to a Tableau dashboard.

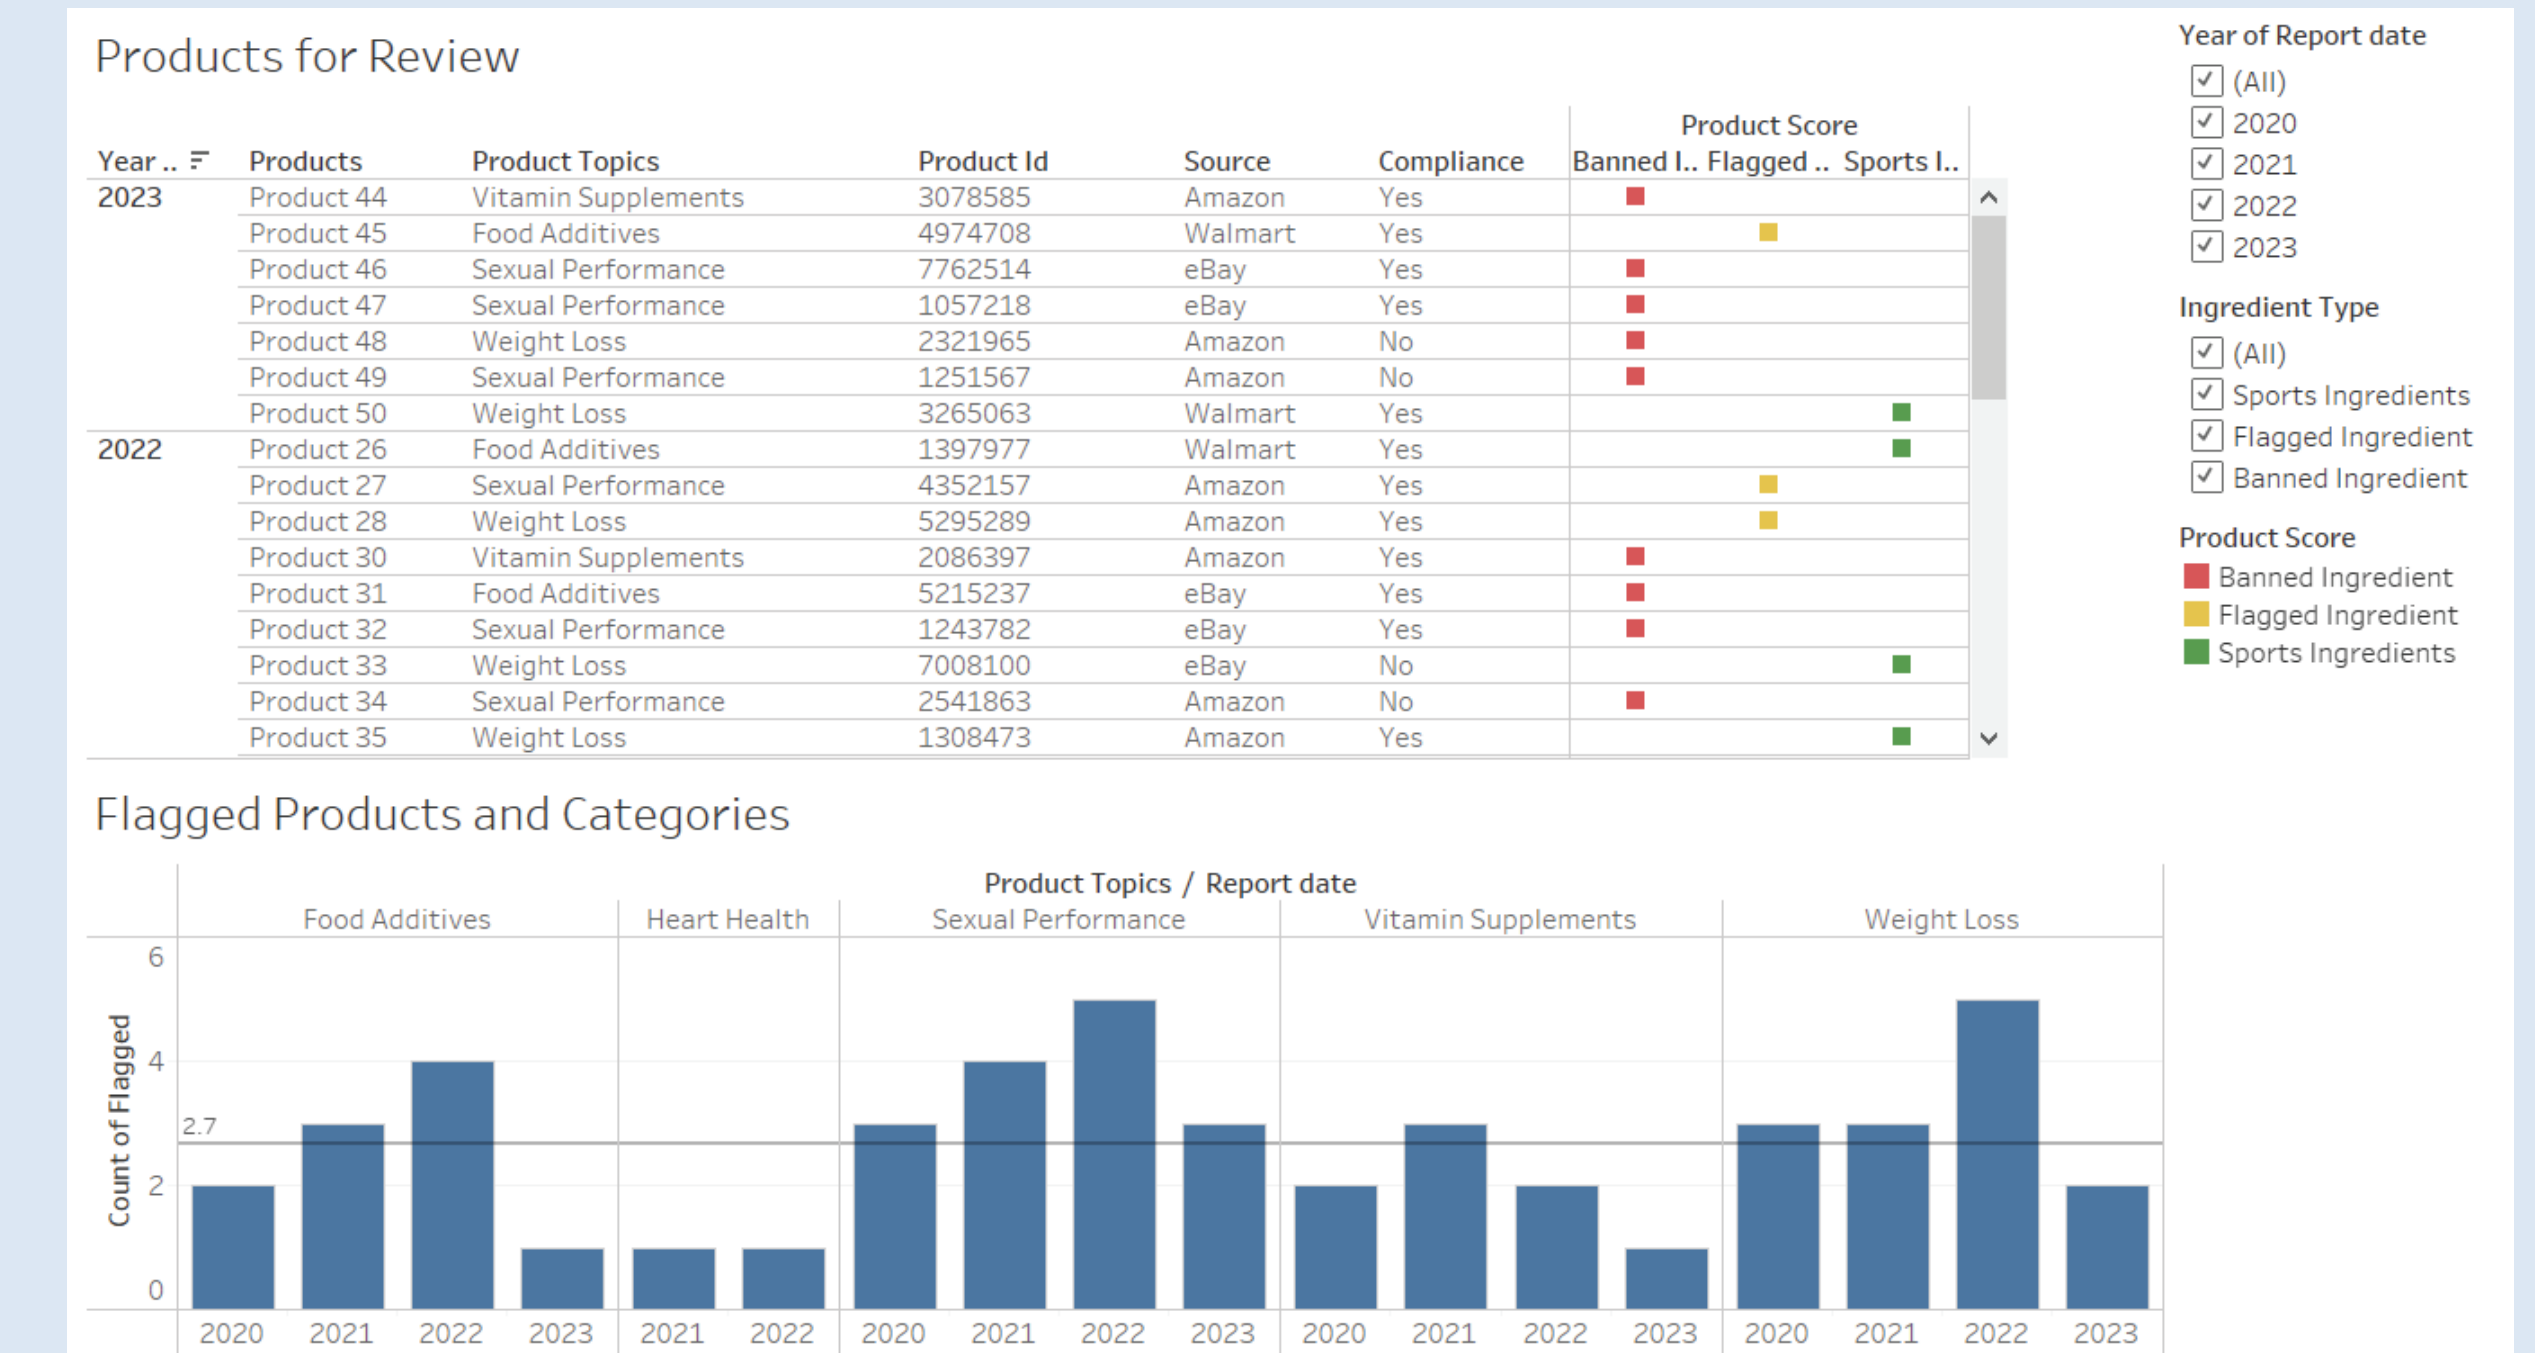


Figure 5. Tableau results dashboard

Conclusion

This Online Marketplace Research & Surveillance Tool enables HFB to proactively identify potentially fraudulent products. Through our marketplace APIs we have access to 52% of all e-commerce in the country, and in the first two months of functionality we analyzed ~39% of listed health supplements on these marketplaces. In a matter of hours, the tool analyzed ~150,000 products and identified 87 with illegal ingredients and 7,846 with at least 1 problematic claim. These results were ordered for ease of review with each signal's weight displayed for transparency, and model explainability features incorporated throughout.

Additionally, we designed with extensibility in mind, and the tool may be applied to use cases beyond health fraud. The pipeline is scalable in a cost-sensitive manner and integrated with FDA's Compliance Management System (CMS), allowing HFB to monitor confirmed fraudulent products to verify removal or compliance with FDA warnings and citations. Ultimately this tool will may assist HFB in identifying problematic products more efficiently to consider building a case for a potential investigation.