Automating Active Surveillance of Online Health Fraud using Artificial Intelligence

Marla Hendriksson[1], Sunita Joshi-Taylor[1], Ben Ballintyn[2], Tyree Giles[2], Alex Hanson[2], Anna Hoffman[2], Caroline Trier[2]

[1]US Food and Drug Administration, ORAHQ, OC [2]Booz 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, total revenue from e-commerce grew by approximately 200% from $340 billion in 2019 to $1.03 trillion in 2022. Considering health fraud as our use case, the current method of identifying of potentially tainted or misbranded products is a time-consuming manual process. The Online Marketplace Research and Surveillance Tool project seeks to leverage AI/ML tools to automate the monitoring of online marketplaces to lessen the burden of online surveillance. Designed with extensibility in mind, our proof of concept facilitates the identification of potentially fraudulent products sold through e-commerce by flagging problematic products more efficiently.

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

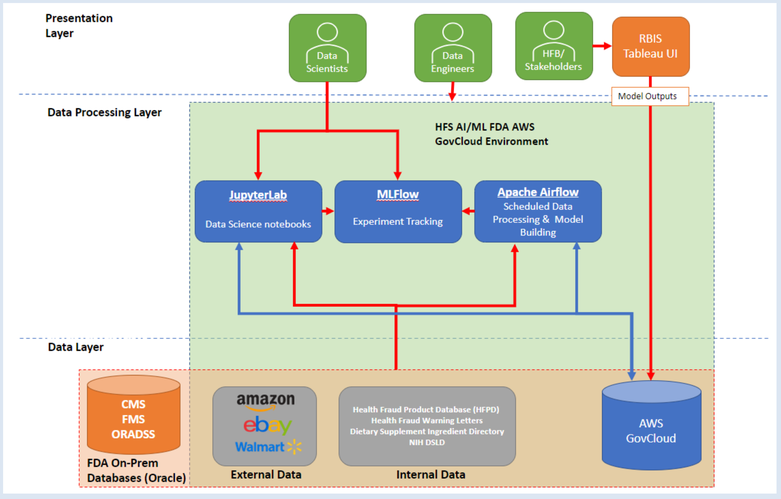

Figure 1. Cloud Architecture of OMRST

As shown in Figure 1, The Online Marketplace Research and Surveillance Tool 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 and action by FDA.

Marketplace Surveillance Coverage

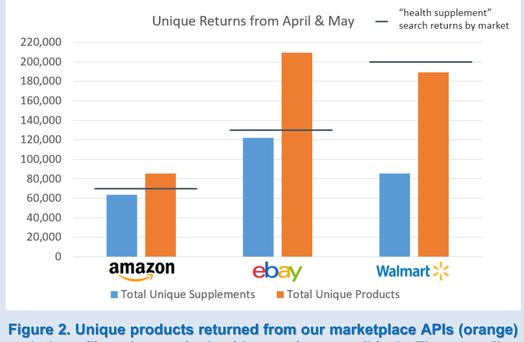

Figure 2. Unique products returned from our marketplace APIs (orange) and then filtered to only health supplements (blue). The gray line depicts the estimated total number of health supplements available per marketplace based on the search "health supplements".

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.

Disclaimer: The information in this poster represents the opinions of the authors and does not necessarily represent the FDA's official position or policy

Online Marketplace Surveillance Tool Models

Claims Model: Automatic identification of problematic claims made by products

Claim Identification

…Turmeric has been considered to be skin food for thousands of years in India and other cultures' Cancer Prevention. Turmeric can help prevent prostate cancer, stop the growth of existing prostate cancer and even destroy cancer cells...

Turmeric has been considered to be skin food for thousands of years in India and other cultures Cancer Prevention.

Turmeric can help prevent prostate cancer, stop the growth of existing prostate cancer and even destroy cancer cells.

Ingredient Turmeric has been considered to be skin food for thousands of years in India and other cultures Cancer Prevention.

Ingredient Turmeric can help prevent prostate cancer, stop the growth of existing prostate cancer and even destroy cancer cells.

Step 1: Preprocessing
- Start with product's description
- Reformat text
- Tokenize into individual sentences

Step 2: Identify claims
- Pass each sentence through finetuned DistilBERT entity model
- Model labels words as 'ingredient', 'condition', or 'Outside'

Step 3: Filter out non-claims
- A claim is defined to be a sentence that contains both an ingredient and a condition

Claim Classification

Binary Classifier

- Classifier is a finetuned version of DistilBERT

Dataset	
Positive Examples	Negative Examples
Claims extracted from FDA warning letters	Claims extracted from front-page products with >500 reviews

Assumptions:
1. Warning letter claims accurately represent fraud
2. Popular product claims represent acceptable language

Initial Results:
- Out of 151,421 products, 7,846 (5.18%) were found to have at least 1 problematic claim

Deep One-class Classification

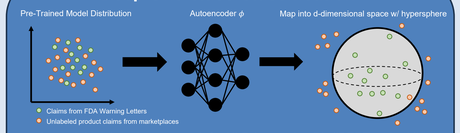

A novel deep learning approach for NLP, leveraging DistilBERT embeddings, contrastive pretraining and Deep SVDD for effective differentiation of nuanced textual data. Assumes few of the unlabeled examples are part of the positive class, allowing them to become part of the learning process.

Model Explainability

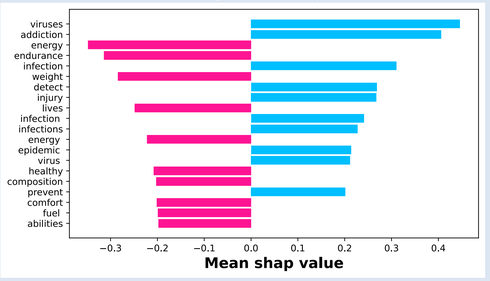

Figure 3. Top 20 most impactful tokens as defined by mean shap value across 5000 problematic claims run through the binary classifier. Blue (+) / Red (-) values increase/decrease output score (where 1 = problematic).

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.

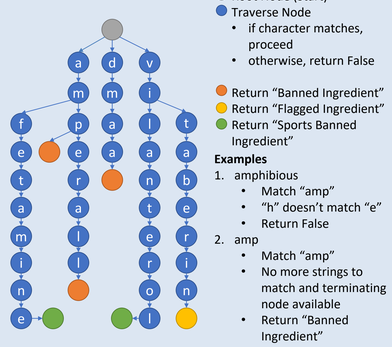

Figure 4. A visualization of a trie with the ingredients amfetamine, amp, amperall, dmaa, vilanterol, and vitaberin

Potential 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 ordered products by the following algorithm: *max((0,1) banned ingredients, (0-1) max(claims confidence score)* and display to a Tableau dashboard.

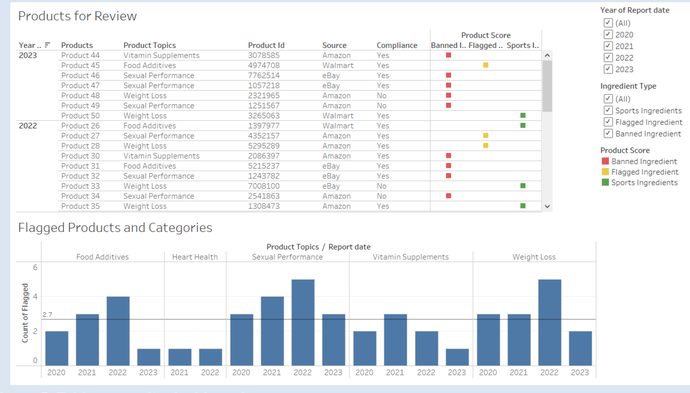

Figure 5. Tableau results dashboard

Conclusion

The Online Marketplace Research and Surveillance Tool is a proof of concept designed to assist FDA in proactively identifying potentially problematic products. Through our marketplace APIs we have access to 52% of all e-commerce in the United States. Our ML models and algorithms were able to analyze ingredients and claims and return promising results for FDA review.

In the first two months of functionality, we analyzed ~39% of health supplements listed 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 were incorporated throughout.

Additionally, we designed the tool with extensibility in mind so it may be applied to other use cases involving online marketplace research and surveillance.

The pipeline is scalable in a cost-sensitive manner and integrated with FDA's Compliance Management System.