Analyzing Public Perception and Filtering Misinformation of Electronic Cigarettes on Twitter: A Natural Language Processing Approach

Alexander Xu

Mentor: Nicole Spinelli

Great Neck South High School, Great Neck, New York

ABSTRACT

Electronic cigarettes (ECs) contain toxic metals that cause neurodevelopmental harm and delivers nicotine at levels comparable to traditional cigarettes. Studies have shown that social media may be perpetuating EC misconceptions among adolescents. This study scraped 9 million EC-related tweets from Twitter, a young-adult-centered platform, to investigate public perception towards ECs using natural language processing. Tweet sentiment, emotion, and topic were classified using the Valence Aware Dictionary for Sentiment Reasoning, Bidirectional Encoder Representations from Transformers, and Latent Dirichlet Allocation, respectively. Structural virality scores were used to analyze the dynamics of EC information dissemination. Overall, positive perceptions of ECs were more prevalent on Twitter, suggesting that EC brands may have successfully created a positive image of their product among Twitter users. Marketing, EC flavors, social appeal, echo chambers, lack of central authorities, and unimplemented Tobacco 21 legislation were found to be potential contributors to this phenomenon. Significant changes in tweet patterns were observed during headline events such as the EVALI outbreak in August 2019. This study elucidates factors affecting positive sentiment and tweet dissemination dynamics and will guide policymakers in implementing more effective preventive and quitting strategies. A browser plugin was developed to filter tweets with misleading EC information from users' feeds and provide credible sources of information and can be easily integrated into social media platforms to impose corrective actions.

INTRODUCTION

Electronic cigarette epidemic

Electronic-cigarettes (ECs) are electronic devices that simulate traditional combustible cigarette (CC) smoking. ECs heat an oil cartridge, or pod, to create vapor, containing flavoring and other chemicals. These particles are breathed into the lungs, which is commonly referred to as vaping. ECs resemble small USB flash drives and can be used discreetly at home, work, or school.

Many Americans, particularly youth and young adults, have grown addicted to vaping. In 2020, approximately 3.6 million, or one in five American high schoolers have used ECs before, which was a 78% increase since the prior year, and 27.6% of these high school EC users vape daily [5][14]. And Pierce et al. (2022) estimated that youth and young adults younger than 25 years of age were over twice as likely to vape daily compared to adults over 25 years of age [27]. Due to this rapid and widespread EC usage, ECs are classified by the CDC as an ongoing epidemic.

EC growth has been driven primarily by teen and young adult sales, which may have been by design [20]. The products are meant to mimic other sleek and trendy high-tech devices that are social status symbols highly coveted by young adults. EC pods also came in many flavors that appeal to youth, such as mango, creme, and cucumber. These flavors are much more appealing to adolescents than tobacco CCs. EC brands also employ youth models, social media influencers, and celebrities to advertise their products to younger audiences and become viral on social media.

JUUL is of special interest because it is the most popular brand of ECs and accounts for 55% of all EC sales in the US and is sometimes synonymous with ECs in youth vocabulary and culture [19][24].

Health risks of vaping

EC vapor contains as much nicotine as CCs [29]. Concentrations of urinary cotinine, a metabolite of nicotine, are sometimes even higher in vapers than in traditional smokers [16]. Nicotine primes youth for psychoactive substances and makes them more vulnerable to addiction throughout adulthood [7]. In 2017, Demissie [12] found that adolescent EC use was associated with a 3x increase in alcohol consumption, a 3x increase in prescription drug abuse, and a 4x increase in marijuana use. EC use was also associated with a 3x increase in future CC use [35]. Nicotine also impairs the neurodevelopment and executive function (e.g. memory, focus) of adolescents.

EC batteries and heating coils also expose users to carcinogenic and toxic heavy metals [15]. Vitamin E acetate (VEA), an additive found in ECs, produces carcinogenic compounds and becomes a highly toxic gas when heated. VEA was associated with thousands of cases of novel lung disease in the 2019 Vaping Product Use Associated Lung Injury outbreak (EVALI) [39].

EC misinformation on social media

Despite these health implications, most youth and young adults on social media have a positive perception of ECs and JUUL. In 2020, Visweswaran et al. [37] found that 62.4% of non-commercial EC-related tweets on the social media platform Twitter expressed positive sentiments toward vaping. In 2015 Cole-Lewis et al. [11] manually categorized most of these tweets as personal experiences and opinions posted by common people (as opposed to news or marketing sources).

One reason for this positive perception may be the prevalence of EC misinformation on social media platforms. Misinformation is widespread and increasing with the growth of social media and the internet [8][38]. EC brands advertise their products as smoking cessation tools and a safer alternative to CCs [29]. Frequent exposure to EC ads on Facebook doubles the likelihood of subsequent e-cigarette use one year later among the EC naive [6]. In 2018, Pepper and Farrelly [26] found that over half of youth who used electronic cigarettes believed that they don't contain nicotine. Instead, these youth believed that the ECs contained synthetic forms of nicotine and were, therefore, safer than CCs.

Social bots may also perpetuate these misconceptions, as they generate more tweets than the average human user, and are more likely to post hashtags that reference smoking cessation and new EC products [3]. Moreover, health organizations are not the primary sources of health information on social media, making reliable sources of information scarce [11][25].

Twitter for public health surveillance

Modern public health surveillance is the process of collecting user-generated content on social media platforms, like Twitter, using keywords, hashtags, and other parameters to monitor public awareness of epidemics or reactions to legislation. Twitter in particular has several advantages over other social media platforms and traditional forms of surveillance like surveys. 1) It is a large corpus of real-time data (500+ million tweets per day), 2) Twitter users are less conservative and more likely to express their true experiences and emotions when formulating their answers [25], 3) Approximately 60% of Twitter users were between 18 and 30 years of age, making Twitter the ideal platform for investigating youth and young adult populations [34].

Natural language processing

Researchers commonly employ Natural Language Processing (NLP), a field of artificial intelligence that can analyze textual data, to monitor the millions of daily tweets [13][17][21]. NLP has been employed in other epidemics like Ebola to find the underlying sentiments, topics, and emotions [22][30].

Sentiment analysis models classify a text into positive, negative, or neutral sentiments toward a topic. Several models have been used for this task in the past, namely Linguistic Inquiry and Word Count (LIWC, statistical), Affective Norms for English Words (ANEW, rule-based), General Inquirer (rule-based), Valence Aware Dictionary for Sentiment Reasoning (VADER, rule-based) Naive Bayes (supervised learning), and Support Vector Machines (SVM, supervised learning) [25]. While these models were able to classify tweet sentiment with upwards of 80% accuracy, VADER outperformed all of these models with an accuracy of 0.96, precision of 0.99, recall of 0.94, and r=0.9 correlation with human annotators on a set of tweets [18]. VADER determines the sentiment of a text by using its dictionary of assigned 'valences', or sentiment weights, of common words, phrases, social media slang, emoticons, punctuation, and capitalization. It calculates the sum of valences in a body of text while using predefined language rules and patterns (e.g. parts of speech, negation keywords) to manipulate the polarity and weighting of phrases. This value is normalized between -1.0 and 1.0 with -1 representing a completely negative sentiment and 1 representing a completely positive sentiment tweet.

Topic modeling is used to determine the latent themes, or topics, in a text. Latent Dirichlet Allocation (LDA), an unsupervised statistical model, is the most frequently used model for this task [28][36]. LDA works on the assumption that topics are composed of a set of words and 'documents' (i.e. tweets) are composed of a set of topics. LDA cannot understand the meaning behind topics but can unveil the latent patterns connecting groups of similar words. The model initially assigns random topics to each word. At each iteration, for each word, the model finds the probability of the word belonging to each topic (based on each topic's word distribution). Then, the word is reassigned topic with the highest probability and updates the topic distributions with the new assignment. This process is performed until a stable state is reached, or when topic distributions no longer change significantly from one iteration to the next. Now, the topic distributions can be used to identify the hidden topics in the document. While large-scale topic modeling on EC-related tweets has not yet been performed, it has been done on a smaller scale through manual annotation with tobacco products. In 2013, Myslín et al. [25] found that the most common themes among tobacco-related tweets were hookah, cessation, and pleasure. Topics related to recreation, social interaction, marketing, discreet vaping, and flavoring were correlated with positive sentiment [32], while topics related to cessation, health, 'big tobacco' ulterior motives, and social image correlated with negative sentiment [2].

Emotion analysis is usually used to classify a text into Ekman's emotion categories: joy, anger, fear, and neutral. Because emotions are more nuanced than sentiment and require a deeper understanding of semantics, neural network models are generally used. In the past, Long Short-Term Memory Networks (LSTM) and Recurrent Neural Networks (RNN) were used. However, these models did not remember complex and distant (separated within large bodies of text) relationships between words well due to their model architecture and could not be parallelized (i.e. models had to be trained on small datasets due to time constraints). In recent years, the Bidirectional Encoder Representations from Transformers (BERT) models have become popular because of their transformer model architecture [31]. Transformer models use attention mechanisms, allowing them to be parallelized and trained on terabytes of data, learn the 'meaning' of words by their context, and perform better than LSTMs and RNNs in almost every NLP task. BERT is a state-of-the-art Transformer model because of its application of bidirectional training (as opposed to left-to-right or right-to-left training), giving BERT a deep understanding of language and context. This training technique is called Masked Language Modeling (MLM), where random parts of a text are masked and the model predicts the missing words. The cased-BERT model classified six emotion categories with an accuracy of 0.90 and an

F1 score of 0.91 [31]. BERT frequently missed the happiness class, with a recall of 0.85, but recognized sadness well with a recall of 0.96.

Information dynamics

Several studies have attempted to quantify the spread of information or misinformation, on social media [22][30]. The spread of information is measured on a social network of connected retweeters (aka. An information cascade or diffusion tree). The method of information dissemination falls between two ends of a spectrum: broadcast, where information quickly reaches many users from a single source (e.g. news organizations, public influencers, government), and viral, where information is spread farther in a peer-to-peer process (long retweet chains reaching many communities). Understanding whether EC information spreads primarily in a broadcast or viral fashion can help stop the spread of misinformation and ensure public health campaigns reach a wide audience. In 2019, Liang et al. [22] determined that Ebola-related information spread primarily in a broadcast fashion with over 50% of retweet chains in information cascades less than 1 retweet in depth. This may be because of major health organizations' mainstream coverage of the Ebola pandemic. However, JUUL has been hypothesized to spread in a viral peer-to-peer fashion on social media by a landmark JUUL study by Stanford Research into the Impact of Tobacco Advertising [20]. Users with positive EC sentiment were also observed to be more likely to interact with users of similar sentiment (i.e. echo chambers, where misconceptions are reinforced) [10].

Research objectives

One objective was to uncover patterns in sentiment, emotions, and topics, in EC-related tweets over time. How these trends change based on EC-related events and state legislation was also of interest. Another objective was to determine the method of EC information dissemination and how sentiment affects its dynamics.

Currently, social media platforms like Twitter do not show warnings on posts containing potentially misleading EC information. One solution is to attach warning labels to a social media post that indicates to users that it may contain misinformation [8]. Warnings that include explanations of why the post had been marked or links providing additional information are found to be more effective. This is known as corrective communication, which directly refutes misinformation and provides an alternative truth. This type of software is effective due to how frequently teens check social media. Frequent exposure to age-appropriated quitting advice has proven to be effective in helping teens quit ECs (e.g. the text message-based program developed by the Truth Initiative). So, another objective is to create a browser-based plugin that applies these strategies to filter EC misinformation on Twitter.

METHODS

Data sources

Tweets were collected using the Social Network Scrape Python library (SNScrape) to access historical Twitter data without the rate limitations and tweet cap of the official Twitter API. All tweets matched all of the following criteria: 1) English text that included “juul", "juuls", "vape", "vapes", "juuling", "vaping", "juuled", "vaped", "juuler", "vaper", "juulers", or "vapers” using regular expression matching [37] 2) Posted between January 1st, 2015 and November 30th, 2022 (due to resource limitations) 3) Had unique content, in the case of multiple tweets containing the same text, only the tweet that was posted first was included to exclude potential spam or bot-generated content. The latent tweet place attribute was only queried for a random sample of tweets. Quote tweets, which are retweets that contain added information by the retweeter, could not be extracted using SNScrape and were extracted using the official Twitter API. For each quote tweeter, their user followee lists were also extracted using the official Twitter API.

Tweet preprocessing

To remove noise from the data for statistical model analysis, stop words, irrelevant words (e.g. articles), non-Latin characters, punctuation, links, and user tags were removed [4][23][25]. Next, all letters were converted to lowercase [1][28]. Words were then lemmatized using the Python NLTK-WordNet package, which removed inflectional endings and converted words to their roots (e.g. “connected” and “connecting” were lemmatized to “connect”).

Sentiment analysis

VADER was used to determine whether a tweet exhibited positive, negative, or neutral sentiments [18]. Sentiment intensity ranged from -1 (outspoken and negative) to 0 (completely neutral) to 1 (outspoken and positive), so tweets with intensities between -0.25 and 0.25 were classified as neutral.

Topic modeling

LDA was used to identify the topics among EC tweets [28]. The Python Gensim LDA model was run to identify 20 latent topics and the optimal lambda relevance value of 0.6 [33]. Topic terms were then manually vetted for relevance and generalizability and named with an identifying term. Each tweet had one topic, so the most probable topic by LDA was assigned to the tweet.

Emotion analysis

BERT was used for emotion analysis [31]. To apply transfer learning, a linear transformation layer was added to the BERT-base-cased model from Hugging Face. The linear transformation had 4 out-features to fine-tune the model for a dataset annotated by Sailunaz and Alhajjab in 2019 [30], which contains labeled tweets expressing anger, fear, joy, and neutrality (Ekman's emotion model). The model reached the minimum validation loss after training for 6 epochs at a learning rate of 10^-6 and with the Adam optimizer. The model performed with an overall testing accuracy and f1-score of 0.92.

Emotion	Precision	Recall	F1-score
Joy	0.95	0.93	0.94
Anger	0.93	0.84	0.88
Fear	0.89	0.90	0.90
Neutral	0.98	1.00	0.99

Information dissemination

Diffusion trees were constructed by adapting the information cascade method (Liang et al., 2019). A tree was created by joining each quote retweeter node to the original tweet poster node with a directed edge. Then for each retweeter node A, its outward edge connected to node B was switched

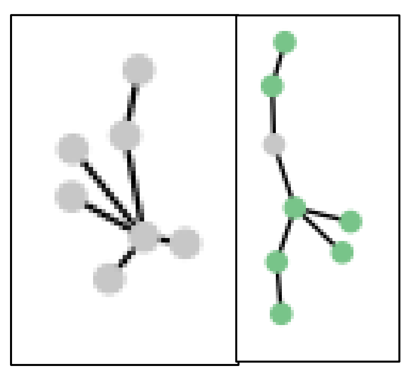

from node B to another node C if A is in the follower list of C and C's tweet was posted after B's tweet and before A's tweet. Next, using the breadth-first-search tree traversal algorithm the structural virality score (average pairwise distance between nodes), cascade height (max tree depth), cascade scale (number of direct descendants of the root), and cascade size (number of nodes) were calculated. A normalized structural virality score of 0 means that the diffusion tree is purely broadcast (similar to the figure on the left) while a score of 1 means the diffusion tree is purely viral (similar to the figure on the right).

Statistical analyses

The t-test was used to compare continuous variables. The chi-square test for independence was used to assess the association between categorical variables. Structural virality scores were min-max normalized before analysis. Model performance was assessed using accuracy, precision, recall, and f1-score. All statistical analyses were done in Python SciPy with an alpha level of $p<0.05$ unless otherwise specified.

Application

The front-end Chrome extension was developed using JavaScript and the back-end server was developed using the Python Flask framework. The frontend application scrapes the user's Twitter feed for EC-related keywords. The textual information is then sent to the server's classification API endpoint via the body of an HTTP request. The backend then uses the VADER model and pre-determined topics by LDA to return the overall tweet classification and misleading terms in the HTTP response. The front-end application then uses this information to blur and highlight the user's feed accordingly.

RESULTS

Data extraction with SNScrape yielded 8,077,408 unique tweets (after excluding 1,007,224 duplicates).

Greater positive sentiment

VADER was used to classify all tweets into positive, negative, and neutral sentiments, as well as calculate sentiment intensity scores. There were 2,888,729 (35.8%) tweets expressing positive EC sentiment, 1,788,097 (22.1%) tweets expressing negative sentiment, and 3,400,582 (42.1%) tweets expressing neutral sentiment. So overall, 1.6 times more tweets expressed positive rather than negative sentiment towards ECs. The vast majority of neutral sentiment tweets were also classified as slightly positive (Figure 1), though the level of intensity was negligible and therefore likely due to inherent bias in the VADER model or speech.

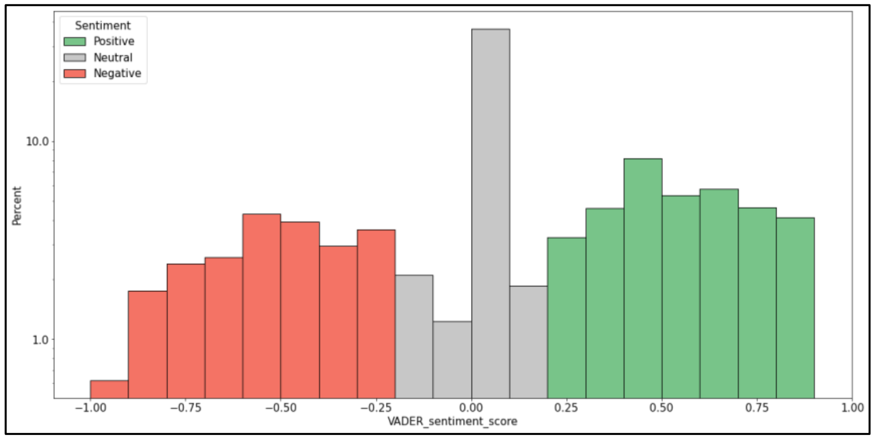

Fig 1. *Tweet sentiment intensity distribution. Histogram shows n=8,077,408 tweets (y-axis is log scale), with hues showing sentiment classified by VADER. Negative sentiment is red, neutral is gray, and positive is green.*

Sentiment polarization

The mean sentiment intensity score among positive sentiment tweets was 0.568 (±0.00012) compared to a mean sentiment intensity score of 0.527 (±0.00015) among negative sentiment tweets (Figure 2). Positive sentiment tweets had significantly greater (albeit a small increase in) sentiment intensity scores than negative sentiment tweets (p<0.001).

There were significant changes in tweet patterns during the E-cigarette or Vaping Use-Associated Lung Injury (EVALI) outbreak in 2019. Before this event, Twitter users posted 1,157,283 EC-related tweets per year, with only 54.8% of tweets expressing strong positive or negative sentiments. During the EVALI outbreak, the average number of tweets per day increased by 30% from 4,546 to 5,954, and the percentage of tweets expressing negative sentiment towards ECs increased by 128% from 16.2% to 36.9% (Figure 3).

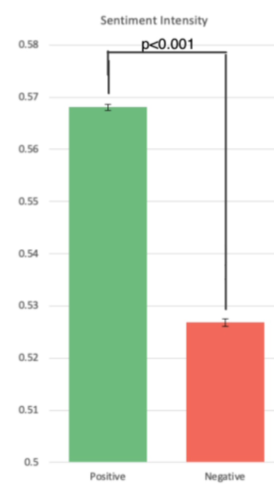

Fig 2. *Mean sentiment intensity by sentiment class. The bar plot shows n=4,676,826 tweets, with mean normalized sentiment intensity scores calculated using VADER (y-axis starts from 0.50) and error bars showing ± 1 SE.*

During the outbreak, the mean sentiment intensity score of negative sentiment tweets decreased to 0.60, a noticeable difference from the average of 0.53. Interestingly, the mean sentiment intensity of positive sentiment tweets also slightly decreased. Following this event, the average number of daily EC-related tweets decreased to just 1,685 tweets per day, but the percentage of tweets expressing strong opinions rose to 63.8%. In 2022, JUUL's marketing denial order (MDO) also temporarily spiked EC conversation on Twitter with similar increases in negative sentiment tweets and overall volume.

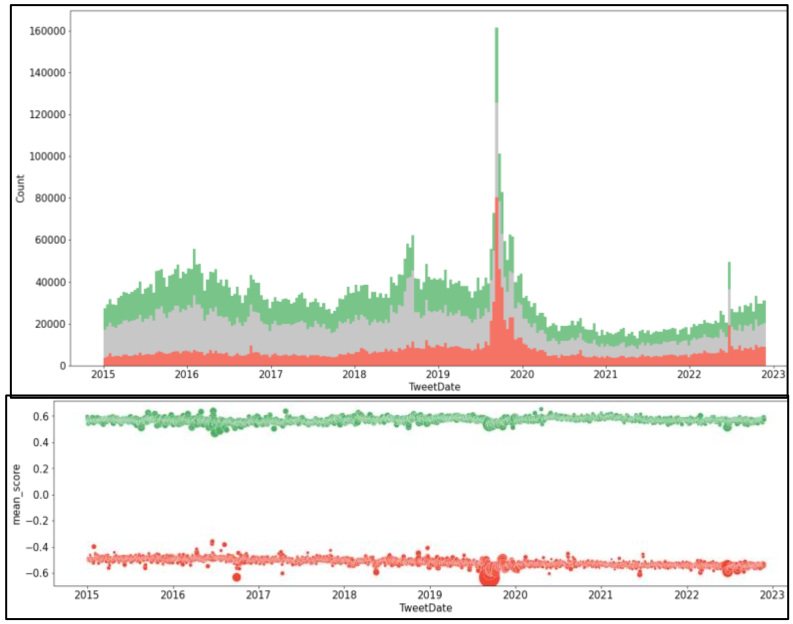

Fig 3. *Tweet sentiment intensity breakdown over time. Histogram shows n=8,077,408 tweets. Scatterplot shows n=4,676,826, and mean intensity scores from VADER.*

Trends in topics

LDA was used to determine the most common topics, including anti-EC legislative action (i.e. banning), social and community appeal (e.g. #vapelife), smoking cessation, vaping products (e.g. juices), usage in school, health risks, marijuana, flavors, marketing (e.g. giveaway), and addiction to ECs. Next, LDA was used to determine the most probable topic for each tweet. The five most common topics associated with positive sentiment tweets were vaping products (9.3%), social and community appeal (8.3%), marketing (7.6%), marijuana (5.5%), and smoking cessation (3.6%) (Figure 4).

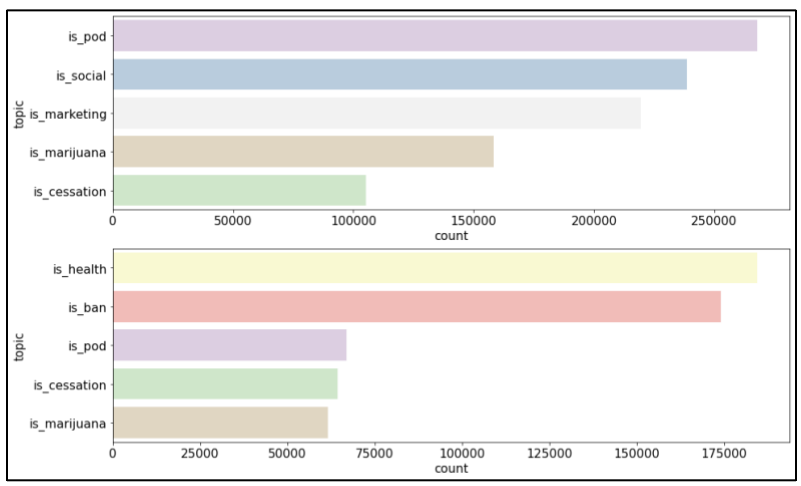

Fig 4. *Popular topics by sentiment. Top bar plot shows top 5 topics obtained by LDA through n=2,888,729 positive sentiment tweets. Bottom shows top 5 topics for n=1,788,097 negative sentiment tweets.*

The five most common topics associated with negative sentiment tweets were health risks (10.3%), banning (9.7%), vaping products (3.7%), smoking cessation (3.6%), and marijuana (3.5%). It is important to note that health risks and banning were the most common topics by far, which may be by-products of the EVALI and MDO events.

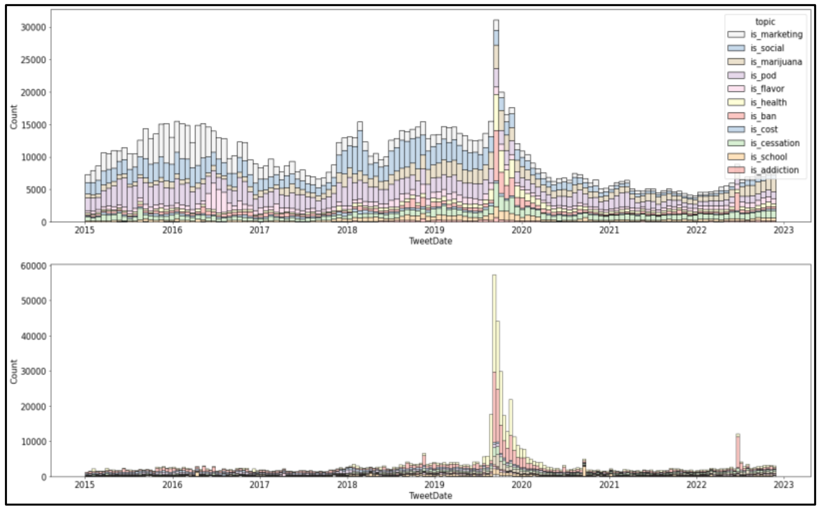

Fig 5. *Topic distribution over time. Top histogram shows topics for n=2,888,729 positive sentiment tweets. Bottom histogram shows topics for n=1,788,097 negative sentiment tweets.*

The distribution of topics changes over time. Positive sentiment tweets, flavors, and marketing are especially pronounced in JUUL's early years from 2015 to 2017 (Figure 5). For both positive and negative sentiment tweets, there is an increase in health risks, banning, and smoking cessation during the EVALI and MDO events. Following EVALI, the only topics that become more prominent are smoking cessation and marijuana.

Tobacco-21 is linked with sentiment and emotion

The latent tweet place attribute (user location) was queried for a random sample of 300k EC-related tweets, yielding 29,189 (9.7%) tweets with a set location. The tweet's origin state was determined through the location.

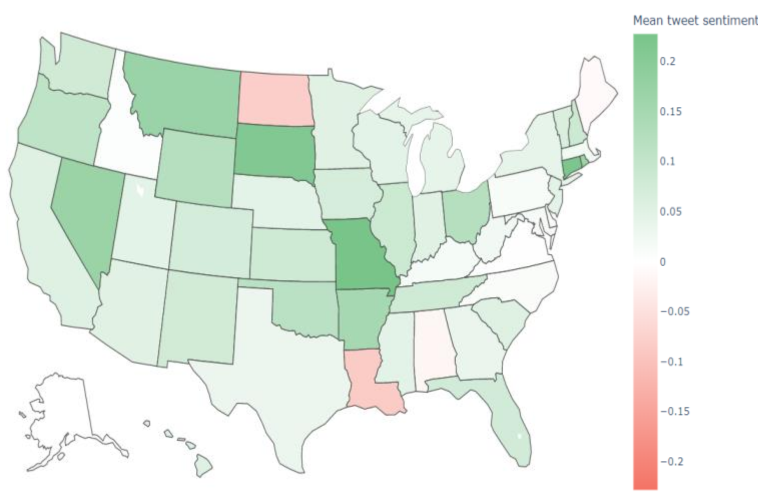

Fig 6. *Mean sentiment intensity by American state. Geoplot shows n=29,189 tweets. Intensity scores were obtained through VADER.*

Only four states had a negative mean sentiment score intensity: North Dakota, Louisiana, Alabama, and Maine (Figure 6). The state with the highest positive sentiment intensity was Missouri with a mean score of

0.227 and the state with the highest absolute negative sentiment intensity score was Louisiana with a mean score of 0.083. Three out of these four states with negative mean sentiment intensity scores have passed Tobacco-21 (T21) legislation, a campaign aimed to raise the legal age of purchase of tobacco/nicotine products to 21. In general, states that have acted upon T21 show significantly lower mean sentiment intensity scores compared to states that have not ($p<0.0001$) but are not different between the stages of legislation (Figure 7). 20 states have not acted upon T21 with a mean intensity of 0.10 (±0.0034), 15 have introduced legislation with a mean of 0.05 (±0.0081), 9 have passed with a mean of 0.06 (±0.0061), 6 have enacted with a mean 0.05 (±0.0073).

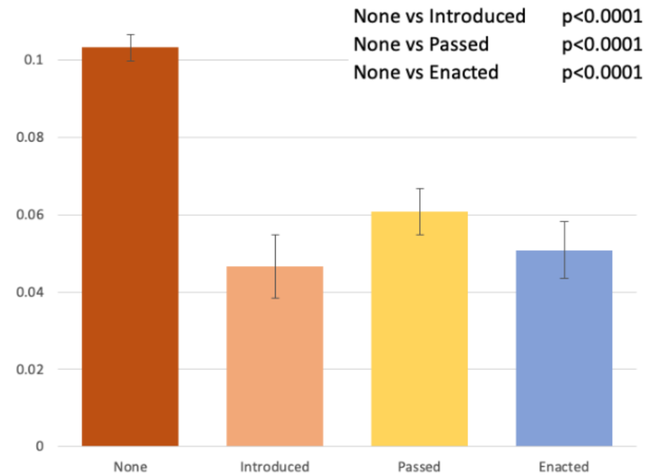

Fig 7. *Mean tweet sentiment intensity by Tobacco-21 status. Graph shows n=29,189, mean normalized intensity scores obtained by VADER, and error bars showing ± 1 SE.*

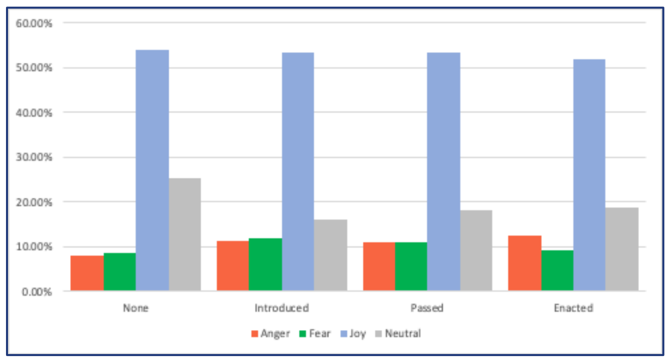

Fig 8. *Emotion classification by Tobacco-21 status. Graph shows n=29,189 tweets, with emotion classified by BERT.*

BERT was used to classify each tweet into the emotion categories of joy, anger, fear, and neutral. There was an association between emotions and T21 status ($\chi2=128.2$, $p<0.001$). In states that have acted upon T21, the proportion of tweets that express neutral emotions is 20% less, while the proportions of tweets that express anger and fear are 20% greater (Figure 8).

Broadcast model prevails

Quote tweets, which are retweets that contain added textual information by the retweeter, could not be extracted using SNScrape and were instead extracted using the official Twitter API. 12,600 Tweets were queried for retweets. In total, 245 (1.9%) diffusion trees were created (i.e. 245 of 12,600 tweets had any quote retweets) and had a total of 967 total quote retweets.

Diffusion trees overwhelmingly spread in a broadcast fashion. Raw structural virality had a mean of 1.289 (±0.034). Normalized structural virality had a mean of 0.104 (±0.012) and a median of 0.00 [0, 0.119]. However, it is more nuanced when broken down by sentiment intensity. Negative sentiment diffusion trees spread with higher structural virality (p=0.003), cascade size (p=0.033), cascade scale (p=0.041), and cascade depth (p=0.007) compared to positive sentiment trees (Figure 9). While the positive mean structural virality was 0.078 (±0.016), the negative mean structural virality was 0.174 (±0.031). And while the positive mean cascade size was 3.415 (±0.658), the negative mean cascade size was 6.625 (1.523). So overall, negative sentiment diffusion trees reached a larger audience and tended to spread more virally compared to positive and neutral sentiment diffusion trees. Additionally, further analysis showed that 90% of all direct quote retweets (depth=1) shared the sentiment of the original tweet. Each tweet further removed from the original discussion had a 1.2 times increase in the likelihood of disagreement with the original tweet.

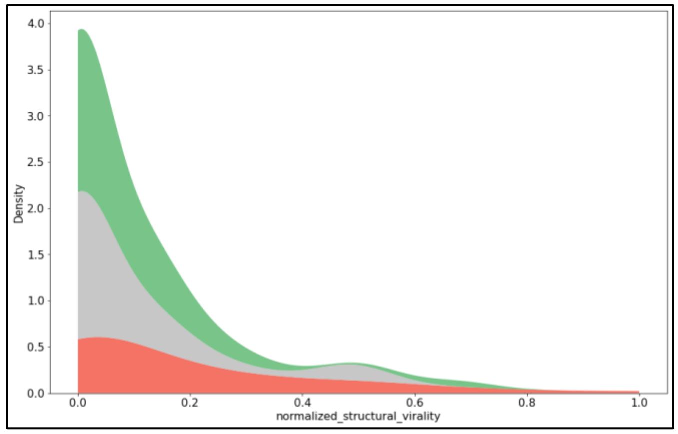

Fig 9. *Distribution of normalized structural virality scores by sentiment class. Graph shows n=245 diffusion trees, with structural virality scores calculated using breadth-first-search.*

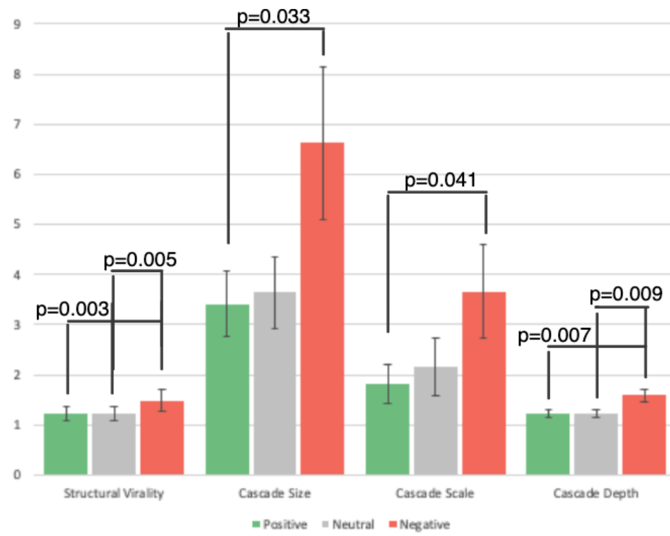

Fig 10. *Diffusion tree parameters by sentiment class. The bar plot shows n=245 diffusion trees, with all tree parameters calculated using breadth-first-search and error bars showing ± 1 SE.*

Application

The application was able to successfully filter potentially misinformed tweets using VADER and highlight misleading or enticing terms using the results of the LDA analysis.

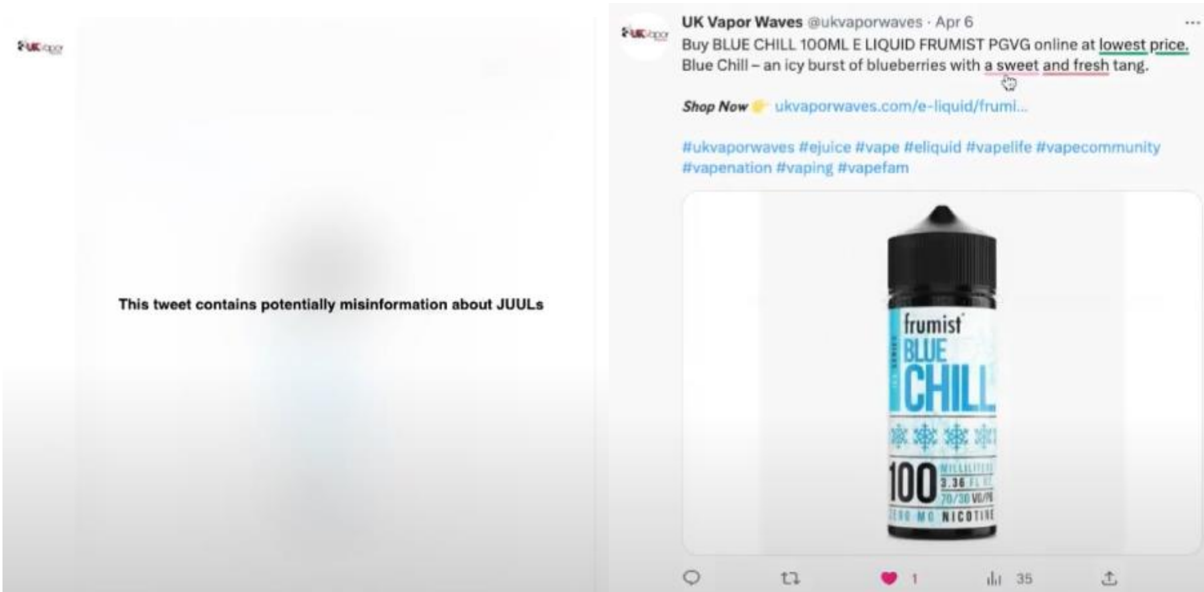

Fig 11. *Browser-based application. Left image shows how a tweet is hidden from the user's feed. This filter is removed once the user hovers over the tweet. Right image shows how potentially misleading, enticing, or dangerous terms are highlighted and linked to credible sources.*

DISCUSSION

State of current EC perception on Twitter

The findings of this study confirm several important observations documented in previous literature and greatly expand the current understanding regarding the state of EC information trends and dynamics on Twitter. The finding that approximately 60% more tweets expressed positive rather than negative sentiment is consistent with Visweswaran et al. from 2020 [37]. This suggests that EC brands have effectively created a positive image of their products among Twitter users, further evidenced by the greater vocalization of individuals with positive sentiments. This study sheds light on several potential factors contributing to this phenomenon. The first is the dominance of the vaping product, social and community appeal, and marketing themes across tweets. Over 25% of all EC-related tweets expressing positive sentiment revolve around these topics. Continuous exposure to new vaping products and EC-related marketing tactics may contribute to user addiction or catalyze future EC use among those who are initially unfamiliar with the product. The prominence of flavors and marketing during JUUL's early days on Twitter further validates Jackler's 2019 hypothesis [20], which proposed that JUUL's rapid growth was primarily driven by its youth appeal through social media influencer promotions and enticing flavor options. Furthermore, these findings suggest that the absence of T21 legislation may play a role in shaping EC perceptions. States that have introduced, passed, or enacted T21 legislation exhibit more negative overall perceptions of ECs, which may indicate a heightened awareness of the product's health risks among adolescents. In addition, the observed mean normalized structural virality score of 0.104 is twice as viral as Liang’s reported value of 0.050 in the context of Ebola [22]. This suggests that information about ECs is disseminated more extensively through peer-to-peer channels compared to other epidemics, driven by users' personal experiences and opinions, rather than through constant mainstream coverage by health organizations like the CDC in the case of Ebola. Moreover, the high proportion of agreement with direct retweets, low cascade depth, and high cascade scale indicate the presence of echo chambers within EC-related discussions. These structures may reinforce users' existing beliefs and make them less receptive to alternative perspectives. Given these findings, health organizations must expand their influence on social media platforms and increase educational initiatives targeting these users.

Escalating polarization in EC-related discussions

One important trend observed in this study is the increasing polarization in EC-related discussions on Twitter. As depicted in Figure 1, the distribution of sentiment intensities deviates from a bell-shaped curve, instead displaying a bimodal distribution that delineates two distinct communities. Over 8 years, a higher proportion of users expressed either positive or negative sentiments, accompanied by a noticeable decline in neutral sentiment (Figure 2). The geographical analysis indicates that the introduction, passage, or enactment of T21 laws is associated with a decrease in neutral emotions and an increase in anger, fear, and overall negative sentiment. In the diffusion tree assay, the prevalence of both positive and negative sentiment echo chambers serves as additional evidence of this polarization. The implications of this increasing polarization remain unclear and require further investigation. This polarization may reflect a growing awareness of the risks associated with EC use among individuals who were previously neutral or uninformed. However, because this is a retrospective study, it cannot be determined whether factors like T21 legislation effectively contributed to the rise in negative sentiment and emotions, or if the legislation was enacted in response to these pre-existing differences in sentiment. Considering the observed polarization, effective preventive and quitting applications should prioritize facilitating communication between these divergent groups and encouraging the sharing of viewpoints. Fostering dialogue and promoting the exchange of experiences and information can potentially foster a better understanding of the risks and consequences associated with EC use.

Study limitations

The Official Twitter API was a major limiting factor in this study. Deleted tweets are not archived by Twitter and could not be included in the analysis. This problem may affect tweets that were posted earlier more than recent ones. Twitter's tweet length policy also changed from 140 to 280 characters in 2017, which was not controlled for in this analysis. This may have biased sentiment and topical analysis to overcount tweets posted before this change as a user may have posted several tweets expressing the same theme and sentiment in the past, but only tweeted once after the change. Twitter's database also does not store all retweets (aka. statuses) together. So, the API caps the number of retweets that can be pulled at 100, which is why quote tweets were analyzed exclusively in this study (as there are rarely over 100 quote retweets), which is a limitation as quote tweet dynamics may differ from retweet dynamics. Twitter's database also does not store a reference directly to the tweet that a retweet originated from and instead stores a reference to the original retweet. This is why this study adapted the methods from Liang et al. [22], which entails reconstructing the diffusion tree using heuristics and may not always construct the real diffusion tree. This method also biases the broadcast method as links are formed only when the retweeter follows the tweeter when in reality retweeters can retweet the post of any user. Therefore the dissemination conclusions made in this study are primarily made through comparisons to prior studies that have used this method.

Future directions

Several important questions remain outside the scope of this investigation and warrant further exploration. Investigating the relationship between EC use and marijuana, as well as other forms of substance abuse is of great importance. This study identified marijuana as a prominent topic in EC-related discussions. However, the proportion and sentiments, and emotions of users who potentially engage in both EC and marijuana use remain unknown. Future research in this area should also aim to examine the behavioral and psychological characteristics of these people through their tweets. Additionally, changes in EC users should be examined in a longitudinal style and control for age and gender. This approach would elucidate whether positive sentiment towards EC at an early age is linked to subsequent discussions of other substances and suicidal ideation.

Expanding upon the current application, there is a need for a more general-purpose application that isn't limited to EC-related keywords and themes. This application should feature an adaptable interface capable of monitoring and analyzing any emerging epidemic across various social media platforms. And while this application relies on historical data, future applications should leverage real-time data from Twitter to identify and analyze emerging trends more promptly. Developing an application that can capture and analyze Twitter trends in real time would enable more immediate responses to public health concerns. The integration of models capable of identifying bot-generated content may also be beneficial.

CONCLUSION

The primary objective of this study was to provide a comprehensive report on the trends and dynamics of EC-related discussions on Twitter over the past eight years since the product's launch. As the first study to examine changes in EC perception and topics over time (including before and after major events), the findings shed light on the evolving landscape of EC-related discussion, including growing polarization between pro and anti-EC communities. Additionally, this is the first study to identify state legislation as a significant factor influencing EC sentiment. Furthermore, this study also quantified the dynamics of EC conversations on social media.

Overall, this analysis revealed that positive perceptions of ECs were more prevalent among Twitter users, indicating that the company has been successful in creating a positive image of its product on Twitter. Several factors were found to potentially contribute to this phenomenon, including marketing strategies, EC flavors, social appeal, the presence of echo chambers, the absence of central authorities, and the (lack of) implementation of T21 legislation. This study also observed significant changes in tweet patterns during headline events, such as the EVALI outbreak in August 2019. This understanding of the dynamics surrounding EC conversations will guide policymakers and health organizations in implementing more effective preventive and cessation strategies to address the EC epidemic. In addition to the insights gained from this study, a plugin that can be integrated into social media platforms to facilitate corrective actions was also successfully developed.

BIBLIOGRAPHY

[1] Aljedaani, W., Rustam, F., Mkaouer, M. W., Ghallab, A., Rupapara, V., Washington, P. B., Lee, E., & Ashraf, I. (2022). Sentiment analysis on Twitter data integrating TextBlob and Deep learning models: The case of the US airline industry. Knowledge-Based Systems, 255. https://doi.org/10.1016/j.knosys.2022.109780

[2] Allem, J.-P., Dharmapuri, L., Unger, J. B., & Cruz, T. B. (2018). Characterizing JUUL-related posts on Twitter. Drug and Alcohol Dependence, 190, 1–5. https://doi.org/10.1016/j.drugalcdep.2018.05.018

[3] Allem, J. P., Ferrara, E., Uppu, S. P., Cruz, T. B., & Unger, J. B. (2017). E-Cigarette Surveillance With Social Media Data: Social Bots, Emerging Topics, and Trends. *JMIR Public Health and Surveillance*, *3*(4), e98. https://doi.org/10.2196/publichealth.8641

[4] Basáñez, T., Majmundar, A., Cruz, T. B., & Unger, J. B. (2018). Vaping is associated with healthy food words: A content analysis of Twitter. Addictive Behaviors Reports, 8, 147-153. https://doi.org/10.1016/j.abrep.2018.09.007

[5] Besaratinia, A., & Tommasi, S. (2020). Vaping epidemic: challenges and opportunities. Cancer Causes & Control, 31(7), 663–667. https://doi.org/10.1007/s10552-020-01307-y

[6] Camenga, D., Gutierrez, K. M., Kong, G., Cavallo, D., Simon, P., & Krishnan-Sarin, S. (2018). E-cigarette advertising exposure in e-cigarette naïve adolescents and subsequent e-cigarette use: A longitudinal cohort study. Addictive Behaviors, 81, 78–83. https://doi.org/10.1016/j.addbeh.2018.02.008

[7] Casey, B. J., & Jones, R. M. (2010). Neurobiology of the Adolescent Brain and Behavior: Implications for Substance Use Disorders. The Journal of the American Academy of

Child and Adolescent Psychiatry, 49(12), 1189–1201. https://doi.org/10.1016/j.jaac.2010.08.017

[8] Chen, S., Xiao, L., & Kumar, A. (2023). Spread of misinformation on social media: What contributes to it and how to combat it. Computers in Human Behavior, 141, 107643. https://doi.org/10.1016/j.chb.2022.107643

[10] Cinelli, M., Morales, G. D. F., Galeazzi, A., Quattrociocchi, W., & Starnini, M. (2021, February 23). The Echo Chamber Effect on social media | PNAS. Retrieved October 6, 2022. https://www.pnas.org/doi/10.1073/pnas.2023301118

[11] Cole-Lewis, H., Pugatch, J., Sanders, A., Varghese, A., Posada, S., Yun, C., Schwarz, M., & Augustson, E. (2015). Social Listening: A Content Analysis of E-Cigarette Discussions on Twitter. Journal of Medical Internet Research, 17(10), e243. https://doi.org/10.2196/jmir.4969

[12] Demissie, Z., Everett Jones, S., Clayton, H. B., & King, B. A. (2017). Adolescent Risk Behaviors and Use of Electronic Vapor Products and Cigarettes. Pediatrics, 139(2). https://doi.org/10.1542/peds.2016-2921

[13] Edo-Osagie, O., De La Iglesia, B., Lake, I., & Edeghere, O. (2020). A scoping review of the use of Twitter for public health research. Computers in Biology and Medicine, 122, 103770. https://doi.org/10.1016/j.compbiomed.2020.103770

[14] *FDA Continues to Implement Law, Regulate Non-Tobacco Nicotine Products, Warns Retailers and Manufacturers Against Illegal Sales*. U.S. Food and Drug Administration. (2022, July 13). Retrieved January 13, 2023, from https://www.fda.gov/tobacco-products/ctp-newsroom/fda-continues-implement-law-regulate-non-tobacco-nicotine-products-warns-retailers-and-manufacturers

[15] Gaur, S., & Agnihotri, R. (2018). Health Effects of Trace Metals in Electronic Cigarette Aerosols—a Systematic Review. Biological Trace Element Research, 188(2), 295–315. https://doi.org/10.1007/s12011-018-1423-x

[16] Goniewicz, M. L., Boykan, R., Messina, C. R., Eliscu, A., & Tolentino, J. (2018). High exposure to nicotine among adolescents who use Juul and other vape pod systems

(‘pods’). Tobacco Control, 28(6), 676–677. https://doi.org/10.1136/tobaccocontrol-2018-054565

[17] Guiñazú, M. F., Cortés, V., Ibáñez, C. F., & Velásquez, J. D. (2020). Employing online social networks in precision-medicine approach using information fusion predictive model to improve substance use surveillance: A lesson from Twitter and marijuana consumption. Information Fusion, 55, 150-163. https://doi.org/10.1016/j.inffus.2019.08.006

[18] Hutto, C., & Gilbert, E. (2014). VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text. Proceedings of the International AAAI Conference on Web and Social Media, 8(1), 216-225. https://doi.org/10.1609/icwsm.v8i1.14550

[19] Huang, J., Duan, Z., Kwok, J., Binns, S., Vera, L. E., Kim, Y., Szczypka, G., & Emery, S. L. (2018). Vaping versus JUULing: how the extraordinary growth and marketing of JUUL transformed the US retail e-cigarette market. Tobacco Control, 28(2), 146–151. https://doi.org/10.1136/tobaccocontrol-2018-054382

[20] Jackler, R. K., Chau, C., Getachew, B. D., & Ramamurthi, D. (2019). JUUL Advertising Over its First Three Years on the Market. SRITA. https://tobacco.stanford.edu/

[21] Jordan, S., Hovet, S., Fung, I., Liang, H., Fu, K.-W., & Tse, Z. (2018). Using Twitter for Public Health Surveillance from Monitoring and Prediction to Public Response. Data, 4(1), 6. https://doi.org/10.3390/data4010006

[22] Liang, H., Fung, I.CH., Tse, Z.T.H., Yin, J., Chan, CH., Pechta, L.E., Smith, B.J., Marquez-Lameda, R.D., Meltzer, M.I., Lubell, K.M., Fu, KW (2019). How did Ebola information spread on twitter: broadcasting or viral spreading?. BMC Public Health 19, 438. https://doi.org/10.1186/s12889-019-6747-8

[23] Mackey, T. K., Li, J., Purushothaman, V., Nali, M., Shah, N., Bardier, C., Cai, M., & Liang, B. (2020). Big Data, Natural Language Processing, and Deep Learning to Detect and Characterize Illicit COVID-19 Product Sales: Infoveillance Study on Twitter and

Instagram. JMIR Public Health and Surveillance, 6(3). https://doi.org/10.2196/20794

[24] Malik, A., Li, Y., Karbasian, H., Hamari, J., & Johri, A. (2019). Live, Love, Juul: User and Content Analysis of Twitter Posts about Juul. American Journal of Health Behavior, 43(2), 326–336. https://doi.org/10.5993/ajhb.43.2.9

[25] Myslín, M., Zhu, S.-H., Chapman, W., & Conway, M. (2013). Using Twitter to Examine Smoking Behavior and Perceptions of Emerging Tobacco Products. Journal of Medical Internet Research, 15(8). https://doi.org/10.2196/jmir.2534

[26] Pepper, J. K., Farrelly, M. C., & Watson, K. A. (2018). Adolescents' understanding and use of nicotine in e-cigarettes. Addictive Behaviors, 82, 109–113. https://doi.org/10.1016/j.addbeh.2018.02.015

[27] Pierce, J. P., Zhang, J., Crotty Alexander, L. E., Leas, E. C., Kealey, S., White, M. M., Strong, D. R., Trinidad, D. R., McMenamin, S. B., Chen, R., Benmarhnia, T., & Messer, K. (2022). Daily E-cigarette Use and the Surge in JUUL Sales: 2017–2019. Pediatrics. https://doi.org/10.1542/peds.2021-055379

[28] Prier, K., Smith, M., Giraud-Carrier, C., Hanson, L. C. (2011). Identifying Health-Related Topics on Twitter An Exploration of Tobacco-Related Tweets as a Test Topic. Lecture Notes in Computer Science, 6589, 18-25. https://doi.org/10.1007/978-3-642-19656-0_4

[29] *Public health consequences of e-cigarettes*. The National Academies Press. (2018, January 23). Retrieved October 3, 2022, from https://nap.nationalacademies.org/catalog/24952/public-health-consequences-of-e-cigarettes

[30] Sailunaz, K., & Alhajjab, R. (2019). Emotion and sentiment analysis from Twitter text. Journal of Computational Science, 36. https://doi.org/10.1016/j.jocs.2019.05.009

[31] Saravia, E., Liu, H.-C. T., Huang, Y.-H., Wu, J., & Chen, Y.-S. (2018). CARER: Contextualized Affect Representations for Emotion Recognition. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.

https://doi.org/10.18653/v1/d18-1404

[32] Sidani, J. E., Colditz, J. B., Barrett, E. L., Chu, K., James, A. E., & Primack, B. A. (2019). JUUL on Twitter: Analyzing Tweets About Use of a New Nicotine Delivery System. Journal of School Health, 90(2), 135–142. Portico. https://doi.org/10.1111/josh.12858

[33] Sievert, C., & Shirley, K. E. (2014). LDAvis: A method for visualizing and interpreting Topics. https://doi.org/10.13140/2.1.1394.3043

[34] Sloan, L., Morgan, J., Burnap, P., & Williams, M. (2015). Who Tweets? Deriving the Demographic Characteristics of Age, Occupation and Social Class from Twitter User Meta-Data. PLOS ONE, 10(3), e0115545. https://doi.org/10.1371/journal.pone.0115545

[35] Soneji, S. (2017, August 1). E-cigarette use and subsequent cigarette smoking in adolescents and young adults. JAMA Pediatrics. Retrieved October 4, 2022, from https://jamanetwork.com/journals/jamapediatrics/fullarticle/2634377

[36] Sowles, S. J., Krauss, M. J., Connolly, S., & Cavazos-Rehg, P. A. (2016). A Content Analysis of Vaping Advertisements on Twitter, November 2014. *Preventing Chronic Disease*, *13*. https://doi.org/10.5888/pcd13.160274

[37] Visweswaran, S., Colditz, J. B., O'Halloran, P., Han, N.-R., Taneja, S. B., Welling, J., Chu, K.-H., Sidani, J. E., & Primack, B. A. (2020). Machine Learning Classifiers for Twitter Surveillance of Vaping: Comparative Machine Learning Study. Journal of Medical Internet Research, 22(8). https://doi.org/10.2196/17478

[38] Wang, Y., McKee, M., Torbica, A., & Stuckler, D. (2019). Systematic Literature Review on the Spread of Health-related Misinformation on Social Media. *Social Science & Medicine*, *240*, 112552. https://doi.org/10.1016/j.socscimed.2019.112552

[39] Wu, D., O’Shea, D. F. (2020). Potential for release of pulmonary toxic ketene from vaping pyrolysis of vitamin E acetate. PNAS, 117(12), 6349-6355. https://doi.org/10.1073/pnas.1920925117