

Categorising Vaccine Confidence with Transformer-Based Machine Learning Model: The Nuances of Vaccine Sentiment on Twitter

Per Kummervold, Sam Martin, Sara Dada, Eliz Kilich, Chermain Denny, Pauline Paterson, Heidi J Larson

Submitted to: JMIR Medical Informatics
on: April 13, 2021

Disclaimer: © The authors. All rights reserved. This is a privileged document currently under peer-review/community review. Authors have provided JMIR Publications with an exclusive license to publish this preprint on its website for review purposes only. While the final peer-reviewed paper may be licensed under a CC BY license on publication, at this stage authors and publisher expressly prohibit redistribution of this draft paper other than for review purposes.

Table of Contents

Original Manuscript..... 4



Categorising Vaccine Confidence with Transformer-Based Machine Learning Model: The Nuances of Vaccine Sentiment on Twitter

Per Kummervold¹ PhD; Sam Martin² LLB, MSc, PhD; Sara Dada³ MSc; Eliz Kilich³ BMBCh; Chermain Denny⁴ BSc; Pauline Paterson³ PhD; Heidi J Larson³ PhD

¹Centre for Clinical Vaccinology and Tropical Medicine (CCVTM), University of Oxford Oxford GB

²Faculty of Epidemiology and Population Health, London School of Hygiene & Tropical Medicine, London, UK London GB

³Faculty of Science, Vrije Universiteit Amsterdam, Amsterdam, Netherlands Amsterdam NL

Corresponding Author:

Per Kummervold PhD

Abstract

Background: With growing conversations online and less than desired maternal vaccination uptake rates, these conversations could provide useful insight to inform future interventions. Automated processes for this type of analysis, such as natural language processing (NLP), have faced challenges extracting complex stances, like attitudes toward vaccines, from large text.

Objective: In this study, we aimed to build upon recent advances in Transformer-based machine learning methods, and test if this could be used as a tool to assess the stance of social media posts towards vaccination during pregnancy.

Methods: A total of 16,604 Tweets posted between 1 November 2018 and 30 April 2019 were selected by boolean searches related to maternal vaccination. Tweets were coded by three individual researchers into the categories “Promotional”, “Discouraging”, “Ambiguous” and “Neutral” After creating a final dataset of 2,722 unique tweets, multiple machine learning methods were trained on the dataset and then tested and compared to the human annotators.

Results: We received an accuracy of 81.8% (F-score= 0.78) compared to the agreed score between the three annotators. For comparison, the accuracies of the individual annotators compared to the final score were 83.3%, 77.9% and 77.5%.

Conclusions: This study demonstrates the ability to achieve close to the same accuracy in categorising tweets using our machine learning models as could be expected by a single human annotator. The potential to use this reliable and accurate automated process could free up valuable time and resource constraints of conducting this analysis, in addition to inform potentially effective and necessary interventions. Clinical Trial: N/A

(JMIR Preprints 13/04/2021:29584)

DOI: <https://doi.org/10.2196/preprints.29584>

Preprint Settings

1) Would you like to publish your submitted manuscript as preprint?

Please make my preprint PDF available to anyone at any time (recommended).

Please make my preprint PDF available only to logged-in users; I understand that my title and abstract will remain visible to all users.

Only make the preprint title and abstract visible.

No, I do not wish to publish my submitted manuscript as a preprint.


2) If accepted for publication in a JMIR journal, would you like the PDF to be visible to the public?

Yes, please make my accepted manuscript PDF available to anyone at any time (Recommended).

Yes, but please make my accepted manuscript PDF available only to logged-in users; I understand that the title and abstract will remain visible to all users.

Yes, but only make the title and abstract visible (see Important note, above). I understand that if I later pay to participate in [http](#)

Original Manuscript



Categorising Vaccine Confidence with Transformer-Based Machine Learning Model: The Nuances of Vaccine Sentiment on Twitter

Per E Kummervold PhD^{1*}, Sam Martin PhD⁷, Sara Dada MSc², Eliz Kilich BMBCh², Chermain Denny BSc³, Pauline Paterson PhD^{2,6}, Heidi J Larson PhD^{2,4,5,6}

¹ FISABIO-Public Health - Vaccine Research Department, Valencia, Spain

² Faculty of Epidemiology and Population Health, London School of Hygiene & Tropical Medicine, London, UK

³ Faculty of Science, Vrije Universiteit Amsterdam, Amsterdam, Netherlands

⁴ Institute of Health Metrics and Evaluation, University of Washington, Seattle, USA

⁵ Chatham House Centre on Global Health Security, The Royal Institute of International Affairs, UK

⁶ National Institute for Health Research Health Protection Research Unit (NIHR HPRU)

⁷ Centre for Clinical Vaccinology and Tropical Medicine (CCVTM), University of Oxford

Addresses:

FISABIO, Vaccine Research Department, Avda. Cataluña, 21. 46020 València, Spain

London School of Hygiene & Tropical Medicine, Faculty of Epidemiology and Population Health, Keppel St, London WC1E 7HT, United Kingdom

Vrije Universiteit Amsterdam, Faculty of Science, De Boelelaan 1105, 1081 HV Amsterdam, Netherlands

University of Washington, Institute of Health Metrics and Evaluation, 2301 Fifth Ave., Suite 600, Seattle, WA 98121, USA

The Royal Institute of International Affairs, Chatham House, 10 St James's Square, London SW1Y 4LE, United Kingdom

Oxford Vaccine Group, Centre for Clinical Vaccinology and Tropical Medicine (CCVTM), Churchill Hospital, Oxford, OX3 7LE, United Kingdom

Corresponding author:

Per Egil Kummervold

per@capia.no

(+47) 414 35 795

Funding

PK has been funded from the European Commission for the call H2020-MSCA-IF-2017 and the funding scheme MSCA-IF-EF-ST for the project VACMA (Grant agreement ID: 797876).

HL, PP, SM, SD, ELK have been funded by a grant from GlaxoSmithKline Biologicals SA (GSK) to support research on maternal vaccination. The funder had no role in study design, data collection, analysis, interpretation, or writing of this study.

The research is supported with Cloud TPUs from Google's TPU Research Cloud (TRC).

Conflicts of Interest

HL's research group the Vaccine Confidence Project™ (HL, PP, SM, SD, ELK) has received funds from GlaxoSmithKline and Merck. HL has served on the Merck Vaccines Strategic Advisory Board. None of the other authors have conflicts of interest to declare.

HL and PP are affiliated to the National Institute for Health Research Health Protection Research Unit (NIHR HPRU) in Immunisation at LSHTM in partnership with Public Health England (PHE). The views expressed are those of the author(s) and not necessarily those of the NHS, the NIHR, the Department of Health or Public Health England.

Credit Author Statement

Per Egil Kummervold: Conceptualization, Methodology, Formal Analysis, Software Development, Funding Acquisition, Writing- original draft preparation, Writing- Reviewing and Editing.

Sam Martin Conceptualization, Methodology, Formal analysis, Data Annotation, Writing- Reviewing and Editing.

Sara Dada: Conceptualization, Methodology, Formal Analysis, Data Annotation, Writing- Reviewing and Editing, Project administration.

Eliz Kilich: Conceptualization, Methodology, Formal Analysis, Data Annotation, Funding Acquisition, Writing- Reviewing and Editing.

Chermain Denny: Data Annotation, Writing- Reviewing and Editing.

Pauline Paterson: Conceptualization, Funding Acquisition, Writing- Reviewing and Editing.

Heidi J Larson: Conceptualization, Methodology, Supervision, Funding Acquisition, Writing- Reviewing and Editing.

Keywords

Computer Science, Information Technology, Public Health, Health Humanities, Vaccines, Machine Learning

Abstract

Background

With growing conversations online and less than desired maternal vaccination uptake rates, these conversations could provide useful insight to inform future interventions. Automated processes for this type of analysis, such as natural language processing (NLP), have faced challenges extracting complex stances, like attitudes toward vaccines, from large text. In this study, we aimed to build upon recent advances in Transformer-based machine learning methods, and test if this could be used as a tool to assess the stance of social media posts towards vaccination during pregnancy.

Methods

A total of 16,604 Tweets posted between 1 November 2018 and 30 April 2019 were selected by boolean searches related to maternal vaccination. Tweets were coded by three individual researchers into the categories “Promotional”, “Discouraging”, “Ambiguous” and “Neutral” After creating a final dataset of 2,722 unique tweets, multiple machine learning methods were trained on the dataset and then tested and compared to the human annotators.

Results

We received an accuracy of 81.8% (F-score= 0.78) compared to the agreed score between the three annotators. For comparison, the accuracies of the individual annotators compared to the final score were 83.3%, 77.9% and 77.5%.

Conclusion

This study demonstrates the ability to achieve close to the same accuracy in categorising tweets using our machine learning models as could be expected by a single human annotator. The potential to use this reliable and accurate automated process could free up valuable time and resource constraints of conducting this analysis, in addition to inform potentially effective and necessary interventions.

Funding

GlaxoSmithKline. European Commission.

Background

While individuals have been found to share different thoughts, questions and concerns about vaccines via social media [1], studies of vaccine discourse on social media [2] indicate that concerns, and indeed the sharing of misinformation, are particularly amplified [3]. What is of concern is the amount of imprecise and inaccurate articles available with regards to vaccinations.

Multiple studies have been conducted to monitor vaccination discussions on social media [4, 5, 6] Addressing misunderstandings and inaccuracies as early as possible is vital to making sound vaccine policies. However, there is currently insufficient research on how to effectively categorise the nuances in perceptions of sentiment towards vaccines in the large volumes of vaccine data shared daily on social media. Being able to monitor and understand the spread and traction of such rumours in big data on a larger global level is key to mitigating this information.

While data retrieved from social and news media might not be representative for the entire population, it gives a snapshot of discussions and thoughts, and changes observed here are still thought to be of vital importance to understanding emerging issues of concern and the link between misinformation on news and social media and its effect on vaccination confidence and uptake. To detect such changes we need an in-depth understanding of the content of these messages. While qualitative methods might give this insight, the sheer volume of news and social media makes it difficult to apply on conversations from entire populations over time. Machine learning and natural language processing has the potential for handling huge amounts of information. However, the accuracy, especially when dealing with the complexity of language used to express opinions about vaccines, has prevented the method from being very effective.

Sentiment analysis in machine learning refers to automatically determining whether the author of a piece of text is positive, negative, or neutral towards the subject of the statement. This is slightly different from stance detection where the task is automatically determining the author's attitude towards a proposition or target [7]. While a sentiment analysis can look only at the tone in the particular statement, stance detections often refers to a target outside the particular statement.

The author of a tweet could express a positive attitude toward vaccination by expressing negativity toward people opposing vaccines (for instance so called "anti-vaxxers"). This double negation would then be interpreted as "positive" or "promotional". This could be referred to as the "authors sentiment toward vaccination" but since "sentiment" often is used for referring to the "sentiment of the statement", we find it less confusing to refer to this as the "author's stance" toward vaccination. This distinction is particularly important when studying more complex issues like vaccination since many texts often express strong opinions about vaccination without addressing vaccines directly. The distinction is illustrated in Table 1.

Table 1: Difference between sentiment and stance		
Text	Sentiment (subject)	Stance (target)
Vaccines saves lives	Positive (vaccines)	Positive (vaccines)

Anti-vaxxers kill people with their misinformation	Negative (anti-vaxxers)	Positive (vaccines)
Trust your doctor's knowledge regarding vaccines	Positive (doctors' knowledge)	Positive (vaccines)
Anti-vaxxers tell the real truth about vaccines	Positive (anti-vaxxers)	Negative (vaccines)

Historically, natural language processing has often concentrated on ordinary sentiment analysis. This is technically a much easier task, but unfortunately it is less useful from a sociological point of view. In contrast to 'sentiment', a person's 'stance' towards a target can be expressed by using both negative or positive language. People could for instance switch from opposing "abortion" to promoting "pro-life" without changing their basic stance. In a sociological analysis we would usually be more interested in the stance people have toward a topic or target, than the sentiment expressed in a particular statement.

The task chosen for this paper is "maternal vaccination". This in itself creates an additional layer of complexity from just looking at the stance towards vaccination. There are for instance several live vaccines that are not recommended by the health authorities during pregnancy, so it is possible to have a positive stance on maternal vaccination and at the same time express that some vaccines should be avoided. Another semantic complexity is that the maternal vaccine is for the mother, while the effect often is for the unborn fetus.

In this paper we look especially at how well such a complex task of detecting stance toward maternal vaccination can be solved by using multiple machine learning methods. We try to quantify how accurate such tweets can be categorised by trained annotators, and how this compares to modern machine learning methods.

Methods

This research collected 16,605 Twitter messages (tweets) published over 6 months between 1st November 2018 and 30th April 2019 from Meltwater [8], a media intelligence system. This dataset was collected and coded to complement a larger research study on sentiments and experiences around maternal vaccination across 15 countries (Australia, Brazil, Canada, France, Germany, India, Italy, Mexico, Panama, South Africa, South Korea, Spain, Taiwan, United Kingdom and United States). Non-English tweets were translated into English using a Google Translate script. Appendix 1 includes the search queries used. Before coding, all usernames and links were replaced by a common tag. This serves two purposes as it both preserves anonymity and limits potential bias from the annotator's interpretation of the username. The target for the analysis should be what the actual text is telling the reader about the writers' stance toward vaccination.

In this paper, "maternal vaccination" typically refers to the vaccines that are recommended by health authorities for pregnant women.

Tweets were individually manually coded into stance categories. Table 2 shows how stance was categorized across four sentiments towards maternal vaccines: "Promotional" (towards maternal vaccines), "Ambiguous" (uncertainty with mixed sentiment towards maternal vaccines), "Discouraging" (against maternal vaccines) and "No stance" (statements/facts about maternal vaccines that do not reveal the author's stance). Though it can be argued that some of the categories

can be ordered, we are treating them as nominal and not as ordinal variables in the analysis. A tweet saying pregnant women should take the tetanus vaccine, but not the measles vaccine, is therefore considered a promotional post towards maternal vaccines, since it encourages the current health recommendations.

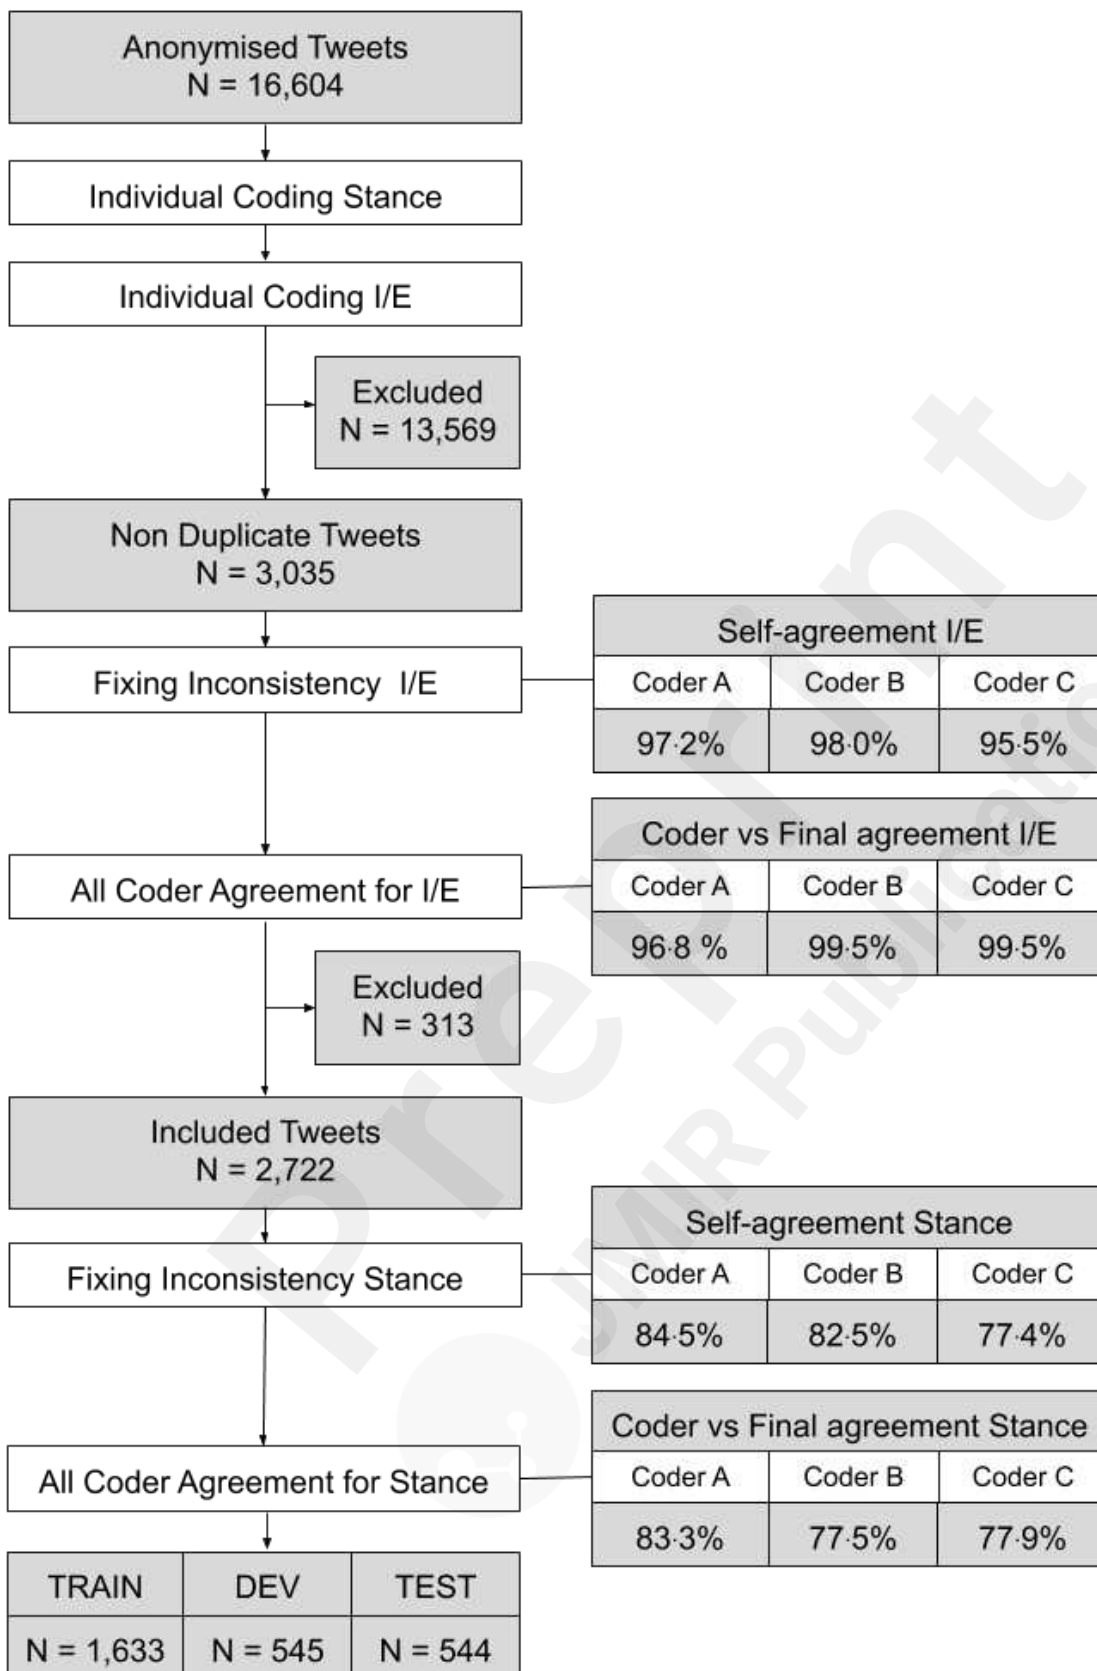
Table 2: Stance categories and definitions

Promotional	Ambiguous	Discouraging
<ul style="list-style-type: none"> ● Posts communicate public health benefits or safety of maternal vaccination. ● Contains positive tones, supportive or encouraging towards maternal vaccination. ● Describes risk of not vaccinating in pregnancy. ● Posts refute claims maternal vaccines are dangerous. 	<ul style="list-style-type: none"> ● Content contains indecision, uncertainty on the risks or benefits of maternal vaccination, or is ambiguous. ● Contains disapproving and approving information 	<ul style="list-style-type: none"> ● Contains negative attitude/arguments against maternal vaccines. ● Contains questions re. effectiveness/safety or possibility of adverse reactions (e.g. links to disability/autism). ● Discourages the use of recommended maternal vaccines
<p>Neutral / No stance</p> <ul style="list-style-type: none"> ● Contains no elements of uncertainty, promotional or negative content. These are often not sentiments online but rather statements, devoid of emotion. This includes factual posts pointing to articles about maternal vaccines e.g. "Study on effectiveness of maternal flu vaccine." 		

After the initial coding, the dataset was cleaned for duplicates and semi-duplicates. Semi-duplicates are tweets where a few characters differ but the meaning is unchanged. Typical examples are when a post is retweeted and is prefixed by 'RT:'. Another example is when a tweet is suffixed (by a user/bot) with some random characters to avoid being recognised (by Twitter detection algorithms) as a mass posting. For detecting semi-duplicates, we used a non-normalised Levenshtein distance of less than 30 for tweets above 130 characters. For shorter tweets the distance was scaled. The validity of the deduplication algorithm was evaluated qualitatively by the annotators. For our use we were aiming for a "greedy" algorithm that identified too many semi-duplicates rather than too few. While this could slightly affect the size of the training set, it was considered to be of greater importance to prevent tweets that looked too similar to be included both in the training and the test data set. We have open sourced the Python code we developed for cleaning and removing duplicates and made it available at our online GitHub repository [9].

In our study, deduplication was conducted after the first round of coding. The annotators were then asked to recode any tweet where they had given inconsistent coding. For example, there were instances where the same annotator coded identical tweets inconsistently. From the tweets that appeared only twice in the material we calculated a self-agreement score both for include/exclude and for stance. This was done to illustrate some of the potential challenges of manual coding (Figure 1).

Figure 1. Flow Chart of our screening and coding procedure



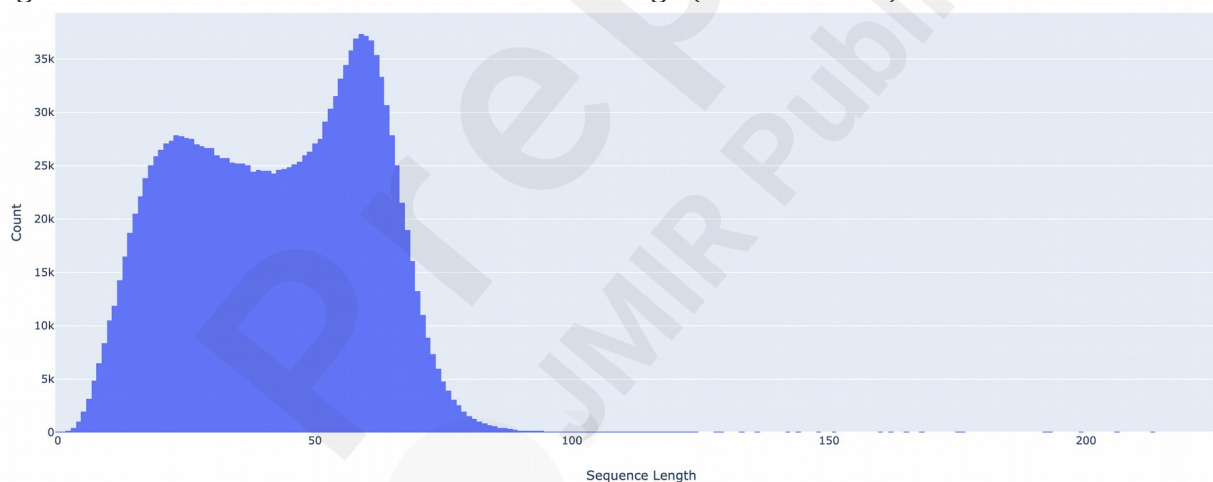
The main model was based on the newest (May 2019) Whole Word Masking variant of Google’s Bidirectional Encoder Representations from Transformers (BERT) [10]. When published in late 2018 the model demonstrated state-of-the-art results on 11 Natural Language Processing tasks, including

the Stanford Question Answering Dataset (SQuAD v1.1) [11]. BERT is a bidirectional, contextual encoder built upon a network architecture called Transformer, based solely on attention mechanisms [12]. The main part of the training can be done on unsupervised and unlabelled text corpuses like Wikipedia, and the pre-trained weights [13] are solely trained on this general corpus.

To expose the model to the vocabulary typical for vaccination, we trained it on a domain-specific corpus. We started creating domain-specific pre-training data by downloading 5.9 million tweets acquired by keyword searches related to vaccine and vaccination (Appendix 2). The set was downloaded from Meltwater and pre-processed the same way as the maternal vaccine tweets (i.e. de-duplication and username/link anonymisation). The Bert-architecture depends on doing unsupervised training by using a technique called next-sentence prediction (NSP). This requires each article/tweet to be at least two sentences long. We therefore filtered out all tweets that did not satisfy this criterion, reducing the dataset from 1.6 to 1.1 million tweets. We refer to this dataset as the vaccine-tweet-dataset.

We tokenized the tweets using the BERT vocabulary, and limited the sequence length to 96 tokens. By limiting the sequence length, we are able to increase the batch size which is known to have a positive effect on performance. Figure 2 shows sequence length of downloaded tweets, showing that this trimming would affect less than 0.03% of the tweets. The tweets longer than 96 tokens were manually examined, confirming that these were mainly repetitive sequences, and that the trimming did not affect the meaning (e.g. - a statement followed by strings of varying length of repeated characters, such as “.....” or “?????”).

Figure 2. Number of tokens in each Twitter message (N=1.6 million)



In addition, we acquired a dataset with a total of 201,133 vaccine related news articles from the Vaccine Confidence Project™ media archive. The articles were collected by automated keyword searches from several sources, among them Google News, Healthmap and Meltwater. It is an extensive collection of English vaccine related articles from both the news media and blogs. The search criteria have been developed over the years and so has varied slightly but is very similar to the list in Appendix 2. We refer to this dataset as the vaccine-news-dataset. We chose not to pre-train on a maternal vaccine specific dataset, since we wanted the encoder representations to also be used on other vaccine related topics. All pre-training was done using learning rate to $2e-5$, the batch size to 216 and the maximum sequence length to 96.

These domain specific pre-trained weights were the starting point for the classification of the

maternal vaccination tweets. The manually classified maternal vaccination tweets were pre-processed the same way as the tweets in the vaccine-tweet-dataset, and then divided into a training, development and test dataset in the ratio 60/20/20 (N=1633/545/544).

The pre-training of Transformer-models are very slow, but when these weights are determined, the last final fine-tuning step is fast. To our knowledge the best way of comparing the various pre-trained models is by comparing how they do after fine-tuning. Figure 3 shows that the fine-tuning does not improve after 15 epochs, but that there are considerable variance between each run. For this reason, all pre-trained models are evaluated by an average of 10 fine-tuned runs each at 15 epochs.

Figure 3. Fine tuning accuracy



To get a baseline score for comparative machine learning models, various traditional established networks were trained. The aim was to use well established networks with known performance against standardised datasets for sentiment and stance analysis. The benchmark architectures, the neural network and the long short-term memory networks (LSTM) with and without GloVe word embeddings, were all taken from Chollet's Deep Learning With Python [14].

To verify that the neural network was able to solve other neural network tasks, we tested the network structures on one of the most basic natural language processing tasks: predicting positive/negative sentiments from IMDB movie reviews [15].

The final domain specific pre-training and fine-tuning were done on a Cloud TPU v2-8 node with 8 cores and a total of 64GiB of memory and an estimated performance of 180 teraflops. Domain specific pre-training was done for two weeks but as shown in Figure 3 did not gain measurable improvements after a few days of running. Fine-tuning requires fewer computing resources and usually completes in a few minutes on this platform.

Results

In total, three annotators each coded 2,722 tweets. Of these, 1,559 (57.3%) tweets were coded identically, with a Fleiss agreement score of $\kappa = 0.56$. After meeting and discussing the tweets they disagreed on, the annotators were able to agree on the coding of all remaining tweets. Though the

annotators were able to agree on a final category for every tweet, they also reported that 186 (6.8%) of the tweets could be “open to interpretation”. Comparing the final agreed coding after discussions with the annotators’ original initial coding, the accuracy of the individual annotators were 83.3%, 77.9% and 77.5%.

One of the most basic neural networks for natural language processing consists of two fully connected layers. For our dataset this only gives an accuracy of 43.7%. The network is therefore not able to get a better result than simply predicting the overrepresented task “Promotional” for all data points. Adding pre-trained GloVe word-embeddings to this structure performs slightly better with a maximum accuracy on the test dataset of 55.5%.

To evaluate the reason for this low accuracy, we also tested the same network on the IMDb dataset setting the number of training examples to be the same (N=1,633). In this case the network got an accuracy of above 80% even without the GloVe embeddings, showing that the low accuracy is related to the difficulty of the maternal vaccine categorisation.

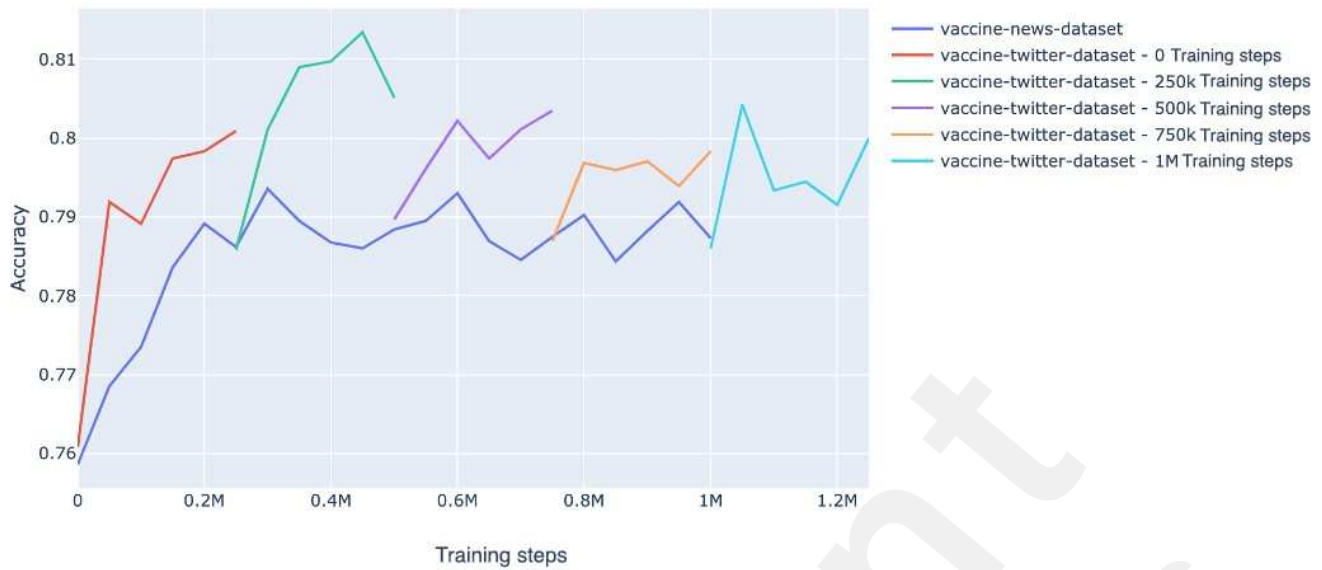
A more modern model is called long short-term memory (LSTM). This is a recurrent neural network (RNN) with a memory module. This model architecture was considered state-of-the-art a couple of years ago. We are able to get an accuracy of 63.1% here and can improve this to 65.5% by adding pre-trained GloVe-embeddings.

Our main research target was to investigate if state-of-the-art NLP models could be improved by using the BERT architecture. Using the original pre-trained weights, we are able to achieve an average accuracy of 76.7% when fine-tuning for 15 epochs.

Starting from the original weights, the model weights were pre-trained on the larger vaccine-news-dataset for 1 million training steps. At various checkpoints (0; 250,000; 500,000; 750,000 and 1,000,000), the model was forked and then trained on the smaller and more specific vaccine-tweet-dataset.

At each of the checkpoints, the network was fine-tuned on the manually labelled tweets for 15 epochs, and the average of 10 runs are reported in Figure 4. Using pre-training on domain specific content, the accuracy peaked around 79% when training only on the vaccine-news-dataset. However, by training first on the vaccine-news-dataset for 250,000 training steps, and then on the vaccine-news-dataset for an additional 200,000 training steps, we are able to get an accuracy of 81.8% (See Table 3).

Figure 4. Evaluation of domain specific pre-training using both the vaccine-news-dataset and the vaccine-twitter-dataset



Results are average from 10 fine-tunings for 15 Training steps evaluated on the development dataset.

Table 3 - Accuracy

	Accuracy	F-score
Annotator average	0.795861	0.739396
Annotator 1 - ELK	0.833211	0.796272
Annotator 2 - SCM	0.775165	0.710356
Annotator 3 - SD/CD	0.779206	0.711559
Neural network - no embeddings	0.436697	0.436697
Neural network - GloVe Word Embeddings	0.544954	0.457813
LSTM - No Embeddings	0.631193	0.549997
LSTM + GloVe Embeddings	0.655046	0.593942
Bert - default weights	0.766972	0.718878
Bert - domain specific	0.818349	0.775830

The final accuracy scores for the Bert-based models are based on selecting the best network from the results based on the results from the development dataset. The reported number are from evaluating the training dataset.

Discussion

The categories chosen in this study went through several revisions to ensure that they could clearly be understood. The annotators were fluent English speakers with a postgraduate degree, as well as several years of work experience within the field.

Some of the nuances contained in the tweets meant that it was difficult to categorise them as

definitively one stance. Thus, even when the same annotators are asked to code a nearly identical tweet at a later time, the annotators choose a different code one out of five times. After being given a second chance to code all duplicate tweets with inconsistencies, the annotators met and discussed the categories they disagreed on. The final accuracy was then on average 79.6%.

Ideally, the correct coding should be the coding that an average of a large number of experienced annotators would have chosen. Limiting the number of annotators to three, opens up for cases where all annotators by chance coded the same tweet identically and erroneously, and cases where none of the annotators chose the categories that a larger number of annotators would have chosen. It is therefore reasonable to assume that the accuracy of 79.6% might be slightly optimistic as what to expect from an average human annotator, even one that has long experience in the area.

The task of annotation is challenging since it is open for interpretation, a similar challenge which natural language processing also struggles with. Our tests show that a simple neural network that had no problem achieving above 80% on an IMDb movie review task, was unable to predict anything better than the most prevalent category when tested on maternal vaccination tweets.

Unsurprisingly, long-short term memory (LSTM) networks perform better than ordinary neural networks. Pre-trained embeddings help in all cases. Using GloVe embeddings increases accuracy. With the LSTM we are able to get an accuracy of 63.1% and are able to improve this to 65.5% by adding pre-trained GloVe-embeddings. However, they still lag behind what could be considered human accuracy.

Transformer-based architectures do perform significantly better. By using the pre-trained openly available BERT weights, we were able to get an accuracy of 76.7%. This is around the same level of accuracy as the lowest of our three annotators.

Domain specific pre-training also shows potential. While the pre-training does require some computing power, it does not require manual coding that could entail high time and resource costs. It is also worth noticing that we deliberately trained only on general vaccine terms. We did not make optimisations specifically for the domain of maternal vaccines. The main reason for this is that we wanted weights that were transferable to other tasks within the field of vaccines.

In our setting, the best results of 81.8% was achieved after initially training on news articles about vaccines and then training on vaccine-related tweets. This accuracy is above the average of the three annotators, even after the annotator has done multiple codings of the same tweet, and been given the opportunity to re-code any inconsistencies.

Limitations

We used a limited dataset, especially for the tweet dataset containing only 1M vaccine related tweets. It is also reasonable to assume that pre-training on a larger dataset of non-vaccine specific tweets could have a positive effect, since the language of tweets are quite different from other texts. Increasing the datasets are an easy way of potentially increasing the accuracy.

After Google released BERT late in 2018, there have been multiple general improvements by both Facebook, Microsoft and Google to the Transformer-based architecture BERT and how to improve the base models [16, 17, 18]. These have not been implemented in the current study. There is currently significant research activity in the field, and it is reasonable that implementing these

improvements to the base model and restarting the domain specific pre-training checkpoint would lead to higher accuracy in our categorisation.

Conclusion/Recommendations

Being able to categorise and understand the overall stance of social media conversations, especially in terms of identifying clusters of discouraging or ambiguous conversations - will make it easier to spot activities that may signal vaccine hesitancy or a decline in vaccine confidence, with greater speed and more accuracy. To manually, and continually, monitor this in our information society is near impossible. In that respect it has always been obvious that NLP has a huge potential since it can process an enormous amount of textual information.

However, so far NLP has only been able to solve very easy tasks, and unable to handle the nuances in language related to complicated issues (for example, attitudes toward vaccination). The new advances in transformer-based models indicates that this is about to become a useful tool in this area, and therefore opens up a new area for social research.

We have been able to demonstrate that with a training dataset of around 1600 tweets, we are able to get at least the accuracy that should be expected by a trained human annotator in categorising the stance of maternal vaccination discussions on social media. While there are benefits of increasing this accuracy even more, the main research challenge is in reducing the number of training samples. So far this has been an under-prioritised area of research, and an area where we should expect advances in the future. The real benefit from the technology will first be apparent when we are able to do this kind of categorisation with only a few initial examples.

References

- [1] Wilcox, C.R., et al., Influenza and pertussis vaccination in pregnancy: Portrayal in online media articles and perceptions of pregnant women and healthcare professionals. *Vaccine*, 2018. 36(50): p. 7625-7631.
- [2] - Kang, G. J., Ewing-Nelson, S. R., Mackey, L., Schlitt, J. T., Marathe, A., Abbas, K. M., & Swarup, S. (2017). Semantic network analysis of vaccine sentiment in online social media. *Vaccine*, 35(29), 3621-3638.
- [3] - Broniatowski, D. A., Jamison, A. M., Qi, S., AlKulaib, L., Chen, T., Benton, A., ... & Dredze, M. (2018). Weaponized health communication: Twitter bots and Russian trolls amplify the vaccine debate. *American journal of public health*, 108(10), 1378-1384.
- [4] Assessing Vaccination Sentiments with Online Social Media: Implications for Infectious Disease Dynamics and Control. Marcel Salathé, Shashank Khandelwal. *PLOS* 13 October 2011. <https://doi.org/10.1371/journal.pcbi.1002199>
- [5] Majumder N, Hazarika D, Gelbukh A, Cambria E, Poria S. Multimodal sentiment analysis using hierarchical fusion with context modeling. *Knowledge-Based Syst.* 2018;161:124-33.
- [6] Smith MC, Dredze M, Quinn SC, Broniatowski DA, editors. *Monitoring real-time spatial public health discussions in the context of vaccine hesitancy.* 2017.
- [7] Mohammad, Saif M. et al., Stance and Sentiment in Tweets. *ACM Transactions on Internet Technology*, 2016.

- 17(3). <https://arxiv.org/pdf/1605.01655.pdf>
- [8] Meltwater. <https://www.meltwater.com/>
- [9] VACMA GitHub repository. <https://github.com/peregilk/VACMA-PUBLIC>
- [10] BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. Devlin J., Chang M-W., Lee L, Toutanova K. (Submitted on 11 Oct 2018 (v1), last revised 24 May 2019 (v2)). <https://arxiv.org/abs/1810.04805>
- [11] Stanford Question Answering Dataset. <https://rajpurkar.github.io/SQuAD-explorer/>
- [12] Attention is all you need. Vaswani A. et al. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc. 2017. <https://arxiv.org/pdf/1605.01655.pdf>
- [13] Bert. Tensorflow code and pre-trained models for BERT. <https://github.com/google-research/bert>
- [14] *Deep Learning with Python*. 1st edition. 2017 Francois Chollet Manning Publications Co. Greenwich, CT, USA.
- [15] Analyzing Sentiment of Movie Review Data using Naive Bayes Neural Classifier. *International Journal of Emerging Trends & Technology in Computer Science (IJETTCS)*, Volume 3, Issue 4. July-August 2014.
- [16] Liu, Y. et al. RoBERTa: A Robustly Optimized BERT Pretraining Approach. Facebook AI. 2019. <https://arxiv.org/abs/1907.11692>
- [17] Liu, X. et al. Multi-Task Deep Neural Networks for Natural Language Understanding. 2019. Microsoft. <https://arxiv.org/abs/1901.11504>
- [18] Lan, Z et al. ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. 2019. Google. <https://arxiv.org/abs/1909.11942>

Appendix 1 - Maternal vaccination keyword search

ENGLISH

((("vaccin*" OR "immuniz*" OR "immunis*" OR "Tdap" OR ("vaccin*" NEAR/3 "pertussis") OR ("vaccin*" NEAR/3 "whooping cough") OR ("vaccin*" NEAR/3 "Tetanus") OR ("vaccin*" NEAR/3 "Influenza") OR ("vaccin*" NEAR/3 "flu") OR "flu shot*" OR "tetanus shot*" OR "whooping cough shot*" OR "pertussis shot*" OR ("vaccin*" NEAR/3 "Group B streptococcus") OR ("vaccin*" NEAR/3 "Respiratory Syncytial") OR ("vaccin*" NEAR/3 "GBS") OR ("vaccine*" NEAR/3 "RSV")) NEAR/8 (matern* OR pregna* OR antenatal) AND ("during pregna*" OR "while pregna*" OR "whilst pregna*" OR "when pregna*" OR "in pregna*" OR "are pregna*" OR "pregnant wom*")) OR "maternal immuniz*" OR "maternal Vaccin*" OR "maternal immunis*") NOT ("a vet" or veterinary OR dog* OR cat* OR horse* OR mouse* OR pig* OR cow* OR (financ* near/3 stock*) OR "immunoglobulin*" OR "LON:" OR "NYSE:" OR url:www.clinicaltrials.gov OR ((child* NEAR/1 vaccin*) AND (child* NEAR/1 vaccin*) AND (child* NEAR/1 vaccin*)))

PORTUGUESE (with Brazilian text included)

((("vacin*" OR "imuniz*" OR ("Tdap") OR ("vacin*" NEAR/3 "coqueluche") OR ("vacin*" NEAR/3 "tosse convulsa") OR ("vacin*" NEAR/3 "tétano") OR ("vacin*" NEAR/3 "influenza") OR ("vacin*" NEAR/3 "gripe") OR " injeção gripe *" OR " injeção tétano *" OR " injeção tosse convulsa *" OR " injeção coqueluche *" OR ("vacin*" NEAR/3 "estreptococo do grupo B") OR ("vacin*" NEAR/3 "sincicial respiratório") OR ("vacin*" NEAR/3 "SGB") OR ("vacin*" NEAR/3 "VSR")) NEAR/8 (matern* OR gravid*) AND ("durante gravidez*" OR "enquanto grávida*" OR "quando grávida*" OR "na gravidez*" OR "estão grávidas*" OR "mulher* grávida*")) OR "imuniz* matern*" OR "vacinação materna*" OR " imunização* matern*") NOT ("um veterinário" or veterinário OR cão* OR canídeo* OR gato* OR felino* OR cavalo* OR equino* OR rato* OR porco* OR suíno OR vaca* OR bovino OR (financ* near/3 gado*) OR "abastecimento" OR " imunoglobulina*" OR "LON:" OR "NYSE:" OR url:www.clinicaltrials.gov OR ((criança* NEAR/1 vacin*) AND (criança* NEAR/1 vacin*) AND (criança* NEAR/1 vacin*)))

GERMAN

((("vakzin*" OR "immunis*" OR "impf*" OR ("Tdap") OR ("impf*" NEAR/3 "Pertussis") OR ("impf*" NEAR/3 "Keuchhusten") OR ("impf*" NEAR/3 "Tetanus") OR ("impf*" NEAR/3 "Influenza") OR ("impf*" NEAR/3 "Grippe") OR "Grippe-Impfung*" OR "Grippeimpfung*" OR "Influenza-Impfung*" OR "Influenzaimpfung*" OR "Tetanus-Impfung*" OR "Tetanusimpfung*" OR "Keuchhusten-Impfung*" OR "Keuchhustenimpfung*" OR "Pertussis-Impfung*" OR "Pertussisimpfung*" OR ("impf*" NEAR/3 "Streptokokken Gruppe B") OR ("impf*" NEAR/3 "B-Streptokokken") OR ("impf*" NEAR/3 "Respiratorische Synzytial") OR ("impf*" NEAR/3 "GBS") OR ("impf*" NEAR/3 "RSV")) NEAR/8 (matern* OR schwanger*) AND ("während schwanger*" OR "solange schwanger*" OR "sobald schwanger*" OR "als schwanger*" OR "in der schwanger*" OR "sind schwanger*" OR "schwängere frau*")) OR "maternale Impfung*" OR "maternale Immunisierung*" OR "maternale Vakzination*" NOT ("ein Tierarzt" OR Tierarzt OR Hund* OR Katze* OR Pferd* OR Maus* OR Schwein* OR Kuh* OR Kühe OR (finanz* near/3 Aktien*) OR "Immunglobulin*" OR "LON:" OR "NYSE:" OR url:www.clinicaltrials.gov OR ((Kind* NEAR/1 impfung*) AND (Kind* NEAR/1 impfung*) AND (Kind* NEAR/1 impfung*)))

AFRIKAANS

((entstof* OR ent OR inenting OR immunisering OR immuniseer OR (entstof* NEAR/3 pertussis) OR (entstof* NEAR/3 kinkhoes) OR (entstof* NEAR/3 tetanus) OR (entstof* NEAR/3 griep) OR griepinspuiting OR "griep spuit*" OR "griep spuite" OR "tetanus-inenting*" OR (entstof* NEAR/3 Groep B streptokokke) OR (entstof* NEAR/3 Respiratoriese sincytiale virus) OR (entstof* NEAR/3 GBS) OR (entstof* NEAR/3 RSV)) NEAR/8 (moeder* OR swanger*) AND ("terwyl swanger*" OR "tans swanger*" OR "wanneer swanger*" OR in swangerskap OR "is swanger*" OR "swanger vrou*" OR "verwagte vrouens" OR "swanger vrouens" OR "swanger vroue")) OR (moederinenting* OR "moeder inenting*" OR "moeder immunisering*") NOT ("n veearts" OR veeartsenykundige OR hond* OR kat* OR perd* OR muis* OR vark* OR koei* OR (finansies* near/3 aandele*) OR immunoglobulien* OR url:www.clinicaltrials.gov OR ((kind* NEAR/1 entstof*) AND (kind* NEAR/1 entstof*) AND (kind* NEAR/1 entstof*)))

TRADITIONAL CHINESE

((("疫苗" OR "疫苗" OR "疫苗" OR ("疫苗") OR ("疫苗" NEAR/3 "疫苗") OR ("疫苗" NEAR/3 "疫苗") OR ("疫苗" NEAR/3 "疫苗") OR ("疫苗" NEAR/3 "疫苗") OR "疫苗" OR "疫苗" OR "疫苗" OR "疫苗" OR "疫苗" OR "疫苗" OR ("疫苗" NEAR/3 " B 疫苗") OR ("疫苗" NEAR/3 "疫苗") OR ("疫苗" NEAR/3 "GBS") OR ("疫苗" NEAR/3 "RSV")) NEAR/8 (疫苗 OR 疫苗 OR 疫苗) AND ("疫苗" OR "疫苗" OR "疫苗" OR "疫苗" OR "疫苗" OR "疫苗" OR "疫苗" OR "疫苗")) OR "疫苗" OR "疫苗" OR "疫苗") NOT ("疫苗" OR 疫苗 OR 疫苗 OR 疫苗 OR 疫苗 OR 疫苗 OR 疫苗 OR 疫苗 OR (疫苗 near/3 疫苗) OR "疫苗" OR "疫苗:" OR "疫苗:" OR url:www.clinicaltrials.gov OR ((疫苗 NEAR/1 疫苗) AND (疫苗 NEAR/1 疫苗) AND (疫苗 NEAR/1 疫苗)))

ITALIAN

("vaccin*" OR "immuniz*" OR ("vaccin*" NEAR/3 "Tdap") OR ("vaccin*" NEAR/3 "pertosse") OR ("vaccin*" NEAR/3 "tosse convulsiva") OR ("vaccin*" NEAR/3 "tosse asinina") OR ("vaccin*" NEAR/3 "tetano") OR ("vaccin*" NEAR/3 "influenza") OR ("vaccin*" NEAR/3 "influenza") OR "vaccinazione antinfluenzale" OR "vaccino antinfluenzale" OR "vaccinazione antitetanica" OR "vaccin* antitetanica" OR "vaccino anti-pertosse" OR "vaccin* anti-tosse convulsiva" OR "vaccin* antidifterite-tetano-pertosse" OR ("vaccin*" NEAR/3 "streptococco Gruppo B") OR ("vaccin*" NEAR/3 "virus respiratorio sinciziale") OR ("vaccin*" NEAR/3 "GBS") OR ("vaccin*" NEAR/3 "RSV")) NEAR/8 (matern* OR gravid*) AND ("durante la gravidanza" OR "mentre gravidanza" OR "in gravidanza" OR "incinta" OR "sono gravid*" OR "é gravid*" OR "donne gravide" OR "donna gravida" OR "donna in gravidanza" OR "immunizzazione materna" OR "immunità materna" OR "vaccinazione materna") NOT (vet* OR veterinario* OR can* OR gatt* OR cavall* OR top* OR maial* OR (finanz* NEAR/3 bestiame) OR immunoglobulin* OR ((bambin* NEAR/1 vaccin*) AND (bambin* NEAR/1 vaccin*) AND (bambin* NEAR/1 vaccin*)))

SPANISH - SPAIN

((("vacuna*" OR "inmuniz*" OR ("TDaP") OR ("vacuna" NEAR/3 "DPT") OR ("vacuna" NEAR/3 "difteria-tétanos-pertussis") OR ("vacuna" NEAR/3 "DTPa") OR ("vacuna" NEAR/3 "DT") OR ("vacuna" NEAR/3 "triple bacteriana") OR ("vacuna" NEAR/3 "pertussis") OR ("vacuna" NEAR/3 "tos convulsiva") OR ("vacuna" NEAR/3 "tos ferina") OR ("vacuna" NEAR/3 "coqueluche") OR ("vacuna" NEAR/3 "tétanos") OR ("vacuna" NEAR/3 "tétano") OR ("vacuna" NEAR/3 "difteria") OR ("vacuna" NEAR/3 "Influenza") OR ("vacuna" NEAR/3 "gripe") OR "vacuna antigripal" OR "vacuna antiinfluenza" OR "vacuna antitetánica" OR "toxoides tétanico" OR "vacuna antipertussis" OR "vacuna antipertusis" OR "vacuna antipertúsica" OR "vacuna anticoqueluchosa" OR "vacuna

antidiftérica" OR "toxoides diftérico" OR ("vacuna" NEAR/3 "streptococcus del grupo B") OR ("vacuna" NEAR/3 "estreptococos del grupo B") OR ("vacuna" NEAR/3 "respiratorio sincicial") OR ("vacuna" NEAR/3 "sincicial respiratorio") OR ("vacuna" NEAR/3 "GBS") OR ("vacuna" NEAR/3 "EGB") OR ("vacuna" NEAR/3 "SGB") OR ("vacuna" NEAR/3 "RSV") OR ("vacuna" NEAR/3 "VRS"))

NEAR/8 (matern* OR embaraz* OR gestación OR preñez OR madres) AND ("durante el embarazo" OR "mientras embarazadas" OR "cuando están embarazadas" OR "durante la preñez" OR "durante la gestación" OR "en el embarazo" OR "en gestantes" OR "estas embarazada" OR "muje* embaraz*" OR "muje* gestantes" OR "la gestante") OR "inmunización matern*" OR "vacuna* matern*") NOT (veterinari* OR perr* OR canin* OR gat* OR felin* OR equin* OR caballo* OR yegu* OR rat* OR roedo* OR cerd* OR porcin* OR marrana* OR vacas OR bovin* OR vacuno OR aves OR aviar OR (financia* NEAR/3 valores) OR "abastecimiento" OR "immunoglobulin*" OR "LON:" OR "NYSE:" OR url:www.clinicaltrials.gov OR ((niñ* NEAR/1 vacuna*) AND (niñ* NEAR/1 vacuna*)) AND (niñ* NEAR/1 vacuna*))

SPANISH – MEXICO

((("vacuna*" OR "inmuniz*" OR ("TDaP") OR ("vacuna" NEAR/3 "DPT") OR ("vacuna" NEAR/3 "difteria-tétanos-pertussis") OR ("vacuna" NEAR/3 "difteria-tétanos-tos ferina") OR ("vacuna" NEAR/3 "DTPa") OR ("vacuna" NEAR/3 "DT") OR ("vacuna" NEAR/3 "triple bacteriana") OR ("vacuna" NEAR/3 "pertussis") OR ("vacuna" NEAR/3 "tos convuls*") OR ("vacuna" NEAR/3 "tos ferina") OR ("vacuna" NEAR/3 "coqueluche") OR ("vacuna" NEAR/3 "tétanos") OR ("vacuna" NEAR/3 "tétano") OR ("vacuna" NEAR/3 "difteria") OR ("vacuna" NEAR/3 "influenza") OR ("vacuna" NEAR/3 "gripe") OR ("vacuna" NEAR/3 "gripa") OR "vacuna antigripal" OR "vacuna antiinfluenza" OR "vacuna antitetánica" OR "toxoides tetánico" OR "vacuna antipertussis" OR "vacuna antipertusis" OR "vacuna antipertúsica" OR "vacuna anticoqueluchosa" OR "vacuna antidiftérica" OR "toxoides diftérico" OR ("vacuna" NEAR/3 "streptococcus del grupo B") OR ("vacuna" NEAR/3 "estreptococos del grupo B") OR ("vacuna" NEAR/3 "respiratorio sincicial") OR ("vacuna" NEAR/3 "sincicial respiratorio") OR ("vacuna" NEAR/3 "GBS") OR ("vacuna" NEAR/3 "EGB") OR ("vacuna" NEAR/3 "SGB") OR ("vacuna" NEAR/3 "RSV") OR ("vacuna" NEAR/3 "VRS")) NEAR/8 (matern* OR embaraz* OR gestación OR gravidez OR madres) AND ("durante el embarazo" OR "mientras embarazadas" OR "cuando están embarazadas" OR "durante la gravidez" OR "durante la gestación" OR "en el embarazo" OR "en gestantes" OR "estas embarazada" OR "muje* embaraz*" OR "muje* gestantes" OR "la gestante" OR "durante la gravidez") OR "inmunización matern*" OR "vacuna* matern*") NOT (veterinari* OR perr* OR canin* OR gat* OR felin* OR equin* OR caballo* OR yegu* OR rat* OR roedo* OR cerd* OR porcin* OR marrana* OR vacas OR bovin* OR vacuno OR aves OR pájaros OR aviar OR (financia* NEAR/3 valores) OR "abastecimiento" OR "immunoglobulin*" OR "LON:" OR "NYSE:" OR url:www.clinicaltrials.gov OR ((niñ* NEAR/1 vacuna*) AND (niñ* NEAR/1 vacuna*)) AND (niñ* NEAR/1 vacuna*))

SPANISH – PANAMA

((("vacuna*" OR "inmuniz*" OR "TDaP" OR "DPT" OR "DTP" OR ("vacuna contra difteria tétanos y tosferina") OR ("vacuna" NEAR/3 "pertussis") OR ("tosferina" OR "coqueluche" OR "tos convuls*")) OR ("vacuna" NEAR/3 "tosferina" OR ("tos convuls*" OR "coqueluche")) OR ("vacuna" NEAR/3 "tétan*") OR ("vacuna" NEAR/3 "influenza") OR ("resfriado*" OR "catarro*") OR ("vacuna" NEAR/3 "gripe") OR ("resfriad*" OR "catarro*" OR "influenza")) OR "vacuna antigripal" OR "vacuna antiinfluenza" OR "vacuna antitetánica" OR ("toxoides antitetánico") OR "vacuna contra la tosferina" OR ("vacuna contra pertussis ") OR ("vacuna" NEAR/3 "estreptococos del grupo B" OR "EGB") OR ("vacuna" NEAR/3 "respiratorio sincicial") OR ("sincicial respiratorio") OR ("vacuna" NEAR/3 "SGB") OR ("vacuna" NEAR/3 "VRS")) NEAR/8 (matern* OR embaraz* OR gestación OR preñez OR prenatal) AND ("durante el embarazo" OR "mientras el embarazo" OR "cuando están embarazadas" OR "durante la preñez" OR "durante la gestación" OR "en el embarazo" OR "en gestantes" OR "estas embarazada" OR "muje* embaraz*" OR "muje* gestantes" OR "la gestante") OR ("inmunización matern*" OR "vacuna* matern*") NOT (veterinari* OR perr* OR canin* OR gat* OR felin* OR caball* OR equin* OR rat* OR roedo* OR cerd* OR porcin* OR marran* OR vaca* OR bovin* OR vacuno OR (financ* OR finanz* NEAR/3 valores) OR abastecimiento OR surtid* OR inventario* OR suministro* OR existencia* OR immunoglobulin* OR ((niñ* OR chic* OR pequeñ* OR crí* OR infant* OR menor*) NEAR/1 vacuna*) AND ((niñ* OR chic* OR pequeñ* OR crí* OR infant* OR menor*) NEAR/1 vacuna*) AND ((niñ* OR chic* OR pequeñ* OR crí* OR infant* OR menor*) NEAR/1 vacuna*))

FRENCH

((("vaccin*" OR "immunis*" OR "dTP" OR ("vaccin*" NEAR/3 "coqueluche") OR ("vaccin*" NEAR/3 "dTPca") OR ("vaccin*" NEAR/3 "tétanos") OR ("vaccin*" NEAR/3 "grippe") OR ("vaccin*" NEAR/3 "antigrippal") OR

"vaccin contre la grippe*" OR "vaccin contre le tétanos*" OR "vaccin contre la coqueluche*" OR ("vaccin*" NEAR/3 "streptocoque du Groupe B") OR ("vaccin*" NEAR/3 "virus respiratoire syncytial") OR ("vaccin*" NEAR/3 "SGB") OR ("vaccin*" NEAR/3 "VRS")) NEAR/8 (enceinte* OR grossesse*) AND ("pendant la grossesse" OR "lors de la grossesse" OR "chez les femmes enceintes" OR "chez la femme enceinte")) OR "vaccination maternelle") NOT (vétérinaire OR chien* OR chat* OR cheva* OR souris* OR cochon* OR vache* OR (financ* near/3 stock*) OR "immunoglobuline*" OR "LON:" OR "NYSE:" OR url:www.clinicaltrials.gov OR ((enfant* NEAR/3 vaccin*) AND (enfant* NEAR/3 vaccin*) AND (enfant* NEAR/3 vaccin*)))

FRENCH - CANADA

((("vaccin*" OR "immunis*" OR ("dTTP") OR ("vaccin*" NEAR/3 "coqueluche") OR ("vaccin*" NEAR/3 "dTTPca") OR ("vaccin*" NEAR/3 "Tétanos") OR ("vaccin*" NEAR/3 "grippe") OR ("vaccin*" NEAR/3 "antigrippal") OR ("vaccin*" NEAR/3 "streptocoque B") OR ("vaccin*" NEAR/3 "SGB") OR ("vaccin*" NEAR/3 "virus respiratoire syncytial") OR ("vaccin*" NEAR/3 " streptocoque du Groupe B") OR ("vaccin*" NEAR/3 "VRS")) NEAR/8 (enceinte* OR grossesse*) AND ("pendant la grossesse" OR "lors de la grossesse" OR "chez les femmes enceintes" OR "chez la femme enceinte")) OR "vaccination maternelle") NOT (vétérinaire OR chien* OR chat* OR cheva* OR souris* OR cochon* OR vache* OR (financ* near/3 stock*) OR "immunoglobuline*" OR "LON:" OR "NYSE:" OR url:www.clinicaltrials.gov OR ((enfant* NEAR/3 vaccin*) AND (enfant* NEAR/3 vaccin*) AND (enfant* NEAR/3 vaccin*)))

ZULU

((("ukugoma" OR "Umgomo" OR " Ukugonywa" OR " Ukugoma" OR ("Tdap") OR TT OR Td OR ("ukugoma*" NEAR/3 "pertussis") OR ("ukugoma*" NEAR/3 " Ukukhwehlela ") OR ("ukugoma*" NEAR/3 "Tétanus") OR ("ukugoma*" NEAR/3 " Umkhuhlane") OR ("ukugoma*" NEAR/3 " Umkhuhlane") OR "umjovo womkhuhlane" OR "tetanus womkhuhlane*" OR "whooping cough womkhuhlane*" OR "pertussis womkhuhlane*" OR ("ukugoma*" NEAR/3 "Group B streptococcus") OR ("ukugom*" NEAR/3 "Respiratory Syncytial") OR ("ukugoma*" NEAR/3 "GBS") OR ("ukugomo*" NEAR/3 "RSV")) NEAR/8 (matern* OR pregna*) AND ("during pregna*" OR "ngenkathi ekhulelwe" OR " uma ekhulelwe" OR " uma ukhulelwe" OR "ekukhulelweni" OR "are pregna*" OR " owesifazane okhulelwe")) OR " ukugoma komama" OR "ukugoma komama" OR " ukugoma komama*")

NOT ("udokotela wezilwane" OR "udokotela wezilwane" ORinja OR ikati OR Ihashi OR igundane OR ingulube OR inkomo OR (Izimali near/3 isitoko) OR " izakhamzimba*" OR "LON:" OR "NYSE:" OR url:www.clinicaltrials.gov OR ((umntwana OR ebunganeni OR izingane) NEAR/1 Umgo) AND ((umntwana OR ebunganeni OR izingane) NEAR/1 Umgo) AND ((umntwana OR ebunganeni OR izingane) NEAR/1 Umgo))

KOREAN

((("□□*" OR "□□*" OR "□□*" OR "Tdap" OR ("□□*" NEAR/3 "□□□") OR ("□□*" NEAR/3 "□□□") OR ("□□*" NEAR/3 "□□□") OR ("□□*" NEAR/3 "□□") OR ("□□*" NEAR/3 "□□") OR "□□ □□*" OR "□□□ □□*" OR "□□□ □□*" OR "□□□ □□*" OR ("□□*" NEAR/3 ("B □□ □□□□" OR "□□ B □□□□□")) OR ("□□*" NEAR/3 "□□ □□□□") OR ("□□*" NEAR/3 " B □□ □□□□") OR ("□□*" NEAR/3 "□□ □□□□ □□□□")) NEAR/8 (□□* OR □□* OR □□) AND ("□□*" OR "□□ □□*" OR "□□□□*" OR "□□□*" OR "□□□□*" OR "□□□□*" OR "□□□□□*" OR "□□□□□□*" OR "□□□□□□□□*" OR "□□□□□□□□*" OR "□□□□□□□□□*" OR "□□□□□□□□□□*" OR "□□□□□□□□□□□□*") NOT ("□□□□" □□ OR □□ OR □* OR □□□* OR □* OR □* OR □□* OR □* OR ((□□* OR □□* OR □□*) near/3 □□*) OR "□□ □□□□*" OR "LON:" OR "NYSE:" OR url:www.clinicaltrials.gov OR ((□□□□* NEAR/1 □□*) AND (□□* NEAR/1 □□*) AND (□□* NEAR/1 □□*)))

HINDI

("वैक्सीन" OR "टीका*" OR "टीका लगाना" OR "वैक्सीन" "प्रतिरक्षण" OR "रोगक्षम करना" OR "Tdap" OR "TT" OR "Td" OR "DT" OR ("वैक्सीन" NEAR/3 "पर्दुसिस") OR ("वैक्सीन" NEAR/3 "इंफ्लुएंजा") OR ("वैक्सीन" NEAR/3 "फ्लू") OR ("वैक्सीन" NEAR/3 "टिटनेस") OR "टी टी" OR "फ्लू का टीका*" OR "टीटी*" OR "काली खांसी का टीका" OR "पर्दुसिस का टीका*" OR ("टीका*" NEAR/3 "ग्रुप B स्ट्रेप्टोकोक्कस") OR ("टीका*" NEAR/3 "श्वसन संबंधी") OR ("टीका*" NEAR/3 "गिओन बार") OR ("टीका*" NEAR/3 "आरएसवी")) NEAR/8 ("गर्भवती" OR "गर्भिणी" OR "गर्भवती होने" OR "जबकि गर्भवती" OR "गर्भावस्था में" OR "गर्भवती हैं" OR "गर्भवती महिला" OR "गर्भवती महिलाओं" OR "मातृ टीकाकरण") NOT ("पशु चिकित्सक" OR "पशुचिकित्सा" OR कुत्ता OR बिल्ली OR

Appendix 2 - Vaccination keyword search terms

"vaccin*" OR "vax" OR "vaxer" OR "vaxers" OR "vaxx" OR "vaxxer" OR "vaxxers" OR "immuniz*" OR "immunis*" OR ("shot" OR "jab") NEAR/4 ("influenza" OR "influenta") OR ("shot" OR "jab") NEAR/4 "flu") OR ("shot" OR "jab") NEAR/4 "virus") OR ("shot" OR "jab") NEAR/4 "whooping") OR ("shot" OR "jab") NEAR/4 "pertussis") OR ("shot" OR "jab") NEAR/4 "strepto*" OR ("shot" OR "jab") NEAR/4 "respiratory") OR ("shot" OR "jab") NEAR/4 "syncytial") OR ("shot" OR "jab") NEAR/4 "rzv") OR ("shot" OR "jab") NEAR/4 "rsv") OR ("shot" OR "jab") NEAR/4 "h1n1") OR ("shot" OR "jab") NEAR/4 "zvl") OR ("shot" OR "jab") NEAR/4 "pcv13") OR ("shot" OR "jab") NEAR/4 "ppsv23") OR ("shot" OR "jab") NEAR/4 "menb") OR ("shot" OR "jab") NEAR/4 "varicella") OR ("shot" OR "jab") NEAR/4 "syncytial") OR ("shot" OR "jab") NEAR/4 "tdap") OR ("shot" OR "jab") NEAR/4 "dtap") OR ("shot" OR "jab") NEAR/4 "pertussis") OR ("shot" OR "jab") NEAR/4 "tetanus") OR ("shot" OR "jab") NEAR/4 "measles") OR ("shot" OR "jab") NEAR/4 "cholera") OR ("shot" OR "jab") NEAR/4 "rota*" OR ("shot" OR "jab") NEAR/4 "ebola") OR ("shot" OR "jab") NEAR/4 "hepat*" OR ("shot" OR "jab") NEAR/4 "hepa") OR ("shot" OR "jab") NEAR/4 "hepb") OR ("shot" OR "jab") NEAR/4 "hib") OR ("shot" OR "jab") NEAR/4 "heamophilous") OR ("shot" OR "jab") NEAR/4 "hvp") OR ("shot" OR "jab") NEAR/4 "papilloma*" OR ("shot" OR "jab") NEAR/4 "cervical") OR ("shot" OR "jab") NEAR/4 "encephalitis") OR ("shot" OR "jab") NEAR/4 "cholera") OR ("shot" OR "jab") NEAR/4 "malaria") OR ("shot" OR "jab") NEAR/4 "meningitis") OR ("shot" OR "jab") NEAR/4 "polio*" OR ("shot" OR "jab") NEAR/4 "pnemonia*" OR ("shot" OR "jab") NEAR/4 "pneumonia*" OR ("shot" OR "jab") NEAR/4 "mmr") OR ("shot" OR "jab") NEAR/4 "varicella") OR ("shot" OR "jab") NEAR/4 "var") OR ("shot" OR "jab") NEAR/4 "chickenpox") OR ("shot" OR "jab") NEAR/4 "chicken pox") OR ("shot" OR "jab") NEAR/4 "zoster") OR ("shot" OR "jab") NEAR/4 "yellow fever") OR ("shot" OR "jab") NEAR/4 "zika") OR ("shot" OR "jab") NEAR/4 "gbs") OR ("shot" OR "jab") NEAR/4 "rubella*")