

# Dense Feature Memory Augmented Transformers for COVID-19 Vaccination Search Classification

**Jai Gupta**

Google Research  
jaigupta@google.com

**Yi Tay**

Google Research  
yitay@google.com

**Chaitanya Kamath**

Google Research  
ckamath@google.com

**Vinh Q. Tran**

Google Research  
vqtran@google.com

**Donald Metzler**

Google Research  
metzler@google.com

**Shailesh Bavadekar**

Google Research  
shaileshb@google.com

**Mimi Sun**

Google Research  
mimisun@google.com

**Evgeniy Gabrilovich**

Google Research  
gabr@acm.org

## Abstract

With the devastating outbreak of COVID-19, vaccines are one of the crucial lines of defense against mass infection in this global pandemic. Given the protection they provide, vaccines are becoming mandatory in certain social and professional settings. This paper presents a classification model for detecting COVID-19 vaccination related search queries, a machine learning model that is used to generate search insights for COVID-19 vaccinations. The proposed method combines and leverages advancements from modern state-of-the-art (SOTA) natural language understanding (NLU) techniques such as pretrained Transformers with traditional dense features. We propose a novel approach of considering dense features as memory tokens that the model can attend to. We show that this new modeling approach enables a significant improvement to the Vaccine Search Insights (VSI) task, improving a strong well-established gradient-boosting baseline by relative +15% improvement in F1 score and +14% in precision.

## 1 Introduction

Though COVID-19 continues to be a challenge worldwide, vaccines have provided the much needed hope. Countries and governments have significantly ramped up their efforts to improve the reach of the vaccines including booster shots. As such, it is important to understand how users search for vaccine related information such as vaccine efficacy, safety, and regional availability as this information can be very useful to inform policy decision making, create effective public service announcements, implement more efficient distribution of vaccines, etc. To this end, we released a public tool for COVID-19 Vaccination Search Insights, an interactive report that provides insights on user searches for COVID-19 vaccinations. For an example please see Figure 1. On the backend, this tool performs privacy preserving classification of

user search queries and clusters them based on the region they were issued from to present a timeline of how these searches have changed over time. A core task in the VSI tool is the problem of classifying search query intent, and one of them is whether users are seeking information on *vaccine access*, i.e. queries related to the eligibility, availability, and accessibility of COVID-19 vaccines.

The task is a challenging one since simply matching for COVID related terms is insufficient, as queries such as *fully vaccinated travel* or *proof of covid vaccination* are negative classes. Hence, this problem is nuanced and may benefit from a coalition of advanced language understanding systems and traditional search-related feature engineering methods. As such, the problem at hand crosses two main modalities, i.e., text and traditional dense features. In this problem, the text features are user search queries that are short and/or lack context. The dense features (discussed in more detail in section 5) are hand crafted features from recent and past activities, previous clicks, and named entities that play a critical role in understanding the user’s intent. However, dense features alone fail to capture important contextual language cues, such as those that state-of-the-art natural language understanding systems (Devlin et al., 2018; Raffel et al., 2019) have been shown to handle well. We find that both modalities are highly complementary and it is difficult to achieve strong performance using only a single modality.

**Our Contributions** The overall contributions of this paper can be summarized as follows:

- We propose a new model and framework for search query intent classification for our COVID-19 Vaccination Search Insight tool.
- We propose a paradigm of exploiting the benefits of text inputs through state-of-the-art NLU models, along with traditional dense features found in large-scale systems.

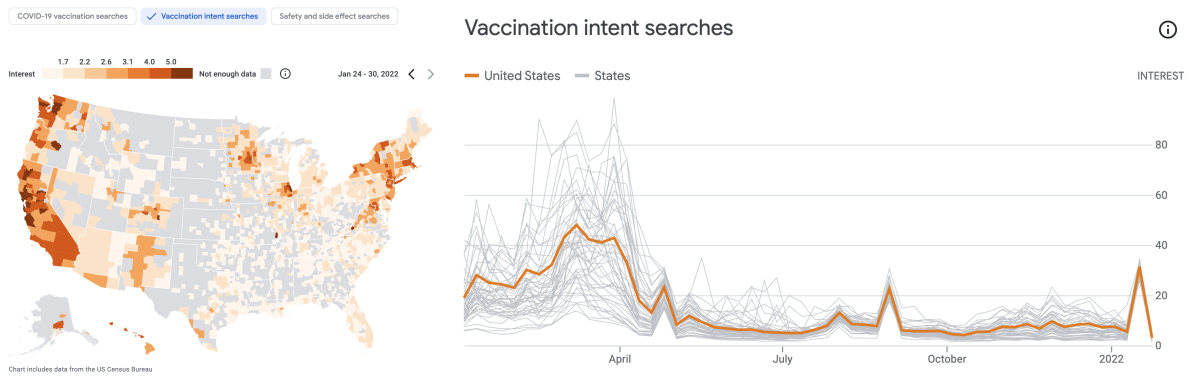


Figure 1: VSI tool presenting vaccination intent search query statistics segmented by location and time respectively.

- We propose a novel method of fusing dense features with Transformers that enables queries to retrieve from a dense memory store, in similar spirit to a contextual key-value store. Memory tokens here are used in similar manner as memory-based methods described in (Tay et al., 2020) and models such as Set Transformer (Lee et al., 2019), Memory Transformers (Sukhbaatar et al., 2019; Burtsev et al., 2020; Wu et al., 2020), and Global Memory tokens in ETC (Ravula et al., 2020) and BigBird (Zaheer et al., 2020). Notably, this is the first proposal for constructing global memory tokens using dense features.
- We conduct extensive experiments on real production data from three geographical regions. Our experiments show that the proposed method significantly outperforms a strong gradient boosting baseline by +15% and outperforms a SOTA Transformer by +5% on F1 score achieving very strong F1 score of over 98% on the US dataset with similar strong performance for other regions.

## 2 Related Works

This section presents related works and the background for this paper.

*Classification with Feature-based ML.* Building machine learned (ML) systems that operate across dense hand-crafted features is a well-established method. It is common to consider this class of ML methods as *tabular* machine learning whereby tree-based methods (Chen and Guestrin, 2016; Wikipedia, 2022) are dominant<sup>1</sup>. Within the context of text classifiers, feature engineering typically

<sup>1</sup><https://www.kaggle.com/shivamb/data-science-trends-on-kaggle>

leverages stemming, lemmatization, part-of-speech tags, tf-idf vectors, entities, salient terms, and other features that are relevant to the task at hand. It is also popular to use semantic representations from Glove (Pennington et al., 2014) or BERT (Devlin et al., 2018) as input dense features to a model.

*NLU with Pretrained Transformers.* Transformers (Ashish Vaswani, 2017), characterized by interleaved self-attention and MLP blocks, have become the dominant sequence model for language processing and understanding (NLU) (Devlin et al., 2018; Raffel et al., 2019; Brown et al., 2020). The key idea behind self-attention is to perform token-to-token alignment where the joint interaction of queries and keys retrieve from a memory store (value). To this end, it is also common for advanced Transformer architectures to leverage global memory tokens (Zaheer et al., 2020; Lee et al., 2019; Jaegle et al., 2021) that act as a parameter store for the query to attend to (Tay et al., 2020). A cornerstone of these systems is the pretraining task that learns general purpose language representations, which have been shown to be extremely beneficial (Tay et al., 2021) due to the gains from transfer learning (Pan and Yang, 2010).

*Joint Learning of Textual and Dense Features.* Performing feature extraction to convert text into dense features (e.g. TF-IDF (Salton and Buckley, 1988), word2vec (Mikolov et al., 2013), etc.) for the purpose of learning classifiers jointly with other non-textual (numerical) dense features has been common practice in machine learning for some time (Kowsari et al., 2019; Macskassy et al., 1999; Richardson et al., 2007). Previous research in the multi-modal domain adopt a strategy of *early fusion* or *late fusion*: joining the two modalities (in this case textual and numerical) either in feature space early in the architecture or in seman-

tic/decision space late in the architecture, respectively (Snoek et al., 2005). Perhaps most related to our approach is a recent work on joint representation of text and tabular data (Yin et al., 2020; Zhu et al., 2021) that pass a flattened representation of a table alongside text during encoding. Our work instead provides a method for jointly encoding text with dense features of a more generic, unstructured form.

*COVID-19 Vaccine Search Insights.* An important distinction of Query classification from generic text classification is that the former are significantly shorter, and may be underspecified (He et al., 2000; Beitzel et al., 2005). Due to this nature, some previous works in this area have augmented queries with additional context to improve query classification performance (Broder et al., 2007; Shen et al., 2006; Li et al., 2008; Jiang ).

Analyzing user queries and social interactions in health settings has been well studied. Sadilek et al. (2020) presents an analysis of user search queries for Lyme disease forecasting. Similarly, Sadilek et al. (2018) uses a machine-learned model for real-time detection of foodborne illness using web search and location data.

### 3 Problem Description

Given a search query  $q$ , the objective is to classify whether the query was issued with the intent of seeking information related to *vaccine access*. Notably,  $q$  is short in length and may or may not contain all the information needed to make the correct prediction.

Additionally, each  $q$  is supplemented with numerical features in the form of a dense feature vector  $X_f \in \mathbb{R}^{d_{features}}$ , where  $d_{features}$  can be any number of features. Details about how  $X_f$  is constructed for our setup is present in section 5, but in short  $X_f$  represents the topicality scores of phrases related to the query with  $d_{features}$  being 60k. However, it is important to note that all methods described in this paper are agnostic to the source of  $X_f$  and can be applied to any vector of numerical features.

## 4 VSI Transformer

This section describes the proposed method.

### 4.1 Pretrained Transformer Encoder

The main backbone of the proposed architecture is a Transformer (Ashish Vaswani, 2017) encoder.

We leverage the state-of-the-art T5 (Raffel et al., 2019) model as a starting point. Since T5 is a seq2seq model, we only utilize the T5 encoder as the Transformer model and discard the decoder of the pretrained model for our classification task. This is done by pooling the output of the encoder stack followed by a dense classification layer.

### 4.2 Input formulation

Given a query  $q$ , the input to the model is a discrete integer sequence representing the tokens (Google, 2021) of  $q$ , i.e.,  $X_\ell$  where  $\ell$  is the number of tokens in the query from the subword vocabulary  $V$ . The input sequence is selected from an embedding matrix of  $\mathbb{R}^{|V| \times d_{model}}$  to form a tensor of  $\mathbb{R}^{\ell \times d_{model}}$ .

### 4.3 Dense Feature Memory

For each input-target example, the input to the Dense Feature Memory module is a dense feature  $X_f \in \mathbb{R}^{d_{features}}$ . Given this dense feature of dimensions  $d_{features}$ , we transform it into memory tokens of dimension  $d_{model}$  via:

$$M_i = ReLU(W_i X_f + b_i)$$

where  $W_i \in \mathbb{R}^{d_{features} \times d_{model}}$  and  $M_i$  is the  $i$ -th memory token. We consider the number of memory tokens to be a hyperparameter. To this end, we then concatenate  $[M_1; \dots; M_{N_{memory}}]$  to the input query sequence  $X \in \mathbb{R}^{\ell \times d_{model}}$ . Along with the input query, we also pass in dense features corresponding to the query to the main body of the network. We note that memory tokens participate in the rest of the computation in a similar spirit to query tokens, i.e., they go through the same MLP and self-attention layers. The dense features are of  $d_{features}$  dimensions and are passed into the dense memory module. The dense features used in our setup is explained in section 5.

### 4.4 Attention Blocks: Querying & Retrieving from Dense Feature Memory Tokens

The dense feature memory token is appended to the input sequence and participates in the self-attention mechanism of the Transformer model. Concretely, the QK matrix of the Transformer can be now written as:

$$A_{\ell,h} = Softmax([Q_{\ell,h}; m_{\ell,h}][K_{\ell,h}; m_{\ell,h}]_{\ell,h}^\top)$$

$$Y_{\ell,h} = A_{\ell,h}[V_{\ell,h}; v_{\ell,h}m_{\ell,h}]$$

where  $Y_{\ell,h}$  is the  $h$ -th head of the output at layer  $\ell$ ,  $Q, K, V$  are the standard transformations of the

query input sequence, and  $m_{\ell,h}$  is the dense feature memory token for layer  $\ell$ . Since  $Q$  and  $V$  are both augmented with dense features, this provides an opportunity for both the dense features to align with query tokens and vice versa.

#### 4.5 Output layer and Optimization

The final output layer of the Transformer stack is then passed into a pooling and MLP layer.

$$Y_{out} = MLP(\psi(Y_L)) \quad (1)$$

where  $\psi(\cdot)$  is a pooling operator that maps  $\mathbb{R}^{n \times d_{model}} \rightarrow \mathbb{R}^{d_{model}}$ . Our  $MLP(\cdot)$  function maps  $\mathbb{R}^{d_{model}} \rightarrow \mathbb{R}^{N_{class}}$  to the number of classes. Our model then optimizes the Softmax cross entropy loss between the true classes and the predicted values.  $L = \sum^L \sum_{i=1}^n y_i \log(\pi_i) + (1 - y_i) \log(1 - \pi_i)$ , where  $\pi_i$  is the prediction of class  $i$  and  $y_i$  is the ground truth label of the class  $i$ .

### 5 COVID-19 Vaccine Access Dataset

Due to the novel nature of COVID-19, no previous datasets exist to accurately learn a model for the purpose of vaccine access query classification. Thus, in this section we outline the process we used to create this dataset.

*Collection.* To collect a dataset of queries to be labeled for vaccine access, we sample anonymized queries from real search traffic. Since a small minority of our search queries are for COVID-19 vaccination topics, we leveraged Google’s Knowledge Graph entities to find queries that included high confidence positives, potential positives, and close negatives. For example, for high precision candidates we sample top and random queries associated with the entity “COVID-19 Vaccination”, while for high recall low precision candidates, we sample queries that are only associated with the entity “COVID-19” or with “Vaccination”.

*Labeling.* To label this dataset for the specific purpose of vaccine access, we rely on a large pool of search quality raters who have deep experience with how health-related information needs are reflected in search queries. These raters were unknown to and independent of the developers of the classifiers. Each query is rated by three independent raters.

*Label Expansion.* We expand our dataset using label propagation to queries that are very similar to labeled queries. We include examples of positive and negative vaccine access queries in Table 5.

*Dense Feature Augmentation.* We augment our dataset by supplementing each query using dense features. To generate these dense features we use a combination of (1) entities mentioned in the query via a proprietary library analogous to Google Cloud Entity Analysis (goo, 2019) and (2) related search queries determined by a proprietary algorithm. We pool these two sources and use the 60,000 most common words and phrases to create a dense feature representation with dimension 60,000. At each dimension, we assign a relevance score for the phrase. For mentioned entities, this is analogous to salience in the Google Cloud Entity Analysis API. Table 6 shows some of the top features that are generated for each classification split.

## 6 Experiments

Below, we dive into experiment setups and results.

### 6.1 Baselines

We compare our proposed approach with three competitive baselines. The choice of baselines serves two primary purposes, i.e., (1) to show our method is competitive against well-established methods, and (2) to confirm certain scientific hypothesis by ablation-like studies. Please see Section A.1 in the Appendix for further implementation details on how we configure and train our models.

**Adaboost** A technique used to create an ensemble of weak learners that begins by fitting an estimator on the original dataset and then repeatedly fits additional estimators focusing more on examples that are misclassified by the combination of all the existing set of estimators. Mathematically, the ensembled AdaBoost classifier can be represented as:  $F_N(x) = \sum_{n=1}^N f_n(x)$  which consists of  $N$  weak learners ( $f_n(x)$ ) that are combined to create the ensemble model  $F_N(x)$ .

**Query-only Transformer** This baseline, simply a VSI Transformer without any dense features, is added to evaluate the upper bound of a language only state-of-the-art classifier.

**Late Fusion Transformer** This is an ablative baseline for the VSI Transformer. Instead of employing dense feature memory, we combine the dense features with the Transformer output at the final layers. Hence, we call this baseline *Late Fusion*, representing how the fusion of modalities is done at the final stages. Concretely, we concatenate the dense features to the pooled output from the transformer layer stack and then add a few layers

of MLP before adding the classification head. See figure 2 for setup details.

## 6.2 Results & Analysis

Table 2 presents the F1 and precision metrics from the vaccine intent classification tasks. VSI Transformer outperforms Adaboost models by relative +15.1% gain on the US locale, +17.7% gain on the CA locale, and +7.6% gain on the GB locale on F1 metric. The gains against traditionally strong ML models are substantial and compelling. When compared with NLU-only approaches (e.g., query-only Transformer), VSI Transformer again strongly outperforms the baseline. Finally, there are modest (but consistently strong) gains against the best and strongest baseline considered of up to +2.1% F1 score.

### 6.2.1 Importance of Text and Dense Features

We study text-only NLU models and feature-only state-of-the-art ML based Adaboost models. Generally, it is not clear if NLU-only models outperform Adaboost models (or vice versa). Both modalities have their fair share of wins and losses across the three datasets and six metrics. To this end, we find that the well-established version of combining text and dense feature to outperform both Adaboost and the Query-only Transformer, signifying the importance of having both modalities for building a successful model.

### 6.2.2 Dense Feature Memory vs Late Fusion

Late fusion is a well-established way to combine end-to-end deep learning with real world tabular features (Severyn and Moschitti, 2015; Tay et al., 2017). Our results show that, while the Transformer with Late Fusion performs the best out of all baselines, the VSI Transformer still comfortably outperforms the Late fusion method, whereby we show that our proposed integration is more effective. In regards to the problem space and domain, this also seems to imply that a deeper fusion of text and dense features can be key in obtaining better model quality.

### 6.2.3 Increasing depth of MLP network

The experiments described above used a single layer in the MLP networks in both the architectures present in Figure 2. Adding a single layer means that we have almost the same number of additional parameters added in both the architectures for a fair comparison.

In this ablation, we study the impact of using multiple layers in the MLP network. Note that each of these layers has a dimension size of 768 and uses GeLU activation function. We see a consistent increase in performance as we increase the number of layers in the MLP network. This increase is more prominent in the Late Fusion architecture which starts to catch up (at the cost of increasing the depth of the model) but still performs worse than the VSI Transformer indicating that for the same number of model parameters, VSI Transformer is a better architecture.

### 6.2.4 Using multiple memory tokens

Table 3 presents the results on increasing the number of memory tokens ( $N_{memory}$ ) to up to 4 tokens. Though intuitively, it might seem that the performance will improve, experiments show that the correlation is not that straightforward.

The metrics seem to improve slightly but there is a consistent degradation observed with the F1 metric as we increase  $N_{memory}$  to 4 tokens. When increasing  $N_{memory}$  from 3 to 4 tokens, all three regions show a degradation. Given the sequence length is constant, our hypothesis is that as we increase the number of memory tokens allocated to the dense features, we are using up tokens that could have otherwise been allocated to the query text. This reduces the length of query that can be seen by the model for large queries thereby impacting the overall performance of the model. Hence, in general, a single memory token might be sufficient, but the optimal number might differ from task to task.

### 6.2.5 Analyzing improvement patterns

Given the improvement in metrics, we looked at some of the examples where VSI Transformer model provides large gains over a query only model. An example pattern is "covid vaccine" followed by some noun. If the noun represents a location, the intent is to search for vaccine availability in that region. The query only Transformer model can only guess whether this represents a location unless it can recall from locations memoized during pretraining. Given that VSI transformer also uses dense features, it has additional signals that provide information like whether the query includes a location by sources like Google Cloud Entity analysis. Hence, using this additional signal, the model is able to cut down on a lot of false positives of this pattern.

Model	Input Features	US	CA	GB
Query-only Transformer	Q only	0.9395 / 0.9060	0.8715 / 0.8059	0.8896 / 0.8159
AdaBoost 20 estimators	DF only	0.8288 / 0.8399	0.7830 / 0.7918	0.8909 / 0.9019
AdaBoost 50 estimators	DF only	0.8570 / 0.8598	0.8315 / 0.8678	0.9132 / 0.9128
Transformer Late Fusion	Q + DF	0.9780 / 0.9698	0.9585 / 0.9289	0.9719 / 0.9519
VSI Transformer	Q + DF	<b>0.9868 / 0.9809</b>	<b>0.9784 / 0.9655</b>	<b>0.9824 / 0.9730</b>
% Improvement (vs Query-only)	-	+5.0% / 8.3%	+12.3% / +19.8%	+10.4% / +19.3%
% Improvement (vs Adaboost)	-	+15.1% / +14.1%	+17.7% / +11.3%	+7.6% / +6.6%
% Improvement (vs best)	-	+0.8% / +1.1%	+2.1% / +3.9%	+1.1% / +2.2%

Table 1: F1 and Precision metrics on COVID-19 vaccination access search intent prediction. VSI Transformer outperforms best Transformer baseline by +0.8% to +2.1% and Adaboost by up to +17.7% F1 score.




Figure 2: Depiction of adding transformed dense features to Transformer based models for a sample query. Architecture on the left depicts late fusion of the dense features to the query embeddings while the one on the right depicts addition of dense features to the query embeddings as a single memory token. Assumes that both the models have sequence length of 8 and the tokenizer produces one token for every word.

Model	$N_\ell$	US	CA	GB
Late Fusion	1	0.9780	0.9585	0.9719
Late Fusion	2	0.9827	0.9690	0.9817
Late Fusion	3	0.9846	0.9788	0.9822
VSI	1	0.9868	0.9784	0.9824
VSI	2	0.9870	0.9785	0.9840
VSI	3	<b>0.9872</b>	<b>0.9789</b>	<b>0.9848</b>

Table 2: Impact of increasing the number of layers ( $N_\ell$ ) of the MLP networks on F1.

### 6.2.6 Using smaller Transformer models

The above experiments are performed with T5 1.1 Base model, but models of this size are often prohibitive for online applications due to their resource requirements and latency where smaller and shallower models are preferred. But shallower models tend not to perform as good and One possible reason is that shallower models are limited in their capability to extract complex features.

$N_{memory}$	US	CA	GB
0	0.9395	0.8715	0.8896
1	0.9868	0.9784	0.9824
2	0.9875	0.9784	0.9850
3	0.9871	0.9788	0.9849
4	0.9853	0.9785	0.9839

Table 3: Trend of F1 metric on increasing the number of chosen memory tokens  $N_{memory}$  in VSI transformer architecture.

Given that dense features can provide complex features as inputs to the model in a preprocessed format, using dense features can provide a large boost in quality for shallow models. Additionally, such features can provide exclusive information that is not available in the text features (query). Hence, we performed an ablation study on usefulness of dense features on the model size.

Table 4 presents results on running the same experiments as above, but using T5 1.1 Small instead of the Base model. Overall, we see larger relative improvements compared to the Base model. For example, the VSI Transformer Small model has a relative F1 gain of 6.0% and precision gain of 10.4% over query-only Transformer Small model, much higher than the relative F1 gain of 5% and precision gain of 8.3% observed with the Base model.

Model	Small	Base
Query-only	0.9308 / 0.8866	0.9395 / 0.9060
Late Fusion	0.9539/0.9336	0.9780 / 0.9698
% Imp. (vs Query-only)	+2.5%/+5.3%	+4.1% / 7.0%
VSI Transformer	0.9862/0.9788	0.9868 / 0.9809
% Imp. (vs Query-only)	+6.0%/+10.4%	+5.0% / 8.3%

Table 4: Comparison of F1 and precision metric between T5 1.1 small and base models on US dataset.

Even through shallow, the VSI Transformer’s architecture is able to make better use of the attention layers for the dense features leading to pretty high boost even over the Late Fusion architecture, affirming that it is a better architecture for shallow models as well.

## 7 Limitations

Though we have shown that just one memory token is sufficient, assigning tokens to dense features means less number of tokens are available in the sequence for the text input. Another limitation is the use of locale specific vocabularies for dense features for each regional dataset as present in table 5, but that is not a limitation of the VSI Transformer but instead how the dense features are generated.

## 8 Conclusion

This paper presented an important task of classifying search queries with COVID-19 vaccination access intent. With an extensive set of experiments and comparing with strong baselines, we presented VSI Transformer, a novel and generic approach that consistently and strongly outperforms all existing baselines that operate on either of the two modalities, or late fusion of the modalities. With an ablation study on model size, we show that for online applications where shallower models need to be deployed primarily due to latency constraints, making use of dense features can help bridge the gap in performance compared to deeper models. Future work in this direction can help further understand how to choose the optimal number of memory to-

kens, and explore more architectures to efficiently combine sequential and dense features.

## Ethics Statement

In our ongoing fight against the COVID-19 pandemic, understanding whether search queries exhibit an intent to seek vaccine access is an important problem to study to be able to collect search statistics about COVID-19 vaccination efforts. These statistics are made public via an online website, an interactive report that is accessible by anyone. All data used to train these models are anonymized and sampled, and labeled by a large pool of search quality raters who are trained to assess health-related information needs in search queries. These raters are unknown to and independent from the developers of the classifiers. This dataset is not made public or used in any other context.

## References

2019. Google Cloud natural language api basics. ["https://cloud.google.com/natural-language/docs/basics#entity\\_analysis](https://cloud.google.com/natural-language/docs/basics#entity_analysis).
- Niki Parmar Jakob Uszkoreit Llion Jones Aidan N. Gomez Lukasz Kaiser Illia Polosukhin Ashish Vaswani, Noam Shazeer. 2017. *Attention is all you need*. 31st Conference on Neural Information Processing Systems (NIPS).
- Steven M. Beitzel, Eric C. Jensen, Ophir Frieder, David A. Grossman, David D. Lewis, Abdur Chowdhury, and Aleksander Kolcz. 2005. Automatic web query classification using labeled and unlabeled training data. In *SIGIR '05*.
- Andrei Z. Broder, Marcus Fontoura, Evgeniy Gabrilovich, Amruta Joshi, Vanja Josifovski, and Tong Zhang. 2007. [Robust classification of rare queries using web knowledge](#). In *Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '07*, page 231–238, New York, NY, USA. Association for Computing Machinery.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Mikhail S. Burtsev, Yuri Kuratov, Anton Peganov, and Grigory V. Sapunov. 2020. [Memory transformer](#).

- Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [Bert: Pre-training of deep bidirectional transformers for language understanding](#).
- Github Google. 2021. [Sentencepiece Tokenizer](#).
- Daqing He, Daqing, Göker, and Ayse Goker. 2000. Detecting session boundaries from web user logs.
- Andrew Jaegle, Sebastian Borgeaud, Jean-Baptiste Alayrac, Carl Doersch, Catalin Ionescu, David Ding, Skanda Koppula, Daniel Zoran, Andrew Brock, Evan Shelhamer, et al. 2021. Perceiver io: A general architecture for structured inputs & outputs. *arXiv preprint arXiv:2107.14795*.
- Daxin Jiang () and Jian-Tao Sun. 2009. [Context-aware query classification](#). In *SIGIR' 09, The 32nd Annual ACM SIGIR Conference*. Association for Computing Machinery, Inc.
- Kowsari, Jafari Meimandi, Heidarysafa, Mendu, Barnes, and Brown. 2019. [Text classification algorithms: A survey](#). *Information*, 10(4):150.
- Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosior, Seungjin Choi, and Yee Whye Teh. 2019. Set transformer: A framework for attention-based permutation-invariant neural networks. In *International Conference on Machine Learning*, pages 3744–3753. PMLR.
- Xiao Li, Ye-Yi Wang, and Alex Acero. 2008. [Learning query intent from regularized click graphs](#). In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '08, page 339–346, New York, NY, USA. Association for Computing Machinery.
- Sofus Macskassy, Aynur Dayanik, Haym Hirsh, and Morrocroft Ii. 1999. Emailvalet: Learning user preferences for wireless email.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient estimation of word representations in vector space](#).
- Sinno Jialin Pan and Qiang Yang. 2010. [A survey on transfer learning](#). *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Anirudh Ravula, Chris Alberti, Joshua Ainslie, Li Yang, Philip Minh Pham, Qifan Wang, Santiago Ontanon, Sumit Kumar Sanghai, Vaclav Cvicek, and Zach Fisher. 2020. [Etc: Encoding long and structured inputs in transformers](#). In *2020 Conference on Empirical Methods in Natural Language Processing (EMNLP 2020)*.
- Matthew Richardson, Ewa Dominowska, and Robert Ragno. 2007. [Predicting clicks: Estimating the click-through rate for new ads](#). In *Proceedings of the 16th International Conference on World Wide Web, WWW '07*, page 521–530, New York, NY, USA. Association for Computing Machinery.
- Adam Sadilek, Stephanie Caty, Lauren DiPrete, Raed Mansour, Tom Schenk, Mark Bergholdt, Ashish Kumar Jha, Prem Ramaswami, and Evgeniy Gabrilovich. 2018. Machine-learned epidemiology: real-time detection of foodborne illness at scale. *NPJ Digital Medicine*, 1.
- Adam Sadilek, Yulin Hswen, Shailesh Bavadekar, Tomer Shekel, John Brownstein, and Evgeniy Gabrilovich. 2020. [Lymelight: forecasting lyme disease risk using web search data](#). *npj Digital Medicine*.
- Gerard Salton and Christopher Buckley. 1988. [Term-weighting approaches in automatic text retrieval](#). In *Information Processing Management*, 24(5):513–523.
- Aliaksei Severyn and Alessandro Moschitti. 2015. Learning to rank short text pairs with convolutional deep neural networks. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*, pages 373–382.
- Noam Shazeer. 2020. Glu variants improve transformer. *arXiv preprint arXiv:2002.05202*.
- Noam Shazeer, Youlong Cheng, Niki Parmar, Dustin Tran, Ashish Vaswani, Penporn Koanantakool, Peter Hawkins, HyoukJoong Lee, Mingsheng Hong, Cliff Young, et al. 2018. Mesh-tensorflow: Deep learning for supercomputers. *Advances in neural information processing systems*, 31.
- Noam Shazeer and Mitchell Stern. 2018. Adafactor: Adaptive learning rates with sublinear memory cost. In *International Conference on Machine Learning*, pages 4596–4604. PMLR.
- Dou Shen, Jian-Tao Sun, Qiang Yang, and Zheng Chen. 2006. [Building bridges for web query classification](#). In *Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '06, page 131–138, New York, NY, USA. Association for Computing Machinery.



Cees G. M. Snoek, Marcel Worring, and Arnold W. M. Smeulders. 2005. [Early versus late fusion in semantic video analysis](#). In *Proceedings of the 13th Annual ACM International Conference on Multimedia*, MULTIMEDIA '05, page 399–402, New York, NY, USA. Association for Computing Machinery.

Sainbayar Sukhbaatar, Edouard Grave, Guillaume Lample, Herve Jegou, and Armand Joulin. 2019. Augmenting self-attention with persistent memory. *arXiv preprint arXiv:1907.01470*.

Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. 2020. Efficient transformers: A survey. *arXiv preprint arXiv:2009.06732*.

Yi Tay, Mostafa Dehghani, Jai Gupta, Dara Bahri, Vamsi Aribandi, Zhen Qin, and Donald Metzler. 2021. Are pre-trained convolutions better than pre-trained transformers? *arXiv preprint arXiv:2105.03322*.

Yi Tay, Minh C Phan, Luu Anh Tuan, and Siu Cheung Hui. 2017. Learning to rank question answer pairs with holographic dual lstm architecture. In *Proceedings of the 40th international ACM SIGIR conference on research and development in information retrieval*, pages 695–704.

Wikipedia. 2022. [AdaBoost](#).

Qingyang Wu, Zhenzhong Lan, Jing Gu, and Zhou Yu. 2020. [Memformer: The memory-augmented transformer](#).

Pengcheng Yin, Graham Neubig, Wen-tau Yih, and Sebastian Riedel. 2020. [TaBERT: Pretraining for joint understanding of textual and tabular data](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8413–8426, Online. Association for Computational Linguistics.

Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. 2020. Big bird: Transformers for longer sequences. *Advances in Neural Information Processing Systems*, 33:17283–17297.

Fengbin Zhu, Wenqiang Lei, Youcheng Huang, Chao Wang, Shuo Zhang, Jiancheng Lv, Fuli Feng, and Tat-Seng Chua. 2021. [TAT-QA: A question answering benchmark on a hybrid of tabular and textual content in finance](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3277–3287, Online. Association for Computational Linguistics.

## A Appendix

### A.1 Implementation Details

The VSI Transformer is implemented in Mesh Tensorflow (Shazeer et al., 2018), a Tensorflow-like

API that supports distributed model parallelism. We initialize our model with the T5.1.1 Base checkpoint, comprised of 12 encoder layers,  $d_{model}$  size of 768,  $d_{ff}$  of 2048. The model uses GEGLU-based feed-forward layers as described in (Shazeer, 2020). The model has 12 heads. The overall number of non-embedding parameters is approximately 100M parameters. The model also utilizes the standard 32K SentencePiece that was trained on the C4 corpus. Our model is trained with 16 TPUv3 chips. We finetune all models using a sequence length of 32 subword tokens using the Adafactor optimizer (Shazeer and Stern, 2018). The learning rate is a constant learning rate of  $10^{-3}$  and the batch size is 128.

### A.2 Top Search Queries

Figure 3 presents a view of the tool listing top search queries associated with vaccination intent.

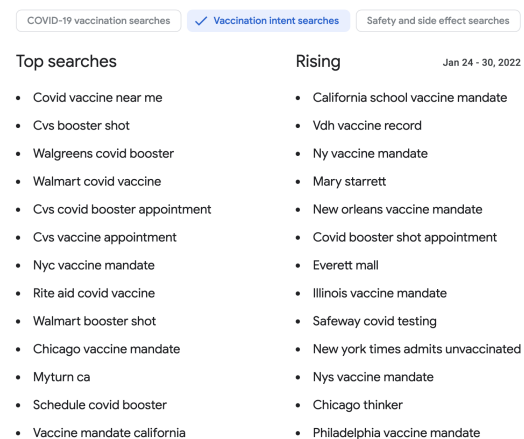


Figure 3: VSI tool presenting top search queries with vaccination intent

### A.3 Sample Queries

Table 5 presents a sample of some positive and negative queries for vaccine acces present in our dataset. The examples in the tabe shows how nuanced the problem is and why simply looking for termsatches didn’t perform as good and a more sophisticated approach was needed.

### A.4 Dense Feature Vocabulary

As mention in section 5, the dataset consists of dense features created from a vocabulary of 60,000 phrases. Table 6 presents a list of top phrases present in the vocabulary of each of the 3 regions.

Class	Examples
Positive	<i>covid vaccine appointment where can i get covid vaccine book covid jab walk-in covid vaccine near me nhs book covid vaccine</i>
Negative	<i>covid stats by country covid vaccine effectiveness fully vaccinated travel proof of covid vaccination how long does the vaccine last</i>

Table 5: Examples of positive and negative COVID-19 vaccine access queries.

Category	Top Features
Vaccination access (US)	<i>pharmacy, pfizer, vaccine appointment, appointment, pharmacies, moderna, dose, appointments, pfizer vaccine, cvs, walgreens, second dose, vaccine appointments, cvs pharmacy, doses, shot, cvs covid, walgreens pharmacy, vaccine eligibility, moderna vaccine</i>
Vaccination access (GB)	<i>appointment, appointments, book, booking, vaccination centre, clinic, vaccination centres, vaccine appointment, clinics, nhs uk, nhs, coronavirus covid, walk in, coronavirus vaccination, covid vaccination, vaccination clinic, vaccine clinic, centres, vaccination appointment, vaccine centre, centre, book covid, pfizer, astrazeneca</i>
Vaccination access (CA)	<i>clinic, appointment, clinics, vaccine clinic, vaccine appointment, vaccination clinic, pharmacy, appointments, book, booking, pharmacies, registration, pfizer, drug mart, vaccination center, walk in, vaccination appointment, vaccination clinics, vaccine pop up</i>

Table 6: Top features associated with a query in the dense feature vector, for each split.