Ophelos

USING MACHINE LEARNING TO DETECT VULNERABILITY

Contents

Introduction

It's not easy dealing with debt. And it can be especially difficult for those who are more vulnerable or in vulnerable circumstances.

At Ophelos, caring for vulnerable customers is of utmost importance. It goes far beyond a need for regulatory compliance - it is central to our aim of providing each customer with the right care and the best customer experience.

Understanding and identifying vulnerability is not a simple task; many individuals are shy or embarrassed about their vulnerability, may not feel comfortable sharing what is often a very personal subject, or are even unaware of their vulnerability.

This paper unpacks how we have harnessed the power of AI and Natural Language Processing to flag up vulnerability and better cater for every customer.

“Around 1 in 5 customers have an additional vulnerability on top of their financial hardship”

There is no set criteria for what constitutes a ‘vulnerable’ customer when it comes to dealing with debt.

According to the FCA (the UK financial regulator) a vulnerable customer is “someone who, due to their personal circumstances, is especially susceptible to harm, particularly when a firm is not acting with appropriate levels of care”. This definition of course is relatively broad-brush, meaning that anyone who finds it challenging because of their situation or health could be considered vulnerable.

Some examples of common drivers of vulnerability are:

Recent bereavement

Mental health problems

Long-term or terminal illness

Learning disability

Relationship breakdown

Addiction

This list is not exhaustive and there are many other reasons why someone could find dealing with debt difficult. According to StepChange, 1 in 5 of their customers in 2017 had an additional vulnerability on top of their financial hardship.

Vulnerability is never solely about the individuals' characteristics or circumstances. It is also strongly influenced by the way an organisation treats their customers.

Everyone is potentially vulnerable to harm, and it is the responsibility of the organisation to prevent customers from becoming vulnerable, whilst also protecting those already in vulnerable circumstances by providing extra levels of support.

The first step to helping customers in vulnerable situations is identifying them. Customers communicate in different ways and so it is important to be on the lookout for any indication that a customer may need extra support.

To ensure we have the greatest success in identifying vulnerable customers we have developed OLIVE (Ophelos Linguistic Identification of VulnErability).

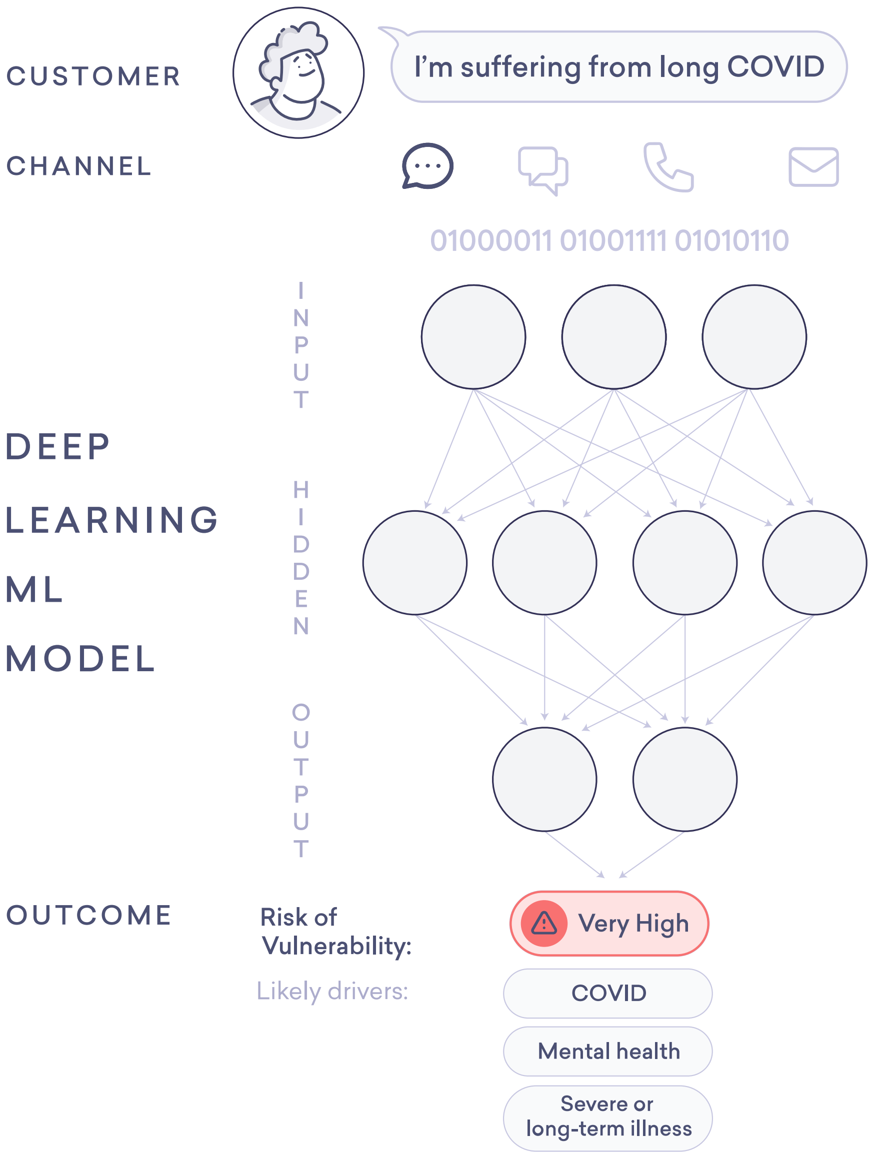

MEET OLIVE

OLIVE is a cutting-edge natural language processing (NLP) model capable of predicting the likelihood that a customer is vulnerable and identifying the possible causes. OLIVE reads customers messages and flags up potential vulnerability to customer support agents in real-time, helping to ensure that the customer receives a necessary level of support.

OLIVE is a supervised learning model. Supervised learning is the most common form of machine learning used to solve real world problems and is the industry standard for high accuracy text classification. The model learns to detect vulnerability by being shown thousands of examples of vulnerable and non-vulnerable messages. When the model is trained it can be used to make predictions on unseen messages.

We trained OLIVE using a dataset of over 33,000 sample messages that we could expect to receive from customers, which were labelled by our in-house vulnerability team.

Model Architecture

OLIVE's predictions are made using two models with similar designs. The first model is a binary classifier which predicts the risk of the customer's vulnerability from their message. The second model is a multi-label classifier which predicts the likely causes of the vulnerability, such as 'Alcohol' or 'Gambling' for example. The second model is only used if the first model has determined that there is sufficient risk that the customer is vulnerable.

Both models were trained independently, allowing them to specialise for their particular task. The two-model approach was found to achieve a higher accuracy in both tasks over training one model.

Each model is a deep learning model built using RoBERTa. RoBERTa is Facebook's larger and improved version of Google's BERT model. BERT (Bidirectional Encoder Representations from Transformers) is a machine learning technique and paper released by Google in 2018 which achieved industry-leading performance on many benchmark NLP tasks. RoBERTa builds on BERT's success but leverages more training data and computational resources to outperform it.

The basic structure of each model is an encoder and a classifier head. The encoder is responsible for converting the message's raw text to an array of 1024 numbers which encapsulates the semantic meaning of the sentence. The classifier head then takes this array and outputs a prediction e.g. "high risk of vulnerability" for the binary classifier model, or "Gambling" & "Alcohol" for the categoriser. In total, each model contains over 355 million trainable parameters.

OLIVE leverages transfer learning. Transfer learning is a machine learning technique in which a model is pre-trained extensively on a large dataset for a general task and then fine-tuned for a more specific task. Transfer learning has become the industry standard for NLP since understanding language at a general level is required for any text classification task. The encoder in OLIVE was pre-trained on a very large dataset (160GB) of general text which provided it with a strong understanding of English language.

We then add the classifier head on top and fine tune the model by training it with the vulnerability data. This not only trains the classifier head to make accurate predictions but also modifies the encoder so that the model now understands English in a way specialised to debt collection and vulnerability.

Dataset description

The data used to train OLIVE consisted of 33,663 phrases and sentences similar to those we could expect to receive from customers. These phrases were extracted by web scraping a popular UK-based debt discussion forum. To create a dataset which mimics customer conversations, the forum sentences were filtered by comparing them to sample customer messages.

Our vulnerability and customer support expert created 200 short sample messages designed to mimic a customer talking to a support agent. These sample phrases covered various message topics including vulnerability, disputes and general queries. For example: 'I lost my job', 'My wife left me', 'I have made a payment', 'This isn't my debt'.

Forum phrases were compared to these sample phrases using an embedding and cosine similarity. This technique allows us to calculate how close semantically-speaking a forum phrase is to the sample messages. Any phrase which was too far away from the sample phrases was discarded. This helps to remove any phrases which are wildly off-topic or do not represent the style of message we would expect to receive.

Example messages

"I used to think bankruptcy was the only option available to fools like me but I'd much rather pay what I owe"

"I'm paying the mortgage on the property I'm living in with the children"

"I unlocked a gambling site that has some of my winnings"

"Lockdown feels so tough this time"

Each phrase was either labelled as non-vulnerable or labelled with as many vulnerability drivers as required. For example "I lost my job and I'm now depressed" would be labelled as "Job loss", "Income shock" and "Mental health" whereas "This isn't my debt" would be labelled as non-vulnerable.

We based the vulnerability drivers used to categorise messages on the FCA's list of likely drivers:

Learning difficulties	**Severe or long-term illness**
Low knowledge in managing finances	**Visual impairment**
Poor digital skills	**Bereavement**
Poor English language skills	**Domestic abuse**
Addiction	**Income shock**
Alcohol	**Job loss**
Cognitive disability	**Relationship breakdown**
Drug abuse	**Retirement**
Gambling	**Erratic or inadequate income**
Hearing impairment	**Over-indebtedness**
Mental health	**Covid**
Physical disability	**Other**

Of the 33,663 phrases, 7459 showed indicators of vulnerability

The most common drivers were:

- **over-indebtedness**
- **mental health**
- **low knowledge in managing finances**

Model training & results

The binary classifier and categorisation models were trained separately. The binary classifier was trained on the entire dataset and tasked with predicting whether the phrase was vulnerable or non vulnerable.

For the categoriser, only the phrases labelled as vulnerable were used and the model was tasked with predicting all the categories assigned to the vulnerable phrase. For both models their datasets were split into 60% training with 20% held back for validation and 20% for testing.

To ensure OLIVE had the greatest predictive power we tested several different model architectures. The candidate models consisted of a mixture of light and deep learning models with the objective to find the model with the greatest performance. Since OLIVE is intended to be used for real-time inference it must be able to make fast predictions. It should also not require extensive computational resources such as a GPU for making predictions. These requirements rule out extremely large models such as Google's T5 and also rule out excessive model ensembling.

The light candidate models used TF-IDF (Term Frequency Inverse Document Frequency) with bigrams to convert each message into an array of its vocabulary. Using these features we trained classifiers using logistic regression and LightGBM (Gradient Boosted Decision Trees).

The deep learning models consisted of 4 BERT-based models of varying size. These should display increased performance as the models increase in size but at the detriment of inference speed and computational resources. The final deep learning candidate model was Electra. Electra was released in 2020 by Google and uses a more efficient version of BERT's pre-training method. It has achieved very high benchmark scores whilst requiring significantly less training time than its counterparts.

Model	AUC	F1	Accuracy
TF-IDF + logistic regression	0.91	0.59	0.85
TF-IDF + lightgbm	0.91	0.65	0.86
distilRoBERTa-base (82M parameters)	0.94	0.76	0.89
RoBERTa-base (125M parameters)	0.94	0.76	0.89
BERT-large (335M parameters)	0.94	0.74	0.88
RoBERTa-large (335M parameters)	0.95	0.77	0.90
Electra-large-discriminator (335M parameters)	0.94	0.77	0.89

Once a phrase has been labelled with some risk of vulnerability it is passed to the categorisation model to predict the likely drivers. Whilst the overall vulnerability of a message is on a spectrum, we wish to make a hard prediction when it comes to predicting the likely drivers. Therefore, the more traditional classification metrics: F1 and Accuracy are the most important.

Similarly to the binary classification task, the deep learning models significantly outperform the light models. RoBERTa-large also achieved the greatest scores on all three metrics and so we chose this for the final model.

Message	Risk of vulnerability	Drivers
I have been furloughed and I'm struggling	Very high	Covid Income shock
This isn't my debt, please don't contact me	Low	None
I have no idea what to do	High	Low knowledge in managing finances Mental health

Conclusion

In this whitepaper we have shown how machine learning and the latest research in natural language processing can be applied to debt collection to better help customers in vulnerable circumstances. Using OLIVE we can greatly increase our chances of identifying vulnerable customers and ensure they get the support they need. We can also report back to our clients data on the vulnerabilities of their customers to help them improve their product.

The next steps for upgrading OLIVE are to move from pure natural language processing to a more general classification model. OLIVE would then make a prediction on the customer as a whole rather than just using their messages. Pieces of information such as the value of the debt, socio-economic data, credit reference agency data and other customer interactions would provide the model with greater predictive power in identifying vulnerable customers.

The final goal for OLIVE is to allow vulnerable customers the opportunity to self-serve. Often customers do not know they are vulnerable or would rather get support without having to go through a customer support agent. For example, if we detect that the customer is struggling with gambling issues we could provide them with resources and support immediately without having to go through a customer support agent. Of course, the customer would always be able to speak with an agent quickly if they prefer.

> **“Using Olive we can greatly increase our chances of identifying vulnerable customers”**

The process for building OLIVE can be applied to other conversational topics in debt collection such as dispute handling, complaints and promise to pay. Disputes are one of the most time-consuming tasks for a customer support agent and being able to triage customers and allow them to self-serve would drastically reduce the associated time and cost.

At Ophelos, we are focused on changing the narrative in debt collection and putting customers first. By using industry-leading machine learning techniques in all aspects of debt collection we can transform a traditionally lagging industry and bring it into the digital age.

About The Author

Jacob Goss is a Data Scientist at Ophelos. He builds and deploys end-to-end machine learning models with a focus on business operations optimisation and natural language processing. Alongside his work in AI, Jacob develops fully automated and scalable cloud based data pipelines and data lakes.

About Ophelos

Ophelos is a London-based technology company building products that nurture financially healthy and trusted relationships between customers and businesses. Our first product is a customer-centric debt management platform that digitises and automates the debt collection process, to allow customers to resolve their debts on their own terms and improve their financial health.

Ophelos is authorised and regulated by the FCA and is a Pending B Corporation. For more information, please visit our website.

sales@ophelos.com

020 3318 2823

68 Hanbury Street
London
E15JL
United Kingdom

www.ophelos.com