

Natural Language Processing for Financial Regulation

Ixandra Aчитouv

*Department of Mathematics, Imperial College London and CNRS
Complex Systems Institute of Paris Île-de-France*

ixandra.achitouv@cnrs.fr

Dragos Gorduza[†]

Oxford Man Institute of Quantitative Finance

Antoine Jacquier[†]

*Department of Mathematics
Imperial College London and the Alan Turing Institute*

Corresponding Author: Ixandra Aчитouv

[†]These authors contributed equally to this work.

Copyright © 2024 Ixandra Aчитouv, et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Abstract

This article provides an understanding of Natural Language Processing techniques in the framework of financial regulation, more specifically in order to perform semantic matching search between rules and policy when no dataset is available for supervised learning. We outline how to outperform simple pre-trained sentences-transformer models using freely available resources and explain the mathematical concepts behind the key building blocks of Natural Language Processing.

Keywords: Language models, Financial Regulation, Natural Language Processing

1. INTRODUCTION

Over the past ten years, modern natural language processing models have revolutionised the field of artificial intelligence, transforming how computers understand and generate language. ChatGPT is one key example of how AI technology is becoming increasingly important in generating more and more precise human-like responses to a wide variety of problems, including business ones. These advancements have significant implications for the financial industry, where vast amounts of regulatory data and policies must be analysed and understood in order to comply with laws and regulations.

Indeed, Natural Language Processing (NLP) has gained traction in financial regulation as a solution for managing and interpreting complex regulatory information. The ability to effectively process and analyse large amounts of regulatory text is crucial for compliance officers, risk managers, and other financial professionals. However, traditional methods of automatic text processing including keyword searches and dictionaries are both inefficient to run at the scale required by the amount of

files to treat, and costly to set up since they require system experts to design dictionaries on a case-by-case basis. Moreover, they are also prone to errors due to lack of coverage which can expose the user to costly regulatory violations. Modern NLP techniques have emerged in recent years, mostly based on deep neural networks and after the seminal breakthrough [1], by Vaswani, Shazeer, Parmar, Uszkoreit, Jones, Gomez, Kaiser and Polosukhin, who incorporated the attention mechanism to learn latent semantic links between words in a sentence more accurately. These improvements, along with orders of magnitude increases in the availability of datasets and computational power have led to large language model (LLM), a type of model known for its ability to achieve general-purpose language understanding and generation such as OpenAI’s GPT (used in ChatGPT), Meta’s LLaMa or Google’s PaLM. One may however wonder whether LLMs are actually required for simple NLP tasks such as semantic search, especially due to their costs, instead of free open resources.

In the present work, we perform a semantic search applied to laws and regulation using a so called domain adaptation technique with limited amount of data, bearing in mind limited available resources from small and medium enterprises. We also recall the historical NLP approaches and explain the leading concepts in this field. In particular, we focus on semantic search, which naturally leads to the encoding concepts of a sentence. We thus aim to fill the gap between well-known NLP concepts and the mathematical frameworks behind.

This paper is organised as follows: Section 2 is a short historical review on encoding methods of a sentence into a mathematical object as well as key machine learning advances. In Section 3 we describe the task of semantic textual similarity, the different techniques at hand and illustrate them with pseudo codes. In Section 4, we apply these techniques concretely to financial regulation problems to improve semantic search matching. Section 5 draws conclusion from this analysis.

2. HISTORICAL OVERVIEW OF WORD EMBEDDING APPROACHES

Semantic vectors of words are high-dimensional vectors, and pairwise distances between two vectors reflect the similarity of words. We briefly review the main literature on this topic, clarifying the notions of ‘similarity’, ‘meaning’ and ‘correlation between words’.

2.1 Deterministic Approaches Based on Word Counts

The roots of major modern NLP techniques can be traced back Firth’ quote [2], ‘*You shall know a word by the company it keeps*’. Word co-occurrence statistic is a simple deterministic approach to obtain a semantic vector representation of words. Latent Semantic Analysis (HAL) [3], and Hyperspace Analogue to Language (HAL) [4], are historically the two methods that built embedded semantic vectors derived from the statistics of word co-occurrence. LSA derives its vectors from collections of segmented documents, while HAL makes use of unsegmented text corpora. For segmented documents, a document-term matrix is used to describe the frequency of terms that occur in a collection of documents such that rows correspond to the documents and columns correspond to the terms. Each (i, j) cell, then, counts the number of times word j occurs in document i . Later, Rohde, Gonnerman and Plaut [5], improved the HAL method introducing the Correlated Occurrence Analogue to Lexical Semantics (COALS) algorithm, which achieves considerably better performance through improvements in normalisation and other algorithmic details.

2.1.1 Co-occurrence & singular value decomposition

The deterministic approach proposed in [5], is one of the first methods to achieve good and consistent results in predicting human similarity judgements in words embedding. It is a two-step approach:

Step 1: produce a co-occurrence matrix. For each word w_1 , count the number of times another word w_2 occurs in close proximity to w_1 . Here 'close proximity' (or 'windows size') is a hyperparameter indicating the number of words to consider before/after the central word w_1 . HAL suggested to use a 4-word window while COALS proposes 10-word windows. The counting is done using a weighting scheme whereby if w_2 is adjacent to w_1 it receives the maximal weight, and the weight is proportional to the length separation of the central word w_1 (ramped window). This results in an $N \times N$ co-occurrence matrix where N represents the number of words in the whole corpus. One could in principle stop here, where each row of the co-occurrence matrix can thus be associated with a word. With this procedure, HAL [4], demonstrated semantic similarity between any desired pair of words. A few tricks were also proposed in the computation of the co-occurrence matrix, for instance putting a threshold to the number of co-occurrence count to remove frequent words such as 'the' or 'a'.

Step 2: Dimension reduction of word vectors. From the co-occurrence count of words the resulting embedded vectors should have at least the size of the words in the corpus¹. In HAL, the authors proposed to keep a relatively small number of principal components of the co-occurrence matrix. For instance reducing the vector size by eliminating all but the k columns with the highest variance. In [5], the reduction of dimension is based on Singular Value Decomposition (SVD) of the normalised co-occurrence matrix (this was also proposed in for latent semantic analysis in [3]).

2.2 Machine Learning

Over the last decade, Machine Learning algorithms contribution to NLP have brought significant progress. Among these, *Word2vec* in 2013 by Mikolov, Chen, Corrado and Dean [6], *Glove* in 2014 by Pennington, Socher and Manning [7], *BERT* in 2018 by Devlin, Chang, Lee and Toutanova [8], and *SBERT* in 2019 by Reimers and Gurevych [9], are the most popular ones, which we shall use in this project. The key idea is to perform a non-linear regression where the coefficients are adjusted by minimising a loss function (or maximising a likelihood function). In the case of NLP, the coefficients/weights are the embedded vectors.

2.2.1 SKIP-GRAM and CBOW

To understand the *Skip-gram* model, we briefly review some fundamental concepts of probabilistic language modelling where the goal is to compute the probability $\mathbb{P}(w)$ of a sequence of words

¹ In fact it has a $2N$ dimension in HAL where for every word in the target vocabulary, there is both a row and a column containing relevant values. For instance, the row may contain co-occurrence information for words appearing before the word under consideration, while the column contains co-occurrence information for words following it. This (row,column) pair may be concatenated so that, given an $N \times N$ co-occurrence matrix, a co-occurrence vector of length $2N$ is available. In COALS the distinction between left/right of the central word is ignored.

$w := (w_0, \dots, w_n)$ occurring; we also write $w_{i:j} := (w_i, \dots, w_j)$ and omit i whenever $i = 0$. A similar goal is to compute the conditional probability of an upcoming word $\mathbb{P}(w_j|w_{j-1})$. The computation of these probabilities rely on conditional probabilities where the general chain of rule can be expressed as

$$\mathbb{P}(w) = \mathbb{P}(w_0)\mathbb{P}(w_1|w_0)\mathbb{P}(w_2|w_{:1}) \cdots \mathbb{P}(w_n|w_{:n-1}).$$

To estimate this probability we count from a corpus

$$\mathbb{P}(w_j|w_{:j-1}) = \frac{\#w_{:j}}{\#w_{:j-1}}.$$

For sentences with more than a few words, the probability that a complex sentence repeats itself often enough to be able to compute statistical properties is unlikely. To overcome this, one usually assumes a pseudo-Markovian approximation (the k -gram model)², whereby $\mathbb{P}(w_j|w_{:j-1}) \approx \mathbb{P}(w_j|w_{j-k:j-1})$ for some $k \leq j$. For instance for $k = 1$ we get the properly Markov bi-gram model $\mathbb{P}(w_j|w_{:j-1}) \approx \mathbb{P}(w_j|w_{j-1})$. Given a sentence, computing for each word the probability that it occurs given the previous words is similar to filling a normalised co-occurrence matrix.

Skip-gram models are generalisations of n -gram models, in which the components do not need to be consecutive, but may leave gaps that are skipped over [10, 11]. For instance, the adjacent words for the sentence: ‘*I hit the blue ball*’ are $\{I \text{ hit}, \text{hit the}, \text{the blue}, \text{blue ball}\}$ for a bi-gram and $\{I \text{ hit}, I \text{ the}, I \text{ blue}, \text{hit the}, \text{hit blue}, \text{hit ball}, \text{the blue}, \text{the ball}\}$ for a 2-skip-bi-gram. Finally, in Word2vec [12], the definition of the Skip-gram architecture is to predict surrounding words given a specific word, which formally corresponds to computing $\mathbb{P}(w_{n-j:n+j}|w_n)$. while the Continuous Bag of Words (CBOW) architecture predicts the current word based on the context: $\mathbb{P}(w_j|w_{:j-1})$. These two models are illustrated in [12, FIGURE 1].

2.2.2 Link with the embedding vectors

In the Skip-gram originally presented in [12], the training objective is to learn word vector representations that are good at predicting the *close-by* words $w(t-2), w(t-1), w(t+1), w(t+2)$ from an input word $w(t)$. This is achieved by maximising the average log likelihood probability, equivalently by minimising the loss function

$$\mathcal{J}(\Theta) := -\frac{1}{T} \sum_{t=1}^T \sum_{-c \leq j \leq c, j \neq 0} \log \mathbb{P}(w_{t+j}|w_t), \quad (1)$$

where c is the size of the training context, (w_1, \dots, w_T) is a sequence of training words and Θ is the vector representation of words that the algorithm tries to optimise. Note that every word w appears both as a central word (input w_I) and as a context word (output w_O) so that $\Theta \in \mathbb{R}^{2w}$ where w is the size of the vocabulary³. In the basic Skip-gram formulation, a softmax function is proposed to model the loss function, namely

$$\mathbb{P}(w_O|w_I) = \frac{\exp(v_{w_O}^\top v_{w_I})}{\sum_{w=1}^w \exp(v_w^\top v_{w_I})}, \quad (2)$$

² strictly speaking it is not Markovian for $k \neq 1$ because we keep a memory of the k previous words

³ the final embedded word w is simply the average of its input w_I and output w_O vector representations.

where w_O and w_I are the output and input words, and, given a word w , v_{w_I} and v_{w_O} are its input and output vector representations. The loss function is then minimised via gradient descent: the gradient of the likelihood (with respect to v_{w_I}) can be computed easily and is equal to zero if and only if

$$v_{w_O} = \left[\sum_{w=1}^{\mathfrak{w}} \exp(v_w^\top v_{w_I}) \right]^{-1} \sum_{w=1}^{\mathfrak{w}} v_w \exp(v_w^\top v_{w_I}), \quad (3)$$

which is the expectation of v_{w_O} occurring given the context word v_{w_I} .

In practice, to improve computation efficiency, a hierarchical softmax is proposed in [12, Equation (3)], which uses a binary tree representation such that the computing time cost reduces from \mathfrak{w} to $\log(\mathfrak{w})$. An alternative technique to improve efficiency is to sub-sample k random words and compute their probabilities to appear in the context of the central words. This is known as negative sampling where $\log \mathbb{P}(w_{t+j}|w_t)$ in (1) is replaced by

$$\log \sigma \left(v_{w_O}^\top v_{w_I} \right) + \sum_{j=1}^k \mathbb{E}_{w_j \sim P_n(w)} \left[\log \sigma \left(-v_{w_j}^\top v_{w_I} \right) \right], \quad (4)$$

where $\sigma(\cdot)$ is the sigmoid function and the sum runs over a random sample of k words (rather than \mathfrak{w}) based on their frequencies around the center word. Thus, the task is to distinguish the target word w_O from draws from noise distribution $P_n(w)$ using logistic regression. The noise distribution $P_n(w)$ in (4) is a free parameter and [12] found that $P_n(w) = u(w)^{3/4}/Z$, where $u(w)$ is a unigram and Z is just the normalisation constant, outperformed computational efficiency significantly over the uniform distributions. This avoids summing over all words in the denominator of (2). In [12], typical values for k lie in the range [5, 20] for small training datasets, while for large datasets k can be as small as [2, 5]. An implementation of this model can be found at https://github.com/chrisjmccormick/word2vec_commented/blob/master/word2vec.c. Improvements of the CBoW and Skip-gram models have been proposed but still rely on the same fundamental idea of co-occurrence statistics of words from [7].

2.2.3 Transformer and attention mechanism

Another major step was achieved using multi-layer bidirectional Transformer encoders [1] for the training architecture, and was used in BERT [8] (*Bidirectional Encoder Representations from Transformers*). A main difference with Word2vec is that BERT is a non-static model with not just one hidden layer. Another key difference is that there is no window-size limitation of n -words in BERT to consider the context of the word to encode. The technique employed instead is known as the *attention mechanism* and is the building block of a transformer architecture. Transformers were originally designed for translation and they are now the state of the art architecture in NLP. The transformer model architecture can be seen in [1, FIGURE 1].

Transformers can be summarised as encoder-decoder tools, where a sentence is fed to the encoder and the decoder generates another sentence; the whole system is built by stacking several layers or encoder/decoder. Each encoder has two sub-layers, a multi-head self-attention layer and a position-wise fully connected feed-forward network, and receives a list of vectors as input, passes it into a self-attention layer, then into a feed-forward neural network, and finally sends out the output

upwards to the next encoder. The decoder itself has three layers: in addition to the two sub-layers of the encoder layer, the decoder inserts a third sub-layer, which performs multi-head attention over the output of the encoder stack and helps the decoder focus on relevant parts of the input sentence.

The Self-attention mechanism is the building block of a Transformer and captures the relationship between the different words in a sentence. All attention mechanisms in a Transformer are split into independent heads whose outputs are concatenated before being further processed. A self-attention mechanism computes a score value between a query with keys, and reweighs the value. All heads' values will be used through a scale dot architecture to update the attention weight matrix. A self-attention head (assuming one single head) can be summarised as follows [13]:

- For each word $i \in \{1, \dots, L\}$ in a sentence of length L , return its positional embedding vector w_i of length d_e , which is one of the hyperparameters. In [1] the base case is $d_e = 512$ while for example in GPT2 it is 768. The size of the embedding impacts the probability that two random embedded words are correlated or not (if their scalar product is close to ± 1 or to 0 respectively).
- Stack each embedding into a matrix $\mathbf{X} \in \mathbb{R}^{L \times d_e}$. For instance if we consider each English word as a token, the sentence 'I hit the blue ball' has $L = 5$ tokens, each of them encoded in a 512-dimensional embedding vector.
- Compute the query, key and value matrices \mathbf{Q} , \mathbf{K} and \mathbf{V} by projecting \mathbf{X} in the (subspace) representations of the packed queries, keys and values:

$$\mathbf{Q} := \mathbf{X}\mathbf{W}_q, \quad \mathbf{K} := \mathbf{X}\mathbf{W}_k, \quad \mathbf{V} := \mathbf{X}\mathbf{W}_v,$$

where \mathbf{W}_q , \mathbf{W}_k , \mathbf{W}_v are weights matrices of length $d_e \times d_w$ trained during the training process (randomly initialised) and their dimensions depend on the architecture. For simplicity, in one single-head attention $d_w = d_e$, while for multiple-head (8 in [1]), $d_w = d_e/N_{\text{head}} = 512/8 = 64$. Multiple-head attention algorithm splits multiple query, key and value 'heads' in order to improve the performance of the algorithm. This split reshapes weight matrices into \mathbf{W}_q^i , \mathbf{W}_k^i , \mathbf{W}_v^i , of dimensions $d_e \times d_w$. In the base case of [1], $d_w = 64$ was selected after hyperparameter tuning as was the length of the embedded vectors d_e . Larger values of the hyperparameters such as d_w or d_e provide better performance [1, Table 4], but require more computational resources.

- Compute a self-attention score of each word of the input sentence against the others, by taking the dot product between query and key vectors of the respective words. These scores are then normalised and passed into a softmax function. The resulting attention-weighting matrix of size $L \times L$ represents the correlation between word i and word j . Finally, the attention function or self-attention head takes the attention-weighting matrix and multiplies it by the value matrix \mathbf{V} , defined as

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax} \left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_w}} \right) \mathbf{V}, \quad (5)$$

resulting in an updated vector representation of the contextualised token \mathbf{V} .

The whole transformer itself uses h attention heads ($h = 8$ in [1]) and is summarised in Algorithm 1 to allow the model to jointly represent information from different representation subspaces at different positions. The reweighing “values” from the heads are passed through another densely connected layers. The outputs of these attention heads are concatenated into a matrix \mathbf{O} of dimensions $L \times d_e$. This matrix is multiplied by an output weight matrix \mathbf{W}_o of dimension $d_e \times d_e$. The result is $\tilde{\mathbf{X}}$, an updated representation of \mathbf{X} the input matrix of the embedded tokens with $\tilde{\mathbf{X}} \in \mathbb{R}^{L \times d_e}$.

Algorithm 1 Multiheaded Attention Layer

Require: $\mathbf{X}, \mathbf{W}_q, \mathbf{W}_k, \mathbf{W}_v, \mathbf{W}_o$
for $i = 1$ to h **do** (Split $\mathbf{W}_q, \mathbf{W}_k$, and \mathbf{W}_v into h heads $(\mathbf{Q}_i, \mathbf{K}_i, \mathbf{V}_i)_{i=1, \dots, h}$)
 $\mathbf{Q}_i \leftarrow \mathbf{XW}_q^i$; $\mathbf{K}_i \leftarrow \mathbf{XW}_k^i$; $\mathbf{V}_i \leftarrow \mathbf{XW}_v^i$
 output $_i \leftarrow$ Attention($\mathbf{Q}_i, \mathbf{K}_i, \mathbf{V}_i$);
 $\mathbf{O} \leftarrow$ Concatenate(output $_1, \dots, \text{output}_h$)
 $\tilde{\mathbf{X}} \leftarrow \mathbf{O}\mathbf{W}_o$
Ensure: $\tilde{\mathbf{X}}$

The transformer model has been very successful and is at the core of many machine learning applications: for NLP, the BERT family [1], in biology, the Generative Pre-trained Transformer (GPT) family to predict molecular structures [14], in computer vision, to extract image information [15]. Recent models differ from their predecessors primarily by their size (BERT has 310 million parameters, GPT3 has 175 billion [16], PALM has 540 billion). The other chief difference lies in which part of the encoder-decoder framework of the transformer architecture is leveraged. BERT uses the entire encoder-decoder while models like the GPT family use the only decoder. The classical BERT architecture is trained on two tasks [8]. The first is sentence entailment, where the input is one sentence and the model is asked to predict whether the question proposed as the output is logically entailed by the first or not. The second task is Masked Language Modelling, where a random 15% subset of tokens in the entire training corpus is replaced by a [MASK] token and the BERT model is asked to predict the most likely candidate replacement. Both these pre-training approaches yield a very powerful model that can be fine-tuned on a downstream task.

3. SEMANTIC TEXTUAL SIMILARITY SEARCH

To perform semantic similarity search, a classical approach would be to feed an input sentence or text to the BERT transformer network which produces contextualised word embeddings for all input tokens in the text. Then, via a pooling layer (such as mean-pooling), the average of the contextualised word embeddings would return a fixed-sized output representation vector. In the BERT-base model the dimension is 768; similarly to the previous discussion, larger values provides better accuracy but at the cost of computational memory and time. This sentence/text could be directly compared with another pair using for instance cosine similarity between the two vectors. However this would often yield poor sentence embedding compared to GloVe [9] and to expensive computational time. For instance in [9], computing the most similar sentence pairs from 1,000 sentences took 65 hours with BERT. We have also tried to perform some fine-tuning and further train BERT on our corpus and indeed found the computational time to be significantly larger and the semantic search score high even for non matching (policy, rule) pair.

To palliate this, SBERT (Sentence-BERT) uses a BERT architecture with siamese and triplet networks [17] that is able to derive semantically meaningful sentence embeddings. This adds a pooling operation to the output of BERT and fine-tune the latter by adding a training dataset with pairs of sentences. The cosine similarity between a query vector \mathbf{Q} and a policy vector \mathbf{P} can be defined as

$$\mathcal{S}_c(\mathbf{Q}, \mathbf{P}) := \frac{\mathbf{Q}^\top \mathbf{P}}{\|\mathbf{Q}\| \|\mathbf{P}\|}, \quad (6)$$

which converges to one as the similarity increases. This allows SBERT networks to be fine-tuned and to recognise the similarity of sentences. Different versions of loss functions can be used; for example, we later use the `all-MiniLM-L6-v2`, trained on a dataset of over one billion training pairs. These models can be freely accessed using Sentences Transformers (www.sbert.net) to perform text embeddings directly from a large collection of pre-trained models tuned for various tasks.

Our goal here is to link rules issued by an organisation (such as governments or financial regulators) to financial institutions policies (banks in particular). Financial rules are often updated and financial institutions are required to issue new policies addressing these changes. It is extremely challenging for large institutions to keep track of regular updates and to ensure their policies match the latest rules. Using NLP, we perform here semantic search between query sentences (rules) and answers (policies) written in a corpus as follows:

- Clean corpus and rules from meaningless symbols (bracket, hashtag, special character, ...).
- For each query, split the corpus into sentences or paragraphs, depending on the data and the query size. Sentence-by-sentence comparison similarity will be more precise than pooling many sentences into paragraphs.
- Use a pre-trained model such as BERT to tokenise each word of the sentence, and map sentences to embedded vectors, usually via mean-pooling of each embedded token. The size of any embedded token depends on the pre-trained model. One can also use pre-trained model such as `all-MiniLM-L6-v2` and embed the sentence directly using sentence-transformer library.
- Use cosine similarity (6) as a metric between an embedded rule query and the embedded policies to find the best match. By requiring that a match is positive above a threshold (say $\mathcal{S}_c > 0.7$), we can easily extract from a policy corpus which sentences match the rule (if any).

This powerful pipeline can however be improved. Using pre-trained models to encode our sentences to vectors may sound sensible, but they are trained on generic corpus and are not familiar with domain-specific targets—such as financial regulation—with high accuracy. One may instead completely re-train BERT on a domain-specific corpus, as done in [18]; this however requires huge amount of training data (3.3 billion tokens), which is extremely costly. Instead—and this is the state of the art—training can be performed using an Adaptive Pre-Training, described at www.sbert.net; this was successfully applied in various domains, such as for the bio-medical language [19]. It can then be archived with Sentence-Transformers frameworks and the Hugging Face Transformers library, specifically built for NLP applications huggingface.co.

3.1 Classical Domain Adaptation (DA) Pre-Training method

The DA Pre-Training method can be split into two steps:

- Step 1: further train a pre-trained model for words embeddings such as BERT, SBERT or all-MiniLM-L6-v2 using a context specific corpus.
- Step 2: fine-tune the resulting model based on an existing training dataset (such as paired sentences), similarly to SBERT but on the domain-specific dataset.

Before performing Step 1, one can also add up domain-specific tokens to the existing model, although given the limited amount of datasets at hand, this may not always be possible. To perform Step 1, known as Pre-Training on Target Domain, we use the Mask Language Modelling (MLM) approach [20], which masks a random fraction of tokens in a sentence and the transformer then tries to guess what is missing. Alternative approaches include Transformer-based Denoising AutoEncoder (TSDAE) [20]. While MLM outputs a token vector (mask word), the TSDAE encoder is fed with noisy sentences which the decoder uses to predict the full original sentences. To perform Step 2, we need to specify the loss function and the training set. Many generic training sets with labelled paired sentences are available (huggingface.co/datasets), but we do not use them since wish to target our fine-tuning on (rules, policies) pairs. For the loss function, we use the Multiple Negatives Ranking (MNR) Loss detailed in [21, Section 4] and reviewed below.

The idea of Multiple Negatives Ranking (MNR), introduced in [21], is to suggest responses to a question embedded in a vector \mathbf{X} maintaining some kind of memory throughout the course of a dialogue, which translates to finding the probability $\mathbb{P}(\mathbf{Y}|\mathbf{X})$, where \mathbf{Y} is the embedded vector answer. This probability is used to rank possible answers \mathbf{Y} given an input question \mathbf{X} . Bayes' theorem and the total law of probability indicate that this probability can be expressed as

$$\mathbb{P}(\mathbf{Y}|\mathbf{X}) = \frac{\mathbb{P}(\mathbf{X}, \mathbf{Y})}{\sum_k \mathbb{P}(\mathbf{X}, \mathbf{Y}_k)}, \quad (7)$$

and the joint probability $\mathbb{P}(\mathbf{X}, \mathbf{Y})$ is estimated using a neural network scoring function \mathcal{S} such that

$$\widehat{\mathbb{P}}(\mathbf{X}, \mathbf{Y}) := \exp \{ \mathcal{S}(\mathbf{X}, \mathbf{Y}) \}.$$

In practice the denominator in (7) (equal to $\mathbb{P}(\mathbf{X})$) is approximated by sampling K responses from the training corpus with $\widehat{\mathbb{P}}(\mathbf{X}) := \sum_{k=1}^K \mathbb{P}(\mathbf{X}, \mathbf{Y}_k)$, leading to the approximate probability used to train the neural network:

$$\widehat{\mathbb{P}}(\mathbf{Y}|\mathbf{X}) := \frac{\exp \{ \mathcal{S}(\mathbf{X}, \mathbf{Y}) \}}{\sum_{k=1}^K \exp \{ \mathcal{S}(\mathbf{X}, \mathbf{Y}_k) \}}.$$

Having a training dataset with (question, answer) pairs, each embedded question \mathbf{X}_i is paired with the embedded answer \mathbf{Y}_i and all other \mathbf{Y}_j for $i \neq j$ is treated as a negative candidate for question \mathbf{X}_i (so $K - 1$ negatives). The goal of the training is to minimise the loss function

$$\mathcal{J}(\mathbf{X}, \mathbf{Y}, \Theta) := -\frac{1}{K} \sum_{i=1}^K \log \widehat{\mathbb{P}}(\mathbf{Y}_i|\mathbf{X}_i),$$

where Θ gathers the word embeddings and the neural network parameters. To compute the score, the authors in [21] represent the input question tokenised into a word sequence \mathbf{X} and responses \mathbf{Y} as fixed-dimensional input features, extracting n -gram features from each. During training, they learn a d -dimensional embedding for each n -gram jointly with the other neural network parameters. To represent sequences of words, they combine n -gram embeddings by summing their values. This

bag of n -grams representation is denoted as $\Psi(\mathbf{X}) \in \mathbb{R}^d$. Then a feedforward scoring model takes the n -gram representation of a question and a response, and computes a score. For instance in the dot-product architecture from [21, FIGURE 3], $\Psi(\mathbf{X})$ and $\Psi(\mathbf{Y})$ go to two separate tanh activation layers returning the encoded \mathbf{X} and \mathbf{Y} as \mathbf{h}_x , \mathbf{h}_y and performing $\mathcal{S}(\mathbf{X}, \mathbf{Y}) = \mathbf{h}_x^\top \mathbf{h}_y$. Using the MNR Loss, the score is computed using at first the chosen pre-trained model and the default scoring is computed via cosine similarity in (6).⁴ The key issue and challenge for this work lies in the absence of a training dataset for policies and rules. While unsupervised text embedding learning performs rather poorly in learning on domain-specific concepts without fine tuning, we tested an approach to perform an Adaptive Pre-Training without using the training dataset we generated in Section 4.1. This method generates pseudo-labeling pairs for fine tuning and we describe it in the next section.

3.2 Domain Adaptation with Generative Pseudo-Labeling

Alternatively to the approach described above, we also tested the unsupervised fine-tuning with Generative Pseudo-Labeling (GPL) [22]. This method can be used to perform Step 2 from Section 3.1. In our context we wish to investigate how it improves Step 2 over a fine-tuning fitting method that uses labelled (rules, policies) pairs. The idea behind GPL can be condensed into four steps, that we now describe. Below, we shall use calligraphic letters to denote sets (of paragraphs for example) as opposed to standard letters for elements, for example $\mathcal{P}^- := \{P_1^-, P_2^-, \dots\}$.

3.2.1 GPL Step 1: Generating queries from a domain specific corpus

For each paragraph of the corpus, we generate queries using a T5-encoder-decoder model [23] similarly to the architecture of [1] discussed in Section 2.2.2. This query generator, which can easily be called again from Sentences-Transformer models, yields positive (query, answer) pairs and is summarised in Algorithm 2.

Algorithm 2 Generator

Require: input paragraph P^+ , vocabulary size V , length of the output query paragraph L_q

$G \leftarrow \text{T5}(P^+)$ ($G \in \mathbb{R}^{V \times L_q}$: matrix of probability distributions)

$Q \leftarrow (\arg \max(G_i))_{i=1, \dots, L_q}$

$\tilde{Q} \leftarrow \text{Decode}(Q)$

Ensure: (Q, P^+)

3.2.2 GPL Step 2: Negative mining using dense retrieval

This step finds sentences in the text that share many similar words. For each query or paragraph, there is now a positive pair. In addition, for each of the generated queries, GPL retrieves M negative passages that are similar but not a match to the query, which is known as a negative mining process. At the end of this step, we obtain triplets of positive and negative passages associated with a query. This step makes use of a function `DenseRetrieval` which is detailed in Algorithm 3.

⁴ see github.com/UKPLab/sentence-transformers/blob/master/sentence_transformers/losses/MultipleNegativesRankingLoss.py for details

Algorithm 3 Dense Retrieval

Require: Remaining set $\mathcal{P} = \{P_1, \dots, P_{|\mathcal{P}|}\} \setminus \{P^+\}$ of paragraphs; user query Q ; $M \leq |\mathcal{P}|$: number of negative passages;
 $\mathcal{P}^-, \mathcal{D} \leftarrow \{\}, \{\}$
 $E \leftarrow (T5(P_i))_{i=1, \dots, |\mathcal{P}|}$ (Calculate an embedding E for each paragraph)
 $E_Q \leftarrow T5(Q)$ (Calculate an embedding E_Q for from the query Q)
 $\mathcal{D} \leftarrow (\mathcal{S}_c(E_i, E_Q), E_i)_{i=1, \dots, |\mathcal{P}|}$ (with $\mathcal{S}_c(\cdot)$ defined in (6))
 $\mathcal{D} \leftarrow \text{sort}(\mathcal{D})$ (Sort by cosine similarity)
 $\mathcal{P}^- \leftarrow (D_1, \dots, D_M)$
Ensure: (Q, P^+, \mathcal{P}^-)

3.2.3 GPL Step 3: Pseudo labelling

Use a crossencoder [24] to score all query-positive passage and query-negative passages. Note that some negative passage might actually be positive pairs. This step scores the triplets composed of our initial paragraph P^+ , our generated query based on the initial paragraph Q , one output P_i^- from our set of negative samples \mathcal{P}^- .

3.2.4 GPL Step 4: Training/Tuning the Transformer model to identify the differences between positive and negative passages, using a crossencoder model

The latter compare the embeddings of the passages by generating similarity scores for both positive and negative pairs. Given a transformer model that we wish to train using GPL and the output tuples from the previous step, we will train the said transformer by backpropagation between the margin computed in Algorithm 4 and the predicted margin we extract. This process is optimised using an MSE loss function applied to the margin

$$S_{\text{MSE}} := \left| X_Q^T (X_{P^+} - X_{P_i^-}) \right|, \quad (8)$$

with $X_Q, X_{P^+}, X_{P_i^-}$ the embedded vectors of the query, positive passage, and negative passage i .

Algorithm 4 Pseudo Labelling

Require: (Q, P^+, \mathcal{P}^-) , crossencoder, M
for $i = 1$ to M **do**
margin $_i \leftarrow \text{crossencoder}(Q, P^+) - \text{crossencoder}(Q, P_i^-)$
Ensure: $(Q, P^+, P_i^-, \text{margin}_i)_{i=1, \dots, M}$

3.3 To summarise

Both supervised and unsupervised domain adaptation methods use a pre-trained model and further train it by comparing (query, answer) pairs or triplets (negative answer). These pairs and triplets are in the classical supervised method given by the dataset used to perform the domain adaptation (corpus of a specific domain), but it may also be generated. That is the GPL approach for a fully

Algorithm 5 Training loop of Transformer using GPL

Require: $(Q, P^+, P_i^-, \text{margin}_i)_{i=1, \dots, M}$ **for** $i = 1$ to M **do** $(X_Q, X_{P^+}, X_{P_i^-}) \leftarrow \text{Transformer}(Q, P^+, P_i^-)$ $\text{predicted}_{\text{sim}^+} \leftarrow \mathcal{S}_c(X_Q, X_{P^+})$ (with $\mathcal{S}_c(\cdot)$ defined in (6)) $\text{predicted}_{\text{sim}^-} \leftarrow \mathcal{S}_c(X_Q, X_{P_i^-})$ $\text{predicted}_{\text{margin}_i} \leftarrow \text{predicted}_{\text{sim}^+} - \text{predicted}_{\text{sim}^-}$ $\text{Loss}_i \leftarrow \mathcal{S}_{\text{MSE}}(\text{margin}_i, \text{predicted}_{\text{margin}_i})$ defined in (8)) $\text{Output}_i \leftarrow \text{backpropagate}(\text{Loss}_i)$ **Ensure:** Output

unsupervised domain adaptation. Below, we explore the results of performing an Adaptive Pre-Training method and a purely unsupervised approach, using GPL. To compare our results and to perform supervised fine-tuning we need a pseudo training/validation dataset. We describe in the next section how we create it.

4. APPLICATION TO SEMANTIC MATCHING FOR FINANCIAL REGULATION

4.1 Creation of a Pseudo Training and Validation Dataset

To create our validation dataset we use a mean assemble average of N pre-trained models (in what follows $N=10$)⁵ for the task of semantic search between rules and policies. We run these N sentence-transforming models on a catalogue of rules and financial policies provided by FinregE (<https://finreg-e.com>), a company providing clients in financial services with a software focusing on current and future regulations, using ML and AI tools to identify and interpret regulatory requirements and to integrate compliance workflows for action and compliance management.

The dataset of rules is publicly available from Financial Conduct Authority (FCA) Rulebook 2022. After cleaning and splitting the rules into sentences of length 200, we obtain about 50,000 sentences, composed of 21,914 rule IDs. The policies' dataset was obtained from FinregE and is composed of 2,374 policies that may not cover all of the newly edited FCAR rules. Our approach is the following:

- For each model, we keep the rule/policy match sentences if the cosine similarity (6) is above 0.7. This is an empirical choice that provides fair matches.
- We keep pairs that have been identified by a number of models greater than \sqrt{N} . Given that the standard deviation of shot noise is equal to the square root of the average number of events N , hence if models are uncorrelated, selecting a signal to noise match above one sigma should correspond to a \sqrt{N} cut.

⁵ multi-qa-mpnet-base-cos-v1, sentence-t5-xl, multi-qa-distilbert-cos-v1, msmarco-bert-base-dot-v5, all-distilroberta-v1, all-MiniLM-L12-v2, distiluse-base-multilingual-cased-v2, all-mpnet-base-v2, stsb-distilbert-base, bert-base-nli-mean-tokens. These models can be found at www.sbert.net/docs/pretrained_models.html

We end up with 1,760 matching (rule, policy) pairs out of which we keep 1,408 pairs to fine-tune our model (Section 4.2) and 352 for the validation dataset. In FIGURE 1 we can see a sample of our validation dataset results. This dataset will be good to identify pairs that share high similarities but will not be able to pick up on subtle pairs. This is why it is not an unbiased dataset for fine-tuning purposes but we proceed anyway due to the absence of hand labelled pairing between rules and policies.

rule txt	policy txt	Score2
a publicly available research report accompanying a change compared to the previous sovereign rating or related rating outlook shall include at least the following a detailed evaluation of the changes to the quantitative assumption justifying the reasons for the rating change and their relative weight	additionally when the rating determination involves either a change from the current rating review and or outlook status or an affirmation of a rating or a confirmation of a rating or the assignment of an initial rating the accompanying publicly disseminated information shall include a detailed evaluation of the quantitative and qualitative factors and their relative weighting that were taken into account when making the rating change determination	90.00
a credit rating agency shall ensure that fees charged to its clients for the provision of credit rating and ancillary services are not discriminatory and are based on actual costs	price discrimination under the eu regulations relating to credit rating agencies fees charged by credit rating agencies must be nondiscriminatory and based on costs	80.00
a credit rating agency shall state clearly and prominently when disclosing credit ratings or rating outlooks any attributes and limitations of the credit rating or rating outlook	these elements should be addressed in the rating rationale or similar section of the credit rating announcement	80.00

Figure 1: Sample of the semantic text search output between rules and policies. The score corresponds to the percentage of models that have identified the same match.

4.2 Domain Adaptation Pre-Training method

The corpus we use for further pre-training (Section 3.1-Step 1) are rules from the FCA Rulebook. We use a full word mask and we use the `DataCollatorForLanguageModeling` which is a class of transformers to perform the MLM training⁶. Then we run Step 2 of Section 3.1 with the MNR loss using our training dataset matching (rule, policy) pairs. We perform Step 1 and Step 2 using both BERT and `all-MiniLM-L6-v2` pre-train model to further train in our domain-specific corpus.

4.2.1 Results

In the absence of an unbiased validation dataset, it is difficult to perform proper model comparison. We do so by rating two scores.

- Score 1: margin between matching and random pairs. For all the rules in our validation datasets we use the model considered to compute the cosine similarity $\mathcal{S}_c(R_i, P_i)$ between rule i and policy i (the matching pair) and the cosine similarity $\mathcal{S}_c(R_i, P_j)$ between rule i and a random policy j , with $j \neq i$. Then we compute

$$\mathcal{S}_1 := \frac{1}{N} \sum_{i=1}^N \left\{ \mathcal{S}_c(R_i, P_i) - \mathcal{S}_c(R_i, P_j) \right\}, \quad (9)$$

where N is the length of the validation dataset (here $N = 352$).

⁶ see huggingface.co/docs/transformers/main_classes/data_collator and www.sbert.net/examples/unsupervised_learning/MLM/README.html

- **Score 2:** For each rule R_i , we compute the highest similarity scoring policy $P_i^* := \arg \max_{j=1, \dots, N} \mathcal{S}_c(R_i, P_j)$. If $P_i^* = P_i$, we then increment \mathcal{S}_2 by one and then divide by N to get the fraction of ‘correct’ matches per model, namely

$$\mathcal{S}_2 := \frac{1}{N} \sum_{i=1}^N \mathbf{1}_{\{P_i = \arg \max_{j=1, \dots, N} \mathcal{S}_c(R_i, P_j)\}}.$$

For a perfectly tuned model and for an unambiguous validation dataset, \mathcal{S}_1 and \mathcal{S}_2 should be close to 1 since $\mathcal{S}_c(R_i, P_i)$ approaches 1 and $\mathcal{S}_c(R_i, P_j)$ approaches zero. In practice, since we split rules and policies in sentences, it may happen that the same sentence occurs in different policies. Therefore a match between R_i and P_j (with $j \neq i$) is not necessarily a wrong match. In addition in each rule, some sentences are quite generic and are addressed by many policies. Our score does not account for these and is therefore just a benchmark to test the impact of the domain adaptation and not an actual score of the quality of semantic matching. We also note that the margin score (Score 1) is a relatively noisy metric as only one randomly-selected other policy serves as the negative for the purpose of scoring. However, we find that on average, the cosine similarity of the matching pair is higher than between any non-matching pair, as such, by averaging the margins across our sample as in Equation ??, we argue \mathcal{S}_1 is a good benchmark indicator to illustrate the impact of domain adaptation in conjunction with \mathcal{S}_2

Using BERT as the initial pre-trained model our results are shown in TABLE 1.

Table 1: Result for the Domain adaptation with BERT and steps defined in Section 3.1

Model	Score 1	Score 2
Regular BERT	0.09	0.32
BERT + Step 1	0.10	0.33
BERT + Step 1 & 2	0.32	0.33

The key improvement of the Domain Adaptation comes from fine-tuning Step 2 although we know that regular BERT is not really suited to perform semantic matching. In fact the low **Score 1** for regular BERT comes from the fact that both positive and negative matches are high. Only fine-tuning can help BERT distinguish the negative pairing. That is why we also tested the all-MiniLM-L6-v2 model using it as is, performed Step 1 and Step 2, and quote the improvement of performing an adaptive pre-training method defined as $\frac{\text{score fine-tune} - \text{score baseline}}{\text{score baseline}}$ TABLE 2.

Table 2: Result for the Domain adaptation with all-MiniLM-L6-v2 and steps defined in Section 3.1

Model	Score 1	Score 2
all-MiniLM-L6-v2	0.21	0.46
all-MiniLM-L6-v2 + Step 1 & 2	0.27	0.56
Improvement	29%	22%

This shows again the improvement over the baseline. In addition we tested the impact of the quality of the training/validation dataset by decreasing the score threshold from 0.7 down to 0.6 when generating the training/validation dataset in Section 4.1 and found that it decreases both scores by

10%. Finally we find that the fraction of masked words in the Mask Language Modelling does not impact significantly the results. By default the fraction is set to 0.15 and we tried 0.2 in the BERT Step 1 and did not find a change in the scoring. Overall, using limited data for Step 1 and Step 2 we find that the key effect comes from the fine-tuning as was previously stated in [9].

4.3 Fully Unsupervised Training With the Gpl Method

This time we perform fine tuning solely by running the GPL method on the FCA Rulebook corpus and train the GPL method on 10,000 and 20,000 triplets. To do so we follow the steps described in github.com/UKPLab/gpl using the default pre-trained model MSMARCO. We output samples at each step to check the pseudo-labelling. On Step 1, to generate queries from a passage of the corpus to get positive pairs, such as

Example 1 Selected passage: *'The investor providing the capital may choose not to be involved in the running of the venture'*.

Generated queries:

- *how are investors involved in a venture*
- *who is involved in the running of the venture*
- *what are the investors of a venture*

Example 2 Selected passage: *'If a circular submitted for approval is amended a copy of amended drafts must be resubmitted marked to show changes made to conform with FCA comments and to indicate other changes'*.

Generated queries:

- *how to amend FCA circular*
- *what must FCA amendments show*
- *how to revise circular for approval*

Clearly some of these positive pairs can be noisy. A sample of our generated training datasets containing our triplets of query, positive, negative passages is displayed in FIGURE 2.

4.3.1 Results

Similarly to the scoring analysis in Section 4.2, we see that the GPL fine-tuning does not impact the scores as much as the adaptative pre-training method. The scorings themselves are again to be considered as a benchmark from the no-fine-tuning scoring which translate into an improvement of

Query	Passage	Negative Mining	Margin
who accesses financial promotion	At the same time it provides very little control over who is able to access the financial promotion	A firm may also wish to approve a financial promotion that it communicates itself	6.398293
what should an information referred to in a document be	The information referred to in should be provided in plain and intelligible language on paper or in another durable medium	This will involve advice not just information	15.899925
what can fscs recover in error	If a relevant unused levy limit in the table in is exceeded by a subsequent levy imposed on the corresponding funding class by the FSCS under the PRA s or FCA s rules in the same financial year the FSCS must recover any previous contributions by the relevant category or class in the way set out in but only to the extent necessary to correct that relevant unused levy limit excess	The FSCS may postpone payment of compensation only in strictly limited circumstances	6.397892
what regulates a trustee change	b The appropriate regulator should be consulted about or pre notified of a potential change of trustee	Entering into investment transactions concerning securities or relevant investments on behalf of the trustees	12.635609

Figure 2: Sample of the training dataset generated by the GPL approach from the FCA Rulebook. The Margin is defined in (8). This data is used to fine-tune the msmarco pre-trained model.

the ratio $\frac{\text{score fine-tune} - \text{score baseline}}{\text{score baseline}}$. As expected, GPL fine-tuning is not as efficient as mapping rules to policies, most likely because it did not perform the training over a set of policies but solely on rules. We also find that performing the training on 10,000 or 20,000 triplets does not impact the results by more than a percent TABLE 3.

Table 3: Result for the GPL fine-tuning on the pre-trained MSMARCO model.

Model	Score 1	Score 2
msmarco	0.14	0.79
msmarco + GPL	0.15	0.84
Improvement	7%	6%

5. DISCUSSION AND COMPARISON BETWEEN GPL AND ADAPTATIVE PRE-TRAINING

When looking at the validation dataset for the rule and policy match as well as the policy identifies by the Adaptative pre-training (DA) and the GPL fine-tuning, there is no clear way to identify which pairing is the more accurate. In fact they all return meaningful pairings. In FIGURE 3 we show a sample where either the DA or the GPL policy match does not agree with the validation policy. It seems that the validation policy is not always the best match to address the rule. In addition, not all the rules have been addressed in our policy sample. Therefore it is not possible to conclude which approach performs the best matching to our rules. The three approaches we investigate in this work are:

- Mean ensemble of N pre-trained models selecting policies if they are returned by at least \sqrt{N} models with cosine similarity > 0.7 . This is our validation policy defined in Section 4.1.
- Model from an Adaptative pre-training (DA), defined in Section 3.1
- Model from GPL fully unsupervised training from rules corpus in Section 4.3.

rule	policy validation	policy DA	policy GPL
With a view to ensuring transparency, disclosure of any material modification to the methodologies and practices, procedures and processes of credit rating agencies should be made prior to their coming into effect, unless extreme market conditions require an immediate change in the credit rating.	The transparency of both Methodologies and Credit Rating decisions with Methodologies being published and Credit Rating Announcements disclosing which Methodology has been applied as well as the publication of research on Credit Ratings supports consistent application by creating a public forum in which feedback from issuers and investors can uncover potentially inconsistent application of Methodologies.	The precise format and information content requirements differ between rating groups and will generally reflect the relevant Credit Rating Methodology, with the relevant Analyst responsible for compiling, performing or directing the necessary analysis and performing any quality checks required.	The transparency of both Methodologies and Credit Rating decisions with Methodologies being published and Credit Rating Announcements disclosing which Methodology has been applied as well as the publication of research on Credit Ratings supports consistent application by creating a public forum in which feedback from issuers and investors can uncover potentially inconsistent application of Methodologies.
The names, job titles and contact details for the lead rating analyst and the name and position of the person primarily responsible for approving the credit rating together with the name and address of the legal entity responsible for the credit rating.	a. a clear and prominent disclosure of the name and title of the lead Analyst and person primarily responsible for approving the Credit Rating; b. an indication of when the Credit Rating was first distributed and when it was last updated; and c. the extent to which MIS Singapore has examined the quality of information used in the Credit Rating process and whether it is satisfied with the quality of information on which it bases its Credit Rating.	The names, job titles and contact details for the lead rating analyst and the name and position of the person primarily responsible for approving the credit rating together with the name and address of the legal entity responsible for the credit rating Requirement currently being met.	The names, job titles and contact details for the lead rating analyst and the name and position of the person primarily responsible for approving the credit rating together with the name and address of the legal entity responsible for the credit rating Requirement currently being met.
(2) A credit rating agency shall list and provide a detailed explanation of the following points with regard to the credit rating methodologies used regarding: (a) each qualitative factor, including the scope of qualitative judgement for that factor; (b) each quantitative factor, including key variables, data sources, key assumptions, modelling and quantitative techniques.	The precise format and information content requirements differ between rating groups and will generally reflect the relevant Credit Rating Methodology, with the relevant Analyst responsible for compiling, performing or directing the necessary analysis and performing any quality checks required.	This document provides the procedures to be applied to ensure that the technical specifications of Credit Rating Models (CRMs), Credit Rating Scorecards (CRSs) and Other Credit Rating Quantitative Tools (OCRQTs) -collectively, the Credit Rating Quantitative Tools (CRQTs)- used in the Credit Rating process are consistent with the requirements of the Credit Rating Methodologies (Methodologies) with which they are associated in accordance with the Credit Rating Quantitative Tool Technical Specifications Review - Policy.	However, as appropriate, select those principal methodologies that apply to the subject credit rating.

Figure 3: Sample comparison of pairing based on cosine similarity from our pseudo validation dataset where we add the best match of the DA and GPL approaches.

It is likely that a combination of these three approaches should provide the best pairing. Overall it is quite remarkable that even in the absence of a training dataset, we manage to perform an Adaptive pre-training that improves the scores. This can be helpful as an alternative or a complement to the GPL approach.

6. DATA AND CODES

The codes used for fine tuning can be found at https://github.com/IxandraAchitouv/NLP_semantic_matching. The data are not public. Please contact Finreg-E (<https://finreg-e.com>) to access the data.

7. ACKNOWLEDGEMENTS

The authors would like to thank Rohini Gupta and Amit Madhar from Finreg-E (<https://finreg-e.com>), without whom this project would not have been possible. IA and AJ are supported by the Innovate UK Smart Grant ‘*Finreg-E / Natural Language Processing for Financial Regulation*’. DG is supported by the ESRC Grand Union Doctoral Training Partnership Grant and the Oxford Man Institute of Quantitative Finance.

References

- [1] Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, et al. Attention Is All You Need. *Adv Neurol.* 2017:30.
- [2] Firsth JR. *Studies in Linguistic Analysis.* Oxford: Blackwell Publishing; 1957.
- [3] Deerwester S, Dumais ST, Furnas GW, Landauer TK, Harshman R. Indexing by Latent Semantic Analysis. *J Am Soc Inf Sci.* 1990;41:391-407.
- [4] Lund K, Burgess C. Producing High-Dimensional Semantic Spaces From Lexical Co-occurrence. *Behav Res Methods Instrum Comput.* 1996;28:203-208.
- [5] Rohde DL, Gonnerman LM, Plaut DC. An Improved Model of Semantic Similarity Based on Lexical Co-occurrence. *Commun ACM* 8. 2006;116:627-633.
- [6] Mikolov T, Chen K, Corrado G, Dean J. Efficient Estimation of Word Representations in Vector Space; 2013. Available from: arXiv:1301.3781.
- [7] Pennington J, Socher R, Manning CD. Glove: Global Vectors for Word Representation. In: *Proceedings of the 2014 conference on empirical methods in NLP*; 2014:1532-1543.
- [8] Devlin J, Chang MW, Lee K, Toutanova K. Bert: Pretraining of Deep Bidirectional Transformers for Language Understanding. In: *Proceedings of the 2019 conference of the north American Chapter of the Association for Computational Linguistics: Human Language Technologies.* Vol. 1; 2018:4171-4186.
- [9] Reimers N, Gurevych I. Sentence-Bert: Sentence Embeddings Using Siamese Bertnetworks; 2019. Available from: arXiv:1908.10084.
- [10] Rosenfeld R. A Maximum Entropy Approach to Adaptive Statistical Language Modelling. *Comput Speech Lang.* 1996;10:187-228.
- [11] Guthrie D, Allison B, Liu W, Guthrie L, Wilks Y. A Closer Look at Skip-Gram Modelling. In: *Proceedings of the Fifth International Conference on Language Resources and Evaluation.* Vol. 6; 2006:1222-1225.
- [12] Mikolov T, Sutskever I, Chen K, Corrado GS, Dean J. Distributed Representations of Words and Phrases and Their Compositionality. *Adv Neurol.* 2013:26.
- [13] Phuong M, Hutter M. Formal Algorithms for Transformers; 2022. Available from: arXiv:2207.09238.
- [14] Maziarka L, Danel T, Mucha S, Rataj K, Tabor J, Jastrzebski S. Molecule Attention Transformer; 2002. Available from: arXiv:2002.08264.
- [15] Dosovitskiy A, Beyer L, Kolesnikov A, Weissenborn D, Zhai X, Unterthiner T, et al. An Image Is Worth. 2010;16×16 Words: Transformers for Image Recognition at Scale. Available from: arXiv:2010.11929.
- [16] Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P et al. Language Models Are Few-Shot Learners. *Adv Neurol.* 1877-1901;33.

- [17] Schroff F, Kalenichenko D, Philbin J. Facenet: A Unified Embedding for Face Recognition and Clustering. In: Proceedings of the Ieee Conference on Computer Vision and Pattern Recognition; 2015:815-823.
- [18] Beltagy I, Lo K, Cohan A. Scibert: A Pretrained Language Model for Scientific Text. In: Proceedings of the 2019 Conference on Empirical Methods in NLP; 2019:3613-3618.
- [19] Lee J, Yoon W, Kim S, Kim D, Kim S, So CH et al. Biobert: A Pre-trained Biomedical Language Representation Model for Biomedical Text Mining. *Bioinformatics*. 2020;36:1234-1240.
- [20] Wang K, Reimers N, Gurevych I. Tsdæ: Using Transformer-Based Sequential Denoising Auto-Encoder for Unsupervised Sentence Embedding Learning. In: Findings of the Association for Computational Linguistics, pp. 671-88 2021.
- [21] Henderson M, Al-Rfou R, Strophe B, Sung YH, Lukács L, Guo R et al. Efficient Natural Language Response Suggestion for Smart Reply; 2017. Available from: [arXiv:1705.00652](https://arxiv.org/abs/1705.00652).
- [22] Wang K, Thakur N, Reimers N, Gurevych I. Gpl: Generative Pseudo Labeling for Unsupervised Domain Adaptation of Dense Retrieval. In: Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies; 2022:2345-2360.
- [23] Raffel C, Shazeer N, Roberts A, Lee K, Narang S, et al. Exploring the Limits of Transfer Learning With a Unified Text-To-Text Transformer. *J Mach Learn Res*. 2020;21:5485-5551.
- [24] Humeau S, Shuster K, Lachaux MA, Weston J. Poly-Encoders: Transformer Architectures and Pretraining Strategies for Fast and Accurate Multi-Sentence Scoring; 2019. Available from: [arXiv:1905.01969](https://arxiv.org/abs/1905.01969).