
BERTaú: Itaú BERT for digital customer service

Paulo Finardi **José Dié Viegas** **Gustavo T. Ferreira** **Alex F. Mansano**
paulo.finardi jose.barros-viegas gustavo.tino-ferreira alex.mansano

Vinicius F. Caridá
vinicius.carida@itau-unibanco.com.br

MaLS Data Science Team - Digital Customer Service
Itaú Unibanco - São Paulo, Brazil

Abstract

In the last few years, three major topics received increased interest: deep learning, NLP and conversational agents. Bringing these three topics together to create an amazing digital customer experience and indeed deploy in production and solve real-world problems is something innovative and disruptive. We introduce a new Portuguese financial domain language representation model called BERTaú. BERTaú is an uncased BERT-base trained from scratch with data from the Itaú virtual assistant chatbot solution. The novelty of this contribution lies in that BERTaú pretrained language model requires less data, reaches state-of-the-art performance in three NLP tasks, and generates a smaller and lighter model that makes the deployment feasible. We developed three tasks to validate our model: information retrieval with Frequently Asked Questions (FAQ) from Itaú bank, sentiment analysis from our virtual assistant data, and a NER solution. All proposed tasks are real-world solutions in production on our environment and the usage of a specialist model proved to be effective when compared to Google BERT multilingual and the Facebook’s DPRQuestionEncoder, available at Hugging Face. BERTaú improves the performance in 22% of FAQ Retrieval MRR metric, 2.1% in Sentiment Analysis F_1 score, 4.4% in NER F_1 score. It can also represent the same sequence in up to 66% fewer tokens when compared to "shelf models".

1 Introduction

In recent years, creating and managing digital customer experiences has appeared as to be a key area for many companies on “leveraging digital advancement for the growth of organizations and achieving sustained commercial success” (1). The idea of interacting with computers through natural language dates back to the 1960s, but recent technological advances have led to a renewed interest in conversational agents such as chatbots and digital assistants. In the customer service context, conversational agents promise to create a fast, convenient, and cost-effective channel for communicating with customers (2).

In the area of digital customer service, there is a growing demand for assertiveness and specialization, since a more assertive service allows the customer to be served more quickly, thus increasing the efficiency of the entire system. In this context, three main standard NLP (natural language processing) tasks stand out, as follows: 1- Named Recognition Entity (NER): recognizing the entities during the conversation is fundamental to understanding customer’s demands. 2- Information Retrieval (IR) with Frequently Asked Questions (FAQ): once the demand is understood, it is necessary to present the most relevant information. 3- Sentiment Analysis (SA): managing customer sentiment/satisfaction has become crucial(3).

Machine learning has received increased interest both as an academic research field and as a solution for real-world business problems. In benchmarks of NLP tasks, such as squad(4) and glue(5), machine learning models can achieve better performance than humans. The BERT(6) algorithm and variants are considered state-of-the-art at solving NLP tasks. However, the deployment of machine learning models in production systems can present several issues and concerns (7).

Given the context described so far, the goal of this paper is to test these hypothesized advantages of using and fine-tuning pretrained language models for a Brazilian Portuguese financial domain. To address these problems, we developed BERTaú, a BERT base uncased pretrained with data from AVI (Itaú Virtual Assistant). Given that the raw text is at the core of digital service data, it is necessary to use increasingly robust models to achieve the best possible digital service.

Our choice for using BERT (6) is based on the idea that it is an encoder model that is powerful and widely established in the Natural Language Processing (NLP) field and programming libraries, such as Hugging Face Transformers (8). Also, Open NeuralNetwork Exchange (ONNX) (9) has presented good deployment solutions and tools for reducing inference time, which allows BERTaú to be deployed.

The key point of this research is the data. For better results, it is mandatory to have a big and high-quality dataset - a large number of tokens with good semantic text. From the fact that the AVI has an average of 2 million monthly digital visits, few months of data would be necessary to enable it to train the model from scratch. However, we used a larger data window to capture semantic text before and after the pandemic – approximately 18 months of data. Further details about the data and dataset is described in section 3.

Once the model has been trained, we validate the model in three main tasks in our environment: Information Retrieval (IR) with Frequently Asked Questions (FAQ), Sentiment Analysis (SA) on phrases about our service, with three classes: positive, neutral and negative, and Named Recognition Entity (NER), which recognizes entities during the conversation between the customer and AVI.

The article is organized as follows: in section 3 we describe the process of the training from scratch; in section 2 we detail the baselines and related work; in section 4 we describe the setup for our experiments; section 4 presents the results; and finally, in section 5 we describe our evaluation results and conclusions.

2 Baselines and Related Work

Although the BERT model is open source and offers pretrained weights in English language and has a multilingual version available (mBERT), many studies have been developed to create specialist BERTs. A natural advantage of having a model pretrained in a specific vocabulary domain is that it can represent sequences using fewer tokens and performs the training stage with more computational efficiency. Table 1 directly compares the sequence lengths between our model and the baselines in NER and Information Retrieval tasks. Besides short token representations, many specialized models have reached the state of the art (sota) when trained from scratch, the BERTimbau¹ (10) is sota on NER task at public Portuguese MiniHAREM dataset. In the financial domain, the FinBERT (11) achieves sota on financial sentiment analysis.

Table 1: Comparison of token lengths between models, where the models in FAQ dataset vary between [15%, 66%] of sequences sizes greater than BERTaú and in the NER dataset between [20%, 30%].

Model	FAQ _{dataset}	NER _{dataset}
DPRQuestionEncoder	130 tokens	-
mBERT uncased	97 tokens	13 tokens
BERTimbau - Portuguese BERT	90 tokens	12 tokens
BERTaú	78 tokens	10 tokens

In the NER task we compared our model with BERTimbau base and mBERT uncased. In Information Retrieval task we conducted experiments where we compared our model with BM25+ algorithm (12),

¹BERTimbau results at: <https://github.com/neuralmind-ai/portuguese-bert>

Sentence Transformer (SBERT) - `distiluse-base-multilingual` (13), Dense Passage Retrieval (DPR) QuestionEncoder (14) and the mBERT model (6).

2.1 BM25+

The BM25 is a probabilistic model meant to estimate the relevance of a document based on the idea that the query terms have different distributions in relevant and non-relevant documents. BM25+ has a better way of scoring long documents with the addition of a parameter δ that has a default value of 1. The formula is

$$score(D, Q) = \sum_{i=1}^n IDF(Q_i) \left[\frac{TF(Q_i, D) \cdot (k_1 + 1)}{TF(Q_i, D) + k_1 \cdot \left(1 - b + b \cdot \frac{|D|}{avgdl}\right)} + \delta \right], \quad (1)$$

where Q_i is term frequency in the document D , $|D|$ is the length of the document D in words, and `avgdl` is the average document length in the text collection. k_1 and b are free parameters. More details about BM25+ are shown in (12) section 3.3.

2.2 SBERT - `distiluse-base-multilingual`

We use the `distiluse-base-multilingual`² model proposed by (13) and measure the similarity between the question and answer by cosine distance. The `distiluse` model supports 50+ languages including Portuguese.

2.3 Dense Passage Retrieval (DPR) QuestionEncoder

The DPR (14) has the same structure as BERT and tackles the problem with the objective of getting better representations of dense embeddings. This model was trained in pairs of questions and answers. The authors also argue that the best way to obtain better representations of dense embeddings for Information Retrieval task is by maximizing the inner products of the question and vectors of relevant answers in a batch.

2.4 mBERT uncased

The best-known and was used as a baseline for Information Retrieval, NER and Sentiment Analysis tasks, this model was trained on the XNLI: Cross-Lingual NLI corpus (15) with 102 languages.

3 BERTaú From Scratch

In this work, we chose to train BERT uncased, since the vast majority of the dataset that constitutes AVI conversations are in this format. We have 14,5GB of data (22,500,000 words). Each line in the data is an AVI session, that is: A complete iteration between customer and chatbot. For obvious reasons, we swapped all numbers in the dataset with random numbers, in order to avoid any possibility of sensitive information, document number, or monetary values being exposed. Regarding the training configuration, we follow the guidelines of the BERT article in a straight line, with a maximum sequence length of 512 tokens and batch size of 256 sequences. The model was trained for 1,000,000 steps with a learning rate of 5e-5 and a warm up of 20,000 steps. The accuracy and loss at the end of training is shown in Table 2.

The dataset for the vocabulary creation was also extracted from the AVI, but with a different period³. This dataset has 2GB and 34100 tokens. We create our vocabulary using the Hugging Face tokenizer `BertWordPieceTokenizer`, performing NFKC text normalization, and setting `strip_accents` to `false`.

²Model available at: <https://www.sbert.net/examples/training/multilingual>

³without interception to the data used in BERT training

Table 2: Evaluation training accuracy and loss of Masked Language Model (MLM) and Next Sentence Prediction (NSP) tasks.

Eval training on step 1,000,000	Accuracy	Loss
Masked Language Model	0.785	0.891
Next Sentence Prediction	0.863	0.311

4 Experiments

Another feature of BERT is that it is a versatile model. The structure of the Transformer encoder allows the model to perform different tasks with few adaptations. Given its versatility, we perform three different experiments with BERTaú: 1. FAQ retrieval; 2. Sentiment Analysis (SA); 3. Named Entity Recognition (NER).

For all experiments, we used Nvidia GPU’s with half-precision FP16, which allows for less memory use during the training phase. Using FP16 with the same batch size, we were able to perform experiments twice as fast when compared to the FP32 full precision. A minor drawback of FP16 is its limited range: large numbers can explode and small numbers are truncated to zero. To avoid this behavior it is common to scale the objective function by a large number in the gradients. If this number is so large that it generates an overflow, the update is rejected and a new scale factor is obtained automatically with PyTorch Automatic Mixed Precision.

4.1 FAQ Retrieval

The task of FAQ retrieval can be described as follows: *Given a question and a set of candidate answers we want to identify which are the correct answers / true labels.* This is a classic problem in the Information Retrieval field and many solutions have been developed in the last 3 years. A complete survey with this kind of solutions can be found here (16). Regarding metrics, we use the Reciprocal Rank (RR). The RR of a single query is given by $\frac{1}{rank_K}$, where rank is a descending ordered list of size K : where answers with greater probabilities appear at the beginning of the list. The Reciprocal Rank is meant to answer: "Where is the first relevant answer on the list?", for example: if the relevant answer gets the seventh place, the RR for this query is $1/7$ and if no answer is found, the RR is 0. To calculate the RR for multiple queries we use the MRR (Mean Reciprocal Rank) metric, which is the RR average and given by the formula:

$$MRR = \frac{1}{Q} \sum_{j=1}^Q \frac{1}{rank_j} \quad (2)$$

We also measure the experiments with Average Precision AP@1 in the first position, which determines whether the first answer on the list is correct or not. The dataset was obtained from the FAQ Itaú bank⁴. Our data has 1427 questions with 2118 answers and some questions have more than one correct answer, i.e. the question \mathbf{Q}_1 can be an answer of $\{\mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_n\}$ $n \leq 5$ and the average sequence length \mathbf{QA} pair is 192 tokens.

Our approach for this problem is pretty simple: we tackle the problem as a binary classification problem, where the samples in the dataset are triples $\{(\mathbf{Q}_1, \mathbf{C}_1, \mathbf{A}_1), (\mathbf{Q}_2, \mathbf{C}_2, \mathbf{A}_2), \dots, (\mathbf{Q}_n, \mathbf{C}_n, \mathbf{A}_n)\}$, where \mathbf{C}_k with $k = 1, 2, \dots, n$ is a list of candidates answers chosen randomly. For each triple $(\mathbf{Q}_k, \mathbf{C}_k, \mathbf{A}_k)$ we create M training samples, where $M \in \{15, 30, 45\}$ and this method follows the same idea introduced in (17). After training the binary classifier, we rank the answers according to the Algorithm 1. For the same dataset and ranking proposal we use two approaches: The pointwise and pairwise (18).

4.1.1 Pointwise - Label Smoothing

Our pointwise method uses label smoothing (19) as an objective function to regularize the neural network by penalizing confident output distribution. The label smoothing mitigates the symptom of overfitting by penalizing output distributions with low entropy. Low entropy occurs when the network

⁴A sample set of FAQs used to build the dataset can be found here: <https://www.itaubank.com.br/contas/conta-corrente/>

Algorithm 1 Ranking FAQ

```
1: Input: logits from output model
2: Output: dict  $q_{id} \rightarrow \text{list}(\text{predict\_rank})$ 
3: for cand in cands do
4:   logits = output-model
5:   pred_values = softmax(logits)
6:   pred_index = argsort(pred_values)
7:   predict_rank = doc_ids[pred_index]
8: end for
```

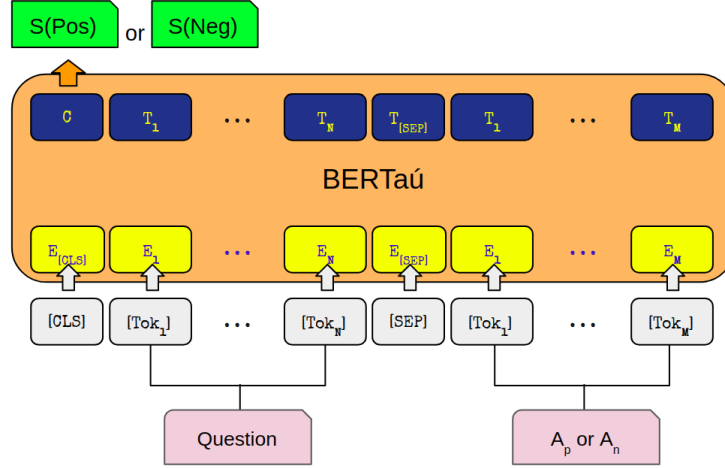


Figure 1: Starting from the bottom: The input is given by a `Question` and a positive answer `Ap` or negative answer `An`. The Bert’s tokenization process take care to separate the question and answer tokens to produce a score `S(Pos)` or negative score `S(Neg)`.

places all probability on a single class during the training phase (20). We set up (without any grid search) the confidence penalty to 0.1.

Alternatively, we have also experimented with Cross Entropy loss (when setting the confidence penalty to 0). This experiment had slightly lower results in MRR@10 and AP@1.

4.1.2 Pairwise

The pairwise approach follows the same one adopted in (21) and shown in Figure 1 which takes a couple of candidate answers and learns the most relevant answer for the question. Explicitly, he have a triple (Q, A_p, A_n) where Q is the question, A_p and A_n are the positive and negative answer respectively. This triple is broken into two pairs (Q, A_p) and (Q, A_n) and each pair is sent individually to a BERTaú. For the loss function, we used only the hinge loss pairwise function and not the combined cross entropy loss function with hinge loss function proposed in (21) because the two strategies obtained similar results and we explicitly choose the simplest loss. The loss used in the pairwise model is:

$$\max\{0, M - \hat{y}_\theta(Q, A_p) + \hat{y}_\theta(Q, A_n)\} \quad (3)$$

where $\hat{y}_\theta(Q, A_p)$ and $\hat{y}_\theta(Q, A_n)$ denote the predicted scores of positive and negative answer, whereas M is the margin parameter, in our case fixed in $M = 0.2$.

For the FAQ retrieval experiment we used AdamW optimizer, learning rate of $5e-5$, linear scheduler with warmup of 2% of total steps⁵ for 1 epoch.⁶ We also varied the number of cands in the dataset build, by increasing the number of cands for the same question, the imbalance of the dataset increases, we testing three levels of cands $\{15, 30, 45\}$, the Figure 2 shows the performance of the pointwise and

⁵total steps = $(\text{len}(\text{train_loader}) * \# \text{ epochs})$

⁶We have tested for 2 and 3 epochs, obtaining similar results but with more overfit symptoms

pairwise models. Table 3 shows the main FAQ retrieval results. Note that the pairwise model gets the best results, however this model takes twice as long in training when compared with pointwise.




Figure 2: Comparison between models, varying the imbalance of the FAQ dataset. x-axis legend means the number of possible candidates for one sample and xx%-tl is the percentage of unbalance in the dataset where tl = true-label.

Table 3: Evaluation result on IR FAQ dataset for each sample with 30 candidates unbalanced dataset: 97.60/2.40 true labels.

Model	MRR@10	AP@1
BM25+	0.345	0.246
distiluse-base-multilingual	0.417	0.285
mBERT uncased	0.458	0.505
DPRQuestionEncoder	0.526	0.482
BERTaú pointwise label smoothing	0.544	0.511
BERTaú pairwise	0.552	0.521

4.2 Sentiment Analysis

The study of Sentiment Analysis, brings valuable feedbacks for any kind of problems during the digital customer service . When we identify the feeling of a sentence, We can act proactively and guide customer service more efficiently.

In this sense, it is essential to promptly identify customers who demonstrate dissatisfaction, so that they can count on the proper assistance for their respective issues. The institution must also understand topics that have good acceptance so that it can ensure suitable maintenance and continuous improvement. Finally, its technicians must also fix processes and functionalities that are often related to negative sentiments. Thus, offline models of sentiment analysis were developed with the help of machine learning techniques.

Given the considerable amount of interactions that may not express any sentiment, the classification metrics proved to be more robust for multiclass models with three classes of sentiment: positive, negative and neutral:

1. Positive class - Praise and thanks.
2. Neutral Class - Doubts, requests and tracking of status of preview demands.
3. Negative Class - Cursing, complaints and indication of complaints in consumer protection agencies, reports of systemic errors, indication of improper service, disagreement with products, services and fees.

These classes follow the pattern of classical literature for the problem of sentiment analysis (22). Also, aiming to expand the labeled dataset reliably, databases from other sources and service channels

were used, such as the bank’s official social networks, examples of previously labeled interactions from human chat, and also telephone transcriptions.

4.2.1 Pre-processing

During the development of these models, a relatively traditional pre-processing text approach was used.

1. Removing accents, punctuation and special characters.
2. Conversion to lowercase.
3. Removing numbers.
4. Removing stopwords from a validated dictionary in the business context.

For comparison purposes, it is also worth noting that the models were trained by alternating stemming.

For a representation of the term-document matrix, we decided to use n-grams, more precisely unigrams, bigrams and trigrams, where, each term t_j (n-gram) of each document d_i , the product $tfidf(t_j, d_i) = TF(t_j, d_i)IDF(t_j)$ was computed. The $TF(t_j, d_i)$ can be interpreted as the frequency of occurrence of the n-gram t_j in a d_i document, and $IDF(t_j)$ as a measure of occurrence of the term t_j in the corpus as a whole.

Thus, from all n-grams $t_j, i = (1, \dots, M)$ and all documents $d_i, i = (1, \dots, N)$, we compute the products $TF(t_j, d_i).IDF(t_j)$ and form the matrix T .

To address issues related to the sparsity and high dimensionality of the corpus, we performed a *Singular Value Decomposition* on the matrix T (or SVD, for short).

Therefore, T can be represented as:

$$T' = U\Sigma V^t, \quad (4)$$

where U contains the eigenvectors of the correlations of the terms t , V contains the eigenvectors of the correlations of the documents d and Σ contains the singular values of the decomposition.

To find the projection of the vector representation of a document \mathbf{d}_j in the new space, we can write:

$$\hat{\mathbf{d}}_j = \Sigma^{-1}U^t\mathbf{d}_j \quad (5)$$

We applied a grid-search to find the best hyperparameters and ideal number of components used in SVD.

In the best model obtained, the initial 64,686 features were reduced via SVD decomposition to 650 components. The stage of decomposition into singular values was the most computationally expensive, even though it was performed using parallelization.

4.2.2 Semi-Supervised Learning

As it is an intrinsically human activity, the labeling phase by the data quality team is, in general, a moment in the process of creating models that naturally demands a longer execution time. For this step, each text example (document) selected to be labeled was examined by 2 different analysts, so that labeling errors due to individual bias were mitigated.

Only those documents that had been labeled equally by different analysts were selected to compose the training set. This way of composing the training set, although guaranteeing higher quality, also reduced drastically the example labeling rate.

In order to minimize this drawback and improve the labeling time, it was decided to use semi-supervised learning algorithms, in particular, the Co-training method (23).

The main idea is based on the cooperation of two supervised learning algorithms. Be matrix T with N documents and M features.

For the implementation of the Co-training method, the matrix T was divided into two other matrices T_1 and T_2 , as follows:

$$T = T_1 \cup T_2 \quad (6)$$

$$\emptyset = T_1 \cap T_2 \quad (7)$$

If we think of each term \vec{t}_k as a column vector of T , we can define T_1 and T_2 as:

$$T_1 = (\vec{t}_1, \vec{t}_2, \dots, \vec{t}_j) \quad (8)$$

$$T_2 = (\vec{t}_{j+1}, \vec{t}_{j+2}, \dots, \vec{t}_M), \quad (9)$$

so both T_1 and T_2 represent, somehow, the same documents, however with different sets of features.

For each of the sets, labeled examples were selected and the models h_1 and h_2 were trained. The techniques for training h_1 and h_2 do not need to be exactly the same. However, for convenience of implementation, we chose the LightGBM (24) method for both h_1 and h_2 .

From h_1 , unlabeled instances were classified into T_1 and the same process was repeated for unlabeled instances in T_2 . At each step k of the model, the sets T_1 and T_2 can be subdivided into $(L_{T_1}^k, U_{T_1}^k)$ and $(L_{T_2}^k, U_{T_2}^k)$, respectively, where the set $L_{T_i}^k$ contains the labeled samples from T_i of the k_{th} interaction and $U_{T_i}^k$ the unlabeled samples from T_i in the k_{th} interaction.

Given some objective criteria (see (23)), with certain degree of confidence, once the combined predictions from h_1^k and h_2^k indicate that some unlabeled document d_h belongs to a certain class, in the next interaction of the method, d_h , originally unlabeled, now comprises both $L_{T_1}^{k+1}$ and $L_{T_2}^{k+1}$.

This method is repeated until all instances are labeled or any stop criteria are met.

It is possible to show (23) that from h_1 and h_2 a third classifier can be formed, which works similarly to the Naive Bayes one.

The number of examples obtained from other service channels, and manually labeled was about $5k$ training instances. With the co-training method described above, it was possible to expand the dataset at least 5 times.

4.2.3 Model

Different types of models were tested and trained to classify sentiments for the AVI interactions expanded dataset.

The model that had the best F_1 score was a random forest with 450 trees and a depth of 5. The F_1 score for this configuration was 0.76.

4.2.4 Further considerations

Other types of vector representations for AVI text interactions were also tested, in particular a 300-component pretrained Portuguese skip-gram word2vec embedding that was extracted from a public repository.

The results, however, were relatively worse than those obtained by means of aforementioned method, a fact that seems to be related to the banking context present in the training dataset.

As previously emphasized, some criteria were used for labeling data under human supervision. Thus, when classification elements were found that were consistent with the negative class, it was prioritized, although - for example- the interaction could also bring a simple question (such as the status of a credit card request), praise or thanks. This orientation aimed to ensure that interactions with a negative sentiment were not labeled in other classes because they also could satisfy some of the rules contemplated by these other classes, a fact that implies a certain hierarchy between the classes.

It is worth noting that in order to achieve even more confidence in the classification of sentiments, some interactions are previously identified via an exact match and also by (high) similarity methods that use Bag-of-Words representations and compute both the distances, cosines and Levenstein. The main results of the SA task are shown in Table 4.

Table 4: Evaluation result on SA task.

SA task evaluated on F_1 score	SA_{trinary}	SA_{binary}
Random Forest + SVD + n-grams	0.760	-
mBERT uncased	0.838	0.901
BERTaú	0.850	0.920

4.3 Named Entity Recognition

Named Entity Recognition (NER) is one of the most common and fundamental sequence labeling tasks. Given a text, the objective is to correctly identify and extract a set of general-purpose and domain-specific named entities on a token level. While previous works focus on combining LSTM (Long Short-Term Memory) (25) and CRF (Conditional Random Fields) (26) models, our approach is to attach a dense layer at the end of the BERTaú model and train it to predict each token’s entity class.

Our NER dataset consists of 18370 manually annotated examples. Each label is, then, further divided following the BILOU schema (B - ‘beginning’ I - ‘inside’ L - ‘last’ O - ‘outside’ U - ‘unit’) with 16 different classes. The classes include specific banking products, services, functionalities, organizations, companies, places, and documents.

In our experiments, we tested three different ways of using BERTaú ’s weight vectors: concatenating the last 4 layers, summing the last 4 layers, and using the last layer only. The results were very similar, with the summing approach being slightly better. The best configuration was a dense layer trained with AdamW optimizer using a learning rate of 5×10^{-5} with linear scheduler, and a warm-up of 2% of total steps for 5 epochs. We measured the performance of the models using sequeval’s (27) implementation of the F_1 score, due to some classes being severely unbalanced. The results are shown in Table 5.

Table 5: Evaluation result on NER task

NER task evaluated with sequeval	F_1 score
mBERT uncased	0.840
BERTimbau - base	0.853
BERTaú	0.877

4.4 Quantization

In order to use our model as a feature extractor, i.e. to generate embedding representations of words and sentences from BERTaú we conducted our model to a simple quantization, although there are more sophisticated solutions with greater compression power such as the Open Neural Network Exchange (ONNX) (9) we chose a more traditional and simple way, the PyTorch quantization in INT8, which converts the FP32 tensors to INT8. A quantized model allows the storing tensors at lower bitwidths than floating point, which implies less time of inference. We compared the inference time and sizes on three devices: GPU-V100, CPU and Quantized. The results are in Table 6.

5 Conclusion

The field of Deep Learning and NLP has developed rapidly and has achieved good results in several tasks in the academic universe. However, applying these solutions in the industry to solve real-world problems is still a challenge. In order to improve the efficiency of our digital customer service, it is necessary to seek better solutions and algorithms such as BERT, but also customized solutions that are feasible to be implemented in production ensuring good results. In addition, such results translate into actual improvements in the users’ experience.

In this work, we present BERTaú , a specialist pretrained BERT model in the AVI data and fine-tuned for three tasks: FAQ Retrieval, NER and Sentiment Analysis. when compared with the mBERT model, BERTaú improves the performance in 22% of FAQ Retrieval MRR@10, 2.1% in Sentiment Analysis F_1 score and 4.4% in NER F_1 score. BERTaú can also represent the same sequence in up to 66% fewer

Table 6: GPU and CPU model size is 428.95Mb, Quantized model is 183.76Mb. Inference time in FAQ Retrieval experiment with BERTaú pointwise - label smoothing in one sample with 15 cand and inference time of NER and SA experiments in one sample.

Experiment	Device	Inference time
FAQ Retrieval	GPU - V100	0.2129s
	Quantized	7.2901s
	CPU	18.1826s
SA	GPU - V100	0.01644s
	Quantized	0.09357s
	CPU	0.15289s
NER	GPU - V100	0.01944s
	Quantized	0.03457s
	CPU	0.10064s

tokens when compared to "shelf models". This new approach, in addition to being applied to the three NLP tasks already carried out, we intend to apply our model as well as extracting resources from words and phrases.

Due to responsibilities related to sensitive data, privacy and business strategy, we cannot disclose BERTaú, NER and Sentiment Analysis training datasets, as well as the complete FAQ dataset. However, we have assembled a small dataset of public FAQs from Itaú Unibanco bank website and we tackle this dataset with BERTaú pairwise model.

6 Conflicts of Interest

Any opinions, findings, and conclusions expressed in this manuscript are those of the authors and do not necessarily reflect the views, official policies nor position of Itaú Unibanco.

References

- [1] C. Bones and J. Hammersley, "Leading digital strategy: Driving digital strategy: Driving business growth through effective e-commerce," *The Journal of Decision Makers*, vol. 42, pp. 128–132, 2017.
- [2] U. Gnewuch, S. Morana, and A. Maedche, "Towards designing cooperative and social conversational agents for customer service," *Proceedings of the International Conference on Information Systems (ICIS)*, 2017.
- [3] R. Mousavi, M. Johar, and V. S. Mookerjee, "The voice of the customer: Managing customer care in twitter," *Information Systems Research*, vol. 31, no. 2, pp. 340–360, 2020.
- [4] P. Rajpurkar, J. Zhang, K. Lopyrev, and P. Liang, "Squad: 100,000+ questions for machine comprehension of text," 2016.
- [5] A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, and S. R. Bowman, "Glue: A multi-task benchmark and analysis platform for natural language understanding," 2019.
- [6] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *Proceedings of the 2019 Conference of the North*, 2019.
- [7] A. Paleyes, R.-G. Urma, and N. D. Lawrence, "Challenges in deploying machine learning: a survey of case studies," 2020.
- [8] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, and A. M. Rush, "Transformers:

- State-of-the-art natural language processing,” in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Online: Association for Computational Linguistics, Oct. 2020, pp. 38–45. [Online]. Available: <https://www.aclweb.org/anthology/2020.emnlp-demos.6>
- [9] J. Bai, F. Lu, K. Zhang *et al.*, “Onnx: Open neural network exchange,” <https://github.com/onnx/onnx>, 2019.
- [10] F. Souza, R. F. Nogueira, and R. de Alencar Lotufo, “Portuguese named entity recognition using BERT-CRF,” *CoRR*, vol. abs/1909.10649, 2019. [Online]. Available: <http://arxiv.org/abs/1909.10649>
- [11] D. Araci, “Finbert: Financial sentiment analysis with pre-trained language models,” *CoRR*, vol. abs/1908.10063, 2019. [Online]. Available: <http://arxiv.org/abs/1908.10063>
- [12] A. Trotman, A. Puurula, and B. Burgess, “Improvements to bm25 and language models examined,” in *Proceedings of the 2014 Australasian Document Computing Symposium*, ser. ADCS '14. New York, NY, USA: Association for Computing Machinery, 2014, p. 58–65. [Online]. Available: <https://doi.org/10.1145/2682862.2682863>
- [13] N. Reimers and I. Gurevych, “Making monolingual sentence embeddings multilingual using knowledge distillation,” in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 11 2020. [Online]. Available: <https://arxiv.org/abs/2004.09813>
- [14] V. Karpukhin, B. Oguz, S. Min, P. Lewis, L. Wu, S. Edunov, D. Chen, and W.-t. Yih, “Dense passage retrieval for open-domain question answering,” in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Online: Association for Computational Linguistics, Nov. 2020, pp. 6769–6781. [Online]. Available: <https://www.aclweb.org/anthology/2020.emnlp-main.550>
- [15] A. Conneau, R. Rinott, G. Lample, A. Williams, S. R. Bowman, H. Schwenk, and V. Stoyanov, “Xnli: Evaluating cross-lingual sentence representations,” in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2018.
- [16] J. Lin, R. Nogueira, and A. Yates, “Pretrained transformers for text ranking: Bert and beyond,” 10 2020.
- [17] T. Lai, T. Bui, and S. Li, “A review on deep learning techniques applied to answer selection,” 08 2018.
- [18] T.-Y. Liu, *Learning to Rank for Information Retrieval*. Springer, 2011.
- [19] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, “Rethinking the inception architecture for computer vision,” *CoRR*, vol. abs/1512.00567, 2015. [Online]. Available: <http://arxiv.org/abs/1512.00567>
- [20] G. Pereyra, G. Tucker, J. Chorowski, L. Kaiser, and G. Hinton, “Regularizing neural networks by penalizing confident output distributions,” 01 2017.
- [21] D. Li, Y. Yu, Q. Chen, and X. Li, “Bertsel: Answer selection with pre-trained models,” *CoRR*, vol. abs/1905.07588, 2019. [Online]. Available: <http://arxiv.org/abs/1905.07588>
- [22] W. Medhat, A. Hassan, and H. Korashy, “Sentiment analysis algorithms and applications: A survey,” *Ain Shams Engineering Journal*, vol. 5, no. 4, pp. 1093–1113, 2014. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2090447914000550>
- [23] A. Blum and T. Mitchell, “Combining labeled and unlabeled data with co-training,” in *COLT'98: Proceedings of the eleventh annual conference on Computational learning theory*, 1998, pp. 92–100.

- [24] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu, “Lightgbm: A highly efficient gradient boosting decision tree,” in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017, pp. 3146–3154. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/file/6449f44a102fde848669bdd9eb6b76fa-Paper.pdf>
- [25] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [26] J. D. Lafferty, A. McCallum, and F. C. N. Pereira, “Conditional random fields: Probabilistic models for segmenting and labeling sequence data,” in *Proceedings of the Eighteenth International Conference on Machine Learning*, ser. ICML ’01. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2001, pp. 282–289. [Online]. Available: <http://dl.acm.org/citation.cfm?id=645530.655813>
- [27] H. Nakayama, “seqeval: A python framework for sequence labeling evaluation,” 2018, software available from <https://github.com/chakki-works/seqeval>. [Online]. Available: <https://github.com/chakki-works/seqeval>
- [28] J. Johnson, M. Douze, and H. Jégou, “Billion-scale similarity search with gpus,” *arXiv preprint arXiv:1702.08734*, 2017.

A Drawbacks

Here we briefly describe unsuccessful experiments

1. Train BERTaú from scratch starting from the final checkpoint of the mBERT: After 1,000,000 steps the model performed very poorly when compared with the training from scratch without any initial checkpoint.
2. FAQ retrieval experiment – combine the `distiluse` and use `Faiss` (28): Did not go well when compared to the cosine distance solution.
3. FAQ retrieval experiment – use the BM25+ to build the dataset and then use BERTaú as a re-ranking: There was no performance improvement and the solution only increased in complexity.
4. FAQ retrieval experiment – use BERTaú’s predictions for BM25+ and `Faiss`: results were comparable with BM25+ and worse with `Faiss`, our hypothesis is that the phrases ranked by BERTaú for a given query are more similar than a random choice, which may present greater difficulty for `Faiss`.
5. NER experiment - A popular way to NER is using a CRF layer at the end of BERT, following the idea of (10). Our results with this idea weren’t as good as the previously mentioned experiments, resulting in a slightly lower F_1 score and more symptoms of overfitting. Our hypothesis is that this result was influenced by our relatively small NER dataset.