

# Augmenting Customer Support with an NLP-based Receptionist

André Barbosa<sup>2</sup>, Alan Godoy<sup>1</sup>

<sup>1</sup> QuintoAndar, São Paulo, Brazil

<sup>2</sup>University of São Paulo, São Paulo, Brazil

andre.barbosa@ime.usp.br

alan.godoy@quintoandar.com.br

**Abstract.** *In this paper, we show how a Portuguese BERT model can be combined with structured data in order to deploy a chatbot based on a finite state machine to create a conversational AI system that helps a real-estate company to predict its client's contact motivation. The model achieves human level results in a dataset that contains 235 unbalanced labels. Then, we also show its benefits considering the business impact comparing it against classical NLP methods.*

## 1. Introduction

Effective customer support is an essential piece for any service company. On the one hand, as a business grows, it can hire more agents and create specialized departments focusing on each possible problem a customer may face, increasing the availability of support for clients. On the other hand, growth brings new challenges: knowledge about client issues gets scattered across the company, information regarding processes may not diffuse properly to all agents, clients may have trouble finding the specific department able to help them, to name a few. As a real estate company, [QuintoAndar](#)<sup>1</sup> has to deal with a wide range of issues involving individuals as distinct as tenants (and prospective tenants), landlords, agents (realtors, inspectors, photographers, etc) and building administrators. The channel of choice for most contacts is [WhatsApp](#), a messaging app widely adopted in Brazil and other countries. This introduces even more complexity in terms of user identification and routing, as any WhatsApp user can start a conversation with our customer support team simply by sending an open text message.

To improve the routing quality, speed, and automate the collection of all necessary information before a human analyst is allocated to the ticket, our team developed a chatbot system that acts as a receptionist for customer support. This chatbot (presented in section 4) uses a state-machine dialog manager to integrate multiple machine learning (ML) based classifiers and business rules to define the messages to be sent and actions to be taken. The implementation of such service brought a lot of positive results to QuintoAndar, allowing the full automation of chat triage in customer support, a task performed previously by a dedicated team, which was reallocated to other departments.

Beyond presenting a chatbot architecture that combines multiple ML models and discussing how we merged textual information with structured data in order to better

---

<sup>1</sup>A quick presentation about QuintoAndar and its business model can be found [here](#).

understand an user’s context, the main contribution of this work is to bring data about the results achieved by state-of-the-art models in a real-world application. Much has been discussed about the fantastic results achieved by recent Natural Language Processing (NLP) models for the English language both in academic literature and in industry. How these outcomes transfer to Brazilian Portuguese is not as clear though. We hope that, by sharing our experience, we may shed some light on how Portuguese-speaking companies and research groups may use these technologies outside common benchmark problems.

In summary, the main goal of this paper is to show how we evolve until we achieve human-level results in a real-world application through combining Portuguese BERT with structured metadata from our clients with a complete end-to-end system.

## 2. Related Work

Customer support is a natural candidate to be one of the areas most benefited by the recent boost in NLP research and application [Wan and Chen 2018, Liu et al. 2020, Molino et al. 2018]. By having a proper system for customer triage, a company can match its user to the agent most likely to solve their problem and also provide the agent with contextual information — e.g.: selected user information, procedure suggestion and a list of similar issues — to make their work more effective. Alexandra DeLucia and Elisabeth Moore [DeLucia and Moore 2020] used IT support tickets in a high-performance computing laboratory to study automatic categorization and similar ticket recommendation. They combined classical NLP pre-processing, including steps as stop-word removal, stemming and topic modeling, with a random-forest classifier to select the most appropriate label among 93 possible categories and retrieve other similar tickets. A similar work was done by Fotso *et al.* [Fotso et al. 2018], that created a system to classify client emails in 12 possible categories, using it to select relevant articles to be automatically sent as response. The team combined textual information with data regarding the user’s prior interactions with their products, using word-embeddings, a BiLSTM network and an attention-based mechanism to predict the most adequate class.

Rather than applying a single machine learning model to classify the user demand, it’s also possible to create a dialog system that interacts with the customer to try solving their problem or to ensure that all relevant information was provided before assigning the ticket to an agent. Such system demands not only a model that extracts information from user messages but also a dialog management module to track the current dialog state and define which action to perform in each moment and also a module to generate the messages sent to the user [Dai et al. 2020, Jurafsky and Martin 2021, Zhang et al. 2020].

When considering Brazilian Portuguese, however, there are few works reporting applications of NLP in industrial applications. Finardi *et al.* [Finardi et al. 2021] built a BERT language model [Devlin et al. 2019] to be used for customer support in a large Brazilian bank, evaluating its performance for sentiment analysis, question answering and named entity recognition (NER) using datasets extracted from user interactions with the bank’s chatbot. Azevedo *et al.* [Azevedo et al. 2020] created a system to automatically route customer emails to four different boxes using a Support Vector Machine (SVM). Works have also been published regarding applications in law, as [Bonifacio et al. 2020] — which extensively studies the impacts of fine-tuning of transformer-based language models using legal texts in NER tasks — and [Dal Pont et al. 2020] — that provides an

empirical study of word embeddings in legal domain.

On the academic front, however, many works were produced recently related to dialog systems that can be applied for customer support in Portuguese. Santos *et al.* [Santos et al. 2020] and Melo and Coheur [Melo and Coheur 2020], for instance, built retrieval-based conversational agents trained to answer specific data. Carvalho *et al.* [Carvalho et al. 2020], on the other hand, created an LSTM-based common-sense module to augment interactions in a dialog system. A recent significant contribution was also made by Souza *et al.* [Souza et al. 2020] that made open-source a Brazilian Portuguese pre-trained BERT language model, showing how such mono-lingual model was able to surpass the performance achieved by a multi-lingual BERT for named entity recognition (NER).

### 3. Problem Definition and Challenges

A typical customer support flow at QuintoAndar (Fig. 1) starts with a customer contacting our team through the channel of their choice — namely, phone, WhatsApp or email/support form. Whichever channel is chosen, the user is required to state what kind of support is needed, an information that is used to define to which department the contact should be directed. After the department is defined, a task allocation system selects among all online agents who should be responsible for handling the customer’s issue. The role of such agent is to comprehend the user’s demand and get all relevant information to solve it — if it is simple enough, it can be resolved immediately; if it requires more complex actions, a task is created for the appropriate back-office team.

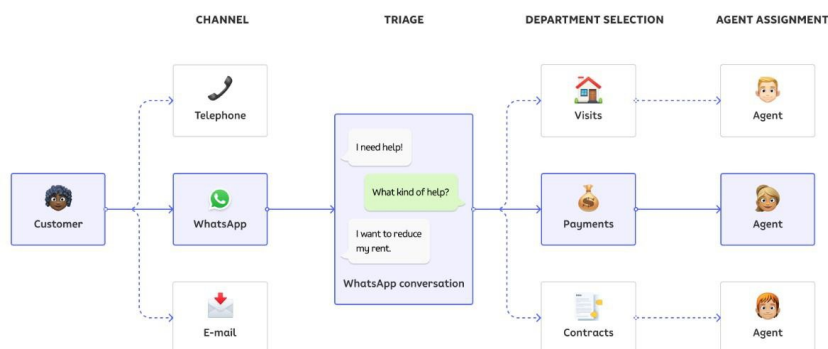


Figure 1. A typical flow of customer support at QuintoAndar.

In a digital real-state company, however, there are hundreds of different types of contacts, going from simple questions about using our site to search for a house to requests for intermediation of complex tenant-owner issues like negotiating a temporary rent reduction. This means that matching a user to the right specialist is easier said than done. On the one hand, presenting a menu with a multitude of customer support departments may be time-consuming and ineffective, as users may not be familiar with each department’s responsibilities and are likely to make mistakes when choosing an option. On the other hand, letting the users freely declare what they need has its own drawbacks, such as the possibility that the user’s explanation does not have sufficient context regard-

ing their demand and the issue of how to convert that explanation into a decision about the destination department of the contact.

For WhatsApp — our main support channel and the focus of the present work — this issue gets trickier, though: the app is ubiquitously used in Brazil as a means of communication with family and friends. Added to this is the fact that Brazilian clients are used to small real estate companies in which commercial interactions are based on a person-to-person relationship. Contacts, thus, may start with a simple “*Hi!*”, with a complete account of the user’s issue or with a partial description that refers to past contacts.

Prior to this work, triage and department selection was carried out by a team dedicated solely to verifying that sufficient context was provided, asking users to give more information if necessary, and directing each chat to the appropriate department. Of course, this was a far from ideal situation. A large number of skilled agents were diverted from customer troubleshooting to perform a highly repetitive task that added latency to proper chat routing and stiffened any processes changes that incur in changes in departmental responsibilities, as any such changes would require retraining the triage team.

To sustain company’s scalability, our team developed a “receptionist” chatbot that should be able to:

1. Determine whether the user has provided enough information in their initial message, requesting a better description if deemed necessary.
2. Combine this textual information with data about the user’s relationship with QuintoAndar to predict the reason for the current contact.
3. Automate the collection of basic data, such as information that helps to identify users who could not be identified based on their phone number and data to help the agent verify the authenticity of the user.
4. Use business rules together with the predicted contact reason to determine to which department the chat should be directed.
5. Provide the selected agent with the predicted reason for the contact in order facilitate access to the appropriate guidelines for that service and speed up the mandatory annotation after the conversation is finished.

This was not without its hurdles. First, there were more than 300 standardized possible contact reasons with a substantial gray area between them, ensuring a great deal of noise in labels. Classes were highly unbalanced, with some containing only a handful of chats in a given year, while others classes totaling thousands in a single week. Finally, as a Portuguese-speaking company, there were no open resources available related to customer service, so we could rely only on pre-trained models and our own private datasets to develop the chatbot.

## 4. System Architecture

Following the dialog-state architecture [Jurafsky and Martin 2021], the bot engine we created has the following basic components:

- Dialog memory: a flexible memory to store all relevant information produced by handlers during the conversation, including special field  $S$  that indicates the automaton’s current state.

- **Handlers:** processors used by the chatbot to extract relevant information from input messages or make decisions about what to do next.
- **Finite-state automaton:** a recipe that indicates which handler should be called at each state  $S$  and how to use a policy handler results to select the next state.

Handlers, in turn, can be of two different types:

- **Message processing handlers:** responsible for extracting information from input messages and storing it on dialog memory. The present work focuses on proposing two machine learning models that act as message processing handlers.
- **Policy decision handlers:** responsible for deciding what the system should do given the information stored in dialog memory. Examples of possible actions are performing some calculations and storing their results in the memory, sending a response to the user or transferring the chat to a certain department. When a policy handler decides to send the user a response, it forwards asks the language generation unit to select a template to use to create the message.

Each handler has its own implementation, which may range from simple application of business rules to complex ML models and third-party services. This work makes use of two different NLP-based handlers that will be explained on the next subsections.

#### 4.1. Context evaluation model

A common behavior by our users was to start the conversation by saying a greeting, only stating what she needed after being asked by the support agent. This posed a significant challenge for our receptionist bot, as it would predict contact reasons and select destination departments based on messages that were not at all explanatory. A simple rule of thumb to avoid this problem would be to always discard the first message in a conversation. However, this would hurt the client's experience, which is only worsened by the fact that many customer support contacts are made by already stressed users.

To alleviate this issue, we created a model to predict whether or not a particular message was sufficiently explanatory. Based on such prediction our chatbot may decide to ask the client for more input. A dataset with 5348 chats randomly sampled from QuintoAndar's ticketing system was manually annotated by two customer support agents. They could classify a message as one of the following: *has context*, *no context*, *returning client* and *low value* (a message that is not intended to create a new conversation, as "Thanks!"). The final dataset contained 2926 chats labeled as its content has enough context and 2422 labeled as not enough context (*no context* and *low value*; *returning client* messages were discarded), endorsing the relevance of this first model in an efficient reception.

Our goal was to have a simple model that would introduce only a small delay in response time, so we decided to use bag-of-words (BoW) representation with a logistic regression (LR) as classifier. Text pre-processing included accent stripping, uncasing and stop-word removal; no stemming or lemmatization was used. Bag-of-words representation was then computed for 3-grams with a vocabulary of 5000 items. LR hyperparameters ( $C$ , penalty type and class weight) were selected through Bayesian optimization — the solver used was 'lbfgs' [Zhu et al. 2011]. The resulting model achieved accuracy of 85% on the testing split, a value considered good enough for production as most errors in this step could be fixed by human agents after chat allocation.

## 4.2. Contact reason prediction model

The main model for the chatbot is the contact reason prediction model. As explained above, its main purpose is to combine the user’s message with tabular information indicating their relationship with QuintoAndar to predict how likely the contact is to be related to each of 306 possible standardized contact reasons (see Table 1 for some examples).

cr_pg	Payments of real-estate brokers.
ft_ag_alteracao	Schedule changes for photographers.
iq_pr_reserva	Issues related to house reservation by a tenant after a proposal is accepted by an owner.
pp_cm_venda_imovel	An owner informs that she intends to sell the house she currently rents.

**Table 1. Examples of contact reasons defined by QuintoAndar’s process team.**

The association of text and tabular data is important as a way to avoid requiring users to explain their whole history to the chatbot. Consider the phrase “*I need to cancel the visit tomorrow.*”. It is a complete yet very ambiguous message: is it related to a photographer that wants to reschedule a photo shoot, to a potential tenant who is not interested on visiting an apartment anymore, or to a real estate agent that will not be able to present some house and wants to leave it to another colleague? Only by having access to the user’s relationship with QuintoAndar the model would be able to accurately predict the contact reason without requesting further context. Therefore, we used 66 handcrafted features available from our feature store [Marques 2020, Hermann and Balso 2017]. Some examples are the type of the last automatic message sent to the user (and time since it was sent), the contact reason for the last ticket (and time since it was created), whether the user is a registered agent, number of rented houses (as owner) and number of ended contracts (as tenant). These features are pre-processed, performing one-hot encoding on all categorical columns and scaling the numeric ones to have zero mean and unit variance.

We treat both feature groups (textual and tabular) as different modules that are fed into a separate classifier. In a first version (V1 - bag-of-words), we used a simple unigram bag-of-words to extract features from messages. After analyzing this model errors, we noticed that, despite showing good results, it had trouble with synonyms and complex sentences. In a second version of the model (V2 - BERT) we addressed this issue by using for textual feature extraction a representation extracted by a Portuguese version of a BERT model [Devlin et al. 2019, Souza et al. 2020].

## 5. Experiments

We briefly discuss the results achieved for version 1 and version 2 of contact reason prediction model, taking into consideration different setups as well as analyzing overall model’s performance.

### 5.1. Dataset

The full dataset used to train the contact reason model contains data for 639159 chats between May 2019 and August 2020 manually annotated by support agents, selecting only user messages sent before an agent entered the conversation. We used out-of-time

split to partition the between training (511327 chats), validation (63916 chats) and testing (63916 chats) sets. With respect to classes, originally the data contained 306 distinct ones — after filtering all tickets belonging to classes with less than 50 samples on training set the final dataset contained 235 different labels. To fine-tune BERT on our data, since it is a considerably heavy model, we have used a smaller portion of dataset, with 178578 samples for training, 19843 for validation and 66141 for testing.

## 5.2. Metrics

To analyze the model performance, we have calculated top-1 and top-3 test accuracy both for contact reason and for department<sup>2</sup>.

## 5.3. V1 - Bag-of-words

For this first solution approach, we have evaluated three possible classifiers: logistic regression (LR), random forest (RF) and multilayer perceptron classifier (MLP). Hyperparameters were selected through Bayesian optimization with 100 evaluated configurations [Biewald 2020]. Classifiers comparison are presented in Table 2.

**Table 2. Test accuracy for contact reason and department. Higher is better.**

Model	Top-1 accuracy	Top-3 accuracy	Top-1 dept. accuracy	Top-3 dept. accuracy
LR	42.8%	63.6%	77.8%	84.6%
RF	40.7%	61.2%	75.2%	82.7%
MLP	<b>44.1%</b>	<b>65.1%</b>	<b>78.2%</b>	<b>85.0%</b>

### 5.3.1. V2 - BERT

As observed in Table 2, the MLP model obtained the best results according to all metrics. Given the high cost of using a large Transformer-based network and the fact this new version is very similar to V1, we decided to keep the same architecture.

The first step we took was to fine-tune the BERT-Large model made available by Souza *et al.* [Souza et al. 2020] to predict contact reason. Sentences with more than 64 tokens were truncated. Once this fine-tuned procedure was completed, we tested different methods in order to replace the bag-of-words of V1 with a BERT module:

- Use the logits of BERT classification head.
- Use the last layer of BERT language model head as suggested by Devlin *et al.* [Devlin et al. 2019].
- Concatenate last four layers of BERT language model head as suggested by Devlin *et al.* [Devlin et al. 2019]

The results are shown in Table 3. First of all, it is relevant to notice the large gains achieved by combining both textual and tabular features. Also, we can see that using BERT as textual feature extractor provided better results than using bag-of-words. Considering engineering constraints we have decided to use the logits of BERT classification head + tabular data as version 2 of the contact reason model.

<sup>2</sup>Department selection is performed by summing predicted probabilities for all contact reasons associated to a given department. The department with highest score is selected.

**Table 3. Comparison of contact reason accuracy for multiple models.**

Model	Top-1 accuracy
BoW alone	38.14%
BoW + tabular data (V1)	44.11%
BERT classifier	45.57%
BERT classifier logit heads + tabular data	53.10%
<b>Last layer LM BERT + tabular data</b>	<b>53.20%</b>
Last 4 layers LM BERT concat + tabular data	53.05%

## 6. Business impacts

To assess the results in production, we have collected data from a triage human team and for a set of heuristics rules that routed simply based on the last automatic message sent to the user (e.g., if the message was related to a visit, then the user was directed to the visits department). The business metric that we decided to follow to compare was the transference rate — *i.e.*, the rate in which a chat already routed to a given support department should be transferred to another department.

We have run tests in production comparing human triage and both chatbot versions. As a security measure to avoid deterioration of user experience, we only routed automatically the 80% of tickets with highest department score leaving all low-confidence tickets to human triage. As the results were good enough (*i.e.* similar to human performance), we expanded it to 100% to our clients base. Results are presented in Table 4.

**Table 4. Production results. Data for human triage and heuristic rules refer were collected prior the tests. Lower the better for both columns.**

Model	Transf. rate	Avg. msg. per ticket
Human triage	12.8%	18.2
Heuristics rules	18.3%	11.2
V1 - Bag-of-words (80%)	13.9%	13.2
V2 - BERT (80%)	10.3%	13.7
V2 - BERT (100%)	13.2%	14.2

Considering these results, we can easily say that the BERT embedding combined with tabular data achieved human-level performance. The number of message exchanges until the conversation was ended (avg. message per ticket) was also substantially smaller for clients routed by the chatbot than for those routed by humans. An hypothesis regarding the reduction in message numbers is the fact that by using tabular features the model has access to a large amount of information not easily consumed by humans.

Such results in a real-world application using business metrics endorses the potential of modern NLP techniques for Brazilian Portuguese. We hope that results like these help companies and governments to see that it's now feasible to go much beyond the widespread rigid conversational interfaces based on buttons or simple keywords.

## Acknowledgments

We would like to thank Muriel Dias for producing the images in our paper, as well as Marco Antonio Rocha Vinha for reviewing the system architecture shown in section 4.



## References

- Azevedo, R. F. d., Rodrigues Pereira de Araujo, R., Guimarães Araújo, R., Moreira Bitencourt, R., Ferreira Alves da Silva, R., de Melo Vaz Nogueira, G., Marques Franca, T., Otharan Nunes, J., Ralff da Silva, K., and Regiane Cunha de Oliveira, E. (2020). Screening of email box in portuguese with svm at banco do brasil. In Quaresma, P., Vieira, R., Aluísio, S., Moniz, H., Batista, F., and Gonçalves, T., editors, *Computational Processing of the Portuguese Language*, pages 153–163, Cham. Springer International Publishing.
- Biewald, L. (2020). Experiment tracking with weights and biases. Software available from wandb.com.
- Bonifacio, L. H., Vilela, P. A., Lobato, G. R., and Fernandes, E. R. (2020). A study on the impact of intradomain finetuning of deep language models for legal named entity recognition in portuguese. In Cerri, R. and Prati, R. C., editors, *Intelligent Systems*, pages 648–662, Cham. Springer International Publishing.
- Carvalho, C. S., Pinheiro, V. C., and Freire, L. (2020). A deep learning model of common sense knowledge for augmenting natural language processing tasks in portuguese language. In Quaresma, P., Vieira, R., Aluísio, S., Moniz, H., Batista, F., and Gonçalves, T., editors, *Computational Processing of the Portuguese Language*, pages 303–312, Cham. Springer International Publishing.
- Dai, Y., Yu, H., Jiang, Y., Tang, C., Li, Y., and Sun, J. (2020). A survey on dialog management: Recent advances and challenges.
- Dal Pont, T. R., Sabo, I. C., Hübner, J. F., and Rover, A. J. (2020). Impact of text specificity and size on word embeddings performance: An empirical evaluation in brazilian legal domain. In Cerri, R. and Prati, R. C., editors, *Intelligent Systems*, pages 521–535, Cham. Springer International Publishing.
- DeLucia, A. and Moore, E. (2020). Analyzing hpc support tickets: Experience and recommendations. *ArXiv*, abs/2010.04321.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding.
- Finardi, P., Viegas, J. D., Ferreira, G. T., Mansano, A. F., and Caridá, V. F. (2021). Bertaú: Itaú bert for digital customer service.
- Fotso, S., Spanoudes, P., Ponedel, B. C., Reynoso, B., and Ko, J. (2018). Attention fusion networks: Combining behavior and e-mail content to improve customer support. *ArXiv*, abs/1811.03169.
- Hermann, J. and Balso, M. D. (2017). Meet michelangelo: Uber’s machine learning platform.
- Jurafsky, D. and Martin, J. H. (2021). *Speech and Language Processing (3rd Edition)*. Prentice-Hall, Inc., USA.
- Liu, C., Jiang, J., Xiong, C., Yang, Y., and Ye, J. (2020). Towards building an intelligent chatbot for customer service: Learning to respond at the appropriate time. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Dis-*

- covery & Data Mining*, KDD 20, page 3377–3385, New York, NY, USA. Association for Computing Machinery.
- Marques, A. (2020). Butterfree: A spark-based framework for feature store building.
- Melo, G. and Coheur, L. (2020). Towards a conversational agent with “character”. In Quaresma, P., Vieira, R., Alúcio, S., Moniz, H., Batista, F., and Gonçalves, T., editors, *Computational Processing of the Portuguese Language*, pages 420–424, Cham. Springer International Publishing.
- Molino, P., Zheng, H., and Wang, Y.-C. (2018). COTA: Improving the speed and accuracy of customer support through ranking and deep networks. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD 18, page 586–595, New York, NY, USA. Association for Computing Machinery.
- Santos, J., Alves, A., and Gonçalo Oliveira, H. (2020). Leveraging on semantic textual similarity for developing a portuguese dialogue system. In Quaresma, P., Vieira, R., Alúcio, S., Moniz, H., Batista, F., and Gonçalves, T., editors, *Computational Processing of the Portuguese Language*, pages 131–142, Cham. Springer International Publishing.
- Souza, F., Nogueira, R., and Lotufo, R. (2020). BERTimbau: pretrained BERT models for Brazilian Portuguese. In *9th Brazilian Conference on Intelligent Systems, BRACIS, Rio Grande do Sul, Brazil, October 20-23*.
- Wan, M. and Chen, X. (2018). Beyond “how may i help you?”: Assisting customer service agents with proactive responses.
- Zhang, Z., Takanobu, R., Zhu, Q., Huang, M., and Zhu, X. (2020). Recent advances and challenges in task-oriented dialog system.
- Zhu, C., Byrd, R., Nocedal, J., and Morales, J. L. (2011). L-BFGS-B — software for large-scale bound-constrained optimization.