

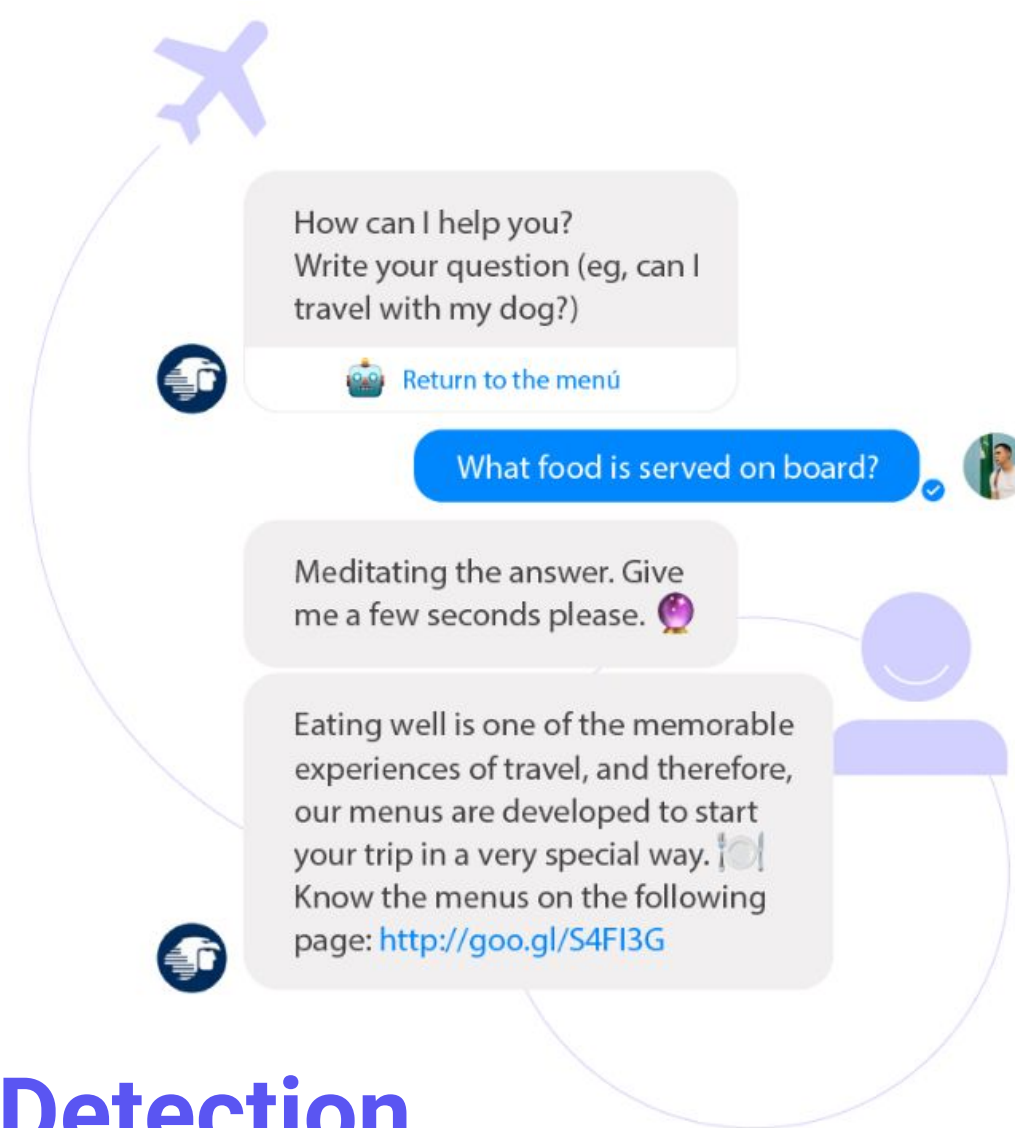
Yalo Enhancing Conversational Experiences with Artificial Intelligence

Introduction

The main medium of information exchange in our product is text. Therefore understanding human natural language in text is essential. We face interesting challenges continuously. In this poster we describe some of these challenges, the approaches we took to solve them and the lessons we learned along the way.

What is Yalo

- Yalochat is a Customer Relationship Management (CRM) solution for massive messaging channels such as WhatsApp and Messenger. We want to make the day of +1 billion people by delivering delightful conversational experiences.
- Our customers can create and grow meaningful and personalized 1-on-1 relationships with their users.



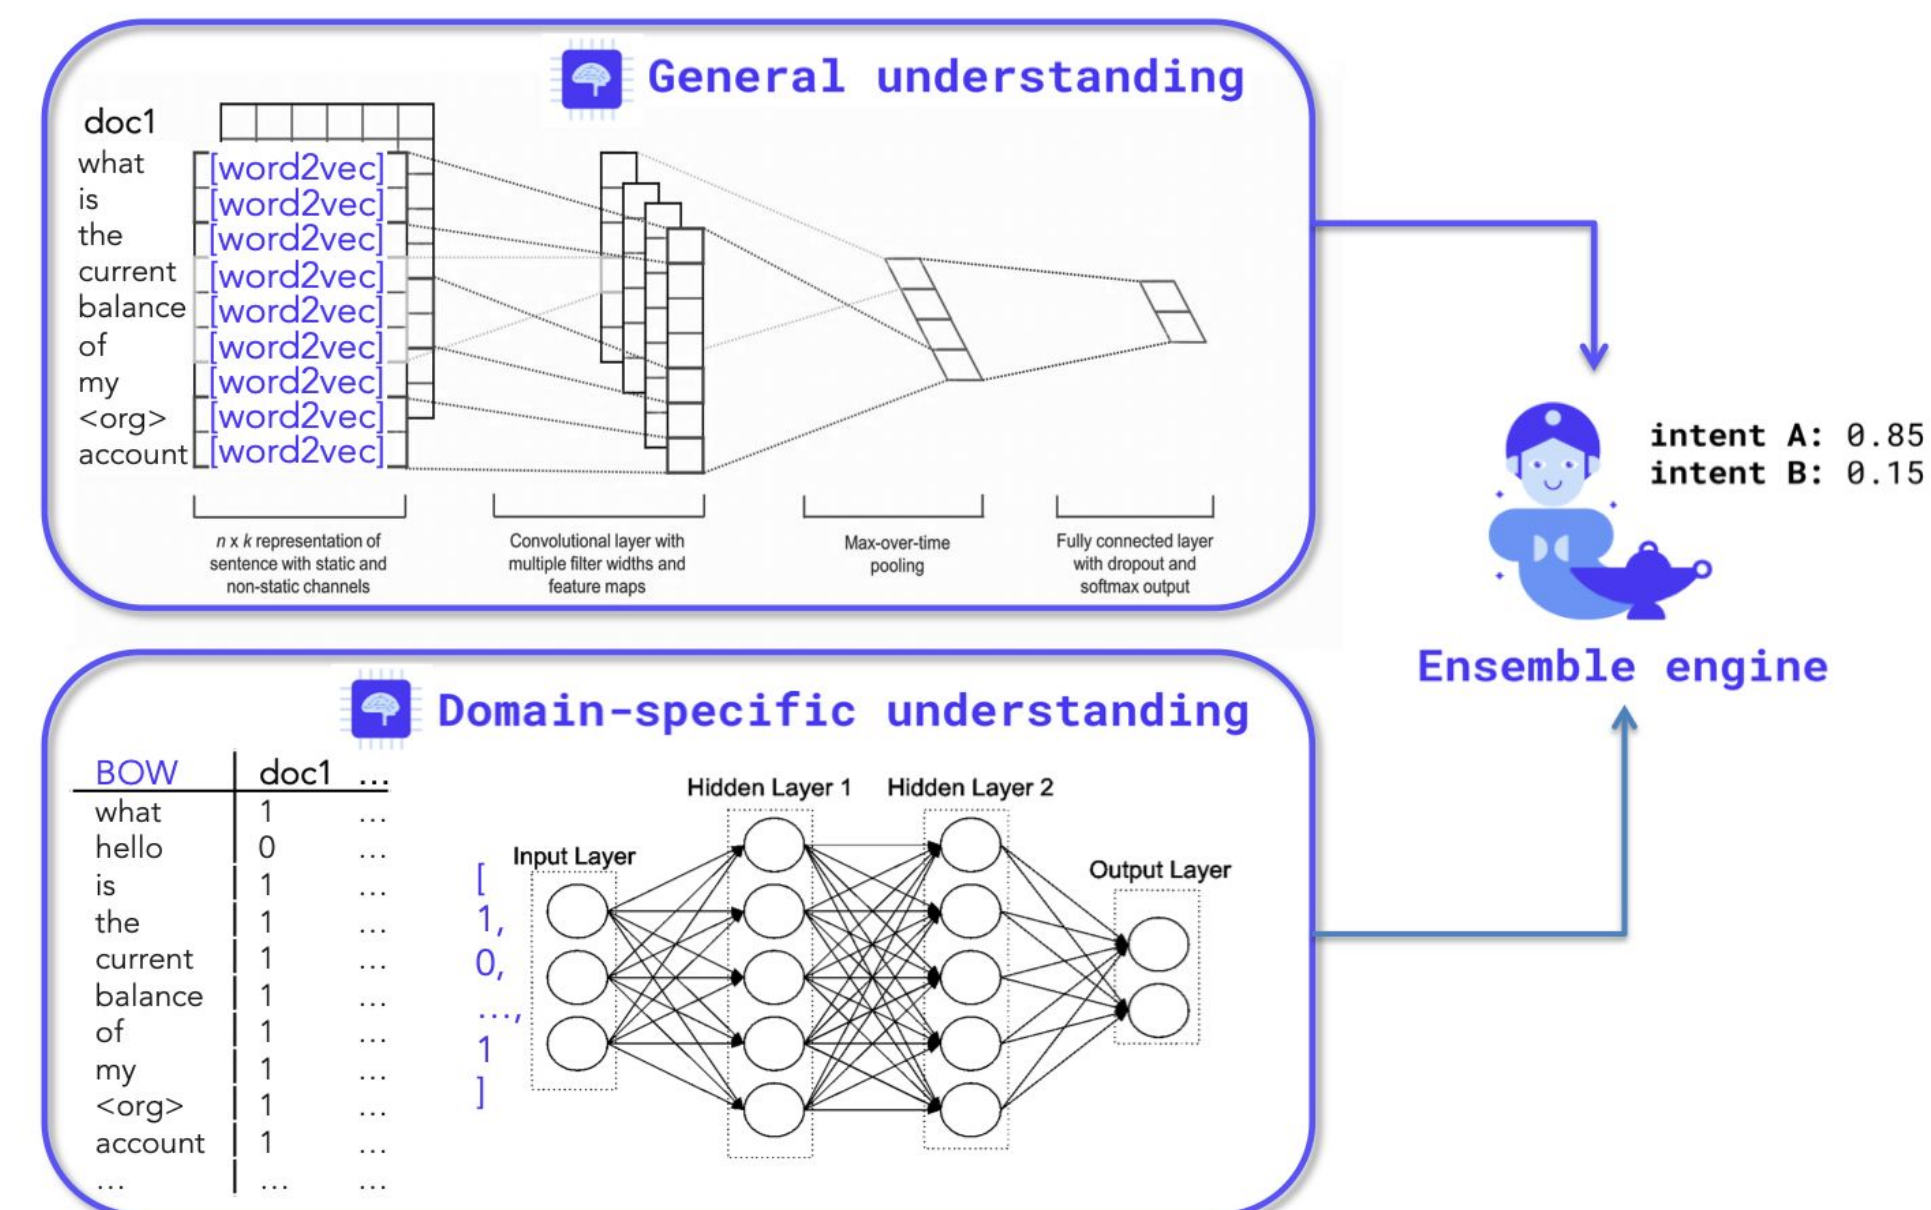
User Intent Detection

- Users have an intent when they communicate with a company.
- We cast user intent detection as a text classification problem.
- We leverage supervised learning. We train various generic and domain-specific classifiers and then combine them in with ensemble learning.

Intent Detection Learnings

Pre-computed embeddings with deep learning for general understanding, BOW approaches and shallow models for specific understanding and noisy data

- Conversational data is very noisy and very domain specific. (misspellings, slang, product/services names, brands)
- We use Deep Neural Networks with word-level embeddings (Word2Vec, GloVe, Facebook Starspace) trained with general corpora.
- We also train with domain data, overfitted shallow neural networks with BOW-based encodings (one-hot vector, TF-IDF, LDA) to understand the intent in noisy or domain-specific data.
- We obtain an aggregated result with an ensemble engine.



FAQs Resolution

- Our customers want our product to answer Frequently Asked Questions (FAQs) automatically.
- FAQs resolution is more challenging than intent detection because the data is semantically richer.

FAQs Resolution Learnings

It is important to quantify the semantic richness of data

- This is crucial to define the data life cycle of your product.
- The semantically richer your data is, the more powerful your model has to be. You may need to gather more data. You need to spend more time and money in data annotation.

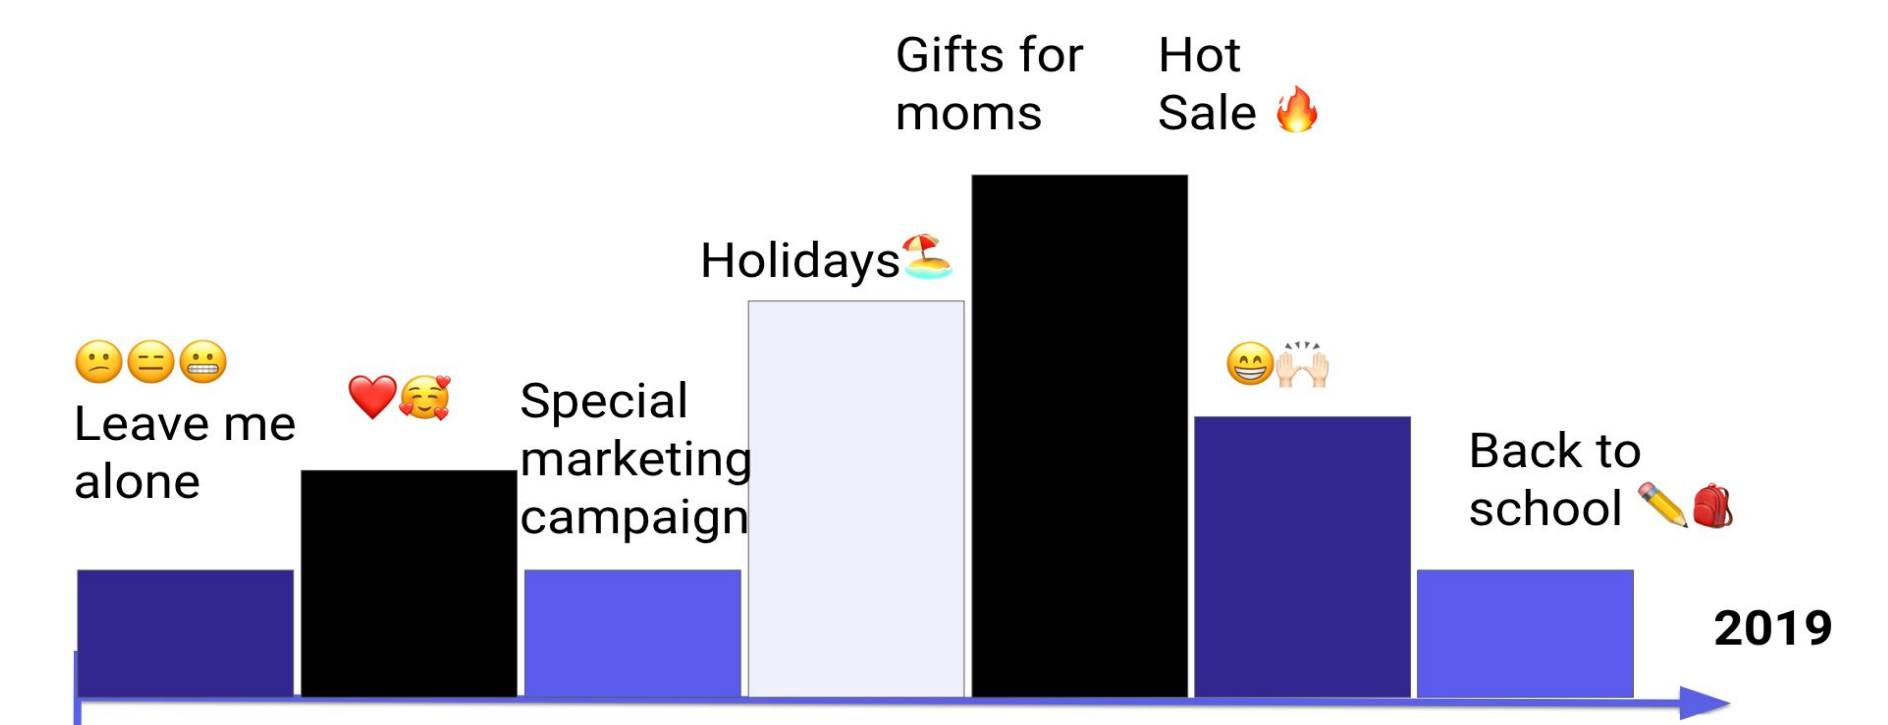
Unsupervised first, supervised later.

- Performance of unsupervised methods is good enough for a lot of use cases.
- Start with an unsupervised approach to gather data to train a supervised model for better results.

	Unsupervised	Supervised
Level of language	Semantic	Pragmatic
1st answer is accurate	60%	90%
Misspellings	✗	✓
Regional slang	✗	✓
Business technicalities	✗	✓
Variation in compound words	✗	✓

Mapping data variance through time is essential.

- You need to know how your data changes through time. There is seasonality and periodicity in the data you need to handle to get good performance.
- Use topic modeling to qualitatively describe the data. Describe your data using metrics like vocab size and dataset size.



Conclusion

At Yalo, challenges are ever-changing, numerous and fun to work with. We are exploring things like generative chat, semantic image search and real-time insights and trends. Scaling AI is a challenging issue. We hope to present some of this in further venues.