

Extracting Latent Beliefs and using Epistemic Reasoning to Tailor a Chatbot

Amit Sangroya, C. Anantaram, Pratik Saini, Mrinal Rawat
 TCS Innovation Labs, Tata Consultancy Services Limited, Gurgaon, India
 (amit.sangroya, c.anantaram, pratik.saini, rawat.mrinal)@tcs.com

Abstract

During dialog with a customer for addressing his/her complaint the chatbot may pose questions or observations based on its underlying model. Sometimes the questions or observations posed may not be relevant given the nature of complaint and the current set of beliefs that the customer holds. In this paper we present a framework to build conversation system that addresses customer complaints in a meaningful manner using domain understanding, opinion analysis and epistemic reasoning¹. Extraction of latent beliefs assists in performing epistemic reasoning to maintain a meaningful conversation with the customer.

1 Introduction

Dialog based interaction between a customer and a bot may become tedious and irrelevant if the system’s beliefs are not consistent with the set of beliefs of the customer. Beliefs are cognitive representational states that represent the presumed context of the conversation perceived by each agent. In other words, the presumed context that each agent has should match with the evolving context of the dialog for a meaningful conversation to take place. For example, if a chatbot fails to understand the customer’s emotional situation and responds mechanically with an irritated customer, then the chatbot may fail to achieve its primary objective.

Traditional dialog systems are specialized for a domain [Henderson *et al.*, 2014; Serban *et al.*, ; Weston, 2016; Williams *et al.*, 2016; Young *et al.*, 2013] and rely on slot-filling driven by a knowledge base and a finite-state model [Lemon *et al.*, 2006; Wang and Lemon, 2013]. Deep learning based dialog systems [Miller *et al.*, 2016] use memory networks to learn the underlying dialog structure and carry out goal-oriented dialog. Another approach is to build a system with a complex set of hand-crafted rules that may address some specific instances. Both approaches may be impractical in many real-world domains. In this paper, we propose a methodology that uses a combination of machine-learning mechanisms and domain specific knowledge extraction to understand the severity of customer’s complaint. This

helps to understand the customer’s latent emotional beliefs while giving the complaint. Our model then evaluates the beliefs to tailor the dialog and make it consistent with the set of beliefs of the customer. This process then helps drive the conversation in a meaningful way.

2 Proposed Framework

A customer’s complaint description is usually a set of complex natural language sentences describing the problem faced by the customer. Sometimes such descriptions contain background information, other superfluous information and may also have customer emotions embedded in it. We first categorize the customer complaint description into a set of categories and then extract relevant information. It is then evaluated by a hand-crafted knowledge base to assert facts about the probable beliefs held by the human agent in the complaint.

We assume that we have a hand-crafted complaint-management finite-state-machine (FSM) to carry out the dialog with the customer. The FSM operates on slots that are filled by extracting information from the complaint and subsequent interaction. The probable beliefs of the customer that were asserted as facts, and the category of the complaint are then evaluated by the epistemic rules encoded in a knowledge base for the domain. The rules make assertions about the states in the FSM that need to be skipped and the states that need to be evaluated in order to be consistent with the beliefs of the customer. The subsequent dialog is carried out and the next set of beliefs are then asserted. The cycle then continues. Algorithm 1 presents the overall process.

2.1 Categorizing the Human Interactions

As we discussed, a customer’s initial description of his/her complaint may be a verbose set of natural language sentences. Since descriptions about the same type of complaint can be vastly different across customers, we use machine-learning ways to categorize the initial complaint made by a customer. Recurrent neural networks (RNNs) have been widely studied and used for various machine learning tasks which involve sequence modeling, especially when the input and output have variable lengths. For example, in the domain of car-complaints we use RNN to classify customer complaints into six categories such as Transmission Problems, Gear Problems, Windows-Windshield Problems,

¹<https://youtu.be/vTJCoZltiGs>

Algorithm 1 Algorithm for Tailoring Chatbot using Latent Beliefs and Epistemic Reasoning

Require: Complaints dataset D
Require: ML_1model of complaint_classes $C_1, C_2, C_3, \dots, C_n$
Require: ML_2model of opinions $O_1, O_2, O_3, \dots, O_m$
Require: BELIEFS_rulebase consisting of latent belief rules
Require: EPISTEMIC_rulebase consisting of epistemic rules
Require: FSM for dialog
for all complaints c **do**
 Classify Compliant using ML_1model (*engine, transmission* etc.)
 Extract facts (information) using Compliant and Class C_i
 Extract customer opinions using $ML_2modelO_m$ (*neg, strongneg* etc.)
 Using Information Extraction, Domain Knowledge, BELIEFS_rulebase, assert facts F_i
 for all states s_i of FSM **do**
 Using Domain Ontology and Nearest Neighborhood approach, update w_i
 Find Latent Beliefs using complaint class C_i and facts
 Using facts F_i and EPISTEMIC_rulebase, fill slots of Dialog manager Frames
 update weight w_i of state s_i using Latent Beliefs and Epistemic Reasoning
 end for
 while consumer_complaint not addressed **do**
Ensure: Meaningful dialog states
 Using threshold and weights, build meaningful states of FSM (ask_states, skip_states)
 end while
end for

Engine failure Problems, Wheels-Hubs Problems and AC-Heater Problems. A complaint such as: *Slowly the noise from my wheel bearings grew and grew and is now unbearable* is a customer complaint that belongs to the category of wheel problems. The categorization helps in further information extraction and belief processing.

2.2 Extracting Facts from Complaints

We parse a complaint description through Dependency parsers (such as Stanford-CoreNLP, GATE, MITIE etc.) and extract triples from the description by focusing on the dependencies identified among nouns and verbs [Cunningham *et al.*, 2011; Stanford, 2011]. For example, for a description *“my car just died on me”*, triples such as (my-car, just-died-on, me) are extracted.

Once the triples are extracted, we use a hand-crafted fact-assertion rulebase to assert facts implied by the triples. This is done by evaluating the triples in the context of a car ontology, synonym-and-slang dictionary, information-extraction patterns that are relevant for the category of the complaint, and by triggering the fact-assertion rules. For instance, the fact asserted for the above example would be

“car(not-running)”. The Information Extractor also fills slots in a Frame relevant for the category that has been determined.

2.3 Extracting Latent Beliefs

Once the facts from the complaints are asserted, a hand-crafted belief knowledge base evaluates the facts, the slot values together with the category of the complaint, and asserts the probable beliefs held by the human agent in the complaint. For example, in a car-complaints domain, given the initial customer complaint *“Drove vehicle out of town. While on highway vehicle engine just quit. Made an emergency stop because car lost power. While on side of the road vehicle would not start! Vehicle was towed in.”*, beliefs such as “B1) Car was working before B2) Engine not working now B3) Cannot use car” can be asserted.

2.4 Evaluating the Epistemic Rules

The probable beliefs and the facts extracted trigger the epistemic rules. For example in the car complaints domain rules such as “R1) Belief (Car was working before) and Category(Complete Engine Failure) => Knows-Agent (Car does not start now), Knows-Agent (Engine does not crank), Knows-Agent (Engine is dead)” asserts facts about the current epistemic state of the agent.

2.5 Tailoring the Dialog

The facts asserted by the epistemic rules help tailor the dialog. For the above example we have the tailoring as follows: “ A1) Skip car starting FSM states, A2) Skip engine cranking FSM states, A3) Enquire when the engine stopped working, A4) Enquire noise and other signs that occurred when the car stopped working, A5) Enquire on the state of the battery, A6) Enquire when the servicing was done last”. This makes the dialog more relevant for the customer and his/her belief set.

2.6 Bot’s Response and Further Dialog

The Bot then asks relevant questions to the customer. The customer response is “yes/no” type or a set of new sentences. If the response is a new set of sentences then the fact-extraction and belief revision followed by epistemic rules evaluation is repeated.

3 Results and Discussion

In our experiments we consider 61,252 complaints about car faults from <http://www.carcomplaints.com/>. We consider complaints across six categories. The test accuracy of complaint categorization with Vanilla LSTM (Long Short Term Memory) was 68.47%. The beliefs and epistemic rules helped tailor the dialog to the customer expectations. In this work, we demonstrate the overall architecture where we use RNN based classification of customer complaints as an input to the epistemic rule engine. Our approach is generic and can be applied easily in any other domains such as complaints about hardware/software issues etc.

References

- [Cunningham *et al.*, 2011] Hamish Cunningham, Diana Maynard, Kalina Bontcheva, Valentin Tablan, Niraj Aswani, Ian Roberts, Genevieve Gorrell, Adam Funk, Angus Roberts, Danica Damljanovic, Thomas Heitz, Mark A. Greenwood, Horacio Saggion, Johann Petrak, Yaoyong Li, and Wim Peters. *Text Processing with GATE (Version 6)*. 2011.
- [Henderson *et al.*, 2014] Matthew Henderson, Blaise Thomson, and Steve Young. Word-based dialog state tracking with recurrent neural networks. In *Proceedings of the 15th SIGDIAL*, pages 292–299, 2014.
- [Lemon *et al.*, 2006] Oliver Lemon, Xingkun Liu, Daniel Shapiro, and Carl Tollander. Hierarchical reinforcement learning of dialogue policies in a development environment for dialogue systems: Reall-dude. In *BRANDIAL’06*, pages 185–186, 2006.
- [Miller *et al.*, 2016] Alexander Miller, Adam Fisch, Jesse Dodge, Amir-Hossein Karimi, Antoine Bordes, and Jason Weston. Key-value memory networks for directly reading documents. *arXiv preprint arXiv:1606.03126*, 2016.
- [Serban *et al.*,] Iulian V Serban, Alessandro Sordoni, Yoshua Bengio, Aaron Courville, and Joelle Pineau. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Proceedings of the 30th AAAI-16*.
- [Stanford, 2011] Stanford. Corenlp, 2011.
- [Wang and Lemon, 2013] Zhuoran Wang and Oliver Lemon. A simple and generic belief tracking mechanism for the dialog state tracking challenge: On the believability of observed information. In *Proceedings of the SIGDIAL 2013 Conference*, pages 423–432, 2013.
- [Weston, 2016] Jason Weston. Dialog-based language learning. *CoRR*, abs/1604.06045, 2016.
- [Williams *et al.*, 2016] Jason Williams, Antoine Raux, and Matthew Henderson. The dialog state tracking challenge series: A review. *Dialogue & Discourse*, 2016.
- [Young *et al.*, 2013] Steve Young, Milica Gašić, Blaise Thomson, and Jason D Williams. Pomdp-based statistical spoken dialog systems: A review. *Proceedings of the IEEE*, 2013.