Human–Computer Interaction Series

The Human–Computer Interaction Series, launched in 2004, publishes books that advance the science and technology of developing systems which are effective and satisfying for people in a wide variety of contexts. Titles focus on theoretical perspectives (such as formal approaches drawn from a variety of behavioural sciences), practical approaches (such as techniques for effectively integrating user needs in system development), and social issues (such as the determinants of utility, usability and acceptability).

HCI is a multidisciplinary field and focuses on the human aspects in the development of computer technology. As technology becomes increasingly more pervasive the need to take a human-centred approach in the design and development of computer-based systems becomes ever more important.

Titles published within the Human–Computer Interaction Series are included in Thomson Reuters' Book Citation Index, The DBLP Computer Science Bibliography and The HCI Bibliography.

More information about this series at http://www.springer.com/series/6033

Boris Galitsky

Artificial Intelligence for Customer Relationship Management

Solving Customer Problems

Springer

Boris Galitsky
Oracle Labs
Redwood Shores, CA, USA

ISSN 1571-5035 ISSN 2524-4477 (electronic)
Human–Computer Interaction Series
ISBN 978-3-030-61640-3 ISBN 978-3-030-61641-0 (eBook)
https://doi.org/10.1007/978-3-030-61641-0

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface to Volume 2

The second volume of the book on Artificial Intelligence for Customer Relationship Management presents a broad spectrum of application domains with a focus on solving customer problems. We design a chatbot to specifically address issues customer experiences, as well as a system that tries to understand the customer complaint, his mood, and what can be done to resolve an issue with the product or service.

In the first volume of the book, we addressed the key issues of modern CRM systems, which frequently cannot handle the high volume of calls and chats. The main limitations of CRM today is a lack of machine understanding of what a customer is trying to say, what are the issues he is trying to communicate or attempts to conceal. To attack these limitations, we proposed a number of Natural Language Understanding (NLU) techniques with the focus on deep language analysis and reasoning.

Nothing repels customers as much as irrelevant, annoying, intrusive answers and low-quality content. The inability to find information users need is a major reason they are unsatisfied with a producer of a provider. In the first volume, we shared the evaluations of question-answering components, verified their explainability features and observed a satisfactory lab performance for major CRM tasks. We concluded that keeping the customers informed is key to their acquisition and retention. For example, in Fig. 1, fast access to information on who owns the bird is critical for a smooth transaction.

In this second volume, we make the next step to the heart of the customer retention issue: solving a customer problem. Now just understanding what the customer is saying is not enough anymore: we need to involve a formalized knowledge about a product and reason about it. To solve a customer problem efficiently and effectively, we focus on dialogue management. A dialogue with the customer is maintained in such a way so that the problem can be clarified and multiple ways to fix it can be sought.

If the CRM is not intelligent enough, it is not very usable. In particular, CRMs are used by people who sell, and as such, they are often traveling to customer sites, away from their desk, in front of clients, and are paid to generates sales, so communication with a CRM is often an afterthought. Salespeople are not hired for their computer or data entry skills but instead for their ability to sell. If they close a significant deal and

Fig. 1 Communicating specific features of products with a buyer

Fig. 2 Charles Fillmore's example for joint sentence discourse

never entered it into the CRM, something may be said after the congratulations on the sale, but no one would complain about not using the CRM. The general consensus among sales representatives is simple: no one likes dumb CRM. They are frequently overly-complex, clunky systems without a conversational interface that do not help them sell and even obstruct the sales process. So the personnel may use unintelligent CRM as little as possible. More than a third of businesses face low adoption rates of CRM systems.

Our dialogue management is based on discourse analysis, a systematic linguistic way to handle the thought process of an author. Discourse analysis is a method for studying the natural language in relation to its social context. Discourse analysis tracks how language is used in real-life situations. A well-known American linguist Charles Fillmore demonstrates that two sentences taken together as a single discourse can have meanings different from each one taken separately (Fig. 2).

The objects of discourse analysis (texts, conversations, communicative events) are defined in terms of coherent sequences of sentences, propositions and speech acts. Sometimes discourse analysis even helps to get a consumer out of trouble by making a conversation convincing (Fig. 3). A dialogue structure can take a peculiar form for a conversation between two apple maggots (Fig. 4).

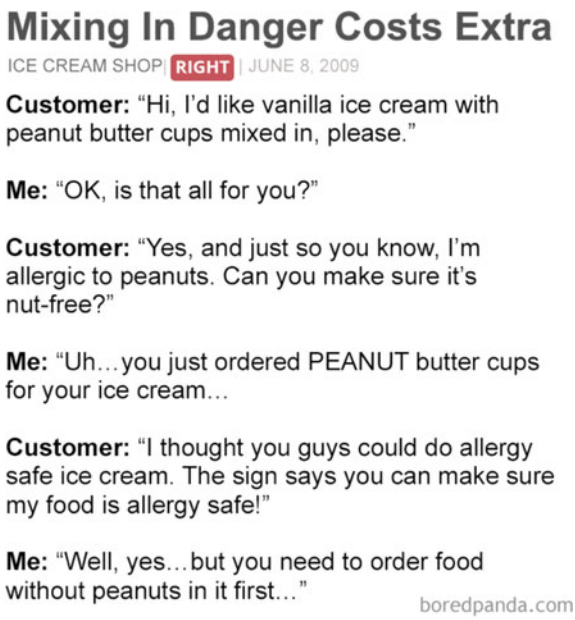

Fig. 3 Finding a contradiction in customer's request and communicating it

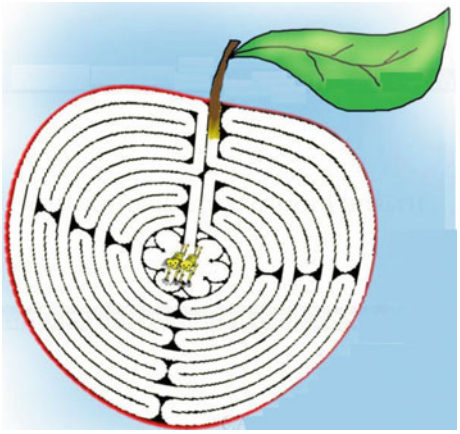

Fig. 4 A conversational structure of two apple maggots

We demonstrate how a dialogue structure can be built from an initial utterance. We also introduce real and imaginary discourse trees as a means to represent an explicit and implicit discourse of text. A problem of involving background knowledge on-demand, answering questions is addressed as well. We outline the *Doc2Dialogue* algorithm for converting a paragraph of text into a hypothetical dialogue based on an analysis of a discourse tree for this paragraph. This technique allows for a substantial extension of chatbot training datasets in an arbitrary domain.

Fig. 5 An example of hypocrisy

We compute user sentiments and personal traits to tailor dialogue management and content to individual customers. We also design several dialogue scenarios for CRM with replies following certain patterns and propose virtual and social dialogues for various modalities of communication with a customer.

To detect fake content, deception and hypocrisy (Fig. 5), we need to analyze associated discourse patterns, which turned out to differ from the ones for genuine, honest writing. With discourse analysis, we can zoom in further and characterize customer complaints with respect to the best way to resolve them. We simulate the mental states, attitudes and emotions of a complainant and try to predict his behavior. Having suggested graph-based formal representations of complaint scenarios, we machine-learn them to identify the best action the customer support organization can choose to retain the complainant as a customer.

Customer complaints are classified as valid (requiring some kind of compensation) or invalid (requiring reassuring and calming down) the customer. Scenarios are represented by directed graphs with labeled vertices (for communicative actions) and arcs (for temporal and causal relationships between these actions and their parameters). The classification of a scenario is computed by comparing a partial matching of its graph with graphs of positive and negative examples. We illustrate machine learning of graph structures using the Nearest

Neighbor approach as well as concept learning, which minimizes the number of false negatives and takes advantage of a more accurate way of matching sequences of communicative actions.

Redwood Shores, CA, USA Boris Galitsky

Acknowledgements

The author is grateful to Dmitri Ilvovsky, Tatyana Machalova, Saveli Goldberg, Sergey O. Kuznetsov, Dina Pisarevskaya and other co-authors for fruitful discussions on the topics of this book.

The author appreciates the help of his colleagues from the Digital Assistant team at Oracle Corp. Gautam Singaraju, Vishal Vishnoi, Anfernee Xu, Stephen McRitchie, Saba Teserra, Jay Taylor, Sri Gadde, Sundararaman Shenbagam and Sanga Viswanathan.

The author acknowledges substantial contribution of the legal team at Oracle to make this book more readable, thorough and comprehensive. Kim Kanzaki, Stephen Due, Mark Mathison and Cindy Rickett worked on the patents described in this book and stimulated a lot of ideas which found implementation in this book.

Dmitry Ilvovsky contributed to the book Chap. 1 Volume 2 within the framework of the National Research University Higher School of Economics Basic Research Program and funded by the Russian Academic Excellence Project '5–100'.

Contents