

Chapter 5

Deep Learning for Sentiment Analysis: An Overview and Perspectives

Vincent Karas

University of Augsburg, Germany

Björn W. Schuller

University of Augsburg, Germany

ABSTRACT

Sentiment analysis is an important area of natural language processing that can help inform business decisions by extracting sentiment information from documents. The purpose of this chapter is to introduce the reader to selected concepts and methods of deep learning and show how deep models can be used to increase performance in sentiment analysis. It discusses the latest advances in the field and covers topics including traditional sentiment analysis approaches, the fundamentals of sentence modelling, popular neural network architectures, autoencoders, attention modelling, transformers, data augmentation methods, the benefits of transfer learning, the potential of adversarial networks, and perspectives on explainable AI. The authors' intent is that through this chapter, the reader can gain an understanding of recent developments in this area as well as current trends and potentials for future research.

INTRODUCTION

In recent years, the amount of information available on the Internet has grown rapidly. At the beginning of 2019, Twitter had 326 million monthly active users, and 500 million tweets were sent per day Cooper (2019). Facebook, the largest social media platform, reported 2.41 billion monthly active users for the second quarter of 2019 Facebook (2019). Every minute, 4.5 million YouTube videos and 1 million Twitch videos are viewed, and the Google search engine processes 3.8 million queries (Desjardins, 2019). This trove of online content constitutes a valuable resource for business applications, e.g. for providing the users with personalised search recommendations and tailored advertisements. If the data is harnessed

DOI: 10.4018/978-1-7998-4240-8.ch005

properly, it may deliver new insights that can help improve existing products and services and inspire future business models. Among the available content, text, in particular, is rich in information, as it can contain nuanced emotions, multiple layers of meaning and ambiguities. However, this complexity also results in it being challenging to analyse. Natural Language Processing (NLP), which addresses this challenge, has become a popular field of research.

Sentiment Analysis (SA), which is often also referred to as opinion mining or comment mining in the literature, is a discipline of NLP-based text analysis whose goal is to determine the writer's feelings about a particular topic. Emotions have been shown to play an essential role in human decision making (Bechara, Damasio, & Damasio, 2000) and behaviour in general. Consequentially, SA has many conceivable applications in business and academia. Examples include companies looking to improve their services by automatically assessing customer reviews (Hu & Liu, 2004), (Zvarevashe & Olugbara, 2018), comparing products online, or analysing newspaper headlines (Rameshbhai & Paulose, 2019).

Sentiment also plays an important role in the financial market. Ranjit, Shrestha, Subedi, and Shakya (2018) used SA to predict the exchange rates of foreign currencies. Shah, Isah, and Zulkernine (2018) predicted stock prices in the pharmaceutical industry based on the sentiment in news coverage. C. Du, Tsai, and Wang (2019) classified financial reports in terms of expected financial risk using SA.

In addition, there are medical applications for SA. Müller and Salathé (2019) introduced an open platform for tracking health trends on social media. Luo, Zimet, and Shah (2019) created an NLP framework to investigate sentiment fluctuation on the subject of HPV vaccination, expressed by Twitter users between 2008 and 2017.

Furthermore, political analysts and campaigns can benefit from mining the opinions and emotions expressed towards candidates, issues and parties on social media. Jose and Chooralil (2016) used an ensemble classifier approach to predict results of the 2015 election in Delhi. Joyce and Deng (2017) applied SA to tweets collected in the run-up to the 2016 US presidential election and compared them to polling data. They found that automatic labelling of tweets outperformed manual labelling.

Many tools used in sentiment analysis are designed for a specific application, which negatively impacts their diffusion. Joshi and Simon (2018) introduced a cloud-based open-source tool which provides various APIs in order to perform SA on data from arbitrary sources.

While SA has attracted considerable attention, the field still faces challenges. These include domain dependence, negations, handling fake reviews (Hussein, Doaa Mohey El-Din Mohamed, 2018), as well as incorporating context, dealing with data imbalance and ensuring high-quality annotations (Boaz Shmueli & Lun-Wei Ku, 2019).

This chapter introduces the reader to selected methods used for sentiment analysis, with a focus on techniques based on deep learning. Its contribution consists of a discussion of the latest advances in the state of the art, as well as an outlook concerning ongoing trends in the field and recommendations on future research directions.

The rest of the chapter is structured as follows. In the next section, the fundamentals of SA and select machine learning concepts are presented. Topics covered include a categorisation of analysis approaches by level of granularity, how to measure sentiment, traditional sentiment analysis methods employing lexica and machine learning, as well as tools for word embedding and sentence modelling such as autoencoders, GloVe, fastText and Word2vec. The chapter will then continue with its main section, focusing on current developments in deep learning-based SA. Topics include popular neural network architectures and their combination into hybrid models, capturing contextual information by adding attention, Transformer networks and the challenges and benefits of transfer learning. In the fol-

34 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:
www.igi-global.com/chapter/deep-learning-for-sentiment-analysis/259785

Related Content

Location Extraction to Inform a Spanish-Speaking Community About Traffic Incidents

Alejandro Requejo Flores, Alejandro Ruiz, Ricardo Marand Raúl Porras (2021). *Handbook of Research on Natural Language Processing and Smart Service Systems* (pp. 347-367).

www.irma-international.org/chapter/location-extraction-to-inform-a-spanish-speaking-community-about-traffic-incidents/263110

Revolutionizing Conversational AI: Unleashing the Power of ChatGPT-Based Applications in Generative AI and Natural Language Processing

C. V. Suresh Babu and P. M. Akshara (2024). *Advanced Applications of Generative AI and Natural Language Processing Models* (pp. 228-248).

www.irma-international.org/chapter/revolutionizing-conversational-ai/335840

Emotion Mining Using Semantic Similarity

Rafiya Jan and Afaq Alam Khan (2020). *Natural Language Processing: Concepts, Methodologies, Tools, and Applications* (pp. 1115-1138).

www.irma-international.org/chapter/emotion-mining-using-semantic-similarity/239981

Sentiment Analysis of Arabic Documents: Main Challenges and Recent Advances

Hichem Rahab, Mahieddine Djoudi and Abdelhafid Zitouni (2021). *Natural Language Processing for Global and Local Business* (pp. 307-331).

www.irma-international.org/chapter/sentiment-analysis-of-arabic-documents/259795

Social Commerce Recommendation Systems: Leveraging User Behaviour and Preferences

Aanchal Taliwal, Mimansa Pathania and Mitali Chugh (2024). *Advanced Applications of Generative AI and Natural Language Processing Models* (pp. 265-285).

www.irma-international.org/chapter/social-commerce-recommendation-systems/335842