MARKETING
ZFP – Journal of Research and Management

Executive Summary

Sentiment Analysis in Marketing – From Fundamentals to State-of-the-art

By *Andreas Karasenko*

This study aims to guide researchers and practitioners in their sentiment analysis (SA) model selection and usage. The paper provides a pipeline for SA with step-by-step instructions for all relevant parts of the SA pipeline such as: web-scraping, preprocessing, evaluation methods, hyperparameter optimization and estimation, and optionally model interpretability. To validate this pipeline, we subsequently conduct case study in the form of a quantitative method comparison. We use both scraped data from the Google Play Store, as well as benchmark datasets commonly used in SA literature and validate our results against established literature. Finally, we used two popular frameworks (SHAP and LIME) to provide model interpretability to the otherwise black-box machine learning models. Based on this we derive model strengths and weaknesses as well as practical guidelines for model selection.

A quantitative analysis was conducted using more than 300.000 user generated texts from diverse backgrounds such as Amazon, Twitter, Google Play Store. We use the presented pipeline to show the usage of traditional machine learning models (e.g.: Naïve Bayes, Random Forests, Logistic Regression), deep learning models (e.g.: Convolutional Neural Networks), transfer learning models (e.g. Bidirectional Encoder Representations from Transformers), and state of the art few-shot learning approaches using Large Language Models (e.g.: ChatGPT). The empirical evaluation showed results that are in line with established literature: (1) modern transfer learning models, especially those further trained on a broad SA task, outperform all other approaches; (2) deep learning methods outperform traditional approaches; (3) few-shot learning and prompting approaches perform significantly worse than all other approaches which is corroborated by the computer science literature and contradicts current marketing research.

We then demonstrated how frameworks such as SHAP and LIME can be used to interpret the sentiment classification decision of a machine learning model for a given text. Both approaches show that words such as “savings” and “recommend” indicate a more positive sentiment for a positive review. And similarly sentence pieces in a negative review like “forced update” correctly indicate a more negative review.

To conclude, our paper presents a practical introduction into SA, showcasing application potential both in literature and industry. We explicated every step necessary to conduct thorough model comparisons and provide code snippets as well as end-to-end examples that can be used to jump-start SA research projects.