Abstract

Aspect-Based Sentiment Analysis (ABSA) aims to identify the sentiment polarity associated with specific aspects mentioned in review texts, such as service, food, or price in restaurant reviews. Unlike conventional sentiment classification that focuses on overall opinion at the sentence or document level, ABSA requires fine-grained analysis to capture how evaluative expressions relate to individual aspects. By uncovering aspect-level opinions, ABSA enables more precise understanding of user attitudes, supporting applications such as product improvement, market monitoring, and decision-making based on detailed customer feedback. Despite the success of deep neural architectures and pre-trained language models, ABSA systems remain highly domain-dependent. When an ABSA model trained on a source domain such as restaurant, its performance often degrades on another different domain like laptop, because two different domains usually vary significantly in vocabulary, sentiment expressions, and aspect semantics.

To address this issue, domain adaptation has been extensively studied to enable knowledge transfer across domains. However, conventional approaches typically assume the availability of labeled target-domain data during training, which is often unrealistic in practical scenarios. Furthermore, they overlook the fact that aspect-level variation itself can serve as an additional dimension of domain shift. Different aspects within a single domain, such as service and price, follow distinct linguistic patterns and sentiment distributions. Consequently, achieving robust adaptation requires a framework capable of transferring sentiment–aspect knowledge not only across domains but also across aspects within a domain.

In this dissertation, we investigate two complementary adaptation scenarios designed to overcome these challenges, each reflecting a different level of data availability and generalization difficulty. The first, referred to as Unsupervised Domain Adaptation (Scenario A), assumes that labeled target data are not accessible during training, but unlabeled target data are available. In this scenario, the domain boundary is redefined from the conventional dataset-based perspective to an aspect-oriented one, in which each aspect (such as service, food, or price) is treated as a domain. The key challenge lies in transferring sentiment knowledge from labeled source aspect to unlabeled target aspect, where linguistic expressions and sentiment tendencies vary considerably. The second scenario is called Domain-Agnostic Adaptation (Scenario B), which removes the dependence on target-domain data entirely. Here, the model is trained only on labeled data from multiple

source domains, and generalize to unseen target domains. This formulation represents a more extreme yet realistic condition where no target data, even unlabeled, are available in advance. The focus of this setting is to build a system that can internalize aspect-level knowledge and apply it universally. To tackle these two scenarios, this dissertation proposes two complementary frameworks tailored to each setting.

In Scenario A, where unlabeled target-aspect data are available, each aspect is regarded as a distinct domain, and the goal is to transfer sentiment knowledge from labeled source domain to unlabeled target domain. The proposed framework is divided into two methods for training data construction: Pseudo-Label assignment for unlabeled target domain data and Cross-Aspect Review Generation (CARG). Specifically, a base classifier that fine-tuned on the labeled source domain is trained, and is used to predict the sentiment of unlabeled target domain. Reviews whose maximum class probability exceeds an empirically determined threshold are accepted as reliable pseudo-labeled instances. In the CARG stage, labeled target reviews are automatically generated by replacing sentiment words and keywords in sentences of the source domain. Finally, the labeled data constructed by these two methods are used to train a classifier for determining the polarity of the aspects in the target domain. Experiments on several aspect domains from restaurant and laptop dataset demonstrate improved accuracy and macro-F1 scores of polarity classification. Additional analyses are conducted on threshold selection for pseudo-labeling and sentiment word scoring, and qualitative examples of valid and invalid generations.

In Scenario B, where no target-domain data are available during training, this dissertation proposes an Aspect-Enhanced Prompting (AEP) framework for domain-agnostic adaptation in aspect sentiment classification. The proposed AEP framework is based on two generative language models: one generates a prompt from a given review, while the other follows the prompt and classifies the sentiment of an aspect. The first model extracts Aspect-Related Features (ARFs), which are words closely related to the aspect, from the review and incorporates them into the prompt in a domain-agnostic manner, thereby directing the second model to identify the sentiment accurately. Our framework incorporates an innovative rescoring mechanism and a cluster-based prompt expansion strategy. Both are intended to enhance the robustness of the generation of the prompt and the adaptability of the model to diverse domains. The results of experiments conducted on five datasets (Restaurant, Laptop, Device, Service, and Location) demonstrate that our method outperforms the baselines, including a state-of-the-art unsupervised domain adaptation method. The effectiveness of both the rescoring mechanism and the cluster-based prompt expansion is also validated through an

ablation study.

KeyWords: aspect-based sentiment analysis, unsupervised domain adaptation, multi-source domain adaptation, data argumentation, text generation model, prompt engineering