School of Informatics

Informatics Research Review
Creating Targeted Advertisements Using Social Media Data

[redacted]

January 2021

Abstract

Over the last few decades, we have gradually shifted our social lives to online networks. This entails making available massive amounts of personal data for social media companies to effectively create targeted advertisements based on user activity and behaviour. The review aims to discuss the data mined from user profiles, the techniques and tools used to analyse the data, and how to improve the likelihood of interaction of the users with the suggested ads. The review also mentions recent advances in methodologies through appropriate case studies.

Date: Friday 22nd January, 2021

Supervisor: [redacted]

1 Introduction

1.1 Motivation & Aim

Social media can be seen as an extension of social lives online. From staying in touch with friends and family to applying to jobs, the impact of social media on our daily activities is irrefutable. While social media would be formally defined as a group of Internet based applications which allow users to create and share content [1], one would more commonly interpret it as an umbrella term for websites like Facebook, Twitter, Instagram, etc. Each of these sites primarily generate different kinds of user content: videos (YouTube), pictures (Instagram), or a combination (Facebook). Despite these differences, there are two main things in common: connecting people and providing a platform for online marketing through advertisements.

As the world grows more connected while using these platforms, so does the potential of online marketing and advertising. These advertisements often act as a source of revenue for the social media sites. This acts as a major incentive to track user activity and analyse user data in order to create and deliver relevant advertisements to the user [2]. The more the relevance, the more likely would the user be to engage in the product or service offered by the advert.

It is critical to understand how the ad targeting mechanisms work because of the involvement of large scale personal data aiding monetary gain for these networks. This review aims to provide a consolidated overview on advertising in social media, focusing on methods, tools, techniques, and recent case studies. In particular, we cover the following questions:

- What are the main methods used in recent years by to display relevant advertisements?
- What techniques and tools, particularly with respect to Natural Language Processing, are used? Are they used in isolation or combination?
- Apart from text in a user's posts, what other aspects of the user's data can be mined to create better ads?
- What persuades a user to interact with an advertisement given it is relevant? Are there ways to increase this likelihood?

The next two subsections focus on the structure of the paper and related topics which will not be discussed in this review. The final part mentions the background required for a comprehensive understanding of the review.

1.2 Structure

We divide the paper into four main parts: Introduction, Literature Review, Summary and Conclusion, and Future Work.

Section 2 is further divided into four parts:
2.1 opens with different approaches used for online advertising, how these can be extended to social media, and what aspects of user data are needed to implement these;
2.2 discusses some novel NLP techniques used for Sentiment Analysis and Opinion Mining;
2.3 looks at two implementations of advertisement generation using the aforementioned NLP methods;
2.4 introduces the concept of *social context*, emphasizing that simply suggesting relevant adverts

to users is not sufficient, and social context is needed to improve the effectiveness of targeted ads.

Section 3 concludes the review and summarizes the key findings.

Section 4 briefly mentions future work which can be done in relevant, and relatively unexplored subdomains.

1.3 Related Topics

Significant efforts have been made to narrow down the scope of this review. Here we list some topics closely related to the review which will not be covered:

- Though advertisements are used to generate revenue, we do not discuss the business models used by social media companies
- We refrain from discussing Natural Language Processing (NLP) techniques in depth; we only mention them in the context of online advertising but do not discuss their inner workings
- Creating advertisements using user data poses many legal and ethical questions regarding data collection and usage. We also keep such discussions out of the review

1.4 Background

The intended audience for the review is undergraduate students studying fields related to Computer Science or Linguistics. The review should still be fairly comprehensible to those outside of these fields. We also aim this review at those without a strong mathematical background, and have generally summarized any technical aspects of research findings. Here we list some key topics essential to the review which the reader should be familiar with.

- Familiarity of social media and basic terms like *users*, *posts*, *likes*, *shares*, etc is essential
- A good understanding of foundation level Natural Language Processing is needed; an introductory course should suffice
- Conceptual level knowledge of Sentiment Analysis and Opinion Mining; understanding of detailed implementation not required
- Familiarity of basic Machine Learning algorithms used for text processing and representations is expected
- Mathematical concepts related to the algorithms are not necessary but would be recommended

2 Literature Review

2.1 Understanding the Data

Advertisements are not restricted to social media and can be found on most websites. The types can vary significantly from simple text and links to pictures and videos.

Fan & Chang (2010) mention two main categories for text-based advertisements: *Sponsored Search* and *Content-based*. Sponsored Search relies on the user specifically querying certain keywords based on which a list of advertisements are triggered and displayed along with results of the user's query. Content-based ads, on the other hand, analyse the contents of a given page and display ads based on what is already present [3]. This is often related to what the user has queried, but is independent of the actual query. A notable example of Content-based advertising would be Google AdSense[1], which analyses multiple pages to display relevant advertisements.

Though the authors made this categorization for text-based ads, we believe this can be extended to other forms such as pictures and videos which are more common on social media platforms. Google AdSense already achieves this on other websites. Comparing the two, we find Content-based ads to be a more appropriate choice for such platforms because of two main reasons:

1. The primary goal of social media is to build and sustain connections online, not to shop. The search bar is used to find people rather than products. Thus using user queries would not be very effective. [4]
2. (a) Users often use these platforms to express their opinions and share their experiences about products and services as posts or comments [5].
   (b) Posts that are displayed to the user on their *feed* or *timeline* are already customized to their preferences.

   This means the on-screen content can directly be used to show them relevant ads.

Since targeted ads must come from the user's feed, we must identify what data can be used from the feed. A feed is generally a stream of posts. Posts themselves comprise of text, media (such as pictures and videos), comments, likes, and shares. The feed is determined by algorithms (such Facebook's EdgeRank[2]) which rank the possible posts to be displayed by their importance to the user [6]. Companies usually have a Facebook page which acts much like a user profile. Posts made using these pages show up on a user's feed as part of it[6].

The question, then, is what data can we use from the feed to tailor the ads? Many elements such as likes, shares, user interaction with other users, comments, etc can be mined to extract a user's behavioural patterns [4]. For the moment, we restrict ourselves to only textual data, i.e., comments and the text from posts.

We do this, firstly, because different platforms offer different kind of media as the primary element of a post. For example, YouTube is entirely video based, Instagram is entirely picture based, Twitter is primarily text based, etc. Despite these differences in the primary content, all social media sites allow some form of textual data. Secondly, textual data is often easier to scrape and analyse as opposed to photos and videos; these require more complex processes and techniques.

From the text, we can identify various products and services that are of interest to the user, perhaps by using Named Entity Recognition. Fan & Chang (2010) argue that displaying ads merely relevant to the context is not sufficient. A user writing a long post about a fast-food joint may actually be talking about health concerns instead of promoting it. In this case, it would not make sense to suggest more fast-food brands.

The sentiment associated with content must be considered so as to not make suggestions based on negative correlations [3]. The framework proposed by Fan & Chang (2010) to achieve this is

[1] https://www.google.com/adsense/
[2] http://edgerank.net/

called SOCA - Sentiment Oriented Contextual Advertising. Before discussing the results of this paper, we introduce two closely related NLP techniques used heavily in web pages to generate ads: Sentiment Analysis and Opinion Mining.

2.2 Natural Language Processing Techniques and Recent Advances

Often used to determine how the author of a text *feels*, there is much debate in academia about the difference in definitions of Sentiment Analysis and Opinion Mining. Since both fields use Natural Language Processing and Data Mining techniques, they are often used interchangeably [7][8]. Pang & Lillian (2008) use the two terms to mean the same while Cambria et al. (2013) argue that Opinion Mining is used for polarity detection and Sentiment Analysis is used for emotion recognition.

Opinion Mining and Sentiment Analysis can be seen as subcategories of Subjectivity Analysis, and often share common works in their fields [8]. In this review, we will mostly use the two terms interchangeably while recognizing the subtle differences. This is to account for the different interpretations of the terminologies across different literature.

Though we claimed that analysing text is easier, it is by no means easy. Social Media in particular suffers from heavy slang usage, abbreviations, ungrammatical sentences, etc. This problem further worsens when a character limit is enforced, such as on Twitter [5]. Gokulakrishnan et al. (2012) discuss many social media specific data cleaning techniques on a Twitter data stream. We refrain from talking about such issues, however, since we assume appropriate pre-processing would have been done.

Extracting sentiments or opinions requires extracting meaning from text. One of the simplest representations of textual data is the Bag-of-Words Model where the frequency (or occurrence) of each word in a given document (or text) is captured. The problem with this model is that it is far too simple; we essentially lose all information about the structure of a sentence. Word order and grammar which would otherwise help us to resolve ambiguity or provide context can no longer be used.

Though there are other advanced representations, a novel paradigm was suggested by Poria et al. (2014): the Bag-of-Concepts Model. Instead of keeping a track of word frequency, this keeps track of the different sentiments in a text. They achieve this using *Sentic Computing*, which uses affective and common sense computing by merging disciplines like linguistics, computer science, psychology, sociology [9].

The pipeline proposed by Poria et al. (2014) for detecting polarity is as follows:

1. Text is first converted to a Bag-of-Concepts representation using a semantic parser.
2. Concepts are checked if they are present in the SenticNet[3], which can be defined as a concept-level knowledge base.
3. (a) If there is a match against a SenticNet entry, sentic patterns are applied. Sentic patterns leverage syntactic dependency relations and the SenticNet to compute polarity. Section 4 of Poria et al. [9] explores this depth. We however, refrain from doing so due to the scope of this review.
   (b) If no match is found, an Extreme Learning Machine (ELM) classifier is used to compute the polarity.

[3] https://sentic.net/

Though we do not look at models suggested in other papers in depth, it is imperative to note the model outperformed state-of-the-art methods mentioned in Socher et al. (2012) and Socher et al. (2013) . It is, however, limited by the dependency-based rules and the richness of the knowledge base [9], which leaves much scope for future work.

This pipeline is particularly useful for dealing with words that with different polarities based on the context. Bollegala et al. (2013) define the problem as cross-domain sentiment classification, where each a domain is a different context. For example, the word 'loud' would be considered positive in context of a music player but negative in context of a vacuum cleaner.

In social media, there are no fixed number of contexts or topics that users share. From posting cat videos to recommending movies, the contexts between even two consecutive posts maybe hugely different. A generic classifier which does not consider contexts may fail to extract the true sentiment of the text. We must therefore first identify some context for each text.

Some platforms like Twitter show heavy usage of *hashtags (#)*, which are user defined tags for a post. This makes the task slightly easier since tags can be attributed to different contexts. Where *hashtags* or other tags are not available, we can use some form of a clustering algorithm to find these contexts. Since there are millions of users, we are not limited by either the number of posts or the number of topics (or contexts). The clusters should therefore be fairly accurate.

Bollegala et al. (2013) focus on the problem of training a sentiment classifier on one domain and applying it on another. This is not a particular problem in social media because one can always assume millions of posts to exist in a single domain due to sheer number of posts and the nature of clustering algorithms. Yet, new topics and products emerge on social media from time to time as new trends. Where a context of a text is not inherently clear, it is best to treat it as a new domain.

Bollegala et al. (2013) constructed a Sentiment Sensitive Thesaurus which captures the relatedness of words used in different domains. Words in one domain do not necessarily appear in other. They solve this issue by proposing a method of feature expansion. These approaches are elucidated in sections 4 and 5 of Bollegala et al. (2013). We, once again, leave this out of the review due to space restrictions. The Sentiment Sensitive Thesaurus was compared against SentiWordNet[4] which is a lexical resource used for Opinion Mining and is based on Princeton University's WordNet[5]. The thesaurus produced desirable results, accurately grouping words that expressed similar sentiments.

Both Poria et al. (2014) and Bollegala et al. (2013) define sentiments as polarity. Although most emotions tend to be positive, negative, or neutral, by ignoring individual emotions like joy and anger, this makes the problem of finding user preferences far too generic. Let us consider the example of a user expressing relief towards passing a difficult exam. This would be classified as positive, causing us to suggest textbooks related to the subject of the exam. This is obviously not desirable because even though the user expressed a positive sentiment about a topic, based on the emotion expressed, they may not necessarily be interested in related products.

Unfortunately, most literature only explore sentiments as polarity. We continue this review with the polarity based model mentioned above.

[4]https://github.com/aesuli/SentiWordNet
[5]https://wordnet.princeton.edu/

2.3 Notable Experiments

2.3.1 Sentiment Oriented Contextual Advertising

The Sentiment Oriented Contextual Advertising framework suggested by Fan & Chang (2010) detects user opinions or sentiment by processing the contents of the page, searching a collection of ads and displaying the best possible by ranking them according to relevance. It was designed and carried out using 150 blog pages and around 100,000 advertisements sampled from Google AdSense.

Three processes were carried out, namely Sentiment (or polarity) Detection, Term Expansion, and Page-Ad Matching.

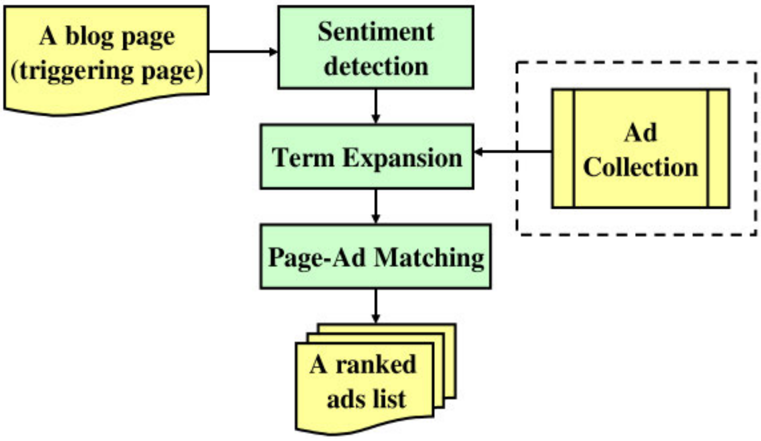

Figure 1: The SOCA Framework [3]

Sentiment Detection was done using two approaches: an SVM algorithm (commonly used for classification and regression) and C4.5 algorithm (a decision tree algorithm which is generally more interpretable). The SVM model only considered the polarity while the C4.5 model also considered the strength of each word and its impact on the overall sentence. The authors claim that the SVM performed better. Although true, the margin is too small too make any decisive conclusions.

The keywords in a page are unlikely to trigger any ads because of small amount of intersection between terms on the page and in ads. Thus the idea of Term Expansion was introduced, which involves finding other words relevant to a given word. For this, only nouns on the page were considered. Subsequently, three approaches were taken:

1. A given term was looked up on WordNet for a synonym
2. Where synonyms could not be found, a relevant entry was looked up on Wikipedia[6]
3. Co-occurences of word pairs to determine their similarity

We believe choosing three different approaches for Term Expansion would be very useful in finding similar words because these approaches all function differently. Moreover, if one fails, one can easily adopt another.

[6] http://en.wikipedia.org

Pages and ads were matched using two measures: the Cosine Similarity measure and Ontology Mapping. The former calculates the angle between the vector representations of two words. The latter refers to the process of determining the mapping between the semantically related concepts of two ontologies. This step was used to create the rankings for the ads. It was found that a combination of Cosine and Ontology yielded the best results.

To evaluate how well the ads were matched to the pages, a gold-standard was developed using human annotators. The top ten ads for each page were compared for three different frameworks: SOCA, CA (Contextual Advertising without sentiment detection), and Google AdSense. The SOCA approach performs the best, followed by CA, and finally Google AdSense.

It is commendable to have outperformed Google AdSense. Fan and Chang chose not emphasize this since Gooogle's ad pool is much larger and diverse. Regardless, the SOCA framework shows significant potential to be adopted in social media. Blogs may be easier to work on due to their static nature in content and longer texts available for the same context. However, we believe that if the key terms are carefully selected and Term Expansion appropriately implemented, ads could be displayed with at least a similar accuracy on social media.

2.3.2 Advertising in Social Networks Using Opinion Mining

The SOCA framework was tested exclusively on blogs, which under many definitions are not considered social networks. This led to the speculation of whether the framework was substantially robust to be used in social media. First suggested in Dragoni (2017) and later elucidated in Dragoni (2018), the author proposes a model of extraction of opinions from reviews on product sites and using this to present the user relevant ads. It can be seen as an extension of the cross-domain sentiment classification problem mentioned by Bollegala et al. (2013).

Dragoni defines computational advertising as a research field combining techniques from Text Mining, Information Retrieval, and Machine Learning. The paper integrates Opinion Mining with user engagement and interests to display relevant ads. Twitter was chosen as the network to experiment due to the ease of generation of advertising messages and user profiling through timeline analysis [14]. We argue that timeline analysis is not unique to Twitter; it can be carried out on Facebook, Instagram, etc as well. Twitter might actually yield less accurate results due to 140 character limit enforced on all Tweets [5].

The proposed approach was divided into three phases:

1. Extraction of interesting aspects from user reviews clustering semantically correlated labels
2. Computation and aggregation of polarities for a cluster
3. Creation of user profiles by detecting interesting aspects and displaying adverts

The first phase, Aspect Extraction, was performed by obtaining a large corpora (about 35 million) of user generated product reviews from Stanford's SNAP Dataset[7]. Each review was tagged with Part-of-Speech (POS) tags, and dependencies were extracted. Nouns connected with opinion words were labelled as *aspects*, which were then extracted. Vectors containing semantically related terms were clustered together. The accuracy of Aspect Extraction was tested against the standard set by SemEval 2015 Task 12[8]. While this is an unsupervised

[7] http://snap.stanford.edu/

[8] https://alt.qcri.org/semeval2015/task12/

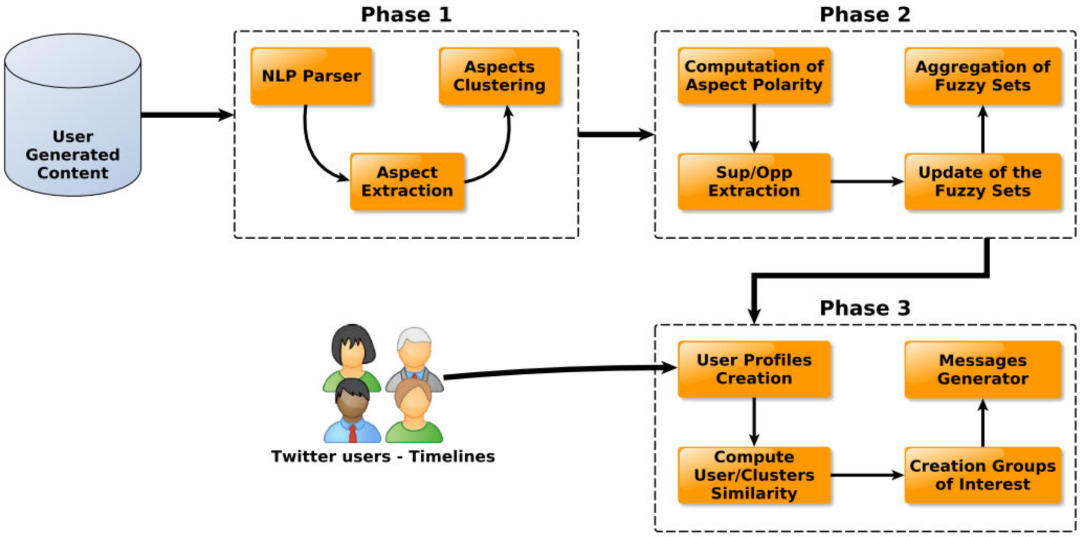

Figure 2: Workflow of the Three-Phase Approach [14]

technique, the performance was comparable to that of the best systems based on supervised methods.

For Polarity Computation, three resources were used: SenticNet, the MPQA vocabulary[9], and the Harvard General Inquirer[10]. The polarity of all aspects were combined to compute the polarity of a single cluster. Polarity of individual aspects was calculated by aggregating fuzzy polarities of opinion words related to that aspect. It was observed that the polarities were more accurate on the Laptop Dataset due to simpler language. This implies that there is a huge scope of improvement for this part of the system. Perhaps Poria et al.'s (2014) Bag-of-Concepts model would yield better results.

User profiles for Twitter users were created by constructing word-frequency paired vectors for nouns identified in the user's posts. The similarities between user profiles and clusters formed from user reviews were computed to determine the most likely preference of products for a given user. Finally, a degree of interest in the range of [0,1] was computed, with 1 being the *most interesting*.

Based on fixed-English templates, messages containing links to products were generated. Around 20,000 accounts were profiled and presented with these advertising messages. The efficiency was tracked by a variation of the Click-through Rate (CTR), which is the ratio of clicks obtained per instances of an advert displayed [4].

Using fixed textual templates for the ads instead of using samples from Google AdSense like Fan & Chang (2010) would probably make it less likely for users to click on the ads. Considering the experiment was run over five days, the wording of the advertisements would become repetitive after a point. The experiment showed favourable outcomes with the *interesting* messages getting a higher number of clicks than the others. However, the author chose to report number of clicks rather than the widely used CTR which may lead to questioning the evaluation mechanism. Nevertheless, the approach shows significant potential to be adopted in mainstream social media, especially in combination with pictures and videos used with the generated text.

[9] http://mpqa.cs.pitt.edu/

[10] http://www.wjh.harvard.edu/ inquirer/homecat.htm

2.4 Social Context

So far we have looked at methods to create and display relevant advertisements to the user. Even if ads are perfectly matched with user preferences, the users are still skeptical to click on an ad assuming that it is spam [4]. Users are also concerned about fraud and breach of privacy [4]. Lee et al. (2018) document that humorous and emotional ads are associated with higher user engagement. Since social media is inherently about building connections, it would be fitting to assume that people would be more likely to approach the advertisements if the ads contained human characteristics [6] or if their friends and followees endorsed the product or service. Li et al. (2014) deduce from their experiments that social context also helps improve the brand image.

Traditional targeted advertisements are based on analysing human behaviour and grouping users into clusters using recommender systems. We believe they can be very powerful in combination with contextual advertising. In addition to profiling users based on their activity, Facebook displays ads that include *social context* in the form of 'Brand X has been liked by 3 of your friends: A, B, and C' [4]. Li et al. (2014) define social context as data related to a user's social circle such as friends and likes received and given on posts.

Who can best influence a user varies vastly. Some are more likely to be influenced by friends while for others it is famous people and celebrities. Analysing the likes, comments and friend circles of users helps us determine whose posts they would be most likely to engage in [6]. One way of finding such influential users is discussed in Erlandsson et al. [15] which uses Association Rule Learning and compares it to other prominent methods. We do not analyse this paper in depth. This power to persuade a user to interact with an advertisement can be seen as a type of social influence [4].

For the rest of the paper, Li et al. (2014) develop a complex framework comprising of the following:

1. User Targeting: Determining the degree of fit of the advertisements and respective social contexts

2. Preference Analysis: Determining the similarity between user preferences and suggested advertisements. The method adopted here is similar to those adopted in other related studies.

3. Quality Analysis: Determining the persuasiveness of the model based on the emphasis of the opinion and profession of the influencer

4. Influence Analysis: Determining the social influence of a user

5. Priority Ranking: Determining which users are mostly likely to make an impactful influence on a target user

We limit our analysis due to the heavy mathematical notation and length limitations.

The Social Context model was compared with other models (Random, In-degree, and Content-based) and outperformed each of them. The authors thus conclude that social context endorsement can significantly increase the Click-though Rate, and improve the user's experience and impression of a brand.

Since this field is relatively new, there is no extensive collection of literature to compare. We believe the proposed framework could be used in combination with the experiments mentioned in

Section 2.3 and NLP techniques proposed in Section 2.2 to yield a state-of-the-art advertisement system. There are limited experiments carried out with respect to advertising in social media that encompass majority of the methodologies suggested here. There is, therefore, significant scope of development and expansion in the future.

3 Summary & Conclusion

As our lives get intertwined with social media, the more personal data is made available for social media sites to mine. This data is used to show advertisements to users which generates revenue for these companies.

We compare *Sponsored Search* ads to *Content-based* ads to gauge what kind of data could possibly be mined. This includes text and media from posts and comments, likes, shares, and list of friends. We focus primarily on the textual content available from a user *feed* because of the simpler models compared to images and videos.

Opinion Mining and Sentiment Analysis, often used to mean polarity detection, are applied to extract meaning from text. In context of advertising, polarity detection is used to understand a user's preference (or lack thereof) to a specific product or a group of products.

We discuss two novel NLP techniques: the Bag-of-Concepts Model and the Cross-Domain Sentiment Sensitive Thesaurus. The Bag-of-Concepts model, as proposed by Poria et al. (2014), captures the concepts in a text instead of word frequencies in a traditional Bag-of-Words Model. These concepts are used in combination with SentiNet to compute polarities. The model is however limited by the dependency-based rules and knowledge base.

The Sentiment Sensitive Thesaurus captures relations of words from different domains. The model gave desirable outcomes tested against SentiWordNet. This would be useful because of the large number of domains in social media and the different meanings associated with the same word across domains.

We then discuss the Sentiment Oriented Contextual Advertising (SOCA) framework in detail, which generates ads from Google AdSense based on contents of a blog page. Compared against Contextual Advertising (without sentiment detection) and Google AdSense, SOCA performed the best. A possible drawback is the fact this was tested on blogs, which are often not considered to be social media. The efficiency of this approach must be tested on social media for future work.

Dragoni's Opinion Mining approach tackles this issue by analysing user reviews and deploying the ads on users on Twitter. The evaluation was done by comparing the number of clicks for each message type instead of using the standard Click-through Rate (CTR). Further, the approach would possibly benefit from not using generic message templates to provide a larger variety of advertisements.

Finally, we discuss the aspect of *social context* in advertising. We find that a user may refrain from interacting with an ad even if it is very relevant over privacy and fraud concerns. Li et al. (2014) introduce the concept of *social context endorsement* where suggested ads are shown to be liked by other users. Testing this against other frameworks which do not use social context, Li et al. find that it is the most effective form of advertising.

In general, we find that most of the literature discussed are quite independent of each other. There is little overlap between the different approaches mentioned. They should possibly be tested in combination on some social network to see if they produce better results. This could be further explored as future work.

4 Future Work

- The Sentic patterns proposed by Poria et al. (2014) is limited by the dependency-based rules and knowledge base. Expanding the knowledge base would help improve the overall model.
- There is no literature comparing static and interactive forms of advertisements like image-based vs gameplay-based. The preference and effectiveness of the two could be demographic based. This would be an interesting idea to explore to further customize ads.
- Majority of advertising related literature strictly deal with polarized sentiments. Future work could be done on the entire sentiment spectrum specifically for the application to advertising.

References

[1] Andreas M. Kaplan and Michael Haenlein. Users of the world, unite! The challenges and opportunities of Social Media. *Business Horizons*, 53(1):59–68, 1 2010.

[2] Danah Boyd and Kate Crawford. Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. *Information Communication and Society*, 15(5):662–679, 6 2012.

[3] Teng Kai Fan and Chia Hui Chang. Sentiment-oriented contextual advertising. *Knowledge and Information Systems*, 23(3):321–344, 6 2010.

[4] Yung-Ming Li, Lienfa Lin, and Shih-Wen Chiu. Enhancing Targeted Advertising with Social Context Endorsement. *International Journal of Electronic Commerce*, 19(1):99–128, 10 2014.

[5] Balakrishnan Gokulakrishnan, Pavalanathan Priyanthan, Thiruchittampalam Ragavan, Nadarajah Prasath, and Ashehan Perera. Opinion mining and sentiment analysis on a Twitter data stream. In *International Conference on Advances in ICT for Emerging Regions, ICTer 2012 - Conference Proceedings*, pages 182–188, 2012.

[6] Dokyun Lee, Kartik Hosanagar, and Harikesh S. Nair. Advertising content and consumer engagement on social media: Evidence from Facebook. *Management Science*, 64(11):5105–5131, 11 2018.

[7] Erik Cambria, Bjorn Schuller, Yunqing Xia, and Catherine Havasi. New avenues in opinion mining and sentiment analysis. Technical Report 2, 2013.

[8] Bo Pang and Lillian Lee. Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, 2(1-2):1–135, 2008.

[9] Soujanya Poria, Erik Cambria, Grégoire Winterstein, and Guang Bin Huang. Sentic patterns: Dependency-based rules for concept-level sentiment analysis. *Knowledge-Based Systems*, 69(1):45–63, 10 2014.

[10] Richard Socher, Brody Huval, Christopher D Manning, and Andrew Y Ng. Semantic Compositionality through Recursive Matrix-Vector Spaces. Technical report, 2012.

[11] Richard Socher, Alex Perelygin, Jean Y Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank. Technical report, 2013.

[12] Danushka Bollegala, David Weir, and John Carroll. Cross-Domain Sentiment Classification Using a Sentiment Sensitive Thesaurus. *IEEE Transactions on Knowledge and Data Engineering*, 25(8):1719–1731, 8 2013.

[13] Mauro Dragoni. A three-phase approach for exploiting opinion mining in computational advertising. *IEEE Intelligent Systems*, 32(3):21–27, 5 2017.

[14] Mauro Dragoni. Computational advertising in social networks: An opinion mining-based approach. In *Proceedings of the ACM Symposium on Applied Computing*, pages 1798–1804, New York, New York, USA, 4 2018. Association for Computing Machinery.

[15] Fredrik Erlandsson, Piotr Bródka, Anton Borg, and Henric Johnson. Finding Influential Users in Social Media Using Association Rule Learning. *Entropy*, 18(5):164, 4 2016.