International Journal of Electronics, Communication & Soft Computing Science and Engineering
ISSN: 2277-9477, Volume 5, Issue 2

Sentiment Analysis: A Comprehensive Review of Traditional, Machine Learning, Deep Learning and Transformer-Based Approaches

Mr. Vaibhav Gajananrao Ingole
Lecturer in Information Technology, Government Polytechnic, Amravati, India-444603
Email:- vingole42@gmail.com

Received on: 25 Oct 2020 **Revised on: 6 Nov 2020** **Published on: 9 Nov 2020**

Abstract- Sentiment analysis has emerged as an important search area in Natural Language Processing (NLP), machine learning, computational linguistics, and artificial intelligence. The objective of sentiment analysis is to automatically identify opinions, emotions, attitudes, and polarity expressed in textual data. The rapid expansion of social media, e-commerce platforms, online reviews, blogs, discussion forums, and digital communication has generated enormous volumes of opinionated content, making automated sentiment analysis increasingly important for organizations and researchers. This paper presents a comprehensive review of sentiment analysis techniques developed up to 2020. The review covers lexicon-based methods, traditional machine-learning algorithms, feature-engineering techniques, deep-learning models, aspect-based sentiment analysis, social-media sentiment analysis, and transformer-based approaches. Particular attention is given to Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM) networks, attention mechanisms, contextual word representations, and Bidirectional Encoder Representations from Transformers (BERT). The paper also reviews commonly used datasets and evaluation measures and discusses challenges involving negation, sarcasm, ambiguity, domain adaptation, implicit aspects, multilingual text, code-mixed language, and noisy social-media content. The review demonstrates the transition from manually engineered features toward contextual representation learning and identifies important research directions for multilingual, explainable, robust, and aspect-aware sentiment analysis.

***Keywords*:** Sentiment Analysis, Opinion Mining, Natural Language Processing, Machine Learning, Deep Learning, LSTM, CNN, BERT, Aspect-Based Sentiment Analysis, Social Media.

I. INTRODUCTION

The widespread adoption of Internet-based communication has resulted in an enormous increase in user-generated textual content. Customers publish product reviews, users share opinions through social networking platforms, patients discuss healthcare experiences, and individuals express attitudes toward brands, organizations, products, and events. This information provides valuable knowledge for decision-making but is difficult to analyze manually because of its volume, velocity, and linguistic diversity.

Sentiment analysis, also referred to as opinion mining, is concerned with automatically determining the sentiment or opinion expressed in textual content. The output may be a simple polarity such as positive, negative, or neutral, or may involve more detailed emotional or aspect-level information.

Early research concentrated on traditional machine-learning algorithms such as Naïve Bayes, Maximum Entropy, and Support Vector Machines (SVMs). Pang et al. demonstrated the applicability of these algorithms to movie-review sentiment classification and showed that sentiment classification presents characteristics different from conventional topic classification [1].

analysis. Pang and Lee provided an influential survey of opinion mining and sentiment analysis, establishing many of the fundamental concepts used in later research [2].

The emergence of distributed word representations and deep learning significantly changed the field. Word2Vec provided dense vector representations of words [3], while Paragraph Vector extended distributed representations to documents [4]. Neural models subsequently learned increasingly complex semantic representations directly from text.

The development of BERT in 2018–2019 represented another important transition. Devlin et al. introduced BERT as a bidirectional transformer-based language representation model that could be pretrained on large amounts of unlabeled text and subsequently fine-tuned for downstream NLP tasks [5]. BERT-based approaches were also investigated for aspect-based sentiment analysis, including out-of-domain ABSA [6].

II. FUNDAMENTALS OF SENTIMENT ANALYSIS

Sentiment analysis can be considered at different levels of granularity.

A. Document-Level Sentiment Analysis

Document-level sentiment analysis assigns a single sentiment label to an entire document.

For example:

“The laptop is excellent and provides outstanding performance.”

can be classified as **Positive**.

A basic formulation is:

S(D)∈{Positive,Negative,Neutral}S(D) \in \{Positive, Negative, Neutral\}

where DD represents the document and S(D)S(D) represents its sentiment.

Document-level classification is appropriate when a document expresses an overall opinion. However, it becomes less suitable when multiple opinions occur in the same document.

B. Sentence-Level Sentiment Analysis

Sentence-level sentiment analysis determines the polarity of individual sentences.

For example:

“The camera quality is excellent.”

is positive, whereas:

“The battery performance is disappointing.”

is negative.

Sentence-level analysis provides finer granularity than document-level analysis but still may not identify the particular entity or aspect toward which the sentiment is directed.

C. Aspect-Level Sentiment Analysis

Aspect-Based Sentiment Analysis (ABSA) identifies both the aspect and the associated sentiment.

For example:

“The phone has an excellent camera but poor battery life.”

can be represented as:

Aspect	Sentiment
Camera	Positive
Battery life	Negative

ABSA has become an important research direction because users frequently express different

sentiments toward different characteristics of the same entity. Surveys published by 2019–2020 identify aspect extraction and aspect sentiment classification as major components of ABSA research. (ScienceDirect)

III. EVOLUTION OF SENTIMENT ANALYSIS

The evolution of sentiment analysis can broadly be divided into the following stages:

1. Lexicon-based approaches
2. Traditional machine-learning approaches
3. Distributed word representations
4. Deep neural networks
5. Attention-based models
6. Transformer-based models

TABLE I

EVOLUTION OF SENTIMENT ANALYSIS

Period	Major Approach	Representative Techniques
Early stage	Lexicon-based	Sentiment dictionaries, polarity scoring
2000–2010	Machine Learning	Naïve Bayes, SVM, Maximum Entropy
2010–2014	Feature-based NLP	TF-IDF, N-grams, linguistic features
2013–2016	Distributed representations	Word2Vec, Paragraph Vector
2014–2017	Deep learning	CNN, RNN, LSTM
2016–2018	Attention/memory models	Attention networks, memory networks
2018–2020	Transformers	BERT and BERT-based sentiment models

IV. LEXICON-BASED SENTIMENT ANALYSIS

Lexicon-based sentiment analysis uses predefined collections of words associated with sentiment polarity.

For example:

Positive words: excellent, wonderful, good, outstanding

Negative words: poor, terrible, disappointing, bad

A simple sentiment score can be represented as:

Score(T)=∑i=1nPolarity(wi)Score(T)=\sum_{i=1}^{n} Polarity(w_i)

where wiw_i represents a sentiment-bearing word.

Taboada et al. developed the Semantic Orientation CALculator (SO-CAL), which incorporates lexical polarity, negation, and intensification [7].

Advantages

- Does not require large labeled datasets.
- Relatively simple to implement.
- Computationally inexpensive.
- Easy to interpret.

Limitations

- Domain dependence.
- Difficulty handling sarcasm.
- Difficulty understanding context.
- Limited handling of new expressions.
- Poor performance on multilingual and code-mixed text.

V. TRADITIONAL MACHINE-LEARNING APPROACHES

Traditional machine-learning systems generally follow:

Text → Preprocessing → Feature Extraction → Classification → Sentiment

A. Naïve Bayes

Naïve Bayes is a probabilistic classifier frequently used for text classification.

P(C|D)=P(D|C)P(C)P(D)P(C|D)=\frac{P(D|C)P(C)}{P(D)}

where CC is a sentiment class and DD represents a document.

The algorithm is computationally efficient and performs well in high-dimensional text-classification problems.

B. Support Vector Machine

SVM attempts to identify an optimal hyperplane separating different sentiment classes.

f(x)=wTx+bf(x)=w^Tx+b

SVM became one of the most widely investigated approaches for sentiment classification because text representations are generally sparse and high-dimensional.

C. Maximum Entropy

Maximum Entropy models estimate the probability of a class based on multiple features:

P(c|x)=ewcTx∑c′ewc′TxP(c|x)=\frac{e^{w_c^Tx}}{\sum_{c'}e^{w_{c'}^Tx}}

Pang et al. evaluated Maximum Entropy along with Naïve Bayes and SVM in their influential sentiment-classification study [1].

VI. TEXT REPRESENTATION AND FEATURE ENGINEERING

Traditional sentiment-analysis methods rely heavily on text representation.

A. Bag-of-Words

The Bag-of-Words model represents text using word frequencies while ignoring word order.

B. TF-IDF

TF-IDF assigns importance to a term according to its frequency within a document and its rarity across the corpus.

TFIDF(t,d)=TF(t,d)×IDF(t)TFIDF(t,d)=TF(t,d)\times IDF(t)

where:

IDF(t)=log Ndf(t)IDF(t)=\log\frac{N}{df(t)}

TF-IDF was frequently combined with SVM, Naïve Bayes, and logistic regression.

C. N-Grams

N-grams represent sequences of words.

Examples:

- Unigram: excellent
- Bigram: very good
- Trigram: not very good

N-grams can capture local sentiment expressions but significantly increase the dimensionality of the feature space.

VII. WORD EMBEDDINGS

A major limitation of Bag-of-Words and TF-IDF is their inability to capture semantic relationships effectively.

Mikolov et al. introduced Word2Vec, which produces distributed vector representations of words [3].

For example, semantically related words such as:

excellent – wonderful – outstanding

can occupy nearby regions of the embedding space.

Word embeddings can therefore provide richer representations than sparse word-frequency features.

Le and Mikolov extended distributed representations to sentences and documents through Paragraph Vector [4].

VIII. DEEP LEARNING FOR SENTIMENT ANALYSIS

Deep learning reduced dependence on manually engineered features by allowing models to learn representations directly from data.

A. Convolutional Neural Networks

Kim demonstrated the effectiveness of CNNs for sentence classification [8].

A typical architecture is:

Word Embedding → Convolution → Pooling → Fully Connected Layer → Classification

CNNs can identify local sentiment patterns such as:

- "very good"
- "extremely poor"
- "not recommended"

B. Recurrent Neural Networks

RNNs model sequential dependencies between words.

For a sequence:

x1,x2,...,xnx_1,x_2,...,x_n

the hidden representation at time tt can be expressed as:

ht=f(xt,ht−1)h_t=f(x_t,h_{t-1})

This enables the model to incorporate preceding context.

C. LSTM

Long Short-Term Memory networks were introduced to address difficulties associated with learning long-term dependencies.

LSTM models use memory cells and gating mechanisms to control information flow.

LSTM became particularly important for sentiment analysis because sentiment can depend on words appearing earlier in a sentence.

IX. ATTENTION MECHANISMS

Attention mechanisms allow a model to assign different importance to different words.

For example:

"The battery is extremely disappointing but the camera is excellent."

For battery sentiment, the model should focus more heavily on **battery** and **disappointing**.

For camera sentiment, it should emphasize **camera** and **excellent**.

Attention therefore provides a mechanism for selectively focusing on relevant contextual information.

Attention-based models became particularly important for ABSA because the sentiment associated with an aspect may depend on distant words in a sentence.

X. ASPECT-BASED SENTIMENT ANALYSIS

ABSA is a fine-grained form of sentiment analysis.

It generally includes:

1. Aspect extraction
2. Aspect-category detection
3. Opinion extraction
4. Aspect-sentiment classification

Do et al. reviewed more than 40 deep-learning models for ABSA and highlighted the importance of aspect extraction and target-dependent sentiment classification. (ScienceDirect)

Liu et al.'s 2020 survey further categorized ABSA methods and reviewed benchmark datasets, evaluation metrics, deep-learning architectures, challenges, and future research directions. (Digital Commons)

TABLE II

MAJOR ABSA TASKS

Task	Objective	Example
Aspect Extraction	Identify aspects	"camera", "battery"
Opinion Extraction	Identify opinion expressions	"excellent", "poor"
Aspect Polarity	Determine sentiment	camera → positive
Aspect Category	Assign category	camera → hardware
Target-Dependent Sentiment	Determine sentiment toward target	Target → polarity

XI. BERT AND TRANSFORMER-BASED SENTIMENT ANALYSIS

The introduction of Transformer architectures significantly changed NLP.

Vaswani et al. introduced the Transformer architecture based primarily on attention mechanisms [9]. The architecture eliminated the need for recurrence in the core sequence-processing mechanism.

Devlin et al. subsequently introduced BERT, which uses bidirectional transformer representations pretrained on large text corpora [5].

BERT can be fine-tuned for sentiment classification by adding a classification layer above the contextual representation.

A simplified architecture is:

Input Text → Tokenization → BERT → [CLS] Representation → Classification Layer → Sentiment

Unlike static embeddings such as Word2Vec, BERT generates contextual representations. Consequently, the representation of a word can vary according to its surrounding context.

Hoang et al. investigated BERT for aspect-based sentiment analysis and reported its potential for out-of-domain ABSA tasks [6]. (ACL Anthology)

XII. SOCIAL-MEDIA SENTIMENT ANALYSIS

Social media provides an important source of sentiment data.

However, social-media text differs considerably from formal writing.

Typical characteristics include:

- Abbreviations
- Hashtags
- Emojis
- Emoticons
- Misspellings
- URLs
- User mentions
- Slang
- Repeated characters
- Code-mixing

Go et al. demonstrated the use of distant supervision for Twitter sentiment classification by exploiting emoticons as noisy sentiment indicators [10].

Hutto and Gilbert introduced VADER, a rule-based model specifically designed for sentiment analysis of social-media text [11].

XIII. SARCASM AND SENTIMENT ANALYSIS

Sarcasm represents a major challenge because literal word polarity may contradict the intended sentiment.

For example:

"Great, my flight has been delayed for another five hours."

The word **great** is lexically positive, but the intended sentiment is negative.

Bouazizi and Ohtsuki investigated pattern-based sarcasm detection on Twitter, demonstrating the importance of identifying sarcastic content separately from conventional sentiment classification [12].

XIV. MULTILINGUAL AND CODE-MIXED

Most early sentiment-analysis research focused heavily on English. However, Internet users increasingly communicate in multiple languages.

Code-mixed text can combine languages within the same sentence.

For example:

"Movie khup awesome aahe."

A system must identify:

- English words
- Marathi words
- Sentence context
- Sentiment-bearing expressions

Multilingual and code-mixed sentiment analysis remains challenging because of limited annotated datasets, language-specific morphology, vocabulary variation, transliteration, and lack of high-quality sentiment resources.

Akhtar et al. investigated cross-lingual and multilingual methods for addressing data sparsity in ABSA [13].

XV. DATASETS USED IN SENTIMENT ANALYSIS

TABLE III

REPRESENTATIVE SENTIMENT DATASETS

Dataset	Domain	Main Task
Movie Review Dataset	Movies	Binary sentiment
Stanford Sentiment Treebank	Movies	Fine-grained sentiment
Twitter Dataset	Social media	Tweet sentiment
SemEval-2014	Restaurants/Laptops	ABSA
SemEval-2015	Reviews/Tweets	ABSA
SemEval-2016	Multiple domains	ABSA
Amazon Reviews	E-commerce	Product sentiment

Benchmark datasets are essential for comparing algorithms under standardized conditions. The SemEval ABSA tasks were particularly influential in evaluating aspect extraction and sentiment classification.

XVI. EVALUATION METRICS

The most common metrics include accuracy, precision, recall, and F1-score.

Accuracy

Accuracy=TP+TNTP+TN+FP+FNAccuracy=\frac{TP+TN}{TP+TN+FP+FN}

Precision

Precision=TPTP+FPPrecision=\frac{TP}{TP+FP}

Recall

Recall=TPTP+FNRecall=\frac{TP}{TP+FN}

F1-Score

F1=2Precision×RecallPrecision+RecallF1=2\frac{ Precision\times Recall} {Precision+Recall}

For imbalanced datasets, macro-F1 is often more informative than accuracy because it gives equal importance to individual classes.

XVII. COMPARISON OF SENTIMENT ANALYSIS APPROACHES

TABLE IV

COMPARISON OF MAJOR APPROACHES

Approach	Feature Engineering	Context Modeling	Training Data	Interpretability	Main Limitation
Lexicon-based	Low	Low	Not required	High	Domain dependence
Naïve Bayes	High	Low	Required	Medium	Independence assumption
SVM	High	Low–Medium	Required	Medium	Manual features
Word2Vec + ML	Medium	Medium	Required	Low	Static embeddings
CNN	Low	Medium	Required	Low	Limited long-range context
LSTM	Low	High	Required	Low	Training complexity
Attention-LSTM	Low	High	Required	Medium	Computational cost
BERT	Very Low	Very High	Pretraining + fine-tuning	Low–Medium	Resource requirements
ABSA	Variable	High	Usually required	Medium	Aspect identification

XVIII. CHALLENGES

Despite significant progress, sentiment analysis continues to face several challenges.

A. Negation

“Good” and “not good” have different polarity.

B. Sarcasm

Literal word polarity may contradict intended sentiment.

C. Implicit Sentiment

The sentiment may not be directly expressed.

For example:

“The phone stopped working after two days.”

This implies a negative experience without explicitly using a negative sentiment word.

D. Domain Adaptation

A sentiment model trained on restaurant reviews may not perform equally well on financial or healthcare text.

E. Multilingual Sentiment

Different languages require different linguistic resources and modeling strategies.

F. Code-Mixed Language

Mixed-language communication creates vocabulary and grammatical challenges.

G. Explainability

Deep neural and transformer models can produce accurate predictions while providing limited explanations for individual decisions.

H. Computational Requirements

Transformer-based models generally require substantially more computational resources than lexicon-based and traditional machine-learning approaches.

XIX. RESEARCH GAPS

The literature up to 2020 indicates several research gaps:

1. Limited resources for low-resource languages.
2. Insufficient datasets for code-mixed sentiment analysis.
3. Difficulty identifying implicit aspects.
4. Difficulty interpreting sarcasm and irony.
5. Domain-transfer problems.
6. Limited explainability of deep-learning models.
7. Difficulty integrating linguistic knowledge with neural representations.
8. High computational requirements of transformer models.
9. Challenges in real-time sentiment analysis.
10. Limited unified approaches combining aspect, sentiment, emotion, and contextual information.

The 2020 ABSA survey specifically identifies problems involving datasets, evaluation, deep-learning architectures, and future research directions. (Digital Commons)

XX. FUTURE RESEARCH DIRECTIONS

Future sentiment-analysis research can focus on the following areas.

1. Multilingual Models

Development of models capable of handling multiple languages within a unified architecture.

2. Code-Mixed Sentiment Analysis

Specialized models for multilingual social-media communication.

3. Explainable Sentiment Analysis

Integration of explainability methods to identify which words, aspects, or contextual features contributed to a prediction.

4. Aspect-Aware Transformers

Transformer architectures specifically designed to model relationships between aspects and sentiment expressions.

5. Multimodal Sentiment Analysis

Integration of:

Text + Image + Audio + Video

for sentiment and emotion recognition.

6. Low-Resource Sentiment Analysis

Transfer learning and multilingual representation learning can reduce dependence on large manually annotated datasets.

7. Real-Time Analytics

Sentiment models can be integrated with streaming platforms to analyze continuously generated social-media and customer-feedback data.

CONCLUSION

Sentiment analysis has evolved substantially from lexicon-based polarity scoring and conventional machine-learning algorithms to deep-learning and transformer-based architectures. Early studies demonstrated the effectiveness of Naïve Bayes, Maximum Entropy, and SVM for sentiment classification. Lexicon-based methods provided

interpretable alternatives that do not require large labeled datasets.

The introduction of distributed word representations improved the ability of models to capture semantic relationships. CNN, RNN, and LSTM architectures subsequently enabled automated learning of increasingly sophisticated textual representations. Attention mechanisms provided additional capability for focusing on sentiment-relevant contextual information.

Aspect-Based Sentiment Analysis extended conventional sentiment classification by identifying sentiment toward individual aspects. This development is particularly important for product reviews and customer feedback, where a single document may contain both positive and negative opinions about different characteristics.

The emergence of Transformer architectures and BERT around 2018–2019 represented another major transition. BERT introduced contextual bidirectional representations that could be fine-tuned for downstream sentiment-analysis tasks, including aspect-based sentiment analysis. Research published by 2020 had therefore established a progression from manually engineered features toward contextual pretrained representations. (ACL Anthology)

Nevertheless, important challenges remain in sarcasm detection, implicit sentiment, domain adaptation, multilingual processing, code-mixed text, explainability, and low-resource languages. These challenges provide substantial opportunities for future research in multilingual, explainable, aspect-aware, and resource-efficient sentiment-analysis systems.

REFERENCES

[1] B. Pang, L. Lee, and S. Vaithyanathan, "Thumbs up? Sentiment Classification Using Machine Learning Techniques," *Proceedings of EMNLP*, pp. 79–86, 2002.

[2] B. Pang and L. Lee, "Opinion Mining and Sentiment Analysis," *Foundations and Trends in Information Retrieval*, vol. 2, no. 1–2, pp. 1–135, 2008.

[3] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed Representations of Words and Phrases and Their Compositionality," *Advances in Neural Information Processing Systems*, pp. 3111–3119, 2013.

[4] Q. V. Le and T. Mikolov, "Distributed Representations of Sentences and Documents," *Proceedings of the 31st International Conference on Machine Learning*, pp. 1188–1196, 2014.

[5] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," *Proceedings of NAACL-HLT*, pp. 4171–4186, 2019.

[6] M. Hoang, O. A. Bihorac, and J. Rouces, "Aspect-Based Sentiment Analysis using BERT," *Proceedings of the 22nd Nordic Conference on Computational Linguistics*, pp. 187–196, 2019. (ACL Anthology)

[7] M. Taboada, J. Brooke, M. Tofiloski, K. Voll, and M. Stede, "Lexicon-Based Methods for Sentiment Analysis," *Computational Linguistics*, vol. 37, no. 2, pp. 267–307, 2011.

[8] Y. Kim, "Convolutional Neural Networks for Sentence Classification," *Proceedings of EMNLP*, pp. 1746–1751, 2014.

[9] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention Is All You Need," *Advances in Neural Information Processing Systems*, pp. 5998–6008, 2017.

[10] A. Go, R. Bhayani, and L. Huang, "Twitter Sentiment Classification Using Distant Supervision," *CS224N Project Report*, Stanford University, 2009.

[11] C. J. Hutto and E. Gilbert, "VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text," *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 8, no. 1, pp. 216–225, 2014.

[12] M. Bouazizi and T. Ohtsuki, "A Pattern-Based Approach for Sarcasm Detection on Twitter," *IEEE Access*, vol. 4, pp. 5477–5488, 2016.

[13] M. S. Akhtar, P. Sawant, S. Sen, A. Ekbal, and P. Bhattacharyya, "Solving Data Sparsity for Aspect Based Sentiment Analysis Using Cross-Linguality and Multi-Linguality," *Proceedings of NAACL-HLT*, pp. 572–582, 2018.

[14] H. Ha Do, P. W. C. Prasad, A. Maag, and A. Alsadoon, "Deep Learning for Aspect-Based Sentiment Analysis: A Comparative Review," *Expert Systems with Applications*, vol.

118, pp. 272–299, 2019. (ScienceDirect)

[15] H. Liu, I. Chatterjee, M. Zhou, X. S. Lu, and A. Abusorrah, "Aspect-Based Sentiment Analysis: A Survey of Deep Learning Methods," *IEEE Transactions on Computational Social Systems*, vol. 7, no. 6, pp. 1358–1375, 2020. (Digital Commons)

[16] J. Bao, P. Lambert, and T. Badia, "Attention and Lexicon Regularized LSTM for Aspect-Based Sentiment Analysis," *Proceedings of the ACL Student Research Workshop*, pp. 253–259, 2019.

[17] Q. Jiang, L. Chen, R. Xu, X. Ao, and M. Yang, "A Challenge Dataset and Effective Models for Aspect-Based Sentiment Analysis," *Proceedings of EMNLP-IJCNLP*, pp. 4579–4589, 2019.

[18] H. Huang, X. Zhang, Z. Zhao, J. Xu, and Z. Li, "Image–Text Sentiment Analysis via Deep Multimodal Attentive Fusion," *Knowledge-Based Systems*, vol. 167, pp. 26–37, 2019.

[19] N. Jakob and I. Gurevych, "Extracting Opinion Targets in a Single and Cross-Domain Setting with Conditional Random Fields," *Proceedings of EMNLP*, pp. 1035–1045, 2010.

[20] B. Liu, *Sentiment Analysis: Mining Opinions, Sentiments, and Emotions*. Cambridge, U.K.: Cambridge University Press, 2015.