



Identifying Time-varying Drivers of Social Media Issues and Conversations

Gurpreet Singh Bawa^{1*}, Suresh Kumar Sharma¹ and Kanchan Jain¹

¹Department of Statistics, Panjab University, Sector 14, Chandigarh, India.

Authors' contributions

This work was carried out in collaboration between all authors. All authors read and approved the final manuscript.

Article Information

DOI: 10.9734/JAMCS/2019/47118

Editor(s):

- (1) Dr. Oguntunde, Pelumi Emmanuel, Department of Mathematics, College of Science and Technology, Covenant University, Ota, Ogun State, Nigeria.
- (2) Dr. Thomas L. Toulas, Department of Biomedical Sciences, University of West Attica, Greece.
- (3) Dr. Kai-Long Hsiao, Associate Professor, Taiwan Shoufu University, Taiwan.

Reviewers:

- (1) Margarita Išoraitė, Vilniaus kolegija-The University Applied Sciences, Lithuania.
- (2) Husein Mohamed Cirbad, Annamalai University, India.

Complete Peer review History: <http://www.sdiarticle3.com/review-history/47118>

Received: 12 November 2018

Accepted: 31 January 2019

Published: 06 February 2019

Original Research Article

Abstract

Successfully understanding of social media conversation growth, dissemination and extinction is a challenging task that relies on identifying groups, group influence, diffusion models, forecast models, social dynamics and text analytics. In this problem, we concentrate on the description of a novel approach for identifying drivers of the direction and momentum of social conversations, including the spread of mood, sentiment and issues. The approach first groups potential drivers of conversation based on variability. The primary driver in each group is then selected. Finally, the relationship between the selected drivers and the topic outcome is calculated and displayed visually. This enables the quick identification of the form and structure of the conversation and allows us to predict momentum, direction, contagion risks, potential responses and interventions.

There is a huge amount of data in the form of text available today in the internet across various channels – social media, news articles, blogs, e-commerce websites. Most of this data is a part of some “conversation” or the other where real-world entities discuss, analyze, comment, exchange information in the form of written expressions in textual format. Driver Modeling on textual data can be useful in observing the key drivers which are driving the “conversation” coupled with the associated sentiments and mood states for the observed key drivers. These insights about the key conversational drivers are often used in a variety of domains such as tracking news cycles, stock movements, legislation developments, brand image, viral breakouts and much more.

*Corresponding author: E-mail: gurpreetsinghbawa@yahoo.co.in;

Keywords: Driver analysis; textual topic; latent Dirichlet allocation; sentiment analysis; NLP; sub-net identification; Hypernyms prediction.

1 Introduction

In most e-commerce/connected commerce websites, the consumer base has the authority to comment or provide feedback pertaining to the services received. E-commerce/retail websites have the historical feedback/reviews received open for public viewing. 90% of consumers say buying decisions are influenced by online reviews; More than 50% of online purchasers said they had read online reviews prior to hitting the buy button, while nearly 40% of consumers who made purchases in-store did so (*Report published by Google in 2012*). For connected commerce business, ratings & reviews play a critical role in driving brand/product choice and ultimately conversion to purchase [1]. There is an opportunity to mine the ratings & reviews data and convert it into a competitive advantage to drive conversion & loyalty, win vs. competitors, and ultimately drive incremental sales. There is an opportunity to mine the ratings & reviews data and convert it into a competitive advantage to drive conversion & loyalty, win vs. competitors, and ultimately drive incremental sales. The goal is to utilize technology (Natural Language Processing) in and machine learning to understand drivers of ratings that can then be leveraged to reach right consumers/right time/right place with right product at right price. Through NLP/ML, we would be able to better understand drivers of Ratings & Reviews, cluster consumers; deploy learnings across consumption generating vehicles and drive higher conversion/sales/loyalty. With over 60% of all outlet sales being digitally influenced, this would help drive both online and offline sales. Learnings will also influence innovation and communications strategy. The approach is illustrated in Fig. 1, where we show the process flow.

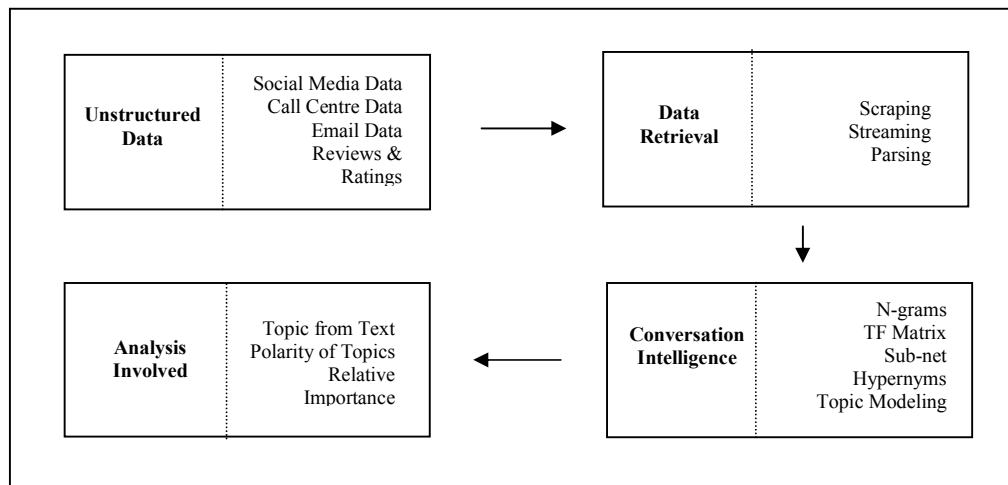


Fig. 1. Illustration of the high-level process flow

1.1 Preliminary pre-processing

To perform Driver Analysis, it is necessary to process the natural language and convert it into a model-ingestible format. The major preliminary processes involved are as follows (Table 1):

- Stemming
- Lemmatization
- POS Tagging
- Dictionary and Taxonomy construction
- Acronym resolution
- N-gram Tokenization

Table 1. Preliminary pre-processing of textual data to make it ingestible

Unstructured Data	Natural Language Text Pre-processing	Tokenization
Totally love the product	Stemming & Lemmatization – to bring the words to its base form	Tokenization for the input of the topic model can be taken as n-grams, where n can be set by user as required
Good for a gift set	Part of Speech Identification – to observe semantic context	
...		
A little expensive		

2 Key Components

The key functional unique components of the driver model are as follows:

- **Community or Sub-net identification** and their respective hypernyms using ‘Relatedness’ matrix (Explicit Semantic Indexing).
- **N-gram Tokenization for Topic Modeling** - For each hypernym, we would be identifying the topics using n-grams.
- **Weighing the Document-Term Matrix** with Inverse Document Frequency to provide higher weights to low-frequency terms and lower weights to high-frequency terms.
- **Topic Modeling** on each sub-net to identify top n-grams (hypernyms) [2,3], using Latent Dirichlet Allocation/Latent Semantic Indexing.
- **Polarity and importance** of n-grams in the topics.

2.1 Community or Sub-net identification

The prime advantage of **Explicit Semantic Indexing (ESI)** during sub-net creation is that it provides greater flexibility. It also does not rely on a specific taxonomy to ingest and categorize documents (Fig. 2).

ESI works at a semantic and contextual level rather than a superficial vocabulary of a word or document. It represents the meaning of a piece text, as a combination of the concepts found in the text. The context or meaning of a piece of text, as a composition of the concepts embedded, is best represented by ESI. It has extensive utilities in document classification, semantic relatedness and information retrieval [4,5].

In document classification, for example, documents are labeled to make them easier to organize. Labeling a document with keywords/concepts makes it easier to pinpoint. However, keyword labeling alone has its cons; searches carried out using vocabulary with similar semantic context, but different actual terms may not obtain relevant documents. However semantic classification of text reduces the effect on specific terms and can greatly improve the understanding of the text [6,7].




Fig. 2. Community or Sub-net identification process flow

2.1.1 Example of community or Sub-net

Wikipedia is an extensive source of information where each article can be assumed to be a distinct concept. In ESA based on Wikipedia, a concept is initiated for each article. A vector composed of the terms which prevail in the article then represents the concepts, weighted by their tf-idf (term frequency - inverse document frequency) measure.

The context of any given term can then be projected as a vector of the “association weighting” to the

previously established concepts. By comparing two word-vectors using association measures, we can identify their semantic relation to the underlying concepts (Table 2).

Table 2. Community or Sub-net identification example

Economy (Category)				
Market	Trade	Property	Word n
Bazaar (0.72)	Economy (0.61)	Real estate (0.81)	$\langle c_k w_k \rangle$
NASDAQ (0.65)	Trade Union (0.51)	Theft (0.75)	$\langle c_i w_i \rangle$
Economy (0.53)	Silk road (0.3)	Copyright (0.58)	$\langle c_i w_i \rangle$
...

2.1.2 Representation of Community or Sub-net

Larger documents are projected as a contrast of word vectors obtained from the document vocabulary. Resultant document vectors are called “concept” vectors [8]. For example, a concept vector might look something like the following (Table 3).

Table 3. Community or Sub-net identification representation

‘Concept’ Vectors				
“Mars”	$\langle \text{planet}, 0.90 \rangle$	$\langle \text{Solar system}, 0.85 \rangle$	$\langle \text{Jupiter}, 0.30 \rangle$
“Explorer”	$\langle \text{adventurer}, 0.89 \rangle$	$\langle \text{pioneer}, 0.70 \rangle$	$\langle \text{vehicle}, 0.20 \rangle$
::	::	::	::	::
“word n ”	$\langle \text{concept}, \text{weight } b \rangle$	$\langle \text{concept}, \text{weight } d \rangle$	$\langle \text{concept}, \text{weight } p \rangle$..

Geometrically, a concept vector can be represented as the centroid of the word vectors it is composed of. The image below illustrates the same. The blue dots in the image represent the ‘word’ vector and green dot represents ‘centroid’ of the words (Fig. 3).




Fig. 3. Community or Sub-net identification representation, Distance of nodes from its centroid, Image taken from R software

Hence, to compare the similarity between two phrases, we can generate their individual concept vectors from their underlying word vectors and then compare these two using association or similarity measures.

2.2 N-gram tokenization

In NLP and computational linguistics, an n-gram is an adjoining progression of n items from a fragment of text or speech in question. These items can be either of these – words, letters, syllables, phonemes or any base object according to the requirement (Table 4).

Sentence: “A swimmer likes swimming thus he swims.”

Table 4. Tokenization example

N-gram	Tokenization
Unigram (1-gram)	Swimmer, likes, swimming, thus, he, swims, ...
Bigram (2-gram)	A swimmer, swimmer likes, likes swimming, swimming thus, ...
Trigram (3-gram)	A swimmer likes, swimmer likes swimming, likes swimming thus, ...

Most topic model algorithms, like Latent Dirichlet Allocation, are dependent on bag-of-words approach. However, ordering and placement of phrase and words are often useful in understanding the context (Table 5). Topical n-grams, a topic model that establishes topics as well as topical phrases instead of just bag of words, would be able to capture more context – like the ability to differentiate “white house” in real estate or political context [9].

Table 5. Unigram and bigram tokenization

Text Data	N-gram	Tokenization
The packaging was not good	Unigram	The, packaging, was, not, good
	Bigram	The packaging, Packaging was, was not, not good
The product is not expensive	Unigram	The, product, is, not, expensive
	Bigram	The product, product is, is not, not expensive

2.2.1 Handling redundancy and noise

The approach of using n-gram tokenization also has some drawbacks. The huge number of n-grams generated from the raw texts often has few occurrences of n-grams that are extremely rare or occur a lesser number of times. This number is often very high and leads to unnecessary complications and clutter during further model building. Hence a method to select only those n-grams that are semantically sound or have some context helps in solving the above problem.

Detecting tokens is an important problem in the field of computational linguistics, where they’re usually called collocations. Most approaches carry out statistical tests that attempt to measure the variance with respect to a random bag-of-words model. **Ted Dunning’s G^2** is one such likelihood ratio test which compares the estimated probability of term B occurring after term A to the marginal distribution of term A and B (Table 6). If the occurrence of A is independent of B, then the probabilities will be almost same. If B is much more frequent after A, then the probabilities will be significantly different.

Table 6. Ted Dunning’s G^2

	Event A	Everything but not A
Event B	A and B together (k_{11})	B, but not A (k_{12})
Everything but not B	A without B (k_{21})	Neither A nor B (k_{22})

Computing the log-likelihood ratio score (also known as G^2) is defined as

$$LLR = 2 \sum(k) (H(k) - H(row\ sum(k)) - H(column\ sum(k))).$$

2.3 Weighting the document term matrix

Term Frequency (TF) weights are proportional to the occurrence of terms in the document whereas Inverse document frequency (IDF) is an inverse function of the number of documents containing a term thereby quantifying the specificity.

If we are planning to use LDA as the method to observe Topic Models, it must be kept in mind that the

weighting cannot be TF-IDF because the basic assumption for Dirichlet distribution demands the support of the Random Variable to be a fractional value and sum of all values should be 1. Hence the only solution is to use Term Frequency approach. But using TF approach can lead to loss of information of rarely used words in the documents, which even though small, are significant drivers of conversation. Hence, we plan to use TFIDF at a first level screening to screen words and drop ones that are highly frequent and less specific. Next, with the filtered list of terms, we propose to use its Document-Term Matrix with its Normalized Term Frequency weighting as the input for the Topic Model.

2.4 Topic modeling

Topic Modeling is an effective method of modeling textual corpora and other sets of discrete data in text format.

The primary aim is to find portrayals of the components of a group that allows dynamic processing of larger sets while sustaining the necessary statistical relationships that can be utilized further for tasks such as text classification, document summarization, relevance judgment prediction etc.

In ML and NLP, a topic model is a type of statistical model which can identify the hypothetical underlying "topics" that are present in a group of documents thus observing the latent semantic structures in text. Intuitively, if a given document is related to a topic, one can expect a set of words from the vocabulary to exist in the given document more frequently.

2.4.1 Latent semantic indexing

While the tf-idf measure provides some useful features such as identification of sets of terms/tokens that are discerning for documents, the approach does not provide a reasonable reduction in length of its description. It also provides very little information pertaining to intra or inter-document statistical structure [10].

To address these drawbacks, several dimensionality reduction techniques were proposed, one important such approach being Latent Semantic Indexing (LSI). A singular value decomposition of the document term matrix results in a linear vector space in the tf-idf feature space. This helps to explain a significant amount of variance in the collection of documents [11].

This approach can be used to obtain a significant reduction in size within large groups. Further, the derived features, which are linear contrasts of the original features, can explain some of the characteristics of synonymy and polysemy.

On selecting the top k singular values, and their analogous singular vectors from U and V , we get the approximation to the document term matrix with rank k with a minimal error. Most importantly, now the term and document vectors can be treated as a "semantic space".

We write this approximation as:

Let (i, j) be an element of matrix X denote the frequency of occurrence of i^{th} term (t_i) in j^{th} document (d_j). Matrix X looks like:

$$\begin{array}{c}
 d_j \\
 \downarrow \\
 \begin{array}{c}
 \begin{array}{c}
 x_{1,1} \quad \cdots \quad x_{1,j} \quad \cdots \quad x_{1,n} \\
 \vdots \quad \ddots \quad \vdots \quad \ddots \quad \vdots \\
 x_{i,1} \quad \cdots \quad x_{i,j} \quad \cdots \quad x_{i,n} \\
 \vdots \quad \ddots \quad \vdots \quad \ddots \quad \vdots \\
 x_{m,1} \quad \cdots \quad x_{m,j} \quad \cdots \quad x_{m,n}
 \end{array}
 \end{array}
 \end{array}
 \begin{array}{l}
 \\
 \\
 \xrightarrow{t_i^{Transpose}}
 \end{array}
 \end{array}$$

Now, as per linear algebra concepts, matrix X can be factorized such that there exists U and V which are

orthogonal and a matrix Σ which is diagonal. This is known as the factorization of singular value decomposition:

$$X = U\Sigma V^T$$

$$\begin{array}{ccccc}
 & X & & U & \Sigma & & V^T \\
 & (d_j) & & & & & (\widehat{d}_j) \\
 & \downarrow & & \downarrow & \downarrow & & \downarrow \\
 (t_i^T) \rightarrow & \begin{bmatrix} x_{1,1} & \cdots & x_{1,j} & \cdots & x_{1,n} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ x_{i,1} & \cdots & x_{i,j} & \cdots & x_{i,n} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ x_{m,1} & \cdots & x_{m,j} & \cdots & x_{m,n} \end{bmatrix} & = & (\widehat{t}_i^T) \rightarrow & \begin{bmatrix} \left[\begin{array}{c} u_1 \end{array} \right] & \cdots & \left[\begin{array}{c} u_l \end{array} \right] \end{bmatrix} & \times & \begin{pmatrix} \sigma_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \sigma_l \end{pmatrix} & \times & \begin{bmatrix} \left[\begin{array}{c} v_1 \end{array} \right] \\ \vdots \\ \left[\begin{array}{c} v_l \end{array} \right] \end{bmatrix}
 \end{array}$$

Now if we consider the top k singular values which has highest magnitude and their corresponding singular vectors from U and V , we get an approximation to X .

We can now use the decomposition to:

- Observe how related two documents are by analyzing the vectors by similarity measures
- Compare two terms within a document
- Vector representations of documents and terms can be clustered using different similarity measures.
- Given a query, using it as a reference, and compare it to existing documents in a space of lower dimension.

2.4.2 Advantages of latent Dirichlet allocation

Most dimensionality reduction approaches are governed by the fundamental “bag-of-words” assumption where the ordering of words and their occurrence in the text can be neglected [12].

A Dirichlet random variable of k dimensions, say θ , can assume values in the $(k-1)$ -simplex (a k -dimensional θ lies within the $(k-1)$ -simplex if $\theta_i \geq 0$, $\sum_{i=1}^k \theta_i = 1$), and the associated probability density is of

the form:

$$p(\theta | \alpha) = \frac{\Gamma\left(\sum_{i=1}^k \alpha_i\right)}{\prod_{i=1}^k \Gamma(\alpha_i)} \theta_1^{\alpha_1-1} \dots \theta_k^{\alpha_k-1}$$

where the parameter α is a k -dimensional vector with components $\alpha_i > 0$, and $\Gamma(x)$ represents the Gamma function. The rationale behind selecting Dirichlet is because it is a suitable distribution on the simplex. It belongs to the exponential family with enough statistics of finite dimensions and is conjugate to the multinomial distribution. These key properties help in the statistical inferential analysis and parametric estimation procedures for Latent Dirichlet Allocation. Given the parameters α and β , a set z of N topics and a set w of N words, topic mixture θ has a joint distribution which can be written as

$$p(\theta, z, w | \alpha, \beta) = p(\theta | \alpha) \prod_{n=1}^N p(z_n | \theta) p(w_n | z_n, \beta)$$

where $p(z_n | \theta)$ is θ_i for unique i such that $z_n^i = 1$. Integrating over the range of θ and summing over values of z , the marginal distribution of a document can be represented as:

$$p(w | \alpha, \beta) = \int p(\theta | \alpha) \left(\prod_{n=1}^N \sum_{z_n} p(z_n | \theta) p(w_n | z_n, \beta) \right) d\theta$$

Eventually, probability for the corpus is attained by computing the product of the marginal probabilities of single documents:

$$p(D | \alpha, \beta) = \prod_{d=1}^M \int p(\theta_d | \alpha) \left(\prod_{n=1}^{N_d} \sum_{z_{dn}} p(z_{dn} | \theta_d) p(w_{dn} | z_{dn}, \beta) \right) d\theta_d$$

In the jargon of probability theory, this spells out into the premise of exchangeability for the terms in a document. Although less often formally stated, the assumption of these methods is that the documents are exchangeable, that is, specific ordering of documents in the corpus can be ignored [13].

De Finetti's classic representation theorem establishes that any set of random variables which are exchangeable should have a depiction as a mixture distribution - in general – an infinite mixture. Hence, if we consider exchangeable depictions for terms and documents, it is justified to assume the mixture models that would capture this exchangeability of both terms and documents. This line of thought led to the Latent Dirichlet Allocation (LDA) model (Fig. 4).

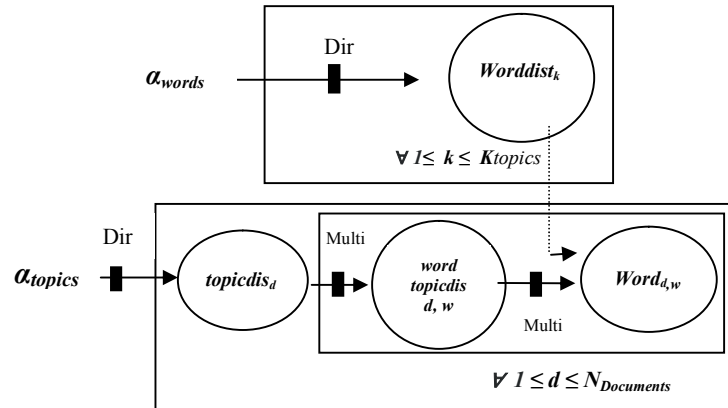


Fig. 4. Latent Dirichlet Allocation process flow

The purpose of using LDA/LSI and performing a Topic Model on the set of documents is to have a set of tokens from the top topics which would act as a base for our next step – the driver modeling. Topic modeling would provide us a bag of tokens as the top conversational topics and in the next step, we would be statistically observing the top tokens that are significantly driving the conversation and how it is impacting the ratings.

2.5 Driver analysis

In general, a key driver analysis is the study of the relations among many factors to identify the primary and significant ones. In our case, we are interested in the relationship between the text tokens which form the

documents (here reviews) and the associated ratings that the user provided along with the review/feedback. Such an analysis results, for example, explores and estimates the relationship between the top conversational topics from the review texts and the categorical rating variable [14,15].

Now to model the tokens from reviews and the categorical variable of ratings, we can use a multinomial logit model where the target/independent variable is ‘Rating’. We can consider either of the following two cases for defining the dependent variable measures:

- TF-IDF values of the tokens with respect to the reviews
- Similarity measure (Jaccard, Levenshtein and Cosine) between the tokens and reviews

The rationale behind using the Multinomial Logit Model is because of the presence of more than two categories of the independent variable. One more advantage behind this being it would make easier to analyze the drivers between reviews having a very low difference in rating (for example – why a set of reviews are rated 5 and another set being rated 4).

Table 7. Snapshot of how drivers and associated importance looks after sentiment integration, Screenshot of results from worksheet

Negative (Neg)/Positive (Pos)	Bigram	Weight	Abs Weight	% Importance
Pos	helps calm nap	-0.005103523	0.005103523	18.4%
Neg	expensive products	-0.005103523	0.005103523	18.4%
Pos	relieve migraines	-0.003645644	0.003645644	13.2%
Pos	overpowering silky	-0.001973021	0.001973021	7.1%
Neg	smell overpowering	-0.001973021	0.001973021	7.1%
Pos	calming smell	-0.001973021	0.001973021	7.1%
Neg	sensitive ingredients	-0.000967017	0.000967017	3.5%
Neg	child rash	-0.000967017	0.000967017	3.5%
Pos	soothing powering	-0.000847478	0.000847478	3.1%
Pos	congestion soothing	-0.000847478	0.000847478	3.1%
Pos	tear free	-0.000806603	0.000806603	2.9%
Pos	sinus headaches	-0.000608513	0.000608513	2.2%
Pos	smell good	-0.000498742	0.000498742	1.8%
Pos	hair smell	-0.000362972	0.000362972	1.3%
Pos	soothing effects	-0.000362972	0.000362972	1.3%
Neg	bad seller	-0.000314663	0.000314663	1.1%
Neg	soap expensive	-0.000314663	0.000314663	1.1%
Neg	caused irritation	-0.00022629	0.00022629	0.8%
Pos	feeling great	-0.00022629	0.00022629	0.8%
Pos	minor congestion	-0.00022629	0.00022629	0.8%
Pos	smell amazing	-0.00008310	0.00008310	0.3%
Pos	menthol doesn't	-0.00007230	0.00007230	0.3%
Neg	disappointed bought	-0.00007230	0.00007230	0.3%
Neg	dont recommend	-0.00007230	0.00007230	0.3%
Pos	great smell	-0.00003060	0.00003060	0.1%

The rationale behind using TF-IDF values of tokens as dependent variables is that it represents the frequency as well as specificity of the tokens for each document. However, one drawback might be the extreme sparseness of the matrix since the chances of a token occurring in a document (review) decreases when an n-gram of higher order is used. As an alternative to the TF-IDF approach, we can use the Similarity approach. This method takes into consideration how a token is semantically like each of the review. The setback this method might face is – if the tokens are created without stop words being removed from the review text and then compared to the entire document, there could be noise and misleading measures of similarity caused by the presence of common words in the reviews. To avoid this over or underestimation of the similarity scores, we can compare the tokens with the pre-processed review texts which contain no stop words.

Once we have the logit model, we can observe the standardized beta estimates of the dependent variables or the tokens to have an idea of their contribution in explaining the target or independent variable – ratings.

2.6 Sentiment and mood-state integration

This component is a further refinement on the drivers observed from the above step. Here we try to classify the drivers into either being a positive or a negative sentiment driver [16]. This would bring deep insights into why a token is driving the conversation in a positive/negative direction on the rating scale (Table 7). Also, if we layer in the notion of mood states into our analysis, then we can have a further idea of how and where a driver is affecting the mood state of consumers. The sentiment or mood state integration can be easily implemented using lexicon-based word-emotion association.

3 Conclusion

This research proposed an analytical methodology which enables to understand the drivers of textual data that can be leveraged to drive competitive advantage. The proposed methodology involves various components namely Sub-net Identification, tokenization based on N-grams, calculating the document term matrix using inverse document frequency, topic modeling on each sub-net to identify top n-grams and finally polarity & importance of n-grams in the topics – all these components to be performed in a sequential manner resulting in identifying the key drivers along with their polarity (positive or negative) from textual data.

By using the above components in the paper, we have shown the drivers of reviews which impacts the ratings in e-commerce/connected commerce platforms, which intern helped the business to achieve more traffic on their websites.

As a generalized use case, if there arises a scenario where we have a textual data under study and want to understand the drivers of same, we can implement this approach effectively to attain variable importance of hypernyms (which are drivers of textual data). If there is no target variable per se, we can use the overall sentiment polarity scores to define the target variable of text in question and perform the analysis using the proposed methodology in a similar fashion.

Competing Interests

Authors have declared that no competing interests exist.

References

- [1] Huang X, Huang YR, Wen M, An A, Liu Y, Poon J. Applying data mining to pseudo-relevance feedback for high performance text retrieval. In Proceedings of the Sixth IEEE International Conference on Data Mining (ICDM'06). IEEE Computer Society, Hong Kong. 2006;295–306.
- [2] Hearst M. Automatic acquisition of hyponyms from large text corpora. In Proceedings of 14th International Conference on Computational Linguistics; 1992.
- [3] Snow R, Daniel J, Andrew YN. Learning syntactic patterns for automatic hypernym discovery. In Advances in Neural Information Processing Systems; 2004.
- [4] Gabrilovich E, Markovitch S. Feature generation for text categorization using world knowledge. In Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI-05). Morgan Kaufmann Publishers Inc., Edinburgh, Scotland. 2005;1048–1053.

- [5] Gupta R, Ratnov LA. Text categorization with knowledge transfer from heterogeneous data sources. In Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence. AAAI Press, Chicago, IL. 2008;842–847.
- [6] Chang MW, Ratnov L, Roth D, Srikumar V. Importance of semantic representation: Dataless classification. In Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence. AAAI Press, Chicago, IL. 2008;830–835.
- [7] Manfred S. The hyperonym problem revisited: Conceptual and lexical hierarchies in language generation. In Association for Computational Linguistics. 2000;93-99.
- [8] Deville B, Bawa GS. United States Patent 9,317,594. Social community identification for automatic document classification; 2016.
- [9] Yan X, Guo J, Lan Y, Cheng X. A biterm topic model for short texts. Proceedings of the 22nd International Conference on World Wide Web. ACM. 2013;1445– 1456.
- [10] Dumais ST. Latent semantic indexing (lsi) and trec-2. In Proceedings of the Second Text REtrieval Conference (TREC-2). NIST, Gaithersburg, MD. 1994;105–116. Egozi O, Gabrilovich E, Markovitch S. Concept-based feature generation and selection for information retrieval. In Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence. AAAI Press, Chicago, IL. 2008;1132–1137.
- [11] Ville BD. A tale of two SAS technologies, generating maps of topical coverage and linkages. In SAS User Conference Papers Denise Bedford, Kent State University, Cary, NC. 2013;102-112.
- [12] Blei DM, Ng AY, Jordan MI. Latent Dirichlet allocation. Journal of Machine Learning Research. 2003;3:993–1022.
- [13] Zhao WX, Jiang J, Weng J, He J, Lim EP, Yan H, Li X. Comparing twitter and traditional media using topic models. Advances in Information Retrieval: 33rd European Conference on IR Research, ECIR. 2011;338–349.
- [14] James AC, Zheng Z. United States Patent 8832015 B2 Fast binary rule extraction for large scale text data; 2014.
- [15] John GH, Kohavi R, Pflieger K. Irrelevant features and the subset selection problem. In Proceedings of the 11th International Conference on Machine Learning. New Brunswick, NJ. 1994;121–129.
- [16] Kaptein R, Kamps J, Hiemstra D. The impact of positive, negative and topical relevance feedback. In Proceedings of the 17th Text Retrieval Conference (TREC-17). NIST, Gaithersburg, MD; 2008.

© 2019 Bawa et al.; This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Peer-review history:

The peer review history for this paper can be accessed here (Please copy paste the total link in your browser address bar)

<http://www.sdiarticle3.com/review-history/47118>