Paper SAS1951-2015

Using SAS® Text Analytics to Examine Labor and Delivery Sentiments on the Internet

Michael D. Wallis, SAS Institute Inc.

ABSTRACT

In today's society, where seemingly unlimited information is just a mouse click away, many turn to social media, forums, and medical websites to research and understand how mothers feel about the birthing process. Mining the data in these resources helps provide an understanding of what mothers value and how they feel. This paper shows the use of SAS® Text Analytics to gather, explore, and analyze reports from mothers to determine their sentiment about labor and delivery topics. Results of this analysis could aid in the design and development of a labor and delivery survey and be used to understand what characteristics of the birthing process yield the highest levels of importance. These resources can then be used by labor and delivery professionals to engage with mothers regarding their labor and delivery preferences.

INTRODUCTION

Ranging from early signs and symptoms of pregnancy to postpartum health, a plethora of pregnancy-related topics are discussed today in online social media forums. Determining what are the most significant and meaningful topics from the seemingly endless supply of patient-provided, unsolicited responses can be challenging. Nevertheless, providing a direct "voice of the patient" to researchers and medical personnel provides extremely valuable insight as to what matters most to mothers.

Social media and other related unstructured text-based sources serve as both informational and educational content providers for patients who seek guidance and advice for medical and health-related matters. Labor and delivery Internet-based resources enjoy a large online presence via forum websites. Such forums enable mothers to engage in dialogue with other mothers in order to gain a better understanding of their personal experiences with various childbirth and delivery types. A comparative analysis between different delivery types can provide useful information that can help mothers better understand their options in determining which method of delivery best suits their needs and preferences.[1]

This paper shows how SAS Text Analytics was used to find differences, in terms of topics and sentiment. It describes how the data were collected and prepared for analysis, followed by processing of those data in topic discovery and phrase extraction and association. Results from these steps were then used in sentiment modeling and taxonomy development. Topics help to navigate you toward patient areas of interest while associated phrases aid in describing specific topics. Sentiment analysis results are then used as input into survey design for measuring patient satisfaction. Figure 1 highlights the methodology flow used in this study.

---

[1] All results are considered to be experimental, are not authoritative, and do not represent the opinions of SAS or the authors.

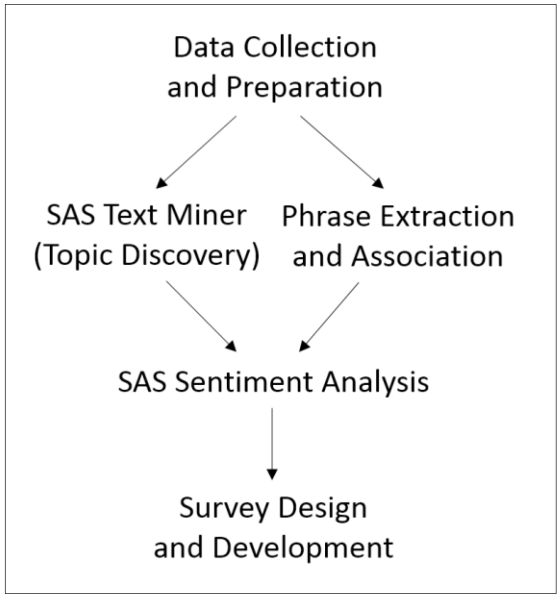

Figure 1. Methodology Flow Diagram

DATA COLLECTION AND PREPARATION

Data in this analysis come from unsolicited forum posts that patients provided on the Internet about three delivery types: vaginal (normal) delivery, cesarean section (often called C-section), and vaginal birth after cesarean (VBAC). Other delivery methods were also considered for analysis but ultimately excluded in order to simplify the study. All data are from January 1, 2014, to December 31, 2014.

Data collection and preparation usually consist of the following steps:

1. **Data collection**: In order to locate documents relevant to each delivery type, a data aggregation service was queried as follows: queries first specified what should be contained in each document, and then specified what should not be contained in the document. Quoted terms and phrases can be specified using syntax that impacts the way a match occurs. String literals can be used to perform exact string matching as in the example “vaginal delivery.” Negation (prefixing a hyphen in front of one or more quoted strings) can be used to exclude matching documents. Wildcard characters (*!?) can be used to generalize queries for allowing variations in matching documents. Terms that contain wildcard characters enable character substitution and expansion for various forms of a particular string. For example, the following query locates documents that are specific to vaginal delivery and do not mention cesarean sections or VBACs:

```
(“vaginal delivery”)(-“cesarean”|”section!”|”c*section!”|
   “vaginal delivery after cesarean”|”VBAC!”)
```

   Similarly, the following two queries locate documents that are specific to cesarean section and VBAC, respectively:

```
(“cesarean”|”section!”|”c*section!”)
   (-“vaginal delivery”|”vaginal delivery after cesarean”|”VBAC!”)
(“vaginal delivery after cesarean”|“VBAC!”)
   (-“vaginal delivery”|”c*section!”)
```

   Each query results in a collection of documents, called a document set, for that delivery type.

2. **Data cleaning and segmentation**: For this study, documents that exist in more than one document set are excluded from the analysis. For example, a document that discusses both vaginal and cesarean section delivery types is removed from further analysis. An exception is made for the VBAC delivery type because the term “vaginal” and “cesarean” are part of the delivery type name. However, the phrases “vaginal delivery” and “c*section!” are kept so that documents that contain these strings continue to be excluded.

   In social media and forum-based environments, it is common for multiple participants to post messages about a central topic and for their posts to include excerpts or entire copies of previous posts. This duplicated text can skew results and provide a false representation to further analysis. Documents that are considered to be either duplicates or near-duplicates of other documents are excluded from the final data set.

3. **Deleting non-relevant documents from the document sets**: A final data preparation step is to remove off-topic, nonrelevant documents from the document sets. When you specify queries that attempt to generalize in order to capture the greatest amount of relevant content, the possibility of matching documents that are “noise” increases. In nearly all data collection processes, one goal is to select as many documents as possible (high recall) that are relevant (high precision).

The queries in step 1 have terms and phrases that are unambiguous, so they result in documents that are relevant to their corresponding delivery type and step 3 is not necessary. Table 1 shows the final number of documents.

Delivery Type	**Total**
Vaginal delivery	15,596
Cesarean section	24,267
Vaginal birth after cesarean	23,794
Total	63,657

Table 1. Number of Documents by Delivery Type, after Filtering

DATA ANALYSIS METHODS

When you have prepared your data, you can analyze them by discovering the main topics, extracting and associating significant terms and phrases, and analyzing sentiments. These tasks are described in greater detail in the following sections.

PREPARING THE DATA AND DISCOVERING THE MAIN TOPICS

SAS® Text Miner, a component of SAS Text Analytics, enables you to perform fully automated, context-free topic discovery. For each delivery type in this study, a process flow diagram is created in SAS Text Miner to perform topic analysis as shown in Figure 2.

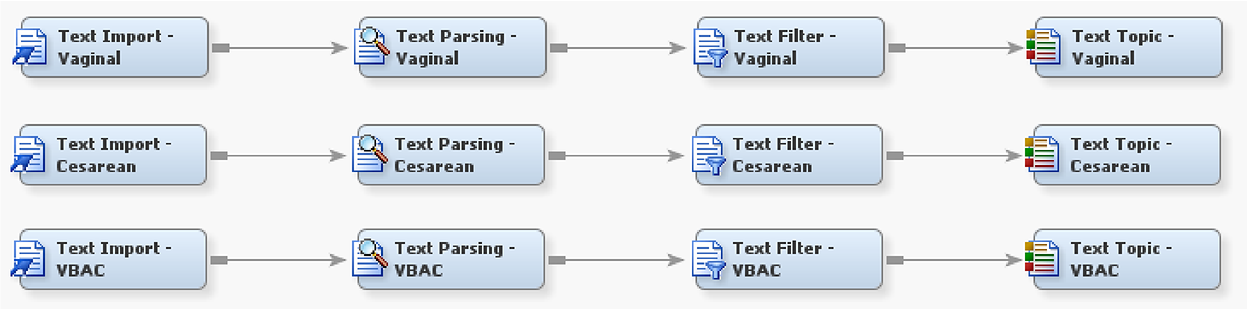

Figure 2. SAS Text Miner Text Topic Process Flow Diagram

Each delivery type uses the same node layout: (1) a node for importing text, (2) a node for parsing text, (3) a node for filtering text, and (4) a node for text topic discovery. The Text Import, Text Parsing, and Text Filter nodes serve as data preparatory steps, and the Text Topic node is used for topic discovery.

Each node can be configured to operate in specific way, thus enabling you to customize and gain further control over your analysis.

The following list describes what happens in each node:

1. In the Text Import node, you can specify the maximum number of characters that will be used for the TEXT variable in the output data set. For this study, the maximum value (32,000) of this parameter is recommended in order to capture as much of the patient posting as possible.

2. The Text Parsing node enables you to configure different parts of speech and noun groups to help you determine how best to parse the unstructured data and how to handle domain-specific terminology. You can also enable term stemming, a process that reduces the different morphological expansions of a term to their root form to help decrease the total number of terms used for analysis. For example, the terms “delivering” and “delivered” reduce to the root term “deliver.” And you can specify additional synonyms in order to align terms that have the same meaning. For example, terms such as “labor” and “labour” are treated the same and are grouped under the parent term “labor.” In this study, the Text Parsing nodes are configured with synonyms that are relevant to their respective delivery type, and multiterm entities specify terms that should be processed together. For example, the phrase “vaginal birth after cesarean” is treated as a single term and is subsequently mapped to the parent term “VBAC.”

3. The Text Filter node enables you to activate spelling correction for misspelled words, apply frequency and term weighting, and term and document filtering. Forums tend to impose little to no restriction on grammatical structure and spelling, and it is quite common for unstructured data to contain misspelled terms. Spelling correction enable you to reduce the total number of parsed terms and documents. In addition, the Text Filter node provides an interactive viewer that enables you to view which terms are preserved and which terms are dropped from the analysis.

   Another useful feature of this node is concept linking, which enables you to view other terms that are strongly associated with a particular term. For example, when you examine the results that the Text Filter node produces for the VBAC document set, you can gain a better understanding of what other terms are the most strongly associated. Figure 3 shows an expansion of the most strongly associated terms for the term “know.” You might infer from this expansion that the patient is seeking information about recovery times or hospitalization. The term “water” might indicate that the patient is seeking information about water births or whether their water is to be broken (rupturing of the amniotic sac) for inducing labor. Figure 4 show an expansion of the term “second,” which has a strong association with the term “VBAC” and might suggest that women are discussing a second childbirth after previously having a cesarean section delivery.

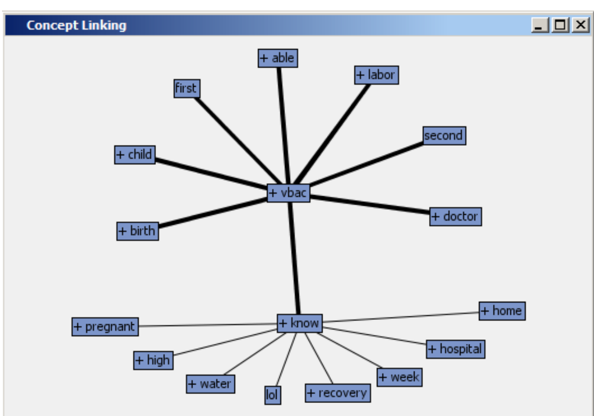

Figure 3. Concept Linking for the Term “*know*”

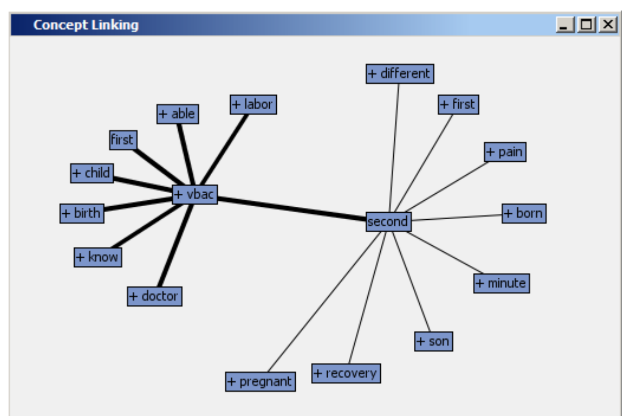

Figure 4. Concept Linking for the Term “*second*”

4. The Text Topic node enables you to explore document collections by automatically associating documents and terms with a main theme or idea. Documents and terms are automatically associated with these main themes or ideas under topics. The Text Topic node assigns a score for each document and term to each topic. Thresholds are then used to determine whether the relationship is

strong enough to consider that the document or term belongs to each topic. Thus, documents and terms can exist in more than one topic (or in none at all).

The Text Topic node generates a term-by-document frequency matrix from all the documents that are contained in the document collection. When you have large numbers of documents, this matrix can grow to contain hundreds of thousands of words, making it too costly in computing time and memory to process effectively. To improve performance, you can apply a dimensionality technique known as singular value decomposition (SVD) to reduce the matrix to a lower-dimensional form, making it easier to process. Multiterm topics are created by taking the rotated SVD of the weighted term-by-document matrix, where similar documents are grouped together. These topics are important because they represent the most discussed main themes or ideas, and they provide the basis for topics to be modeled for sentiment in subsequent analysis. This study requests that a maximum of five multiterm topics be discovered for each delivery type. The tables for each delivery type are combined in Table 2.

Delivery Type	Topic	Document Count
Vaginal delivery	+week,+bleed,+pp,+tear,+stitch	2,410
	+hour,+labor,+epidural,+contraction,+induce	2,067
	+breech,+head,+baby,+week,+deliver	2,035
	+woman,+pregnancy,+risk,+child,+year	1,540
	+pay,+insurance,+hospital,+bill,+cost	1,266
Cesarean section	+scar,+month,+year,+pain,+heal	3,608
	+birth,+natural,+natural birth,+emergency,+emergency c-section	3,456
	+hour,+room,+hospital,+pain,+recovery	3,448
	+schedule,+date,+week,+due,+due date	3,325
	+woman,+doctor,+risk,+medical,+write	2,975
Vaginal delivery after cesarean	+section,+birth,+consultant,+recovery,+vaginal	3,794
	+cesarean,+week,+successful,+induce,+due	3,759
	+hospital,+midwife,+find,+supportive,+vbacs	3,632
	+risk,+rupture,+uterine,+uterine rupture,+woman	2,996
	+contraction,+hour,+water,+start,+push	2,786

Table 2. Top Five Discovered Text Topics by Delivery Type

The Text Topic node assigns a score for each document and term for each topic. Threshold values are applied to determine whether the scored documents and terms have associations that are strong enough to warrant their inclusion in each topic. After the documents are scored, the top five most strongly associated terms in each topic are displayed in a row of the topic table and are also used as the topic name. In addition to having the strongest association in each topic, these five terms serve a secondary purpose: to represent what is contained in the topic. For example, in Table 2, in the “cesarean section” delivery type, the topic named “+hour,+room,+hospital,+pain,+recovery” indicates that patients are discussing matters related to hospital rooms, pain, and recovery. It could be that patients are expressing concerns over hospital room and environmental conditions, along with pain management and recovery time.

The “vaginal delivery after cesarean” delivery type contains topics that indicate patient discussions about care providers, including consultants and midwives. For example, the topic named “+hospital,+midwife,+find,+supportive,+vbacs” indicates that patients are discussing matters related to medical facilities, care providers, and support. A possible explanation for this topic is that patients are attempting to locate hospitals or midwives that support a vaginal delivery after a cesarean section has been performed. Vaginal birth after a cesarean section is known to have additional risks for mothers and their newborns (Cahill et al. 2006). Risks include, but are not limited to, repeated cesarean sections (also known as failed trial of labor after cesarean), uterine infection, and uterine rupture. From this, you might infer that mothers are assessing the risks associated with VBAC and are seeking consultation with medical facilities or birthing personnel for electing to undergo a VBAC.

EXTRACTING AND ASSOCIATING SIGNIFICANT TERMS AND PHRASES

Sets of individual terms provide a limited picture of what is contained in a document collection. The use of multiterm phrases can provide a more precise representation and is preferred when you perform sentiment analysis. As noted in Reckman et al. (2014), "expressions of sentiment often consist of more than one word and can contain negations or other modifiers." As a result, phrase extraction involves capturing sequences of terms that often occur together. When such sequences are detected, they are treated as a single entity rather than as individual terms. Such a single entity has more relevance than any of the individual terms in the sequence or any other combination of them. In addition, phrases can be related. Because the terms *vaginal delivery* and *delivery procedure* have a strong association, *vaginal delivery procedure* is also treated as a phrase, assuming it occurs within the document collection. This study adopts the methodology used in Reckman et al. (2014).

Multiterm phrases are identified when the actual frequency of collocated terms exceeds their expected frequency by a specified chi-square (X^2) statistic. For example, for the phrase *natural birth*, you can count the number of times the individual terms occur: *natural* occurs 7,060 times and *birth* occurs 17,042 times. The total number of words for the document collection is 7,303,435. Using these frequencies and assuming that terms are randomly distributed, you can calculate the expected frequency (how often you expect to see the two terms occur together) by multiplying the frequencies of the individual terms and then dividing by the total number of terms in the document collection. Thus, the expected frequency of the *phrase natural* birth is 16 (7,060 * 17,042 / 7,303,435). The actual frequency for this phrase is much higher (2,381 times higher) than the expected frequency. You can use the chi-square (X^2) statistic to show that the difference between the observed and expected frequencies is statistically significant. Observed and expected frequencies are shown in Table 3 and Table 4, respectively.

Observed	+birth	-birth	Total
+natural	5,597	1,463	7,060
-natural	11,445	7,284,930	7,296,375
Total	17,042	7,286,393	7,303,435

Table 3. Observed Frequencies

Expected	+birth	-birth	Total
+natural	16	7,044	7,060
-natural	17,026	7,279,349	7,296,375
Total	17,042	7,286,393	7,303,435

Table 4. Expected Frequencies

For the phrase *natural birth*, X^2 = 1,896,648.197. This value surpasses any reasonable threshold, so the phrase is considered to be one of the strong collocations in the document set. As a result, the terms are taken together as a phrase rather than as two single words.

Adjusting the threshold can yield different phrases, enabling you to see which phrases you want to keep in your analysis. You can ensure that the associations between terms in a phrase are statistically significant by making sure that the X^2 value of the phrase is less than chi-square threshold values that correspond to various levels of significance shown in Table 5.

p-value	Chi-Square Threshold
0.01	6.63
0.005	7.88
0.001	10.83
0.0005	12.12

Table 5. Minimal Chi-Square Values for Statistical Significance

A common problem in processing unstructured data is dealing with low-frequency terms. Many terms occur only once throughout a document collection, making it difficult to accurately determine whether the terms have any special relationship with any of their neighboring terms. Because chi-square does not account for this problem, you need to take additional steps to ensure adequate frequency estimation. Terms that occur only once throughout the document set have a low probability of occurring together with other terms. Nevertheless, such combinations can appear much more frequently than expected. Frequency smoothing enables you to allocate probability mass both to terms in the document set and to terms that have yet to be seen. The frequency smoothing technique used in this study is called "Simple Good-Turing" (Gale and Simpson 1995).

Textual Extraction and Tokenization

Textual extraction and tokenization is performed to extract the strongest terms and phrases that are associated with a particular delivery type. In order to obtain these associations, you can apply the TGFILTER and TGPARSE procedures in SAS Text Miner to your document collection to extract plain text from documents and to reveal information they contain, respectively. Each document is processed in isolation. Isolating the processing simplifies the separation of per-document results and subsequent *n*-gram analysis (a probabilistic language model that follows an (n-1)-order Markov model and is used to predict the next item within a sequence of terms).

The following SAS code shows how to invoke PROC TGFILTER and PROC TGPARSE to perform text extraction and parsing on one document at a time in a directory in order to facilitate the *n*-gram processing:

```
proc tgfilter
 out=documents
 srcdir="C:\labor_and_delivery\data\source"
 numchars=32767;
run;

proc tgparse
 data=documents
 key=parsekey
 out=parseout
 outoffset=parseoutoffset
 stemming=yes;
 var text;
select "punct" / drop;
run;
```

PROC TGFILTER generates a data set that contains the results of the textual extraction process, which are used later in text parsing. The OUT= option specifies the name of the output data set, and the SRCDIR= option specifies a directory from which to read single temporary files that are created from the document collection. Each document in the document collection is copied into this directory for each invocation of the preceding code.

Tokenization of delivery type data is performed using PROC TGPARSE. This procedure enables you to specify options, including term stemming and punctuation, that direct how to tokenize the data. Term stemming reduces a term down to its root form. For example, the term "ruptures" is stemmed to "rupture."

Punctuation can be preserved or removed from parsing by using the SELECT statement. In this study, punctuation is dropped from the results.

After the PROC TGFILTER and PROC TGPARSE steps, you can use the SORT and EXPORT procedures (not shown here) to sort and export the results. PROC SORT is used to sort the ParseKey data set in ascending order by the variable name Key, and to sort the ParseOutOffset data set in ascending order by the variable name _OFFSET_. The Key variable uniquely identifies a term in the ParseKey data set, and the _OFFSET_ variable provides the byte offset of the term in the document. Both data sets are then exported to a comma separated value (CSV) file by using PROC EXPORT. From these exported results, you can construct a document that uses the terms produced by PROC TGPARSE and that preserves the original term order by the byte offset value.

The top 10 highest-ranked phrases for each delivery type are shown in Table 6.

Ranking	Vaginal Delivery	Cesarean Delivery	Vaginal Birth after Cesarean
1	vaginal delivery	c section	vbac
2	pay	staple	csection
3	stitches	pain	rupture
4	bill	numb	rcs
5	bleed	skin	supportive
6	head down	tug	labor
7	estradiol	surgery	midwife
8	out of pocket	the spinal	provider
9	vertex	recovery	uterine rupture
10	down there	scar	hospital

Table 6. Highest-Ranked Phrases for Delivery Types, by Association Strength

Phrase Association

To find out which phrases are related with each delivery type, you can follow a method similar to the one described in the section “Extracting and Associating Significant Terms and Phrases.” First, you count how often phrases occur in each document that is specific to a particular delivery type. Next, you count how often the same phrases occur through the entire document set. Finally, you compare these two frequencies. Because you have three delivery types, you will have three lists of phrases. The first list contains phrases that occur much more frequently about vaginal delivery, a second list contains phrases that occur much more frequently about cesarean section, and a third list contains phrases that occur much more frequently about vaginal birth after cesarean. Each phrase has a corresponding association strength. You order each list by this association strength in decreasing order, so that the most frequent phrases are positioned at the beginning of the list.

ANALYZING SENTIMENTS

SAS® Sentiment Analysis enables you to analyze documents for expressions of sentiment (positive, negative, neutral, or unclassified). Sentiment models are constructed using the concept of products (which correspond to any object, person, event, or experience) and features (which correspond to attributes about the product). Sentiment is expressed and analyzed at both the product and feature levels. In this study, delivery types serve as the products, and the discovered topics serve as the features. The sentiment model tracks sentiment about each delivery type and tracks mentions of the most significant topics that SAS Text Miner produced: care providers, hospitalization, locations, medications, recovery, and risks.

The sentiment model for the three delivery types is customized by assembling a list of potential names, noun phrases, and abbreviations that could be used to refer to each delivery type. This list is then used

as input into sentiment rules that identify delivery types in document texts. Rules specify matching criteria for a particular product or feature, and they can be made up of different rule types. Two rule types are classifier rules and concept rules.

Classifier rules specify strings that are used to match text exactly as specified in the rule body. You can create classifier rules for specifying strings that you expect, such as variations in spelling. Classifier rules are also efficient because they do not require further expansion to cover additional morphological expansion.

Using concept rules, you can supply the ‘@’ symbol to enable term expansion (also known as morphological expansion). For example, “caesarean section@” expands to cover both “caesarean section” and “caesarean sections.”

The following SAS Sentiment Analysis rules include classifier and concept rules:

```
CLASSIFIER:caesarean
CLASSIFIER:cesarean
CLASSIFIER:caesarian
CONCEPT:c section@
CONCEPT:c-section@
CONCEPT:caesarean section@
CONCEPT:abdominal delivery@
```

Additional rule types exist and can include ways to match terms and phrases by using order and distance, both of which are used in other rules in this sentiment model. Rules are specified to capture positive and negative sentiment expressions for both products and features. For example, the following excerpts of positive sentiment rules are defined for the cesarean section delivery type:

```
PREDICATE_RULE:(SENT, "_def{CESAREAN} _a{_def{PositiveSubject}}")
PREDICATE_RULE:(UNLESS, "_def{DEFCESAREANBlockers}",
       (ORDDIST_10, "_a{_def{Positive}}", "_def{CESAREAN}"))
```

The first rule captures positive sentiment in a sentence by looking for any of the definitions that are specified for cesarean sections followed by a positive subject (a positive statement). The second rule indicates that a positive sentiment results if a positive match occurs before a cesarean section match and is within an ordered distance of 10 terms away from it unless a match that would block a cesarean section occurs between the two. The ordered distance operator, ORDDIST_N, specifies the order and distance between the terms or concepts that you want to match. Blocking a match refers to any terms or concepts that could prevent a match from occurring. This rule enables you to associate mentions of positive sentiment with a particular delivery type while ensuring that no other delivery types occur between the positive mention and the delivery type being examined. The “Positive” concept corresponds to general, context-free mentions of positive sentiment and includes (but is not limited to) positive words, phrases, emoticons, and slang. The “DEFCESAREANBlockers” concept identifies phrases that block the other two matches specified in the ORDDIST_10 operator. Blocking definitions consist of both context-free phrases (such as “as compared to” and “in contrast to”) and domain-specific phrases (which are definitions of the other delivery types: “vaginal delivery” and “vaginal delivery after cesarean”).

Negative sentiment rules are identical to positive sentiment rules in syntactic and linguistic structure. For each positive rule this study, there exists a corresponding negative rule that is parameterized with the names of equivalent negative concepts. For example, where a positive rule might have used the concept named “PositiveSubject,” the negative rule contains a concept named “NegativeSubject.” The following excerpts of negative sentiment rules are defined for the cesarean section delivery type:

```
PREDICATE_RULE:(SENT, "_def{CESAREAN} _a{_def{NegativeSubject}}")
PREDICATE_RULE:(UNLESS, "_def{DEFCESAREANBlockers}",
       (ORDDIST_10, "_a{_def{Negative}}", "_def{CESAREAN}"))
```

The methodology described in the section “Extracting and Associating Significant Terms and Phrases” provides additional terms and phrases that can be used to associate sentiment with a product or feature. Such phrases yield language patterns that are frequently found in labor and delivery dialogue and can supplement the out-of-the-box sentiment taxonomy for improving the matching of mentions of sentiment against a product or feature. Figures 5 through 7 show the sentiment distribution for each delivery type.

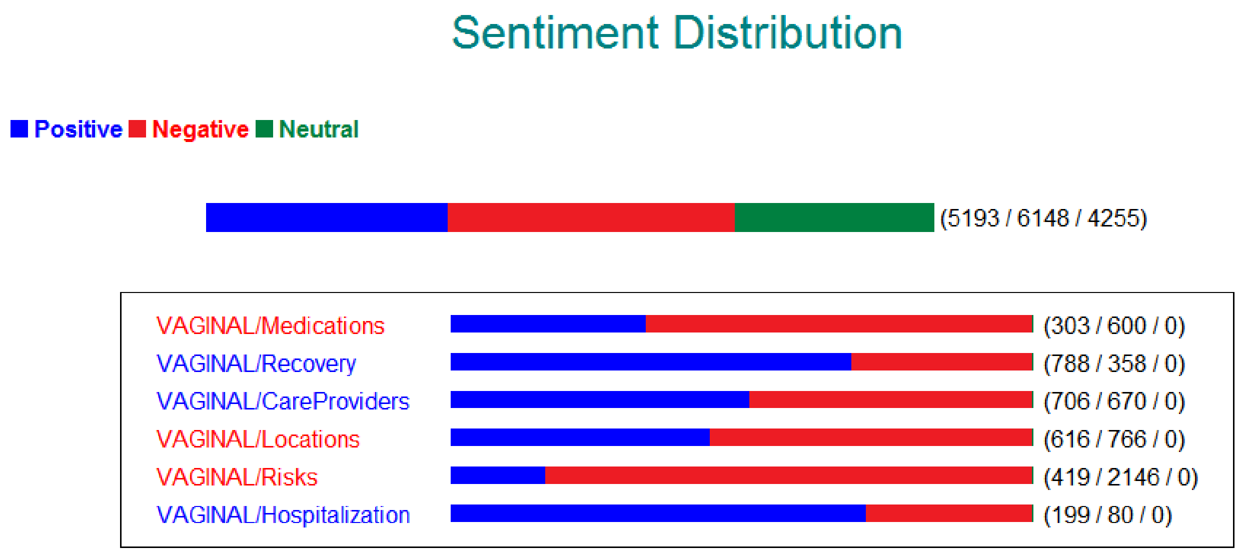

Figure 5. Sentiment Analysis of Vaginal Delivery

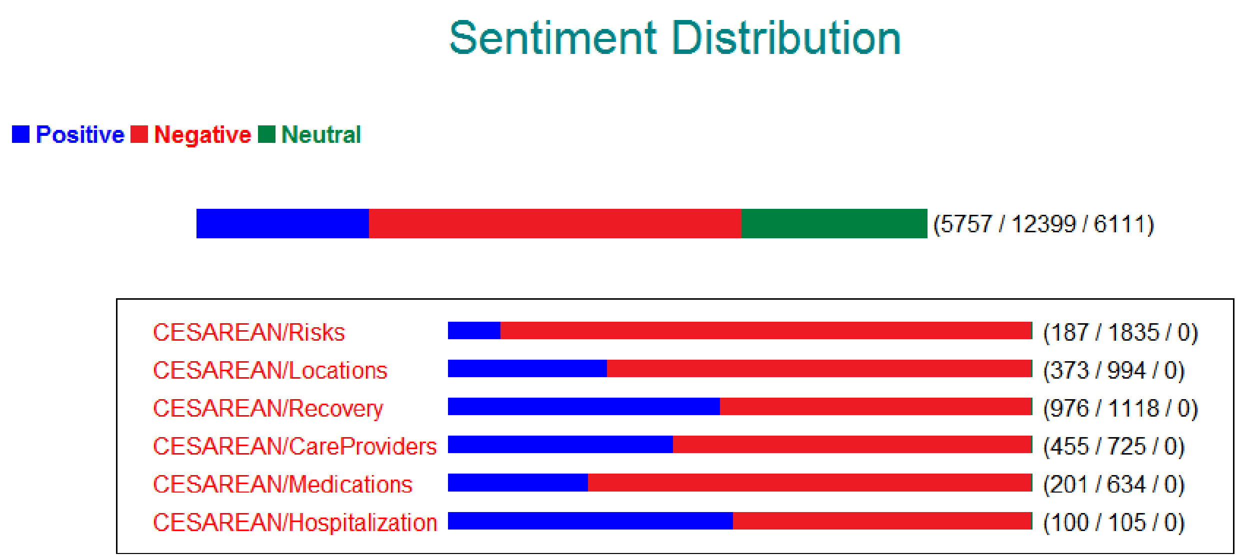

Figure 6. Sentiment Analysis of Cesarean Delivery

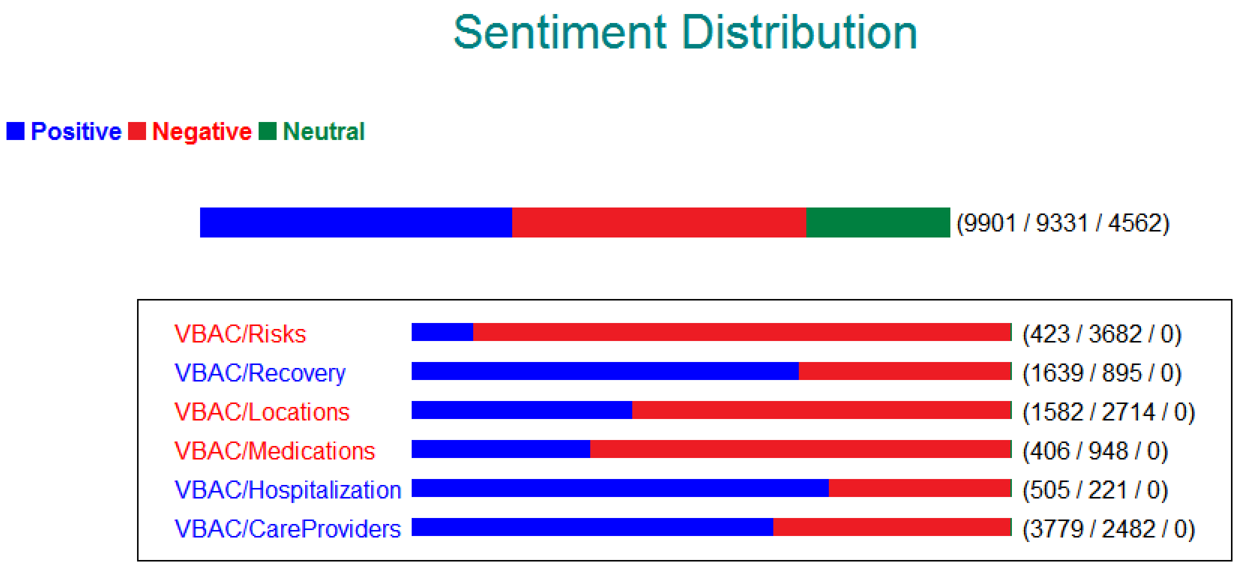

Figure 7. Sentiment Analysis of Vaginal Birth after Cesarean

Table 7 shows topics and their associated positive and negative sentiments for each delivery type.

Vaginal Delivery			Cesarean Section			Vaginal Birth after Cesarean			Topics
POS	**NEG**	**TOT**	**POS**	**NEG**	**TOT**	**POS**	**NEG**	**TOT**	
706	670	1,376	455	725	1,180	3,779	2,482	6,261	*CareProviders*
199	80	279	100	105	205	505	221	726	*Hospitalization*
616	766	1,382	373	994	1,397	1,582	2,714	4,296	*Locations*
303	600	903	201	634	835	406	948	1,354	*Medications*
788	358	1,146	976	1,118	2,094	1,639	895	2,534	*Recovery*
419	2,146	2,565	187	1,835	2,022	423	3,682	4,105	*Risks*

Table 7. Feature Mentions by Delivery Type

Risks is the most discussed topic for vaginal delivery, the second-most discussed topic for cesarean section, and the third-most discussed topic for VBAC. All three delivery types have predominantly negative sentiment for the *Risks* topic: 83.6% for vaginal delivery, 90.8% for cesarean section, and 89.7% for VBAC. Looking closer at the document collections for each of the delivery types yields more insight into what risks are being discussed. For example, the cesarean section document set mentions risk surrounding both mother and baby and the impacts of major surgery. According to Narins (2013), George Macones, MD, a representative of the American College of Obstetricians and Gynecologists (ACOG), a nonprofit women's advocacy group, notes that cesarean section deliveries are a riskier form of delivery because abdominal incisions expose the mother to bacteria, thus increasing the chances for infection. Additional risks include excessive bleeding and blood clots, and in some rare cases, death. Along with the physical toll on the mother's body and long recovery time, women who have cesarean sections face an increased risk of complications in future pregnancies. The vaginal delivery document set mentions risk associated with breech birth (feet-first position), baby positioning, head entrapment, and umbilical cord prolapse. The VBAC document set mentions breech birth, along with bleeding, infection, failure to progress, and higher risks to both mother and baby.

Recovery yields opposing sentiment types for cesarean section versus vaginal delivery and VBAC: 68.8% positive for vaginal delivery and 64.9% positive for VBAC, but 53.4% negative for cesarean. Explanations for this result could be that cesarean section procedures require surgical incisions in a mother's abdomen and uterus, are more invasive, and have recovery times estimated to be from six to eight weeks (Narins 2015). Compared to a vaginal delivery recovery time of one to two weeks (Narins 2015), you can see that cesarean sections clearly have a much greater impact on recovery time. As in the vaginal delivery and VBAC document sets, mothers express positive sentiment about their delivery experience much more often than mothers who underwent a cesarean section. According to the National Vital Statistic Reports (NVSS) 2012 report by the Centers for Disease Control and Prevention (CDC) (Martin et al. 2012), cesarean birth accounted for 32.8% of all US births and has remained unchanged since 2010. Prior to 2010, the cesarean delivery rate had increased every year since 1996, when approximately one-fifth of births were delivered by cesarean section. For births at less than 39-weeks gestation, cesarean sections peaked in 2009 at 38.3% and have declined every year since. You could hypothesize that after mothers become more familiar with the cesarean section procedure and the impact that it has on the body, they decide that the recovery process is too difficult and prefer a vaginal delivery.

The *Locations* topic yields negative sentiment for all three delivery types: vaginal delivery 55.4%, cesarean section 71.2%, and VBAC 63.2%. In all three document sets, expressions of negative sentiment addressed complications in delivery process, feelings of lack of attention by medical staff, and perceived delays in medical care.

Hospitalization for vaginal delivery and VBAC yields positive sentiment with percentages of 71.3% and 69.6%, respectively. Cesarean section has a nearly even split with 51.2% negative sentiment. Vaginal delivery and VBAC positive sentiment show mothers have hospitalization preferences that are non-eventful, equating to a planned or expected labor and delivery process that is free of complications or interference.

The *Medications* topic yields negative sentiment for all three delivery types: 66.4% for vaginal delivery, 75.9% for cesarean section, and 70% for VBAC. Negative sentiment expressions for all three delivery types address failure or partial effectiveness of medications (for example, sensation present in some part of the body) or lack of pain relief provided by the medication. VBAC reports include negative sentiment surrounding discussions on forced epidurals.

CareProviders are viewed more positively for both vaginal delivery at 51.3% and VBAC at 60.4%, whereas 61.4% are viewed more negatively for cesarean section. Mothers who chose to undergo a VBAC expressed positive sentiments about the use of midwives and doulas, specifically focusing on the themes of support, care, guidance, and reassurance.

CareProviders, *Hospitalization*, and *Recovery* were the only topics that had more positive sentiment for both vaginal delivery and VBAC versus cesarean section.

LABOR AND DELIVERY SURVEY

The main objective of this study is to apply SAS Text Analytics to the unsolicited, patient-provided Internet forum postings in order to reveal insight about preferences surrounding the subject of labor and delivery. Results from this study can aid personnel who work in the area of labor and delivery to navigate the most significant and most often discussed topics among a large collection of patient-provided forum postings.

Satisfaction is one of the most reported outcome measures reported for quality of care (Sawyer et al. 2013). Women's satisfaction for maternity care is of great importance to medical providers, administrators, and policy makers (Hodnett 2002 and Redshaw 2008). Surveys are one of the most common forms of assessing patient satisfaction. As research shows, part of women's satisfaction with childbirth is partially related to the health and well-being of the mother and her child (Sawyer et al. 2013). As a result of these findings, you could suggest that on a per-delivery type, the discovery of the most prominent topics in a body of documents, coupled with yielding the most significant phrases contained in those documents, could be used as a basis for the design and construction of a survey.

Measures of sentiment parallel levels of satisfaction. Strong positive sentiment corresponds to a high level of satisfaction, whereas strong negative sentiment corresponds to a low level of satisfaction. This suggests that different levels of satisfaction not only can be aligned closely with a specific type of

sentiment (for example, positive, negative, neutral) but also can have a measure of association with that type of sentiment (for example, strongly positive or slightly negative). SAS Sentiment Analysis aids in these measures of association by providing scoring results for overall measures of sentiment, probability of being a certain level of sentiment, and a level of confidence associated with each measure of probability.

Many surveys have been developed that consist of a multi-item scale for measuring patient satisfaction for labor and delivery and that include psychometric content such as information about survey construction, reliability, and validity. Part of survey construction consists of item generation, which is regarded as the portion of the design where you research and include important items relevant to satisfaction research. Those items include, but are not limited to, opinions from focus groups of mothers (Sawyer et al. 2013). Patient forums can be considered to be a pseudo form of a focus group. Input that women provide directly represents what is most important to them and should be treated as contributions to item-level design and development of surveys.

The study in this paper indicates that mothers are predominantly interested in risks and recovery that are associated with labor and delivery. Further examination of these topics, in addition to topics that are not as dominant in the overall document collection, provide insight as to what preferences women have regarding the childbirth process.

CONCLUSION

This study finds that using text mining, combined with sentiment analysis, provides a more comprehensive view and understanding of which topics and sentiment are associated with patient labor and delivery preferences. By sampling a larger number of unsolicited, patient-provided forum-based postings, this study attempts to identify and capture the most significant and meaningful topics that are associated with specific delivery types. These results can be further paired with an application of sentiment analysis at the document, product, and feature-levels, ultimately to yield insight into what preferences patients have surrounding their labor and delivery experiences.

REFERENCES

Bane, M. 2011. Python implementation of Gale and Sampson's "Simple Good-Turing" algorithm. Available at https://raw.github.com/maxbane/simplegoodturing/master/sgt.py. GitHub. June 21. Accessed February 12, 2015.

Cahill, A., Stamilio, D., Odibo, A., Peipert, J., Ratcliffe, S., Stevens, E., Sammel, M., and Macones, G. 2006. "Is Vaginal Birth after Cesarean (VBAC) or Elective Repeat Cesarean Safer in Women with a Prior Vaginal Delivery?" *American Journal of Obstetrics and Gynecology* 195:1143–1147.

Gale, W., and Sampson, G. 1995. "Good-Turing Frequency Estimation without Tears." *Journal of Quantitative Linguistics* 2:217–237.

Hodnett, E. 2002. "Pain and Women's Satisfaction with the Experience of Childbirth: A Systematic Review." *American Journal of Obstetrics and Gynecology* 186:S160–S172.

Martin, J., Hamilton, B., Osterman, M., Curtin, S., and Matthews, T. 2012. "Births: Final Data for 2012." Center for Disease Control and Prevention. *National Vital Statistics Report* 62:8.

Narins, E. 2013. "What You Don't Know about C-Sections." *Women's Health*, July 23. Available at http://www.womenshealthmag.com/health/what-you-dont-know-about-c-sections. Accessed February 2, 2015.

Reckman, H., Baird, C., Crawford, J., Crowell, R., Micciulla, L., Sethi, S., and Veress, F. 2013. "Teragram: Rule-Based Detection of Sentiment Phrases Using SAS Sentiment Analysis." Paper presented at International Workshop on Semantic Evaluation at SEM, NAACL 2013, Atlanta.

Reckman, H., Wallis, M., Crowell, R., Micciulla, L., and Baird, C. 2014. "Summarizing and Highlighting Differences in Senate Race Data Using SAS Sentiment Analysis." *Proceedings of the SAS Global Forum 2014 Conference*. Cary, NC: SAS Institute Inc. Available http://support.sas.com/resources/papers/proceedings14/SAS349-2014.pdf.

Redshaw, M. 2008. "Women as Consumers of Maternity Care: Measuring 'Satisfaction' or 'Dissatisfaction'?" *Birth* 35:73–76.

Sawyer, A., Ayers, S., Abbott, J., Gyte, G., Rabe, H., and Duley, L. 2013. "Measures of Satisfaction with Care during Labour and Birth: A Comparative Review." *BMC Pregnancy and Childbirth* 13:108.

ACKNOWLEDGMENTS

The author would like to acknowledge Dr. Julie Ivy, Associate Professor, and Karen Hicklin, Ph.D. student, of the Edward P. Fitts Department of Industrial and Systems Engineering, North Carolina State University, for their expertise in the subject areas of delivery preferences and labor and delivery.

The author would also like to recognize Cheyanne Baird for her guidance and expertise in sentiment model design and development. Cheyenne's in-depth understanding of modelling sentiment and subsequent knowledge transfers provided the foundation for constructing the labor and delivery sentiment model used in this study.

The author would also like to recognize Anne Baxter for her exhaustive and thorough editorial assistance, as her input has significantly improved the content and quality of this paper. Anne's insightful contributions have been vital throughout the time of this writing and are very much appreciated.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Michael D. Wallis
919-531-5886
michael.wallis@sas.com
http://www.sas.com