

How good are ideas identified by an automatic idea detection system?

Journal:	<i>Creativity and Innovation Management</i>
Manuscript ID	1587-CIM-A-12-2016.R3
Manuscript Type:	Special Issue Article
Keywords:	Idea quality, Text mining, Machine learning, Natural language processing, Big data, Online communities, Support vector machines, Crowdsourcing

SCHOLARONE™
Manuscripts

Review

How good are ideas identified by an automatic idea detection system?

Abstract

Online communities can be an attractive source of ideas for product and process innovations. However, innovative user-contributed ideas may be few. From a perspective of harnessing “big data” for inbound open innovation, the detection of good ideas in online communities is a problem of detecting rare events. Recent advances in text analytics and machine learning have made it possible to screen vast amounts of online information and automatically detect user-contributed ideas. However, it is still uncertain whether the ideas identified by such systems will also be regarded as sufficiently novel, feasible and valuable by firms who might decide to develop them further.

A validation study is reported in which 200 posts from an online home brewing community were extracted by an automatic idea detection system. Two professionals from a brewing company, evaluated the posts in terms of idea *content*, idea *novelty*, idea *feasibility* and idea *value*. The results suggest that the automatic idea detection system is sufficiently valid to be deployed for the harvesting and initial screening of ideas. And that the profile of the identified ideas (in terms of novelty, feasibility and value) follows the same pattern identified in studies of user ideation in general.

Introduction

Big data has been predicted to revolutionise innovation and how firms will create value for themselves, their customers and society (e.g., see McAfee & Brynjolfsson, 2012). Artificial intelligence systems that leverage big data allow more and more tasks to be solved in an automatic manner. Whilst in the past, these were predominantly tasks of a mundane and repetitive nature, advances in text analytics and machine learning have also made it possible to solve more complex problems (Christensen, Nørskov, Frederiksen, & Scholderer, 2017).

A problem that continues to occupy scholars and practitioners of new product development is how to obtain and select ideas for new products (e.g., di Gangi, Wasko, & Hooker, 2010; van den Ende, Frederiksen, & Prencipe, 2015; Frederiksen & Knudsen, 2017). In the context of inbound open innovation, Ooms, Bell and Kok (2015), for example, argue that firms can enhance their receptivity—i.e., their capacity to absorb more diverse external knowledge from more varied sources—by engaging with social media. Whilst this can in theory expand a firm's boundaries for information absorption, the extent of engagement with social media is still constrained by available staff time. Such constraints can to some degree be overcome if companies develop or adopt systems that automate parts of the absorption process.

The aim of the research presented here is to show how the performance of automated systems in areas such as inbound open innovation can be evaluated. On one hand, the study should be seen as a feasibility study of whether automated detection of ideas for product- and process innovations is actually possible. On the other hand, it should also be regarded as a validation study that probes the “veracity” and “value” aspects of big data (Gandomi & Haider, 2015) in the context of a specific application case.

Literature review

Big data

Big data has received much attention in recent years, but it is not a new concept as such. Big data can be seen as a product of digitalisation, the “digital footprint” of an electronically mediated reality (Zwitter, 2014). Others stress instrumental aspects, regarding big data as a tool for generating insights (Boyd & Crawford, 2012; Nunan & Di Domenico, 2013). We prefer to see big data as a *resource* whereas *tools* (such as text analytics, machine learning and other artificial intelligence techniques) help create value from the resource. As an analogy, one might think of big data as the *oil* and of artificial intelligence as the *combustion engine* that makes the oil useful. In technical terms, big data refers to databases that are too big to be handled by conventional data warehousing systems. The “bigness” of big data is often characterised in terms of three parameters: variety, velocity and volume (Hsinchun Chen, Chiang, & Storey, 2012). Variety refers to the heterogeneity of data types: part of the database content may be structured and numeric (e.g., transaction data from retail channels) but other parts may have different forms, for example free text, image and video files exchanged on social media networks. Velocity refers to how fast new data is being generated. In order to utilise newly generated data, it be retrieved on a continuous basis. Volume refers to the amount of data, measured in terms rows and columns, complexity of databases, and total storage volume.

Online communities as idea reservoirs

One domain where big data has been generated since the early days of the Internet is online communities on message board systems and social media networks. Online communities where users exchange experiences and discuss potential improvements for new products and processes have been identified as rich reservoirs of ideas that can fuel the

1
2
3 innovation processes of firms (Van de Ven, 1986; Ekvall, 1997; Vandenbosch, Saatcioglu, &
4 Fay, 2006; van den Ende, Frederiksen, & Prencipe, 2015). Ideas do not have to originate from
5 the creative mind of the firm's employees but can also originate from the users of its products,
6 services and technologies (Kristensson, Gustafsson, & Archer, 2004; Magnusson, 2009; von
7 Hippel, Ogawa, & de Jong, 2011; Poetz & Schreier, 2012; Majchrzak & Malhotra, 2013;
8 Magnusson, Wästlund, & Netz, 2014).

9
10
11
12
13
14
15
16 Prominent examples of the role of user communities in open inbound innovation are
17 the communities hosted by Dell (di Gangi, Wasko, & Hooker, 2010; Poetz & Schreier, 2012),
18 Lego (Antorini, 2007; Antorini, Muñiz, & Askildsen, 2012; Nørskov, Antorini, & Jensen,
19 2015), Propellerhead (Jeppesen & Frederiksen, 2006) and IBM (Mahr & Lievens, 2012).
20 Firm-hosted communities such as these have the advantage that the hosting firm can retain a
21 certain degree of control. The communities are typically based on software that allows
22 registered users to post ideas, comment on and vote for ideas posted by other users in a highly
23 structured manner. The downside of this approach is that it requires an extensive base of
24 committed product users or firm-loyal customers who have an intrinsic interest in suggesting
25 ideas to the firm.
26
27
28
29
30
31
32
33
34
35
36
37
38

39 However, users do not only gather in firm-hosted communities. A vast amount of
40 online communities exist that are firm-free (Füller, Bartl, Ernst, & Mühlbacher, 2006; Füller,
41 Jawecki, & Mühlbacher, 2007). The most prominent cases include open-source software
42 development communities such as those responsible for the Linux kernel, R and Python.
43 These are examples of firm-free “products” and platforms that have developed in a distributed
44 manner, utilising online collaboration tools such as GitHub and Sourceforge. The fact that the
45 resulting products are now perfectly able to compete with their commercial counterparts (such
46 as the products ranges of the SAS Institute or Microsoft) is a clear demonstration of the
47
48
49
50
51
52
53
54
55
56
57
58
59
60

1
2
3 potential of such communities (von Krogh, Spaeth, & Lakhani, 2003; von Krogh & von
4
5 Hippel, 2006)

6
7
8 The problem with firm-free communities is that they, unlike most firm-hosted
9
10 communities, are usually *not* based on a crowdsourcing architecture that would enable easy
11
12 harvesting and collaborative filtering of the community-generated ideas. Assigning employees
13
14 to manual monitoring of community contributions is often the only viable solution if firms
15
16 want to benefit from the ideas generated in firm-free communities. This is time-consuming
17
18 and expensive; online communities may contain several hundred thousand posts and
19
20 comments. The sheer amount of information in which the ideas are hidden is a practical
21
22 barrier to finding the ideas and utilising them for innovation (Lin, Hsieh, & Chuang, 2009;
23
24 Thorleuchter & Van den Poel, 2013).

25 26 27 28 *Automatic idea detection*

29
30
31 A new and efficient way of solving the needle-in-a-haystack problem is to use classifi-
32
33 cation algorithms that can screen arbitrary amounts of community posts and comments and
34
35 identify those that are likely to contain ideas. Using text analytics and machine learning meth-
36
37 ods, Christensen, Nørskov, Frederiksen, & Scholderer (2017) develop such an algorithm and
38
39 demonstrate its classification performance and efficiency for the case of extracting new prod-
40
41 uct ideas from an online community related to Lego. Christensen, Liland, et al. (2017) show
42
43 that the same principles can be applied to extract ideas for innovations from a community
44
45 related to craft brewing. The authors argue that their method is applicable across different
46
47 technological areas and product categories because most people use a specific set of words
48
49 and expressions when they communicate ideas to each other. That is, we humans have a very
50
51 special discourse for talking about our ideas and problems. We humans recognize ideas, when
52
53 we read and hear them, in the same manner as we recognize a car, when we see a car, and we
54
55
56
57
58
59
60

1
2
3 recognize a ship when we see a ship. Since the presence of such linguistic markers can easily
4
5 be detected in a given online community post or comment, it can also potentially be exploited
6
7 in the screening of arbitrarily large collections of posts, comments or other types of semi- or
8
9 unstructured text. If implemented as a screening tool in a company's R&D or marketing de-
10
11 partment, it can significantly reduce the labour costs that would arise if R&D staff were as-
12
13 signed to manual monitoring of community activity.
14
15

16 *Aims of the study*

17
18
19 We believe that the method introduced by Christensen et al. (2017) shows potential for
20
21 aiding firms in their search for innovative ideas and may thus serve as a tool for extending the
22
23 boundaries of inbound open innovation. Still, some questions must be addressed before such a
24
25 method should be implemented in a firm's innovation processes. Ideas identified by the Chris-
26
27 tensen et al. (2017) method, for example, have not yet been evaluated by company-internal
28
29 R&D or marketing staff and it remains to be investigated if ideas detected by such an auto-
30
31 mated system will also recognized as ideas by company staff. In addition, the ideas must be
32
33 seen as sufficiently novel, feasible and valuable by the R&D or marketing staff who would be
34
35 responsible to take the identified ideas further (e.g., development into concepts or prototypes).
36
37 The aim of the present paper is to fill these two gaps. Specifically, we would like to contribute
38
39 in two respects to the literature:
40
41
42
43

- 44 • Our first contribution is to assess whether ideas from an online community, identified by
45 an artificial intelligence system such as the one described by Christensen et al. (2017),
46 will also be perceived as ideas by company-internal staff.
47
48
- 49 • Our second contribution is to investigate if the ideas that are detected by the system will
50 also be perceived as *good* ideas by company-internal staff.
51
52
53
54
55
56
57

1
2
3 There is a reason why believe these issues are important. We address potential
4
5 acceptance problems that were also in the general innovation literature initially seen as
6
7 barriers for the uptake of user-contributed ideas by companies. Since then, many studies have
8
9 demonstrated that user-contributed ideas can often compete with the ideas generated by
10
11 company-internal staff (see e.g. Kristensson et al., 2004; Magnusson, 2009; Magnusson et al.,
12
13 2014) and therefore deserve to be given a fair chance. As a consequence, dedicated idea
14
15 crowdsourcing systems have gained widespread acceptance in the business community. Poetz
16
17 and Schreier (2012), for example, investigate if users ideas posted in a *firm-hosted, closed*
18
19 idea-crowdsourcing community can compete with ideas generated by company professionals.
20
21 Our study extends this question to the mode of automated idea-harvesting. The study we
22
23 report asks if user-contributed ideas posted in an *open* online community, identified by an
24
25 artificial intelligence-based system, can reach sufficient recognition among company
26
27 professionals. The answer to this question will provide guidance for research and practice in
28
29 the open innovation domain. If the answer is positive, artificial intelligence systems for idea-
30
31 harvesting can be implemented in practice, and research can focus on optimising the methods.
32
33 If the answer is negative, research can focus on unresolved problems and acceptance barriers
34
35 related to automatic idea detection systems.
36
37
38
39

40 An online community related to craft brewing was used as the idea base for our study.
41
42 A dataset consisting of 200 automatically extracted online community posts was generated for
43
44 addressing the two aims. It is the first time results based on this particular dataset are reported
45
46 in the literature. Employees of Norwegian craft brewery *Nøgne Ø* evaluated the automatically
47
48 extracted ideas.
49
50
51
52
53
54
55
56
57
58
59
60

Method

Machine learning for idea detection

The machine learning system we used for extracting the 200 texts is described in detail in Christensen, Nørskov et al. (2017) and Christensen, Liland, et al. (2017). Although the technical properties of the system are not the central focus of the present paper, we will give a brief description of the system and how it was employed in our study.

The machine learning system takes as input idea texts and non-idea texts that have been identified by human raters. The texts used for this study originate from *alt.beer.homebrewing*, a Usenet-based online community related to craft brewing. In this community people from all over the world discuss brewing-related issues. At the time the texts were extracted, the community contained altogether 10582 posts. 3000 of these were selected at random and extracted for the development of the training of the system (detailed results based on these 3000 texts have been reported in Christensen, Liland, et al. (2017)). Those that contained ideas were identified via crowdsourcing, using the *CrowdFlower* platform (a service similar to Amazon's *Mechanical Turk*). Five raters were assigned to each text and instructed to label the text as an idea text if it contained at least one idea.

Before the texts could be used for machine learning, several text pre-processing steps were performed. In this process, the raw text content was turned into a row-column format, where each text was represented as a row and each term (i.e., each unique word or expression) as a column. All numbers, punctuation marks and stop words were removed. Uni-grams, bi-grams and tri-grams were generated. All terms that did not occur in at least 0.2% of the texts were omitted from the analysis (this is a standard text cleaning step; e.g., see Antons, Klerer, & Salge (2016)). This process resulted in a dataset consisting of 10514 terms representing 10582 texts.

1
2
3 From the 3000 texts in the database, we excluded all texts where not all five
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

From the 3000 texts in the database, we excluded all texts where not all five
CrowdFlower raters had agreed on the class membership. After excluding these, the new
database contained 1393 texts. 405 of the texts were idea texts and 988 were non-idea texts.
The texts were partitioned at random into three separate data sets: a training set (consisting of
70% of the texts), a validation set (15% of the texts) and a hold-out or test set (15% of the
texts). Such a partition is essential for training a machine learning system (in the training set),
the fine-tuning of its paramters (in the validation set) and for an unbiased evaluation of its
performance on previously unseen data (hold-out). Based on the training set, validation set
and hold-out, the automatic idea detection system was trained and tested. The system is based
on a linear support vector machine classifier (for details, see Christensen, Liland, et al., 2017).
Key performance statistics are reported in Table 1.

--- Table 1 ---

From the remaining 7582 texts ($10582 - 3000 = 7582$) which had not been involved in
the training, validation and testing of the system, 200 texts were extracted for the present
study by using the linear support vector machine classifier: 100 which the classifier had
labelled as idea texts and 100 which the classifier had labelled as non-idea texts. A histogram
of the posterior probability scores underlying these classifications is shown in Figure 1. These
200 texts were used in the present study as the idea and non-idea texts to be classified and
rated by two brewing professionals.

--- Figure 1 ---

Measuring idea quality

The perceived quality of an idea can depend on the perspective of the person
evaluating the idea. This topic has received much attention in the creativity and innovation

1
2
3 management literature. In principle, idea quality can be measured on a “good idea” to “bad
4
5 idea” scale, but in most research it is decomposed into several attributes that represent
6
7 conceptually distinct dimensions of quality. Dean, Hender, Rodgers and Santanen (2006)
8
9 provide a comprehensive review of the idea quality literature published between 1990 and
10
11 2005. Based on the altogether 90 identified studies, they suggest that four dimensions of idea
12
13 quality can be distinguished: novelty, workability, relevance and specificity. An idea is novel
14
15 if it contains something that is new. An idea is workable if it is easy to implement and does
16
17 not violate known constraints. An idea is relevant if it satisfies pre-defined goals. An idea is
18
19 specific if it has been worked out in detail.
20
21
22

23 Comparable sets of sub-dimensions have been suggested in the user innovation
24
25 literature. Kristensson, Gustafsson and Archer (2004) compared the ideation performance of
26
27 ordinary users, expert users and professionals. They used three quality attributes: originality
28
29 (comparable to the novelty dimension suggested by Dean et al. (2006)), realisability
30
31 (comparable to the feasibility dimension) and value (comparable to the relevance dimension).
32
33 In a similar study, Magnusson (2009) compared the ideation performance of professionals,
34
35 technically skilled users, ordinary users, consulting users and creativity-trained ordinary users.
36
37 He used the quality attributes originality (comparable to novelty), producibility (comparable
38
39 to feasibility) and user-value (comparable to relevance). Using the same attributes,
40
41 Magnusson et al. (2014) compared technically skilled users with technically naïve users.
42
43 Poetz and Schreier (2012) compared the ideas of users and professionals in terms of the
44
45 attributes novelty, feasibility and customer benefit (comparable to value). Based on the four
46
47 studies that have a product-user ideation focus, we chose novelty, feasibility and value as the
48
49 quality attributes for our study.
50
51
52
53
54
55
56
57
58
59
60

Procedure

We established contact with Norwegian craft brewery *Nøgne Ø*. The brewery was founded in 2002 by two Norwegian home brewers and is nowadays part of Norwegian brewery group Hansa Borg Bryggerier. In 2015, *Nøgne Ø* produced 30 different styles of ales and exported to more than 40 markets. Two company professionals were recruited as expert raters. Expert 1 was 29 years old, female and had a business school background. Her responsibilities at *Nøgne Ø* were sales and logistics. At the time the study was conducted, she had been working for the brewery for 12 years. Expert 2 was 40 years old, male and had an engineering background. His responsibilities at *Nøgne Ø* were related to marketing and the web shop. At the time the study was conducted, he had been working for the brewery for 4.5 years.

The experts evaluated the 200 texts one-by-one and independently from each other. First, the experts were instructed to read the respective text carefully. Then, they were asked: *“Please evaluate if you think that the text contains one or more ideas”* and to respond on a binary *“yes”* versus *“no”* scale. If the expert responded *“yes”*, three rating scales were presented on which the expert was asked to evaluate the quality of the idea in terms of the three attributes novelty, feasibility and value. The scales were horizontally aligned ranging from very low (1) to very high (10). The instruction for the novelty attribute was: *“Please evaluate the novelty of the idea(s) in the text (by this we mean: to what degree does the idea suggest something new)”*. The instruction for the feasibility attribute was: *“Please evaluate the feasibility of the idea(s) in the text (by this we mean: to what degree is it possible to implement the idea)”*. The instruction for the value attribute was: *“Please evaluate the value of the idea(s) in the text (by this we mean: to what degree does the idea solve the underlying problem)”*.

Inter-rater reliability

To assess the inter-rater reliability of the idea/non-idea classification task, we calculated Cohen's kappa, normalised for differences between raters in their marginal distributions. The normalised version of kappa takes on values between 0 and 1 where a value of 0 stands for chance-level agreement and a value of 1 for the theoretical maximum of agreement, given the marginal distributions of the raters. Expert 1 identified 41 texts as containing ideas and 159 as not containing ideas. Expert 2 identified 87 texts as containing ideas and 113 as not containing ideas. They agreed on 35 texts as containing ideas and 107 as not containing ideas (See Table 2 for examples). These counts correspond to a normalised kappa of 0.74, suggesting that there was substantial agreement between the two experts as to whether a given text did or did not contain an idea (Cohen, 1960; Landis & Koch, 1977; von Eye & von Eye, 2008).

--- Table 2 ---

To assess the inter-rater reliability of the idea quality rating task, we calculated reliability measures based on generalisability theory (Cronbach, Gleser, Nanda, & Rajaratnam, 1972; Brennan, 2001). Only the 69 texts which the machine learning classifier had classified as an idea and which at least one of the brewery professionals had identified as an idea were included in the analysis. The design was a two-facet crossed design with tasks (the three quality attributes) and raters (the two brewery professionals) treated as fixed effects. The reliability (generalisability coefficient) of the averaged rating of a randomly picked idea text on the three attributes by the two raters was $E\rho^2 = 0.71$.

Results

Presence of ideas

Since our two company professionals had not perfectly agreed with each other on the presence or absence of ideas in the texts, we defined two validation criteria: a lenient criterion (Boolean OR: at least one professional had identified the respective text as containing an idea) and a strict criterion (Boolean AND: both professionals had identified the respective text as containing an idea).

Using the lenient criterion as a gold standard (where 47% of the 200 texts would be defined as true idea texts), the automatic idea detection system performed well. The classifier agreed with the company professionals in 77% of the cases as to whether a text did or did not contain an idea (accuracy). 75% of the texts which the classifier had identified as idea texts were also identified as idea texts by the company professionals (precision, also referred to as positive predictive value in the literature). The classifier correctly identified as idea texts 74% of the texts the professionals had identified as ideas (recall, also referred to as sensitivity or true positive rate in the literature). Since precision and recall always represent a trade-off, we also calculated their harmonic mean, the F_1 measure, as a compromise. Using the lenient criterion, it reached a very respectable value of $F_1 = 0.75$. For comparison see Christensen, Nørskov, et al. (2017) who obtained $F_1 = 0.54$, $F_1 = 0.55$ and $F_1 = 0.81$. Classification accuracy statistics are reported in Table 3.

Using the strict criterion as a gold standard (where only 18% of the 200 texts would be defined as containing ideas), the automatic idea classification system still agreed with the company professionals in 67% of the cases as to whether a text did or did not contain an idea (accuracy). Due to the much stricter criterion as to what defined an idea text, the precision of the classifier was lower: only 33% of the texts which the classifier had identified as idea texts

1
2
3 were also identified as idea texts by the company professionals. For the same reason, recall
4
5 was higher: the classifier correctly identified as idea texts 86% of the texts the professionals
6
7 had identified as ideas. The F_1 measure, as a compromise between precision and recall,
8
9 reached a value of 0.47.
10

11
12 Taken together, the criterion validity of the automatic idea detection system can be
13
14 regarded as satisfactory as long as it is used for the screening of potential ideas. Deployed in a
15
16 company as a tool for filtering out candidate ideas for product and process innovations, it may
17
18 significantly reduce the time and effort that would otherwise have to be spent by company
19
20 staff on manual screening and preliminary evaluation of a number of user contributions in
21
22 potentially relevant online fora.
23

24
25
26 --- Table 3 ---
27

28
29 *Quality of automatically detected ideas*
30

31
32 Figure 2 shows the distribution of the quality ratings of the ideas (i.e., those texts that
33
34 had been identified as ideas by the automatic idea detection system and which had been also
35
36 been identified as ideas by at least one of the two company professionals). For texts which
37
38 both company professionals had classified as an idea, the values on the novelty, feasibility
39
40 and value attributes are the averaged ratings of both company professionals. For texts which
41
42 only one of the company professionals had identified as an idea, the values are the ratings
43
44 given by that professional. The overall quality values were calculated as unweighted averages
45
46 of the ratings on the novelty, feasibility and value attributes.
47
48
49

50 --- Figure 2 ---
51

52
53 The distribution of the novelty ratings was concentrated in the lower range of the
54
55 response scale (which had a minimum of 1 and a maximum of 10), the distribution of the
56
57

1
2
3 feasibility ratings in the upper range of the response scale, and the distributions of the value
4
5 ratings and overall quality in the middle of the response scale. The results suggest that, on
6
7 average, the ideas which the automatic idea detection system extracted from the
8
9 *alt.beer.home-brewing* community appeared rather feasible to brewery professionals, were not
10
11 particularly novel, but had medium value and medium overall idea quality. Although the
12
13 results are generally consistent with the findings of Kristensson et al. (2004), Magnusson
14
15 (2009), Poetz and Schreier (2012) and Magnusson et al. (2014), the algorithmically identified
16
17 ideas in the present study were on average slightly less novel but also slightly more feasible
18
19 than the ideas generated by the human users who directly participated in the above-cited
20
21 studies.
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Discussion and conclusion

Implications for researchers and practitioners

The age of “big data” has generated opportunities and challenges for companies. In the present study, we focused on the case of data generated by users of social media and online communities, which can pose a challenge due the semi- or unstructured nature of such data (Olsen & Christensen, 2015). If people’s thoughts and ideas expressed on social media can be captured in a systematic manner, inbound open innovation processes can be accelerated (McAfee & Brynjolfsson, 2012; Jin, Wah, Cheng, & Wang, 2015) and thereby make companies more receptive (Ooms et al., 2015). The first aim of the present study was to investigate if ideas for product- and process innovations detected by an artificial intelligence system (in this case, the one developed by Christensen, Nørskov, et al. (2017) would also be regarded as ideas by company-internal staff who would be responsible for taking the ideas further in the innovation process.

Our results suggest that this is to a considerable extent the case: the performance of the system can be regarded as sufficient for an initial screening of potential ideas. Deployed in a company as a tool for selecting candidate ideas for product and process innovations, it can significantly reduce the time and effort that would otherwise have to be spent by company staff on wading through a large number of user contributions in potentially relevant online communities. The exact level of criterion-related validity that our system could achieve depended on several factors. The most important of these are (a) the definition of the “gold standard” against which the predictions are validated and (b) the cut-off used for transforming the continuous posterior probability score generated by the system into a binary prediction. In our analysis, we used two of the possible gold standards: a lenient criterion (at least one of the company professionals had rated the respective text as containing an idea) and a strict

1
2
3 criterion (both company professionals had rated the text as containing an idea). The lenient
4
5 criterion led to an implied base rate of 47% for the target event (i.e., the probability that a
6
7 randomly chosen text from among the 200 used in the present study would contain an idea),
8
9 whereas the strict criterion reduced the implied base rate to 18%. It is not possible to define
10
11 on purely statistical grounds what the right base rate should be. This is complicated by the fact
12
13 that the two company professionals who served as experts in our study did not have the same
14
15 base rates in their individual classifications: Expert 1 appeared to use a more conservative
16
17 standard of judgment, rating 21% of the 200 texts as containing ideas, whilst Expert 2
18
19 appeared to use a more liberal standard, rating 44% of the texts as containing ideas.
20
21
22

23 Whether it makes more sense for a given company to use a stricter or more lenient
24
25 criterion for further filtering of the automatically identified ideas may depend more on
26
27 strategy and available resources: a lenient criterion may be more appropriate if a company
28
29 wants to cast its net wide and thereby reduce the risk of missing certain ideas which might not
30
31 yet be able to achieve full cross-functional consensus. However, the company would also
32
33 have to be prepared to assign the necessary resources for dealing with the larger number of
34
35 ideas that would enter the innovation funnel. If, on the other hand, a company wants to limit
36
37 its resource expenditure and focus on ideas that can already in the early phases achieve cross-
38
39 functional consensus, a stricter criterion would be appropriate.
40
41
42

43 A similar objective can be achieved by tuning the cut-off value of the SVM classifier
44
45 underlying the Christensen et al. (2017) system. The algorithm yields a posterior probability
46
47 score that is continuous on the (0,1) interval. A traditional way of transforming the posterior
48
49 probability score into a binary classification rule is to use the value 0.50 as a cut-off such that
50
51 a text is classified as an idea text if the probability that the text contains an idea, given the
52
53 support vectors, is larger than 0.50, and classified as a non-idea text otherwise. However, the
54
55 traditional way of setting the cut-off may not always be the most useful way. Another
56
57
58
59
60

1
2
3 heuristic that is typically more useful is to set the cut-off equal to one minus the base rate of
4 the target even, either on the posterior probability scale or on the empirical percentile scale.
5 This heuristic would match the prior probability of classifying a text as an idea to the base rate
6 of the event. A third way of setting the cut-off is to estimate how many additional ideas a
7 company would be able to absorb into its innovation funnel and to use an appropriate absolute
8 cut-off, selecting the right number of ideas from the top of the posterior probability ranking.
9
10
11
12
13
14
15

16 The second aim of the present study was to investigate if the automatic idea detection
17 system developed by Christensen, Nørskov et al. (2017) would extract *good* ideas from the
18 online community that served as an example here. For the online community under
19 investigation, our answer is a qualified yes: the distribution of the overall idea quality score,
20 calculated as the average rating of each idea on the three quality attributes (novelty,
21 feasibility, value) by the two company professionals, was concentrated in the middle of the
22 response scale (mean = 4.8, 25th percentile = 3.8, 50th percentile = 5, 75th percentile = 5.7) and
23 ranged from a minimum of 1 (the lower end of the response scale) to a maximum of 8 (two
24 points below the maximum of the response scale). Overall, the ideas extracted by the
25 automatic detection system appear to have made a reasonable impression on the company
26 professionals.
27
28
29
30
31
32
33
34
35
36
37
38
39
40

41 Another important aspect for the evaluation of innovative ideas is their *timing*.
42 Although we did not explicitly focus on this aspect in the present study, many of the ideas
43 identified by our system could serve as good cases here. Take the idea about gluten-free beer
44 in Table 2 as an example. It was identified as an idea by the automatic idea detection system
45 and by both company professionals. Notably, the idea was posted to the community in 2005,
46 one year before the world's first gluten-free beer was launched ("New Grist" by *Lakefront*
47 *Brewery Inc.*, Milwaukee, WI, launched in 2006). An interesting avenue for future research
48 would be to follow up more systematically on how often, and with which lead time, user-
49
50
51
52
53
54
55
56
57
58
59
60

1
2
3 contributed ideas in online communities precede the development and market launch of
4
5 commercial products.
6

7
8 An interesting detail related to the quality evaluations is that the identified ideas
9
10 tended to be regarded as more feasible and valuable by our company professionals than they
11
12 were regarded as novel. This finding reflects results obtained by Kristensson et al. (2004).
13
14 However, as already observed, agreement between our experts was not perfect here either. As
15
16 an example, consider the text shown in Table 4: a community member suggests a new mead
17
18 recipe. Overall, the idea was rated as one of the best by the two company professionals.
19
20 Expert 1 assigned a rating of 2 on the novelty attribute, 7 on feasibility and 4 on value. Expert
21
22 2 rated it 9 on novelty, 9 on feasibility and 9 on value. In the additional, qualitative responses
23
24 we obtained from the two professionals, it became clear that Expert 1 evaluated the idea in
25
26 terms of its quality as an idea for process innovation whereas Expert 2 evaluated it in terms of
27
28 its quality as idea for product innovation. Different perspectives, either due to the functional
29
30 specialisation of our company professionals or due to their different levels of experience with
31
32 the product category, seem to have led to different standards of judgment.
33
34
35

36 --- Table 4 ---
37

38 39 *Limitations and suggestions for future research* 40

41
42 The results presented here are an evaluation of a particular automatic idea detection
43
44 system (the one developed by Christensen, Nørskov et al., 2017) to a particular case (the craft
45
46 brewing community *alt.beer.home-brewing*), evaluated from the point of view of two brewing
47
48 professionals connected to a particular craft brewing company (*Nøgne Ø*). Naturally, this
49
50 poses limits to the generalisability of our findings. The ideas detected by an automated system
51
52 can only be as good as the ideas voiced by the users in the online community under
53
54 investigation. Furthermore, the 200 texts we selected for evaluation were only a sample and
55
56
57

1
2
3 therefore unlikely to reflect the whole range of ideas discussed in the community. It is an open
4
5 question whether similar results will be achieved when automatic idea detection systems are
6
7 applied to other technology domains or product categories.
8

9
10 This question can only be answered by follow-up research. However, we do believe
11
12 that we have demonstrated the potential of automatic idea identification systems: they can be
13
14 a powerful technique for the harvesting and initial screening of user ideas from online fora
15
16 that do not conform, and are not limited to, the highly restrictive architecture and user basis of
17
18 dedicated crowdsourcing systems. We hope that studies such as ours can also also make a
19
20 contribution to a wider discussion: which business tasks of a more complex nature can
21
22 credibly be solved by artificial intelligence-based systems? We are convinced that the answer
23
24 does not only lie in what is technically possible but also in what is acceptable to the
25
26 prospective users of the information generated by such systems. More user evaluations of the
27
28 performance of artificial intelligence-based systems are needed.
29

30
31 The presented method adds a new channel for feeding ideas into the innovation
32
33 processes of firms, complementing company-hosted idea crowdsourcing communities such as
34
35 those studied by di Gangi et al. (2010) and Poetz and Schreier (2012). But can the method
36
37 completely substitute company-hosted idea crowdsourcing communities? In our opinion, the
38
39 answer to this question is a qualified “no”. A company-hosted crowdsourcing community
40
41 generates more than just ideas for the company. It also serves as an arena for cultivating
42
43 customer relations, enabling the company to interact directly with its most dedicated
44
45 customers. The method evaluated in this paper is strictly “one-way” and does not offer such
46
47 opportunities. The method can, however, dramatically reduce the costs of crowdsourcing new
48
49 ideas, which is particularly relevant for companies who do not enjoy such large and loyal
50
51 customer bases as Dell (di Gangi et al., 2010) or Lego (Antorini et al., 2012). This we see as a
52
53
54
55
56
57
58
59
60

1
2
3 key advantage of the method, and it may even level the competitive playing field between
4
5 large companies and SMEs.
6
7
8
9

11 **References**

- 12
13
14 Antons, D., Klerer, R., & Salge, T. O. (2016). Mapping the Topic Landscape of *JPIM*, 1984-
15
16 2013: In Search of Hidden Structures and Development Trajectories: Mapping the
17
18 Topic Landscape of *JPIM*, 1984-2013. *Journal of Product Innovation Management*,
19
20 33(6), 726–749.
21
22
23 Antorini, Y. M. (2007). *Brand Community Innovation: An Intrinsic Case Study of the Adult*
24
25 *Fans of LEGO Community*. Copenhagen Business School, Frederiksberg: Center for
26
27 Europaforskning.
28
29 Antorini, Y. M., Muñiz, J., Albert M., & Askildsen, T. (2012). Collaborating With Customer
30
31 Communities: Lessons from the Lego Group. *MIT Sloan Management Review*, 53(3),
32
33 73–95.
34
35
36 Boyd, D., & Crawford, K. (2012). Critical Questions for Big Data: Provocations for a cultural,
37
38 technological, and scholarly phenomenon. *Information, Communication & Society*,
39
40 15(5), 662–679.
41
42 Brennan, R. L. (2001). *Generalizability theory*. New York: Springer-Verlag.
43
44 Christensen, K., Liland, K. H., Kvall, K., Risvik, E., Biancolillo, A., Scholderer, J., ... Næs,
45
46 T. (2017). Mining online community data: The nature of ideas in online communities.
47
48 *Food Quality and Preference*, 62, 246–256.
49
50
51 Christensen, K., Nørskov, S., Frederiksen, L., & Scholderer, J. (2017). In search of new
52
53 product ideas: Identifying ideas in online communities by machine and text mining.
54
55 *Creativity and Innovation Management*, 26(1), 17–30.
56
57
58
59
60

- 1
2
3 Cohen, J. (1960). A Coefficient of Agreement for Nominal Scales. *Educational and*
4
5 *Psychological Measurement*, 20(1), 37–46.
6
7 Cronbach, L. J., Gleser, G. C., Nanda, H., & Rajaratnam, N. (1972). *The dependability of*
8
9 *behavioral measurements*. New York: Wiley.
10
11 Dean, D. L., Hender, J. M., Rodgers, T. L., & Santanen, E. L. (2006). Identifying Quality,
12
13 Novel, and Creative Ideas: Constructs and Scales for Idea Evaluation. *Journal of the*
14
15 *Association for Information Systems*, 7(1), 646–698.
16
17 di Gangi, P. M., Wasko, M. M., & Hooker, R. E. (2010). Getting Customers' Ideas To Work
18
19 For You: Learning from Dell How To Succeed With Online User Innovation
20
21 Communities. *MIS Quarterly Executive*, 9(4), 213–228.
22
23 Ekvall, G. (1997). Organizational conditions and levels of creativity. *Creativity and*
24
25 *Innovation Management*, 6(4), 11.
26
27 Frederiksen, M. H., & Knudsen, M. P. (2017). From Creative Ideas to Innovation
28
29 Performance: The Role of Assessment Criteria. *Creativity and Innovation*
30
31 *Management*, 26(1), 60–74.
32
33 Füller, J., Bartl, M., Ernst, H., & Mühlbacher, H. (2006). Community based innovation: How
34
35 to integrate members of virtual communities into new product development.
36
37 *Electronic Commerce Research*, 6(1), 57–73.
38
39 Füller, J., Jawecki, G., & Mühlbacher, H. (2007). Innovation creation by online basketball
40
41 communities. *Journal of Business Research*, 60(1), 60–71.
42
43 Gandomi, A., & Haider, M. (2015). Beyond the hype: Big data concepts, methods, and
44
45 analytics. *International Journal of Information Management*, 35(2), 137–144.
46
47 Hsinchun Chen, Chiang, R. H. L., & Storey, V. C. (2012). Business intelligence and analytics:
48
49 From big data to big impact. *MIS Quarterly*, 36(4), 1165–1188.
50
51
52
53
54
55
56
57
58
59
60

- 1
2
3 Jeppesen, L. B., & Frederiksen, L. (2006). Why Do Users Contribute to Firm-Hosted User
4
5 Communities? The Case of Computer-Controlled Music Instruments. *Organization*
6
7 *Science*, 17(1), 45–63.
8
- 9 Jin, X., Wah, B. W., Cheng, X., & Wang, Y. (2015). Significance and Challenges of Big Data
10
11 Research. *Big Data Research*, 2(2), 59–64.
12
- 13 Kristensson, P., Gustafsson, A., & Archer, T. (2004). Harnessing the Creative Potential
14
15 among Users*. *Journal of Product Innovation Management*, 21(1), 4–14.
16
17
- 18 Landis, J. R., & Koch, G. G. (1977). The Measurement of Observer Agreement for
19
20 Categorical Data. *Biometrics*, 33(1), 159.
21
- 22 Lin, F.-R., Hsieh, L.-S., & Chuang, F.-T. (2009). Discovering genres of online discussion
23
24 threads via text mining. *Computers & Education*, 52(2), 481–495.
25
- 26 Magnusson, P. R. (2009). Exploring the Contributions of Involving Ordinary Users in
27
28 Ideation of Technology-Based Services*. *Journal of Product Innovation Management*,
29
30 26(5), 578–593.
31
- 32 Magnusson, P. R., Wästlund, E., & Netz, J. (2014). Exploring Users' Appropriateness as a
33
34 Proxy for Experts When Screening New Product/Service Ideas*. *Journal of Product*
35
36 *Innovation Management*, 33(1), 4–18.
37
38
- 39 Mahr, D., & Lievens, A. (2012). Virtual lead user communities: Drivers of knowledge
40
41 creation for innovation. *Research Policy*, 41(1), 167–177.
42
43
- 44 Majchrzak, A., & Malhotra, A. (2013). Towards an information systems perspective and
45
46 research agenda on crowdsourcing for innovation. *The Journal of Strategic*
47
48 *Information Systems*, 22(4), 257–268.
49
- 50 McAfee, A., & Brynjolfsson, E. (2012). Big data. The management revolution. *Harvard*
51
52 *Business Review*, (October), 3–9.
53
54
55
56
57
58
59
60

- 1
2
3 Nørskov, S., Antorini, Y. M., & Jensen, M. B. (2015). Innovative brand community members
4 and their willingness to share ideas with companies. *International Journal of*
5 *Innovation Management*.
6
7
8
9 Nunan, D., & Di Domenico, M. (2013). Market research and the ethics of big data.
10 *International Journal of Market Research*, 55(4), 2–13.
11
12
13 Olsen, N. V., & Christensen, K. (2015). Social media, new digital technologies and their
14 potential application in sensory and consumer research. *Current Opinion in Food*
15 *Science*, 3, 23–26.
16
17
18
19
20 Ooms, W., Bell, J., & Kok, R. A. W. (2015). Use of Social Media in Inbound Open
21 Innovation: Building Capabilities for Absorptive Capacity: Use of Social Media in
22 Inbound Open Innovation. *Creativity and Innovation Management*, 24(1), 136–150.
23
24
25
26 Poetz, M. K., & Schreier, M. (2012). The value of Crowdsourcing: Can Users Really
27 Compete with Professionals in Generating New Product Ideas?*. *Journal of Product*
28 *Innovation Management*, 29(2), 245–256.
29
30
31
32
33 Thorleuchter, D., & Van den Poel, D. (2013). Web mining based extraction of problem
34 solution ideas. *Expert Systems with Applications*, 40(10), 3961–3969.
35
36
37
38 Van de Ven, A. (1986). Central Problems in the Management of Innovation. *Management*
39 *Science*, 32(5), 590–607.
40
41
42 van den Ende, J., Frederiksen, L., & Prencipe, A. (2015). The Front End of Innovation:
43 Organizing Search for Ideas. *Journal of Product Innovation Management*, 32(4), 482–
44 487.
45
46
47
48 Vandebosch, B., Saatcioglu, A., & Fay, S. (2006). Idea management: a systemic view.
49 *Journal of Management Studies*, 43(2), 259–288.
50
51
52
53 von Eye, A., & von Eye, M. (2008). On the marginal dependency of Cohen's κ . *European*
54 *Psychologist*, 13(4), 305–315.
55
56
57
58
59
60

- 1
2
3 von Hippel, E., Ogawa, S., & PJ de Jong, J. (2011). The age of the consumer-innovator.
4
5 von Krogh, G., Spaeth, S., & Lakhani, K. R. (2003). Community, joining, and specialization
6
7 in open source software innovation: a case study. *Research Policy*, 32(7), 1217–1241.
8
9 von Krogh, G., & von Hippel, E. (2006). The Promise of Research on Open Source Software.
10
11 *Management Science*, 52(7), 975–983.
12
13 Zwitter, A. (2014). Big Data ethics. *Big Data & Society*, 1(2).
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

For Peer Review

Table 1 - Performance of the automatic idea detection system

Partition	True positives (TP)	True negatives (TN)	False positives (FP)	False negatives (FN)	Classification accuracy	Precision	Recall	F_1
Validation set	27%	70%	1%	2%	0.97	0.97	0.92	0.94
Hold-out set	25%	70%	1%	3%	0.96	0.96	0.88	0.92

For Peer Review

1
2
3 *Table 2 - Example of an idea text and a non-idea text on which both raters agreed*

4

Idea text	Non-idea text
'Buckwheat has been used as an adjunct for a long time in a few beers. It also is used to make gluten free beers. It has a high gelatinisation temp so need to be boiled first. Extract potential is about 1.032. Can be used lightly roasted to add colour to gluten free beers, or use Kasha (a roasted buchwheat). I think Rogues make a buckwheat ale'	'Thanks for the help. My internet is screwy or I would have replied sooner. I re- pitched and it is going crazy. A load off my mind! now i can concentrate on getting another cider and a wit going. Anyone have any suggestions for a good belgian style ale like duvel? I am an extract with specialty grains level brewer, so whole grain is out for now. Thanks again for all the help!'

5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Table 3 - Presence of ideas: classification accuracy of the automatic idea detection system, validated against the judgments of two company professionals

Validation criterion	True positives (TP)	True negatives (TN)	False positives (FP)	False negatives (FN)	Classification accuracy	Precision	Recall	F_1
Lenient criterion:								
Classified as idea by Expert 1 OR Expert 2	35%	42%	12%	12%	0.77	0.75	0.74	0.75
Strict criterion:								
Classified as idea by Expert 1 AND Expert 2	15%	52%	31%	3%	0.67	0.33	0.86	0.47

1
2
3 *Table 4 - Idea text identified by classifier, Expert 1 and Expert 2*

4 'I've made several batches. Below is my recipe The love of my life I love Mead as you can probably
5 tell. Please note, this is Mead but I do not use any water. I use apple juice as the base. You can use
6 water but I find the apple juice makes it a bit nicer for those of you who love apples and like a high
7 alcohol content. No citric acid needed. This is called Apple Honey Melonomel Meade You will need...
8 1 Package Red Star wine yeast 4 Gallons apple juice from concentrate 2-5 pounds of pure honey, the
9 more the better. This shit is expensive though. 1 cup table sugar 5 Fuji apples Siphon hose, any small
10 tube will work. A 5 gallon carboy or tub 1 balloon Step one, crush your apples or use a blender. Step
11 two, boil apples in large pot with apple juice. Step three, set aside to cool Step Four, boil honey in
12 large pot of apple juice Step five, set aside to cool. Step six, dump mixture into large 5 gallon carboy
13 and add activated yeast. Step six, allow the mead to ferment for 3-4 weeks, once fermentation
14 begins to slow prime with table sugar by diluting the 1 cup of table sugar in 1/2 gallon of apple juice
15 then pour this directly into the carboy. A balloon can be placed over the mouth of the carboy to
16 monitor the fermentation. Simply pierce a small hole in the balloon to allow CO2 to escape. Once the
17 Meade has cleared (meaning you can read a newspaper through it) transfer it into a secondary (Save
18 the sediment for use as the Yeast in your next batch of Meade) and let it clarify for 2-3 weeks. After
19 this bottle the meade and let fermentation finish off. Total process about 70 days and its ready to
20 drink. This will burn going down but is smooth as a whistle. Enjoy...'
21
22




Figure 1 - Histogram of the posterior probability scores generated by the SVM-based automatic idea detection system for the 200 texts used in the present study

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60




Figure 2 - Box plots of the distribution of quality ratings (overall quality = unweighted average of novelty, feasibility and value; diamonds represent 95% confidence intervals around distribution means)