

Understanding the Dynamics of Crowdfunding: Kickstarter Edits

CS224N Final Project

Viswajith Venugopal
viswa@stanford.edu

Sameep Bagadia
sameepb@stanford.edu

Abstract

Crowdfunding is now well-established as a platform for independent artists and entrepreneurs to reach out to the public to finance their projects. There is a fair amount of existing work which seeks to answer the question: how do you convince hundreds of people to contribute hundreds of thousands of dollars? [1][2] Such attempts have been fairly successful, and have been able to predict the success or failure of projects with accuracies from 75%-90%, using both linguistic and metadata-based features. The problem is, such work has typically used one *static* snapshot of the project for prediction – and worse, a snapshot scraped *after* the end of the project. However, such work does not examine the dynamics of the fund-raising process, and ignores the fact that projects are often edited during fund-raising. In this work, we analyse daily snapshots of close to 20,000 projects from the most popular web-based crowdfunding platform, Kickstarter, with the specific view of linguistically characterizing the edits, and seeing if we can predict the impact of the edit, or the ultimate success of the project, by linguistic analysis of its edits.

1 Introduction

In recent years, crowd-funding has proved to be an extremely effective means for independent artists and entrepreneurs to reach out to wealthy donors and finance their projects. Incredible amounts of money are now raised by projects that are successful in appealing to the public, which naturally leads to the question: can we figure out what makes an attempt at crowd-funding successful or unsuccessful? As we proceed to describe in this

section, there is a body of work that seeks to address this question; however, all existing work that attacks this question from a linguistic angle does not take into account the dynamics of the project, or the effect of edits. Our goal in this project is to understand editing behaviour, and to characterize edits linguistically.

1.1 What is Kickstarter?

Kickstarter¹ is arguably the most popular web-based crowd-funding platform. Independent artists and entrepreneurs create projects on the platform. Each project has its own page (an example is shown in Figure 1), which contains the project title and a description (text-based, but most projects include a video as well). Each project has a *funding goal*, a target amount that the project creators seek to raise within a specified *project duration*.




Figure 1: An Example of a Kickstarter project page

A distinguishing feature of Kickstarter (which makes it particularly interesting for research of our kind) is the fact that it follows an *all-or-nothing* model, distinguishing sharply between success and failure: if a project reaches its target, it gets all of the pledged money, but if it falls short by even a dollar, the creators get nothing.

¹<http://www.kickstarter.com>

1.2 Related Work

A whole host of research exists on crowdfunding, analysing how and why it works from various angles.[3][4] See [2] for a detailed summary of work in this direction.

More relevant to our specific research question, there is a significant amount of past work in the domain of Kickstarter projects, which try to predict the success of projects using information from the project pages, and from the social media activity of the project.[2][1] It is fairly well-established that metadata-based features such as the project category, location, target funds, social media activity and presence of a video are good predictors of project success. Prediction accuracy is further boosted when linguistic analysis of the project page is done. The authors of [1] find that, in general, the most successful projects are emotive, colloquial, indirect, descriptive, and framed as collective efforts. The authors of [2] find several commonly used phrases that characterize successful projects, and try to linguistically characterize them, in terms of reciprocity, authority, social identity and several other aspects. The models of [2] and [1], among others, work fairly well, with prediction accuracies ranging from 75% to 90%.

There is also some work in understanding the dynamics of crowdfunding projects [5][6]. The authors of [5] perform time-series analysis on the project reward of Kickstarter and comes up with interesting characterizations.

1.3 Our Contribution

Our goal is, in some sense, to address a gap in the existing work that tries to predict the success of Kickstarter projects: these projects don't account for the fact that project creators often edit their descriptions. (In fact, since they try to predict project success based on features scraped *after* the end of the project, there are concerns that the accuracy of their models is inflated.) Thus, we seek to incorporate these ideas into a study of the dynamics of the fundraising process, by looking at editing behaviour on Kickstarter, and trying to linguistically characterize edits. Our primary interest is in the linguistic insights that we can derive, so we refrain from using other metadata-based predictors in our models.

In section 2, we give an overview of the data we used. Then we do some preliminary analysis on the editing behavior of project descriptions in section 3. We analyze how often do people edit, when do they edit during the project time line, how significant are the edits that they make and what is the impact of edits. In section 4, we explain the linguistic features that we use to analyze the edits. Then, we go through the various experiments we performed in section 5. We performed a linguistic analysis of the edits to answer the questions: What kind of edits are made and Which edits have most impact? We also used supervised machine learning methods to solve the task of predicting success of projects as well as gain from the linguistic features of an edit. We then summarize our results and give some ideas of possible future work in section 6.

2 An Overview of our Data-set

Our data-set² consists of daily HTML snapshots of over 20,000 projects on Kickstarter. Initially, we filtered the data, removing some projects for which we didn't have snapshots for the entire duration, and some other projects which were suspended or canceled in the middle. This left us with a total of 19,299 projects. Out of these, 6,998 projects have at least one edit in them, which make them interesting for us. Most projects are around 30 days long, although there is a fair amount of deviation.

Table 1: Summary of the projects in our data set, with the total count, and the counts of successful and failed projects for edited and unedited projects.

	Total	Succ	Fail	Frac
All	19299	5863	13436	0.304
Edited	6998	3469	3529	0.495
Unedited	12301	2394	9907	0.195

Table 1 presents the summary of the projects in our data-set. As we can see, two-thirds of projects don't have a single edit, but the fraction of successful projects is significantly higher among edited projects: we conclude that this is because unedited projects are more likely to be

²We would like to thank Rob Voigt, PhD student in Computational Linguistics, for collecting this data.

‘abandoned’ projects, whose creators weren’t that serious about the fund-raising. It can be seen that most successful projects are likely to be ‘active’, and have at least one edit. This means that zooming in to just the edited projects still leaves meaningful problems: in particular, note that almost exactly half the edited projects are successful, which means the problem of trying to predict project success from edits is interesting.

3 Understanding Edits

In this section, we focus on projects that have at least one edit, and try to glean a more nuanced understanding of editing behavior on Kickstarter by answering some basic questions.

3.1 How Often Do People Edit?

A natural question to ask is how frequently projects are edited. Figure 2 presents the histogram for number of projects against the percentage of days in which their descriptions are edited. It can be seen that most projects are edited only a handful times, usually once or twice. There are few projects which have been edited very frequently.




Figure 2: Frequency of edits on Kickstarter

3.2 Time line of Edits

Next, we find out how long into the project edits are made, by looking at the edit time line. As Figure 3 shows, a large number of edits are made in the first few days. This is likely because the project creator is still setting up his page and experimenting with different descriptions, trying to add more details or improve the layout. As the project progresses, the number of edits shows a

decreasing trend and then increases again towards the end of the project. (Although there are some spikes in the days in the middle.) The edits near the end could be due to projects reaching their goal, and the creators updating the description to thank their backers, or because of projects being near completion, and the creators petitioning potential pledgers to give what they need to reach their target.




Figure 3: Editing timeline

3.3 Impact of edits

An important question is: does editing help projects or not? Figure 4 presents a histogram of the number of projects against the fraction of the target (1 or greater means success) raised for both edited and non-edited projects. We see that edited projects tend to achieve a higher fraction of the goal as compared to non-edited projects.




Figure 4: Impact of edit

Overall, we also see that the number of projects

that hardly receive any money are very high and this number drops down as we go further and it spikes up at 1. This can be because if projects raise enough money to go near completion, project creators would put in more effort and publicize it more to complete it. It might also be possible for project creators themselves to chip in the remaining amount since if the goal is not reached they would receive nothing.

3.4 Significant Edits

Our findings so far indicate that edited projects are more likely to be successful. Now, we zoom in on the edits, and try to look at which edits are significant, based on textual analysis alone. We calculate edit significance based on the difference between the project descriptions: the more different the descriptions are, the higher the ‘significance’ of our edit. Difference is calculated using Python’s `difflib`, which uses the following metric:

$$\text{Difference} = 1 - \frac{2 \times M}{T}$$

where M is the number of matches between the two sequences and T is the total number of elements in both the sequences. Thus, the difference value is 0 if both the texts are exactly the same, and 1 if they’re completely different.

Figure 5 shows the fraction of successful projects against the edit significance threshold. (That is, the fraction of successful projects among projects that have at least one edit with a difference value higher than the threshold.) We see that projects with more significant edits have a lower success fraction. Thus, counter-intuitively, even though editing boosts project success, the fraction of successful projects goes down as the significance of the edits we look at increases.

3.5 Gain due to edits

In order to study the impact of an individual edit, we define edit gain as the ratio of average daily pledge money received after the edit to the average daily pledge money received before it. More concretely, edit gain due to an edit at day i of a project which runs for a total of n days is given by:

$$\text{Gain}_i = \frac{(\text{raised}[i-1] - \text{raised}[0]) / (i-1)}{(\text{raised}[n] - \text{raised}[i+1]) / (n-i-1)}$$




Figure 5: Result of edit significance

where $\text{raised}[i]$ denotes the money raised by the i^{th} day.

One possible issue with this metric is that multiple edits in the same project can interfere with each other. But for our purposes, we ignore that case: this is good to a first approximation, since, as we saw earlier in Figure 2, most projects are edited only once or twice as we saw.




Figure 6: Edit gain histogram

Figure 6 shows the histogram of the number of edits against their gain value. We can see that, by our metric, edits on average tend to have a negative impact on the project. But there are edits which have huge positive impact too. The average gain is very close to 1.

4 Linguistic Modeling of Project Descriptions

In order to understand and analyze the edits in more detail, we carry out linguistic analysis of project descriptions with respect to their edits to gain more insight on what kind of edits are made and how useful they are. In this section, we explain the various linguistic features that we studied.

4.1 N-grams

Previous work in this domain has shown that n-grams are good features in models that aim to predict project success. To decide which n-grams to throw in as features, we enumerated all uni-, bi- and trigrams seen in our project descriptions, and counted their frequencies. We retained the 5,000 most common n-grams, and incorporated them into our linguistic models. For each edit, we use the *difference* in frequency of each n-gram as a feature to describe the edit.

4.2 LIWC Features

We use the Linguistic Inquiry and Word Count (LIWC) [7] corpus which has a bag of words for 64 linguistic categories, including Affect, Positive and Negative Emotions, Social, Insight and several other categories. For each edit, we look at the number of words in each category that are added or removed due to the edit, so as to gain qualitative insight into the nature of the edits, and the effect of various linguistic categories on edit impact.

4.3 Concreteness

We consider the concreteness or abstractness of the project descriptions and analyze based on this criteria. We use the concreteness score for various words as obtained from the MRC Psycholinguistic Database [8]. We use this database to create a score of concreteness for each project description by calculating an average score of the database words present in the description. As before, we use the difference in concreteness value before and after the edit as a predictor in our models.

4.4 Sentiment

We calculate the sentiment of the project description in order to analyze the how positive or negative sentiment is used and what impact does it have. In order to calculate the sentiment value,

we use TextBlob [9] library for Python. Again, we calculate the sentiment change in an edit by simply taking the difference in values before and after.

5 Experiments Run

5.1 A Linguistic Analysis Of Edits

We performed an analysis of the linguistic features described in Section 4. We wish to answer two broad questions. What kinds of edits are generally made? What linguistic features have the greatest impact on whether the edit helps the project, and on the ultimate success of the project?

5.1.1 What kind of edits are made?

Looking closely at edits, we see that, on average, content is added to the project descriptions. There is an average increase of 25% to the project description length across all edits. Therefore, we also see, on average, an increase in counts of words in all the LIWC categories. Instead of studying absolute change, therefore, we study the relative change in LIWC category words to account for the fact that some category words would naturally occur more and some would occur less.




Figure 7: Relative change in LIWC category words

Figure 7 shows the relative change in LIWC category words for the five most and least changed categories. We see that all the values are positive which means that for each of the categories we see an increase in words of that category. Categories like ‘You’, ‘I’, ‘We’ are among the most changed categories. This shows that edits have a higher addition of words of first and second

person personal pronouns. We also see significant increase in functional and exclusive words. On the other hand, extreme categories like swear words, death, anxiety have a very low increase. We notice that words related to religion are also very uncommon additions in edits.

When we look at how edits change concreteness, we notice that there are significant numbers of edits of both kinds: that is, those that increase concreteness and those that decrease it (and make the description more ‘abstract’). But, on average, there is a decrease in the concreteness of project description due to the edits. Similarly, when we look at the sentiment of project descriptions, there are significant amount of edits that increase and decrease the sentiment. But, on average, we see that the sentiment value becomes more positive due to edits.




Figure 8: Best bigrams and trigrams

Figure 8 shows the most changed bi-grams and tri-grams in the project descriptions. We see a lot of expected phrases involving words like role, support, donations, use etc. to convince people to donate for their project.

5.1.2 Which edits have most impact?

In order to study the impact of an edit, we consider the gain of an edit as defined in section 3.5. We divide the edits into high gain and high loss edits. An edit with a gain greater than 2 is a high gain edit while one with a gain of less than 0.5 is a high loss edit. We then analyzed our linguistic features with respect to high gain and high loss edits.

Figure 9 shows the relative change of project description length for high gain, high loss and all the edits. We see that high loss edits have significantly higher relative change in description length




Figure 9: Relative change in length of project description

as compared to high gain edits. High gain edits increase the length by about 10% whereas high loss edits increase the length by more than 50% in average. This suggests that even though editing is good, the edit should not be adding too much of content as that would actually lead to decrease in money generated. This corroborates our earlier findings that projects with more significant edits actually did worse.




Figure 10: Percent change in sentiment of project description

Figure 10 shows the percentage change in sentiment for high gain, high loss and all projects. We see that, in all of them, there is a net increase in sentiment value. But the increase is higher for high loss projects as compared to high gain projects. Thus, good edits increase the sentiment value, but don't overdo it.

Figure 11 shows the relative change in LIWC features for the five most changed and the five least changed categories for both high gain and high loss edits. Overall, we see that change values are higher for high loss edits as compared to high gain. This corroborates the description length result seen in Figure 9.

We see that the most changed categories are quite different for the two cases except for category 'You'. Personal pronoun categories




Figure 11: Most changed and least changed LIWC categories for high gain and high loss edits

like ‘You’, ‘I’, ‘SheHe’ and other categories like ‘Excl’ and ‘Percept’ are among the most changed in high gain edits. Categories like ‘Funct’, ‘You’, ‘CogMech’, ‘Social’, ‘Relativ’ have the highest change for high loss edits.

Similarly, if we look at the least changed categories, there is some difference in the categories there. Categories ‘Sad’, ‘Nonflu’, ‘Assent’ are the least changed in high gain edits, whereas categories like ‘Friends’, ‘Relig’, ‘Anx’ are the least changed in high loss edits. Categories like ‘Death’ and ‘Swear’ are least changed in both.

5.2 Building Classifiers to Predict Edit Impact and Project Success

Next, we set up a supervised learning problem with our data: can we use linguistic analysis of edits to predict the success of the project, or to predict whether the edit will result in a high or low gain? We split our data into train and test, and analyse the performance of our classifiers on the test data. Note that we only use the linguistic features described in Section 4 to generate features for each edit. We refrain from using meta-data features since our interest is in the impact of linguistic features of edits. (Further, the meta-data features characterize the project, and not the edit.)

5.2.1 Predicting project success

In order to predict project success, we label each edit with the final success of the project, and train our classifiers on this two-class classification problem.

Table 2: Prediction Results: Success due to edits

Data	Clf	0 F1	1 F1	Avg F1	Acc
AD	Adb	0.434	0.739	0.627	0.643
AD	LR	0.371	0.692	0.574	0.586
AD	RF	0.407	0.739	0.617	0.638
MR	Adb	0.349	0.755	0.619	0.644
HG	Adb	0.558	0.633	0.597	0.599
HD	Adb	0.724	0.476	0.631	0.639
LD	Adb	0.421	0.746	0.630	0.647
LG	Adb	0.160	0.851	0.713	0.747

Table 2 shows the results of our classifiers on our test data. (Note that, due to space constraints, we report only the performance of whichever classifier worked best.) Various data that we have used are all data (AD), Mid range edits (MR), which are the edits that occur in the middle of the project time line (i.e. in the range of (0.3, 0.7) of the project duration), high gain (HG) edits, which are edits that cause a substantial increase or decrease in the gain value, (i.e. $\text{gain} < 0.35$ or $\text{gain} > 3$), high difference (HD) edits, which are substantial edits (i.e. $\text{difference} > 0.5$), low difference (LD) edits, which are smaller edits (i.e. $\text{difference} < 0.3$) and low gain (LG) edits which are edits with smaller impact on the gain value (i.e. $\text{gain} \in [0.75, 1.25]$). The various classifiers we used are Adaboost (Adb), Logistic Regression (LR) and Random Forests (RF). We consider the F1 score for both class 0 (0 F1) and class 1 (1 F1) and weighted average of the two is given by Avg F1.

We see that for all edits data, we can predict the success of the project with Avg F1 score 0.627 using Adaboost which gives prediction accuracy of 64%. We see slightly higher prediction accuracies in some cases when we restrict the data to various conditions. These results aren’t incredibly high, but they’re similar to what existing models can do without the metadata features. In fact, it is actually noteworthy that we can predict the success of projects this well using just the linguistic features of edits. This shows that edits

are actually important determiners of the success of projects.

To figure out which features were most important, we took a look at the importance weights of each feature as output by our classifier. Change in concreteness value was the best feature for this prediction. Other important features were change in sentiment value, LIWC categories ‘I’, ‘See’ and 1-grams ‘fun’, ‘unlocked’, ‘add’ etc.

5.2.2 Predicting gain or loss from an edit

In order to predict gain or loss from the edit, we label each edit with 1 if the gain, as defined in section 3.5, is greater than 1, and 0 otherwise.

Table 3: Prediction Results: Gain due to edits

Data	Clf	0 F1	1 F1	Avg F1	Acc
AD	Adb	0.652	0.364	0.532	0.550
AD	RF	0.700	0.230	0.504	0.568
AD	LR	0.655	0.339	0.523	0.546
MR	LR	0.585	0.421	0.510	0.516
HG	Adb	0.761	0.271	0.609	0.640
HD	RF	0.731	0.361	0.599	0.621
LD	Adb	0.652	0.369	0.535	0.552
LG	Adb	0.609	0.358	0.496	0.514

Table 3 shows the results of our classifier on the test data set. The data and classifier names used are same as those used in Table 2.

The prediction accuracies for gain are much lower as compared to predicting success when we consider all the edits. But, when we put some restrictions on the kind of edits we would look at, in some cases, we do see some improved results. One other thing to note here is that our classifiers are much better at class 0 F1 score rather than class 1 F1 score when predicting gain whereas it was the opposite in the case of predicting success.

Again, an examination of the importance weights shows us that the most important feature for predicting gain was the LIWC category feature of ‘Verbs’. Some other good features are the 1-gram ‘that’ (which is interesting) and LIWC category feature ‘Relativ’.

6 Conclusion and Future Work

We analyzed the editing behavior in Kickstarter project descriptions, with the aim of characteriz-

ing what kind of edits people make, when do they make them and what is the impact of edits. We see that many projects have at least one edit and success probability of edited project is far greater than non-edited ones. Small amount of edits are good but larger edits are harmful.

We studied the linguistic features like change in N-grams, LIWC, concreteness and sentiment values. We analyzed the edits to find out what do edits generally consist of. We also divided the edits into high gain and high loss edits to understand the impact of edits and the kind of changes that lead to high gain and high loss in the projects. We used supervised machine learning methods to create classifiers to predict success and gain due to edits. Even though the performance of the classifiers weren’t very good by using just linguistic features, they give an insight into the problem and also give some important features.

One thing which we haven’t done rigorously is to study the time series of the projects with the amount of money raised over the project time line and how is it different for projects that succeed or don’t succeed. This time series analysis would be useful to analyze the editing behavior and its impact in a better manner since they are inter-linked. One can also add meta-data features like project topic to the classifier and experiment with that.

Another direction to work on would be to analyze more linguistic features like politeness etc. One could also analyze the edits made in rewards and see the impact of that on the project fund-raising. Understanding the reason behind the project creator editing the project might also be another interesting area to look into.

References

- [1] D. J. Rob Voigt, Christopher Manning, “How to convince the crowd: The language of persuasion in crowdfunding.”
- [2] T. Mitra and E. Gilbert, “The language that gets people to give: Phrases that predict success on kickstarter,” in *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*. ACM, 2014, pp. 49–61.

- [3] E. M. Gerber, J. S. Hui, and P.-Y. Kuo, "Crowdfunding: Why people are motivated to post and fund projects on crowdfunding platforms."
- [4] R. P. Fisk, L. Patrício, A. Ordanini, L. Miceli, M. Pizzetti, and A. Parasuraman, "Crowdfunding: transforming customers into investors through innovative service platforms," *Journal of service management*, vol. 22, no. 4, pp. 443–470, 2011.
- [5] V. Kuppuswamy and B. L. Bayus, "Crowdfunding creative ideas: The dynamics of project backers in kickstarter," *UNC Kenan-Flagler Research Paper*, no. 2013-15, 2014.
- [6] E. Mollick, "The dynamics of crowdfunding: An exploratory study," *Journal of Business Venturing*, vol. 29, no. 1, pp. 1–16, 2014.
- [7] "Linguistic inquiry and word count (liwc)." [Online]. Available: <http://liwc.wpengine.com/>
- [8] "Mrc psycholinguistic database." [Online]. Available: http://websites.psychology.uwa.edu.au/school/MRCDatabase/uwa_mrc.htm
- [9] "Textblob library." [Online]. Available: <https://textblob.readthedocs.org/en/dev/>