Plead or Pitch? The Role of Language in Kickstarter Project Success

Nihit Desai and **Raghav Gupta** and **Karen Truong**
Department of Computer Science
Stanford University
Stanford, CA 94305, United States
{nihit,rgupta93,truongk}@stanford.edu

Abstract

We present an analysis of over 26000 projects from Kickstarter, a popular crowdfunding platform. Specifically, we focus on the language used in project pitches, and how it impacts project success/failure. We train a series of discriminative models for binary success/failure prediction of projects with increasing complexity of linguistic features, and show that successful project pitches are, on average, more emotive, thoughtful and colloquial. In addition, we also present an analysis of pledge rewards in Kickstarter, and how these differ across categories and for successful versus unsuccessful projects.

1 Introduction

The advent of web-based crowdfunding in recent years has, by connecting project creators to potential backers all around the world, engendered a healthy and effective environment for the growth of creative as well as entrepreneurial projects created by individuals and small groups, which are otherwise unable to invest the time and effort required for obtaining funding for their projects. While recent work has discussed patterns in what factors correlate well with the success (or failure) of crowdfunding projects, not much has been made of formulating what constitutes a successful and attractive project pitch. While project metadata does assume importance in determining project success, the language used in the project pitch is also a major determinant of success as one of the few variables project creators can tinker with pre-launch to maximize their chances of success.

Briefly, in this work, we explore two dimensions of the problem of formulating attractive project pitches - the description text and the rewards proffered in return for specific pledge amounts. While these dimensions are portable to other online crowdfunding portals, their specifics with regards to Kickstarter are discussed later. The precise questions we try to answer are:

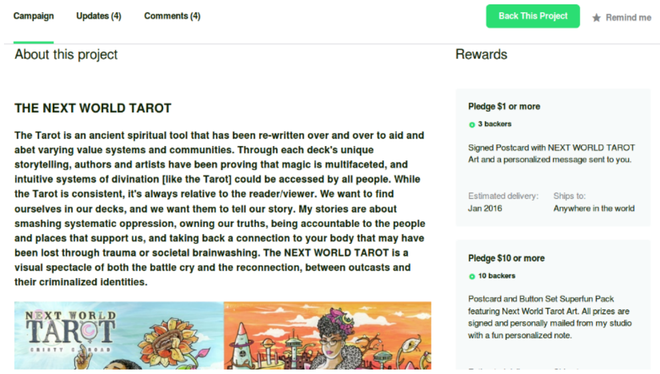

Figure 1: A sample Kickstarter project page showing project description and pledge levels with rewards

- What are some characteristics of the language used in successful projects? Can we use text to effectively predict success of campaigns?
- What kinds of pledge rewards are popular among backers? Do backers desire something in return, or are they altruistic?

The rest of this report is organized as follows. We begin with a discussion of related work. Then, we describe our domain - the Kickstarter model and the dataset extracted from the site. Following this, we detail our experimental design, after which we present our feature set, results and a thorough discussion of the performance of our features. Finally, we present peculiar observations regarding the different kinds of pledge rewards and their role in project success. We end with directions for future work.

2 Related Work

In this section, we give an overview of related work on Kickstarter crowdfunding. Next we discuss work on the social analysis of text. Lastly, we describe persuasion theories will be used to interpret our results.

2.1 Crowdfunding dynamics

There have been several studies published on crowdfunding platforms. (Etter et al., 2013) analyze the social network by constructing a projects/backers graph and monitoring Twitter for Tweets that mention the project, and also present results on prediction based on the time series of early funding obtained. A study by (Mollick, 2014) on Kickstarter dynamics finds that higher funding goals and longer project duration lead to lower chances of success, while inclusion of a video in a project pitch and frequent updates on the campaign increase the likelihood of full funding.

While the above work presents various statistics about the determinant features for success, there has also been work on analyzing the language in project descriptions (Mitra and Gilbert, 2014). However, it only looks at projects when the campaign has passed its last date for collecting funds. In the present work, we instead examine the first day of a campaign when the project creator still has the ability to make changes that increase their likelihood of success.

2.2 Analyzing text for social behavior

In recent years, scientists have been investigating the relationship between social behavior and language in online settings. (Hancock et al., 2007), for instance, have contributed to work on extracting positive and negative emotional tone from text-based communication. Studying the linguistic aspects of politeness, (Danescu-Niculescu-Mizil et al., 2013) show that Wikipedia editors are more likely to be promoted to a higher status when they are polite. In another related work, (Althoff et al., 2014) analyze social features in literature to determine relations that can predict whether an altruistic request will be accepted by a donor.

By analyzing the language and words used in daily conversation, linguists have developed tools like LIWC, a dictionary for inferring psychologically meaningful styles and social behavior from unstructured text (Pennebaker et al., 2001). Other resources like the NTU Sentiment Dictionary and SentiWordNet aid in opinion mining on text-based content (Pang and Lee, 2008).

2.3 Theories of influence

Persuasion can be defined as “a conscious attempt by one individual to change the attitudes, beliefs, or behavior of another individual or group of individuals through the transmission of some message” (Bettinghaus and Cody, 1980). Reciprocity, the tendency to return a favor after receiving one, plays a heavy role in social dynamics in investment situations (Berg et al., 1995). Previous work has demonstrated that the scarcity of a reward is also a significant factor in assessing a reward’s worth, where incentives

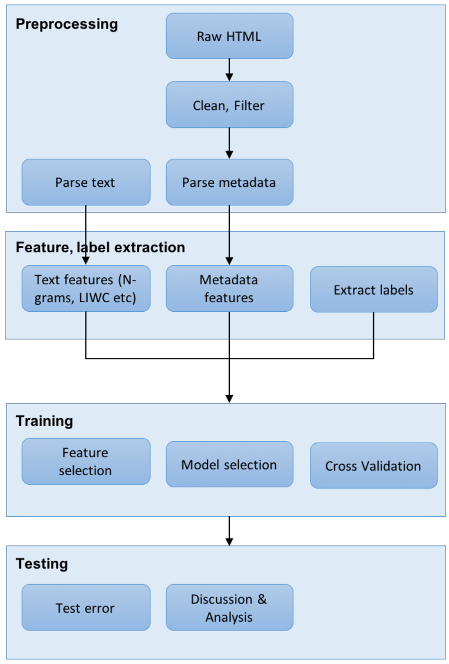

Figure 2: Data processing workflow

that are presented as limited in availability or quantity have an increased perceived value (Cialdini, 1993).

Using this framework of persuasion and reward power, we will delve into our quantitative findings to see what kinds of rewards project authors should offer in exchange for financial backing.

3 Domain and Dataset

We describe here the Kickstarter portal and funding model, followed by a description of our dataset.

3.1 Kickstarter

With over $2bn pledged to a total of 97,311 projects, Kickstarter is the biggest of popular online crowdsourcing portals[1], with projects in diverse domains, including but not limited to performing arts, technology, film, food and clothing. Other than the size and diversity, Kickstarter’s all-or-nothing model makes it even more suitable for a prediction task; projects have a goal amount, and if this funding goal is not met while the project is live, no money is handed to the creator. Furthermore, Kickstarter’s self-reported statistics suggest that projects that fail tend to fail fully: 78% of projects that raise at least 20% of their goal ultimately succeed.[2]

A sample Kickstarter page in Figure 1 shows some relevant components of the project page. We detail the ones we consider.

[1] http://www.forbes.com/sites/chancebarnett/2013/05/08/top-10-crowdfunding-sites-for-fundraising/
[2] https://www.kickstarter.com/help/stats

Metric	Successful	Failed
# Campaigns	7862	18681
%ge campaigns	29.6%	70.4%
Duration in days (avg)	27.1	21.2
Goal amount (avg)	$5747	$19344
# Sentences	24.5	23.6
# Rewards offered	8.73	7.52

Table 1: Kickstarter corpus statistics

- Project description, which can contain text, images and video
- Risk description, detailing the potential reasons the project might fail, even if funded
- All-or-nothing goal amount, and
- Pledge levels - a backer can choose to pay one of many possible amounts; each comes with a 'reward' for the backer from the project creator

3.2 Data

For this task, we use a dataset of 26,543 Kickstarter campaigns whose final outcome (success or failure) is known. The dataset comprises, for each of these projects, one snapshot per day of the project's Kickstarter webpage (in HTML format), while the project is live. The crawl, which was run daily for a period of about a year starting June 2014, contains a daily snapshot of all projects active on that day. The project webpage includes all the aforementioned relevant components, and more metadata that is potentially useful for prediction.

Note that for our precise problem definition, we are interested in the project pages' snapshots only on the first and last days of the project's funding period.

While extracting our features from the raw HTML files, which contained considerable noise, we filtered out campaigns that were:

- Still active on the last day of data collection
- Canceled prematurely by the creator
- Missing data, preventing us from inferring eventual success or failure of the campaign

After cleaning, we parse the HTML documents to extract four types of information: textual description, metadata (including but not limited to the title, start and end date, funding goal, image & video URLs), pledge information and the class label. The overall data processing flow is outlined in Figure 2, and some important statistics about the data are presented in Table 1.

We believe that these corpus statistics raise quite a few interesting questions. Are successful project pitches wordy or brief? Do they try to appear attractive by offering more rewards to backers? We investigate some of these hypotheses in later sections.

4 Experimental Setup

We now detail our experimental setup - mathematical objective and feature sets.

4.1 Problem Definition

We model our task as a binary classification task, treating each project as a training (or testing) example, and with labels being *"success"* and *"fail"*, depending on whether the project met its funding goal or fell short. To analyze the various features quantitatively as well as qualitatively for their impact on classification performance, we use two types of models: l_2-regularized logistic regression and support vector machines with RBF (Gaussian) kernel. We balance class weights since our dataset is skewed towards unsuccessful campaigns (after balancing, a naive "most common label" classifier would have an accuracy of $\sim$ 50%). We provide more details about our model in section 5.

In our experiments, we divide our dataset into training set and dev set (80% and 10% respectively) which we use to run our model selection and feature selection experiments. Finally, we report our performance on the test set (10%).

4.2 Features

We now describe our sets of features for classification, and interesting hypotheses we intend to investigate relating to these features.

4.2.1 Metadata - Baseline

Our first model is a baseline including project metadata attributes that have been explored and shown to be useful in prior work, including (but not limited to)

- Goal amount
- Project category
- Number of videos, images, comments
- Number of projects previously created and previously backed by creator
- Number of pledge levels

4.2.2 N-grams

We next add to our set of features TF-IDF matrices of uni, bi and tri-grams from the project description and the section on risks, in order to discover phrases that are strongly predictive of success or failure. We also experimented with using the plain counts matrix of the n-grams, but that led to a drop in performance. Additionally, we filter our phrases that are very specific to a particular to a given category (e.g. 'virtual reality', 'flying robot', 'arduino' for technology category). This is because we are interested in analyzing the role of language that is employed by successful campaigns. Filtering out category-specific phrases is helpful to filter our irrelevant phrases for this task.

We believe that analyzing the most predictive phrases, judged by the feature weight decided by the classifier, will provide valuable insights into the precise style of language used in successful campaigns, along with psycholinguistic features, which we discuss next.

4.2.3 Psycholinguistic features

We use LIWC: Linguistic Inquiry and Word Count (Pennebaker et al., 2001) to extract interesting psycholinguistic features from the data. LIWC is a lexical database providing sets of words relating to different psycholinguistic categories. Some examples are (*s denote stems)-

- Certainty: invariab*, factual*, absolutely
- Leisure: horseback, dvd*, celebrat*
- Achievement: effect*, persever*, founded

These categories, in conjunction with phrases predictive of success, will help us identify the importance of certain characteristics of language use in successful (and unsuccessful) campaigns - should the language appeal to emotion or logic? Should it be formal and polite, or candid and informal?

4.2.4 Sentiment from user comments

We use the sentiment analyzer in Stanford CoreNLP (Manning et al., 2014) to annotate the comments left on the project webpage by backers and other people. We surmise that the sentiment scores will be indicative of how positively (or negatively) the project is being perceived online. The number of sentences with different sentiment scores ({0,1,2,3,4}, 4 most positive, 0 most negative) form our feature values. It is worth nothing, however, that this feature is not important for success prediction on the first day of the campaign, and is used in our experiments for prediction only on the last day of the project's lifetime.

5 Experiments

In this section, we give details about the experiments we conducted and present results.

5.1 Feature sets

Our first experiment was to test the usefulness of the various features we described previously. Starting with a baseline of predicting based only on the project category, we progressively add more features to our model: project category, project category + metadata, project category + metadata + LIWC categories, and so on.

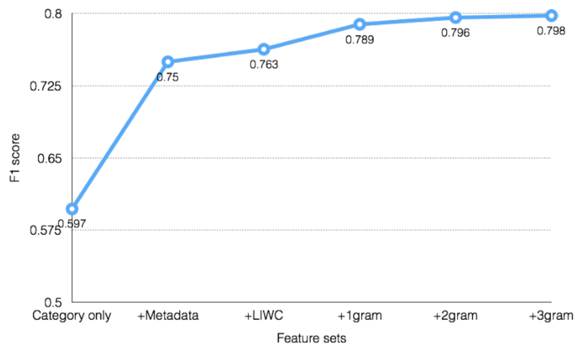

Figure 3: Performance of different feature sets on an l_2-regularized logistic regression model with balanced class weights.

We track performance of these sets of features by training a series of l_2-regularized logistic regression models with balanced class weights on the training set. Figure 3 shows the increase in F1 score on the dev set as more features are added. In subsequent experiments, we will use the entire feature set.

5.2 Models

Next, we experimented with various estimators (model types) to see which one is best suited to our task. As a starting point, we'd like to point out that our prediction task is more naturally modeled using a discriminative approach rather than a generative one. The reason is that all campaigns aspire to succeed. Therefore, it is better to model conditional probability of the label given the data as opposed to their joint probability. Results of some models we tried out are given in Table 2. For notational convenience, we use abbreviations **P** for precision, **R** for recall, **F1** for F-1 score and the subscripts s for label=*successful*, f for label=*failed* and o for overall. To denote different models we use **NB** for multinomial naive bayes, **DT** for decision trees, **LR** for Logsitic regression and **SVM** for support vector machines with Gaussian kernels.

We find that Logistic Regression (with l_2 regularization) and SVM (with Gaussian kernels) give the best results. SVM performs slightly better but also requires more training time compared to Logistic Regression. In subsequent experiments, we will use SVM with Gaussian kernels.

	$\mathbf{P}_s$	$\mathbf{R}_s$	$\mathbf{P}_f$	$\mathbf{P}_f$	$\mathbf{P}_o$	$\mathbf{R}_o$	$\mathbf{F1}_o$
NB	.65	.58	.80	.84	.75	.76	.75
DT	.54	.55	.84	.82	.77	.75	.76
LR	**.74**	.61	.80	**.88**	.78	.78	.78
SVM	.71	**.65**	**.84**	.87	**.79**	**.80**	**.79**

Table 2: Predictive performance of various models

5.3 First day and last day

Having experimented with feature sets and different models, we use these results to see how well can we predict success of a campaign as it progresses. In Table 3, we report the performance of our system on campaign snapshots on the first day (as a lower bound) and on campaign snapshots on the last day (as an upper bound). This is our final experiment for the prediction task and we use the test set which has thus far not been touched. For notational convenience, we use abbreviations as explained in the previous section.

Day	**P**$_s$	**R**$_s$	**P**$_f$	**P**$_f$	**P**$_o$	**R**$_o$	**F1**$_o$
First	.72	.63	.84	.87	.79	.79	.79
Last	.88	.72	.86	.94	.86	.86	.86

Table 3: Predictive performance on the first and last day of a campaign

6 Classifier Errors and Challenges

In this section, we discuss the major sources of the errors made by our classifier.

6.1 Image/video campaign descriptions

Some campaigns describe their campaigns fully or partially through images and videos i.e. the text describing the campaign, and often the rewards, is a part of an image. This is particularly common for projects in the *Fashion* and *Crafts* categories

Since open source frameworks for extracting text from images (we tried Tesseract[3]) are not sufficiently robust for our needs, successful projects falling in this category are often misclassified as unsuccessful (possibly owing to the inherent skew of the dataset).

6.2 Parsing HTML

We believe that the Kickstarter project page format was changed once during the period when the dataset was crawled. As a result, certain attributes for a fraction of the projects were not extracted properly, as the HTML was not parsed completely. These cases may require manual intervention/labeling.

6.3 Data skewed towards good features

Relying only on linguistic features for classification is a bit tricky, since all projects aspire to succeed, and are framed with success in mind. This prevents an ultra-clear demarcation between the language used in successful and unsuccessful campaigns. In other words, there will be false positives and false negatives with any approach relying on features that take values with one class label as the goal, as in our case. This is very different from a problem such as spam classification, where the language found in spammy emails is very different from normal emails.

6.4 Reward quality missing as feature

While we implement certain heuristics for judging whether a reward is deliverable, personalized, etc., a true measure of the quality and attractiveness of a reward is much harder to obtain. LIWC partly overcomes this difficulty with the description text, but we believe that a measure of attractiveness of the rewards would be an important addition to our classifier, while providing interesting hypotheses to evaluate.

7 Discussion

In this section, we analyze results from our experiments so far and present possible conclusions that can be drawn from them. All analysis and discussion in this section uses results from campaign snapshots on first day of the campaign.

7.1 Predictive phrases

Upon taking a closer look at features with highest weights in our model, we find that most of these features are in fact n-grams. In this section, we analyze n-gram features that are highly predictive of success of a Kickstarter campaign and try to offer some insights in this regard:

- **Reciprocity**: Reciprocity is the tendency to return a favor in exchange for receiving one. People often use persuasive appeals when reflecting norms of reciprocity (i.e., If you do X, you will receive Y). In our analysis, we find many predictive phrases that are often used to offer a reward or a gift in return for donation funds, such as *free shipping* and *you receive*. Refer to Table 4 for a more extensive list.
- **Social relationship**: Prior work has shown that networks and community dynamics are particularly important to the success of crowdfunded projects (Kuppuswamy and Bayus, 2015). In our analysis, we find many predictive phrases employ such dynamics and provide the social context in which their projects, if successful, will have a positive impact. For example, *community* and *friends*. Refer to Table 4 for a more extensive list.
- **Emotional appeal**: In our analysis, we find that successful campaigns are emotionally appealing to the readers. For example, we find phrases such as *passion* and *inspired to start* highly predictive of success. This is also indicated by analyzing LIWC categories. Refer to Table 4 for a more extensive list.

[3] https://code.google.com/p/tesseract-ocr/

Group	Phrase list
Reciprocity	free shipping, you receive, early bird, be the first, your reward
Social	friends, friendship, community, and family, his family, people
Emotion	passion, dream, inspired to start, believe that, impact, volunteer
Thankful	thank you, so thankful, thanks, thanks so much, grateful, grateful for your
Pitch	why support, funds will cover, will be used, aiming to, aim to raise
Collective	help us, we can, we raise, we plan to, we need, we found, we created

Table 4: Phrases in project descriptions most predictive of successful Kickstart projects, grouped

- **Gratitude**: In our analysis, we find that successful campaign text conveys gratitude towards the backers. For example, we find that phrases such as *thank you* and *thanks so much* are predictive of success. This, too, indicated by analyzing LIWC categories. Refer to Table 4 for a longer list.
- **Collective phrasing**: In our analysis, we discover that project descriptions making use of the singular first-person pronoun *'I'* tend to belong to unsuccessful projects, while those using the plural pronoun *'we'* are generally successful. This seems to suggest a multitude of things- projects by teams and/or individual projects framed in a subdued manner by the use of *'we'* tend to be more successful, *'I'* may come across as curt to a potential backer, *'we'* may serve as a politeness hedge, making the description more appealing (Danescu-Niculescu-Mizil et al., 2013).

7.2 LIWC categories

In conjunction with the phrases from Section 7.1, we look at weights for the features based on LIWC categories. Our LIWC results in Table 5 reveal some interesting findings: successfully funded campaigns demonstrate use of emotional appeal (LIWC categories *positive emotion*, *negative emotion*) and social processes (LIWC categories *friend, family*). An interesting observation is the high positive weight of the *tentative* category and the high negative weight of the *certainty* category: this leads us to hypothesize that tentativeness in the project description language indicates not a lack of confidence but politeness and hedging overconfidence instead.

Successful (high +ve wt)	Failed (high -ve wt)
emotion (positive, negative), friend, family, tentative, insight, we	anger, affect, death, sad, health, body i, certainty

Table 5: LIWC categories most predictive of successful/unsuccessful projects, judging by feature weights, presented in no particular order

Perhaps not surprisingly, results show that the membership in LIWC categories such as *death, sad, anger* is highly indicative of unsuccessful campaigns. However, the category *affect* is an important negative predictor, which may be due to the fine distinction between affect and emotion (conscious and unconscious expression of emotion, respectively). It is possible that descriptions trying to (consciously) express emotion may turn away potential backers.

Categories *health* and *body* contain words referring to body parts (slang included), which we observe are more often used in socially inappropriate contexts in project descriptions.

Further, an interesting comparison is between categories *we*, which is associated with successful campaigns and *i* which is associated with unsuccessful campaigns. This corroborates our previous finding on collaboratively framed campaigns being more successful.

8 Influence of Rewards on Project Success

In this section we analyze the various rewards project creators offer in return for pledges. Some questions of interest are: what kinds of rewards are commonly present in successful campaigns? Which rewards are popular among donors? Are donors looking for something concrete in return, or are they altruistic while funding projects? We describe our experiments and results in two steps:

- First, to identify keywords in reward texts alone indicative of success or failure, and
- Second, to analyze specific classes of rewards for their popularity and contribution towards the overall funding of successful projects.

8.1 Regression on reward text

We first train an l_2-regularized logistic regression classifier for predicting project success, trained only on the TF-IDF matrix of the text from the pledge rewards for each project. The text for different rewards for the same project is concatenated i.e. no distinction

Successful (high +ve wt)	Failed (high -ve wt)
private, early bird, CDs, PDF, stretch, shipping, print, retail, invitation, edition, unlocked, we'll, download, personalized	autograph, thank email, app, free, shirt, products, hand signed, logo, I'll, shoutout video, poster, coffee mug, hat, named

Table 6: Phrases from pledge reward text most predictive of successful/unsuccessful projects, judging by feature weights, presented here in no particular order

is made between rewards for the same project for this step.

Upon adding this feature to our best model, the performance drops by 0.01 F1, hence it is omitted from the prediction results. However, Table 6 presents phrases with high positive/negative weights, informative of commonly present reward types in successful/unsuccessful campaigns. Some interesting conclusions follow:

- Common **souvenirs** such as mugs, keychains and shirts, often autographed, indicate project failure
- Projects with **deliverable** rewards are relatively popular (as evidenced by *shipping, print, PDF, download, edition*)
- Rewards with a degree of **exclusivity** (*private*, *unlocked*, *early bird*) are possibly popular
- **Personalized** rewards, even if simple thankyou notes/postcards, have a positive weight. While *thank email* has a negative weight, *personalized email* has a positive weight
- This is an interesting observation - as evidenced by the weights of *we'll* and *I'll*, projects framed as **collective efforts** tend to be more successful, corroborating our earlier observations.

8.2 Reward classes

Given the observations in Section 8.1, we explore a couple of questions in more depth. Specifically, we look at individual successful projects and see what fraction of the funding for these comes from rewards that do not include a deliverable (we call these 'feel-good' rewards).

For this, we build a simple rule-based classifier that looks at the parse tree of the reward text's longest sentence, and recursively traverses the head followed by its direct object or clausal component head (dependencies *dobj*, *ccomp*, *xcomp* in that order), checking at each stage whether the current head token indicates a feel-good reward. We currently accomplish this with a hand-crafted lexicon of the most common headwords present in feel-good reward texts (such as *"thank"*, *"shout"*, *"hug"*). On manual inspection on a test set of 50 examples, this classifier has 92% precision and 72% recall.

We next obtain for each successful campaign the fraction of funding and of backers coming from feel-good rewards. We do this once for all successful campaigns, and once for successful campaigns in categories *Technology* and *Games* - among the most likely to have deliverable rewards. The histograms for the data are presented in Figure 4.

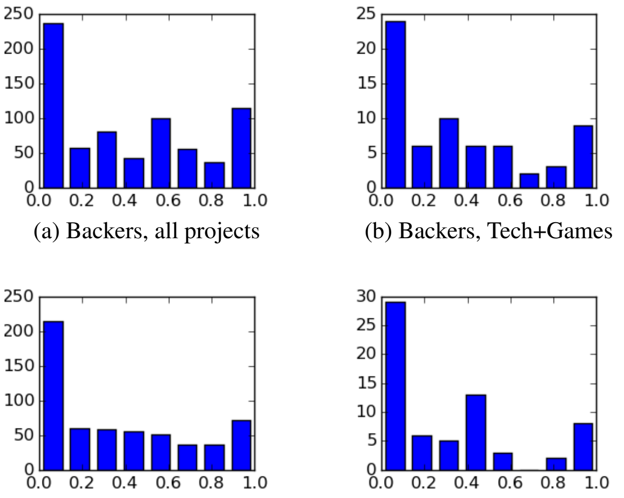

(a) Backers, all projects (b) Backers, Tech+Games

(c) Funding, all projects (d) Funding, Tech+Games

Figure 4: Fraction of funding and backers coming from feel-good rewards

Here, we make the following observations

- Across all projects, most backers and funding seems to come from concrete rewards, judging by the peak around 0 for each of these histograms. As should be expected, projects in categories with deliverables (in our case, *Tech* and *Games*), the proportion of backers and funding from feel-good rewards is slightly less than the average.
- Interestingly, for backers and for funding across all projects, there is a local maximum in the histograms close to 1. This is likely due to the high success rate of projects in categories such as *Dance*, *Theater* and *Film* which generally don't have deliverable rewards, but are not too expensive to fund either.
- A similar maximum for histograms for *Tech* and *Games* projects, we hypothesize, is because *Tech* and *Games* projects are often large scale, with fairly high pledge amounts which a backer may not want to pledge, while still wishing to support the project.

All in all, we gather that feel-good rewards do not make a major monetary contribution to the funding goals of the project; having deliverables wherever possible is of central importance, and for categories with no plausible deliverables (*Dance*, *Theater*), lower production costs could do the trick. There are, indeed, more classes of rewards that we could analyze separately (souvenirs, personalized gifts, and so on). These, we believe, are interesting directions for future research.

9 Conclusion and Future Work

In this report, we present an analysis of the role of language in persuasion, in the Kickstarter domain. We include metadata as well as linguistic features for binary classification of projects as successful or not. We show that linguistic features play a clear positive role in classification, obtaining good results on the initial project description. Further, we investigate the psycholinguistic style of the language used in successful campaigns, and how our findings correspond with existing literature on the role of language in persuasion. We also analyze how reward descriptions are influential in determining project success.

Apart from the feature set extensions discussed in Section 6, we discuss directions for future research. One interesting extension to this work would be to use the trained classifier itself, with or without additional information, to suggest to project creators which portions of a project description are lacking. In addition, a more thorough analysis of backers' preferences for different kinds of rewards, and its impact on project success, would be interesting to explore.

Note : Combining Part of Project with CS229

All three team members are crediting **part** of this project for CS229: Machine Learning. In particular, **parts of sections 3, 4, 5 and 6** will be included in our CS229 submission. However, our main focus with regards to the CS229 project, aside from getting good results on the classification task, will involve being rigorous with **model** and **feature selection** techniques, preventing overfitting, and analyze why certain models work better than others for the precise task at hand. Sections 7 and 8 in this report, providing an in-depth linguistic analysis of our features' performance, will **not** be part of our CS229 submission.

Acknowledgments

We would sincerely like to thank Rob Voigt (Linguistics Dept, Stanford University) for sharing the Kickstarter dataset, which allowed us to work on this project, and for his valuable inputs and suggestions. Additionally, we would like to thank Neha Nayak, for her valuable feedback on the project proposal, and Prof. Manning, for being a wonderful instructor.

References

Bo Pang and Lillian Lee. 2008. *Opinion mining and sentiment analysis*. Foundations and Trends in Information Retrieval:2.1-2:1-135.

Christopher Manning, Surdeanu Mihai, John Bauer, Jenny Finkel, Steven Bethard and David McClosky. 2014. *The Stanford CoreNLP Natural Language Processing Toolkit*. Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations, pp. 55-60.

Cristian Danescu-Niculescu-Mizil, Moritz Sudhof, Dan Jurafsky, Jure Leskovec and Christopher Potts. 2013. *A computational approach to politeness with application to social factors*. Proceedings of the Conference of the Association for Computational Linguistics, 2013.

Erwin P Bettinghaus and Michael J Cody. 1980. *Persuasive Communication*. Holt, Rinehart and Winston.

Ethan Mollick. 2014. *The dynamics of crowdfunding: An exploratory study*. Journal of Business Venturing 29.1:1-16.

James W. Pennebaker, Martha E. Francis, and Roger J. Booth.. 2001. *Linguistic inquiry and word count: LIWC 2001*. Mahway: Lawrence Erlbaum Associates 71.

Jeffrey T. Hancock, Christopher Landrigan, Courtney Silver. 2007. *Expressing emotion in text-based communication*. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 2007.

Joyce Berg, John Dickhaut and Kevin McCabe. 1995. *Trust, Reciprocity, and Social History*. Games and Economic Behavior: 10.1:122-142.

Robert B. Cialdini. 1993. *Influence: The Pyschology of Persuasion*. HarperCollins.

Tanushree Mitra and Eric Gilbert. 2014. *The Language that Gets People to Give: Phrases that Predict Success on Kickstarter*. Proceedings of ACM conference on Computer-Supported Cooperative Work and Social Computing.

Tim Althoff, Cristian Danescu-Niculescu-Mizil, Dan Jurafsky. 2014. *How to Ask for a Favor: A Case Study on the Success of Altruistic Requests*. Proceedings of the International Conference on Weblogs and Social Media.

Vincent Etter, Matthias Grossglauser, and Patrick Thiran 2013. *Launch hard or go home: predicting the success of Kickstarter campaigns*, Proceedings of the first ACM conference on Online social networks.

Vincent Etter, Matthias Grossglauser, and Patrick Thiran 2015. *Crowdfunding Creative Ideas: The Dynamics of Project Backers in Kickstarter*, Proceedings of the 3rd International Conference on Leadership, Technology and Innovation Management