Helpful or Hierarchical? Predicting the Communicative Strategies of Chat Participants, and their Impact on Success

Farzana Rashid
University of North Texas
farzanarashid@my.unt.edu

Tommaso Fornaciari
Bocconi University
fornaciari.tommaso@unibocconi.it

Dirk Hovy
Bocconi University
dirk.hovy@unibocconi.it

Eduardo Blanco
University of North Texas
eduardo.blanco@unt.edu

Fernando Vega-Redondo
Bocconi University
fernando.vega@unibocconi.it

Abstract

When interacting with each other, we motivate, advise, inform, show love or power towards our peers. However, the way we interact may also hold some indication on how successful we are, as people often try to help each other to achieve their goals. We study the chat interactions of thousands of aspiring entrepreneurs who discuss and develop business models. We manually annotate a set of about 5,500 chat interactions with four dimensions of interaction styles (motivation, cooperation, equality and advice). We find that these styles can be reliably predicted, and that the communication styles can be used to predict a number of indices of business success. Our findings indicate that successful communicators are also successful in other domains.

1 Introduction

People are social beings who communicate their feelings, emotions, thoughts, ideas, etc. through verbal and non-verbal interactions. Based on these interactions, we build relationships, and these relationships, in turn, help create and maintain a network of peers. Peers in a network *cooperate* with each other, help each other to *learn*, and exchange ideas. However, they also *compete* for the same resources (Vega-Redondo et al., 2019), not least attention. Peer networks are particularly important for innovation and entrepreneurship (Gonzalez-Uribe and Leatherbee, 2017), as they produce an active exchange of ideas.

People are usually assumed to be altruistic in networks like online social forums. They cooperate with and help one another with answers, advice, and ideas. The motivations behind helping a peer include, but are not limited to, getting pure pleasure from helping, self-advancement, building a reputation, developing relationships, or sheer entertainment (Tausczik and Pennebaker, 2012).

When people interact with each other, their interactions vary along various communicative styles, such as showing *cooperativeness*, *equality*, *business orientation*, etc. (Rashid and Blanco, 2018). Varying these communication styles provides tools to achieve communicative goals. For example, someone trying to build a reputation will tend to use a more *cooperative* style. Someone who tries to be helpful may use more words of *advice* in their interactions. The usage of relationship-establishing styles is more prevalent in certain personalities (Cheng, 2011) and in specific settings. Business-oriented people communicate more independence, tolerance of ambiguity, risk-taking propensity, innovativeness, and leadership qualities (Wagener et al., 2010).

The impact of these styles is, therefore, an essential factor in text analysis. However, due to their complex, decentralized nature, these communication styles have been studied very little in NLP. Cooperativeness is more than just a few keywords—it includes a whole inventory of communicative tools. This property makes it harder to annotate and predict. Part of the reason is the lack of adequate corpora. We provide such a corpus and report encouraging results for the above styles.

Contributions We introduce a new task, predicting the communicative strategies of interlocutors in a real-life setting, and provide a new, multiply-annotated data set of 5k+ instances. We find that the various communicative dimensions can be efficiently predicted. Additional tests suggest that the communicative strategy of a person is somewhat predictive of their business success.

Findings of the Association for Computational Linguistics: EMNLP 2020, pages 2366–2371
November 16 - 20, 2020.

	raw	κ	Krippendorff's α
Cooperative	76%	0.58	0.57
Motivational	91%	0.74	0.74
Equal	77%	0.33	0.34
Advice	91%	0.60	0.60
Average	83%	0.53	0.52

Table 1: Inter-annotator agreements for communication styles. κ values between 0.6 and 0.8 are considered *substantial* agreement, and above 0.8, *nearly perfect* agreement (Artstein and Poesio, 2008).

2 Data

Our ultimate goal is to predict the communicative styles and strategies of aspiring entrepreneurs in an online peer network. The initial corpus is part of a large-scale social science experiment that involved around 5,000 entrepreneurs from 49 African countries (Vega-Redondo et al., 2019). After completing an online business course, those entrepreneurs interacted in groups of sixty through an Internet platform for about two and a half months, resulting in approximately 140,000 chat interactions.

Besides the chat interactions, the original dataset contains background information about the speakers (country of origin, educational background, age, gender, etc.). All the participants submitted business proposals, which were evaluated by a panel to assess their potential.

The original experimental setup was designed to assess how communication among peers affects innovation and entrepreneurship. Vega-Redondo et al. (2019) therefore already applied NLP techniques for semantic analysis of the interactions. They also manually annotated other indicators, i.e., business-relatedness, sentiment, and target audience (i.e., one or several people) on a subset of 10k sentences. They trained classifiers on this data to infer these labels for all remaining 130k instances in the corpus. This dataset provides a perfect starting point for our goals.

The first step to address our goal involves annotating speech styles on several interactions among the participants. We work on the same subset of previously annotated data, and add our own annotations to enrich the data further.

3 Annotating Communication Styles

We sample around 5,500 chat interactions (mostly in English language with traces of other language(s)) which were previously annotated for *business-relatedness*, *sentiment*, and *audience*, and annotate the four communication styles:

1. *Cooperativeness* indicating the friendliness shown towards the target audience, with label values *cooperative*, *competitive*, and *neutral*.
2. *Advice* indicating whether the interaction contains any words of advice with label values *advice* and *neutral*.
3. *Motivational* indicating whether the interaction contains any words of motivation, with label values *motivation* and *neutral*.
4. *Equality* indicating whether there is a display of hierarchy between the speaker and the receiver, with label values *equal* and *hierarchical*.

For all styles, *unknown* is used whenever it is hard to determine any of the other values from context.

3.1 Annotation Process

Three graduate students with experience in NLP tasks annotated the corpus. They were trained with written annotation guidelines consisting of definitions and examples for all the communication styles. They also had an hour-long session carrying out sample annotations to ensure that they properly understood the problem.

For the annotations, the annotators filled out their responses in interactive spreadsheets choosing the correct value for a particular style. Each of the annotators annotated their part of around 2,100 chat interactions. 502 of these were shared among all three annotators so that we can compute agreement measures. We obtain the most probable labels for the shared portion using MACE (Hovy et al., 2013).

We summarize the inter-annotator agreement coefficients in Table 1 (raw agreement: 83%; averaged pairwise Cohen's κ: 0.53; Krippendorff's α: 0.52). The average MACE competence score of these annotators is 0.53.

Table 3 shows the Pearson's correlations between pairs of the styles of interactions and previous annotations. This indicates that ***Motivational*** styles are usually also ***Cooperative*** (0.61), give ***Advice*** (0.56) and are ***Equal*** (0.54). Interestingly, many ***Business-related*** interactions are not very ***Cooperative*** (-0.42).

The label counts are as follows. For *cooperativeness*, 42.3% are labeled *cooperative*, 50.3% are

Cooperative (Rashid and Blanco, 2018; Wish et al., 1976)	
1: Mobile Webshop is a very good concept.	Cooperative
2: You have not done anything yet.	Competitive
Motivational	
3: @NAME1234 well said @NAME456 start small and dream big..welldone	Motivational
4: I meant to say voting contest to be precise.	Neutral
Equal (Rashid and Blanco, 2018; Wish et al., 1976)	
5: Wishing you a very wonderful weekend.	Equal
6: Happy to engage you on this..	Hierarchical
Advice	
7: Think about it.	Advice
8: This is cool Sunday.	Neutral

Table 2: Annotation examples with contrasting values for each communication style. Each chat interaction can be of varying length and is either directed to an individual or others in general.

	S	C	M	E	A
B	-0.17	-0.42	-0.28	-0.27	0.01
S	–	0.29	-0.01	-0.04	-0.10
C		–	**0.61**	0.37	0.40
M			–	**0.54**	**0.56**
E				–	0.26

Table 3: Pearson correlations between pairs of styles of interactions (indicated by the initial letters of ***S**entiment*, ***C**ooperative*, ***M**otivational*, ***E**qual* and ***A**dvice*).

neutral and only 2.14% are labeled *competitive*. For the *motivational* style, 14.1% are *motivational* and 81.2% are *neutral*. For the *advice* style, 9.2% are *advice* and 85.9% are *neutral*. For the *equality* style, 77.3% are *equal* and 8.8% are *hierarchical*.

We release our annotations as stand-alone annotations.[1]

3.1.1 Annotation Examples

Table 2 shows a number of actual chat interactions from the dataset with different values for the styles annotated. In interaction (1), the praising is considered a *cooperative* response, whereas in (2) the speaker is chiding someone, indicating a *competitive*ness. The praise in (3) is *motivational*. Example (4) does not really communicate any motivation, so it is labeled as *neutral* for this style. Example (5) is just a greeting and does not indicate anybody displaying hierarchy over anyone else, so it is *equal*. Example (6) shows that the speaker instructs someone on how to behave (*hierarchical*). In (7), the speaker is *advising* someone to think about a matter whereas example (8) is just another neutral statement.

[1] https://github.com/MilaNLProc/conversationstyle

4 Experiments and Results

We want to predict four styles of interactions (*cooperative*, *motivational*, *advice*, *equality*), and three subsequent indicators of business success: (1) whether the person owns a business (HAS BUSINESS), (2) whether someone has ever owned a business (BUSINESS EVER) and (3) whether they submitted a business proposal to win funding to start a business (BUSINESS PROPOSAL).

We use (1) an SVM classifier with RBF kernel (effective in (Rashid and Blanco, 2018)) to predict both the communicative styles and the business success indicators, and (2) a Multitask Learning (MTL) Convolutional Neural Network to predict the business success indicators.

We divide our annotated dataset into 80-20 stratified train-test splits for predicting communicative styles. For predicting indicators of business success, we use 500 randomly selected instances as test and the rest as training data.

4.1 SVM setup

We use the SVM implementation in scikit-learn (Pedregosa et al., 2011) and tune the hyperparameters

(C and γ) using 10-fold cross-validation within the train split. We train one classifier per style and per indicator of business success to predict the different labels.

Feature Set. After basic preprocessing (removal of stop words), tokenization, and parsing (to get the root verb) using spaCy, we extract features from the chat interactions and sentiment lexica. The feature set relies only on language usage. We extract the first word in a chat interaction, the bag-of-words representations (binary flags and tf-idf scores) of the chat interaction and features from sentiment lexica. Specifically, we extract flags indicating whether the turn has a positive, negative or neutral word in the list by Hamilton et al. (2016), the sentiment score of the chat interaction (summation of sentiment scores per token over number of tokens), and a flag indicating whether the interaction contains a negative word from the list by Hu and Liu (2004). We also extract other features, which include (a) the root verb (b) binary flags indicating the presence of exclamation, question marks and negation cues from Morante and Daelemans (2012).

4.2 Multitask Learning (MTL) setup

We use a standard Convolutional Neural Network over word-embeddings, with one output per task. We preprocess the data (convert to lowercase, removed URLs and stop-words, converted numbers to 0's etc.) and learn a skip-gram embeddings model (Mikolov et al., 2013) trained for 50 epochs. We use an embedding size of 512, choosing a power of 2 for memory efficiency.

In the CNN, the input layer has the word indices of the text, converted via the embedding matrix into word embeddings. We convolve two parallel channels with max-pooling layers, and convolutional window sizes 4 and 8 over the input. The two window sizes account for both short and relatively long patterns in the texts. In both channels, the initial number of filters is 128 for the first convolution, and 256 in the second one. We join the convolutional channels' output and pass it through an attention mechanism (Bahdanau et al., 2014; Vaswani et al., 2017) to emphasize the weight of any meaningful pattern recognized by the convolutions. We use the implementation of Yang et al. (2016). The output consists of 7 independent, fully-connected layers for the predictions, respectively in the form of discrete labels for classification of one of the business success indicators of a person (HAS BUSINESS, BUSINESS EVER or BUSINESS PROPOSAL) as the target task, and the styles of interactions (*business*, *sentiment*, *cooperativeness*, *motivational*, *advice*, *equality*) as the auxiliary tasks. We trained one model per business success indicator.

	Model	P	R	F
Cooperative	majority	0.25	0.50	0.34
	SVM	0.77	0.77	**0.77**
Motivation	majority	0.66	0.81	0.73
	SVM	0.90	0.90	**0.89**
Equal	majority	0.60	0.77	0.67
	SVM	0.78	0.81	**0.78**
Advice	majority	0.74	0.86	0.79
	SVM	0.86	0.88	**0.86**
HAS BUSINESS	majority	0.51	0.71	0.59
	SVM	0.58	0.68	0.59
	MTL	0.61	0.66	0.63
BUSINESS EVER	majority	0.20	0.44	0.27
	SVM	0.54	0.47	0.38
	MTL	0.52	0.51	0.51
BUSINESS PROPOSAL	majority	0.57	0.75	0.64
	SVM	0.66	0.69	0.67
	MTL	0.65	0.75	0.65

Table 4: Results for predicting styles of interactions and three indicators of business success. The F-measures are the weighted averages of the F-measures of the two labels.

4.3 Results

Table 4 compares the results of the different systems to predict the styles of interactions as well as the business success indicators. Our SVM model does much better than the majority baseline for all the styles of interactions (F-measures = 0.77, 0.89, 0.78 and 0.86). For the indicators of business success, either the SVM (F-measures = 0.59, 0.38 and 0.67) or the MTL (F-measures = 0.63, 0.51 and 0.65) model outperforms the majority baseline.

5 Related Work

There have been a few studies analyzing language usage when people communicate. For example, researchers have studied power (or hierarchical) relationships in online communities (Danescu-Niculescu-Mizil et al., 2012), emails (Prabhakaran and Rambow, 2014), and social networks (Bramsen et al., 2011). Some have studied how roles of Wikipedia editors affect their success (Maki et al., 2017). Danescu-Niculescu-Mizil et al. (2013) analyze politeness in online forums using structural and linguistic features derived from the communications between two individuals. Katerenchuk and Rosenberg (2016) develop an algorithm to predict user influence levels in online communities. Rashid and Blanco (2018) characterize interactions between people with dimensions and produce a dataset annotating dimensions on TV scripts. Vega-Redondo et al. (2019) annotate business relevance and sentiment on online chat interactions among aspiring entrepreneurs.

In contrast, we annotate the communicative styles *cooperativeness*, *motivational*, *advice* and *equality* on chat interactions between young aspiring entrepreneurs, and develop machine learning systems to automatically predict these styles and indicators of business success for the participants.

6 Conclusions

We present a data set of 5k+ instances annotated with four communication styles which can effectively be predicted. These communicative styles also influence people's business success. Our results and data set open up interesting new avenues to study the effects of people's communicative strategies on their business success.

References

Ron Artstein and Massimo Poesio. 2008. Inter-coder agreement for computational linguistics. *Comput. Linguist.*, 34(4):555–596.

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. Cite arxiv:1409.0473Comment: Accepted at ICLR 2015 as oral presentation.

Philip Bramsen, Martha Escobar-Molano, Ami Patel, and Rafael Alonso. 2011. Extracting social power relationships from natural language. In *Proceedings of the 49th Annual Meeting of the ACL*, pages 773–782.

Kevin H. C. Cheng. 2011. Further linguistic markers of personality: The way we say things matters.

Cristian Danescu-Niculescu-Mizil, Lillian Lee, Bo Pang, and Jon Kleinberg. 2012. Echoes of power: Language effects and power differences in social interaction. In *Proceedings of WWW*, WWW '12, pages 699–708. ACM.

Cristian Danescu-Niculescu-Mizil, Moritz Sudhof, Dan Jurafsky, Jure Leskovec, and Christopher Potts. 2013. A computational approach to politeness with application to social factors. In *Proceedings of the ACL*, pages 250–259.

Juanita Gonzalez-Uribe and Michael Leatherbee. 2017. The Effects of Business Accelerators on Venture Performance: Evidence from Start-Up Chile. *The Review of Financial Studies*, 31(4):1566–1603.

William L. Hamilton, Kevin Clark, Jure Leskovec, and Dan Jurafsky. 2016. Inducing domain-specific sentiment lexicons from unlabeled corpora. In *Proceedings of EMNLP*, pages 595–605.

Dirk Hovy, Taylor Berg-Kirkpatrick, Ashish Vaswani, and Eduard Hovy. 2013. Learning whom to trust with MACE. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1120–1130, Atlanta, Georgia. Association for Computational Linguistics.

Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '04, pages 168–177, New York, NY, USA. ACM.

Denys Katerenchuk and Andrew Rosenberg. 2016. Hierarchy prediction in online communities. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI'16, pages 4224–4225. AAAI Press.

K. Maki, M. Yoder, Y. Jo, and C. Rosé. 2017. Roles and success in wikipedia talk pages: Identifying latent patterns of behavior. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing*, pages 1026–1035.

Tomas Mikolov, Scott Wen-tau Yih, and Geoffrey Zweig. 2013. Linguistic regularities in continuous space word representations. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT-2013)*. Association for Computational Linguistics.

Roser Morante and Walter Daelemans. 2012. Conandoyle-neg: Annotation of negation cues and their scope in conan doyle stories. In *Proceedings of LREC*, pages 1563–1568. ACL Anthology Identifier: L12-1077.

F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Vinodkumar Prabhakaran and Owen Rambow. 2014. Predicting power relations between participants in written dialog from a single thread. In *Proceedings of the ACL*, pages 339–344.

Farzana Rashid and Eduardo Blanco. 2018. Characterizing interactions and relationships between people. In *Proceedings of EMNLP*, pages 4395–4404.

Yla Tausczik and James Pennebaker. 2012. Participation in an online mathematics community: Differentiating motivations to add. pages 207–216.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *CoRR*, abs/1706.03762.

Fernando Vega-Redondo, Paolo Pin, Diego Ubfal, Cristiana Benedetti-Fasil, Charles Brummitt, Gaia Rubera, Dirk Hovy, and Tommaso Fornaciari. 2019. Peer Networks and Entrepreneurship: a Pan-African RCT. Working Papers 648, IGIER (Innocenzo Gasparini Institute for Economic Research), Bocconi University.

Stephanie Wagener, Marjan Gorgievski, and Serge Rijsdijk. 2010. Businessman or host? individual differences between entrepreneurs and small business owners in the hospitality industry. *The Service Industries Journal*, 30(9):1513–1527.

Myron Wish, Morton Deutsch, and Susan J Kaplan. 1976. Perceived dimensions of interpersonal relations. *Journal of Personality and Social Psychology*, 33(4):409.

Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical attention networks for document classification. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1480–1489, San Diego, California. Association for Computational Linguistics.