Differential expressive-receptive patterns in textual communication in the autism spectrum

Karthik Dinakar
karthik@media.mit.edu
MIT Media Lab

Abstract

Pragmatic language impairment, which has been related to autism and the Asperger syndrome, is a condition that involves difficulties in interpreting the pragmatics of language. It is well reported that people with this condition have difficulties with the semantic aspects of language, especially in the context of social interaction. Recent surveys have shown that though online social networking is an increasingly important means through which people on the spectrum can build relationships, difficulties in interpreting pragmatics of textual communication remains a hindrance in fostering such relationships. In this work, we build an interface tool with affordances for expression and reception of affect, conditioned on a reinforcement-driven affect-reasoning learning engine. The interaction sequence affords users feed-forward mechanisms to semantically mark a textual contribution with affect, as well as feedback mechanisms to help infer the affect of a textual contribution. We hypothesize that our tool will help people on the spectrum with pragmatic language impairment to better understand the pragmatics of conversations in online social networks.

1. Introduction

That people on the autism and Asperger syndrome have difficulties in the reception and expression of visual cues such as facial and gesture recognition pertaining to affect, particularly in the context of social interactions, is well known. Pragmatic language impairment (PLI), previously called semantic language impairment, is an umbrella term for a host of difficulties that individuals on the spectrum have with respect to language, with the distinction that it is considered a pervasive disorder, rather than a developmental one.

Characteristics of the condition include but are not limited to - difficulties understanding contextual cues in a conversation, preferment of facts over storied descriptions, and difficulties understanding satire or jokes. Individuals with PLI often fail to extract key points from a conversation or a story, with a tendency to get absorbed in literals. Contributions to conversations are also often one-sided or off-topic.

According to recent studies [Bishop], children with PLI have difficulties understanding and contributing to discourse, with conversational contributions that are socially irrelevant or contextually tangential. Studies have also shown a co-relation between PLI and behavioral problems in children.

Molecular genetic studies have found a common risk genotype for both ASD and PLI. It is important to distinguish between the two broad categories of problems that are collectively referred to as PLI. Difficulties with the structural aspects of language, such as phonology and syntax are problems that are diagnosed separately from the pragmatic difficulty of using language to communicate with others. Whilst the former kind of PLI has been explored, there is an increasing clamor for more research on the later. The co-morbidity between phonological and syntactic difficulties and difficulties with pragmatics of communication has not been fully established yet.

1.1 Social use of computer-mediated communication

The explosive growth of social networking websites and their defining characteristics of communication – which is essentially nonverbal, makes them an ideal fit for many people on the spectrum who otherwise experience difficulties in real social interactions. As such, studies have shown that for many nonverbal individuals on the spectrum, online social interactions are the primary modes to fostering personal relationships.

However, the lack of extraneous cues in textual communication, with little or no feed-forward or feedback mechanisms for the expression and reception of affect opens possibilities for semantic misinterpretations. In fact, studies have shown that computer-mediated communication intensifies problems of trust, disclosure and semantic misinterpretations that are vital in fostering healthy relationships [Burke].

In this work, we use techniques from natural language processing and commonsense reasoning to power an interface that provides affordances for the expression and reception of affect in textual communication. We hypothesize that awareness of affective valences in the context of a textual social discourse bolsters social interaction and makes it easier for individuals on the spectrum to engage in connected discourse on social networking websites.

This paper is hitherto organized as follows: the second section provides an overview of the related work for this paper. Section three describes the algorithmic approach towards modeling the affect-reasoning engine. Section four describes the interaction sequence of the interface that is powered by the affect-reasoning engine. In section five, we discuss strategies for evaluation of the model and the interface in tandem with an experiment protocol to gauge differential patterns of affect expression and reception by individuals on the autism spectrum. We conclude with section six by summarizing our approach and discussing possible outcomes for our experiment hypothesis.

2. Related work

Prior and related work falls under three categories, namely sentiment analysis and user-interaction design for computer-mediated social interaction, as well as prior tools for addressing PLI in the autism spectrum. Conversational agent interfaces powered by learning algorithms are also related areas for this work.

Sentiment analyses of text can generally involves either unsupervised or supervised learning algorithms or a mixture of both. Support vector machines for classification of text based on affect and opinion has become increasingly popular over recent years. Hidden-topic Markov models (HTMM) and vanilla latent-dirichlet allocation (LDA) are popular unsupervised learning methods that have also grown in popularity for sentiment analysis over the past couple of years. In the space of commonsense reasoning, Liu and Lieberman, using a filtered set of open-mind commonsense knowledge pertaining to affect, devised a society of models for commonsense reasoning as an affect recognition engine [Liu].

The marriage of learning algorithms and user-interaction comes closest in the realm of conversational interfaces. Tools based on conversational agents, such as agent-enabled tutoring systems for collaborative learning that learns from social and communicative factors are also related areas in this field. Most of the work in user-interaction design for social networks tends to focus on learning granular aspects of the dynamics of interaction that can promote easier end-user interaction and growth of the social network.

Much of literature pertaining to PLI agrees that very little has been researched beyond diagnostic approaches, calling for more research in the space of tools for intervention strategies and management of PLI.

2.1 So many approaches for sentiment analysis, so why another one?

In this work, we use a combination of natural language processing and commonsense reasoning approaches for the affect-reasoning engine. We blend relevant in-domain knowledge pertaining to affect with ConceptNet, a knowledge base with more than a million user generated commonsensical statements. Each user first begins with this seeded model. We allow a user to add affect to any contribution in an ongoing discourse and use a user's pattern of adding affect to further enrich the seeded model, thereby producing personalized affect-reasoning models for each user of the system.

Our approach differs from traditional corpora-dependent supervised and unsupervised learning algorithms in two important ways. First, traditional statistical learning methods exploit a bag-of-words approach towards classifying text based on affect, whose performances are often limited by the quality and size of the corpus. For effective reasoning, the inference mechanism must move beyond lexical patterns and words that are predictive of affective classes. Second, interpretation of affective valences in text is unique to each individual, further limiting the role of corpora for affect reasoning as a 'one-size fits all model'. This validates the need for a seeded model that learns from the user's pattern of affect reasoning.

Next, we design the user-interaction with affordances for affect expression and reception to help participants in an ongoing discourse to better understand the pragmatics of the conversation. We hypothesize that the design of the interface, with its feed-forward and feedback mechanisms, makes understanding the pragmatics of conversations far easier for people on the spectrum.

3. A modular affect-reasoning engine

The affect-reasoning engine can be divided into three sub-parts, namely the blending of in-domain knowledge with ConceptNet, preprocessing of text input and the process of inferring affect valences from the blended space and the preprocessed text. In this section we describe each of the aforementioned sub-parts before describing a reinforcement mechanism to refine the blended knowledge space with a user's pattern of affect annotation.

3.1 Blending in-domain knowledge with ConceptNet

ConceptNet is a crowd-sourced aggregation of about a million statements about everyday life that is represented in the form of assertions. Each assertion is a relation that maps two concepts. For example, consider the following assertion:

`wedding` *`Causes`* `happiness`

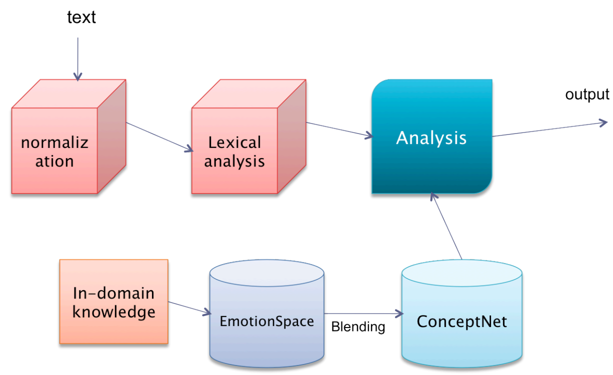

Fig 1: Components of the affect-reasoning engine

In the above example, the concepts wedding and happiness is mapped using one out of a set of twenty relations, in this case 'causes'. An assertion is an atomic unit of knowledge in the ConceptNet database. This representation of knowledge allows it to be transformed into a sparse matrix representation of concepts and relations that is amenable to techniques of dimensionality reduction.

In our work, we use a set of over 53,000 filtered YouTube comments based on unigrams denoting affect from the Ortony lexicon. Using a part-of-speech tagger, we extract event phrases and adjectives denoting affect and connect them with appropriate kind of relation to generate assertions.

We then blend the set of assertions derived from the filtered YouTube corpus with ConceptNet, a database with approximately a million assertions about everyday life using the blending technique invented by Robert Speer [Speer]. The resulting matrix after blending A, is then subjected to singular value decomposition (SVD) for a reduction in dimensionality. We factor A into an orthonormal matrix U, a diagonal matrix Σ, and an orthonormal matrix V^T, giving $A = U \Sigma V^T$. The singular values in Σ are then ordered in descending order, where the larger values correspond to the vectors in U and V that are more significant components of the initial A matrix.

We then subject the resultant matrix to principal component analysis, wherein we discard everything but the first k=100 components, i.e. the principal components of A, which form approximation dense matrix of the original data.

This gives us a model to perform concept comparisons, which is described in section 3.3.

3.2 Text preprocessing & lexical analysis

A given contribution from a user in a discourse is subjected to preprocessing operations before extracting relevant concepts for comparison from each sentence of the contribution. Each contribution is first subjected to operations of stemming and elimination of stop-words (words whose high frequencies make them ineffectual for natural language processing operations) using the NLTK toolkit.

Next, the given contribution is subjected to part-of-speech tagging to extract concepts at the sentence level, for each sentence in the contribution. At this stage we are quipped with the machinery needed to perform affect reasoning using our model A and the set of concepts C.

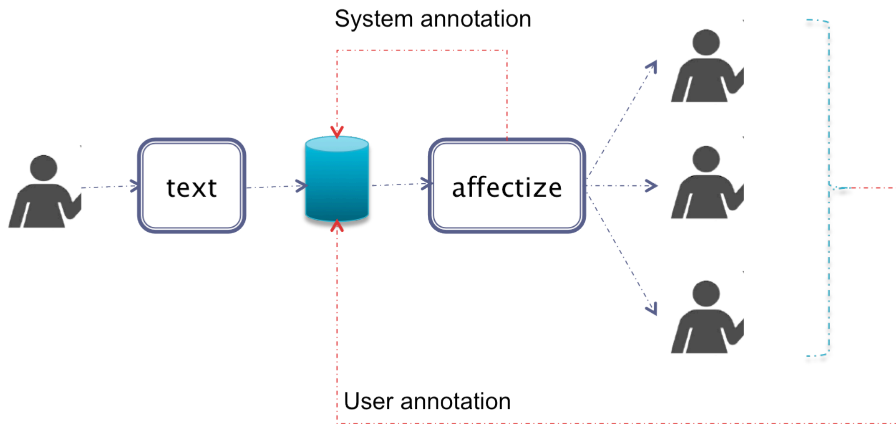

Fig 2: User interaction sequence

3.3 Affect-reasoning through cosine similarity of concepts

For the purpose of affect reasoning, it is necessary to choose a set of affective valences to compare the concepts extracted from the given contribution. We choose two affective valences, namely 'good' and 'bad' as canonical concepts denoting positive and negative connotations respectively.

Given the model **A**, we perform a dot product for each concept in the set **C** with the aforementioned canonical concepts. The result is essentially a *cosine similarity* of the two; thereby giving a measure of how related the two concepts are with each other. This is repeated for all the concepts in the set **C**, before normalizing each comparison with respect to the canonical concepts.

The resultant numbers for each canonical concept can then be compared with each other. We present these two results to the user to denote how much a given contribution tends towards the canonical concepts of 'good' and 'bad' respectively. In the next section, we describe the user interaction sequence that allows for user annotation and affordances for affect expression and reception.

4. An interface for affect expression and reception

The primary goal of the interaction design is to create affordances for affect expression and reception to improve the pragmatic understanding of a discourse for an individual with PLI. In this section, we describe the interaction sequence that meets this goal and also provides the benefit of user annotation for further refining the affect-reasoning engine.

After entering in a contribution to an ongoing discourse, a user can select parts of the contribution over which she or he would like to express affect. This is achieved by highlighting the text concerned and then clicking on the bubble tip. This event then triggers the individual to the next step of expressing affect. The selected text is run through the affect-reasoning engine and the results presented to the user. Should the user still decide to add affect that may or may not be similar to results generated by the engine, the concerned text will then be tagged with the affect chosen by the user.

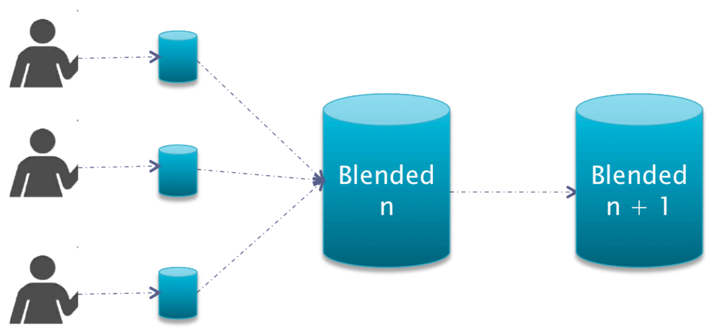

Fig 3: Leveraging personalizes spaces

Every contribution that is affectively tagged by the user is then collected as a part of the user's pattern of affect reasoning, along with what the engine also generated. These two patterns are then used to generate assertion using the machinery described in section 3 for the model to learn the user's patterns. Each user can affectively annotate any piece of text within a discourse, even from another user's contribution.

Upon other users hovering over a contribution in the discourse, the part of the text that was affectively annotated appears as a tooltip, thereby affording a feedback mechanism to infer the pragmatics of what the original author intended. We hypothesize that users with PLI will find both the expression and reception of affect to be helpful in inferring discourse pragmatics

4.2 Leveraging personalizes spaces

Since each user's pattern of affect annotation is collected to further bolster the model, this allows a model to be built for each user, specific to her or his usage patterns. Blending the models of each user together (clustered by some demographic criteria) means that the resultant model is more powerful than the previous one. New users can then be seeded with the latest model, and each merging of individual models results in a version that is more powerful than its predecessor.

5. Proposed evaluation and experiment protocol

The evaluation of this work can be seen under two kinds, namely an evaluation of the model and an evaluation of the user interface. The eventual goal of this work is to setup an experiment protocol for testing the hypotheses mentioned in this paper against groups representing typically developing versus individuals on the autism spectrum.

5.1 Evaluation of the model

Because the model has been trained for reasoning on affective valence, it becomes essential to validate test instances with a set of individuals (n >7) for annotation. After these set of test instances are annotated with affective valence, those whose kappa agreement values are greater than 0.5 can be used to run against the model.

Performance of the models can be ascertained using standard statistical metrics used in supervised learning, namely accuracy, precision & recall values, as well as Cohen's kappa values as a metric against agreement by chance.

It is important to keep in mind that this work is based on reinforcement learning, and the model's accuracy should improve with more users input. For a check on this aspect of the model, we propose the following experiment: have three models built with varying amounts of prior knowledge and run each of them against the evaluation described above. An analysis of each model's performance against varying amounts of blended knowledge will show the technique's improvement with more data. This is very intuitive in the sense that it is akin to adding more training data for a supervised learning method.

5.2 Evaluation of the user interface

The evaluation of the user interface is to be performed to ascertain two parameters, namely the ease and clarity of workflow as a user moves from one step in the annotation process to another, and an evaluation of the visualization of affect vis-à-vis how clearly affect is perceived.

For this part of the evaluation, we propose the setup of a user-evaluation of the interface with participants (n = 5) and a questionnaire about ease of workflow and perception of affect. This would allow for an iterative redesign of the interface if necessary.

Once both the model and the interface have been fully vetted, we now possess the machinery required to conduct an experiment involving typically developing and individuals on the autism spectrum.

5.3 Experiment protocol involving typically developing and individuals with autism

The first step is in the selection of participants. I am not clear on what the established template is to select participants who are 'deemed' to be typically developing. The second group of people will include people on the spectrum. Tests on which of them have PLI should ideally be conducted, but one wonders if questioning a participant if they have PLI like symptoms passes for a test for the purposes of this paper.

Participants can be recruited online since the entire infrastructure so far has already been moved to a server hosted here at the lab. The experiment protocol will involve the following:

a. Each group of users will be given a topic for discussion that tends to have a plethora of strong opinions. Each person in the group will be asked to play a role in the discourse of the topics. An example from our class is that Javier and I were students who liked the class, but there was also a person who didn't like it.
b. Each participant in the experiment will be subjected to a questionnaire at the end of the study, specific to his or her role in the discourse and test the effectiveness of the expression and reception of affect.

Our main hypothesis is that people on the spectrum who would otherwise suffer from PLI would find the tool to be more useful in understanding and modulating their inputs into an ongoing discourse. It is important to note that we make no such hypothesis on the typically developing group. It remains to be seen if even people who don't have PLI find the tool to be useful – in which case we are armed with a bountiful of ideas to talk to companies like Google, Facebook, Microsoft etc.

6. Conclusion and future work

In this work, we build an affect-reasoning engine based on techniques from natural language processing and commonsense reasoning. We also design an interaction sequence that affords affect expression and reception, with the added benefit of fine-tuning the affect-recognition engine.

We propose a robust evaluation of all aspects of this work, from the model to an evaluation of the interaction design. We propose an experiment protocol to study the effectiveness of our tool against two groups of people, namely a typically developing group and an ASD group with individuals who experience PLI. Our main hypothesis is that the ASD group with PLI would find the tool to be very useful in contributing to their understanding of the pragmatics of discourse under the context of online social interaction.

A long list of ideas that could be candidates for future work includes, training of a log-linear maximum entropy model that has HTMM, commonsense reasoning, as well the features of discourse analysis as features, which in many ways is the holy grail of addressing problems of pragmatic discourse understanding and cyberbullying.

This work is futile if all it does is ends up in annals of an ACM conference or journal. There are many opportunities for developing applications for storied understanding, next generation text editors, as well as communication interfaces for people both on and off the autism spectrum.

Acknowledgements

I'd like to express my thanks to Professors Picard and Goodwin for taking the time to meet with me and giving me valuable inputs and advice during the entire trajectory of this class. I'd also like to thank all the people who took this class with me for their valuable ideas, and most of all for their company.

References

[Bishop] Bishop DVM (2000). Pragmatic language impairment: A correlate of SLI, a distinct subgroup, or part of the autistic continuum? In DVM Bishop & LB Leonard (Eds.), *Speech and language impairments in children: Causes, characteristics, intervention and outcome* (pp. 99–113). Hove, UK: Psychology Press.

[Burke] Burke, M., Kraut, R., and Williams, D. (2010). Social use of computer-mediated communication by adults on the autism spectrum. ACM CSCW 2010: Conference on Computer-Supported Cooperative Work, 425-434

[Liu] Hugo Liu, Henry Lieberman, and Ted Selker. 2003. A model of textual affect sensing using real-world knowledge. In *Proceedings of the 8th international conference on Intelligent user interfaces* (IUI '03). ACM, New York, NY, USA, 125-132. DOI=10.1145/604045.604067 http://doi.acm.org/10.1145/604045.604067

[Speer] Robert Speer, Catherine Havasi, and Henry Lieberman. *AnalogySpace: Reducing the Dimensionality of Common Sense Knowledge*, AAAI 2008