IMHO: An Exploratory Study of Hedging in Web Forums

Liliana Mamani Sanchez
Trinity College Dublin
mamanisl@scss.tcd.ie

Carl Vogel
Trinity College Dublin
vogel@tcd.ie

Abstract

We explore hedging in web forum conversations, which is interestingly different to hedging in academic articles, the main focus of recent automatic approaches to hedge detection. One of our main results is that forum posts using hedges are more likely to get high ratings of their usefulness. We also make a case for focusing annotation efforts on hedges that take the form of first-person epistemic phrases.

1 Introduction

Computational linguistics research in hedging, use of linguistic expressions whose contribution to sentence meaning is a modulation of the accuracy of the content they embed, and speculation detection has been done intensively in the domain of scholarly texts. The interest created in this domain has expanded to some extent to other domains such as news and reviews. Automatic processing of speculation requires at some stage the annotation of words or phrases conveying uncertainty (Vincze et al., 2008). More complex endeavours imply the annotation of various elements of context involved in the expression of hedging (Rubin et al., 2005; Wiebe et al., 2005).

In web forums where users' contributions play a vital role in the forum dynamics, such as mutual support forums that are part of the ecosystem of technology company supports for users, exploring the features that make a contributor outstanding is relevant.[1] A user shows a distinctive behavior by writing useful posts that help other users in the problem that first motivated their participation in the forum. This paper emerges from our interest in finding features that predict which contributors will be most appreciated.

Many lexical and grammatical devices aid hedging (expressions such as epistemics verbs, modals, adjectives, etc. name but a few) as do non-lexical devices such as conditionals. We deem singular first person epistemic phrases as hedges that can help to identify the subject of a hedging event. We analyze the correlation between the use of epistemic phrases (vs. other types of hedges) and the probability of posts containing these hedges of being considered useful by the forum community. We also explore whether epistemic phrases constitute a distinctive feature that support user classifications. In §2, we described the function of hedges according to a hedging classification framework and in relation to the domain of web forums. Then §3 describes the profiling work done and discusses the main findings. We conclude in §4.

2 Functions of hedging

The research by Hyland (1998) is one of the broadest studies about hedging functions in scientific articles, and which makes use of categories that have strong relationship, at face value, to the likelihood that the reader of hedged material will find the material sufficiently useful or sufficiently well expressed to prompt the reader to rate highly the message containing the material, whether with an explicit facility to record kudos or otherwise. Hyland proposed a poly-pragmatic classification of hedges based on their indicating function: reader-oriented, writer-oriented, attribute and reliability. Briefly, attribute and reliability hedges both relate to the accuracy of the message conveyed. Attribute hedges relate to the conformity of the described situation with encyclopedic expectations (1), while reliability hedges relate to the level of certainty of the speaker about the propositional content (2). In a different dimension, reader ori-

[1] Throughout, we use "web forum" to refer to such ecosystems: we speculate that their informal nature makes our observations generalize to other sorts of web forum in which solutions to problems are not the focal point; even general discussion forums can be witnessed to trigger community weighting of contributions.

Proceedings of the SIGDIAL 2013 Conference, pages 309–313,
Metz, France, 22-24 August 2013. ©2013 Association for Computational Linguistics

ented hedges are composed with the concern that the "reader" accept the truth of the embedded content (3), thereby presupposing the "writer's" commitment to the content, while writer oriented hedges disclaim commitment to the content (4).

(1) **Protypical** mammals are land-dwellers.

(2) **Probably,** respected ancient Greeks thought whales to be fish.

(3) **I think** that if you reboot, the changes will take effect.

(4) **Based on** what you've said, you seem right.

Applying this classification scheme not to scholarly prose but to web forums, it seems likely that readers in technical forums would prefer the accuracy of attribute hedges (1) over the relative uncertainty of reliability hedges (2), and that the reader oriented hedges (3) supply comfort in the implication of both the quality of the embedded claims and the absence of arrogance. This research is attempting to test these hypotheses by assessing the relationship between the likelihood of posts receiving kudos and the quantity of hedges in these categories that the posts contain.

Unfortunately, answering the question is complex, because it is not in all cases obvious whether a linguistic expression contains a hedge or what function the hedges serve when they do exist. Therefore, we attempt a partial answer to the question by examining those hedge expressions which can be processed with some reliability using automated means. Consider the taxonomy of linguistic expressions in Fig. 1. The boxed regions of this taxonomy are amenable to automatic processing. Further, epistemic hedges with first-person singular subjects relate strongly to reader oriented hedges (3) in Hyland's taxonomy. The non-phrasal hedges are heterogeneous in function.

Figure 1: A taxonomy of linguistic expressions.

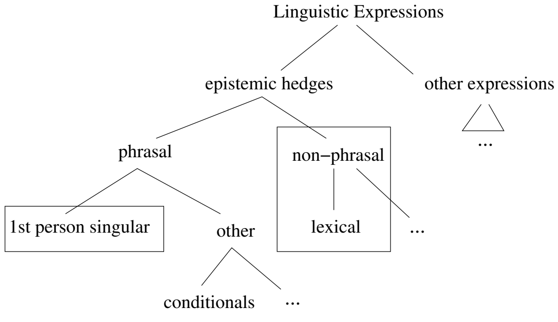

We do not claim this separation of hedging markers can fully account for pragmatic and semantic analysis of hedging in web forums, but we are confident this classification supports reliable annotation for quantificational assessment of certainty and hedging in this informal domain. We base our profiling experiments (§3) on this functional separation of hedging markers.

3 Profiling posts by hedging

3.1 Description of the forum dataset

The dataset we used created out of a forum that is part of customer support services provided by a software vendor company. Although we were not able to confirm the forum demographics, we can infer they are mostly American English speakers as the forum was set up first for USA customers. Some other features are best described by Vogel and Mamani Sanchez (2013). Our dataset is composed of 172,253 posts that yield a total of 1,044,263 sentences. This dataset has been intensively "cleaned", as originally it presented a great variety of non-linguistic items such as HTML codes for URLS, emoticons, IP addresses, etc. These elements were replaced by wild-cards and also user names have been anonymised, although some non-language content may remain.

A forum user can give a post "kudos" if he/she finds it useful or relevant to the topic being addressed in a forum conversation.[2] We counted the number of kudos given to each post. There are four user categories in the forum: {employee, guru, notranked, ranked}.[3] A poster's rank depends, among other factors, on the number of posts they make and their aggregate kudos.

3.2 Epistemic phrases versus other hedges

We created two lexicons, one composed by first person singular epistemic phrases and one by non-phrasal hedges. Initially, a set of epistemic phrases where taken from Kärkkäinen (2010): {I think, I don't know, I know, etc.} and from Wierzbicka (2006). The non-phrasal hedge lexicon was created from words conveying at least some degree of uncertainty: {appear, seem, sometimes, suggest, unclear, think, etc.}, taken from Rubin (2006). Additional hedges were included after the pilot

[2] A user may accord kudos for any reason at all, in fact.

[3] In the forum we studied, there are actually many ranks, with guru as the pinnacle for a non-employee; we grouped the non-guru ranked posters together.

annotation. The lexicons are composed by 76 and 109 items, respectively. There are many other hedge instances that are not included in these lexicons but our experiment restricts to these items. Epistemic phrases include acronyms such as "IMHO", "IMO" and "AFAIK" that we deem meet functions described in §2.

A pilot manual annotation of hedges was conducted on in order to verify the viability of automatic annotation. Our automatic annotation procedure performs a sentence by sentence matching and tagging of both kinds of hedging. The procedure uses a maximal matching strategy to tag hedges, e.g. if "I would suggest" is found, this is tagged and not "suggest". This automatic tagging procedure does not account for distinctions between epistemic and deontic readings of hedges, nor between speculative or non-speculative uses of non-phrasal hedges. 107,134 posts contain at least one hedge: 34,301 posts contain at least one epistemic phrase; 101,086, at least one non-phrasal hedge; 28,253, at least one of each.

3.3 Methods of analysis

In §3.1 we showed there are two ways to characterize a post: 1) By its writer category and 2) by the number of times it gets accorded *kudos*. We devise a third characterisation by exploring epistemic phrases and non-phrasal hedge usage in individual posts as a whole, tracking use of both types of hedge in each post. We devised three discretization functions (DF) for assigning a label to each post depending on the type of hedges contained within. The DFs take two parameters, each one representing either the relative or binarized frequency non-phrasal hedges and epistemic phrases ($nphr$ or $epphr$). DF1 relies on the occurrence of either type of hedge; a post is of a mixed nature if it has at least one of each hedge type. DF2 is based on a majority decision depending on the hedge type that governs the post and only assigns the label `hedgmixed` when both types of hedges appear in the same magnitude. DF3 expands DF1 and DF2 by evaluating whether either majority or only one type of hedge is found, e.g. we wanted to explore the fact that even when non-phrasal hedges domain one post, an epistemic phrase is contained as well, in contrast to when only non-phrasal hedges occur in a post.

DF1		$epphr$==0	$epphr>0$
	$nphr$ ==0	nohedges	epphrasal
	$nphr>0$	nonphrasal	hedgmixed

DF2		
	$nphr$=0 & $epphr$=0	nohedges
	$nphr > epphr$	nonphrasal
	$nphr < epphr$	epphrasal
	$nphr = epphr$	hedgmixed

DF3		$epphr$=0	$epphr>0$	
	$nphr$=0	nohedges	epphronly	
	$nphr>0$	nonphronly	$nphr > epphr$	nonphrmostly
			$nphr < epphr$	epphrmostly
			$nphr = epphr$	hedgmixed

We computed four measures for each post based on these functions, $m1$ is calculated by using $DF1$ having raw frequencies of hedges as parameters, $m2$ and $m3$ result from applying $DF3$ and $DF2$ respectively to frequencies of hedge type averaged by the corresponding lexicon size, and $m4$ is calculated from $DF3$ over hedge frequencies averaged by post word count. Other measures are also possible, but these seemed most intuitive.

We were interested in the extent that hedge-based post categories correlate with a post's kudos and with a post's user category as tests of hypothesis outlined in §2. We want to know which correlations hold regardless of the choice of intuitive measure and which are measure dependent.

3.4 Results and discussion

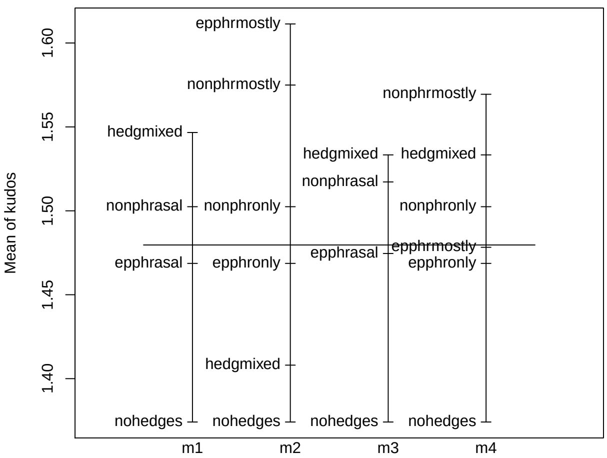

Figure 2: Design plot with the mean of kudos of each kind of post per each measure.

In Fig. 2, we show how the different hedge-based classifications of posts ($m1$, $m2$, $m3$, $m4$) relate to the average kudo counts for posts. Each measure is shown in an individual scale.[4] The horizontal line represents the average of kudos for all posts so we can observe which categories are above/below the mean. Comparison and contrast

[4]For this comparison, we dropped extreme outliers in the number of kudos and hedges, and we calculated these measures only in posts that had at least one kudo attribution.

of the relationship between categorisation of posts with each m^i and mean kudos is interesting. For example, when epistemic phrases dominate a post (`epphrmostly`), there is the greatest mean of kudos visible with the measure $m2$. The second highest positive effect is of non-phrasal hedges dominating a post (`nonphrmostly`) in $m2$ and $m4$. The next strongest effect occurs when both of hedges types appear in a post (`hedgmixed` in $m1$ and $m3$) and when they have about the same average density ($m4$), followed by when non-phrasal hedges appear exclusively in a post. While there is no consensus across the different scales that epistemic phrase-dominated posts are the most likely to obtain kudos, still their occurrence has a positive effect in the average of kudos obtained. There is low probability of kudos when only epistemic phrases appear and the lowest probability when no hedge occurs.[5] Thus, we argue that the four measures are jointly and individually useful.

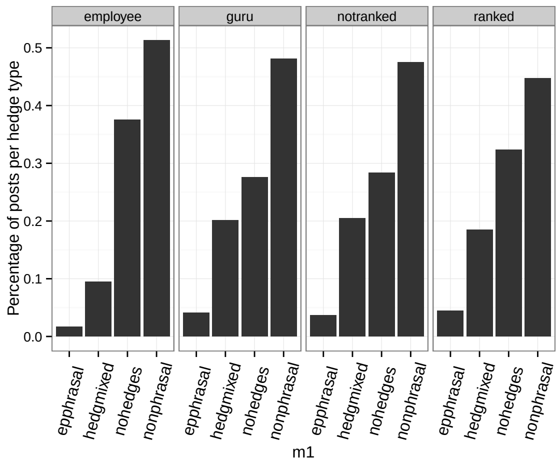

Figure 3: Percentages of $m1$-hedge types in each user category.

The relationship between hedge use and user category is depicted (for $m1$) in Fig. 3. While for all four user roles, epistemic phrases are exclusively present in the lowest percentage of posts, their contribution is shown in posts with mixed hedge types. Posts with only non-phrasal hedges are the most frequent across all user categories. We had predicted no significance in this respect since non-phrasal hedges could map into any of Hyland's functions, however our intuition was wrong as there is a significant difference ($p < 0.05$) in the proportions of posts per hedge type category when making comparisons across user categories one to one. Only when comparing proportions of hedge type posts by gurus and notranked users is there no significant difference in `hedgmixed`, `nonphrasal` and `nohedges` posts.[6] Employees and ranked users have the highest rates of use of mixed hedges. Ranked and guru posts have the highest ratios of exclusively epistemic phrase hedges, meeting expectations. Employees have the lowest ratio of user of epistemic phrases on their own, this presumably since they frequently write posts on behalf of the company so they are least likely to make subjective comments: their posts have the lowest percentage of use of "I" and "my".

These two approaches to assessing associations between different classifications of forum posts reveal that posts using hedges are the most likely to be accorded kudos and that guru and ranked users are the most frequent users of epistemic phrases in general. This lends support to the view that first person singular epistemic phrases, the epitome of reader-oriented hedges, are predictive of coarse grained rank in the forum.

4 Conclusions and future work

We have found that the hedges used contribute to the probability of a post getting high ratings. Posts with no hedges are the ones awarded least kudos. We have still to test the correlation between epistemic phrases and other types of hedges when they both are found in a single post. We think that automatic methods should focus in first person epistemic phrases as they show writer's stance at the same time as softening their commitment or anticipating reader's response. Following the annotation described here, manual annotation work is under way, where epistemic phrases and non-phrasal hedges constitute two distinct categories. Our ongoing work seeks other ways to measure the contribution of these categories to reader expression of appreciation of posts and whether hedge usage creates natural user categorizations. We also study other types of web forum dialogue to explore whether hedging follows similar trends.

[5]The contribution of epistemic phrases to the likelihood of kudos could be due to other factors such as the use of first person in general. We profiled the use of pronouns "I" and "my" and we found a negative correlation between frequency of these pronouns and the number of kudos per post. There is a small but not significant correlation restricting to those posts with non-zero kudos.

[6]A two-sample test of proportions was used to test the significance of differences between amounts of hedge type posts for each category.

Acknowledgements

This research is supported by the Trinity College Research Scholarship Program and the Science Foundation Ireland (Grant 07/CE/I1142) as part of the Centre for Next Generation Localisation (www.cngl.ie) at Trinity College Dublin. This paper has benefited from input from the anonymous reviewers and from Ron Artstein.

References

Ken Hyland. 1998. *Hedging in Scientific Research Articles*. Pragmatics & beyond. John Benjamins Publishing Company.

Elise Kärkkäinen. 2010. Position and scope of epistemic phrases in planned and unplanned american english. In *New approaches to hedging*, pages 207–241. Elsevier, Amsterdam.

Victoria Rubin, Elizabeth Liddy, and N. Kando. 2005. Certainty identification in texts: Categorization model and manual tagging results. In James G. Shanahan, Yan Qu, and Janyce Wiebe, editors, *Computing Attitude and Affect in Text*. Springer.

Victoria L. Rubin. 2006. *Identifying Certainty in Texts*. Ph.D. thesis, Syracuse University, Syracuse, NY.

Veronika Vincze, Gyorgy Szarvas, Richard Farkas, Gyorgy Mora, and Janos Csirik. 2008. The BioScope corpus: biomedical texts annotated for uncertainty, negation and their scopes. *BMC Bioinformatics*, 9(Suppl 11).

Carl Vogel and Liliana Mamani Sanchez. 2013. Epistemic signals and emoticons affect kudos. *Cognitive Infocommunications (CogInfoCom), 2012 IEEE 3rd International Conference on*, pages 517–522.

Janyce Wiebe, Theresa Wilson, and Claire Cardie. 2005. Annotating expressions of opinions and emotions in language ANN. *Language Resources and Evaluation*, 39(2/3):164–210.

A. Wierzbicka. 2006. *English: meaning and culture*. Oxford University Press, USA.