

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Communicative need shapes choices to use gendered vs. gender-neutral kinship terms across online communities

Permalink

<https://escholarship.org/uc/item/5kj3731s>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

Authors

Watson, Julia
Walker, Sarah
Stevenson, Suzanne
[et al.](#)

Publication Date

2023

Peer reviewed

Communicative need shapes choices to use gendered vs. gender-neutral kinship terms across online communities

Julia Watson¹

Sarah Walker¹

Suzanne Stevenson¹

Barend Beekhuizen²

¹Department of Computer Science
University of Toronto

{jwatson, swalker, suzanne}@cs.toronto.edu

²Department of Language Studies
University of Toronto, Mississauga

barend.beekhuizen@utoronto.ca

Abstract

Work has shown that greater need to refer to a semantic domain drives greater lexical precision within that domain, both across languages, and in lexical choice in (within-language) dyadic interactions. We complement this, studying the relation between communicative need and precision across communities of speakers of the same language. Taking kinship as our domain, we find evidence that differences in communicative need between communities within a language contribute to variation in lexical precision in use. We show that this variation is partly due to differences in the kinds of pragmatic contexts the communities talk about, but that community variation in lexical precision exists over and above the factors we study, suggesting that more work is needed to elucidate additional pragmatic influences on the simplicity–informativity trade-off.

Keywords: communicative need; lexical precision; lexical choice; kinship terms; pragmatic factors; within-language variation

Introduction

The way different languages group together or distinguish meanings through their words has been fruitfully studied by research on communicative efficiency, which models how lexical semantic systems optimize a trade-off between simplicity and informativity. One important factor that shapes this trade-off is communicative need, with increased need to refer to a semantic domain leading to a more complex but more informative system, with terms of higher precision. Much work has studied how communicative need contributes to this pattern of cross-linguistic variation in lexical systems (e.g., Khetarpal et al., 2013; Regier et al., 2016; Gibson et al., 2017; Kemp et al., 2018; Zaslavsky et al., 2020; Bradford et al., 2022). Across languages, for example, terms for siblings, which are more frequently referred to than cousins, tend to be more precise than terms for cousins, by specifying gender and/or relative age (Kemp & Regier, 2012).

In addition to work studying how communicative need shapes variation in the words *available* to speakers, experimental work on dyad interactions has studied the relationship between communicative need and *choices to use* more and less lexically-precise forms within a domain. Work in this area has shown that such choices are shaped by a need to disambiguate the intended referent from others in context; for example, speakers are more likely to call a referent *the dog* when there are no other dogs in the contrast set, but *the dalmatian* when the set includes a greyhound (Graf et al., 2016; cf. Xu & Tenenbaum, 2007). A greater need to referentially

distinguish between finer-grained sets of referents has also been shown to lead to the emergence of more precise term systems in artificial language learning experiments (Hawkins et al., 2018; Karjus et al., 2021).

Both kinds of work above – those studies examining the level of precision of a lexical semantic system, and those studying the varying use of more and less precise terms within such a system – have generally conceptualized “informativity” as the ability to accurately convey the speaker’s meaning by identifying the object of reference. However, pragmatic factors beyond referential accuracy are also known to affect lexical precision (Cruse, 1977; Wisniewski & Murphy, 1989), and more generally may affect the make-up of term systems (Enfield, 2014; Kemp et al., 2018). For example, the factors shaping the choice between *daughter* and *child* in responding to the question *who are you picking up from school?* (referring to a particular person) may be different from those shaping the choice between the same words in discussing parenting practices generally (e.g., *if I had a daughter/child...*, referring to a hypothetical person having a certain relation to the speaker). Here, we study the relationship between communicative need and lexical precision ‘in the wild’ – in online communities – which enables us to examine some relevant pragmatic contexts that shape the choice of more or less precise language.¹

To do this, we evaluate how communities of speakers vary in communicative need for a lexical domain, as well as in their choices of more vs. less lexically-precise forms within that domain. This within-language, community-level approach complements work both at the cross-linguistic level and the dyadic interaction level, demonstrating that the connection between communicative need and lexical precision is similarly relevant in understanding variation in lexical choice across communities (cf. Sun & Xu, 2022).

We use the domain of kinship terms for our exploratory case study of community variation in need and precision, because it enables us to build on both sets of work mentioned above. Studies have shown the interaction of communicative need and the make-up of language-specific lexicons for kinship (e.g., Kemp & Regier, 2012). However, such crosslinguistic work, in focusing on mostly mutually ex-

¹Other work on the simplicity–informativity trade-off and pragmatic contexts (Denić et al., 2022, using the corpus of Beekhuizen et al., 2017) has not considered lexical choice across communities.

clusive terms, has not considered the role of lexical choice, although it is prevalent across languages for a domain like kinship:² English, for instance, has choices such as *parent* or *mother/father*, and *child* or *daughter/son*, where the gendered terms are referentially and conceptually more precise.³ While experimental work on lexical precision within a language has studied lexical choice (e.g., Graf et al., 2016), it has not considered factors beyond the paradigm of establishing accurate reference. Kinship is an ideal domain for studying additional pragmatic effects because kinship terms convey social role information, so their dimensions (e.g., gender, relative age) likely have strong pragmatic implications related to these roles. Pragmatic factors are thus likely important both in the way the kinship domain is structured and in how such terms are deployed (cf. Stokoe & Attenborough, 2014).

In investigating this domain, we provide evidence that, within the same language, between-community differences in communicative need explain between-community differences in lexical precision of kinship terms. We further find both that communities vary in pragmatic contexts that talk about kinship, and that this variation helps explain the differences in lexical precision, confirming that explanatory models of lexical semantic systems need to take such factors into account. Finally, we also find that community variation in precision of usage exists over and above the factors we study, suggesting that more work is needed to elucidate additional pragmatic influences on the simplicity–informativity trade-off.

Overview of Approach

To study community variation in the relation between need and lexical precision, we examine the use of more precise (gendered; e.g., *daughter*) and less precise (gender-neutral; e.g., *child*) kinship terms in English, across the online social media platform Reddit. This focus not only provides plentiful data, but data that is separated into identifiable communities (subreddits) distinct in topic or focus (viewing subreddits as “online communities of practice”, Del Tredici & Fernández, 2017). We select pairs of such communities for which we have hypotheses about their difference in need to talk about kinship relations, and assess whether communities that have higher need for kinship terms indeed exhibit greater lexical precision in their usage – that is, have more usages of gendered (over gender-neutral) kinship terms. For space reasons here, we discuss two such pairs of contrasting subreddits:

- AskReddit/AskScience: Both subreddits are broad communities for asking and answering questions. The general

²The corpus work of Bradford et al. (2022) includes variant terms in the calculation of communicative need, but does not consider the impact of such variants on the assessment of lexical precision.

³Other factors beyond lexical precision are known to play a role in the lexical choice between gendered and gender-neutral forms, including, importantly, the cultural shift towards the use of gender-neutral forms to refer to nonbinary or gender-nonconforming individuals (e.g., *they/them* pronouns). Given this cultural shift, some emerging uses of gender-neutral kinship terms may be communicating gender that is not male or female. However, gender-neutral terms remain widely used for reference that is unspecified for gender.

nature of the topics in AskReddit, and the occurrence of many personal stories (in contrast to AskScience), is expected to create *greater communicative need for kinship terms generally* in AskReddit.

- Parenting/EntitledParents: These subreddits are both about parent–child relations, but the former focuses on child-rearing, and the latter on stories about parents. The focus of each of these is expected to create *greater communicative need for particular kinship terms*: child kinship terms in Parenting, and parent kinship terms in EntitledParents.

In Part 1 of the Results below, we first assess, for each pair of subreddits: (a) if their communicative need indeed differs, and (b) if, as predicted, subreddits with greater communicative need use more precise (gendered) kinship terms relative to more general (gender-neutral) terms. To preview our results, we find that greater communicative need is associated with greater lexical precision.

In Parts 2 and 3, we evaluate whether these differences in lexical precision arise because communities talk about kinship relations in different contexts, or because communities make different lexical choices with kinship terms in similar contexts. To do this, we operationalize context as local linguistic properties of comments containing kinship terms, and we examine two types of contexts that may contribute to the community variation we observe in lexical choice. Our goal is to consider discourse/pragmatic factors in lexical choice that go beyond the establishment of accurate reference.

In Part 2, we consider grammatical contexts that are associated with usages of more precise vs. general terms. Generic and plural reference have both been found to coincide with less precise terms in taxonomically-organized lexical domains (Cruse, 1977; Wisniewski & Murphy, 1989). Because gender-neutral kinship terms are less lexically precise than gendered terms, we might expect community variation in use of such contexts to contribute to variation in lexical precision.

In Part 3, we turn to another factor that may influence community variation in English kinship terms. In the kinship domain, “precision” is equated with gender – that is, more precise terms are those with gender encoded in the term. As a highly socially implicative dimension, expectations around gender are expected to affect lexical choice (Stokoe & Attenborough, 2014). If a subreddit talks about kinship relations in contexts that are “more gendered”, this could lead to usage (in English) of more precise kinship terms.

In both Parts 2 and 3, we seek to determine whether these identified factors contribute to the observed patterns in lexical precision – shedding light on whether the different communities talk about kinship in different contexts. Moreover, we test whether subreddit differences in lexical precision persist even when these factors are considered – that is, revealing whether communities make different lexical choices even within the same referential/pragmatic contexts.

While Part 1 of the Results shows that communities that differ in communicative need do indeed differ in their lex-

Kinship term group	Terms in each group
CHILD	<i>child, kid; daughter, son</i>
PARENT	<i>parent; dad, father, mom, mother, mum</i>
PARTNER	<i>partner, s/o, significant other, spouse; bf, boyfriend, gf, girlfriend, husband, wife</i>
SIBLING	<i>sibling; brother, sister</i>

Table 1: Kinship terms used. Bold=more precise (gendered).

ical precision, Parts 2 and 3 together show how referential and discursive functions of lexical items, beyond establishing accurate reference, may vary and may impact the relation between communicative need and precision of lexical choices.

Data and Methods

We make all data and code available at <https://github.com/juliawatson/kinship-reddit-2023>.

Kinship dataset from the Reddit platform We collected data from Reddit using the PushShift API.⁴ To construct our kinship dataset, we extracted comments containing kinship terms with both gendered and gender-neutral variants, including terms for children, parents, siblings, grandchildren, grandparents, niblings, and partners.⁵ From each subreddit under consideration, we collected 100,000 comments containing any of these kinship terms, forming our kinship dataset.⁶ We focus here on the kinship term groups CHILD, PARENT, PARTNER, and SIBLING, since terms for these relations are most frequent in our data; the specific terms we consider in each group are shown in Table 1.

Assessing communicative need as frequency Following much other work (e.g., Kemp & Regier, 2012; Bradford et al., 2022), we operationalize communicative need in terms of the relative frequency of kinship terms in a community. To compute a reliable estimate, we downloaded the most recent 10M words of text from each subreddit and calculated the summed frequency of all the terms in each kinship term group (in words per million), as the estimated communicative need for that kinship relation.

Contextual factors: grammatical contexts As noted, previous work has found that certain discourse/pragmatic factors shape choices between overlapping lexical alternatives; in particular, Wisniewski & Murphy (1989) found that more general terms like *animal* are used more for plural and generic referents, while more lexically-precise terms like *elephant* are used more for specific individuals. In line with this, we may expect the more fine-grained (gendered) kinship terms to be used more in singular and/or specific contexts. We thus consider number (singular or plural) and kind of ref-

erence (specific or generic) as relevant contextual factors in assessing patterns of lexical precision; we refer to these discourse/pragmatic features as “grammatical contexts.” We automatically coded all the contexts of kinship term usages in our dataset for these dimensions. While it is straightforward to automatically identify singular vs. plural, distinguishing specific vs. generic usages is more difficult. We leverage the observation that kinship terms preceded by *my* are very likely to be specific references (cf. Kemp & Regier, 2012), and (in our data) terms preceded by *'s* (e.g., *OP's girlfriend is correct*) are also very likely to be specific. These were automatically labelled as specific, while the remaining usages (which are less likely to be specific) are labelled “other”.⁷ Examples of each usage type are shown in Table 2.

Contextual factors: ‘genderedness’ of contexts For kinship in English, more precise terms are those that specify the gender of the referent. Hence, another contextual factor that may shape usage with respect to lexical precision is the degree to which gender is relevant or highlighted. For example, in *No thief dates my daughter!*, the mention of the activity of dating, where gender may be relevant, may lead to a higher expectation of a gendered term. Here, we operationalize the genderedness of a context as the degree to which a large language model expects a gendered kinship term over a gender-neutral one. We use predictions from the contextual language model BERT (Devlin et al., 2019), which is sensitive to some of the factors shaping acceptability of gendered vs. gender-neutral language (Baumler & Rudinger, 2022). Although BERT may be more conservative than people in such judgments (Brandl et al., 2022), we expect its predictions will capture some relevant contextual information.

We apply BERT to each context of a kinship term usage in our dataset, using a ‘fill-in-the-blank’ task, in which BERT outputs a probability distribution over the vocabulary for the blank, or “masked out” word. For the example above, we would feed *No thief dates my [MASK]!* into BERT, and compute the probability of a gendered variant in this context as:

$$p(\textit{gendered}) = \frac{p(\textit{daughter}) + p(\textit{son})}{p(\textit{daughter}) + p(\textit{son}) + p(\textit{child}) + p(\textit{kid})}$$

(and analogously for usages from other kinship term groups).

Results

We recognize that statistical tests are often significant for small differences in corpus counts as large as ours (mostly in the 1000s, 10,000s, or even 100,000s). With this in mind, we report p-values (which are highly significant, with $p \ll 0.0001$ for all but one test, as noted below)⁸, but we also present data that illustrate the magnitude of effect sizes. We

⁷We validated this approach by manually coding a random sample of kinship terms in our dataset as specific vs. generic. Note that other possessive determiners like *your* were less likely than *my* to be used in specific contexts (e.g., *Well when someone is always trying to tell you how to raise your kid it gets to you*).

⁸Using a simple Bonferroni correction, results are significant at the .05 level if $p < 0.0014$ (for 36 total tests).

⁴<https://api.pushshift.io/reddit/>

⁵Like consanguineal terms, partner terms describe a social role (cf. Service, 1960), and we expect use of their gendered vs. gender-neutral forms to vary in similar ways to other kinship terms.

⁶Filtering to remove non-kinship usages of terms (such as *child* meaning a young person vs. offspring) is not practical due to scale.

singular, specific	<i>I just thought every time I went to bed my daughter decided to rearrange the furniture.</i>
singular, other	<i>There is zero reason to bring a sibling along.</i>
plural, specific	<i>My parents are like this all the time</i>
plural, other	<i>It makes me so sad when parents put new partners before their kids.</i>

Table 2: Examples of the grammatical contexts in Part 2 (number [singular/plural] x specificity [specific/other]).

	CHILD	PARENT	PARTNER	SIBLING	CHILD	PARENT
AR	1525	1252	854	356	P	10715
AS	440	309	72	50	EP	5444
						7874

Table 3: Frequencies per million. The highest frequency among the two subreddits for each kinship group is in bold.

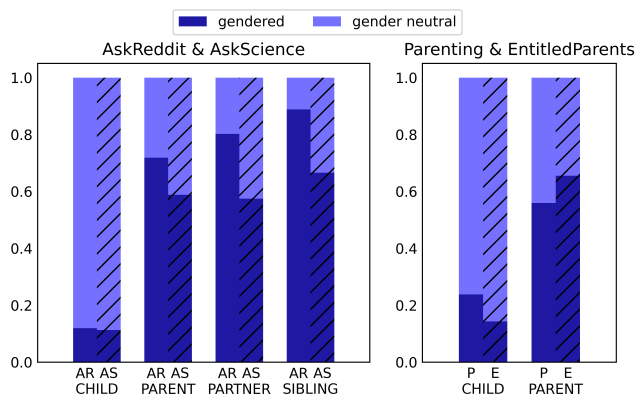


Figure 1: Rates of gendered vs. gender-neutral kinship terms.

abbreviate AskReddit (AR), AskScience (AS), Parenting (P), and EntitledParents (EP) in tables and figures.

Part 1: Do communities with greater communicative need show greater lexical precision?

We first test for differences in communicative need across each of the subreddit pairs. In line with both of our predictions, the between-community variation in communicative need is substantial, and in the expected direction: We find significantly greater rates of kinship term usages on AskReddit than on AskScience for all kinship groups, and we find significantly greater rates of CHILD terms on Parenting and of PARENT terms on EntitledParents (using χ^2 , $p \ll 0.0001$, $N = 20M$, i.e., 10M per subreddit). Table 3 reports the rates of use for each kinship term group.

Having established a difference in communicative need, we explore our central hypothesis – namely that, in cases where there is higher communicative need, there is also greater lexical precision (i.e., increased rate of gendered terms). To test for differences in rates of gendered kinship terms for each subreddit pair, we run logistic regressions in which the dependent variable is whether each usage (from a pair of subreddits) is a gendered term (1) or gender-neutral term (0), and the independent variable is subreddit (dummy coded). We run a regression for each kinship term group separately, since usage patterns may differ across them: for AskReddit and

AskScience, on each of the four kinship term groups; for Parenting and EntitledParents, on the CHILD and PARENT term groups, as these are the ones we have hypotheses about.

Table 4, top panel, shows the results. We find significantly higher rates of gendered terms on AskReddit than AskScience across all kinship groups (though the effect size for CHILD terms is small). For Parenting and EntitledParents, we find significantly higher rates of gendered CHILD terms on Parenting, and of gendered PARENT terms on EntitledParents. These differences in lexical precision for the two pairs of subreddits are shown in Figure 1. (The small effect size for CHILD terms on AskReddit vs. AskScience is apparent in the small difference in the bar plots for this group.)

Thus, for both subreddit pairs, the community with greater communicative need for a kinship term group uses more gendered/lexically-precise language to refer to that relation, supporting our hypothesis linking need to lexical precision in community usage.

Part 2: Do differences in grammatical contexts contribute to differences in lexical precision?

Here we evaluate how the discourse/pragmatic dimensions of number (singular and plural) and kind of reference (specific and other) may play a role in community variation of kinship term usage, and whether they influence lexical precision. (Refer back to Table 2 for examples of the four grammatical contexts – singular–specific, plural–specific, etc.)

We first examine whether the subreddit pairs do indeed vary in the rates they use these contexts (specific vs. other; singular vs. plural) for kinship terms. For each context type, we find substantial variation across subreddits for most kinship term groups considered, although in some cases the differences are highly significant but relatively small (using χ^2 , $p \ll 0.0001$, N ranges from 24,262 to 209,988); Table 5 reports the proportions of use. Figure 2 plots gendered vs. gender-neutral usages across the grammatical contexts for an example kinship term group for AskReddit/AskScience, and for both CHILD and PARENT terms in Parenting/EntitledParents. These results confirm that communities with differing need to discuss kinship relations also exhibit variation in **how** they talk about those relations.

We next assess whether the varying use of these contexts may be (at least partly) responsible for the differences – observed in Part 1 – in lexical precision across communities. To do so, we add the predictors of number and specificity (sum coded) to our logistic regressions; see Table 4, middle panel. We expect to find that singular number and specific reference predict increased rates of gendered terms, in line with find-

		AskReddit vs. AskScience				Parenting vs. EntitledParents	
		CHILD	PARENT	PARTNER	SIBLING	CHILD	PARENT
Part 1	intercept	-2.06	+0.36	+0.30	+0.69	-1.79	+0.64
	subreddit (AR=1; P=1)	+0.06	+0.58	+1.10	+1.38	+0.63	-0.40
Part 2	intercept	-1.46	+0.02	+0.19	+0.97	-1.88	+0.39
	subreddit (AR=1; P=1)	-0.13	+0.62	+0.89	+1.01	+0.42	-0.49
	number (singular=1)	+1.07	+1.67	+0.65	+1.16	+1.38	+1.85
	specificity (specific=1)	+1.11	+0.39	+0.65	+0.40	+0.86	+0.55
Part 3	intercept	-1.88		-1.86		-2.23	
	subreddit (AR=1; P=1)	-0.53		-0.33		+0.40	
	p(gendered) (continuous)	+4.17		+4.74		+3.51	

Table 4: β values from logistic regressions predicting gendered vs. gender-neutral usage for Part 1, Part 2, and Part 3 of Results. All predictors are significant with $p < 0.00001$, except for the subreddit predictor for CHILD terms on AskReddit/AskScience in the Part 1 regression, where $p = 0.0006$. (All intercepts are significant except for the Part 2 regression on AR/AS for PARENT.) The Part 1 and Part 2 regressions include data points from all 4 grammatical contexts, and the numbers of observations in these regressions range from 24,262 to 209,988. The Part 3 regressions include only specific singular usages, and the numbers of observations in these regressions range from 4,838 to 24,558.

	CHILD	PARENT	PARTNER	SIBLING		CHILD	PARENT
AR	0.47	0.71	0.93	0.77	P	0.57	0.66
AS	0.43	0.68	0.80	0.59	EP	0.55	0.71
AR	0.09	0.48	0.45	0.45	P	0.22	0.19
AS	0.05	0.16	0.33	0.20	EP	0.09	0.22

Table 5: Proportion singular usages (out of singular+plural; top panel) and specific usages (out of specific+other; bottom panel) per subreddit and kin group. The highest proportion among the two subreddits for each kinship group is in bold.

ings on how number and specificity shape lexical precision (e.g., Wisniewski & Murphy, 1989). We further predict – based on our overarching hypothesis of greater community need leading to greater lexical precision – that we will continue to find significant effects of the subreddit predictor, even after taking these referential-pragmatic factors into account.

As expected, we find significant, positive effects of singular and specific for both subreddit pairs, showing that speakers use more lexically-precise (gendered) terms in these contexts. This means that, not only do communities vary in the way they talk about kinship, but these differences also help explain observed variation in lexical precision. Moreover, effects of subreddit persist (though are attenuated) in the regressions, even after accounting for these factors. This supports the communicative need–precision connection, but suggests that additional dimensions of variation play a role as well.

Part 3: Do differences in genderedness of contexts contribute to differences in lexical precision?

Here, we assess, as with the grammatical contexts above, whether the use of gendered contexts varies across the communities, and whether genderedness of contexts further predicts lexical precision. In our analyses here, we zoom in on a

single type of context from Part 2, the specific-singular cases, to control for the effect of any interaction with such grammatical contexts.

Analogously to Part 2, we first examine whether the subreddit pairs vary in the genderedness of contexts used with kinship terms. The mean $p(\textit{gendered})$ values across the subreddit pairs for all term groups are significantly different (Mann Whitney U, $p \ll 0.0001$, N ranges from 4,838 to 24,558). Because the genderedness of specific singular contexts is near ceiling (over .98) for PARENTS and SIBLINGS in all subreddits, remaining analyses here focus solely on CHILD and PARTNER terms, for which differences in the mean of $p(\textit{gendered})$ are 2.5 – 3.5%. Figure 3 shows the kernel density estimation plots for CHILD terms in the two subreddit pairs, illustrating the varying use of gendered contexts.

Again as in Part 2, we next assess whether the varying use of gendered contexts may be (at least partly) responsible for the community differences in lexical precision. Here we add the predictor of $p(\textit{gendered})$ to our logistic regressions; see Table 4, bottom panel.⁹ For both subreddit pairs, as predicted, we find increased rates of gendered terms in contexts with higher $p(\textit{gendered})$. Moreover, for Parenting/EntitledParents, we again see a positive effect of subreddit: over and above the predictiveness of genderedness of contexts, we see that Parenting uses more precise (gendered) language when talking about children.

Interestingly, for AskReddit/AskScience, we also see an effect of subreddit, but here it is in the opposite direction to what we predicted: AskScience uses more gendered terms over and above that predicted by $p(\textit{gendered})$. Further analysis will be required to understand this reversal of effect. Qualitatively, we find that many examples on AskReddit use

⁹Because we only look at singular specific cases here, the number and specificity predictors are not included.

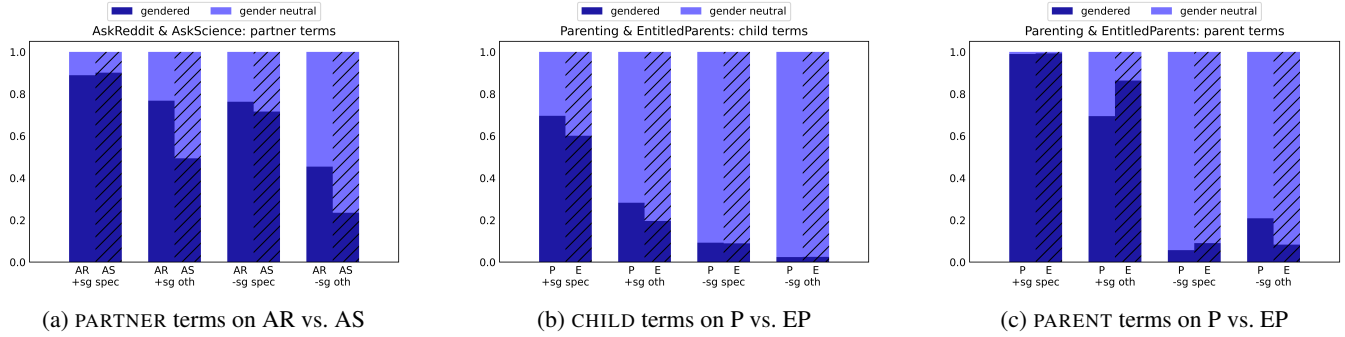


Figure 2: Rates of gendered vs. gender-neutral kinship terms across contexts in Part 2; +/-singular, spec(ific) vs. oth(er).

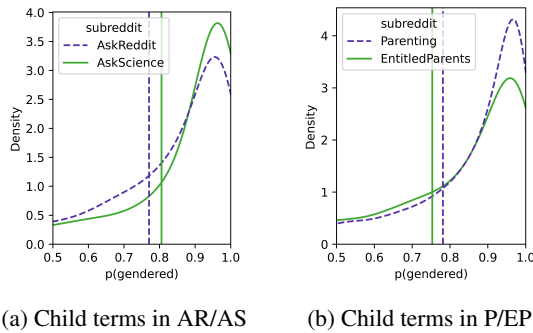


Figure 3: KDEs showing $p(\text{gendered})$ density across usages in the subreddit pairs. Vertical lines are at the means.

a gender-neutral form *child* to emphasize the relation of the parent to their child (e.g., *If you hurt my child, I will end you*), as opposed to focusing on the identity of the child, as is more common on AskScience (e.g., *My daughter is excited about the books*). This is suggestive of the rich contextual factors that can shape lexical choice in this domain, and will need to be explored in future work.

Thus, as in Part 2, we see that communities vary in their use of contexts (here, those with differing degrees of gendered expectations), and that such contexts are predictive of community variation in usage of lexically-precise terms. But also again, we see that the variation in lexical precision cannot be attributed solely to a difference in these contextual usages – in this case, differences in pragmatic expectations for gender of the kinship terms cannot alone account for the variation in lexical precision across communities.

Discussion

A novel contribution of this work is establishing the relationship between communicative need and lexical precision at the community level. We find that English-speaking online communities with greater need to talk about kinship relations use more lexically-precise kinship terms. Additionally, we find that online communities vary not only in how lexically precise they are, but also in the kinds of contexts they use in talking about kinship. Such findings at this “intermediate level” of a community help bridge work on communicative need and lexical precision at the cross-linguistic (e.g., Kemp et al., 2018) and dyadic within-language (e.g., Hawkins et al.,

2018) levels. Our large-scale corpus analysis of community-level variation has enabled us to evaluate the impact of pragmatic influences beyond the case of accurate/specific reference, which much previous work has focused on.

Further, because lexical categories evolve to meet the pragmatic needs of speakers who use them (Enfield, 2014), establishing community-level differences in the kinds of pragmatic contexts speakers talk about has implications for language change. Future research should work to further our understanding of this intermediate level by studying how differences in the pragmatic situations communities talk about might result in different trajectories in language change. Further work should also consider other cultural factors that may shape language change, such as a community’s utility for a given kind of information (e.g., a referent’s gender). While past work has shown that political views shape choices to use gendered vs. gender-neutral forms (Papineau et al., 2022), including over time (CH-Wang & Jurgens, 2021), to our knowledge, past work has not considered how this relates to communicative need and principles of communicative efficiency, which are implicated in language change.

A second contribution is our finding that the relationship between communicative need and lexical precision persists, even when taking into account community variation in the kinds of contexts discussed. This means that communities make different lexical choices in similar pragmatic contexts, which has important implications for how we view lexical concepts. Historically, referential applicability (“can you refer to individual X with this word?”) has been central to our understanding of word meanings. Our results suggest that discursive applicability – “when (in which community, in which discursive context) are you expected to use this word and not another, referentially applicable, word?” – is another facet of our knowledge of word meaning. The existence of community variation shows that knowing when a word is *normatively* used does not follow automatically from knowing when a word can *accurately* be used, and such knowledge must be acquired and represented separately (Goodwin, 1994). Future work will need to consider how referential and discursive applicability interact, and what it means for a lexical system to be efficient across a range of pragmatic functions and discursive contexts.

Acknowledgments

We acknowledge the support of NSERC of Canada (through grants RGPIN-2017-06506 to SS and RGPIN-2019-06917 to BB), as well as the support of the Data Sciences Institute, University of Toronto (through a Catalyst Grant to SS, BB, and JW).

References

- Baumler, C., & Rudinger, R. (2022). Recognition of they/them as singular personal pronouns in coreference resolution. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 3426–3432).
- Beekhuizen, B., Watson, J., & Stevenson, S. (2017). Semantic typology and parallel corpora: Something about indefinite pronouns. In *Proceedings of the 37th Annual Meeting of the Cognitive Science Society*.
- Bradford, L., Thomas, G., & Xu, Y. (2022). Communicative need modulates lexical precision across semantic domains: A domain-level account of efficient communication. In *Proceedings of the 44th Annual Meeting of the Cognitive Science Society*.
- Brandl, S., Cui, R., & Sjøgaard, A. (2022). How conservative are language models? adapting to the introduction of gender-neutral pronouns. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 3624–3630).
- CH-Wang, S., & Jurgens, D. (2021). Using sociolinguistic variables to reveal changing attitudes towards sexuality and gender. In *Proceedings of the 2021 conference on empirical methods in natural language processing*.
- Cruse, D. A. (1977). The pragmatics of lexical specificity. *Journal of linguistics*, 13(2), 153–164.
- Del Tredici, M., & Fernández, R. (2017). Semantic variation in online communities of practice. In *Proceedings of the 12th international conference on computational semantics*.
- Denić, M., Steinert-Threlkeld, S., & Szymanik, J. (2022). Indefinite pronouns optimize the simplicity/informativeness trade-off. *Cognitive Science*, 46(5), e13142.
- Devlin, J., Chang, M., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 4171–4186).
- Enfield, N. (2014). *The utility of meaning: What words mean and why*. Oxford: Oxford University Press.
- Gibson, E., Futrell, R., Jara-Ettinger, J., Mahowald, K., Bergen, L., Ratnasingam, S., ... Conway, B. R. (2017). Color naming across languages reflects color use. *PNAS*, 114(40), 10785–10790.
- Goodwin, C. (1994). Professional vision. *American anthropologist*, 96(3), 606–633.
- Graf, C., Degen, J., Hawkins, R. X., & Goodman, N. D. (2016). Animal, dog, or dalmatian? level of abstraction in nominal referring expressions. In *Proceedings of the 38th Annual Meeting of the Cognitive Science Society*.
- Hawkins, R. X., Franke, M., Smith, K., & Goodman, N. D. (2018). Emerging abstractions: Lexical conventions are shaped by communicative context. In *Proceedings of the 40th Annual Meeting of the Cognitive Science Society*.
- Karjus, A., Blythe, R. A., Kirby, S., Wang, T., & Smith, K. (2021). Conceptual similarity and communicative need shape colexification: An experimental study. *Cognitive Science*, 45(9), e13035.
- Kemp, C., & Regier, T. (2012). Kinship categories across languages reflect general communicative principles. *Science*, 336(6084), 1049–1054.
- Kemp, C., Xu, Y., & Regier, T. (2018). Semantic typology and efficient communication. *Annual Review of Linguistics*, 4(1), 109–128.
- Khetarpal, N., Neveu, G., Majid, A., Michael, L., & Regier, T. (2013). Spatial terms across languages support near-optimal communication: Evidence from peruvian amazonia, and computational analyses. In *Proceedings of the 35th Annual Meeting of the Cognitive Science Society* (pp. 764–769).
- Papineau, B., Podesva, R., & Degen, J. (2022). ‘sally the congressperson’: The role of individual ideology on the processing and production of english gender-neutral role nouns. In *Proceedings of the 44th Annual Meeting of the Cognitive Science Society*.
- Regier, T., Carstensen, A., & Kemp, C. (2016). Languages support efficient communication about the environment: Words for snow revisited. *PLoS one*, 11(4), e0151138.
- Service, E. R. (1960). Kinship terminology and evolution. *American Anthropologist*, 62(5), 747–763.
- Stokoe, E., & Attenborough, F. (2014). Gender and categorial systematics. *Handbook of language, gender and sexuality*, 161–179.
- Sun, Z., & Xu, Y. (2022). Tracing semantic variation in slang. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*.
- Wisniewski, E. J., & Murphy, G. L. (1989). Superordinate and basic category names in discourse: A textual analysis. *Discourse Processes*, 12(2), 245–261.
- Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review*, 114(2), 245–272.
- Zaslavsky, N., Kemp, C., Tishby, N., & Regier, T. (2020). Communicative need in colour naming. *Cognitive neuropsychology*, 37(5-6), 312–324.