



Cognitive Science 45 (2021) e13035

© 2021 Cognitive Science Society LLC

ISSN: 1551-6709 online

DOI: 10.1111/cogs.13035

Conceptual Similarity and Communicative Need Shape Colexification: An Experimental Study

Andres Karjus,^{a,b,c}  Richard A. Blythe,^{c,d} Simon Kirby,^c Tianyu Wang,
Kenny Smith^c

^a*ERA Chair for Cultural Data Analytics, Tallinn University*

^b*School of Humanities, Tallinn University*

^c*Centre for Language Evolution, School of Philosophy, Psychology and Language Sciences, University of Edinburgh*

^d*School of Physics and Astronomy, University of Edinburgh*

Received 25 March 2021; received in revised form 18 June 2021; accepted 19 July 2021

Abstract

Colexification refers to the phenomenon of multiple meanings sharing one word in a language. Cross-linguistic lexification patterns have been shown to be largely predictable, as similar concepts are often colexified. We test a recent claim that, beyond this general tendency, communicative needs play an important role in shaping colexification patterns. We approach this question by means of a series of human experiments, using an artificial language communication game paradigm. Our results across four experiments match the previous cross-linguistic findings: all other things being equal, speakers do prefer to colexify similar concepts. However, we also find evidence supporting the communicative need hypothesis: when faced with a frequent need to distinguish similar pairs of meanings, speakers adjust their colexification preferences to maintain communicative efficiency and avoid colexifying those similar meanings which need to be distinguished in communication. This research provides further evidence to support the argument that languages are shaped by the needs and preferences of their speakers.

Keywords: Colexification; Communicative need; Experimental; Artificial language; Complexity; Cognitive cost; Expressivity; Communicative cost

1. Introduction

When two or more functionally distinct senses are associated with a single lexical form in a given language, expressed by a single word, then these senses can be referred to as being

Correspondence should be sent to Andres Karjus, ERA Chair for Cultural Data Analytics, Tallinn University, 10120 Tallinn, Estonia. E-mail: andres.karjus@tlu.ee

colexified (François, 2008). Colexification is more general than homonymy and polysemy (which also refer to a multiplicity of senses for a word), making no assumptions about the nature of the relations between the senses. Recognizing that a word lexifies more than one sense requires a means to determine what the minimal units of meaning are. This would be difficult to do based on just one language, but can and has been done utilizing systematic cross-linguistic comparison. For example, English has separate monomorphemic words for *leg* and *foot*, while many other languages use a single word to refer to the whole thing, for example, Estonian (*jalg*), Irish Gaelic (*cos*), and Modern Hebrew (*regel*). This does not mean that referring to these concepts separately in these languages is impossible, but rather it may involve more complex compounds or expressions to describe them (e.g., Estonian *jalalaba* “foot,” lit. “wide part of the leg, leg-blade”). It does, however, suggest that LEG and FOOT could potentially be a set of comparable, minimally distinct senses, which some languages colexify, while others do not.

Colexification effectively reduces the complexity of the lexicon, or a subdomain of the lexicon, by reducing the number of words needed to cover that space. A smaller set of words is easier to learn and remember—the associated “cognitive cost” is lower (cf. Kemp, Xu, & Regier, 2018). But if a language, or a subdomain of language such as the body parts lexicon, becomes too simple, then miscommunication becomes more likely. If the subdomain is very complex and precise, then this error, or “communicative cost” is lower—but a language can only be so complex before it becomes unfeasible to learn.

A successful language, therefore, needs to optimize these pressures to be both simple enough to be learnable (Christiansen & Chater, 2008; Gasser, 2004; Kirby, Tamariz, Cornish, & Smith, 2015; Smith, Tamariz, & Kirby, 2013) but also meet the speakers’ requirements for expressive and sufficiently precise communication. This is known as the simplicity-informativeness trade-off (see Carr, Smith, Culbertson, & Kirby, 2020; Kemp & Regier, 2012), illustrated in Fig. 1a.

Kemp et al. (2018) argue that it is social and cultural communicative needs that modulate this perpetual optimization process. If a given domain is less relevant to users of a language, then a language may give way to the pressure of simplicity, and some sense distinctions may be lost or colexified. If something is important and part of frequent discourse, then it will likely be lexified so as to avoid error in communication—increasing informativeness at the expense of increasing complexity (see Fig. 1b). The topics of informativeness, complexity, and communicative need have been discussed in the context of a variety of domains, including expressions of color (Gibson et al., 2017; Lindsey & Brown, 2002; Zaslavsky, Kemp, Tishby, & Regier, 2019a), kinship (Kemp & Regier, 2012), pronouns (Zaslavsky, Maldonado, & Culbertson, 2021), numeral systems (Xu, Liu, & Regier, 2020), natural phenomena (Berlin, 1992; Kemp, Gaby, & Regier, 2019; Regier, Carstensen, & Kemp, 2016; Zaslavsky, Regier, Tishby, & Kemp, 2020), writing systems (Hermalin & Regier, 2019; Miton & Morin, 2021), and various morphosyntactic features (Haspelmath & Karjus, 2017; Mollica, Bacon, Regier, & Kemp, 2020). Similar adaptation and efficiency effects have also been observed in experimental settings with artificial languages (cf. Chaabouni, Kharitonov, Dupoux, & Baroni, 2021; Fedzechkina, Jaeger, & Newport, 2012; Guo et al., 2021; Nölle, Staib, Fusaroli, & Tylén, 2018; Tinits, Nölle, & Hartmann, 2017; Winters, Kirby, & Smith, 2015).




Fig. 1. The simplicity-informativeness or cognitive-communicative cost trade-off (adapted from Kemp et al., 2018). Previous research has shown that natural languages strive to be “as informative as possible for their level of simplicity, and as simple as possible for their level of informativeness” (Kemp et al., 2018). This interplay gives rise to the “optimal frontier,” indicated here by dark blue squares; light gray squares stand for possible but untested languages (see panel a). It has also been argued that where exactly a language or lexical subdomain ends up along the frontier depends on communicative needs (panel b). When a language community has high communicative need to distinguish concepts in some topic or domain of the lexicon, the bias for simplicity may give way to reduce ambiguity and as such communicative cost (bottom right corner of panel b). When some domain is less relevant to a community, it may end up being simplified, at the expense of increased potential for communicative error (top left in panel b).

In a recent study based on a large sample of colexifications from a database (Borin, Comrie, & Saxena, 2013) of about 250 languages, Xu, Duong, Malt, Jiang, and Srinivasan (2020a) demonstrate that similar and associated senses (like FIRE and FLAME) are more frequently colexified than unrelated or weakly associated meanings (like FIRE and SALT, for instance), suggesting that this provides an important constraint on the evolution of lexicons. This work follows a line of research on the variability of lexification patterns across languages of the world (e.g., François, 2008; List, Terhalle, & Urban, 2013; Majid, Jordan, & Dunn, 2015; Malt, Sloman, Gennari, Shi, & Wang, 1999; Srinivasan & Rabagliati, 2015; Thompson, Roberts, & Lupyan, 2018). Xu et al. (2020a) also put forward a hypothesis that, beyond the tendency to colexify similar senses, language- and culture-specific communicative needs should be expected to affect the likelihood of colexification of similar concepts—such as SISTER and BROTHER, or ICE and SNOW—if it is necessary for efficient communication to distinguish them. The latter pair, in particular, was investigated by Regier et al. (2016), who used multiple sources of linguistic and meteorological data to show that languages spoken in colder climates are statistically more likely to distinguish ICE and SNOW, while languages spoken in warmer climates are more likely to colexify them (thus providing an empirical test to ideas going back to Sapir, 1912; Whorf, 1956). They argued this to be an example of lexicons being shaped by local cultural communicative needs, in this case which are in turn shaped by local physical environments.

Languages investigated in cross-linguistic typological studies (like François, 2008; Regier et al., 2016, Xu et al., 2020a) have evolved to be the way they are over time, through

incremental changes in how generations of speakers produce utterances by assigning signals to meanings. In this paper, we employ an artificial language experimental setup to probe lexification decisions by individual speakers, with the aim of investigating the discourse-level generative mechanisms that in natural languages would eventually lead to the observed cross-linguistic patterns.

In Experiment 1, we investigate how colexification choices play out in a dyadic communicative task when communicative needs are uniform and confirm that in this neutral condition speakers do indeed prefer to colexify similar concepts. Comparison of this baseline to a condition where colexification of similar meanings would impede communication shows a reduced tendency to colexify similar meanings, providing evidence in line with the hypothesis proposed by Xu et al. (2020a), that communicative needs of speakers modulate colexification dynamics beyond conceptual similarity.

In Experiment 2, we replicate Experiment 1 with a crowdsourced participant sample. We continue using crowdsourcing to test a variant of the original hypothesis in Experiment 3 and in Experiment 4 relax the constraints on the signal space we provide our participants in order to further explore how communicative need operates under reduced constraints on complexity. The results of these follow-up experiments support the findings of Experiment 1. The implications of the findings and pathways to future research are further discussed in Section 7.

2. Experimental methodology

All our experiments use a dyadic computer-mediated communication game setup (cf. Galantucci, Garrod, & Roberts, 2012; Kirby et al., 2015; Scott-Phillips & Kirby, 2010; Winters et al., 2015) to investigate how similarity and communicative need interact to shape colexification choices by language users. In this section, we first provide a general overview of the methods and analysis we used in all four of our experiments, before setting out each experiment in turn in more detail in later sections; unless stated otherwise, the setup for conducting a given experiment is as described in the general methods section here.

In all our experiments, pairs of participants are faced with the task of communicating single-word messages using a small set of artificial words. In order to successfully do so, they must initially negotiate the meanings for the signals through trial and error. The task was introduced to participants as an “espionage game” where the usage of secret codes is justified as keeping the messages hidden from the enemy. The experiment interface was implemented as a web app in the R Shiny framework (Chang, Cheng, Allaire, Xie, & McPherson, 2020).

2.1. Procedure

All our experiments consist of 135 rounds. At each round, one participant is the *sender* and the other is the *receiver*; these roles switch after every round. The sender is shown two meanings, represented by English nouns (see Section 2.2) and is instructed to communicate one of those meanings to the receiver, using a single word from an artificial lexicon (see Section 2.3). The receiver is then shown the same pair of meanings (in random order) and

Table 1

An example of how the communication game works. Each row corresponds to a change in the screens displayed to the players, the cells in the columns indicate what is currently being shown to each of the two Players. Interactive elements of the graphical user interface are highlighted here in bold font. In Round 1, it is Player 1's turn to signal and Player 2's turn to guess the meaning of the signal. The meanings, here MOTOR and ESSAY, are being displayed to both Players, but in randomized order. After the feedback screen has been shown (informing the players if the guess was correct), the roles are switched and the next round begins

Player 1	Player 2	Comment
MOTOR ESSAY Communicate MOTOR using... nepa qohe lali fuwo nire ruqi lumu	ESSAY MOTOR Waiting for message...	Round 1 starts Player 1 is the sender, clicks on "fuwo"
Sent MOTOR using fuwo Stand by...	ESSAY MOTOR Message: fuwo This means: essay motor	Player 2 (receiver) clicks on "MOTOR"
Correct guess! MOTOR = fuwo	Correct guess! MOTOR = fuwo	Feedback screen
THREAT PURSE Waiting for message...	THREAT PURSE Communicate THREAT using... nepa qohe lali fuwo nire ruqi lumu	Round 2 starts Player 2 is the sender, clicks on "qohe"

the sender's signal and has to guess which of the two meanings the signal represents. After taking a guess, both participants are shown an identical feedback screen which informs them whether the receiver guessed correctly; (see Table 1 for an illustration). The game ends with a screen showing both participants their total score, asking them for optional feedback, and leaving them with instructions on how to claim their monetary reward.

The participants never see more than two meanings on the screen at any time. In each game, there are 10 meanings, among them three "target pairs" consisting of two highly similar meanings, for example, MOTOR and ENGINE. The remaining four meanings serve as distractors that have low similarity scores to all other meanings, including the targets (see below for details). The signal space consists of seven artificial words such as *fuwo* or *qohe*. Since there are fewer signals than meanings, participants must colexify some meanings. We assume it takes a while to establish stable meaning correspondences, so we consider the first one third of the rounds as a "burn-in" phase and only analyze the final 90 rounds of the experiment.

2.2. Stimuli: Meanings

The meanings to be communicated are English common nouns of three to seven characters in length, drawn from the Simlex999 dataset (Hill, Reichart, & Korhonen, 2015), which consists of pairs of words and their crowdsourced similarity judgments. We use Simlex999, as it was built for evaluating models of meaning with the explicit goal of distinguishing genuine similarity (synonymy) from associativity (the dataset incorporates a subset of the USF Free Association Norms for that purpose; cf. Nelson, McEvoy, & Schreiber, 2004). Our target pairs are required to have a Simlex similarity score of at least 8 out of 10, but a free association score below 1 out of 10. This should yield meanings which are near synonymous and not




Fig. 2. The meanings and signals used in one game (dyad no. 21 in Experiment 1). The left-side panel illustrates the meaning space: only the target pairs have high similarities (dark blue), with low similarities between all other meanings. The diagonal (self-similarity) is marked by white squares. The right-side panel shows the similarity of form (as the inverse of edit distance), of both the meanings (in caps) and signals (lowercase). The stimuli in our experiments are generated in a way that ensures only target meanings are semantically similar to one another, and that form similarity remains low across the board.

simply contextually associated. For example, *arm* and *leg* have a high association score of 6.7, that is when people hear *arm* they are likely to think of *leg*—but these two are not synonyms (reflected in the Simlex similarity score of only 2.9). *plane* and *jet* have high similarity (8.1), but they also appear to be highly associated terms (6.6). In contrast, *abdomen* and *belly* have high similarity (8.1) but at the same time low associativity (0.1), therefore, making them suitable for our stimulus set.

Since Simlex does not have scores for all possible word pairs in its lexicon, we also used publicly available pretrained word embeddings (fasttext trained on Wikipedia; cf. Bojanowski, Grave, Joulin, & Mikolov, 2017) to obtain additional computational measures of similarity. We use these to ensure low similarity across the board in the distractor set, which was sampled so that no two distractors and no distractor–target pair would have vector cosine similarity above 0.2 (out of 1). Furthermore, no two nouns were substrings of each other, nor otherwise similar in form (we used a Damerau–Levenshtein edit distance¹ threshold of 3), and targets were not allowed to share the same first letter. This filtering yielded 13 eligible target pairs: ABDOMEN-BELLY, AREA-ZONE, AUTHOR-CREATOR, BAG-PURSE, COAST-SHORE, COUPLE-PAIR, DANGER-THREAT, DRIZZLE-RAIN, ENGINE-MOTOR, FASHION-STYLE, JOB-TASK, JOURNEY-TRIP, NOISE-RACKET. The meaning space for each dyad consists of three target pairs selected at random from this set of 13 possible pairs, plus an additional four distractor meanings, drawn from the remaining 475 possible meanings (see Fig. 2a).

2.3. Stimuli: Signals

For each dyad in Experiments 1–3, we generated a set of seven signals, according to the following constraints (Experiment 4 has an extended signal space of 10 signals). Each signal had a length of four characters and was composed of two consonant–vowel syllables,

constructed from a set of consonants {*qwt ps f h m r l*} and vowels {*aeoui*}. We further constrained the artificial language so that the initial letters of the signals would not overlap with any initial letters of the meanings (English nouns) in a given stimulus set. We used a large English word list to make sure there was no overlap between the artificial signals and actual English words, and furthermore made sure all signals were at least three edits distant from the meanings (the English nouns) in the same game, as well as from other signals in the game (see Fig. 2b).

2.4. *Experimental manipulation of communicative need*

The experiments have two conditions: the *baseline* or control condition, where communicative need is uniform, and the *target* condition where we manipulate communicative need, creating a situation where colexifying certain (similar) concepts would hinder the accurate exchange of messages. This applies to all experiments except for Experiment 3, which only has a (modified) target condition (see details in Section 5).

In the baseline condition, the distribution of meaning pairs (e.g., DRIZZLE-RAIN, STYLE-FASHION, PAYMENT-BULL, RAIN-PAYMENT, RAIN-FASHION) is uniform—each possible combination is shown to the participants exactly three times. Based on cross-linguistic tendencies (cf. Xu et al., 2020a), in this condition we would expect participants to colexify similar meanings. In the target condition, all possible meaning pair combinations still occur in the game, but, crucially, we manipulate the occurrence frequencies so that the target (similar-meaning) pairs occur together more often than the distractor pairs. The target pairs (e.g., DRIZZLE-RAIN, STYLE-FASHION) are shown 11 times each, and the pairs consisting of distractor meanings (e.g., PAYMENT-BULL) five times each. Pairs consisting of a meaning from a target pair plus another meaning are shown two times (e.g., RAIN-PAYMENT or RAIN-FASHION).

The nonuniform distribution of meaning pairs in the target condition entails that to communicate successfully, participants are required to select signals which allow their partner to differentiate between DRIZZLE and RAIN 11 times, but are only required to differentiate between RAIN and PAYMENT two times. The increased co-occurrence of similar meanings in the target condition simulates communicative need. If a pair of similar meanings never or seldom need to be distinguished, then it is efficient to colexify them, both from a learning and communication perspective. In contrast, if the communicative context often requires disambiguating between two similar meanings or referents—such as RAIN and DRIZZLE in a culture obsessed with talking about poor weather—then colexifying them as *rainzle* or blending them into something like *rainzle* would obviously be detrimental to communicative success. We expect this to be reflected in the outcomes of the target condition: participants should avoid colexifying the similar concepts, as that would make it difficult to distinguish them and hinder communicative efficiency.

These pairs are displayed in a randomized order, but randomized separately for the burn-in (first third) and post-burn-in part of the game. In other words, we make sure that if RAIN-FASHION is supposed to appear three times over the course of the game, then it will appear once in the burn-in and twice afterwards. Of course not all values are divisible by 3,

so the distribution of stimuli between burn-in and the rest of the game is not perfect, but we optimize it to be as good as numerically possible.²

The meaning pair frequency distribution used in both conditions ensures that individual meanings all occur exactly the same number of times (which is 27). It is necessary to control for individual frequencies in this manner, as simply making target pairs more frequent would also mean making the individual meanings in those pairs more frequent than the distractors. This would introduce a confound: another reasonable hypothesis could be that it is occurrence frequency that drives colexification (i.e., colexifying frequent meanings is preferred, or avoided; cf. Xu et al., 2020a), over and above communicative need or word similarity.

2.5. *Participants and exclusion criteria based on communicative accuracy*

All our experiments were approved by the Ethics Committee of the School of Philosophy, Psychology and Language Sciences of the University of Edinburgh. All participants provided informed consent on the online game platform prior to participating and were compensated monetarily for their time. We do not include all the data in our analysis, as the communicative accuracy of some dyads is at or near random chance. Low accuracy in guessing the meanings of the signals transmitted indicates a given dyad did not manage to converge on a lexicon they could use with reasonable success to communicate—such data are not informative in terms of our research question.

We calculate the communicative accuracy of a dyad as the percentage of correct guesses out of all guesses made in the game after the burn-in period (see Section 2.1). We make the assumption that dyads with a total accuracy score above 59%, that is guessing correctly in 53/90 of the trials, were unlikely (binomial $p = .036 < .05$) to be signaling or guessing randomly. Dyads scoring below this threshold were excluded from analysis; all excluded dyads still received full payment. The instructions also included a request not to make any notes or write anything down during the experiment, and participants were asked at the end of the experiment whether they had taken written notes; no participants were excluded on the basis of having admitted to taking written notes.

2.6. *Quantifying colexification*

We are interested in what participants do with the target meaning pairs. If they colexify target meanings in the baseline condition (e.g., use the same signal for RAIN and DRIZZLE), that would support the cross-linguistic findings of Xu et al. (2020a), that similar meanings are most often colexified. If participants avoid colexifying target meaning pairs in the target condition, then this would support the hypothesis we intend to test, that given high enough communicative need to distinguish similar meanings, they will not be colexified (see Fig. 3 for an illustration). We, therefore, need a method to operationalize these signal-meaning associations in a manner that would allow us to test our hypothesis via rigorous statistical analysis, while preferably also capturing possible changes in lexification choices over the course of each game.

Data from each game is converted into a new dataset suitable for statistical analysis, consisting only of colexifications involving one of the target meanings. We introduce a new




Fig. 3. Example signal-meaning matrices from two games in Experiment 1. The vertical black bars highlight the similar-meaning target pairs. Counts in the cells indicate how many times in total each signal was used to communicate each meaning in a given game (larger values have darker cells). In the baseline condition game (left), the players have chosen to colexify similar meanings such as SHORE and COAST. In the target condition one (right), similar meanings often need to be distinguished from one another. The players have responded to this pressure by colexifying RAIN with STYLE and DRIZZLE with ORGAN instead.

binomial variable, “colexification with synonym” and assign values to it using the following procedure. We start by iterating through all the messages sent in the post-burn-in part of the game that signal one of the target meanings. We check if the most recent usage of a given signal by the same player in the same game lexified another meaning, and if that meaning belongs to the same target pair or not. For example, given Player 1 signals RAIN using *pami* at round number 53, we check the most recent usage of *pami* by Player 1 in all preceding rounds of the game where Player 1 was the signaler. Let us say this happened at round 39, and Player 1 used *pami* to also signal RAIN. This does not count as a colexification—it has the same meaning—so this case is not added to the new dataset. At some point later in the game, at round 127, Player 1 signals RAIN using *pami* again. The previous usage of *pami* by Player 1 was round 121, where they used it to signal STYLE. This counts as a colexification, and so this case is added to the new dataset. As STYLE is not a synonym of RAIN, the value assigned to the new “colexification with synonym” variable is “no.” If instead of STYLE, the most recent meaning signaled by *pami* had been DRIZZLE—a highly similar meaning belonging to the same target pair with RAIN—then this value would be set as “yes.”

This procedure is repeated for all dyads that score above our previously described accuracy threshold. We do not include the data from the burn-in period (first one third of the game) in the statistical analysis of the results but do take the burn-in into account when checking for the most recent usage of signals, as players do not start from a clean slate after the end of the burn-in, but already have some experience with the language and likely at least some signal-meaning associations in place. This is not the only way to operationalize and analyze such data, but it does provide insights into participant behavior while allowing for explicit comparison of the conditions in terms of the likelihood of target pair colexification, and for modeling changes in lexification over the course of the game.




Fig. 4. The colexification data across all dyads in Experiment 1, operationalized using the procedure described in Section 2.6. Each row is a dyad, labeled by its number. Each tile corresponds to a message that contains a target meaning, which is being colexified with another meaning using the same signal: dark blue if with a related meaning, for example, RAIN-DRIZZLE, light blue if with an unrelated meaning. Trials are shown in order, but those involving distractor meanings, or where the given signal was last used to lexify the same meaning, are excluded (therefore the final number of data points per dyads vary, despite all games having the same 135 rounds). The difference between conditions is visually apparent: the baseline condition on the left has more darker blue, indicating colexification of similar meanings, while the target condition on the right has more light blue tiles, indicating colexification of dissimilar meanings.

The distributions of meanings occurring in the game—the input data for the participants—are carefully balanced, and all games consist of exactly 135 rounds, as described in Section 2.4. Note, however, that under this operationalization, the amount of output data yielded by different dyads varies somewhat (see Fig. 4). The exact number of data points per dyad depends on their colexification behavior: if they converge on a system where most of the target meanings get their own unique signal, then there will be fewer colexifications and as such fewer data points. If they colexify the target meanings, either with their near-synonyms or with unrelated meanings, then there will be more colexifications to analyze. This is by design as we prioritize balancing the input and is controlled for in our statistical analyses below, ensuring that the overall results are not driven by a few data-rich dyads.

2.7. Statistical modeling approach

We used mixed effects logistic regression models (using the `lme4` package in R; Bates, Mächler, Bolker, & Walker, 2015) to model all the datasets derived using the approach described in Section 2.6 above. The model has the same effects structure across all experiments. A condition (baseline or target) is treatment coded, with the baseline condition set as the reference level (with the exception of Experiment 3; this is discussed in the relevant section). A round number (scaled to a range of $[-1, 1]$) is centered at round 68, the middle of the game. The binomial colexification variable is set as the response, predicted by condition, round number, and their interaction. The interaction with round accounts for possible changes in lexification preferences. To account for the repeated measures nature of the data and the fact that the number of data points per dyad is variable, we set random intercepts for meaning and sender (the latter nested in dyad) and a random slope for condition by meaning.³

A full random effects structure would be desirable, but could not be included due to model convergence issues.

3. Experiment 1

3.1. Methods

The setup of our first experiment matches the overview in Section 2: There is a baseline and a target condition, and in both cases participants are provided with seven signals which they can use to lexify the 10 meanings. We operationalize the data for statistical modeling using the procedure as described in Section 2.6 and illustrated in Fig. 4.

3.1.1. Participants

The pool of participants for Experiment 1 consists of students of the University of Edinburgh, recruited through the university's CareerHub portal and departmental mailing lists.⁴ Participants were only allowed to complete the game once. All participants identified as native or near-native speakers of English. 46 dyads finished the experiment, 41 of which were included in the analysis (20 baseline, 21 target condition dyads; 82 participants in total). Data from five dyads were discarded, either because they explicitly admitted to misunderstanding the game instructions in the feedback form (one dyad) or because of low communicative accuracy that likely resulted from random guessing (four dyads; see Section 2.5 above for discussion of our exclusion criterion).

3.2. Results

The data-processing procedure described in Section 2.6 yields a dataset of 1214 cases (597 in the baseline, 617 in the target condition), a median of 30 per dyad (illustrated in Fig. 4). In the mixed effects regression model (see Section 2.7), the dependent variable is the derived binomial colexification measure. The fixed effects consist of the condition, round number, and the interaction of the two. The interaction is included to account for possible changes over the course of the game (cf. Fig. 4). Random effects are included to control for repeated measures of meanings, dyads, and players, as outlined in Section 2.7. In the model described in Table 2, the intercept value of -0.22 stands for the log odds of target meaning pairs being colexified, in the baseline condition, midgame (i.e., a probability of .45; recall that we centered round number). By midgame, the model is not picking up a significant difference between the conditions ($p = .3$). Each passing round does increase the probability of colexification in the baseline condition ($\beta = 1.02$, $p = .0001$). Importantly, the interaction between condition and round is in the opposite direction ($\beta = -1.18$, $p = .0015$), indicating participants were less likely to colexify related meanings in the target condition, where these related meanings frequently co-occurred and needed to be distinguished from one another. By the end of a game, the estimated average probability of colexifying target pairs is only 0.29 in the target condition, compared to 0.69 in the baseline condition. In short, these results support our communicative need hypothesis (for further illustration, see Fig. 5, the leftmost columns).

Table 2

The fixed effects from the mixed effects regression model applied to data from Experiment 1, predicting the value of the colexification variable (reference level: “no”) by the interaction between condition (reference: baseline) and round number. The latter is included to account for progress over the course of the experiment. The results indicate a statistically significant difference in participant behavior between the two conditions, supporting the hypothesis that communicative need can drive lexification preferences above and beyond conceptual similarity

Colexification ~	Estimate	SE	z	p
intercept (baseline condition, midgame)	−0.22	0.37	−0.59	.56
+ condition (target)	−0.52	0.51	−1.03	.3
+ round	1.02	0.27	3.84	<.01
+ condition (target) × round	−1.18	0.37	−3.17	<.01




Fig. 5. Estimated probabilities of colexifying similar-meaning target pairs by the end of the game. Each notch is one dyad, the wider bars are medians (unlike in our full statistical analysis, meanings are collapsed here within dyads for clearer visualization purposes). Target conditions (with manipulated communicative need) are in darker blue. In Experiment 1 (left) as well as its replication (Experiment 2), dyads were more likely to end up colexifying similar meanings like TRIP and JOURNEY in the baseline condition, but less so in the target condition, when faced with communicative need to distinguish them. This graph displays all dyads across all experiments (173 dyads or 346 participants in total) and will be referred back to in later sections.

3.3. Discussion

Human memory is not infinite, and neither is time that can be used for learning. An unlimited number of signals or signal-meaning associations cannot be stored in the brain. We emulated these natural conditions in our artificial communication experiment by providing the participants with a signal space that is smaller than the provided meaning space. The results from the baseline condition provide experimental support for the cross-linguistic findings of Xu et al. (2020a)—that people indeed tend to colexify similar meanings. Yet when faced with a situation where there is elevated communicative need to distinguish certain meaning pairs more often than others, people are more likely to colexify other pairs or clusters of meanings to maintain communicative efficiency—even if this requires colexifying unrelated meanings. This in turn supports the hypothesis suggested by Xu et al. (2020a) and is in line with previous research on communicative need in general (cf. Section 1).




Fig. 6. Communicative accuracy of all dyads across all experiments. Each notch is one dyad, the wider bars are medians. Target conditions (those with manipulated communicative need) are in darker blue. The threshold of 0.59, that we set to filter out dyads which likely played by the random button smash strategy, is shown as a gray segmented line. The student dyads (Experiment 1) had on average higher accuracy than the crowdsourced participants in the rest of the experiments, but all conditions included some dyads that scored very low, as well as some that scored very high.

It should be noted that this setup possibly puts a heavier cognitive load on the participants in the target condition. In the baseline condition, participants can colexify similar meanings (like RAIN and DRIZZLE) without paying an additional communicative cost. It is probably safe to assume similarity-driven pairings are easier to remember. In the target condition, participants are encouraged to colexify meanings which are maximally dissimilar by design (e.g., DENTIST and FASHION).⁵ In that sense, Experiment 1 constitutes a strong test of the communicative need hypothesis—we predict that given high enough communicative need, speakers would even colexify unrelated meanings rather than sacrifice communicative efficiency. However, actual differences in average communicative accuracy in Experiment 1 turn out to be negligible.⁶ Fig. 6 illustrates this, as well as the accuracy levels of the rest of the experiments, which we will discuss further in the next sections. There was also no difference in game length (on average 29 min in both conditions).

Regardless, we will still explore the weaker hypothesis in a follow-up experiment, where participants are provided with less frequently co-occurring but still similar meanings to colexify (Experiment 3 in Section 5). We will also describe an experiment with an expanded signal space (Experiment 4 in Section 6). However, in the next section, we will first replicate the initial study on a different, slightly larger sample of participants.

4. Experiment 2

Experiment 2 is a replication of Experiment 1, where we recruited participants from Amazon Mechanical Turk; all other details of the experimental design and analyses are the same. As our replication below demonstrates, the behavior of these two samples in terms of the research question is very similar, but the samples differ somewhat in terms of average communicative accuracy.

4.1. Methods

4.1.1. Participants

We restricted participation to Mechanical Turk workers based in the United States to have a sample roughly comparable sample to Experiment 1 (i.e., largely English speaking). In this and the following experiments, we only accepted workers with a history of at least 1000 successfully completed tasks and a 97% or higher approval rate. Furthermore, we used the qualifications system of Mechanical Turk to make sure no worker participated more than once (within a single experiment or across Experiments 2–4). As before, all participants provided informed consent on the online game platform prior to participating and were compensated monetarily for their time.

79 dyads finished the experiment, 53 of which were included in the analysis (26 baseline, 27 target condition dyads; 106 participants in total). Data from 26 dyads were discarded, 24 because of low communicative accuracy and two due to suspected cheating (see below). Communicative accuracy turned out to be lower on Mechanical Turk than in our student sample, with many players operating at or even below random chance (see Fig. 6 above in Section 3.3). It is unclear to us why the rejection rate was so high, and this may reflect the difficulty of our communication task relative to other tasks our participants were completing around the same time, although anecdotally we know of other researchers who had problems with data quality on Mechanical Turk in the summer of 2020. There is also a greater difference in communicative accuracy between the baseline and target condition in Experiment 2, compared to Experiment 1, possibly due to the heavier cognitive load of the target condition compared to the baseline, discussed in Section 3.3.

We manually flagged two dyads as suspicious, on the basis that they achieved near-perfect accuracy while sending apparently random signals. This could potentially be one person using two Mechanical Turk accounts to play against themselves, or two workers communicating out with the game interface. After discovering these anomalies, we manually inspected data from high-accuracy dyads in all experiments. It was also not uncommon for participants recruited via Mechanical Turk to simply drop out midgame, effectively canceling the experiment for their dyad partner as well; these dyads were treated as having withdrawn consent, and their data were not analyzed.

4.2. Results

The analysis procedure for Experiment 2 is identical to that of Experiment 1: We operationalize colexification (see Section 2.6; yielding a dataset of 1659 cases, a median of 32 per dyad) and apply the same mixed effects logistic model (Section 2.7). Experiment 2 successfully replicated the results of Experiment 1 with the condition-round interaction coefficient being significantly negative ($\beta = -0.66$, $p = .03$; see Table 3; refer back to Fig. 5 for visual comparison). This value indicates participants were again less likely to colexify related meanings in the target condition which simulated communicative need (average probability of .46 by the end of the game), compared to the baseline condition, where there was no such pressure, and where participants more often colexify related meaning pairs (.72 probability).

Table 3

The fixed effects from the mixed effects regression model applied to data from Experiment 2, predicting the value of the colexification variable (reference level: “no”) by the interaction between condition (reference: baseline) and round number. The results replicate those of Experiment 1

Colexification ~	Estimate	SE	Z	p
intercept (baseline condition, midgame)	−0.2	0.29	−0.68	.49
+ condition (target)	−0.44	0.36	−1.24	.22
+ round	1.14	0.24	4.82	<.01
+ condition (target) × round	−0.66	0.31	−2.11	.03

4.3. Discussion

The successful replication gives us additional confidence in the overarching hypothesis concerning the role of communicative need in lexification choices. We now turn to two more experiments to gain a better understanding of how communicative need and simplicity preferences shape the behavior of our participants, testing a weaker version of the communicative need hypothesis in Experiment 3 and relaxing the signal space constraint in Experiment 4.

5. Experiment 3

As discussed in Section 3.3, the target condition in Experiments 1 and 2 pushes participants to colexify meanings which are highly dissimilar (recall Fig. 4). In that sense, it tested a strong version of our hypothesis, that the need to maintain communicative efficiency would outweigh the awkwardness of colexifying unrelated concepts, for example, BULL and FASHION (which would amount to homonymy in natural languages, something which has been shown to hinder lexical retrieval and processing; cf. Beretta, Fiorentino, & Poeppel, 2005; Klepousniotou, 2002). Here we explore an alternative, possibly more natural target condition, where participants can avoid colexifying target meanings by colexifying nontarget meanings which are similar to one another.

5.1. Methods

5.1.1. Participants

The sample is similar to Experiment 2: We source participants from Mechanical Turk, applying the same restrictions as in Experiment 2. 52 dyads finished the experiment, 27 of which were included in the analysis (54 participants in total). Data from 25 dyads were discarded, 23 because of low communicative accuracy and two due to suspected cheating (see Section 4.1.1). The average accuracy is similar to what we observed in Experiment 2 (see Fig. 6).

5.1.2. Procedure

In this experiment, there is only a single, target-type condition, that is, one where communicative need is manipulated. The meaning space still consists of 10 meanings, with three

Table 4

The fixed effects from the mixed effects regression model applied to combined data from Experiments 2 and 3, predicting the value of the colexification variable (reference level: “no”) by the interaction between condition (reference: new weaker target condition) and round number. Participant behavior does not differ significantly between the original target condition and the new weaker-hypothesis target condition

Colexification ~	Estimate	SE	z	p
intercept (weaker target condition, mid-game)	-0.22	0.25	-0.9	.37
+ condition (baseline)	0.04	0.34	0.12	.9
+ condition (target)	-0.39	0.34	-1.16	.24
+ round	0.14	0.2	0.68	.49
+ condition (baseline) × round	0.95	0.3	3.11	<.01
+ condition (target) × round	0.3	0.29	1.06	.29

high-similarity target meanings pairs and four distractors. Crucially, here the distractors also form two similarity pairs. For example, in Experiment 3, a meaning space might consist of BAG, PURSE, DRIZZLE, RAIN, AUTHOR, CREATOR, DANGER, THREAT, JOURNEY, TRIP (target meaning pairs underlined). Recall that distractor pairs only co-occur (i.e., need to be distinguished) five times over the course of the game, while target pairs co-occur 11 times in the target condition. Colexifying the distractors is now incentivized not only by their lower co-occurrence but also their semantic similarity. Doing so would allow participants to reserve more unique signals to distinguish target meanings, the ones that do co-occur often and need to be distinguished. All other parameters are identical to Experiments 1 and 2.

5.2. Results

To obtain a point of comparison, we combine the data collected in this experiment with the baseline and target condition data from Experiment 2 (totaling 2527 cases, of those 868 from the weaker hypothesis target condition of Experiment 3). We fit a mixed effects logistic regression model with the same structure as before but set the new, “weaker” target condition as the reference level for the condition effect. We find that as with the previous target condition, there is still a significant difference from the baseline condition (the condition × round interaction $p = .002$). However, participant behavior does not differ from the target condition of Experiment 2, that is, our manipulation of the meaning space did not elicit a meaningful difference ($p = .29$; see Table 4 and Fig. 5).

5.3. Discussion

We once again observe the same general trend (echoing the cross-linguistic findings of Xu et al., 2020a) that participants colexify any similar-meaning pairs even if it causes the occasional miscommunication, but change their behavior if that cost becomes too high, as evidenced once more by the significant difference between the baseline condition of Experiment 2 and the new target condition of Experiment 3. We also expected that participants would colexify target pairs even less often than in the target conditions of Experiment 2, since alternative similar pairs were available among the distractors, removing one of the assumed

barriers to avoiding colexifying target pairs. The data do not support this hypothesis; we see the same (relatively low) rate of colexification of target pairs in our new data. One explanation may be that the difference between co-occurrence frequencies is simply not stark enough for participants to see a benefit (consciously or subconsciously) in colexifying the distractors—and there is still the option to colexify target-distractor pairs, which co-occur less often (twice each, vs. the five of the distractor-only pairs).

6. Experiment 4

This follow-up to the initial study explores the effect of removing constraints on the signal space. In all the previous experiments, participants were provided 10 meanings but only seven signals to work with. This meant some meanings would be colexified, either accidentally (leading to lowered communicative accuracy) or systematically (allowing for higher accuracy). Here, we remove this requirement to colexify, with the aim of gaining a better understanding of participants' behavior when this central component of our experimental setup is relaxed. We expect one of two outcomes. Participants may well behave more or less the same as in Experiments 1 and 2, making use of a limited number of signals and colexifying some meanings—after all, seven signals should be easier to learn and remember than 10. Alternatively, if 10 signals are not too unmanageable, participants may forgo colexification altogether: this would eliminate the difference between baseline and target condition outcomes.

6.1. Methods

6.1.1. Participants

We continue using participants from Mechanical Turk as before. 91 dyads finished the experiment, 52 of which were included in the analysis (104 participants in total). Data from 39 dyads was discarded because of low communicative accuracy.

6.1.2. Procedure

This setup is the same as that of Experiment 1 and 2: There is a baseline and a target condition, and the meaning space and its associated occurrence distributions are arranged as discussed in Section 2.1). The single difference is that participants are not forced to colexify any meanings, as the size of the signal space is 10, equal to the meaning space (illustrated in Fig. 7).

6.2. Results

The results of applying the mixed effects logistic regression model to the data from Experiment 4 (1306 cases after operationalizing colexification) show that when the requirement to colexify is removed, by allowing a larger signal space, the difference between the baseline and target condition disappears ($p = .37$; see Table 5; see also Fig. 5 for a visualization across all conditions). Communicative accuracy is roughly the same as in the other Mechanical Turk samples (recall Fig. 6).

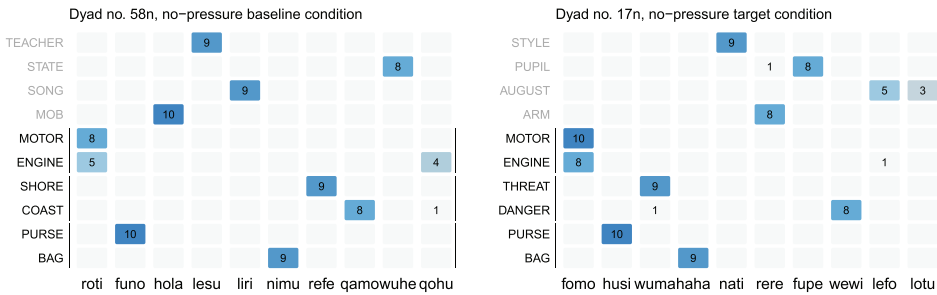


Fig. 7. Signal-meaning matrices illustrating the no-pressure baseline (left) and target condition (right) of Experiment 4. Above average accuracy dyads are used again for visualization purposes, but their behavior is reflective of the average tendencies in this experiment. Here, participants make use of more signals in both conditions and generally colexify less than in previous experiments.

Table 5

The fixed effects from the mixed effects regression model applied to data from Experiment 4, predicting the value of the colexification variable (reference level: “no”) by the interaction between condition (reference: baseline) and round number. Here, in contrast to previous experiments, the conditions yield similar results

Colexification ~	Estimate	SE	z	p
intercept (no-pressure baseline condition, midgame)	-0.73	0.27	-2.74	<.01
+ condition (target)	0.24	0.36	0.68	.5
+ round	0.19	0.24	0.78	.44
+ condition (target) × round	-0.3	0.33	-0.89	.37

6.3. Discussion

When the signal space is larger, the difference between the baseline and the target condition disappears. Participants behave in the baseline condition similarly to participants in the target conditions, avoiding colexification of the target pairs. We also quantified signal usage entropy for all dyads across all conditions. The results show that in the no-pressure conditions with the expanded signal space, in absolute terms more signals were used on average (Fig. 8). In relative terms, signal usage in Experiment 4 was not that different from the previous experiments, with some dyads making use of most of the signal space (notches close to the limit lines in Fig. 8) and others being more conservative and colexifying instead.

It appears participants favored the effort of remembering a few extra signals over having a simpler but more ambiguous language. While the signaling options are visible in the game at all times, the entire meaning space is never revealed all at once (recall Table 1). Participants here seem to have picked up on the relative abundance of the signals; nevertheless, and unlike in the weaker hypothesis condition, participants did not colexify similar meanings to the same extent. These results solidify the original findings of Experiments 1 and 2: the size of the signal space clearly makes a difference in participant behavior and allows for making inferences into speakers’ lexification preferences. Naturally, our signal spaces are tiny compared to real lexicons that humans are able to memorize, not least because our signal space is designed for




Fig. 8. Signaling entropy across all dyads (blue notches), in the post-burn-in part of the game. Low values on the vertical axis indicate fewer signals were consistently used, higher values indicate a larger variety of signals were used by a given dyad. Overlapping values are pushed slightly aside horizontally to ensure visibility. The black bars represent medians. The gray vertical lines at 1.9 and 2.3 indicate maximal entropy given the size of the signal spaces. Dyads in Experiment 4 generally made use of most of the extended signal space (10 instead of 7), setting it apart from the rest of the experiments.

a lexicon that is to be conventionalized and used in the span of about half an hour. And like in natural languages, the complexity of the lexicon, the number of signals, has a limit. It is just that this limit is cognitive in nature in the real world and appears to be largely artificial (i.e., driven by our constraint on the size of the signal space) in our experiments.

7. General discussion

Our research provides evidence that speakers' communicative needs affect their lexification choices, validating the viability of the mechanism hypothesized based on large-scale cross-linguistic studies (cf. Kemp et al., 2018; Xu et al., 2020a). Our research makes this connection explicit and describes an experimental paradigm to test such hypotheses on the level of individual discourse—in comparison to previous research focusing on the level of population consensus based on data such as dictionaries, grammars, and corpora (e.g., Mollica et al., 2020; Ramiro, Srinivasan, Malt, & Xu, 2018, Xu et al., 2020a). Below, we sketch some extensions to the experimental paradigm established here that we feel would be worthwhile to look into in order to gain a better understanding of the role of communicative need, similarity and associativity, and the formation of lexicons in general.

7.1. Extensions and implications

Future research could look into a number of aspects and parameters of the communicative need game, beyond the exploration in our two follow-up studies (Experiments 3 and 4). In terms of setup and procedure, we chose what we assumed would be a reasonably sized meaning space for an experiment of this length, but this, as well as the length itself, is of course arbitrary, as is the chosen co-occurrence distributions that emulate the pressure of communicative need. It would be interesting to see both the effect of stronger and weaker pressures than employed here. Need could also be gradually increased over the course of a longer game

to probe the strength of the pressure required for participants to modify an already established artificial communication system (which would loosely correspond to natural language users changing their language over their lifespan; cf. Sankoff, 2018).

In this study, we chose direct similarity or synonymy as the semantic relationship to explore. Xu et al. (2020a) show that conceptual associativity (i.e., the *car, engine* type) also correlates with cross-linguistic colexification patterns, roughly to the same extent. It would be interesting to use our paradigm to investigate whether dyads preferentially colexify based on associativity or similarity. Another possible predictor, related to associativity, could be co-occurrence probability (or mutual information), which could be inferred from a large text corpus and then tested experimentally. Differences between more fine-grained relationships like register-varying synonymy (*abdomen, belly*) and hyponymy (*bag, purse*) could also be probed, as well as how these preferences may correlate with historical patterns of sense formation and expansion (cf. Ramiro et al., 2018).

Xu et al. (2020a) also discuss a potential role of frequency, namely that more commonly referred-to senses may be more likely colexified. We control for frequency in our experiment by making sure the occurrence distribution of meanings is uniform in all games. Future research could let frequency vary systematically to determine its importance. Previous research (cf. Atkinson, Mills, & Smith, 2018; Blythe & Croft, 2021; Raviv, Meyer, & Lev-Ari, 2019; Raviv, de Heer Kloots, & Meyer, 2021; Segovia-Martín, Walker, Fay, & Tamariz, 2020) has also demonstrated that community size and links between individuals has an effect on (artificial) language formation, learnability, and change. Another potentially interesting extension would be to run the colexification experiment using a larger speaker group than a dyad, perhaps also manipulating the connections between participants to investigate the proposed network effects. In our experiment, the participants were only presented with two possible meanings to guess between, to facilitate convergence on systematic associations in a short-time frame. This also means it was trivial to achieve 50% guessing accuracy. In a longer experiment, the number of choices could be increased.

Our work may also have potentially interesting implications for historical linguistics. Population-level changes large enough to register on historical timescales must also start with differential utterance selection on the individual level, before they can compound over time and larger groups to become the norm (cf. Baxter, Blythe, Croft, & McKane, 2006; Croft, 2000). If the mechanisms discussed and experimented with here are representative of utterance selection dynamics in natural languages, then the next interesting question would be: To what extent does communicative need constitute selection in language change? (cf. Andersen, 1990; Baxter et al., 2006; Newberry, Ahern, Clark, & Plotkin, 2017; Reali & Griffiths, 2010; Steels & Szathmáry, 2018). A comprehensive understanding of lexical evolution, and language evolution in general, would benefit from merging the perspectives of individual mechanisms and population-level consensus changes.

7.2. Complexity, information loss, and communicative need

In more broad terms, our study interfaces with a growing body of work on the interplay between the orthogonal pressures of simplicity and informativeness in language evolution.




Fig. 9. Average communicative cost and cognitive cost (complexity) scores in Experiment 1. Light blue squares are target meaning pairs (such as DRIZZLE-RAIN) used by baseline condition dyads, dark blue ones are pairs as used by target condition dyads. A larger square size indicates multiple overlapping squares at those coordinates. Simulated results are depicted as gray blocks in the background. The few slightly darker gray squares, mostly in the top right, are the few dyads in Experiment 1 that performed below the minimal threshold of 59%. Most dyads communicate at or near the optimal points at (0,1) and (1,0), and none of the dyads (above the accuracy threshold) end up in the suboptimal top right corner. This picture mirrors findings in natural language lexicons, which also tend towards the optimal frontier (cf. Fig. 1).

The former relates to ease of learning; while the latter relates to low information loss or communicative cost (the terminology and foci vary between authors and disciplines, cf. Beckner, Pierrehumbert, & Hay, 2006; Bentz, Alikaniotis, Cysouw, & Ferrer-i Cancho, 2017; Carr et al., 2020; Carstensen, Xu, Smith, & Regier, 2015; Denić, Steinert-Threlkeld, & Szymanik, 2021; Fedzechkina et al., 2012; Gasser, 2004; Haspelmath, 2021; Kemp & Regier, 2012; Kirby & Hurford, 2002; Kirby et al., 2015; Nölle et al., 2018; Smith, 2020; Steinert-Threlkeld & Szymanik, 2020; Uegaki (in preparation); Winters et al., 2015; Zaslavsky, Regier, Tishby, & Kemp, 2019b). These studies have yielded converging evidence that languages which are learned and used in communication—the real-world ones, the artificial ones grown in the lab, as well as those evolved by computational agents—all aspire to balance these two pressures, ending up somewhere along the optimal frontier.

Our results provide support for the argument that culture-specific communicative needs may modulate the location of a language on that frontier (cf. Kemp et al., 2018). The pressure for simplicity in lexicons can be relaxed in favor of more expressivity, given high enough communicative need (which we emulated in the target condition, and the no-pressure conditions)—while informativeness can give way to simplicity when a less expressive lexical subspace does the job (cf. our baseline condition).

To allow for systematic statistical testing of the hypothesis in the previous sections, we applied a transformation to the data (outlined in Section 2.6) that operationalized colexification the most objective manner we could think of. Fig. 9 is a reanalysis of data from Experiment 1 along the axes of complexity and expressivity, using an alternate coding scheme. To simplify things, we ignore the sender here and treat the language of each dyad as a

collaborative effort. Each message containing a target meaning is assigned a simplified cognitive cost score and communicative cost score between $[0,2]$, which depends on the last most recent reference to the same meaning and the last reference to the synonym (target pair member) of the current meaning. Fig. 9 displays the mean scores for each target meaning pair used by each dyad in Experiment 1, across all their respective utterances. We also simulate the full possibility space of the results under this coding scheme, using a simple agent-based model (shown as gray blocks in Fig. 9). This is to provide meaningful dimensionality to the graph: given the number of signals and meanings, it would be impossible to obtain a maximal score simultaneously on both axes. The technical details of both procedures are further described in the Appendix.

In summary, dyads tend towards the optimal frontier between complexity and expressivity (or in this case, the two optimal points at $(1,0)$ and $(0,1)$; positioning near the bottom left diagonal just indicates mixed strategies). Baseline dyads favor simplicity (having the luxury to do so), while target condition dyads sacrifice simplicity to reduce communicative cost (in response to our manipulation driving them to do so).

7.3. *Beyond small subsystems*

Extrapolating the communicative need argument beyond the grammatical and lexical subsystems mentioned in Section 1 to the scale of entire languages, we would expect semantic spaces of different languages to be mostly uniform in density—how many words are used to express shades of any given concept or meaning subspace—but differ in exactly where culture-specific communicative needs of the time either require more detail, or where fewer words will suffice (analogous to uniform information density on the level of utterances; cf. Levy, 2018). Previous research has focused on delimited domains of language like tense or kinship. This makes sense both from a data collection and computational point of view: quality cross-linguistic data are not trivial to acquire, and neither is computing complexity, information loss or expressivity, the larger the system under scrutiny. This is then the next challenge: understanding these pressures and the evolution of lexicons and grammars, over time, and cross-linguistically, on the scale of entire languages (as opposed to isolated domains). This would require combining explicitly quantified metrics of simplicity and expressivity (cf. Bentz et al., 2017; Mollica et al., 2020; Piantadosi, Tily, & Gibson, 2011; Steinert-Threlkeld & Szymanik, 2020; Zaslavsky et al., 2019b), some estimate of communicative need (cf. Karjus, Blythe, Kirby, & Smith, 2020; Regier, Kemp, & Kay, 2015), some measure of density or colexification (see Chapter 5.4 of Karjus, 2020, for one potential approach), and if using a machine learning-driven approach such as word embeddings, either a joint semantic model of multiple languages allowing for direct cross-linguistic comparison of lexical densities and colexification (e.g., Chen & Cardie, 2018; Rabinovich, Xu, & Stevenson, 2020; Thompson et al., 2018), or language-specific diachronic semantic models to observe changes in colexification over time (e.g., Dubossarsky, Hengchen, Tahmasebi, & Schlechtweg, 2019; Rosenfeld & Erk, 2018; Ryskina, Rabinovich, Berg-Kirkpatrick, Mortensen, & Tsvetkov, 2020).

8. Conclusions

We investigated the cross-linguistic tendency of colexification of similar concepts from earlier lexico-typological research using artificial language experiments and tested the hypothesis that colexification dynamics may be driven not only by concept similarity but also the communicative needs of linguistic communities. Our data support both claims: Speakers readily colexify similar concepts, unless distinguishing them is necessary for successful communication, in which case they do not. These results, despite being based on artificial communication scenarios and small lexicons, illustrate the interaction between similarity and communicative need in shaping colexification. We also proposed pathways for future study of these phenomena beyond small word sets and on the scale of entire lexicons.

Language change is driven by a multitude of interacting forces, ranging from random drift to sociolinguistic pressures to institutional language planning, to selection by speakers for more efficient and expressive forms. Our work supports the argument that speakers' communicative needs—a factor balancing and modulating the relative importance of the higher level pressures for simplicity and informativeness—should be considered as one of such forces.

Acknowledgments

We would like to thank Yang Xu, Barbara C. Malt, and Mahesh Srinivasan for useful discussions and comments, Jonas Nölle for advice with the initial experimental design, and the anonymous reviewers of *Cognitive Science* as well as the associate editor, Padraic Monaghan, for their constructive feedback.

Data collection for this research was funded by the Postgraduate Research Support Grant of the School of Philosophy, Psychology and Language Sciences of the University of Edinburgh. This research also received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant Agreement 681942), held by Kenny Smith. Andres Karjus is supported within the CUDAN ERA Chair project for Cultural Data Analytics at Tallinn University, funded through the European Union Horizon 2020 research and innovation program (Project No. 810961), and was also supported by the Kristjan Jaak postgraduate scholarship of the Archimedes Foundation of Estonia during initial data collection.

Andres Karjus designed and carried out the experiments, conducted the analysis, wrote the text, and created the figures. Tianyu Wang carried out additional experiments. Kenny Smith, Richard A. Blythe and Simon Kirby provided advice on the design of the experiment and data analysis, as well as edits and comments on the text. A shorter version of this paper formed a part of a chapter in the doctoral thesis of the first author (Karjus 2020).

Code and data availability

All the experiment data and scripts to replicate the analyses are available in the following Github repository, along with the full codebase of the Shiny game application

we developed to run the dyadic experiments: https://github.com/andreskarjus/colexification_experiment.

Open Research Badges



This article has earned Open Data and Open Materials badges. Data and materials are available at <https://doi.org/10.5281/zenodo.4637166>.

Notes

1. A metric that corresponds to the minimum number of operations required to change one word into the other. The operations consist of insertions, deletions, substitutions of single characters and, as an addition to the original Levenshtein distance, the transposition of two adjacent characters. The edit distance to get from *bag* to *purse* would be 5: three substitutions plus the *s* and *e* inserted at the end.
2. The randomized order could in principle cause a situation where some meaning comes up in multiple successive rounds, or, on the contrary, does not appear for a while. However, since there are only 10 meanings in each game, such runs should be relatively rare and, therefore, have limited systematic effect on our results.
3. In lme4 syntax: `colexification ~ condition * round + (1+condition|meaning) + (1|dyad/sender)`
4. The experiment was initially planned as an in-person lab study that would have commenced in March 2020. The Covid-19 pandemic rendered this impossible, so we reworked it into an online experiment, but continued with the originally intended participant pool of student recruits. We made use of SimpleSignUp (<https://github.com/jwcarr/SimpleSignUp>) to schedule the participants. For the later experiments, we moved to using participants sourced from Amazon Mechanical Turk.
5. However, note that since the guess is always made just between two options, the random baseline is still 50%. This means it is still possible to achieve a reasonably high communicative accuracy score even if the dyad decides to colexify all target meaning pairs—as some dyads indeed do—and just guess when a target pair comes up. Assuming otherwise perfect 100% performance on all other, nontarget pairs (this is, however, never observed) and that 50% of the random guesses hit the mark, it would be technically possible to achieve 82% accuracy in the post-burn-in part of the game, using this strategy.
6. The estimated probability of making a correct guess is 0.84 in the target condition compared to 0.88 in the baseline ($p = .12$), based on a mixed effects logistic model, predicting correctness of guess by condition, with random slopes and intercepts for meaning and dyad (only taking into account the post-burn-in part of each game, and excluding dyads with overall accuracy below 59%, as described above).

References

- Andersen, H. (1990). The structure of drift. In H. Andersen & K. Koerner (Eds.), *Historical linguistics 1987. Papers from the 8th International Conference on Historical Linguistics* (pp. 1–20). Amsterdam, The Netherlands: Benjamins.
- Atkinson, M., Mills, G. J., & Smith, K., (2018). Social group effects on the emergence of communicative conventions and language complexity. *Journal of Language Evolution*, 4(1), 1–18. <https://doi.org/10.1093/jole/lzy010>
- Bates, D., Mächler, M., Bolker, B., & Walker, S., (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. <https://doi.org/10.18637/jss.v067.i01>.
- Baxter, G. J., Blythe, R. A., Croft, W., & McKane, A. J., (2006). Utterance selection model of language change. *Physical Review E*, 73(4), 046118. <https://link.aps.org/doi/10.1103/PhysRevE.73.046118>
- Beckner, C., Pierrehumbert, J. B., & Hay, J., (2017). The emergence of linguistic structure in an online iterated learning task. *Journal of Language Evolution*, 2(2), 160–176.
- Bentz, C., Alikaniotis, D., Cysouw, M., & Ferrer-i Cancho, R., (2017). The entropy of words' learnability and expressivity across more than 1000 languages. *Entropy*, 19(6), 275.
- Beretta, A., Fiorentino, R., & Poeppel, D., (2005). The effects of homonymy and polysemy on lexical access: An MEG study. *Cognitive Brain Research*, 24(1), 57–65. <https://doi.org/10.1016/j.cogbrainres.2004.12.006>.
- Berlin, B. (1992). *Ethnobiological classification: Principles of categorization of plants and animals in traditional societies*. Princeton, NJ: Princeton University Press.
- Blythe, R. A., & Croft, W., (2021). How individuals change language. *Plos One*, 16(6), e0252582. <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0252582>
- Bojanowski, P., Grave, E., Joulin, A., & Mikolov, T., (2017). Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5, 135–146.
- Borin, L., Comrie, B., & Saxena, A., (2013). The Intercontinental Dictionary Series' A rich and principled database for language comparison. *Approaches to Measuring Linguistic Differences*, 285, 302.
- Carr, J. W., Smith, K., Culbertson, J., & Kirby, S., (2020). Simplicity and informativeness in semantic category systems. *Cognition*, 202, 104289. <https://doi.org/10.1016/j.cognition.2020.104289>.
- Carstensen, A., Xu, J., Smith, C. T., & Regier, T. (2015). Language evolution in the lab tends toward informative communication. In D. Noelle et al. (Eds.), *Proceedings of the 37th Annual Meeting of the Cognitive Science Society* (pp. 303–308). Austin, TX: Cognitive Science Society.
- Chaabouni, R., Kharitonov, E., Dupoux, E., & Baroni, M., (2021). Communicating artificial neural networks develop efficient color-naming systems. *Proceedings of the National Academy of Sciences*, 118(12), e2016569118. <https://doi.org/10.1073/pnas.2016569118>.
- Chang, W., Cheng, J., Allaire, J., Xie, Y. & McPherson, J. (2020). Shiny: Web application framework for R. Available at <https://CRAN.R-project.org/package=shiny> (Accessed 01.03.2020)
- Chen, X. & Cardie, C. (2018). Unsupervised multilingual word embeddings. In E. Riloff, D. Chiang, J. Hockenmaier, & J. Tsujii (Eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing* (pp. 261–270). Stroudsburg, PA: Association for Computational Linguistics.
- Christiansen, M. H., & Chater, N., (2008). Language as shaped by the brain. *Behavioral and Brain Sciences*, 31(5), 489–509. <https://doi.org/10.1017/S0140525X08004998>.
- Croft, W. (2000). *Explaining language change: An evolutionary approach*. London: Longman.
- Denić, M., Steinert-Threlkeld, S. & Szymanik, J. (2021). Complexity/informativeness trade-off in the domain of indefinite pronouns. *Proceedings of Semantics and Linguistic Theory*, 30, 166–184. <https://doi.org/10.3765/salt.v30i0.4811>
- Dubossarsky, H., Hengchen, S., Tahmasebi, N., & Schlechtweg, D. (2019). Time-out: Temporal referencing for robust modeling of lexical semantic change. In P. Nakov & A. Palmer (Eds.), *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (pp. 457–470). Stroudsburg, PA: Association for Computational Linguistics. <https://www.aclweb.org/anthology/P19-1044>

- Fedzechkina, M., Jaeger, T. F., & Newport, E. L., (2012). Language learners restructure their input to facilitate efficient communication. *Proceedings of the National Academy of Sciences*, 109(44), 17897–17902. <https://doi.org/10.1073/pnas.1215776109>.
- Frank, M. C., & Goodman, N. D., (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084), 998–998. <https://science.sciencemag.org/content/336/6084/998>
- François, A. (2008). Semantic maps and the typology of colexification: Intertwining polysemous networks across languages. In M. Vanhove (Ed.), *Studies in language companion series* (vol. 106, pp. 163–215). Amsterdam, The Netherlands: John Benjamins. <https://doi.org/10.1075/slcs.106.09fra>.
- Galantucci, B., Garrod, S., & Roberts, G., (2012). Experimental semiotics. *Language and Linguistics Compass*, 6(8), 477–493. <https://doi.org/10.1002/lnc3.351>.
- Gasser, M. (2004). The origins of arbitrariness in language. In K. Forbus, D. Gentner, & T. Regier (Eds.), *Proceedings of the 26th Annual Conference of the Cognitive Science Society* (vol. 26, pp. 434–439). Austin, TX: Cognitive Science Society.
- Gibson, E., Futrell, R., Jara-Ettinger, J., Mahowald, K., Bergen, L., Ratnasingam, S., Gibson, M., Piantadosi, S. T., & Conway, B. R., (2017). Color naming across languages reflects color use. *Proceedings of the National Academy of Sciences*, 114(40), 10785–10790. <https://doi.org/10.1073/pnas.1619666114>.
- Guo, S., Ren, Y., Mathewson, K., Kirby, S., Albrecht, S. V. & Smith, K. (2021). Expressivity of emergent language is a trade-off between contextual complexity and unpredictability. arXiv preprint: 2106.03982.
- Haspelmath, M., (2021). Explaining grammatical coding asymmetries: Form-frequency correspondences and predictability. *Journal of Linguistics*, 57, 605–633. <https://doi.org/10.1017/S0022226720000535>.
- Haspelmath, M., & Karjus, A., (2017). Explaining asymmetries in number marking: Singulatives, pluratives, and usage frequency. *Linguistics*, 55(6), 1213–1235.
- Hermalin, N., & Regier, T. (2019). Efficient use of ambiguity in an early writing system: Evidence from Sumerian cuneiform. In A. K. Goel, C. M. Seifert, & C. Freksa (Eds.), *Proceedings of the 41st Annual Meeting of the Cognitive Science Society, CogSci 2019: Creativity + Cognition + Computation, Montreal, Canada, July 24-27, 2019* (pp. 422–427). Montreal, Canada: Cognitive Science Society. Retrieved from cognitivesciencesociety.org. <https://doi.org/10.1515/ling-2017-0026>.
- Hill, F., Reichart, R., & Korhonen, A., (2015). SimLex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*, 41(4), 665–695. https://doi.org/10.1162/COLI_a_00237.
- Karjus, A. (2020). Competition, Selection and communicative need in language change (Ph.D. thesis). University of Edinburgh, Edinburgh, UK.
- Karjus, A., Blythe, R. A., Kirby, S., & Smith, K., (2020). Quantifying the dynamics of topical fluctuations in language. *Language Dynamics and Change*, 10(1), 86–125. <https://doi.org/10.1163/22105832-01001200>.
- Kemp, C., Gaby, A. & Regier, T. (2019). Season naming and the local environment. In A. K. Goel, C. M. Seifert, & C. Freksa (Eds.), *Proceedings of the 41st Annual Meeting of the Cognitive Science Society, CogSci 2019: Creativity + Cognition + Computation*, Montreal, Canada, July 24-27, 2019 (pp. 539–545). Montreal, Canada: Cognitive Science Society.
- Kemp, C., & Regier, T., (2012). Kinship categories across languages reflect general communicative principles. *Science (New York, N.Y.)*, 336(6084), 1049–1054. <https://doi.org/10.1126/science.1218811>.
- Kemp, C., Xu, Y., & Regier, T., (2018). Semantic typology and efficient communication. *Annual Review of Linguistics*, 4(1), 109–128. <https://doi.org/10.1146/annurev-linguistics-011817-045406>.
- Kirby, S., & Hurford, J. R. (2002). The emergence of linguistic structure: An overview of the iterated learning model. In A. Cangelosi & D. Parisi (Eds.), *Simulating the evolution of language* (pp. 121–147). London: Springer.
- Kirby, S., Tamariz, M., Cornish, H., & Smith, K., (2015). Compression and communication in the cultural evolution of linguistic structure. *Cognition*, 141, 87–102. <https://doi.org/10.1016/j.cognition.2015.03.016>.
- Klepousniotou, E., (2002). The processing of lexical ambiguity: Homonymy and polysemy in the mental lexicon. *Brain and Language*, 81(1), 205–223. <https://doi.org/10.1006/brln.2001.2518>.

- Levy, R. P. (2018). Communicative efficiency, uniform information density, and the rational speech act theory. In *Proceedings of the 40th Annual Meeting of the Cognitive Science Society* (pp. 684–689). Austin, TX: Cognitive Science Society.
- Lindsey, D. T., & Brown, A. M., (2002). Color naming and the phototoxic effects of sunlight on the eye. *Psychological Science*, 13(6), 506–512. <https://doi.org/10.1111/1467-9280.00489>.
- List, J.-M., Terhalle, A., & Urban, M. (2013). Using network approaches to enhance the analysis of cross-linguistic polysemies. In A. Koller & K. Erk (Eds.), *Proceedings of the 10th International Conference on Computational Semantics (IWCS 2013) – Short Papers* (pp. 347–353). Stroudsburg, PA: Association for Computational Linguistics.
- Majid, A., Jordan, F., & Dunn, M., (2015). Semantic systems in closely related languages. *Language Sciences*, 49, 1–18. <https://doi.org/10.1016/j.langsci.2014.11.002>.
- Malt, B. C., Sloman, S. A., Gennari, S., Shi, M., & Wang, Y., (1999). Knowing versus naming: Similarity and the linguistic categorization of artifacts. *Journal of Memory and Language*, 40(2), 230–262. <https://doi.org/10.1006/jmla.1998.2593>.
- Miton, H., & Morin, O., (2021). Graphic complexity in writing systems. *Cognition*, 214, 104771. <https://doi.org/10.1016/j.cognition.2021.104771>.
- Mollica, F., Bacon, G., Regier, T., & Kemp, C. (2020). Grammatical marking and the tradeoff between code length and informativeness. In S. Denison, M. Mack, Y. Xu, & B. C. Armstrong (Eds.), *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Nelson, D. L., McEvoy, C. L., & Schreiber, T. A., (2004). The University of South Florida free association, rhyme, and word fragment norms. *Behavior Research Methods, Instruments, & Computers*, 36(3), 402–407.
- Newberry, M. G., Ahern, C. A., Clark, R., & Plotkin, J. B., (2017). Detecting evolutionary forces in language change. *Nature*, 551(7679), 223–226. <https://doi.org/10.1038/nature24455>.
- Nölle, J., Staib, M., Fusaroli, R., & Tylén, K., (2018). The emergence of systematicity: How environmental and communicative factors shape a novel communication system. *Cognition*, 181, 93–104.
- Piantadosi, S. T., Tily, H., & Gibson, E., (2011). Word lengths are optimized for efficient communication. *Proceedings of the National Academy of Sciences*, 108(9), 3526–3529. <https://doi.org/10.1073/pnas.1012551108>.
- Rabinovich, E., Xu, Y., & Stevenson, S. (2020). The typology of polysemy: A multilingual distributional framework. *arXiv preprint, arXiv:2006.01966*.
- Ramiro, C., Srinivasan, M., Malt, B. C., & Xu, Y., (2018). Algorithms in the historical emergence of word senses. *Proceedings of the National Academy of Sciences*, 115(10), 2323–2328. <https://doi.org/10.1073/pnas.1714730115>.
- Raviv, L., de Heer Kloots, M., & Meyer, A., (2021). What makes a language easy to learn? A preregistered study on how systematic structure and community size affect language learnability. *Cognition*, 210, 104620. <https://doi.org/10.1016/j.cognition.2021.104620>.
- Raviv, L., Meyer, A., & Lev-Ari, S., (2019). Larger communities create more systematic languages. *Proceedings of the Royal Society B: Biological Sciences*, 286(1907), 20191262. <https://doi.org/10.1098/rspb.2019.1262>.
- Reali, F., & Griffiths, T. L., (2010). Words as alleles: Connecting language evolution with Bayesian learners to models of genetic drift. *Proceedings of the Royal Society B: Biological Sciences*, 277(1680), 429–436. <https://doi.org/10.1098/rspb.2009.1513>.
- Regier, T., Carstensen, A., & Kemp, C., (2016). Languages support efficient communication about the environment: Words for snow revisited. *PLOS ONE*, 11(4), 1–17. <https://doi.org/10.1371/journal.pone.0151138>.
- Regier, T., Kemp, C., & Kay, P. (2015). Word meanings across languages support efficient communication. In *The handbook of language emergence* (pp. 237–263). Hoboken, NJ: Wiley. <https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118346136.ch11>
- Rosenfeld, A., & Erk, K. (2018). Deep neural models of semantic shift. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)* (vol. 1, pp. 474–484). Stroudsburg, PA: Association for Computational Linguistics.
- Ryskina, M., Rabinovich, E., Berg-Kirkpatrick, T., Mortensen, D., & Tsvetkov, Y., (2020). Where new words are born: Distributional semantic analysis of neologisms and their semantic neighborhoods. *Proceedings of the Society for Computation in Linguistics*, 3(1), 43–52. <https://doi.org/10.7275/1jra-8m83>.

- Sankoff, G., (2018). Language change across the lifespan. *Annual Review of Linguistics*, 4(1), 297–316. <https://doi.org/10.1146/annurev-linguistics-011817-045438>.
- Sapir, E., (1912). Language and environment. *American Anthropologist*, 14(2), 226–242.
- Scott-Phillips, T. C., & Kirby, S., (2010). Language evolution in the laboratory. *Trends in Cognitive Sciences*, 14(9), 411–417. <https://doi.org/10.1525/aa.1912.14.2.02a00020>.
- Segovia-Martín, J., Walker, B., Fay, N., & Tamariz, M., (2020). Network connectivity dynamics, cognitive biases, and the evolution of cultural diversity in round-robin interactive micro-societies. *Cognitive Science*, 44(7), e12852. <https://doi.org/10.1016/j.tics.2010.06.006>.
- Smith, K., (2020). How culture and biology interact to shape language and the language faculty. *Topics in Cognitive Science*, 12(2), 690–712. <https://doi.org/10.1111/tops.12377>.
- Smith, K., Tamariz, M., & Kirby, S. (2013). Linguistic structure is an evolutionary trade-off between simplicity and expressivity. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.), *Proceedings of the 35th Annual Conference of the Cognitive Science Society* (pp. 1348–1353). Austin, TX: Cognitive Science Society.
- Srinivasan, M., & Rabagliati, H., (2015). How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy. *Lingua*, 157, 124–152. <https://doi.org/10.1016/j.lingua.2014.12.004>.
- Steels, L., & Szathmáry, E., (2018). The evolutionary dynamics of language. *Biosystems*, 164, 128–137. <https://doi.org/10.1016/j.biosystems.2017.11.003>.
- Steinert-Threlkeld, S., & Szymanik, J., (2020). Ease of learning explains semantic universals. *Cognition*, 195, 104076. <https://doi.org/10.1016/j.cognition.2019.104076>.
- Thompson, B., Roberts, S. & Lupyan, G. (2018). Quantifying semantic similarity across languages. In *Proceedings of the 40th Annual Conference of the Cognitive Science Society* (pp. 2554–2559). Austin, TX: Cognitive Science Society.
- Tinitis, P., Nölle, J., & Hartmann, S., (2017). Usage context influences the evolution of overspecification in iterated learning. *Journal of Language Evolution*, 2(2), 148–159. <https://doi.org/10.1093/jole/lzx011>.
- Uegaki, W. (in preparation). *NAND and the Communicative efficiency model. Unpublished manuscript. <https://semanticsarchive.net/Archive/2M0YTUzN>.
- Whorf, B. L. (1956). Science and linguistics. In J. Carroll (Ed.), *Language, thought, and reality: Selected writings of Benjamin Lee Whorf* (pp. 207–219). Cambridge, MA: MIT Press.
- Winters, J., Kirby, S., & Smith, K., (2015). Languages adapt to their contextual niche. *Language and Cognition*, 7(3), 415–449. <https://doi.org/10.1017/langcog.2014.35>.
- Xu, Y., Duong, K., Malt, B. C., Jiang, S., & Srinivasan, M., (2020a). Conceptual relations predict colexification across languages. *Cognition*, 201, 104280. <http://www.sciencedirect.com/science/article/pii/S0010027720300998>
- Xu, Y., Liu, E., & Regier, T., (2020b). Numeral systems across languages support efficient communication: From approximate numerosity to recursion. *Open Mind*, 4, 57–70. <https://doi.org/10.1016/j.cognition.2020.104280>.
- Zaslavsky, N., Kemp, C., Tishby, N., & Regier, T., (2019a). Color naming reflects both perceptual structure and communicative need. *Topics in Cognitive Science*, 11(1), 207–219. <https://onlinelibrary.wiley.com/doi/pdf/10.1111/tops.12395>.
- Zaslavsky, N., Maldonado, M. & Culbertson, J. (2021). Let's talk (efficiently) about us: Person systems achieve near-optimal compression. In *43rd Annual Meeting of the Cognitive Science Society* (pp. 938-944). Austin, TX: Cognitive Science Society.
- Zaslavsky, N., Regier, T., Tishby, N., & Kemp, C. (2019b). Semantic categories of artifacts and animals reflect efficient coding. In A. K. Goel, C. M. Seifert, & C. Freksa (Eds.), *Proceedings of the 41st Annual Conference of the Cognitive Science Society* (pp. 1254–1260). Montreal, Canada: Cognitive Science Society.
- Zaslavsky, N., Regier, T., Tishby, N., & Kemp, C. (2020). Semantic categories of artifacts and animals reflect efficient coding. In *Proceedings of the Society for Computation in Linguistics 2020* (pp. 80–81). Stroudsburg, PA: Association for Computational Linguistics.

Table A1

An example of scoring a message in an experiment on the axes of complexity (cognitive cost) and ambiguity (communicative cost). Given a round where the meaning is TASK, the scores depend on the last most recent lexification of both TASK, its target pair member, JOB, and the new signal used to express TASK. Using different signals to distinguish the two meanings increases complexity but decreases ambiguity. The sum of these scores is the highest if the participants use entirely different signals for the meanings (last row), or if they change the signals without converging on stable associations (second to the last row)

Current Meaning	Last TASK	Last JOB	New Signal	Complexity	Ambiguity
TASK	nopo	nopo	nopo	0	1
TASK	nopo	mumi	nopo	1	0
TASK	nopo	nopo	mumi	1	1
TASK	nopo	mumi	mumi	1	2
TASK	nopo	mumi	fita	2	1

Appendix: Details on the complexity-ambiguity calculation

The procedure to produce Fig. 9 in Section 7 is the following. Each message produced by a dyad after the burn-in period, which contains a target meaning, is assigned a cognitive cost (complexity) score and communicative cost (ambiguity) score. As a simplification, we consider the results by dyads, ignoring who sent a given message within a dyad. Only messages containing target meanings are scored, as distractor meanings lack synonyms in the meaning spaces (of Experiment 1, which we analyze here).

The cognitive cost score is set to 0, if a given utterance does not increase the complexity of the language, within the target pair: if the last reference to the same meaning used the same signal, and the last reference to the synonym (target pair member) of the current meaning also used the same signal. Using a different signal or distinguishing the current meaning from its synonym increases complexity by 1 point each.

Communicative cost is scored as 0 if the last reference to the same meaning used the same signal, and the last reference to the synonym of the current meaning used a different signal. Using a different signal or colexifying the current meaning both increase ambiguity (communicative cost, chance of misinterpreting), so doing either costs 1 point each. Given seven signals and 10 meanings, some meanings are bound to be colexified. The minimal sum of these two scores for any given utterance is 1: It is impossible to be simultaneously maximally simple and maximally informative. The highest sum of scores is 3, which may result in random assignments of signals to meanings, or intentionally misleading behavior (see Table A1 for an example).

In addition to reanalyzing the results of our human experiments, we implement a simple agent-based model that replicates our experimental setup (up to seven signals, equally frequently occurring 10 meanings; 135 rounds). We use the results from this model—analyzed using the same coding procedure—to provide meaningful dimensionality to the human results in Fig. 9. The agent based model is constructed as a single-agent system, representing a dyad, which plays a simple naming game. The model has two parameters: the number of signals allowed (1–7) and the naming strategy. At each round, the agent assigns a signal to a given

meaning. In the most simple case, it only ever produces a single signal, leading to a “degenerate” language (see Kirby et al., 2015). Other strategies include:

- random signaling;
- fixed assignment to as many meanings as possible, with perfect memory (and random assignment for meanings for which there are not enough signals);
- fixed assignment with perfect memory, but colexifies pairs of meanings (others assigned randomly);
- fixed assignment with perfect memory, but colexifies pairs of meanings (others assigned randomly, but avoiding colexification with the pairs where enough signals available);
- rational (in the sense of Frank & Goodman, 2012), keeping the entire lexification history in memory;
- rational, but keeping only the most recent lexification;
- intentionally misleading, that is, the inverse of the rational strategies;
- and also combinations of the above.

Each combination of these parameters is repeated 20 times. The results are analyzed exactly the same way as the human experiments (including the designation of a “burn-in” period) and are graphed as gray background blocks in Fig. 9.