languages

MDPI

Article

Shared Context Facilitates Lexical Variation in Sign Language Emergence

Katie Mudd [1,*], Connie de Vos [2] and Bart de Boer [1]

[1] AI Lab, Vrije Universiteit Brussel, 1050 Elsene, Belgium; bart@ai.vub.ac.be
[2] Tilburg Center for Cognition and Communication, Tilburg University, 5037 AB Tilburg, The Netherlands; c.l.g.devos@tilburguniversity.edu
* Correspondence: katie.mudd@ai.vub.ac.be

Abstract: It has been suggested that social structure affects the degree of lexical variation in sign language emergence. Evidence from signing communities supports this, with smaller, more insular communities typically displaying a higher degree of lexical variation compared to larger, more dispersed and diverse communities. Though several factors have been proposed to affect the degree of variation, here we focus on how shared context, facilitating the use of iconic signs, facilitates the retention of lexical variation in language emergence. As interlocutors with the same background have similar salient features for real world concepts, shared context allows for the successful communication of iconic mappings between form and culturally salient features (i.e., the meaning specific to an individual based on their cultural context). Because in this case the culturally salient features can be retrieved from the form, there is less pressure to converge on a single form for a concept. We operationalize the relationship between lexical variation and iconic affordances using an agent-based model, studying how shared context and also population size affects the degree of lexical variation in a population of agents. Our model provides support for the relationship between shared context, population size and lexical variation, though several extensions would help improve the explanatory power of this model.

Keywords: sign language; social structure; lexical variation; agent-based modeling

check for updates

Citation: Mudd, Katie, Connie de Vos, and Bart de Boer. 2022. Shared Context Facilitates Lexical Variation in Sign Language Emergence. *Languages* 7: 31. https://doi.org/10.3390/languages7010031

Academic Editors: Wendy Sandler, Mark Aronoff, Carol Padden, Juana M. Liceras and Raquel Fernández Fuertes

Received: 4 October 2021
Accepted: 24 January 2022
Published: 10 February 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

In sign language emergence, linguistic variation at the lexical level appears to be the default, where synonyms for a word coexist within a population. However, over time, certain pressures seem to push towards lexical uniformity (Meir et al. 2012). We can thus imagine two extreme cases as languages evolve: one in which the variation present in language emergence is fully retained and a second where all the variation is lost in favor of uniformity. What are the pressures that may drive languages away from linguistic variability? It has been proposed that the communicative context in which languages are used shapes the features of a language (Lupyan and Dale 2010; Trudgill 2011; Wray and Grace 2007). Specifically in this paper, we explore how shared social and psychological information makes it possible to use iconic signs and how this may be a driving factor in retaining the lexical variation present in language emergence.

Traditionally, in the study of lexical variation in spoken languages, it has been assumed that true synonyms do not exist (Clark 1987). Rather, it is accepted that synonyms for a concept coexisting in a population would be conditioned by sociolinguistic and pragmatic factors. However, in the first stage of language emergence, where individuals improvise forms to refer to concepts, it appears that synonyms can coexist. It is possible that the iconic affordances of the manual modality facilitate the coexistence of synonyms in a population. Without data on the emergence of spoken languages, it is unclear how iconic affordances play a role in their emergence. For these reasons, in this paper, we focus on the emergence

of sign languages and how different factors influence the degree of lexical variability across a population.

de Vos (2011) suggests that a high degree of variation at the lexical level may be characteristic of sign languages used in communities with a small population size and a high degree of shared context. Here, we refer to sign languages in such communities as *shared sign languages*, following Nyst (2012). For instance, Ergin et al. (2021) report that the shared sign language Central Taurus Sign Language is "remarkable in its mixture of more or less conventionalized[1] signs or sign sequences, improvised sign sequences, and competing lexical variants".

Similarly, in Kata Kolok, a sign language which emerged in a relatively small, insular village community in northern Bali due to a high incidence of hereditary deafness (de Vos 2012; Marsaja 2008), a high degree of lexical variation has been observed (Mudd et al. 2020); in response to a picture description task, up to nine lexical variants for a stimulus were produced, while other stimuli in the task elicited a uniform response (Mudd et al. 2020). This high degree of lexical variation seems typical of shared sign languages and has also been reported in Al-Sayyid Bedouin Sign Language (ABSL) (Meir et al. 2012), San Juan Quiahije Chatino Sign Language (SJQCSL) (Hou 2016) and Providence Island Sign Language (PISL) (Washabaugh 1986), to name a few.

In contrast, sign languages used predominantly by a large and dispersed group of deaf individuals, most of whom are born to hearing parents, or *Deaf community sign languages* (Meir et al. 2010; Mitchell and Karchmer 2004), appear to exhibit lower levels of lexical variation than shared sign languages (Meir et al. 2012). However, it should be noted that this claim is mostly based on anecdotal evidence (for one exception, see Washabaugh 1986). What can be said is that variation in this category of sign languages is typically structured along different sociolinguistic lines than in shared sign languages, as variation is often the result of schooling practices (Meir et al. 2010). For example, gender-based school segregation in Dublin has resulted in a gendered Irish Sign Language lexicon (LeMaster 2006), and different varieties of American Sign Language (ASL) have emerged due to race-based school segregation (McCaskill et al. 2011).

There undoubtedly also exists structured variation in shared sign languages, such as within families (Sandler et al. 2011) and also along sociolinguistic lines (Mudd et al. 2020). Despite evidence of structured variation, it seems like the degree of lexical variation in shared sign languages is higher within a small community across the board, with frequent interlocutors using different forms to refer to a concept (de Vos 2011). Crucially, despite the existence of multiple forms associated with a concept, signers are able to understand each other. Tkachman and Hudson Kam (2020) posit that a decrease in lexical variation may only be necessary in cases where communication fails. This may be less the case in shared sign language, where pressures for convergence seem to be somewhat alleviated. Meanwhile, in Deaf community sign languages, frequent interlocutors seem to have more synchronized lexical preferences, with higher degrees of variation evident when comparing larger, more dispersed subgroups of the community. What aspects of shared signing communities could reduce the pressure for linguistic uniformity?

One possibility that we explore in the present study is that shared context alone, allowing for the use of iconic forms, may be sufficient to maintain high degrees of lexical variation in a community (Sandler et al. 2011; Tkachman and Hudson Kam 2020). In tight-knit communities, individuals can make use of shared social and psychological information, facilitating the use of strategies such as pointing to concepts and using iconic signs (de Vos 2011). Iconic signs, in which aspects of a sign's form resemble aspects of that sign's meaning (Dingemanse et al. 2015), would only be successfully communicated (if not already conventionalized) when individuals share the same salient features (specific to the individual) associated with a concept (the entity or concept in the real world). For instance, in the shared sign language ABSL, the sign for kettle was shown to differ across families, but within families, members were uniform in their productions (Sandler et al. 2011). Regarding this variation, Sandler et al. (2011) state: "It is likely that all the different versions

would be intelligible across the community, due to iconicity, context, or the existence of synonymy in the signers' mental lexicons—possibly all the of the above". We refer to these as *productive synonyms*, i.e., variants that may be used interchangeably, in contrast to *perceptual synonyms*, i.e., variants which signers may be aware of in a more abstract sense but not use (Mudd et al. 2020).

Figure 1 shows three signs for *pig* used in the Kata Kolok community. Many villagers in this community make a living as farmers, and this is reflected in the iconic motivations underlying the forms produced for PIG-1 and PIG-2. Given that the members of this community share a high degree of cultural context, it is probable that individuals exploit iconic mappings, understanding each other by retrieving the meaning (comprised of culturally salient features) from the form even if they have not seen or produced the form themselves. On the other hand, when shared context is not available (i.e., individuals are from different backgrounds and have different experiences), there is no advantage to using iconic signs, as the culturally salient features of interlocutors are different. Continuing with this example, imagine someone from a different community who does not have experience with farming. The underlying iconic motivations related to this practice will be meaningless to them, and therefore, the meaning comprised of culturally salient features to the individual in the farming community which are expressed in the form (e.g., PIG-1, whose underlying iconic motivation refers to how pigs are killed) would not be understood unless the mapping is learned. As explained by Occhino et al. (2017), iconicity is subjective as it is dependent on one's language and culture-specific experience.



Figure 1. Three variants for *pig* in Kata Kolok produced in response to a picture description task (Lutzenberger et al. 2021; Mudd et al. 2020). The iconic motivation underlying PIG-1 is how a pig is killed, for PIG-2 is how a pig eats and for PIG-3 is the ears of a pig. It is clear that the cultural context of the Kata Kolok community has shaped lexical preferences, illustrated with iconic signs (i.e., mappings between culturally salient features and forms). For instance, the iconic motivations of PIG-1 and PIG-2 stem from farming practices in the community.

Here, we aim to operationalize the relationship between shared context (allowing for iconic mappings) and lexical variation using an agent-based model. In our model, the language representation is adapted from the semiotic triangle (Ogden and Richards 1925). Traditionally, the semiotic triangle consists of a referent (something concrete or abstract referred to in a particular instance of a conversation), a meaning (a representation of that referent by a given individual) and a form (the signal conveyed) (terminology following Steels and Kaplan 1999, definitions following Vogt 2002; Vogt and Divina 2007). The relationship between these components has been used to study the symbol grounding problem (Harnad 1990), i.e., the problem that symbols are internal representations but need to be linked to entities in the real world (Vogt 2002).

In the semiotics literature, there is a heavy emphasis on the conventionalized and/or arbitrary link between the form and referent (see Pierce 1931), which is unsurprising considering the long-held assumption that arbitrariness is a design feature of language (de Saussure 1916; Hockett and Hockett 1960). However, the emphasis on arbitrariness has been reduced due to the overwhelming presence of iconic forms in sign languages as well as in spoken languages (see Perniss et al. 2010 for a review). It should be noted that the

role of iconicity in language emergence may differ in signed and spoken languages, given the different affordances of the modalities, which may have ramifications on the degree of lexical variability.

In the present study, we adapt the semiotic triangle to reflect what we posit is representative of the linguistic situation in sign language emergence. The semiotic triangle presented here consists of three components: (1) a concept, i.e., an abstract notion; (2) culturally salient features, i.e., culturally salient features of a concept; and (3) a form, i.e., the signal conveyed. For example, a hypothetical semiotic triangle from an individual in the Kata Kolok community could consist of (1) the concept *pig*, an abstract representation of the animal; (2) culturally salient features of a pig in this farming community, such as how a pig is killed and how a pig eats; and (3) the form PIG-1 (see Figure 1), whose iconic motivation stems from how a pig is killed. Notably, the inclusion of culturally salient features in the language model allows for the use of iconic mappings between the culturally salient features and the form. As such, the original contributions of this model are the introduction of culturally salient features and the *iconic–inferential pathway* (presented in the right triangle in Figure 2). In addition to the *conventional link* between form and concept, the iconic–inferential pathway goes from form to culturally salient features to concept (or vice versa). Here, an individual can make use of the culturally salient features (unique to them depending on their culture and experiences), which can be retrieved from the form given that cultural knowledge is shared.

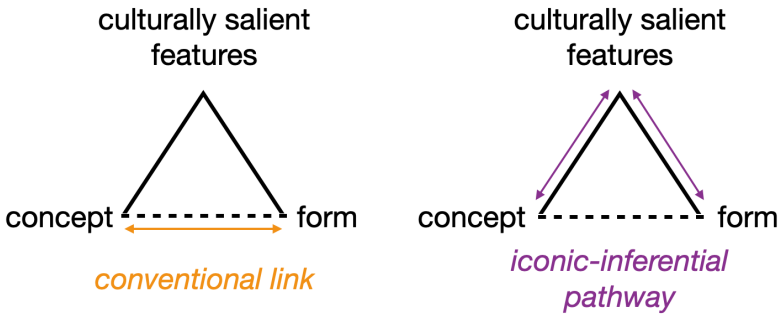


Figure 2. The semiotic triangle used in the current study, consisting of a concept, culturally salient features and a form. The triangle on the left shows the traditional view of the relationship, in which an arbitrary link between the form and concept are made. Depicted in the triangle on the right, we present an alternative route to connecting the form to the concept, through culturally salient features, which we call the iconic–inferential pathway. Figure based on Vogt and Divina (2007), adapted from Ogden and Richards (1925), updated with terminology used in the current study.

Here, we provide an example of how these pathways could be used in interaction with the example of pig again from the Kata Kolok community, using a hypothetical conversation between individual A and individual B, both from this community. In conversation, individual A uses the sign PIG-1 (iconic motivation referring to how a pig is killed). However, individual B is not familiar with this form and, using the conventional link (form to concept), does not know at this stage what individual A is referring to. Subsequently, individual B uses the iconic–inferential pathway to consider if the form produced by individual A overlaps with the culturally salient features of any concept. Because individual B is from the Kata Kolok community, where individuals have knowledge about farming, including the way in which pigs are killed, individual B recognizes that the form PIG-1 produced by individual A refers to how a pig is killed, and thus likely refers to pig. In this way, when individuals share a cultural context, the iconic–inferential pathway can serve as a supporting route in case the conventional pathway fails. In the event that neither of these pathways lead individual B to the concept *pig*, it is probable that these individuals will need to initiate repair in order to understand each other. Though many strategies may be used, one option would be for individual B to learn the form produced by individual A. Although in the operationalization of this model the conventional link has priority over the

iconic–inferential pathway, in the real world, meaning can also undoubtedly be inferred using the iconic–inferential pathway prior to the conventional link or a combination of both.

This theory generates a prediction about the level of iconicity present in different types of communities. Frishberg (1975) showed that in ASL, a Deaf community sign language, signs tend to become less iconic over time. Pleyer et al. (2017) point out that studies from young sign languages and homesign systems show that "signs gradually shed their iconic mapping", potentially in favor of facilitating a larger vocabulary (Gasser 2004). However, what about for shared sign languages? Does the level of iconicity remain high or decrease over time? We predict that in shared sign languages, the level of iconicity will remain relatively high because iconic forms are successfully communicated, as community members share a high degree of cultural context. In contrast, in Deaf community sign languages, we predict that iconicity will decrease, as found by Frishberg (1975) for ASL, because in these larger communities, individuals typically come from diverse backgrounds. Therefore, retrieving culturally salient features from the form will not be useful when communicative partners do not share cultural context. Rather, individuals are more likely to adapt their form moving closer to the form of their communicative partner. This helps them to successfully communicate, as their forms move towards becoming aligned. However, as individuals do not likely share a cultural context (and hence likely have different salient features), adapting one's form would typically result in a move away from its initial highly iconic state. Iconicity is often talked about on a large scale, irrespective of individual experience. While iconic affordances can be grounded in human experience (e.g., men have beards), it must be stressed that iconicity remains subjective (Occhino et al. 2017). Thus, here, iconicity is considered on an individual level, as opposed to across entire communities where individuals may not share much cultural context.

In sum, we propose that in communication individuals may exploit an iconic–inferential pathway, making use of iconic mappings between a form and culturally salient features if a conventional pathway is not available. In communities such as shared signing communities where individuals share psychological and social information, we predict that communicative partners will successfully communicate using the iconic–inferential pathway if the conventional pathway fails. Because communication can succeed using these two routes, lexical variation should remain high, as well as the degree of iconicity in the community. On the other hand, in communities such as those with Deaf community sign languages, because there is less shared information, the iconic–inferential pathway is less useful. Hence, in the case of failure using the conventional pathway, individuals are more likely to proceed to adapt their lexical form in order to be understood. Hence, we predict that communities with little shared context will move towards lexical uniformity and low degrees of iconicity.

In addition to shared context, it has been proposed that population size may affect linguistic features (Lupyan and Dale 2010; Wray and Grace 2007). In sign languages, anecdotal evidence suggests that small populations exhibit a higher degree of lexical variation than large populations (Meir et al. 2012). The relationship between population size and lexical variation has been supported by a recent computational model (Tkachman and Hudson Kam 2020), though previous computational models have found that conventions emerge faster in smaller populations (Baronchelli et al. 2006). Although not the main focus of this study, we also consider the effect of population size on the degree of lexical variability, as typically shared sign languages emerge in smaller populations, and Deaf community sign languages emerge in larger populations. Modeling shared context and population size may help to tease apart the contribution of each on the degree of lexical variation.

In the next section, we describe how this theory is operationalized using an agent-based model. Following this, we begin the results section with two example model runs focusing on the results of the language game component of the model. Then, we study the effect of shared context on lexical variation by altering the number of groups in the model, which determines how many agents share the same cultural context. Concluding the results section, we briefly consider the effect of population size on lexical variation. Finally, in

the discussion section, we first focus on comparing the model results to the evidence from variation in signing communities. Then, we discuss the limitations of this model and how it can be extended to account for these limitations.

2. Model Description

The model description is inspired by the ODD (Overview, Design concepts, Details) protocol for describing agent-based simulations (Grimm et al. 2006; Grimm et al. 2010). The description has been adapted to include links between the model and real world examples, to hopefully make for a more understandable model description. The model was implemented in Mesa, a Python framework for agent-based modeling (Kazil et al. 2020). The model code is available on figshare: https://doi.org/10.6084/m9.figshare.15163872.v1, accessed on 23 January 2022.

Purpose. The purpose of this model is to investigate how shared context affects lexical variation in sign language emergence. As shown in Figure 3, the agent-based model takes the following values as input parameters:

- The number of concepts (*n_concepts*);
- The number of bits (*n_bits*): the number of bits (0 or 1) in the culturally salient features and form (i.e., the length of a word);
- The number of agents in the model (*n_agents*) (i.e., the population size);
- The number of groups (*n_groups*): agents are assigned to a group, which determines which features of a referent are culturally salient to an agent;
- The initial degree of overlap between the culturally salient features and form (*initial_degree_of_overlap*) (the parameter simulating iconicity);
- The number of time steps in the model (*n_steps*).

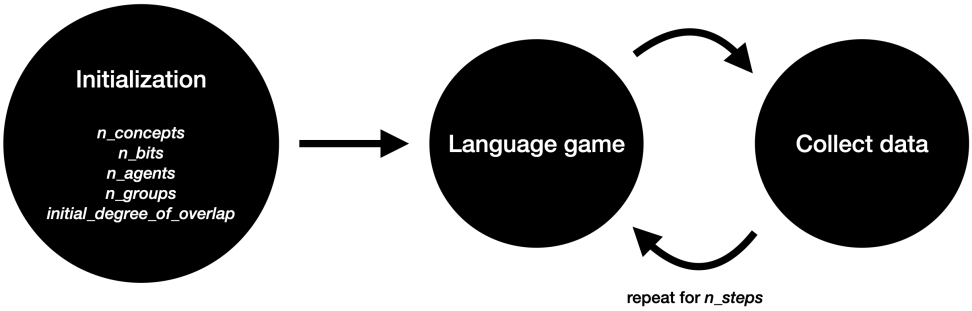


Figure 3. Visualization of the steps and parameters in the agent-based model. During the initialization phase, the number of groups (*n_groups*) determines how many subsets of the population have the same set of identical culturally salient features associated with concepts. Then, a number of agents are created (*n_agents*). Each agent is randomly assigned to a group, and their language representation is set given the following parameters: the number of concepts (*n_concepts*), the number of bits (*n_bits*) and the initial degree of overlap between the culturally salient features and form (*initial_degree_of_overlap*). At each time step, all agents initiate a language game (i.e., they take a turn as the sender). At the end of each time step, data on the mean degree of iconicity and the mean lexical variability are calculated. The model continues for a number of steps (*n_steps*).

Entities, state variables and scales. The only entity in the model is the agent, which is the entity in the model that represents one individual in the real world. Agents consist of a unique id and a group that they are assigned to during the initialization stage (first stage of the model). Furthermore, each agent has a language representation which is explained in the initialization below. Figure 4 shows an example of an agent.

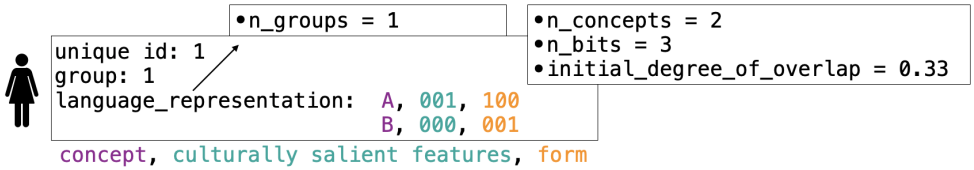


Figure 4. Example of an agent.

The agent's unique id is 1 as it is the first agent created in this run of the model. In this example, there is only one group (*n_groups* = 1), so the agent is assigned to group 1. As there is only one group, all agents in the model would have the same culturally salient features corresponding to each concept. This is akin to individuals of a population having shared social and psychological information, thus they are likely to have similar notions for a given concept. Some examples of concepts in real life are *pig*, *tree* and *destiny*, as discussed in the introduction. In the example in the figure, there are two concepts (*n_concepts* = 2); each concept is associated with culturally salient features and a form, both of which are made up of three bits (*n_bits* = 3). For each bit of the form, the probability that it will have the same value as the corresponding bit in the culturally salient features is determined by *initial_degree_of_overlap*. Hence, the form, corresponding in real life to a sign produced or a word uttered, is determined by the association with the culturally salient features. The idea is that when individuals initially improvise forms, the forms often bear some degree of resemblance to culturally salient features of the concept. For example, in Kata Kolok, signs for *pig* refer to how a pig is killed, how a pig eats or a pig's ears—features that are culturally salient in the Kata Kolok community.

Process overview and scheduling. The set-up of the model is outlined in initialization below. After the initialization phase, each time step consists of the processes outlined in Table 1. For details of these processes, see the Submodels sections. A schematic overview of the order of processes and parameter input is provided in Figure 3.

Initialization. For each group (*n_groups*), a bit vector of length *n_bits* is generated per concept (*n_concepts*). Following the example provided in Figure 4, the culturally salient features associated with concept A is 001 and concept B is 000. In the real world, this could be analogous to two concepts, say, *pig* and *butterfly*, which have different culturally salient features (dependent on the background of a person), such as wings for a butterfly and pigs rolling in mud or how they are killed in farming. Roughly, the string of 0s and 1s representing the culturally salient features can be thought of as a unique representation of the characteristics of that concept, given the group one is in.

Each agent has a language representation which consists of, for each concept, a set of culturally salient features and a form, as shown in Figure 4. *n_concepts* determines the number of concepts in the language representation. This is akin to the number of words in a person's vocabulary. Each concept is associated with culturally salient features and a form, each consisting of a number of bits (0 s and 1 s), determined by *n_bits*. The culturally salient features corresponding to each concept are fixed, based on the group that the agent belongs to. The culturally salient features and concepts are never updated or changed throughout the simulation. Only the forms can be updated. The idea here is a simplification of reality, in which an individual is born in a certain context, determining what features are salient culturally for the entirety of their life (e.g., in communities where farming is practiced, one's concept of an animal is likely related to how that animal is farmed). However, despite this, the form (produced sign or uttered word) can change over the course of one's life (e.g., I may say "rad", "radical" or "cool" to refer to the same concept).

Table 1. Processes, scheduling, pseudo-code and parameters.

Process	Pseudo-Code	Parameters
Initialization	**for each** group G in *n_groups* **for each** concept C in *n_concepts* set the culturally salient features for concept C in group G to a list of length *n_bits* of randomly chosen 0s and 1s create a population with size *n_agents* **for each** agent A in the population randomly assign agent A to a group set all culturally salient features based on assigned group **for each** concept C **for each** culturally salient features F associated with concept C **for each** bit B of culturally salient features F **with** probability *initial_degree_of_overlap* set bit B of corresponding form to bit B of culturally salient features F **else** set bit B of corresponding form to 0 or 1 with equal probability	*n_groups* *n_agents* *n_concepts* *n_bits* *initial_degree_of_overlap*
Step	repeat for *n_steps*	*n_steps*
Language game	**for each** agent A in the population randomly choose another agent and play the language game	
Collect data	calculate the mean iconicity and lexical variability across the population	

Iconicity is represented in the model by the similarity between the forms (sign produced or word uttered) and the culturally salient features for a concept. For example, if a butterfly's wings are salient in one's culture and the sign for butterfly refers to the wings of the insect, then the similarity between the culturally salient features and the form is strong and thus highly iconic for individuals with the same background. In the model, to understand how iconicity affects lexical variability, the parameter determining the degree of overlap between forms and culturally salient features is varied. This is operationalized in the model in the following way: The relationship between each bit of the culturally salient features and the form is determined by the *initial_degree_of_overlap*, such that the probability that the form's bit is the same as the bit from the culturally salient features is equal to the value of *initial_degree_of_overlap*. For a bit of the form that is not chosen to be the same as the bit of the culturally salient features in the initial event, then that bit is randomly[2] assigned a 0 or 1. As such, a non-iconic form does not have a structured relationship between the form and the culturally salient features; rather, its relationship is arbitrary. If the *initial_degree_of_overlap* is set to 1, then there is a 100% chance that each bit of the form will be the same as that of the culturally salient features. If the *initial_degree_of_overlap* is set to 0, then each bit of the form is randomly assigned a 0 or 1.

To illustrate with an example following Figure 4, take concept A, which is associated with the culturally salient features 001. Before assigning the forms, the language representation looks like this: A, 001, NA NA NA, with 001 referring to the culturally salient features and NA NA NA referring to placeholders for each bit of the form. Starting with the first bit of culturally salient features (0), there is a 33% chance that the corresponding bit of the form will be identical to the bit of the culturally salient features in this initial event. The outcome of this event is that the bit of the culturally salient features and of the form are not identical. From here, a new event occurs, randomly assigning a 0 or 1 to this bit; a 1 is randomly assigned (note that at this stage a 0 could also be chosen randomly). Now, the language representation looks like this: A, 001, 1 NA NA. The same process is repeated to determine the second bit of the form, and here the outcome of this event is that this bit is identical, i.e., the second bit of the form is set to 0 as the second bit of the culturally salient features is 0. Finally, this process is repeated a third time, and here the outcome of this event is that the bit of the culturally salient features and of the form are not identical. From here, a new event occurs, randomly assigning a 0 or 1 to this bit; here, it happens to be a 0 that is chosen (note that at this stage, a 1 could also be chosen randomly). Thus, the final language representation of this agent for concept A is: A, 001, 100.

Submodel Language game. The language games consist of two agents interacting—a sender and a receiver, simulating a simplified exchange between two individuals. At each time step in the model, all agents take one turn as a sender in the language game. As shown in Figure 5, the language game consists of four steps. First, the sender randomly chooses a concept and produces the corresponding form. In Figure 5, the sender has randomly chosen concept A. In real life, this would be analogous to an individual wanting to communicate about a given concept and producing the corresponding sign or uttering the corresponding word.

Second, in the language game, the receiver selects the form which is closest to the form of the sender, by calculating the distance between the sender's form to all of the forms of the receiver. Crucially, in this model, the distance is calculated by comparing the bits at the same index. In the event of a tie between two or more forms as having most in common with the form of the sender, a form that tied is randomly chosen. Following the example provided in Figure 5, the sender's form is 100. The distance to the receiver's first form 001 is 2/3 and the distance to the receiver's second form is 100 is 0/3, so the second form is selected. The concept of the selected form of the receiver is then compared to the chosen concept of the sender. If the concept of the sender and receiver are the same, then the language game is over and no update is made. When the language game succeeds here, we refer to this as *form success*. However, if the concepts of the sender and receiver do not match (as is the case in the example presented in Figure 5 where the sender chose concept

A and the receiver's closest match is concept B), then the language game proceeds to the third step. Success at this step of the language game represents the conventional link or memorizing the association between a concept and a form. Typically in language games, it is the conventional link that is modeled.

This next step presents the original contribution of this model, which models the ability of individuals to make use of iconic affordances. In this third step, the form of the sender is compared to all the sets of culturally salient features of the receiver. As performed in step two between forms, the distances between the form of the sender and all of the sets of culturally salient features of the receiver are calculated, and the closest culturally salient features are selected. Again, following the example in Figure 5, the sender's form is 100. The distance to the receiver's first culturally salient features 001 is 2/3, and the distance to the receiver's second culturally salient features is 000 is 1/3, so the second culturally salient features are selected. As in step two, the concept of the receiver's selected culturally salient features is compared to the sender's chosen concept. If these concepts are the same, then the language game is over and no update is made. When the language game succeeds here, we refer to this as *culturally salient features success*. Success at this step of the language game represents the iconic–inferential pathway, where a form and concept are linked via the cultural salient features. Crucially, no memorization is required. However, at this stage, if the concepts of the agents do not match (as is the case in the example presented in Figure 5 where the sender chose concept A and receiver's closest match is concept B), then the language game proceeds to the fourth step.

The last step of the language game represents when communication is unsuccessful via the conventional link and the iconic–inferential pathway. In this case, as is typical in language games, one agent updates their form to hopefully allow for successful communication in the future. In real life, this corresponds to aligning speech with an interlocutor. Concretely, in this fourth step, for the sender's chosen concept (concept A), the receiver updates one bit of the form which is different from the form of the sender. If the language game advances to this stage, we call this *bit update*. In Figure 5, the sender's form corresponding to concept A is 100. The receiver's form corresponding to concept A is 001. The bits that are different between the sender's and the receiver's form are identified (the first and third bits), and one is randomly selected to be changed to correspond to the sender's. In the example, the first bit was chosen and is changed to a 1; now, the receiver's form for concept A is 101.

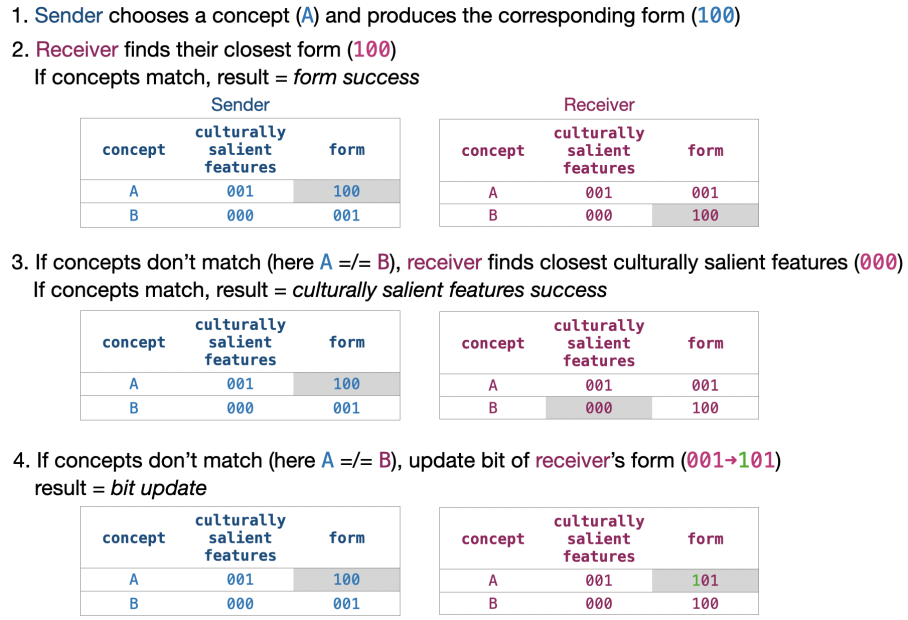


Figure 5. The steps of the language game with an accompanying example.

Submodel Collect data. In the data collection phase of each time step, two calculations are made: the mean degree of iconicity and the mean lexical variability. Calculation examples are demonstrated with the agents in Figure 6.

agent 1

concept	culturally salient features	form
A	001	100
B	000	001

agent 2

concept	culturally salient features	form
A	001	001
B	000	100

Figure 6. Example agents for calculating the mean degree of iconicity and the mean lexical variability.

First, the mean degree of iconicity is calculated for each concept of each agent and averaged across all agents. To calculate the degree of iconicity for a concept, the culturally salient features and the form are compared at each index, with the similarity (or overlap) calculated. For example, for agent 1 in Figure 6, for concept A, the associated culturally salient features are 001 and the form is 100. The similarity between these is 1/3. For concept B, the similarity is 2/3. Thus, the mean degree of iconicity for agent 1 is 1/2.

Next, the mean lexical variability in the population is calculated by comparing all forms for each concept between all pairs of agents in the population. If the agents' forms for a concept are the same, i.e., all bits match at each index, then the distance between the productions is 0. If two agents' forms for a concept are not the same, i.e., the bits differ at one index or more, then the distance between the productions is 1. Thus, the result of the comparison between two agents' forms is binary (distance of 0 or 1)[3]. For each pair, the mean of the distances is taken. We will illustrate this calculation with the agents depicted in Figure 6: For concept A, agent 1's form is 100 and agent 2's form is 001. As these forms differ at the first and last positions, the distance between them is 1. Subsequently, for concept B, agent 1's form is 001 and agent 2's form is 100, which differ at the first and last positions, so the distance between them is 1. Thus, the mean lexical variability between these agents is 1.

3. Results

In this section, we first present results of two single runs in order to explain the measures used and to give an intuition as to what one single run of the model looks like. Here, we explain the results of the language games—that is, for each language game, it is recorded if the game ends at form success (step two from Figure 5), culturally salient features success (step three from Figure 5) or bit update (step four from Figure 5). Additionally, we show the mean degree of iconicity and the mean lexical variability for each run.

Following these examples, the role of shared context is investigated by altering the number of groups (*n_groups*) and the effect of population size is investigated by altering the number of agents (*n_agents*). We consider the effect of these parameters on the mean lexical variability and the mean degree of iconicity. The model simulations presented are of 100 repetitions. The remainder of the parameter explorations can be found in Appendix A, which investigate the effect of the number of concepts (*n_concepts*), the number of bits (*n_bits*) and the initial degree of overlap between the culturally salient features and the form (*initial_degree_of_overlap*).

Additional parameter explorations studying the role of *initial_degree_of_overlap*, *n_concepts* and *n_bits* on lexical variability and iconicity can be found in Appendix A.

3.1. Two Example Runs

To show what one run of the model entails, we present the results from two single model runs. Both model runs differ only in one parameter, the number of groups (*n_groups*), which determines which set of culturally salient features an agent has. The first run

presented consists of one group and the second run presented consists of ten groups. The other parameters are the following:

- The number of concepts (*n_concepts*): 10;
- The number of bits (*n_bits*): 10;
- The number of agents in the model (*n_agents*): 10;
- The initial degree of overlap between the culturally salient features and form (*initial_degree_of_overlap*): 0.9;
- The number of steps in the model (*n_steps*): 2000.

3.1.1. Language Game Results

First, we present a model run consisting of one group (*n_groups* = 1), meaning that all agents belong to the same group. This results in all agents having the same set of culturally salient features.

In the language game step of the model, as shown in Figure 5, there are three ways in which the language game can end: 1. there is a match between the concepts associated with the sender's form and the receiver's closest form to the sender's form (*form success*, step 2 Figure 5); 2. there is a match between the concepts associated with the sender's form and the receiver's closest culturally salient features to the sender's form (*culturally salient features success*, step 3 Figure 5); or 3. for the form of the receiver corresponding to the concept associated with the form communicated by the sender, a bit which does not match the sender's is updated (*bit update*, step 4 Figure 5). These three steps where the language game can end are visualized in Figure 7 for the first 10 stages (left) and over all 2000 model stages (right). To further explain, in each time step, each of the 10 agents initiates 1 language game, which may end in form success, culturally salient features success or bit update. At each time step, the proportions of these language game results are visualized as a barplot. For example, in the run presented in Figure 7, at stage 1 out of the 10 language games played, 8 resulted in form success, 1 resulted in culturally salient features success and 1 resulted in a bit update.

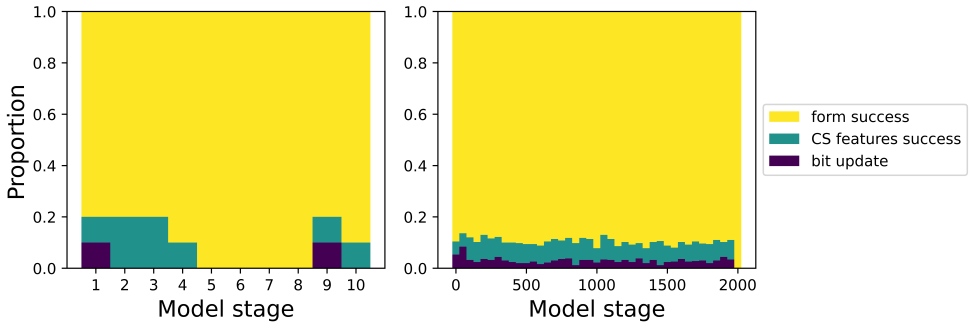


Figure 7. The proportion of language game results (form success, culturally salient features success or bit update) with 10 agents all belonging to the same group (*n_groups* = 1) for the first 10 model stages (**left**) and over the 2000 stages (**right**). At each stage, 10 language games were played. The x-axis starts at stage one because in stage zero there is only model set up and no language game. Across all stages of this run, the majority of the language games end with form success, with a small proportion ending with culturally salient features success (here abbreviated as CS features success). Over 2000 stages, shown on the right, the results were averaged over 50 consecutive model stages (i.e., each bar of the histogram represents the mean of 50 stages).

It is apparent that the vast majority of language games in this run of the model end after form success. In this run of the model, as all agents share the same set of culturally salient features (*n_groups* = 1) and because all agents create their forms to be highly iconic (*initial_degree_of_overlap* = 0.9), the forms of agents will be highly similar at the start of the simulation. The similarity between the agent's forms results in a majority of language games that are ended with form success. Thus, even though the forms stay highly iconic

(they are not changed as there is hardly any bit updating), the agents do not use the iconicity present (language games ending in culturally salient features success) as the language game typically ends with form success. However, throughout the simulation there is still a small proportion (around 10%) of language games ending after culturally salient features success. Few language games end with a bit update.

Second, we present a model run consisting of 10 groups (*n_groups* = 10), meaning that each agent is randomly assigned to 1 of the 10 groups. Because agents are randomly assigned to a group, this does not guarantee that all agents are in a different group. Once assigned to a group, agents are initialized with the set of culturally salient features generated for that group.

Figure 8 shows the results of the language games of 1 model run with 10 agents and 10 groups for the first 10 stages (left) and over 2000 stages (right). For example, in stage 1, 8 language games end with a bit update, 1 ends after culturally salient features success and 1 ends after form success. Over the 2000 stages, it is evident that the proportion of runs ending in a bit update decreases and the proportion of runs ending in form success increases. Over time, form success becomes the most prominent result of the language game, though a considerable amount of language games ending in bit update remains. On the other hand, there are fewer language games ending in *culturally salient features success*; it is clearly the most infrequent result.

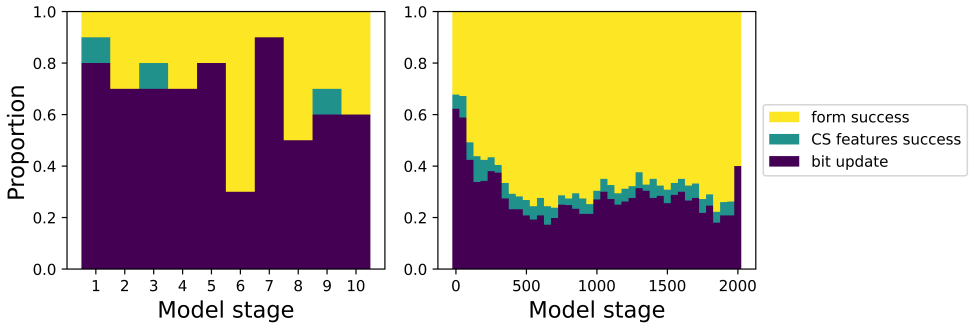


Figure 8. The proportion of language game results (form success, culturally salient features success or bit update) for a model run with 10 agents randomly assigned to 1 of 10 groups (*n_groups* = 10) for the first 10 model stages (**left**) and over the 2000 stages (**right**). At each stage, 10 language games were played. The x-axis starts at stage one because in stage zero there is only model set up and no language game. The majority of the language games at the start of the simulation end with bit update, while later, more end with form success and still a considerable amount end with bit update. Few language games end with culturally salient features success (here abbreviated as CS features success). Over 2000 stages, shown on the right, the results were averaged over 50 consecutive model stages (i.e., each bar of the histogram represents the mean of 50 stages).

In comparing these two example runs, it is evident that the results of the language games with 1 group and 10 groups are different. With 10 groups, bit updates happen much more often than for the run with one group. This is because with one group, if form success is not possible, then culturally salient features success often is as all agents share the same set of culturally salient features. However, with 10 groups, if form success is not possible, agents are likely to end the game with a bit update because it is unlikely that agents share the same culturally salient features, so culturally salient feature success is unlikely to occur. Thus, these two model runs demonstrate how the number of groups (determining the set of culturally salient features of the agents) affect the results of the language games, which in turn affect the degree of lexical variability and iconicity across the population.

3.1.2. Lexical Variability and Iconicity

Figure 9 shows the mean lexical variability and iconicity over the 2000 model stages for the run with 1 group (left) and with 10 groups (right). As previously mentioned, the mean lexical variability is calculated by comparing each bit of each form per pairs of agents (the

distance is 0 if all bits match or 1 if more than 1 bit differs), averaged over all agents at each stage. The mean iconicity is calculated by comparing the degree of overlap between each form and corresponding culturally salient features in an agent's language representation, averaged over all agents at each stage.

First, when all 10 agents belong to 1 group (as can be seen on the left in Figure 9), the degree of iconicity remains constant throughout the run, above 0.9. The mean lexical variability drops slightly and then stabilizes around 0.5. In contrast, when agents are randomly assigned to 1 of 10 groups, the picture is drastically different; as can be seen on the right in Figure 9, both the mean lexical variability and degree of iconicity decrease more than when all agents are assigned to the same group. Initially in this case, the lexical variability across the population is nearly at 1, i.e., the maximum distance possible between the forms of agents. As the forms of agents are initialized on the basis of their culturally salient features, it makes sense that the lexical variability is maximal given that (most) agents are assigned to different groups. From there, the mean lexical variability drops sharply, indicating that there is more lexical similarity across the population over time. The degree of iconicity also drops but stabilizes above 0.5. Given that the degree of iconicity calculation is performed on a bit by bit basis comparing the form to the culturally salient features, 0.5 would represent chance, i.e., an unstructured relationship between the bits of the form and culturally salient features. Though the degree of lexical variability is initially higher when agents are assigned to 1 of 10 groups, the degree of lexical variability decreases much faster and continues to do so, whereas when agents all belong to the same group, the degree of lexical variation (after a short drop in the first 100 stages) remains relatively stable.

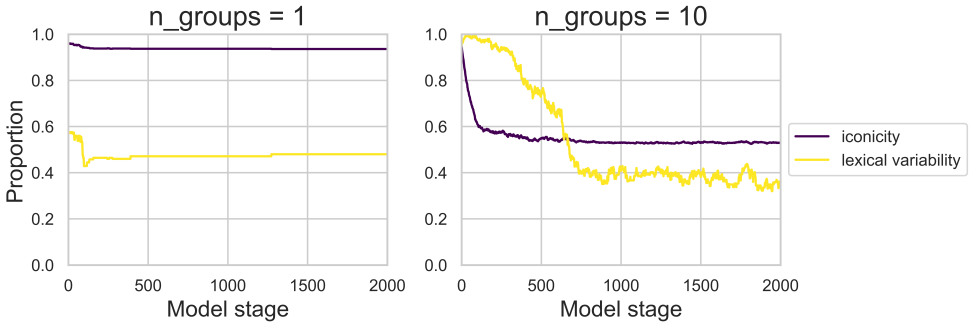


Figure 9. The mean lexical variability and iconicity over the 2000 model stages for one run with 1 group (**left**) and 10 groups (**right**). With all agents belonging to the same group (*n_groups* = 1), the degree of iconicity remains high, and the mean lexical variability across the population remains relatively constant, with more than half of the forms across the population being different. With 10 groups that agents could be assigned to (*n_groups* = 10), the degree of iconicity drops and then stabilizes slightly above 0.5, with 0.5 representing an unstructured relationship between the bits of the form and culturally salient features. The mean lexical variability across the population drops sharply and then continues to drop slowly, indicating that forms become more and more uniform in the population over time.

Now that two examples with just one run have been discussed, we will show the results from 100 repetitions averaged per run with a focus on lexical variation.

3.2. The Effect of Multiple Groups on Lexical Variation

Figure 10 shows different group sizes (*n_groups* = 1, 2, 5 and 10) and the mean degree of lexical variability and iconicity over 100 repetitions. The results from the examples in the previous section are in line with what is shown here; when there is only one group (i.e., all 10 agents have the same set of culturally salient features) at stage 0, there is already some overlap between forms in the population—a lexical variability value of approximately 0.6, indicating that 40% of forms associated with a concept are identical across the population at the start of the run. Over time, the mean lexical variability does not drop below 0.5. The

degree of lexical variability in the population stabilizes more quickly and higher than in the simulations with other groups sizes.

In populations with more groups, the mean lexical variability at the start of the run is high (between 0.8 and 1), as agents belong to different groups and their culturally salient features and hence their forms differ. From this initial point of high lexical variability, there is a sharp decrease in lexical variability. Thus, these populations move quickly towards more uniform form–concept pairings. The number of groups in the population determines at which point the mean lexical variability stabilizes. When there are more groups, the mean lexical variability stabilizes at a lower point. In other words, with more groups, there is more lexical uniformity.

In populations with more groups and hence more culturally salient features, agents cannot rely on shared culturally salient features to communicate. Thus, more often, as shown in the previous section, agents update their forms to be able to successfully communicate with other agents, which results in more uniform form–concept pairings across the population.

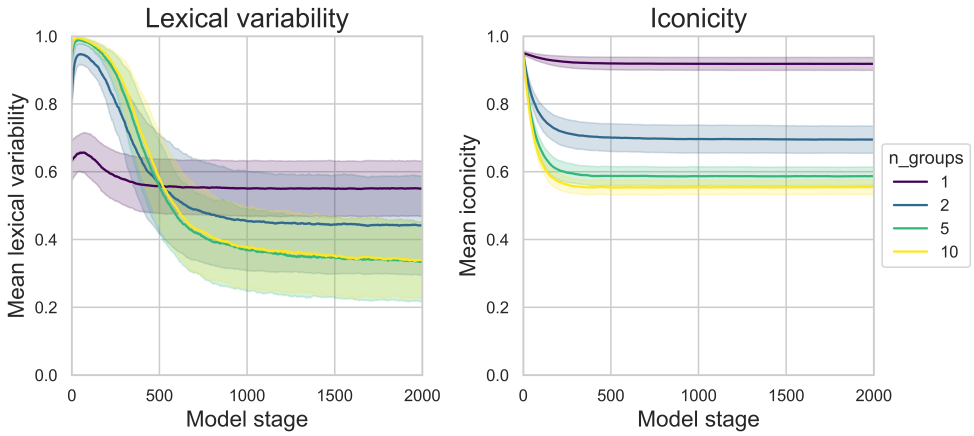


Figure 10. The mean lexical variability over the 2000 model stages for 100 repetitions of a run with 10 agents being assigned to different groups depending on the run. The dark line represents the mean and the shaded area represents the standard deviation of the 100 repetitions. It is evident that there is a relationship between the number of groups and the speed of the decrease of lexical variability, as well as the final amount of lexical variability in the population: The more groups in the population, the higher the initial lexical variability (at stage 0) but the lower the final lexical variability (at stage 2000). In addition, when there are fewer groups, the degree of iconicity is higher.

Additionally, there is a clear relationship between the number of groups and the degree of iconicity: With fewer groups in the population, the degree of iconicity is higher. As predicted, when there are fewer groups, iconic mappings are more useful as more sets of culturally salient features are shared across the population, and therefore the degree of iconicity remains higher. Moreover, as the number of groups increases, the additional difference in lowered iconicity is smaller (e.g., the difference in iconicity between 1 and 2 groups is larger than the difference in iconicity between 5 and 10 groups). In contrast to the lexical variability values, the degree of iconicity quickly stabilizes within the first few hundred stages.

3.3. The Effect of Population Size on Lexical Variation

In this section, we explore the effect of population size on lexical variation for different group sizes. Figure 11 shows different population sizes over time, considering populations consisting of 5, 10, 20, 50 and 100 agents.

In the early stages of the simulation, larger populations exhibit a higher degree of lexical variability than smaller populations. However, over time, larger populations exhibit a steeper decrease in lexical variability compared to smaller populations. In the final stages

of the simulation, the larger population sizes exhibit the lowest degree of lexical variability (i.e., the most lexical uniformity). What can explain this?

In larger populations, there are initially more forms per concept (as forms are generated on an individual level). With agents in a larger population communicating with a larger number of agents, this results in more bit updates. In turn, bit updates typically decrease the degree of iconicity, thereby decreasing the chance of successfully communicating with culturally salient features success. This leads to a feedback loop whereby the frequent bit updates lead to a decrease in the possibility for communicating with culturally salient features success. This process is visualized in Figure 12. On the other hand, in smaller populations, there are initially fewer forms per concept. As agents communicate with a smaller number of agents, less bit updates occur. With fewer bit updates occurring, a higher degree of iconicity is retained, and thus the use of the iconic–inferential pathway (language games ending in culturally salient features success) can be successfully used.

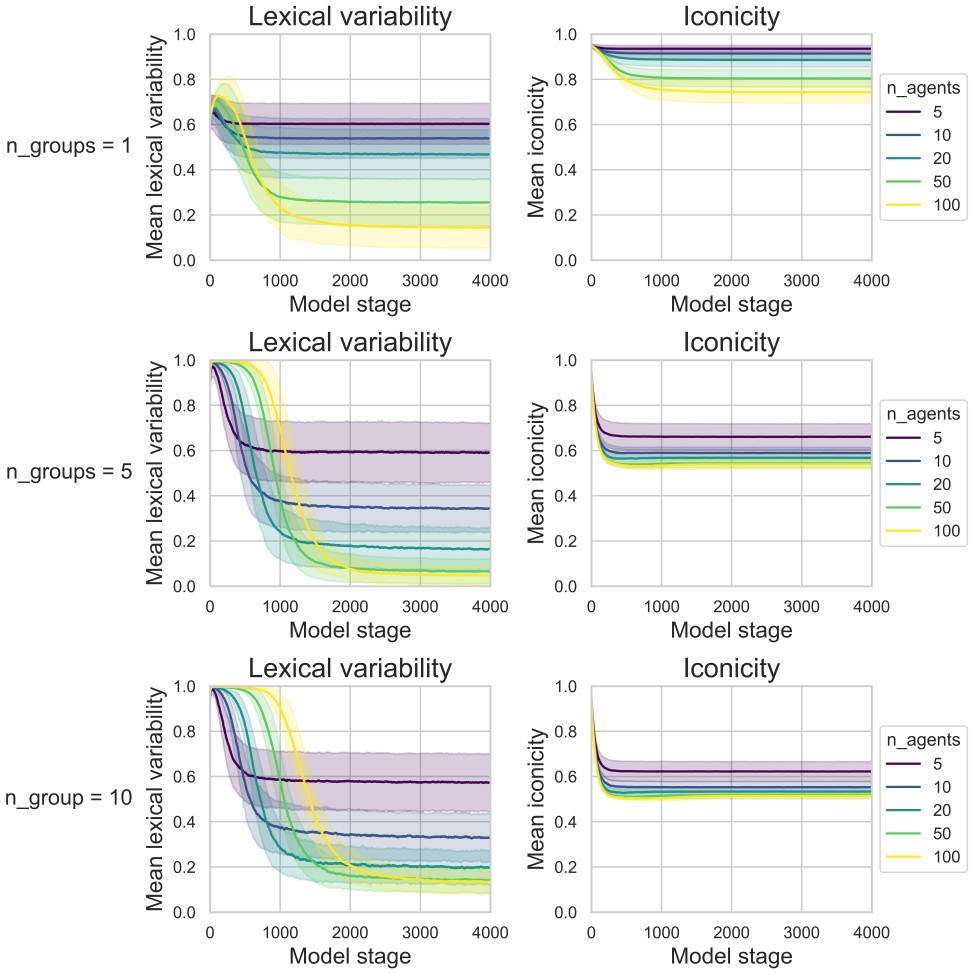


Figure 11. The mean lexical variability over 4000 model stages for different population sizes (*n_agents*), showing three different group values (*n_groups*) determining the sets of culturally salient features of the agents. The dark line represents the mean and the shaded area represents the standard deviation of the 100 repetitions. Regardless of the number of groups, it is clear that the larger population sizes exhibit a lower mean lexical variability than small population sizes. In addition, when there are more agents, the level of iconicity is lower.

Across all population sizes, the more groups, the lower the mean iconicity level is (see Figure 11 iconcity for *n_groups* = 10 vs. *n_groups* = 1), as discussed in the previous section. In addition to this, it is apparent that population size and the number of groups interact in determining iconicity levels. When all agents belong to one group (*n_groups* = 1), there

are larger differences in the mean iconicity level in the population than compared to when agents can be assigned to different groups (*n_groups* = 5 and *n_groups* = 10). The explanation for this relates to the feedback loop mentioned where a lower degree of iconicity stems from more bit updates. When there are more groups, regardless of the population size, agents cannot rely on the iconic–inferential pathway to successfully communicate (language games ending in culturally salient features success) because their sets of culturally salient features differ. With more groups, the feedback loop is present across all population sizes: A lower degree of iconicity stems from more bit updates, here due to the inability of using the iconic–inferential pathway.



Figure 12. The feedback loop from bit updating to the use of culturally salient features success visualized.

4. Discussion

Here, we present a first step in developing a model of how shared cultural context (allowing for the use of iconic mappings) may influence lexical variation in sign language emergence. We have shown that in a model where agents can rely on iconic mappings between a form and culturally salient features in addition to form–concept mappings, populations with a high degree of shared context (operationalized in the model as a smaller number of groups determining the culturally salient features of agents) retain a higher degree of lexical variation. In contrast, populations with many different cultural contexts do not retain the high degree of lexical variation present in language emergence; instead, because these populations cannot rely on iconic mappings between form and culturally salient features, the language becomes more uniform overtime. Overall, these results provide support for the idea that shared context facilitates a high degree of lexical variation (de Vos 2011; Meir et al. 2012).

The main contribution of this model is a novel representation of iconicity, operationalized as a mapping between the bits of the culturally salient features and forms. This has allowed us to consider how iconic properties allow for the retention of lexical variation in culturally homogeneous groups. Crucially, without the iconic–inferential pathway, individuals would need to rely on the conventional link requiring memorizing the association between concepts and forms. Though not tested here, we speculate that a model with only the conventional link would predict a lower degree of lexical variability in communities with more shared context, or at least a comparable degree of lexical variability to communities with less shared context.

In addition to the degree of lexical variability, the model generates predictions about how iconicity is retained in the early stages of language evolution. In populations with a high degree of shared context (i.e., a smaller number of groups), a higher degree of iconicity is exhibited. These populations largely retain the iconicity present in language emergence because agents initially have similar forms (in the model, there was a high degree of initial overlap between forms and culturally salient features), and hence they can typically use the conventional pathway, but if their forms do not match, they can often rely on the iconic–inferential pathway. As agents rarely need to update their forms, a high degree of iconicity is retained. For populations with more diverse backgrounds (i.e., a larger number of groups), the degree of iconicity in the population decreased compared to more

homogenous populations. We are unaware of studies comparing iconicity levels across signing communities with different social structures, but the model generates a prediction which could be empirically tested. It should be noted that in real life, the dynamics of the language game with respect to the two pathways are likely different, as the conventional link has priority over the iconic–inferential pathway in the model. In real life, rather, we assume there is more flexibility with regards to which route is used. We do not expect that the order of the conventional link and the iconic–inferential pathway in the language game has a strong effect on the model results, given that both occur before the form updating step, the step which has ramifications on the degree of lexical variation and iconicity.

We have also explored how population size, in addition to the number of groups, affects lexical variation. We find that larger groups exhibit more lexical uniformity than smaller groups, as found by another computational model in which the lexical variant chosen by the sender depends on their familiarity with the receiver, as agents keep track of individual preferences as well as a group-level preference (Thompson et al. 2020). Interestingly, our model finds the same result without storing information about the frequency of interaction between agents. Instead, the group that agents belong to determines the initial similarity between forms and the ability for agents to rely on the iconic–inferential pathway. All in all, our model provides support for the theories proposing that shared context and population size have an effect on lexical variation in situations of language emergence. Further work must be conducted to determine the precise contribution of each.

The current model is simple—the language model is basic, and there are few model parameters. Simple models permit us to formalize and understand the relationships present in complex systems (Smaldino 2017), such as in the emergence of language. In this way, the relationship between shared context and lexical variability can be studied with minimal confounding factors. However, the model presented here inherently lacks much of the complexity present in signing communities, factors which may have an effect on the degree of lexical variability and iconicity. This model admittedly has several shortcomings, which we discuss and either propose as future model extensions or as general limitations of the model.

One of the biggest shortcomings of this model is that agents only store one form per concept. All sign languages exhibit lexical variation, and while the nature of this variation is still being determined, it is clear that individuals sometimes use multiple forms per concept (i.e., productive synonyms) or understand multiple forms per concept (i.e., perceptual synonyms, see Discussion in Mudd et al. 2020). With regards to productive synonyms, chaining forms has been attested in several shared sign languages, such as in ABSL (Meir et al. 2010), SJQCSL (Hou 2016), in the sign language of Amami Island in Japan (Osugi et al. 1999) and in Kata Kolok (Lutzenberger et al. 2021; Mudd et al. 2020). In addition, compounding is a strategy that has been observed in CTSL (Ergin et al. 2021), SJQCSL (Hou 2016), in the sign language of Amami Island in Japan (Osugi et al. 1999), in ABSL (Meir et al. 2010) and in Kenyan Sign Language (Morgan 2015). For chaining variants together and for compounding, it is necessary that the language representation in the model allows for storing several forms per concept. Hence, the model does not account for productive synonyms. On the other hand, for perceptual synonyms, where an individual can learn a form–concept association even though they might not use it (unless retrieved using the iconic–inferential pathway), it is also necessary to store multiple forms per concept. In the model, perceptual synonyms can be accounted for when agents use the iconic–inferential pathway; the agents have not stored an additional form mapping for a concept, but they may be able to retrieve it. However, in real life, it is much more probable that individuals retain multiple forms associated with a concept even though they have a preference for one form. Thus, in order to account for these different types of synonyms, multiple forms would need to be stored per concept. However, doing so would complicate the dynamics of the language component of the model; it would be necessary to assign weights to each form, as well as assigning a weighing factor for taking iconic affordances into account.

In addition, the update rule in interaction models adapting one's variant in an extremely simplistic, perhaps unrealistic manner. In the case of a bit update (if communication at the form and culturally salient features level has not been successful), the receiver always adapts to the sender. There are many reasons why one individual may adapt their linguistic preferences, such as due to a frequency bias or prestige bias (Boyd and Richerson 1988). Here, agents do not keep track of how many times they have heard a certain variant, nor do agents have varying levels of prestige in the community. The agents simply update if communication has failed. Currently, the language update rule in the model is most akin to explicit feedback from the sender to the receiver. Though explicit feedback is one mechanism used in repair, it is not the only avenue by which individuals come to successfully communicate. Research from the repair sequences in cross-signing, where deaf signers with different native languages meet and communicate, offers an insight into the process of language grounding in its initial stages (Byun et al. 2018). In short, signers anticipate difficulties in communicating and typically produce "try markers" to signal this. The individual producing a try marker essentially asks their communicative partner to produce a grounding sequence, such as an affirmation that their production was understood or a request for clarification. This example highlights that negotiation and repair are complex and nuanced. One way in which the model can be extended is to have more variety in who updates and why exactly, following research from communication in contexts of language emergence and cultural evolution.

Related to this, the update rule dictates that the receiver changes one bit to match the corresponding bit of the sender. In a way, this could be akin to moving phonetically closer to the sender's form. However, this is unrealistic in cases where two forms are very different. Take the example of "sofa" and "couch", both forms referring to the same concept. In the event of communicative failure, it would not make much sense for an individual to adapt only part of the word (e.g., "couch" becomes updated to "souch"). Rather, what would make more sense in this situation is for one individual to learn and potentially use the form sofa from now on. For a more accurate model of human communication, the update rule needs to account for different situations (from learning an entirely new lexical variant to adapting one's existing form phonetically). More research into findings from language acquisition, psycholinguistics and sociolinguistics is necessary in order to adapt this element of the model.

Another unrealistic aspect of this model is that in reality between individuals from different cultures there is likely overlap between the culturally salient features corresponding to concepts, something that is not present in the model as all culturally salient features are generated independently for each group. Returning to the example of *pig*, two individuals from different cultural contexts (e.g., one from a farming community and another from an urban area) are both likely to have salient features comprised of the shape of the pig, the appearance of the animal's face with ears and a snout, the fact that it is an animal, as well as culturally specific points. Though there is undoubtedly overlap in salient features across cultures, for some cultures certain aspects may be more salient than for others. In an urban community with less interaction with pigs, the facial features or the fact that it is food might be more salient, while for a farming community, how it is killed could be more salient. Yet another consideration is how easy it is to represent different facets of culturally salient features. It has been shown in different sign languages that certain semantic categories prompt preferences in production, called patterned iconicity (Padden et al. 2013). For example, across languages signers prefer to use personification (where the culturally salient features are mapped onto the signer's body) for animal signs (Hwang et al. 2017). In the model, as all culturally salient features are generated for specific groups, there is no relationship between the culturally salient features across groups. Given that certain aspects of culturally salient features are typically shared cross-culturally and that patterned iconicity exists, a natural extension of the model would be to model culturally salient features as related, with some degree of overlap between the groups. Better yet, the culturally salient features should even be different for each individual, though more

similar for those in the same group. One final point about the culturally salient features is that in the model only forms can be updated. However, which features are culturally salient in real life become adapted over time, and thus, in the model, this may be important as well. How exactly to model this remains an open question.

Though there are undoubtedly many more ways in which the model can be updated to more closely resemble signing communities and the interaction occurring within them, one final point to address is interaction in the model. In this version of the model, agents all have an equal probability of interacting. This is not the case in real communities—individuals are more likely to interact with some than others. The dynamics in shared signing communities and Deaf community sign language communities with regards to interaction may differ or may be shaped merely because of the size. As shared sign languages are typically small, insular communities, there is more community-wide interaction. On the other hand, in Deaf community sign language communities, which often span entire countries, individuals would typically interact with those in their same city and/or school. This is reflected in the variation observed in these communities; for example, in BSL, a Deaf community sign language, as individuals are more likely to interact with those in their same region, there is substantial regional variation (Stamp et al. 2014). In terms of adding this element of interaction to the model, it would be possible to have agents prefer to select those nearest to them to interact with. This implementation detail may have consequences for the degree and speed of lexical variability and should thus be the subject of future work.

All in all, this research is a first step in developing a model to formalize how shared context affects the degree of lexical variation in sign language emergence. It is unclear to what extent these results may extend to language emergence in our earliest language-using ancestors, who lived in small, insular communities, or esoteric communities (Wray and Grace 2007) and whose communication was likely multi-modal (Levinson and Holler 2014; Perlman 2017). It has been proposed that iconic signs are at the root of proto-language emergence (Számadó and Szathmáry 2012). In addition to the iconic affordances of the manual modality, there is ample evidence that iconicity is also possible and used in the vocal modality and has been shown in spoken languages (Johansson and Zlatev 2013; for a review, see Perniss et al. 2010). As proposed by Meir et al. (2012), it seems plausible that our earliest language-using ancestors residing in small, insular groups had a highly variable lexicon, which may have become more systematic over time. By considering different parameter settings, this model may also provide insights for investigations into what language might have looked like in early human evolution.

5. Conclusions

In this paper, we have presented a model about how lexical variation is affected by societal structure, with an emphasis on shared context. The model validates theories from situations of sign language emergence which find that smaller, insular communities with high degrees of shared context exhibit high degrees of lexical variation, while larger, dispersed communities with more diverse backgrounds (and hence less shared context) exhibit more lexical uniformity. We have shown that shared context, as well as population size, are probable factors influencing lexical variability. Though several additions to the model would yield a more realistic language representation and hence a more valid model, this simple model provides a first step.

Author Contributions: Conceptualization, K.M., B.d.B. and C.d.V.; methodology, K.M., B.d.B. and C.d.V.; software, K.M.; analysis, K.M.; visualization, K.M.; writing—original draft preparation, K.M.; writing—review and editing, K.M, B.d.B. and C.d.V.; funding acquisition, B.d.B. and C.d.V.; All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by FWO-NWO grant number [NWO 326-70-002; FWO G0B4317] "The emergence of phonology within six generations" awarded to Bart de Boer, Paula Fikkert, and Connie de Vos, and by ERC under the ERC Starting Grant [ELISA—852352] "Emergence of Language in Social Interaction" awarded to Connie de Vos, as well as by the Flemish AI plan.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Model code and instructions on how to run the model and produce the plots can be found at: https://doi.org/10.6084/m9.figshare.15163872.v1.

Acknowledgments: We would like to thank Peter Dekker, Marnix Van Soom and Yannick Jadoul for helpful feedback in the initial stages of model development and to Marnix van Soom, Heikki Rasilo, Andrea and Alex Mudd for feedback on and help with producing the plots.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Parameter Exploration

Appendix A.1. Initial Degree of Iconicity

The fixed parameters during this parameter exploration are: *n_concepts* = 10, *n_bits* = 10, *n_agents* = 10, *n_steps* = 4000. The parameters that are varied are *n_groups* and *initial_degree_of_overlap*.

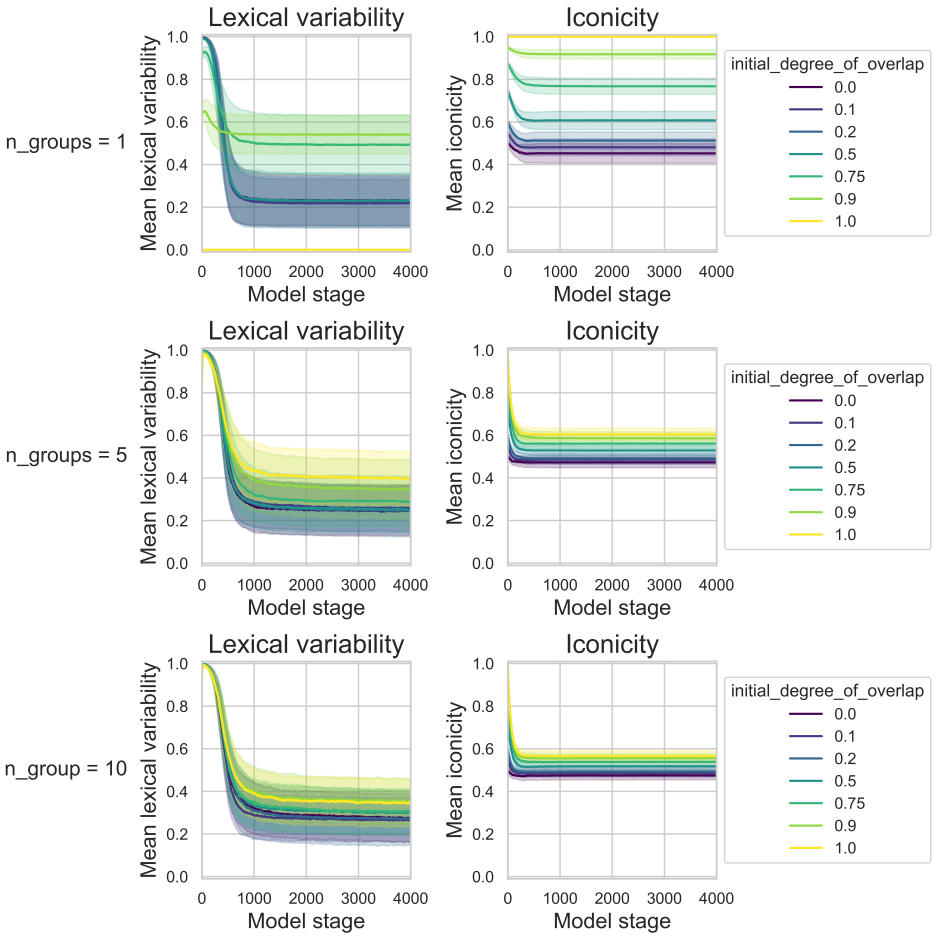


Figure A1. The mean lexical variability over the 2000 model stages for different numbers of groups that agents can be assigned to (*n_groups*): 1, 5 and 10, while varying the degree of overlap between the form and culturally salient features at the start of the run (*initial_degree_of_overlap*). The dark line represents the mean and the shaded area represents the standard deviation of the 100 repetitions. With one exception (*n_groups* = 1 and *initial_degree_of_overlap* = 1), the higher the initial degree of overlap between form and culturally salient features, the lower the mean lexical variability. In addition, the higher the initial overlap between form and culturally salient features, the higher the degree of iconicity.

As shown in Figure A1, when there is only one group where the form and culturally salient features completely overlap, lexical variability is 0 (i.e., all agents have the exact same form for each concept). Except in this case, for all group sizes, smaller overlaps between the form and culturally salient features (*initial_degree_of_iconicity*) result in a lower level of lexical variability. In other words, less lexical similarity initially leads to more uniform productions. With regards to the degree of iconicity, the degree of iconicity is higher in populations with a more initial overlap between form and culturally salient features.

Appendix A.2. The Number of Concepts

The fixed parameters during this parameter exploration are: *n_bits* = 10, *n_agents* = 10, *n_groups* = 5, *initial_degree_of_iconicity* = 0.9, *n_steps* = 2000. The number of concepts is the only parameter that varied.

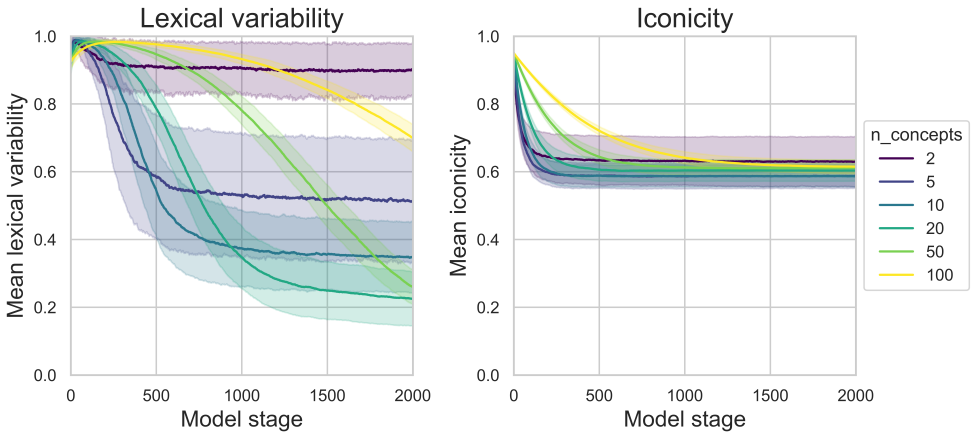


Figure A2. The mean lexical variability over the 2000 model stages for numbers of concepts (*n_concepts*): 1, 5, 10, 20, 50, 100. The dark line represents the mean and the shaded area represents the standard deviation of the 100 repetitions. When there are more concepts, the initial value of lexical variability increases. While the runs with few concepts quickly stabilize at a fairly high degree of lexical variability, the runs with more concepts have a lexical variability value which continues to decrease. The degree of iconicity is comparable across runs with different numbers of concepts, though runs with more concepts retain a higher degree of iconicity longer before stabilizing.

Figure A2 shows the model results of how a different number of concepts affects lexical variation over time. Beginning with runs with a low number of concepts, the lexical variability value quickly stabilizes near the starting lexical variability value. However, in runs with more concepts, the mean lexical variability initially increases before decreasing. The runs with 50 and 100 concepts do not stabilize after 2000 stages. However, it is clear that runs with a larger number of concepts ultimately results in a lower lexical variability value. Why would this be? When there are 2 concepts and 10 bits, because there are only 2 concepts needed to successfully communicate, there is less pressure for the forms to be identical. With few competing concepts, this means that as long as the form from one is different enough from the form of the other, communication will be successful. However, with more concepts, there is more pressure towards uniformity due to the number of competing concepts. In addition, the final degree of iconicity is comparable across runs with different numbers of concepts. However, for runs with more concepts, the degree of iconicity remains higher for longer before reaching a stable point. This is likely because each step of the model only has one language game and hence one chance to update a form. Thus, with more concepts, it takes longer for all forms in the population to update and move away from the initial level of iconicity.

Appendix A.3. The Number of Bits

The fixed parameters during this parameter exploration are: *n_concepts* = 10, *n_agents* = 10, *initial_degree_of_iconicity* = 0.9, *n_steps* = 2000. The parameters that are varied are *n_groups* and *n_bits*.

Figure A3 shows how lexical variability is affected by the number of bits of the forms and culturally salient features. The number of bits affects lexical variability at stage 0, with more bits yielding a higher lexical variability value. Overtime, the more bits there are the more lexical variability is maintained. When there are few bits (*n_bits* = 5), the lexical variability value quickly decreases, before stabilizing above 0. It is probable that these stark differences are the result of using a binary distance measure and the amount of stages that the model was run for: With more bits, more time is needed to make the forms the same. The non-binary distance measure would reveal more similarities across forms. There is, of course, a relationship between the number of bits and the number of concepts in the model; for example, with a large number of bits and few concepts not all bits of a form would need to be identical across the population for the same concept. As long as communication is successful given the pressures imposed by the bits and concepts, there will not be pressure for the population to fully converge on the exact same form for each concept. Thus, for when there are many bits for the form and culturally salient features (when there are few concepts), it would make more sense to have a more nuanced distance measure, taking into account the degree of overlap between forms produced across the population for a given concept.

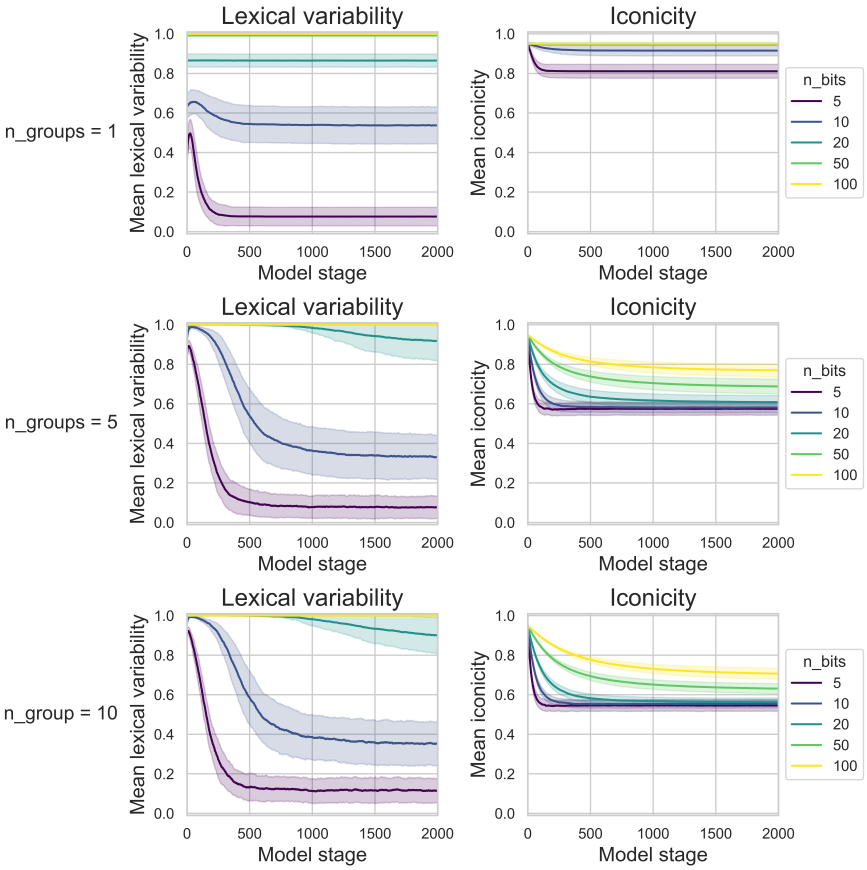


Figure A3. The mean lexical variability over the 2000 model stages for runs with a different number of bits (*n_bits*): 5, 10, 20, 50, 100 and for different numbers of groups that agents can be assigned to (*n_groups*): 1, 5, 10. The dark line represents the mean and the shaded area represents the standard deviation of the 100 repetitions. The lexical variability value is higher when there are more bits. The number of bits also affects the lexical variability value at stage 0. The more bits, the higher the iconicity level.

Finally, with regards to iconicity, the more bits there are, the higher the degree of iconicity. When the language game ends with a bit update, only one bit is updated regardless of the number of bits. Hence, with more bits, it will take longer for the forms to move away from their iconic starting point. In addition, the level of iconicity interacts with the number of groups: With fewer groups, the iconicity remains higher than with more groups. For instance, when all agents belong to the same group, their forms are similar and they are likely to additionally make use of the iconic–inferential pathway. Without the need for bit updating in the case of communicative failure, the level of iconicity in the population remains high.

Notes

[1] We refer to cases with only one form per concept as *uniform*, while some researchers use the term *conventionalized*. Here, we use the term conventionalized for cases when form-concept pairings are generally accepted, so this would apply to cases where there are potentially more than one existing form associated with a concept as long as the form is used and understood. For example, Figure 1 shows three variants for *pig* in Kata Kolok which are conventionalized but not uniform.

[2] We opted not to assign the opposite bit if the initial event fails (and to instead randomly assign 0 or 1) because assigning the opposite bit (0 for 1 and 1 for 0) would still result in a structured relationship. For example, if the culturally salient features are 1111 and the *initial_degree_of_iconicity* = 0, in the current version of the model, this means that the form is comprised of four random bits (ex. 1001), while if the opposite had been assigned, the form would be 0000, where there is still a structured relationship between the culturally salient features and form.

[3] We chose to use a binary distance measure to calculate the lexical variability between agents, as opposed to a continuous distance measure, because lexical variants in the literature about sign languages are often treated categorically. This is true especially in studies of shared sign languages which consider the iconic motivation of signs (e.g., Ergin et al. 2021; Mudd et al. 2020). Unless two forms are identical, in the model, we treat them as different: Two different forms have a lexical variability score of 1, while identical forms have a lexical variability score of 0. However, with larger values of *n_bits*, it may make more sense to use a continuous distance measure for lexical variability. With larger values of *n_bits*, the individual bits may come closer to representing phonetic variation, and hence a continuous distance measure may be more appropriate.

References

Baronchelli, Andrea, Maddalena Felici, Vittorio Loreto, Emanuele Caglioti, and Luc Steels. 2006. Sharp transition towards shared vocabularies in multi-agent systems. *Journal of Statistical Mechanics: Theory and Experiment* 2006: P06014. [CrossRef]

Boyd, Robert, and Peter J. Richerson. 1988. *Culture and the Evolutionary Process*. Chicago: University of Chicago Press.

Byun, Kang-Suk, Connie de Vos, Anastasia Bradford, Ulrike Zeshan, and Stephen C. Levinson. 2018. First Encounters: Repair Sequences in Cross-Signing. *Topics in Cognitive Science* 10: 314–34. [CrossRef] [PubMed]

Clark, Eve. 1987. The principle of contrast: A constraint on language acquisition. In *Mechanisms of Language Acquisition*. Edited by Brian MacWhinney. Mahwah: Lawrence Erlbaum Associates.

de Saussure, Ferdinand. 1916. Nature of the linguistic sign. *Course in General Linguistics* 1: 65–70.

de Vos, Connie. 2011. Kata Kolok Color Terms and the Emergence of Lexical Signs in Rural Signing Communities. *The Senses and Society* 6: 68–76. [CrossRef]

de Vos, Connie. 2012. Sign-Spatiality in Kata Kolok: How a Village Sign Language of Bali Inscribes Its Signing Space. Ph.D. dissertation, Max Planck Institute for Psycholinguistics, Nijmegen, The Netherlands.

Dingemanse, Mark, Damián E. Blasi, Gary Lupyan, Morten H. Christiansen, and Padraic Monaghan. 2015. Arbitrariness, Iconicity, and Systematicity in Language. *Trends in Cognitive Sciences* 19: 603–15. [CrossRef]

Ergin, Rabia, Leyla Kürşat, Ethan Hartzell, and Ray Jackendoff. 2021. Central Taurus Sign Language: On the Edge of Conventionalization. *PsyArXiv*. [CrossRef]

Frishberg, Nancy. 1975. Arbitrariness and Iconicity: Historical Change in American Sign Language. *Language* 51: 696. [CrossRef]

Gasser, Michael. 2004. The Origins of Arbitrariness in Language. Paper presented at the Annual Meeting of the Cognitive Science Society, UC Merced, Chicago, IL, USA, August 4–7; pp. 434–39.

Grimm, Volker, Uta Berger, Donald Lee DeAngelis, Gary Polhill, Jarl Giske, and Steven Floyd Railsback. 2010. The ODD protocol: A review and first update. *Ecological Modelling* 221: 2760–68. [CrossRef]

Grimm, Volker, Uta Berger, Finn Bastiansen, Sigrunn Eliassen, Vincent Ginot, Jarl Giske, John Goss-Custard, Tamara Grand, Simone K. Heinz, Geir Huse, and et al. 2006. A standard protocol for describing individual-based and agent-based models. *Ecological Modelling* 198: 115–26. [CrossRef]

Harnad, Stevan. 1990. The symbol grounding problem. *Physica D* 42: 335–46. [CrossRef]

Hockett, Charles Francis, and Charles D. Hockett. 1960. The Origin of Speech. *Scientific American* 203: 88–97. [CrossRef]

Hou, Lynn Yong-Shi. 2016. "Making Hands": Family Sign Languages in the San Juan Quiahije Community. Ph.D. dissertation, The University of Texas at Austin, Austin, TX, USA.

Hwang, So-One, Nozomi Tomita, Hope Morgan, Rabia Ergin, Deniz Ilkbaşaran, Sharon Seegers, Ryan Lepic, and Carol Padden. 2017. Of the body and the hands: Patterned iconicity for semantic categories. *Language and Cognition* 9: 573–602. [CrossRef]

Johansson, Niklas, and Jordan Zlatev. 2013. Motivations for sound symbolism in spatial deixis: A typological study of 101 languages. *Public Journal of Semiotics* 5: 3–20. [CrossRef]

Kazil, Jackie, David Massad, and Andrew Crooks. 2020. Utilizing Python for Agent-Based Modeling: The Mesa Framework. In *Social, Cultural, and Behavioral Modeling*. Edited by Robert Thomson, Halil Bisgin, Christopher Dancy, Ayaz Hyder and Muhammad Hussain. Cham: Springer International Publishing, pp. 308–17.

LeMaster, Barbara. 2006. Language Contraction, Revitalization, and Irish Women. *Journal of Linguistic Anthropology* 16: 211–28. [CrossRef]

Levinson, Stephen C., and Judith Holler. 2014. The origin of human multi-modal communication. *Philosophical Transactions of the Royal Society B: Biological Sciences* 369: 20130302. [CrossRef]

Lupyan, Gary, and Rick Dale. 2010. Language Structure Is Partly Determined by Social Structure. *PLoS ONE* 5: e8559. [CrossRef]

Lutzenberger, Hannah, Connie de Vos, Onno Crasborn, and Paula Fikkert. 2021. Formal variation in the kata kolok lexicon. *Glossa: A Journal of General Linguistics* 6. [CrossRef]

Marsaja, I. Gede. 2008. *Desa Kolok: A Deaf Village and Its Sign Language in Bali, Indonesia*. Nijmegen: Ishara Press.

McCaskill, Carolyn, Ceil Lucas, Robert Bayley, and Joseph Hill. 2011. *The Hidden Treasure of Black ASL: Its History and Structure*. Washington: Gallaudet University Press.

Meir, Irit, Assaf Israel, Wendy Sandler, Carol A. Padden, and Mark Aronoff. 2012. The influence of community on language structure: Evidence from two young sign languages. *Linguistic Variation* 12: 247–91. [CrossRef]

Meir, Irit, Wendy Sandler, Carol Padden, and Mark Aronoff. 2010. Emerging Sign Languages. In *The Oxford Handbook of Deaf Studies, Language, and Education*. Edited by Marc Marschark and Patricia Elizabeth Spencer. New York: Oxford University Press, vol. 2, pp. 267–80.

Mitchell, Ross E., and Michael Karchmer. 2004. Chasing the mythical ten percent: Parental hearing status of deaf and hard of hearing students in the united states. *Sign Language Studies* 4: 138–63. [CrossRef]

Morgan, Hope E. 2015. When Does a Word Emerge? Lexical and Phonological Variation in a Young Sign Language. *San Diego Linguistic Papers* 2–17. Available online: https://escholarship.org/uc/item/0dg6475 (accessed on 23 January 2022).

Mudd, Katie, Hannah Lutzenberger, Connie de Vos, Paula Fikkert, Onno Crasborn, and Bart de Boer. 2020. The effect of sociolinguistic factors on variation in the Kata Kolok lexicon. *Asia-Pacific Language Variation* 6: 53–88. [CrossRef]

Nyst, Victoria. 2012. Shared sign languages. In *Sign Language: An International Handbook*. Berlin: Mouton de Gruyter, pp. 552–74.

Occhino, Corrine, Benjamin Anible, Erin Wilkinson, and Jill P. Morford. 2017. Iconicity is in the eye of the beholder: How language experience affects perceived iconicity. *Gesture* 16: 100–26. [CrossRef]

Ogden, Charles K., and Ivor A. Richards. 1925. *The Meaning of Meaning: A Study of the Influence of Language upon Thought and of the Science of Symbolism*. San Diego: Harcourt, vol. 29.

Osugi, Yutaka, Ted Supalla, and Rebecca Webb. 1999. The use of word elicitation to identify distinctive gestural systems on Amami Island. *Sign Language & Linguistics* 2: 87–112. [CrossRef]

Padden, Carol A., Irit Meir, So-One Hwang, Ryan Lepic, Sharon Seegers, and Tory Sampson. 2013. Patterned iconicity in sign language lexicons. *Gesture* 13: 287–308. [CrossRef]

Perlman, Marcus. 2017. Debunking two myths against vocal origins of language: Language is iconic and multimodal to the core. *Interaction Studies* 18: 376–401. John Benjamins. [CrossRef]

Perniss, Pamela, Robin L. Thompson, and Gabriella Vigliocco. 2010. Iconicity as a General Property of Language: Evidence from Spoken and Signed Languages. *Front. Psychology* 1. [CrossRef]

Pierce, Charles Sanders. 1931. *Collected Papers*. Cambridge: Harvard University Press.

Pleyer, Michael, Stefan Hartmann, James Winters, and Jordan Zlatev. 2017. Interaction and iconicity in the evolution of language: Introduction to the special issue. *Interaction Studies* 18: 303–13. [CrossRef]

Sandler, Wendy, Mark Aronoff, Irit Meir, and Carol Padden. 2011. The gradual emergence of phonological form in a new language. *Natural Language & Linguistic Theory* 29: 503–43.

Smaldino, Paul E. 2017. Models Are Stupid, and We Need More of Them. *Computational Social Psychology* 311–31. [CrossRef]

Stamp, Rose, Adam Schembri, Jordan Fenlon, Ramas Rentelis, Bencie Woll, and Kearsy Cormier. 2014. Lexical Variation and Change in British Sign Language. *PLoS ONE* 9: e94053. [CrossRef]

Steels, Luc, and Frederic Kaplan. 1999. Situated grounded word semantics. Paper prensented at the IJCAI 99, San Mateo, CA, USA, July 2; Burlington: Morgan Kaufmann, p. 4.

Számadó, Szabolcs, and Eörs Szathmáry. 2012. Evolutionary biological foundations of the origin of language: The co-evolution of language and brain. In *The Oxford Handbook of Language Evolution*. Edited by Kathleen R. Gibson and Maggie Tallerman. Oxford: Oxford University Press. [CrossRef]

Thompson, Bill, Limor Raviv, and Simon Kirby. 2020. Complexity can be maintained in small populations: A model of lexical variability in emerging sign languages. Paper prensented at the the 13th International Conference on the Evolution of Language, Brussels, Belgium, April 14–17. [CrossRef]

Tkachman, Oksana, and Carla L. Hudson Kam. 2020. Measuring lexical and structural conventionalization in young sign languages. *Sign Language & Linguistics* 23: 208–32. [CrossRef]
Trudgill, Peter. 2011. *Sociolinguistic Typology: Social Determinants of Linguistic Complexity*. Oxford: OUP Oxford.
Vogt, Paul. 2002. The physical symbol grounding problem. *Cognitive Systems Research* 3: 429–57. [CrossRef]
Vogt, Paul, and Federico Divina. 2007. Social symbol grounding and language evolution. *Interaction Studies* 8: 31–37. [CrossRef]
Washabaugh, William. 1986. *Five Fingers for Survival*. Ann Arbor: Karoma.
Wray, Alison, and George W. Grace. 2007. The consequences of talking to strangers: Evolutionary corollaries of socio-cultural influences on linguistic form. *Lingua* 117: 543–78. [CrossRef]