

RESEARCH ARTICLE

An iterated learning model of language change that mixes supervised and unsupervised learning

Jack Bunyan¹, Seth Bullock^{1,2}, Conor Houghton^{1,2*}

1 Computer Science, University of Bristol, Bristol, United Kingdom, **2** Intelligent Systems Laboratory, University of Bristol, Bristol, United Kingdom

* conor.houghton@bristol.ac.uk



Abstract

The iterated learning model is an agent model which simulates the transmission of language from generation to generation. It is used to study how the language adapts to pressures imposed by transmission. In each iteration, a language tutor exposes a naïve pupil to a limited training set of utterances, each pairing a random meaning with the signal that conveys it. Then the pupil becomes a tutor for a new naïve pupil in the next iteration. The transmission bottleneck ensures that tutors must generalize beyond the training set that they experienced. Repeated cycles of learning and generalization can result in a language that is expressive, compositional and stable. Previously, the agents in the iterated learning model mapped signals to meanings using an artificial neural network but relied on an unrealistic and computationally expensive process of obversion to map meanings to signals. Here, both maps are neural networks, trained separately through supervised learning and together through unsupervised learning in the form of an autoencoder. This avoids the computational burden entailed in obversion and introduces a mixture of supervised and unsupervised learning as observed during language learning in children. The new model demonstrates a linear relationship between the dimensionality of meaning-signal space and effective bottleneck size and suggests that internal reflection on potential utterances is important in language learning and language evolution.

OPEN ACCESS

Citation: Bunyan J, Bullock S, Houghton C (2024) An iterated learning model of language change that mixes supervised and unsupervised learning. *PLoS Complex Sys* 2(3): e0000030.
<https://doi.org/10.1371/journal.pcsy.0000030>

Editor: Keith Burghardt, University of Southern California, UNITED STATES

Received: May 18, 2024

Accepted: November 24, 2024

Published: March 18, 2025

Peer Review History: PLOS recognizes the benefits of transparency in the peer review process; therefore, we enable the publication of all of the content of peer review and author responses alongside final, published articles. The editorial history of this article is available here: <https://doi.org/10.1371/journal.pcsy.0000030>

Copyright: © 2025 Bunyan et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data availability statement: Code and generated data is available at https://github.com/IteratedLM/2023_12_aim/ and is in the public domain.

Author summary

Languages tend to be expressive: distinct sentences communicate distinct meanings, and compositional: sentence parts consistently correspond to meaning parts. In the iterated learning model a tutor teaches a pupil their language, in this case a mapping from binary vector ‘meanings’ to ‘signals’, by exposing them to a set of meaning-signal pairs. Subsequently, the pupil becomes the tutor for a new naïve pupil. Successive cycles of learning and generalization lead to an expressive, compositional and stable language. However, while an agent’s map from signals to meanings is implemented by its neural network, the

Funding: This work was supported by UKRI (grant EP/Y028392/1 to SB). The funder had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

complementary map from meanings to signals is calculated through ‘obversion’ which becomes impractical for larger languages and is unrealistic, requiring consideration of all possible meaning-signal pairs. Here, we propose a more computationally tractable and more realistic model in which both the encoder and decoder are neural networks, concatenated together as an autoencoder, and pupils are required to learn from a mix of unsupervised and supervised examples, as children do. The model demonstrates the emergence of languages that are expressive and compositional and suggests that internal reflection on potential utterances is important in language learning.

Introduction

In this paper we introduce a new iterated learning model (ILM) for language learning; compared to the existing model, this new model is less computationally demanding and is closer to human language learning.

One of the most important questions in linguistics asks to what degree our linguistic ability prepares us for the specific forms that human languages take [1–3]. It may, however, be possible to sidestep this question by asking two perhaps more pertinent questions: (1) what specific properties do human languages have that make it possible for us to learn and interpret them and to use them to communicate about the diverse and often unforeseen topics we need to discuss and (2) how do languages come to have these properties? Languages are made easier to learn and easier to interpret by their expressivity, compositionality and stability; that is, by the tendency for a language to be able to express many different meanings using different sentences, to compose meaningful sentences from meaningful sub-parts and to change slowly over generations. There are different words for left and right, for example, but the word ‘left’ in ‘you write with your left hand’ means something very similar to the word ‘left’ in ‘turn left at the junction’. Moreover, the difference in meaning between the expressions ‘turn left, then right’ and ‘turn right, then left’ is represented by the different structural organization of their component words. Of course, languages are neither completely expressive nor completely compositional, ‘they are on the left’ typically refers to position but, in some contexts, could refer to politics, and a sailor in some cases would use the word ‘port’ when left might be expected. Nonetheless, expressivity and compositionality are, invariably, important features of languages and might be considered to mark the boundary between a simple system of signs and utterances and a true language [4–6]. The third property, stability, is also important. Since parents need to be able to understand what their children say to them, at least some of the time, it is important that language does not change drastically from generation to generation.

Expressivity, compositionality and stability are only three of a long list of properties a ‘good’ language needs in order that it be easy to learn and effective for communication. However, these three are among the most basic, indeed compositionality is a version of Frege’s principle of compositionality [4,7] which is central to his very influential theory of language. Here, we focus on expressivity, compositionality and stability because of their importance and because they can be measured in the sort of simple language used in our simulations. In ILM-based investigations of how language can acquire important basic properties through iterated learning, these three are intended to stand in for the multiple properties needed for a language to be useful and learnable.

We look at the ILM, [8–11], a model of learning in which a language evolves across generations; this model provides a putative description of how a language might become expressive, compositional and stable. However, as we describe below, the model, requires an ‘obversion’

step which makes it computationally expensive and unrealistic as a simulation of human language learning. Here we propose a solution to this and argue that our new model suggests the importance to language learning and to language evolution of an internal process within the language learner, translating from some meaningful state or idea to an utterance that could be used to express it, and back to how that the meaning of that putative utterance might be interpreted.

The iterated learning model

The ILM models how a language evolves as a tutor teaches their language to a pupil and the pupil, in turn, becomes a tutor and teaches their language to a new pupil. The model demonstrates the gradual emergence of a language that is expressive, compositional and stable, despite the fact that at every iteration the pupil begins in an identical state of ignorance. Here a language is defined as a map from a set of ‘meanings’ to a set of ‘signals’ and, broadly, the goal for the iterated learning model is to shed light on the conditions under which languages that are expressive, compositional and stable arise. We will define measures, x , c and s for expressivity, compositionality and stability; these measures all take values in $[0, 1]$ and we define an XCS language as one where all three exceed $\lambda = 0.95$. In our limited context this will serve as a definition of a ‘good’ language; their precise definitions are given in [S1 Appendix](#).

The ILM was introduced in [8–10]. In its original form, it demonstrated that the emergence of compositionality can be attributed to the influence of a transmission bottleneck on language acquisition taking place over a number of successive generations [9,12]. This original version of the model was used to explain the mixture of irregularity and regularity found in languages [9]. However, this ILM used a language generated by a simple formal grammar and relied on a set of deduction rules that the pupil used to deduce the formal grammar from exemplars. This made it difficult to disentangle the role of the deduction rules from the iterated dynamics of language learning. This was addressed in [11] where a new version of the ILM is presented which did not require these deduction rules.

This ILM has a decoder based on a neural network and incorporated obversion, an algorithm to derive an encoding map from the decoding map, which had been previously defined in a more general context in [13]. This obversion step is problematic in two ways: it is computationally expensive, limiting the ILM to small example languages and it is unrealistic, it makes the transition from pupil to tutor a discrete step during which the pupil must calculate the probability associated with each signal as an encoding of every possible meaning, something very different to what happens in human language learning. Our purpose here is to propose an ILM which does not require obversion.

This is not to suggest that the obverter ILM has not been useful. For example, since it was introduced an obverter-based learning model has been applied to pictures [14]. This is exciting since it simulates the communication aspect of language use and links this class of language-learning model to classic Lewis signalling games [15]. It is hoped that an obversion-free version based on our proposal here would make it computationally possible to investigate this further. Population-based versions of the ILM have been considered [16–18], this is important since language evolution involves space and even geography as well as time, but the obverter ILM has the disadvantage that an agent cannot learn and teach at the same time, something the new model introduced here will address.

The ILM has also been discussed in a broader Bayesian framework, [19,20] and this has been applied to populations [21] and to cultural evolution [22,23]. Similar models of language development include the talking heads experiment, [24], the seeded ILM [25], the neural iterated learning model, [26], and an extension of the restaurant process, [27], to language [28].

These modeling-based approaches also have a counterpart in participant studies using artificial mini-languages, where similar issues of expressivity, compositionality and stability are also considered, see [29–31] for reviews; this approach has been successful in revealing human constraints on language evolution [32,33].

In the ILM, agents can *decode* signals, that is, map them to meanings, and can *encode* meanings, that is map them to signals. A *tutor* agent with its own private encoder and decoder mappings is charged with teaching its language to an initially naïve *pupil* agent, that is, an agent with initial mappings that are empty, random or arbitrary. To do this, the tutor presents a series of meaning-signal pairs to the pupil, where each signal is created by encoding a target meaning using the tutor's encoder. The presentation of each example is an instance of supervised learning, like an adult pointing to a boy stealing biscuits from a cookie jar and saying 'Look! The boy is stealing biscuits from the cookie jar'. The pupil uses these examples to train its own decoder, learning a mapping between signals and meanings. After this training, the pupil 'grows up' and becomes a tutor itself. To act as a tutor it also requires its own encoder, which is derived from its trained decoder by a process of obversion. The set of examples presented to the pupil is only a subset of the full set of possible meaning-signal pairs; this is the transmission bottleneck. The new tutor next trains a new naïve pupil who again initially starts with language mappings that are empty, random or arbitrary, and in this way the process iterates. This cycle of iteration requires the tutor to generalize, in particular, because the transmission bottleneck means that it will, in general, present examples it did not itself learn when it was a pupil, or, because it has not fully learned the exemplars it has been presented by its own tutor.

Despite each pupil starting its training in the same naïve state with no knowledge of the language that the tutor has developed, the ILM demonstrates a tendency for the language that is being learned to change systematically over successive agent 'lifetimes'. Under certain conditions the language stabilizes so that successive agents learn the same language and this language exhibits high levels of expressivity and compositionality.

Whether or not the ILM achieves this kind of outcome depends on the size of its transmission bottleneck: the number of meaning-signal pairs that the tutor presents to the pupil during training. Since this is a subset of the full set of meaning-signal pairs, when the pupil becomes a tutor they must generalize from what they have learned in order to be able to present their pupil with a new set of meaning-signal pairs comprising random meanings, each paired with the signal that encodes it in the tutor's language. When the bottleneck is too small, the language does not stabilize. Where the bottleneck is too large, the result may be a language which is non-expressive and non-compositional, that is, one in which the relationship between each meaning and its associated signal is arbitrary and unstructured.

The neural-network based ILM

Here, the neural-network based ILM presented in [11] will be referred to as the *obverter ILM* to distinguish it from the semi-supervised ILM presented in this paper. In the obverter ILM, and indeed in all the ILMs presented here, the meanings and signals are both binary n -vectors so that $|\mathcal{M}| = |\mathcal{S}| = 2^n$. An individual bit, i , in a meaning vector $m \in \mathcal{M}$ will be called a 'fact'. A bit, j , in a signal vector $s \in \mathcal{S}$ will be called a 'word', in fact, 'morpheme' would be more accurate, since the signal vector is intended to model any aspect of how meaning is encoded, but, for convenience, 'word' will be used. An agent has two maps, an encoding map from meaning to signal:

$$e: \mathcal{M} \rightarrow \mathcal{S} \quad (1)$$

and a decoding map from signal to meaning:

$$d: \mathcal{S} \rightarrow \mathcal{M} \tag{2}$$

A glossary is provided in Table 1 to summarize the notation.

The decoder, $d: \mathcal{S} \rightarrow \mathcal{M}$, is a neural network with input, hidden and output layers each of size n . It is useful to be careful at this point with the notation; the neural network itself maps a meaning, $m \in \mathcal{M}$, to a vector (p_1, p_2, \dots, p_n) where $p_i \in [0, 1]$ is the probability that the i th word in the corresponding signal is one. This map is called \hat{d} :

$$\hat{d}: \mathcal{M} \rightarrow \mathcal{P} \tag{3}$$

where \mathcal{P} is the space of n -vectors whose entries are all probabilities. To add further to the notation, a decision map δ is defined, this maps a vector of probabilities to a binary vector by mapping a probability $p > 0.5$ to one and $p < 0.5$ to zero. In this way $d(m) = \delta(\hat{d}(m))$, or for those who prefer more formal notation: $d = \delta \circ \hat{d}$. As a final note, \hat{d} also defines a map from a meaning and signal pair to a probability:

$$\hat{d}(m, s) = \prod p(m_i) \tag{4}$$

where $\hat{d}(m) = (p_1, p_2, \dots, p_n)$ and $p(m_i) = p_i$ if m_i is one and $1 - p_i$ if m_i is zero. There is a guide to the notation in Fig 1.

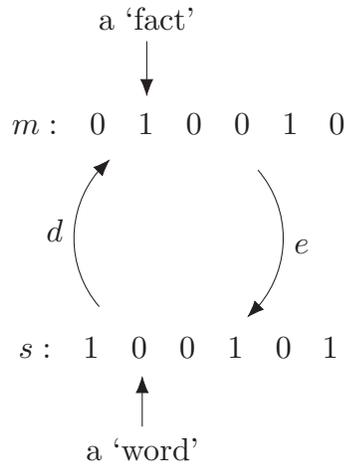
The encoder, e , is more complicated; it takes the form of a list tabulating the preferred signal corresponding to each meaning in \mathcal{M} . This list is calculated by obversion, an inversion-like calculation performed on the decoder map when an agent transitions from pupil to tutor. Obversion is described in more detail in S2 Appendix, but the key point is that it requires the calculation of the 2^{2n} entries in the table of probabilities given by $\hat{d}(m, s)$ for all (m, s) pairs.

Table 1. Glossary. Precise definitions for x , c and s are given in S1 Appendix; all three take values between zero and one, roughly, x measures the degree to which the encoder maps meanings onto distinct signals, c measures the degree to which a given fact is encoded by a particular word and s measures how similar the language in one iteration is to the language in the previous iteration.

n	length of the binary meaning and signal vectors.
\mathcal{M}	set of meanings.
\mathcal{S}	set of signals.
$B \subset \mathcal{M} \times \mathcal{S}$	a set of meaning-signal pairs used for supervised learning.
$\mathcal{A} \subset \mathcal{M}$	a set of meanings used for unsupervised learning.
\mathcal{P}	space of n -vectors (p_1, p_2, \dots, p_n) with $p_i \in [0, 1]$.
meaning	an element of \mathcal{M} (a binary vector of length n)
fact	one bit within a meaning
signal	an element of \mathcal{S} (a binary vector of length n)
word	one bit within a signal
$d: \mathcal{S} \rightarrow \mathcal{M}$	decoder map.
$e: \mathcal{M} \rightarrow \mathcal{S}$	encoder map.
$\hat{d}: \mathcal{S} \rightarrow \mathcal{P}$	decoder neural network.
$\hat{e}: \mathcal{M} \rightarrow \mathcal{P}$	encoder neural network.
δ	$\delta(x) = (x \leq 0.5 ? 0 : 1)$ extended component-wise to vectors.
x	expressivity (blue in graphs).
c	compositionality (orange in graphs).
s	stability (maroon in graphs).

<https://doi.org/10.1371/journal.pcsy.0000030.t001>

A - Words and facts



B - A language

meaning	signal
0 0 0	1 0 1
0 0 1	0 0 1
0 1 0	1 0 0
0 1 1	0 0 0
1 0 0	1 1 1
1 0 1	0 1 1
1 1 0	1 1 0
1 1 1	0 1 0

C - A decoder

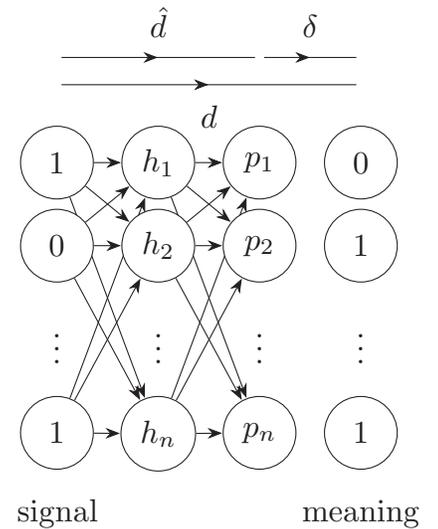


Fig 1. A guide to the notation. A: A meaning is an ordered sequence of n facts, m_1 thru m_n . A signal is an ordered sequence of n words, s_1 thru s_n with $n = 6$ in this example. The use of word is potentially confusing since a word is sometimes thought of as a sequence of letters, but here it is an indivisible component of the signal. A word corresponds to a single bit and a signal can be thought of as corresponding to a phrase; in the same way a meaning can be thought of as corresponding to a state of the world, composed of a set of facts. B: This is an example $n = 3$ language. This is much smaller than the n values actually simulated here but is convenient for illustration. The example here is fully expressive: every possible meaning maps to a different signal, and it is fully compositional: each fact fully determines a unique word. Here, $s_1 = \neg m_3$, $s_2 = m_1$ and $s_3 = \neg m_2$. For illustrative convenience the m_2 and s_3 elements involved in the last of these equivalences have been colored orange. In the example in C the decoder map d maps the signal $(1, 0, \dots, 1)$ to the meaning $(0, 1, \dots, 1)$; the decoder is made up of two parts, the neural network \hat{d} with one hidden layer the same size as the input and output, and the decision map δ which maps probabilities to zero or one.

<https://doi.org/10.1371/journal.pcsy.0000030.g001>

In the obverter ILM the pupil starts as a naïve agent; it is presented with examples of signal-meaning pairs by the tutor agent and uses these to train its decoder map, \hat{d} , using stochastic gradient descent. The set of example meaning-signal pairs, \mathcal{B} , is produced by randomly selecting a subset of meanings and mapping these to signals using e , the tutor's encoder. The set \mathcal{B} is smaller than the set of meanings and so it constitutes a bottleneck in language transmission: in particular the set of examples a tutor presents will generally be different from the examples they learned from and so the tutor is forced to generalize beyond its training corpus. This is an important element of this ILM and the set \mathcal{B} is referred to as the *bottleneck set*. After training, the pupil obverts its decoder map, d , to construct its encoder map, e , and becomes the tutor to a new naïve pupil. This cycle is repeated many times in order to explore how language structure changes over multiple generations. Fig 2 gives a diagram of the generational structure; pseudocode is provided in S3 Appendix.

This model succeeds in simulating the evolution of an XCS language. However, obversion is costly. Already for $n = 8$ the table of probabilities is cumbersome, having 65,536 entries. As n grows it rapidly becomes too large to calculate easily. It is also unrealistic. One of the signature properties of language is its open-ended power to encode novel meanings. It is certainly hard to imagine that when a child is learning language they conjure up all possible subjects of future discussion and estimate the probability that each would be expressed by every possible sentence. The main aim of this paper is to introduce a new ILM which does not require

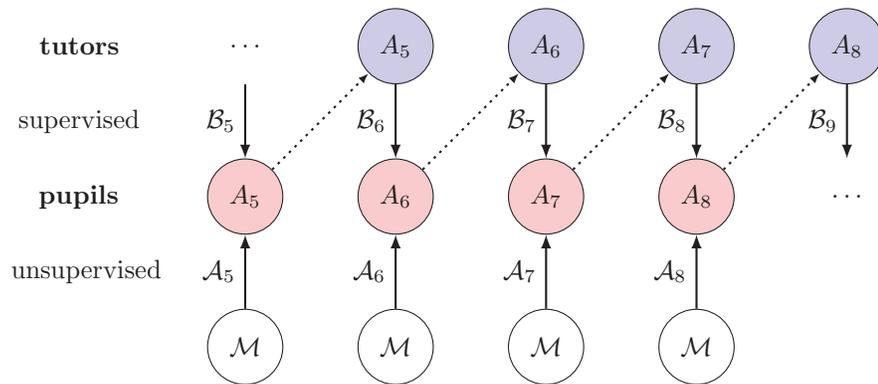


Fig 3. Training the semi-supervised ILM. Each agent, A_i , trains their encoder and decoder during a period combining supervised learning on a set of meaning-signal pairs, B_i , provided by its tutor A_{i-1} and unsupervised autoencoder learning on a set of example meanings, \mathcal{A}_i drawn from \mathcal{M} . Subsequently, the pupil (red) is promoted to become a tutor (blue) to a new pupil agent, A_{i+1} . This is represented by a dashed line because, unlike for the overberter ILM, this promotion is only a change of role, it does not involve any further work for the agent since its encoder is already trained.

<https://doi.org/10.1371/journal.pcsy.0000030.g003>

may make the agent more realistic, from a language learning point-of-view it serves no function and, as far as the language is concerned, this ILM is no different from one with only an encoder. An ILM with only an encoder never becomes expressive; experiments show that the neural network learns to map all meanings to one of the signals represented in the bottleneck set and so the language soon evolves towards one where all meanings are mapped to a single signal. The ILM with only an encoder is discussed further in [S4 Appendix](#).

Motivation from language learning in children To resolve this it is useful to consider the situation with human language. During human language learning, only a portion of a child's language exposure is 'supervised' in the sense that an adult purposefully and explicitly presents a meaning-signal pair. It is clear, on reflection, that the presentation of learning examples to a child is complex, with some 'supervised' child-directed speech including clear indications of meaning, some wholly 'unsupervised' learning based on overheard speech or based on cogitation during observation, and a range of intermediate examples where a child hears utterances related to some activity or situation that they are engaged in, or where some indication of meaning is available or can be inferred, but has not been made explicit or precise.

This is an important and well-studied topic, see [\[34–39\]](#) for example. It is also a complex topic, with child-adult interactions changing as the child becomes older and without any culturally-preserved norms around the structure of a child's language exposure. Crucially, a classic study, [\[40\]](#) replicated in [\[41,42\]](#), demonstrates that during early language acquisition explicit approval or disapproval by a parent of children's utterances depends on whether the utterance is socially appropriate not on whether it is grammatically well formed. Furthermore, in one fascinating quantitative study it is demonstrated that the proportion of speech that is child-directed varies hugely between children in the United States and children in Yucatec Mayan villages [\[43\]](#): despite the difference, children in both environments learn to speak.

Furthermore, the discussion of which exemplars are presented to children during learning typically focuses on the relationship between the child's language exposure and the trajectory of the child's language acquisition, with the suggestion that this is best studied as an example of semi-supervised learning [\[44,45\]](#). Here, however, the hope is to study how languages

develop rather than how children learn, and the mixture of supervised and unsupervised exemplars in learning suggests a solution to the problem of expressivity collapse.

The semi-supervised ILM In the semi-supervised ILM, an agent is considered to have three neural networks: a decoder, d , an encoder, e , and an autoencoder, a , composed of e and d :

$$a = \delta \circ \hat{d} \circ \hat{e} \tag{5}$$

or $a(m) = \delta(\hat{d}(\hat{e}(m)))$, where \hat{e} is defined as the neural network mapping a meaning input to a vector of word probabilities without the final step mapping the probabilities to zeros and ones. This is illustrated in Fig 4. This means that the two neural networks, d and e , are coupled because they are used together to form the autoencoder a .

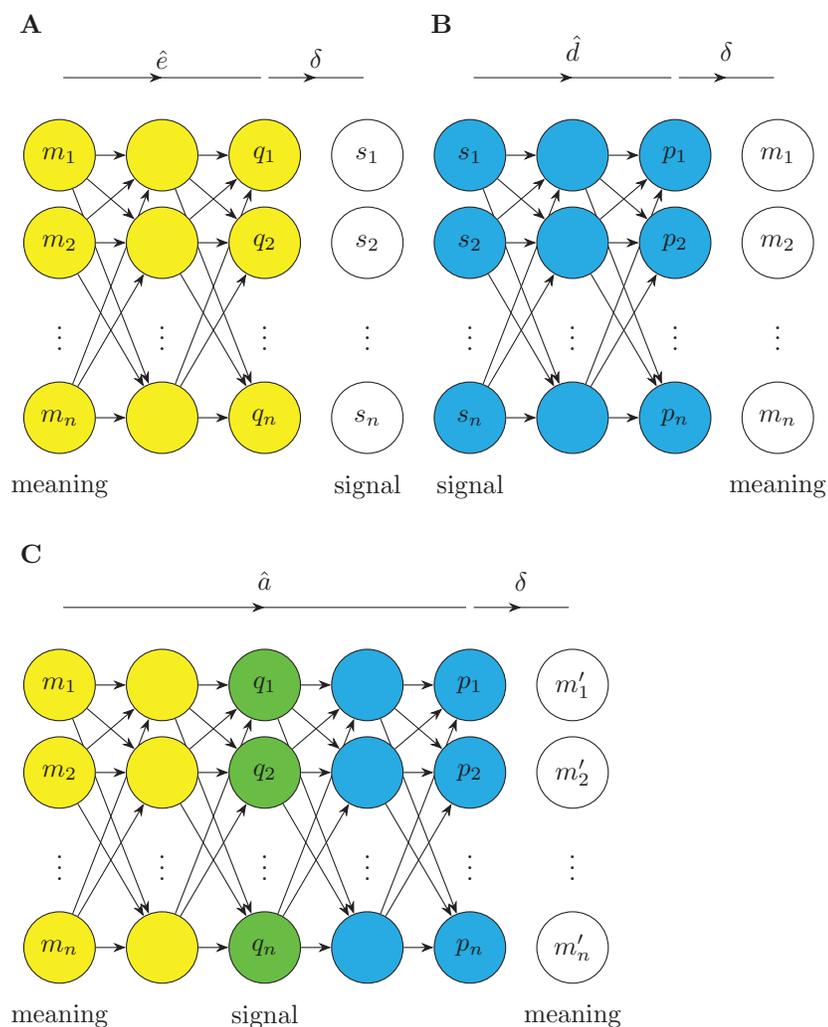


Fig 4. The semi-supervised ILM. A: An encoder, \hat{e} , maps \mathcal{M} to \mathcal{S} . B: A decoder, \hat{d} , maps \mathcal{S} back to \mathcal{M} . C: An autoencoder, \hat{a} , maps \mathcal{M} to \mathcal{M}' , by chaining \hat{e} and \hat{d} . The output of each network is a vector of probabilities (denoted p_i or q_i), which are thresholded by a function δ to deliver a binary vector. In the autoencoder, C, the signal layer is now one of three hidden layers with each node having a value between zero and one.

<https://doi.org/10.1371/journal.pcsy.0000030.g004>

During training, \mathcal{B} is used in three ways: while \mathcal{B} is a set of meaning-signal pairs, it is obviously possible to change the order to get signal-meaning pairs, or to ignore the signal to get a meaning. During each round of training a randomly chosen meaning-signal example from \mathcal{B} is present to e , a randomly chosen signal-meaning is presented to d and $r = 20$ meanings are presented to a . This final training step aims to improve the ability of the autoencoder to associate the meanings with itself and trains the autoencoder to map meanings to signals and back to meanings by measuring how close the input and output are to each other. This is intended to mimic a child observing the world and trying to internally match its observations to self-generated utterances and map these utterances back to what it is observing.

Simulation language evolution in the semi-supervised ILM Fig 5 demonstrates the semi-supervised ILM's ability to produce an XCS language. The performance of the obverter ILM is reprised in Fig 5A–5C. As in [11] it is demonstrated that an XCS language quickly arises within the ILM: Fig 5A–5C show that after 40 generations all three properties are close to their maximum. Fig 5D–5F and, again, in Fig 5G–5I show that the semi-supervised ILM also succeeds in evolving an XCS language. In Fig 5D–5F the meanings used to train the autoencoder are all meanings from \mathcal{B} . This may seem unrealistic: a child's observations are not restricted to the states of the world that it has heard described. In the case of $n = 8$ this does not appear to have a crucial affect on language evolution: using an independently randomly selected and larger set of meanings for autoencoder training produces similar results Fig 5G–5I. However, this no longer holds for larger languages. Using a larger set, \mathcal{A} , of meanings selected at random and independently of \mathcal{B} is required for an XCS language to evolve; this is demonstrated in Fig 6. In Fig 6A–6D the ILM is evolved with $\mathcal{A} = \mathcal{B}$ and $|\mathcal{A}| = |\mathcal{B}| = 160$ and does not reach an XCS language, in Fig 6E–6G, when $|\mathcal{A}| = 480$ and \mathcal{A} and \mathcal{B} are selected independently, it does. The value 480 is picked for illustrative purposes, in fact, lower values work nearly as well; a plot showing how performance changes as the size of \mathcal{A} changes is give in S1 Fig.

In Fig 7 a still larger language is consider: $n = 20$ where the set of meanings and of signals each have over a million elements. This means that the set of expressive and compositional encoders is tiny compared to the vast number of possible maps. Since obversion would require calculating a table with $2^{40} \approx 10^{12}$ entries this is too large to easily test the obverter ILM and so it is unknown whether the obverter ILM produces an XCS language in this case. As seen in Fig 7A–7C, for the semi-supervised ILM the languages become expressive and compositional but the mean stability is just a bit above $s = 0.6$. This might be considered realistic since stability is a harsh measure for compositional languages, swapping just two words gives a stability of $s = 0.5$. However, if this instability is a problem, it is one that can be solved, for example by increasing the number of nodes in the hidden layer of the decoder and encoder networks: Fig 7D–7F shows the behaviour when the hidden layer has 30 nodes while the signal and meanings still have length 20. With this larger hidden layer the language rapidly evolves into an XCS language.

The size of the bottleneck set Fig 8 is perhaps surprising. This set of simulations examines the size of the bottleneck as n , the bit-length of the meaning and signal spaces, is changed. For this a precise, though of course slightly arbitrary, definition of ideal size is used: the model is run until x , c and s all exceed a $\lambda = 0.95$ threshold, indicating that an XCS language has developed. The number of generations is noted, this is repeated 25 times and the mean calculated. The size of the bottleneck that enables the threshold to be reached most quickly, on average, is taken as the best bottleneck size. It turns out that this value increases linearly with n . This is despite the size of the meaning and signal sets increasing exponentially with n ; for $n = 8$ the optimal bottleneck is close to a quarter of the number of meanings, while for $n = 12$ it is only about a 40th. In fact, this is not unique to the semi-supervised ILM:

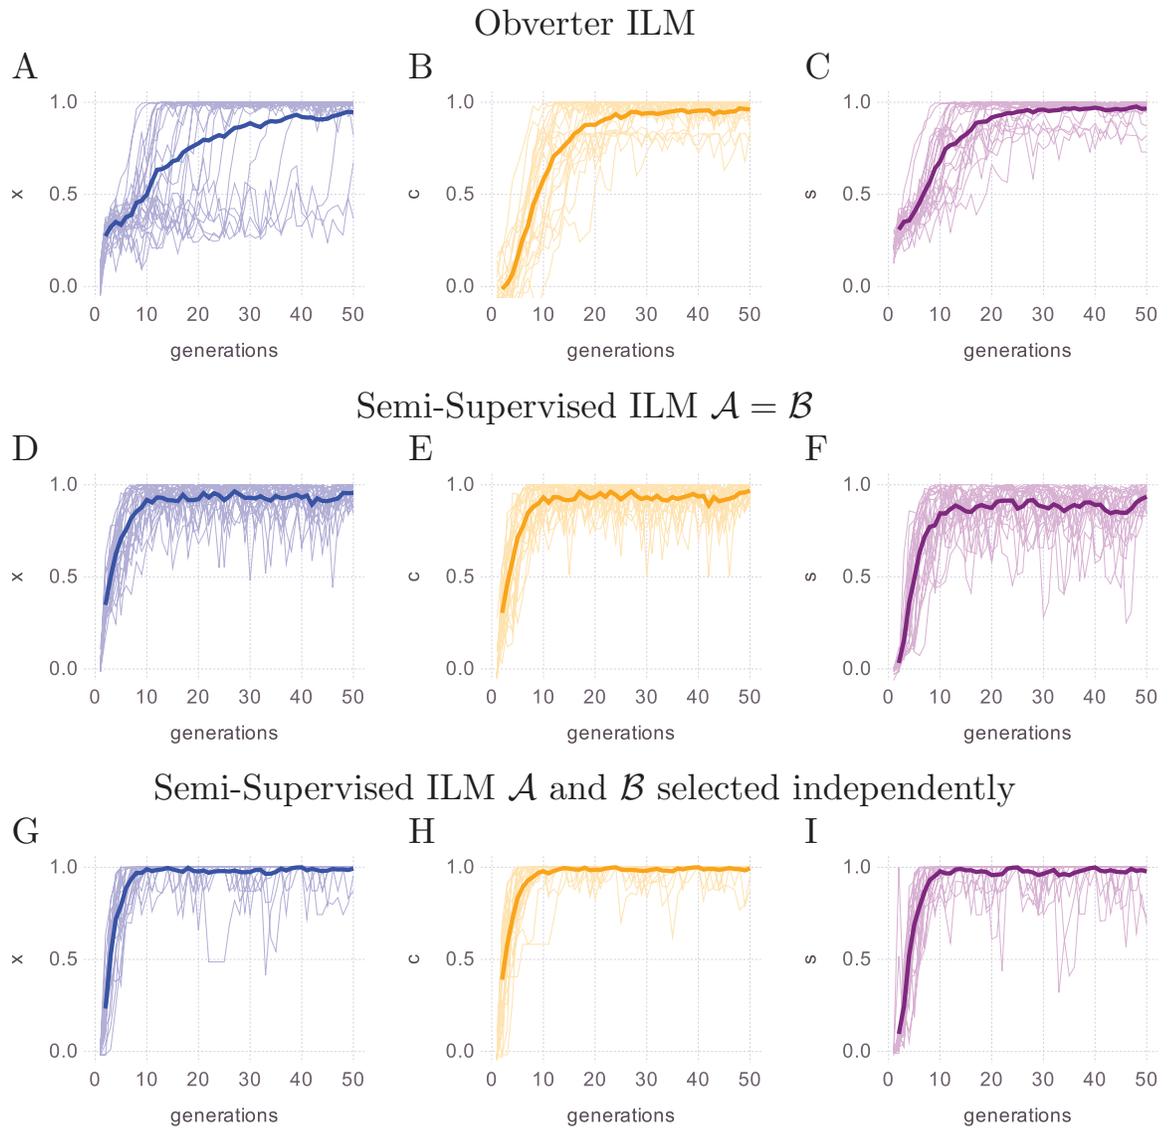


Fig 5. The semi-supervised ILM evolves a stable, expressive, compositional language. This figure describes the performance of the semi-supervised ILM but, for comparison, A–C shows the performance of the obverter ILM. Expressivity (x), compositionality (c) and stability (s) as a function of generation for 25 independent replicates of the $n = 8$ obverter ILM; the thick line is the mean, the thin lines show the individual replicates and the bottleneck size in all cases is 50. D–F plot respectively expressivity (x), compositionality (c) and stability (s) as a function of generation for 25 independent replicates of the $n = 8$ semi-supervised ILM; the thick line is the mean, the thin lines show the individual replicates and the bottleneck size in all cases is 75; the autoencoder is trained using the same meanings as appear in the meaning-signal pairs used to train the decoder and encoder. In G–I \mathcal{A} is selected independently to \mathcal{B} and $|\mathcal{A}| = 225$.

<https://doi.org/10.1371/journal.pcsy.0000030.g005>

although it does not seem to have been previously noted, the obverter ILM demonstrates the same linear dependence on bottleneck size that was observed for the semi-supervised ILM, see S2 Fig.

In both cases the measure used to quantify this effect is somewhat arbitrary and noticeably noisy. Furthermore, as will be discussed in much more detail below, the performance of the semi-supervised ILM does not fall that much as $|\mathcal{B}|$ increases, it produces an XCS

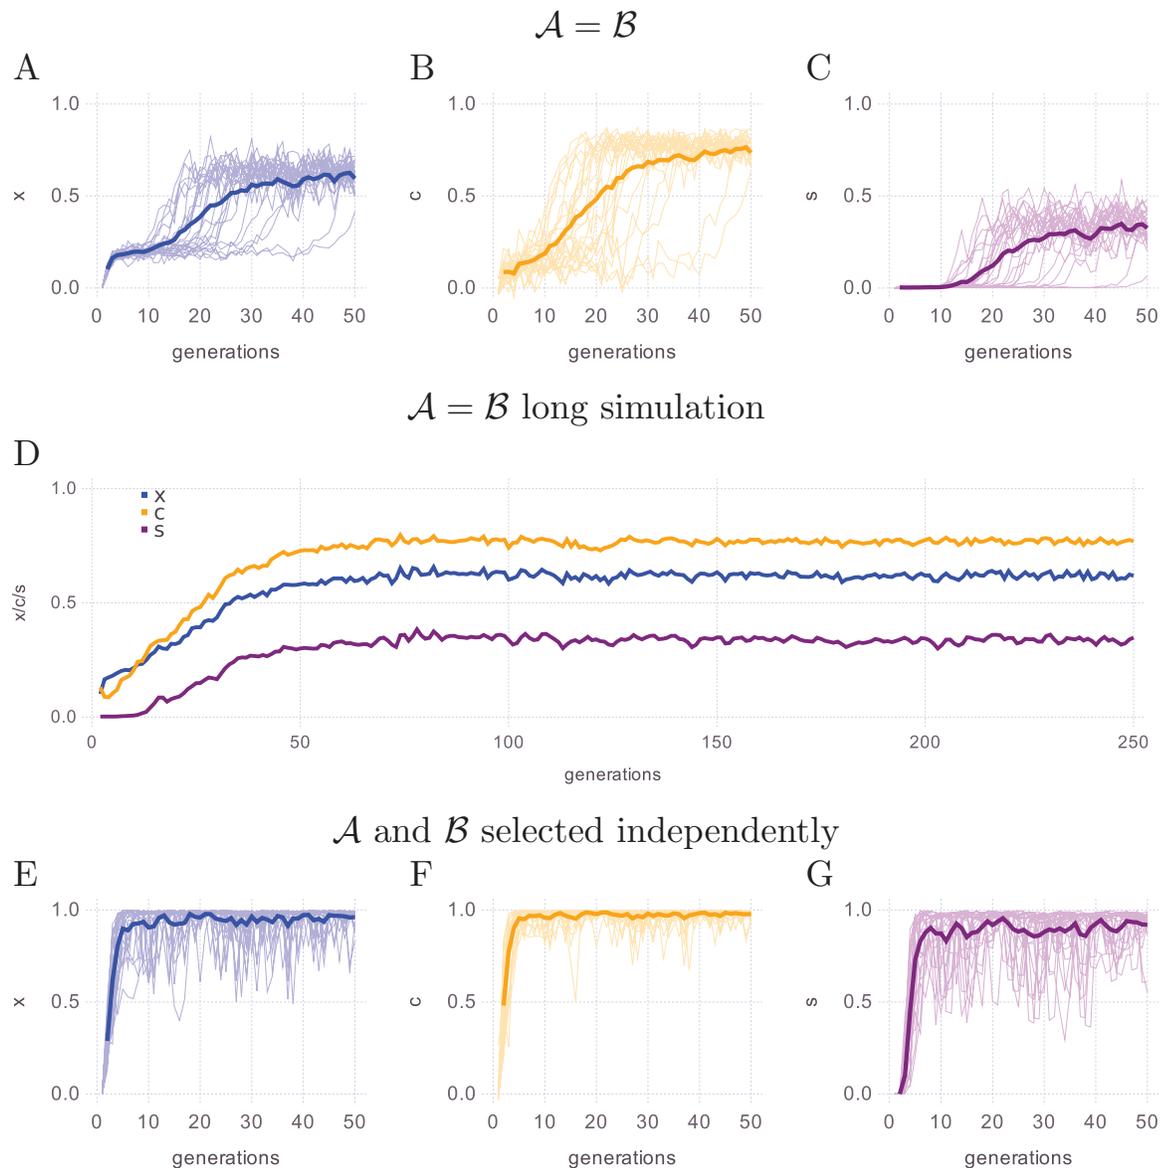


Fig 6. For larger languages \mathcal{A} needs to be different from, and larger than, \mathcal{B} . Here the $n = 16$ semi-supervised ILM is simulated with $|\mathcal{B}| = 160$. In A–C the model is run with $\mathcal{A} = \mathcal{B}$; the language does become somewhat expressive, compositional and stable, but it certainly does not evolve an XCS language. Here results for 25 replicates are shown. It may appear from these simulations that the language would continue to become more expressive, compositional and stable, but, an additional set of simulations run for 250 generations shows, D that this is not the case; here only averages are plotted, but, as before, these are based on 25 replicates. In E–G \mathcal{A} is chosen independently of \mathcal{B} and has size $|\mathcal{A}| = 480$. Here, an XCS language quickly evolves.

<https://doi.org/10.1371/journal.pcsy.0000030.g006>

language even when $|\mathcal{B}| = 2^n$, that is when it is not actually a bottleneck. However, the conclusion is clear, the number of exemplars needed for the ILM to work increases with n , not 2^n : evolving an XCS language relies on matching facts to words, not meanings to signals. In this way, language in the ILM evolves towards one that can be learned for a smaller example set than would be required for non-compositional language, giving a simple example of how the ‘poverty of stimulus’ identified in human language learning, [46,47], could be suggested to

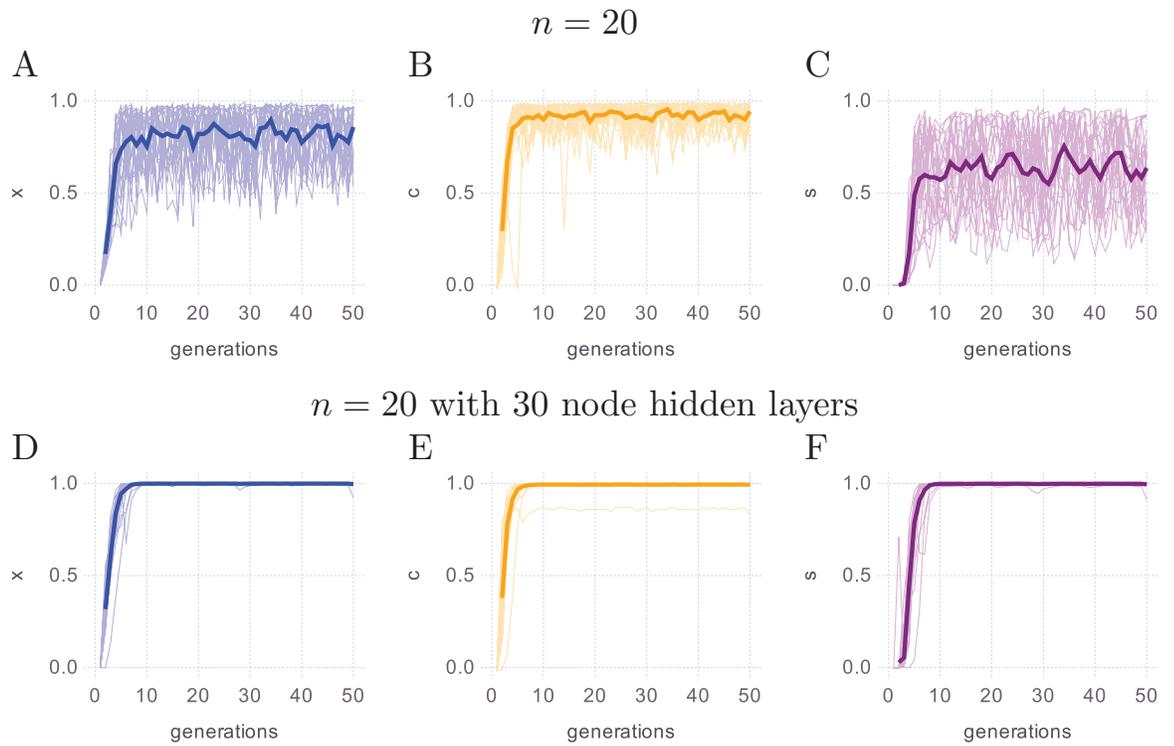


Fig 7. Even larger languages do not become stable unless the hidden layer is made larger. Here $n = 20$, $|\mathcal{B}| = 200$ and \mathcal{A} is chosen independently with $|\mathcal{A}| = 600$. The usual semi-supervised ILM is used in A–C, with as usual 25 replicates. The language evolves some characteristics of an XCS language but does not become stable. In D–F the size of the hidden layer is increased to 30 and the language evolves extremely quickly.

<https://doi.org/10.1371/journal.pcsy.0000030.g007>

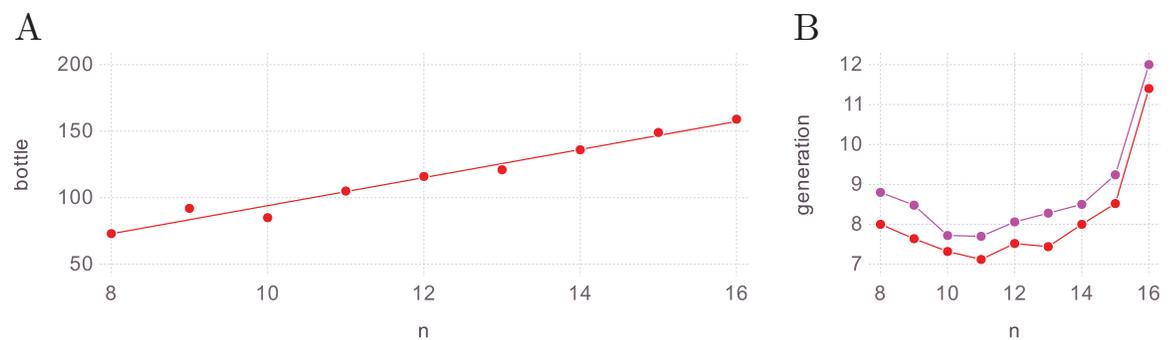


Fig 8. The required bottleneck size increases linearly for the semi-supervised ILM. Here the model is run until x , c and s exceed $\lambda_g = 0.95$ for the first time. This was done for 25 independent replicates. By testing the full relevant range of bottleneck sizes the best value of the bottleneck was estimated and this is plotted in **A**; the line shows the standard best linear fit, this has slope 10.5 and intercept -11.6. In every case the set of signals used to train the autoencoder is three times the size of the set of signal-meaning pairs used to train the encoder and decoder. The average number of generations to reach $\lambda_g = 0.95$ using the optimal value bottleneck size is plotted, in red, as a function of n in **B**. As an indication of how quickly the average number of generation changes as the bottleneck size is changed, the magenta line shows the average number of generations if a value one away from the optimal bottleneck size is used; averaged between adding one and taking one away.

<https://doi.org/10.1371/journal.pcsy.0000030.g008>

impose a structure on language, rather than require a particular instinctive awareness in the human brain.

The rate of learning An expressive language supports a diversity of signals and meanings and as a consequence may be challenging to learn under conditions of limited time and cognitive effort. A compositional language, however, is learnable, even if it is expressive, because a pupil need only learn the rules of composition, rather than a much larger mapping from states of the world to arbitrary, unstructured phrases.

As a simple demonstration of this, we can consider how well the neural networks inside the semi-supervised ILM's agents learn as the initially arbitrary language matures into an XCS language. Of course, the learning implemented here, using stochastic gradient descent, is very different from anything that plausibly happens in the brain, so it would be wrong to suggest that the semi-supervised ILM's learning dynamics align with the learning dynamics of a human language learner. However, the ease with which the neural networks reduce their loss function is likely related in some general way to how easy the language is to learn. In Fig 9 the loss achieved by each of the pupil's neural networks is plotted against training epoch for each generation. Since each pupil always starts in a naïve state, the initial loss is always the same at epoch one, but for higher generations, and hence more mature languages, the loss falls more quickly.

More broadly though, the issue of how well a language supports communication goes well beyond our approach here. If languages are subject to selective pressures to minimize ambiguity or reduce cognitive load on language users, these pressures might be seen to shape language change within an iterated learning model. Given its lower computational demands, the semi-supervised ILM could be extended to incorporate such additional pressures on language, for example related to the cognitive load experienced by a speaker while forming sentences or the ease with which these sentences can be processed or unambiguously understood by a listener.

The size of \mathcal{A} and \mathcal{B} Provided the size of the bottleneck is large enough, the performance of the semi-supervised ILM does not appear to depend in a sensitive way on either $|\mathcal{A}|$ or $|\mathcal{B}|$. This can be seen in Fig 10 which plots the values of x , c and s for both five and 20 generations while $|\mathcal{A}|$ and $|\mathcal{B}|$ are varied. It is striking that increasing the size of the bottleneck set \mathcal{B} does

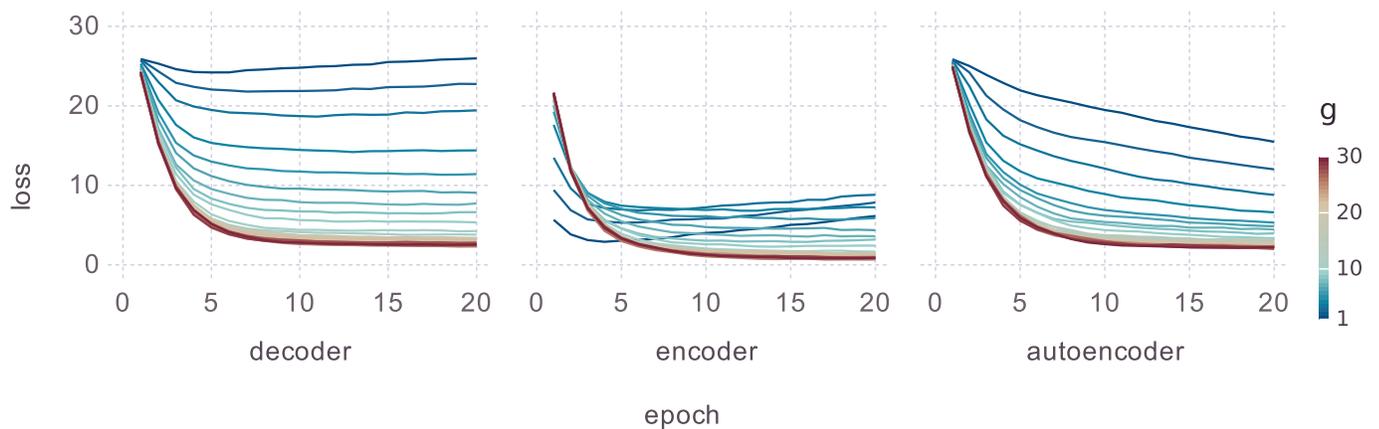


Fig 9. XCS languages are easier to learn. The loss function is plotted against epoch during learning for 30 generations of the semi-supervised ILM; the colour ranges from cold to warm as the generation number, g , increases. This is done separately for each of the three neural networks, the decoder, the encoder and the autoencoder. In each case show is an average over 25 independent replicates. The autoencoder loss has been divided by $r = 8$ to make it comparable to the other two loss values. Here $n = 8$, $|\mathcal{B}| = 50$ and $|\mathcal{A}| = 3|\mathcal{B}| = 150$.

<https://doi.org/10.1371/journal.pcsy.0000030.g009>

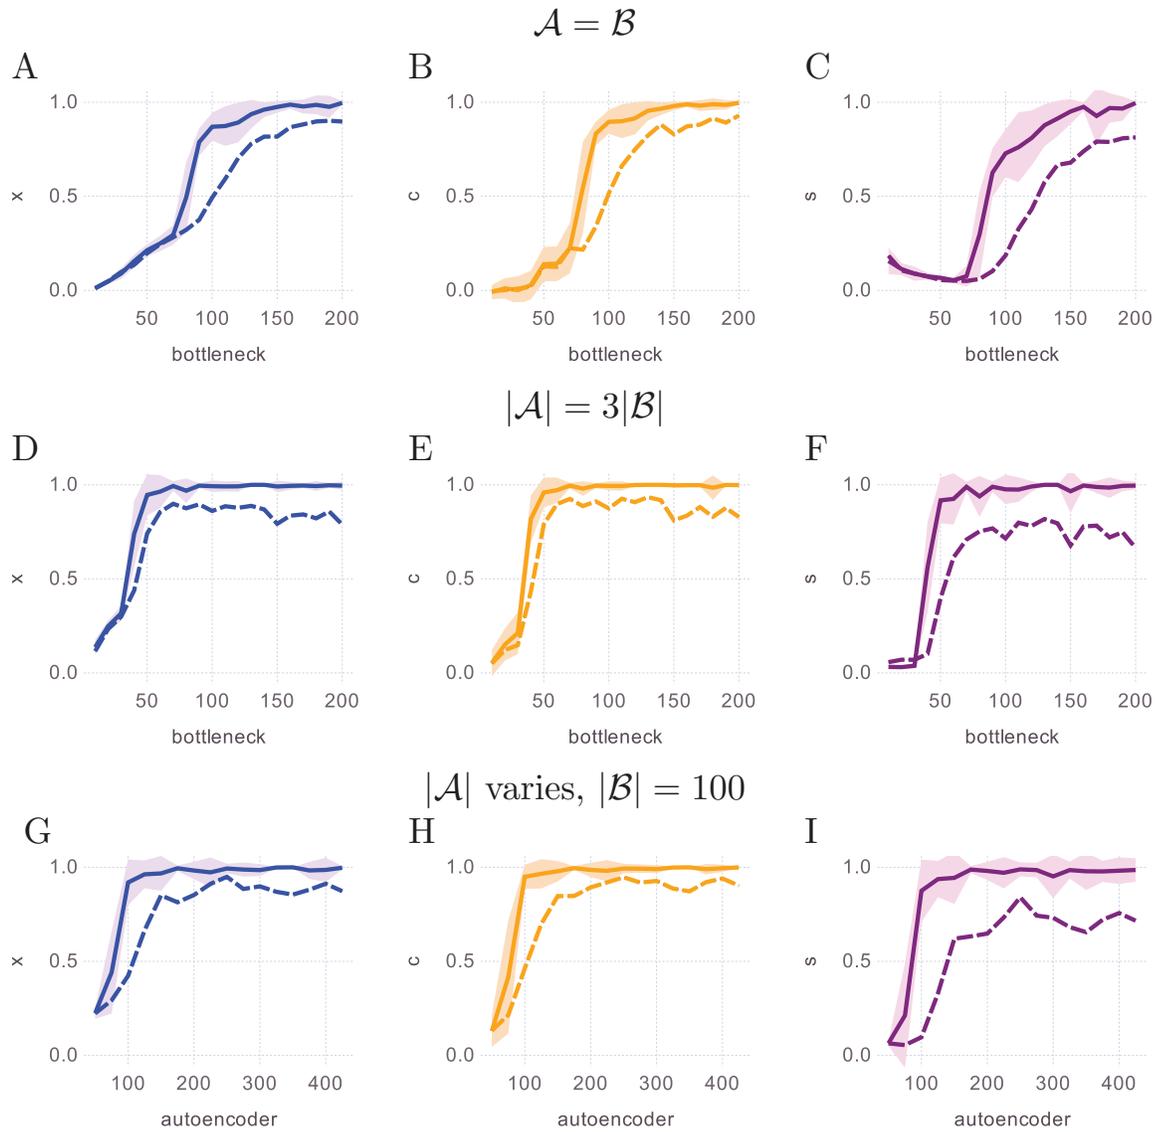


Fig 10. Exploring optimal training set size. A–C: $\mathcal{A} = \mathcal{B}$; D–F: $|\mathcal{A}| = 3|\mathcal{B}|$; G–I: $|\mathcal{B}|$ is fixed at 100. In all cases \mathcal{A} and \mathcal{B} are independent random samples drawn from the space of all possible signals with $n = 10$ ensuring that the number of possible meaning and signals, 2^{10} , far exceeds the largest training sets considered. Mean values for x , c and s are plotted after 20 generations (solid line) and five generations (dashed line) for 25 independent replicates.

<https://doi.org/10.1371/journal.pcsy.0000030.g010>

not substantially reduce performance. This is not true of the oververter ILM, Fig 11 shows that performance is reduced if the bottleneck is too large.

Since the bottleneck is considered essential for the ILM to evolve an XCS language, it is surprising that the performance of the semi-supervised ILM does not appear to fall substantially as the size of the bottleneck increase. To examine this the performance of the ILM with $|\mathcal{B}| = n$, that is, with no bottleneck in the transmission from tutor to pupil, is examined in Fig 12. In Fig 12A–12C with $n = 8$ the semi-supervised ILM develops an XCS language even when the bottleneck size matches the number of meanings and signals and, although the

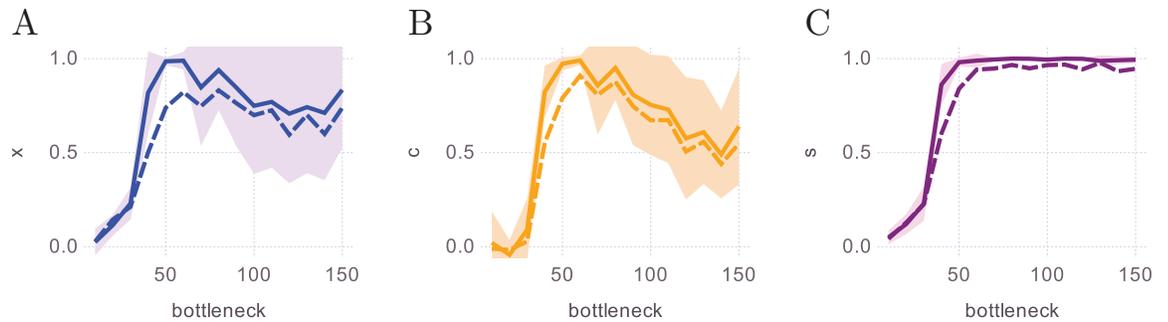


Fig 11. The obverter ILM is sensitive to bottleneck size. A–C plot mean values for x , c and s as a function of bottleneck size after 40 generations (solid line) and 15 generations (dashed line) for 25 independent replicates per data point. The ribbons depict standard deviations around the 40-generation mean.

<https://doi.org/10.1371/journal.pcsy.0000030.g011>

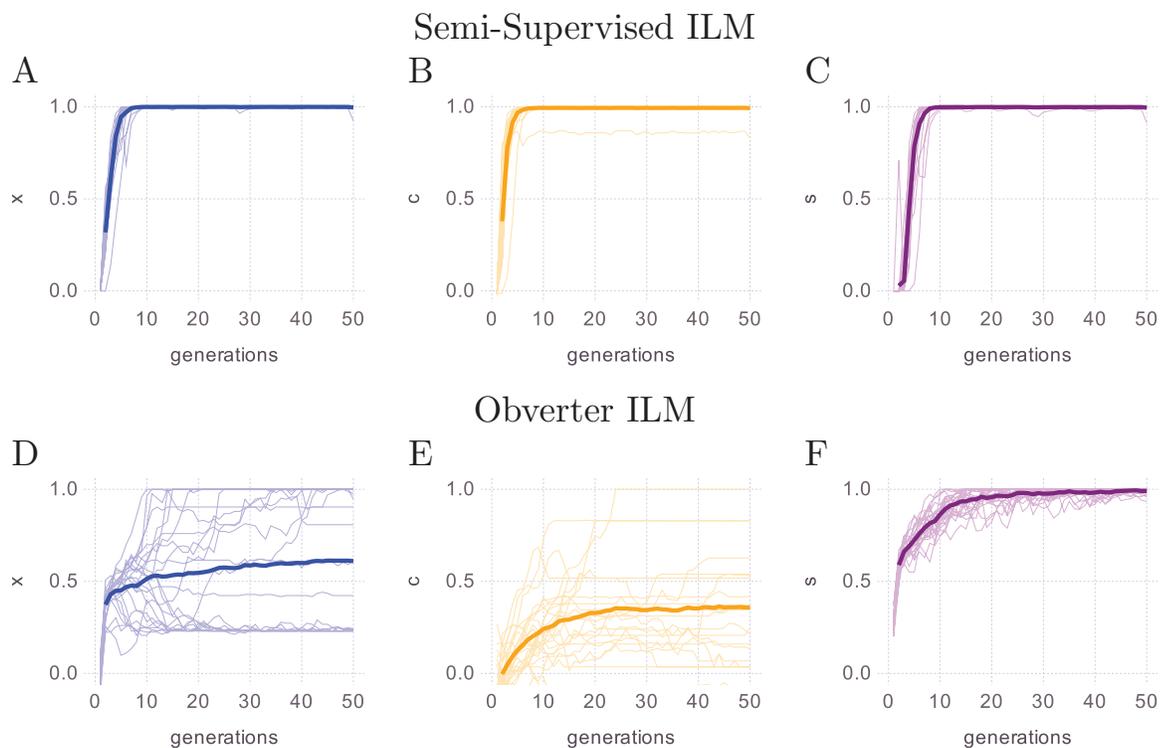


Fig 12. Performance of the ILM when there is no bottleneck. Here the $n = 8$ ILM is run with the bottleneck size set to 256, meaning that all meaning-signal pairs are taught by the tutor to the pupil. A–C show results for the semi-supervised ILM, D–F for the obverter ILM.

<https://doi.org/10.1371/journal.pcsy.0000030.g012>

obverter ILM does not reliably evolve an XCS language, it performs better than might have been anticipated, Fig 12D–12F.

This means that the cycle of generalization from a bottleneck set is not the only aspect of the ILM dynamics that leads to an XCS language. It is notable that for the semi-supervised ILM the stability is not near one for the first few generations, indicating that the pupil has not fully learned the language for the tutor, even though it has been presented a complete set of

examples. This suggests that it is difficult to train the encoder and decoder while also training the autoencoder for languages with low values of x and c . This is not surprising, if the language has poor expressivity it will not be able to map all meanings back to themselves and so while the language remains unexpressive the goal of training the autoencoder is in opposition to the goal of training the decoder and encoder. It appears both that learning in the autoencoder forces a sort of generalization towards a more compositional language and that because the pupil, in early generations, has not fully learned the language of the tutor it is forced to generalize even when there is no bottleneck.

Discussion

One draw back of the ILM is that it involves a large number of meta-parameters and we have only explored a subset of these in detail. We have not presented any description of how performance depends on parameters like learning rate, the number of epochs, the number of presentations to the autoencoder in each epoch, or on any broad consideration of how details of the neural network architecture or training algorithm affect the results. While we have done some preparatory exploration to find values that work, it would be useful to examine these details more formally and more rigorously in future.

The learnability of languages The novel semi-supervised ILM introduced here demonstrates that XCS languages, languages that are expressive, compositional, and stable, arise without the use of an obverter. While natural languages exhibit these properties to a substantial degree, they also represent a trade-off between precision, learnability and ease of production [48,49]. That is, expressivity, compositionality and stability are not the only universal properties that are central to natural language. Fundamentally, however, in addition to being able to support communication, all languages must share one crucial property: they must be learnable. As described above, XCS languages are more easily learned and the ILM demonstrates, in a simple context, how languages might spontaneously evolve to have formal properties that make them easier to learn from a limited corpus of exemplars.

Future work: richer languages The original obverter ILM and the semi-supervised ILM introduced here have a very simple model of a language, consisting entirely of single bit 'words' and 'facts' and with signal and meaning spaces of matching sizes. This is sufficient to explore concepts like expressivity, compositionality and stability but further work should include more of the complexity and structure of language and of the meanings that language is used to communicate. It is hoped that the lower computational complexity of the semi-supervised ILM will make it possible to begin to address this shortcoming in future.

A similar hope is that further elaboration of the model can incorporate more of the *structure* present in language or in putative 'spaces of states of the world'; this includes the complex hierarchical interplay of different types of facts, be they objects, both real and abstract, actions and modifiers along with a parallel set of factors related to biases, ecological salience, novelty and familiarity. In both the obverter ILM and semi-supervised ILM simple rules defining 'bottom up' ingredients in language evolution, such as learning and a generational structure lead to languages that incorporate what at first seems like formal 'top down' requirements like 'a word codes for the same fact, irrespective of context'. As such, it is exciting to ask if applying the ILM to the more complex examples that the semi-supervised ILM will make computationally feasible will help us to understand how additional complex and 'formal' rules can arise. Indeed, the most exciting possibility is that ILMs provide a tool for understanding how simple rules, like the generational structure and the bottleneck, can lead to more complex and

more ‘formal’ rule structures emerging in languages, just as both the obverter ILM and semi-supervised ILM lead to XCS languages, that is languages that incorporate apparently formal requirements like ‘a word codes for the same fact, irrespective of context’.

Future work: communities The ILM has a very simple tutor-pupil structure; in reality, humans do not learn from a single tutor. In some societies, such as those described in [43] or in other earlier, albeit contested, anthropological accounts such as [50], children at the age language learning occurs are embedded in a rich social environment, but are not in any obvious way ‘tutored’. Similarly the wealth of data provided by the difference in experience between children who spend time at a nursery or crèche and those who are cared for more exclusively in the family home, does not seem to produce any clear difference in the speed of language acquisition.

Furthermore, language acquisition and language change continue past early childhood and, indeed, the language of teenagers is notoriously mutable. Interestingly, there is even evidence from human experiments with an artificial language that compositionality can emerge without inter-generational transmission [51]. All this indicates that a model of language development should not be restricted to isolated tutor-pupil interactions. In this, the semi-supervised ILM presents an advantage. In the obverter ILM obversion takes place at a discrete transition from pupil to tutor while for the obverter ILM there is no such transition: an agent can learn and teach at the same time. Consequently, the semi-supervised ILM should lend itself to models where agents learn from each other as a community.

Two different versions of the autoencoder There are two obvious ways to add an autoencoder phase to learning in the ILM, the approach used here with the autoencoder mapping meaning to meaning and the ‘inside-out’ version of this with $a = e \circ d$ mapping signal to signal. The first represents the child describing observations to themselves, while the second represents the child overhearing speech and trying to guess its meaning. It is clear that both types of autoencoder could be justified as representing an aspect of the real-world situation and one approach would be to use both. For simplicity, we have not reported that here, but in simulations using both there appears to be an even faster evolution of an XCS language. However, using only the signal-to-signal autoencoder does not work: it appears that one of the barriers to language evolution, at the initial stage, is that all meanings map to a small subset of the possible signals and so the set of signals generated by the tutor encoder will include only a few signals and training the autoencoder on this small set reduces the expressivity of the language still further.

Supervised and unsupervised learning Currently the pupil learns in part through supervised learning. The tutor supplies input-output pairs and these are used in a classic neural network learning paradigm. However, as discussed above, directly supervised learning, as in ‘here is a cat’ / ‘what is that animal?’, is only a small part of a child’s learning experience. Indeed, if supervised learning was the key to child learning then language supervision would most likely be highly optimized and hence highly preserved across cultures. In fact, this is far from the case, for example, in [52] language learning in three different cultures is compared and, for example, whereas parents in white middle-class Anglo-American families look directly at a child learner and talk to it in the modified speech register often called parentese, in Kaluli families mothers carry the child in such a way that it is facing in the same direction as she is while answering on the child’s behalf in tripartate interactions using an unmodified version of Kaluli. In the semi-supervised ILM, the autoencoder, which is the analogue of reflection using internalized language, is an important part of language learning; it appears to be the key to a language developing with the key properties of expressivity and compositionality, while supervised language on serves to ‘nudge’ that internal dynamics towards a language that matches the tutor’s.

Conclusion

Simple models like the ILM dispense with the semiotics and logic of language; they do not allow for the possibility that an understanding of meaning and of logic might be required to account for the structure of language. The agents are also unreflective, in the sense that they do not attempt to identify deficiencies in their language with a view to improving it. This does not make them uninteresting; it is useful to note how even the simplest of models, such as the ILM, can demonstrate possible mechanisms for the presence of features in language and of the emergence of properties that might otherwise have been attributed to more complex design.

This study introduces a novel ILM that integrates both supervised and unsupervised learning to simulate the evolution of expressive, compositional, and stable languages without relying on observation. By employing neural networks for both encoding and decoding processes, and, crucially, combining them in the form of an autoencoder during unsupervised learning, the model can claim greater ecological validity, mirroring in this respect the complex dynamics of human language acquisition. This underscores the importance of semi-supervised learning in the development of language structures, offering new insights into the mechanisms underlying language evolution and suggesting that languages can evolve properties that make them easier to learn.

Materials and methods

The obverter ILM

The map \hat{d} is an all-to-all feedforward neural network with one hidden layer. All layers have size n and are initialized using the default `flux.jl` Xavier initialization. Both the hidden layer and the output layer have sigmoid non-linearities. The network is trained using stochastic gradient descent with learning rate $\eta = 1.0$ for twenty epochs. In each epoch every pair in \mathcal{B} is presented once; the order of presentation is random for each epoch. Since it has no tutor, the initial tutor calculates its encoder through obverting based on a decoder which is an untrained neural network.

The semi-supervised ILM

The maps \hat{d} and \hat{e} are all-to-all feedforward neural networks with one hidden layer. Except in Fig 7D–7F all layers have size n . All layers are initialized using the default `flux.jl` Xavier initialization. Both the hidden layer and the output layer have sigmoid non-linearities. The map \hat{a} is $\hat{d} \circ \hat{e}$ and so it is an all-to-all feedforward neural network with three hidden layers. The three networks are trained at the same time using stochastic gradient descent with learning rate $\eta = 5.0$. The number of epochs is again 20. At the start of each epoch two copies of \mathcal{B} are made, each in a different random order; at each iteration an item from the first copy is presented to \hat{d} and an item from the other to \hat{e} ; $r = 20$ randomly chosen signals from \mathcal{A} are then presented to \hat{a} . Having no tutor itself, the first tutor uses an untrained encoder to supervise its pupil.

Quantifying the ILM properties

Precise definitions of x , c and s are given in S1 Appendix. These three quantities all have values between zero and one and each has a non-zero bias which depends on n . To make comparison across n easier, in all the graphs presented here the bias has been removed so the

quantity plotted is $(y - y_0)/(1 - y_0)$ where in each case y_0 is a suitable estimate of the background value. For x and c this is the value calculated for a naïve agent, averaged over 40 examples, for s this is the value calculated between 20 pairs of naïve agents.

Software and libraries

All code was written in Julia and run on Julia 1.9.3, neural networks were trained using Flux v0.14.6 and plotting used Gadfly v1.3.4 and Colors v0.12.10. Data was handled using DataFrames v1.5.0 and CSV v0.10.11. The code is available at github.com/IteratedLM/2023_12_ailm.

Supporting information

S1 Appendix. Quantifying language properties.

(PDF)

S2 Appendix. Obversion.

(PDF)

S3 Appendix. Pseudocode.

(PDF)

S4 Appendix. The one-way ILM.

(PDF)

S1 Fig. The semi-supervised ILM performs better for larger \mathcal{A} . A–C plot mean values for x , c and s as a function of the size of \mathcal{A} , with $|\mathcal{B}|$ fixed at 160, after 25 generations (solid line) and five generations (dashed line) for 25 independent replicates per data point with $n = 16$. The ribbons depict standard deviations around the 25-generation mean.

(PDF)

S2 Fig. For the semi-supervised ILM the ideal bottleneck size is smaller than you might expect. For each bottleneck size considered, 25 replicates of the model are run until values of an XCS language develops. The best value of the bottleneck was estimated and this is plotted in **A**; the line shows the standard best linear fit, with slope 8.4 and intercept -11.5. The mean number of generations taken to reach the λ threshold at this bottleneck size is plotted in **B** in red, in magenta is the mean number of generations when a bottleneck is used which is one away from the optimal value.

(PDF)

Acknowledgments

We are grateful to Marcos Oliveira who suggested that we investigate further the surprisingly strong performance of the semi-supervised ILM when there is no bottleneck. Calculations were carried out using the computational facilities of the Advanced Computing Research Centre at the University of Bristol, www.bristol.ac.uk/acrc/. We are grateful to Dr Stewart whose philanthropy provided some of the compute resource used in this project.

Author contributions

Conceptualization: Jack Bunyan, Seth Bullock, Conor Houghton.

Investigation: Seth Bullock, Conor Houghton.

Methodology: Seth Bullock, Conor Houghton.

Project administration: Conor Houghton.

Supervision: Seth Bullock, Conor Houghton.

Writing – original draft: Conor Houghton.

Writing – review & editing: Jack Bunyan, Seth Bullock, Conor Houghton.

References

1. Chomsky N. A review of B. F. Skinner's verbal behavior. *Language*. 1959;35(1):26–58.
2. Jackendoff R. The paradox of language acquisition and the platonic perspective. *Linguist Inquiry*. 1991;22(3):407–37.
3. Christiansen MH, Chater N. Language as shaped by the brain. *Behav Brain Sci*. 2008;31(5):489–508; discussion 509–58. <https://doi.org/10.1017/S0140525X08004998> PMID: 18826669
4. Frege G. On sense and reference. *Zeitschrift für Philosophie und philosophische Kritik*. 1892;100(1):25–50.
5. Chomsky N. *Syntactic structures*. Mouton de Gruyter; 1957.
6. Davidson D. *Inquiries into truth and interpretation: philosophical essays*, vol. 2. Clarendon Press; 2001.
7. Boole G. *An investigation of the laws of thought: on which are founded the mathematical theories of logic and probabilities*. vol. 2. Walton and Maberly; 1854.
8. Kirby S. Syntax without natural selection: how compositionality emerges from vocabulary in a population of learners. In: Knight C, Studdert-Kennedy M, Hurford JR, editors. *The evolutionary emergence of language*. Cambridge, UK: Cambridge University Press; 2000. p. 303–23.
9. Kirby S. Spontaneous evolution of linguistic structure—an iterated learning model of the emergence of regularity and irregularity. *IEEE Trans Evol Computat*. 2001;5(2):102–10. <https://doi.org/10.1109/4235.918430>
10. Kirby S. Learning, bottlenecks and the evolution of recursive syntax. In: Briscoe E, editor. *Linguistic evolution through language acquisition: formal and computational models*. Cambridge, UK: Cambridge University Press; 2002. p. 173–204.
11. Kirby S, Hurford J. The emergence of linguistic structure: an overview of the iterated learning model. In: Cangelosi A, Parisi D, editors. *Simulating the evolution of language*. London: Springer; 2002. p. 121–48.
12. Brighton H. Compositional syntax from cultural transmission. *Artif Life*. 2002;8(1):25–54. <https://doi.org/10.1162/106454602753694756> PMID: 12020420
13. Oliphant M, Batali J. Learning and the emergence of coordinated communication. *Center Res Lang Newslett*. 1997;11(1):1–46.
14. Choi E, Lazaridou A, de Freitas N. Compositional obverter communication learning from raw visual input. In: *Proceedings of the 6th International Conference on Learning Representations (ICLR)*. ICLR; 2018. p. rkn2Be0.
15. Lewis D. *Convention: a philosophical Study*. Wiley; 1969.
16. Brace L, Bullock S, Noble J. Achieving compositional language in a population of iterated learners. In: Andrews P, Caves L, Doursat R, Hickinbotham S, Polack F, Stepney S, et al., editors. *ECAL 2015: Proceedings of the Thirteenth European Conference on Artificial Life*. MIT Press; 2015. p. 349–56.
17. Brace L, Bullock S. Understanding language evolution in overlapping generations of reinforcement learning agents. In: Gershenson C, Froese T, Siqueiros JM, Aguilar W, Izquierdo EJ, Sayama H, editors. *Artificial Life XV: Proceedings of The Fifteenth International Conference on the Synthesis and Simulation of Living Systems*. MIT Press; 2016. p. 492–9.
18. Sains G, Houghton C, Bullock S. Spatial community structure impedes language amalgamation in a population-based iterated learning model. In: Iizuka H, Suzuki K, Uno R, Damiano L, Sychala N, Aguilera M, et al., editors. *Ghost in the Machine: Proceedings of the 2023 Artificial Life Conference*; 2023. p. 377–85.
19. Griffiths TL, Kalish ML. A Bayesian view of language evolution by iterated learning. In: Bara BG, Salvendy G, editors. *Proceedings of the 27th Annual Meeting of the Cognitive Science Society*; 2005. p. 827–32.

20. Kirby S, Dowman M, Griffiths TL. Innateness and culture in the evolution of language. *Proc Natl Acad Sci U S A*. 2007;104(12):5241–5. <https://doi.org/10.1073/pnas.0608222104> PMID: 17360393
21. Kirby S, Tamariz M. Cumulative cultural evolution, population structure and the origin of combinatoriality in human language. *Philos Trans R Soc Lond B Biol Sci*. 2022;377(1843):20200319. <https://doi.org/10.1098/rstb.2020.0319> PMID: 34894728
22. Griffiths TL, Kalish ML. Language evolution by iterated learning with bayesian agents. *Cogn Sci*. 2007;31(3):441–80. <https://doi.org/10.1080/15326900701326576> PMID: 21635304
23. Smith K, Kirby S. Cultural evolution: implications for understanding the human language faculty and its evolution. *Philos Trans R Soc Lond B Biol Sci*. 2008;363(1509):3591–603. <https://doi.org/10.1098/rstb.2008.0145> PMID: 18801718
24. Steels L. The spontaneous self-organization of an adaptive language. *Mach Intell*. 1999;15:205.
25. Lu Y, Singhal S, Strub F, Courville A, Pietquin O. Countering language drift with seeded iterated learning. In: Daum'e III H, Singh A, editors. *International Conference on Machine Learning*. PMLR; 2020. p. 6437–47.
26. Ren Y, Guo S, Labeau M, Cohen SB, Kirby S. Compositional languages emerge in a neural iterated learning model. In: *8th International Conference on Learning Representations, ICLR*; 2020. p. HkePNpVKPB.
27. Aldous DJ, Ibragimov IA, Jacod J, Aldous DJ. *Exchangeability and related topics*. Berlin: Springer; 1985.
28. Reali F, Chater N, Christiansen MH. Simpler grammar, larger vocabulary: how population size affects language. *Proc Biol Sci*. 2018;285(1871):20172586. <https://doi.org/10.1098/rspb.2017.2586> PMID: 29367397
29. Fedzechkina M, Newport EL, Jaeger TF. Miniature artificial language learning as a complement to typological data. In: Ortega L, Tyler AE, Park HI, Uno M, editors. *The usage-based study of language learning and multilingualism*. Washington, DC: Georgetown University Press; 2016. p. 211–32.
30. Smith K, Perfors A, Fehér O, Samara A, Swoboda K, Wonnacott E. Language learning, language use and the evolution of linguistic variation. *Philos Trans R Soc Lond B Biol Sci*. 2017;372(1711):20160051. <https://doi.org/10.1098/rstb.2016.0051> PMID: 27872370
31. Culbertson J, Schuler K. Artificial language learning in children. *Annu Rev Linguist*. 2019;5(1):353–73. <https://doi.org/10.1146/annurev-linguistics-011718-012329>
32. Kirby S. Culture and biology in the origins of linguistic structure. *Psychon Bull Rev*. 2017;24(1):118–37. <https://doi.org/10.3758/s13423-016-1166-7> PMID: 28120320
33. Boeckx C. *Reflections on language evolution: from minimalism to pluralism*. Language Science Press; 2021.
34. Newport EL, Gleitman LR, Gleitman H. Mother, infant, and speech environment: The case for a social-cognitive account of language acquisition. In: Lewis M, Rosenblum LA, editors. *Friendship and peer relations*. Wiley; 1977. p. 137–75.
35. Gopnik A, Choi S. Do linguistic differences lead to cognitive differences? A cross-linguistic study of semantic and cognitive development. *First Lang*. 1990;10(30):199–215. <https://doi.org/10.1177/014272379001003002>
36. Yurovsky D, Doyle G, Frank MC. Linguistic input is tuned to children's developmental level. In: Papafragou A, Grodner D, Mirman D, Trueswell JC, editors. *Proceedings of the 38th Annual Meeting of the Cognitive Science Society*. Austin, TX: Cognitive Science Society; 2016. p. 2093–8.
37. Akhtar N, Behl C. Passive listening during storybook reading enhances vocabulary development in preschoolers. *Reading and Writing*. 2019;32(8):1357–76.
38. Rowe ML, Snow CE. Analyzing input quality along three dimensions: interactive, linguistic, and conceptual. *J Child Lang*. 2020;47(1):5–21. <https://doi.org/10.1017/S0305000919000655> PMID: 31668157
39. Kachergis G, Loukatou G, Frank MC. Identifying the distributional sources of children's early vocabulary. In: Culbertson J, Perfors A, Rabagliati H, Ramenzoni V, editors. *Proceedings of the 44th Annual Meeting of the Cognitive Science Society*; 2022. p. 3202–8.
40. Brown R, Hanlon C. Derivational complexity and order of acquisition in child speech. In: Hayes JR, editor. *Cognition and the development of language*. Wiley; 1970. p. 11–53.
41. Hirsh-Pasek K, Treiman R, Schneiderman M. Brown & Hanlon revisited: mothers' sensitivity to ungrammatical forms. *J Child Lang*. 1984;11(1):81–8. <https://doi.org/10.1017/s0305000900005596> PMID: 6699114
42. Morgan JL, Travis LL. Limits on negative information in language input. *J Child Lang*. 1989;16(3):531–52. <https://doi.org/10.1017/s0305000900010709> PMID: 2808572

43. Shneidman LA, Arroyo ME, Levine SC, Goldin-Meadow S. What counts as effective input for word learning?. *J Child Lang*. 2013;40(3):672–86. <https://doi.org/10.1017/S0305000912000141> PMID: 22575125
44. LaTourrette A, Waxman SR. A little labeling goes a long way: semi-supervised learning in infancy. *Dev Sci*. 2019;22(1):e12736. <https://doi.org/10.1111/desc.12736> PMID: 30157311
45. Bröker F, Love BC, Dayan P. When unsupervised training benefits category learning. *Cognition*. 2022;221:104984. <https://doi.org/10.1016/j.cognition.2021.104984> PMID: 34954447
46. Chomsky N. *On nature and language*. Cambridge University Press; 2002.
47. Crain S, Pietroski P. *Linguist Philos*. 2001;24(2):139–86. <https://doi.org/10.1023/a:1005694100138>
48. Moro A. *The boundaries of babel: the brain and the enigma of impossible languages*. vol. 46. Cambridge, MA: MIT Press; 2015.
49. Hahn M, Jurafsky D, Futrell R. Universals of word order reflect optimization of grammars for efficient communication. *Proc Natl Acad Sci U S A*. 2020;117(5):2347–53. <https://doi.org/10.1073/pnas.1910923117> PMID: 31964811
50. Mead M. *Coming of age in samoa: a psychological study of primitive youth for western civilization*. New York, NY: William Morrow and Co.; 1928.
51. Raviv L, Meyer A, Lev-Ari S. Compositional structure can emerge without generational transmission. *Cognition*. 2019;182:151–64. <https://doi.org/10.1016/j.cognition.2018.09.010> PMID: 30267952
52. Ochs E, Schieffelin B. *Language acquisition and socialization: three developmental stories and their implications*. *Linguist Anthropol: A Reader*. 2001;2001:263–301.