

# Modeling language contact with the Iterated Learning Model

Seth Bullock and Conor Houghton,

Intelligent Systems Laboratory, University of Bristol, United Kingdom conor.houghton@bristol.ac.uk

## Abstract

Contact between languages has the potential to transmit vocabulary and other language features; however, this does not always happen. Here, an iterated learning model is used to examine, in a simple way, the resistance of languages to change during language contact. Iterated learning models are agent-based models of language change, they demonstrate that languages that are expressive and compositional arise spontaneously as a consequence of a language transmission bottleneck. A recently introduced type of iterated learning model, the Semi-Supervised ILM is used to simulate language contact. These simulations do not include many of the complex factors involved in language contact and do not model a population of speakers; nonetheless the model demonstrates that the dynamics which lead languages in the model to spontaneously become expressive and compositional, also cause a language to maintain its core traits even after mixing with another language.

## Introduction

Iterated Learning Models (ILMs) are agent-based models in which languages evolve as they are transmitted from generation to generation. They were proposed as a description of how languages become compositional and expressive (Kirby, 2001; Kirby and Hurford, 2002). Here, the Semi-Supervised ILM, a specific version of the model recently introduced in Bunyan et al. (2024), is used to simulate the effect on a language of contact with another language.

The continuous processes required to produce the vast array of languages currently in use seems to indicate that languages are extremely mutable and unstable. Languages certainly change. For example, since Jonathan Swift's novel *Gulliver's Travels* was published less than three hundred years ago on 28 October 1726, the English lexicon has expanded enormously to provide a vocabulary for the many novel objects, behaviours and social structures that have resulted from the Industrial Revolution, the Age of Revolution and the Computer Age. However, for all that, the language has not changed that much: though unarguably 'old-fashioned' in its phrasing, *Gulliver's Travels* is perfectly comprehensible to the contemporary reader.

This is striking. Between Swift's era and the present day, English has become a *lingua franca*, spoken for periods as the language of administration in countries within which it was not native and used as a second language by more people than use it as a first language. Consequently, it might be expected that it would have changed dramatically with a huge number of words from other languages crowding out older words and rendering the novel incomprehensible. Indeed the idea that English continually adds new words by taking them from other languages is something of a cliché.

However, it is not really true. Taking the example of English and Irish (endonym: Gaeilge, ISO 639-3: gle); Jonathon Swift was among the first prominent Irish writers writing in English and there has been a rich history of interaction between speakers of Irish and of English in the period since *Gulliver's Travels* was written. Ultimately an English language has replaced Irish as the main language of Ireland and the descendants of the substantial Irish diaspora are for the most part English speakers. Nonetheless only a handful of words derived from Irish are in common use in English. Indeed, while it is claimed that the aboriginal Celtic languages of Great Britain and Ireland form a language substrate for English (Filppula et al., 2008), even in Ireland any more recent influence of Irish on the grammar, pronunciation and lexicon of the language now spoken is apparent only in small, though valued, ways. In fact, as in the example of the word 'oxter' for 'armpit', many of the lexical differences between Irish-English and other Englishes represent the preservation in Irish-English of English words that have fallen out of use elsewhere, rather than the influence of Irish.

Of course, the dynamics of language contact is complex and it is certainly not invariably true that one language retreats while leaving the other unchanged. Indeed, in the case of English, the Norman Conquest and the influence of Old French (ISO 639-3: fro) caused substantial changes, marked by the transition from Old English (ISO 639-3: ang) to Middle English (ISO 639-3: enm). Old French in turn showed considerable influences from Old Frankish (ISO 639-3: frk), though this is a salutary example; the main effect of Old

Frankish was to change stress pattern in pronunciation, disrupting the grammatical information in word endings that Vulgar Latin relied on, leading to a shift from a synthetic to an analytic language. In other words, although contact with Old Frankish was important in the change from Vulgar Latin to Old French, these changes did not mean Old French resembled Old Frankish any more than Vulgar Latin did Rickard (2003); Lodge (2013).

The various outcomes of language contact depend on multiple factors such as prestige, economic opportunity and the need for new words for phenomena and objects familiar to speakers of one language but not the other (Thomason and Kaufman, 1988; Mufwene, 2001; Matras and Sakel, 2007; Sakel and Matras, 2008). Our purpose here is to examine the, perhaps surprising, stability of languages during asymmetric contact from the simplest possible point-of-view. We explore whether the mechanism underpinning the development of compositionality and expressivity in the Semi-Supervised ILM could also account for the stability of one of the two languages in a contact situation.


### The Iterated Language Model

Languages are made easier to learn, and easier to use, by their properties of *compositionality* and *expressivity* (Oliphant, 1999). A compositional language is one in which each element of meaning is encoded consistently by the same morpheme within the signal, irrespective of context, so ‘violin’ means much the same thing in ‘she plays the violin’ and ‘it is violin shaped’. While all human languages exhibit compositionality of some kind, they are not completely compositional, so, in other contexts the word ‘fiddle’ can be used: ‘Mairéad Ní Mhaonaigh is a leading exponent in the Donegal fiddle tradition’. An expressive language is one in which a range of different meanings are unambiguously encoded by different signals, so the viola is called a ‘viola’ in English and can be distinguished by name from a ‘violin’. However, just as real languages are not perfectly compositional, they are not perfectly expressive: the ‘fiddle’ in ‘it is tempting to fiddle with it, even though it is perfect as it is’ does not refer to a violin. Along with expressivity and compositionality, a third simple property required from a language is *stability*; while languages change, to support communication they should not change too much from one generation to the next.

ILMs are a family of computational models designed to simulate language change. They were originally intended to describe a mechanism by which languages can become expressive, compositional and stable over multiple generations of use (Kirby, 2001; Kirby and Hurford, 2002). In an ILM there is a language-using ‘tutor’ agent and a language-learning ‘pupil’ agent. The adult teaches language to the pupil by exposing it to a set of utterances. The pupil then becomes a tutor and teaches a new pupil in the same fashion. A language learning bottleneck is a key feature of these

models: since each tutor only presents a finite subset of the language to their pupil, when the pupil in turn becomes an adult it will need to generate utterances that it has not experienced before. The process of generalization inherent in the model is thought to be critical to the behaviour of the ILM.

In the version of the ILM considered here a language consists of a set of meaning,  $\mathcal{M}$ , and a set of signals,  $\mathcal{S}$ , both represented by binary vectors, typically here these will be length ten vectors. In addition the agent has an encoder map  $e : \mathcal{M} \rightarrow \mathcal{S}$  and a decoder map  $d : \mathcal{S} \rightarrow \mathcal{M}$  mapping signals to meanings and vice versa, respectively:



A language, in this context, corresponds to a map from meanings to signals and so an agent’s language can be thought of as being equivalent to its encoder. In Kirby and Hurford (2002) the decoder  $d$  is a simple neural network whereas the encoder  $e$  is a look-up table listing all possible meanings and their corresponding signal. During learning a ‘bottleneck’ subset of meanings  $\mathcal{B} \in \mathcal{M}$  is chosen and the tutor presents this training meaning-signal pairs  $\{(e(b), b) | b \in \mathcal{B}\}$ . This is used to train the pupil’s decoder. At the end of training, when the pupil ‘grows up’, an inversion process, known as *obversion* (Oliphant and Batali, 1997), is used to build an encoder for this agent, after which it can become a tutor for a newly created naïve pupil.

During obversion, the pupil’s decoder is used to assign a probability to every possible signal-meaning pair, the encoder map from meaning to signal is then calculated by a process of ‘picking winners’. This is computationally expensive to the extent that it makes simulations impractical for languages with bit length greater than about ten. It is also unrealistic as it imagines that a pupil is incapable of an utterances until it matures, at which point it must consider all possible signals and all possible meaning and match them to calculate the encoder that it will rely on in order to be able to communicate as an adult tutor.

In a recent paper, Bunyan et al. (2024) introduce a new ILM which is capable of demonstrating the emergence of expressivity, compositionality and stability under bottlenecked conditions without the need for obversion. In this ILM, both the decoder and the encoder are neural networks. During learning the tutor presents pairs  $\{(e(b), b) | b \in \mathcal{B}\}$  to both the pupil’s decoder and to its encoder. In addition, the pupil creates a third neural network,  $a = d \circ e$ , by concatenating the encoder and decoder. This third network is presented with meanings from a set  $\mathcal{A} \in \mathcal{M}$  and it learns as an autoencoder; this is intended to mimic a process of internal learning in which a child sees objects and events and imagines how they might be expressed in its language. This is

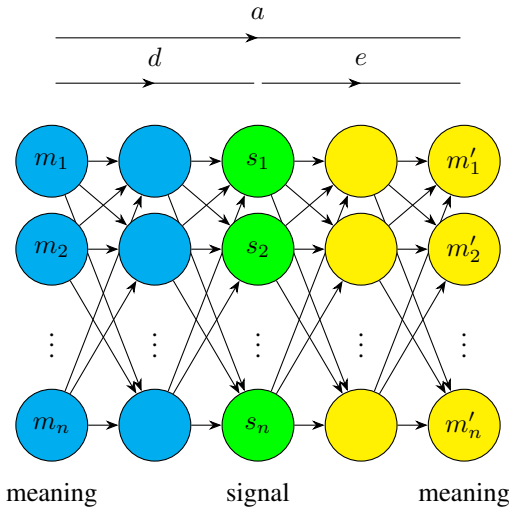


Figure 1: **Neural network.** In the Semi-Supervised ILM both the encoder,  $e$  (yellow), and decoder,  $d$  (blue), are neural networks; in  $e$  the meaning vector forms the input, it is mapped to hidden layer and then to a signal vector; in  $d$  a signal vector forms the input, there is again a hidden layer and the output is a meaning vector. The autoencoder,  $a$ , is formed by concatenating the two neural networks, with the signal layer, green, in this case forming a third hidden layer for the autoencoder.

illustrated in Fig. 1. In Bunyan et al. (2024) it is noted that this ILM is considerably less computationally burdensome and it is argued that it is more realistic as a model of language development. This new version of the ILM is used here and is referred to as the Semi-Supervised ILM.

### Modeling language contact

The ILM was proposed to model the way in which a language can acquire expressivity, compositionality and stability through iterated learning with a bottleneck. Here though, another case is considered in which there are two pre-existing languages; the  $\mathbb{A}$  language and the  $\mathbb{B}$  language, each of which are both expressive and compositional. In our iterated language model the first pupil will be taught a mixture of these two languages, so after the training set of meanings  $\mathcal{B}$  has been selected from  $\mathcal{M}$  some meanings will be encoded according to the  $\mathbb{A}$  language and some according to the  $\mathbb{B}$  language: specifically for each meaning in  $\mathcal{B}$  the  $\mathbb{A}$  language signal will be selected with probability  $p$  and, conversely, the  $\mathbb{B}$  language meaning with probability  $1 - p$ . This is intended to reflect a difference in the number of exemplars presented to a child in a language contact situation,

$n_1$	bit length of the meaning vectors.
$n_2$	number of nodes in the hidden layers.
$n_3$	bit length of the signal vectors.
$n_1 \times n_2 \times n_3$	agent with a $n_1 \times n_2 \times n_3$ encoder and a $n_3 \times n_2 \times n_1$ decoder.
$\mathcal{M}$	set of meanings.
$\mathcal{S}$	set of signals.
$\mathcal{B} \subset \mathcal{M}$	used for supervised learning.
$\mathcal{A} \subset \mathcal{M}$	used for unsupervised learning.
$\mathbb{A}$ and $\mathbb{B}$	the two initial languages.
$p$	the language mixing parameter.
$d : \mathcal{S} \rightarrow \mathcal{M}$	decoder map.
$e : \mathcal{M} \rightarrow \mathcal{S}$	encoder map.
$a : \mathcal{M} \rightarrow \mathcal{M}$	autoencoder map $a = d \circ e$ .
$x$	expressivity (blue).
$c$	compositionality (orange).
$s$	stability (maroon in graphs).
$a$	similarity to language $\mathbb{A}$ (red).
$b$	similarity language $\mathbb{B}$ (red).

Table 1: **Glossary.** Here a *meaning* is a  $n_1$  bit length binary vector whose components are facts or elements of meaning and a *signal* is a  $n_3$  bit length binary vector whose components are morphemes. There is a notation clash,  $a$  is used both for the encoder map and the similarity to language  $\mathbb{A}$ , hopefully which is clear in context. Another near clash is the use of  $\mathbb{A}$  and  $\mathbb{B}$  for languages and  $\mathcal{A}$  and  $\mathcal{B}$  for example sets, here typeface is used to make the distinction.

or, more generally,  $p$  is supposed to represent any weighting, such as the fraction of  $\mathbb{A}$  exemplars or the authority and emphasis with which those examples are presented. After this first initial generation the language is evolved through the usual iterated learning process and its similarity,  $a$  and  $b$ , to the two original languages is quantified. The calculation of these similarities is discussed below.

### Methods

The maps  $d$  and  $e$  are all-to-all feedforward neural networks with one hidden layer. It is convenient to introduce a shorthand for the size of each layer of the encoder  $e$ :  $n_1 \times n_2 \times n_3$ ;  $n_1$  is the bit length of the meaning vectors,  $n_2$  the number of nodes in the hidden layer and  $n_3$  the bit length of the signal vectors, the decoder  $d$ , in turn, is a  $n_3 \times n_2 \times n_1$  neural network. This, and other pieces of notation are summarized in Table 1.

Both the hidden layer and the output layer have sigmoid non-linearities. Since  $d$  and  $e$  map to binary vectors the output values are binarized by setting them to zero or one using a threshold of 0.5. However, the signal layer forms the middle layer of the autoencoder  $a$  and so the rounding step is not performed on this layer in  $a$ : back-propagation requires that the values in the hidden layer remain continuous and, in

any case, the autocoder is intended to model the child contemplating the range of different possible utterances likely to correspond to a meaning.

All weights are initialized using the default `flux.jl` Xavier initialization. The three networks are trained at the same time using stochastic gradient descent with learning rate  $\eta = 5.0$ . The number of epochs is 20. For each generation a bottleneck set,  $\mathcal{B}$ , and an autoencoder set  $\mathcal{A}$  are chosen randomly from  $\mathcal{M}$ . 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  $d$  and an item from the other to  $e$ ; for the  $10 \times 10 \times 10$ ,  $10 \times 12 \times 10$ ,  $9 \times 11 \times 12$  and  $10 \times 15 \times 20$  agents  $r = 15$  randomly chosen meanings from  $\mathcal{A}$  are then presented to  $a$ , for the  $20 \times 30 \times 20$  network  $r$  is increased to 30.

Two initial languages  $\mathbb{A}$  and  $\mathbb{B}$  are formed by randomly shuffling the ‘identity language’, the language in which each meaning vector is mapped to the identical signal vector. The shuffling is done in a way that preserves compositionality, this means that each bit in the meaning determines a different specific bit in the signal, though, unlike in the identity language, it will not, in general, map to a bit in the same position and ‘1’ may map to a ‘0’ rather than to ‘1’. In most simulations the meaning and signal space are the same length. In the example where the signal space is larger the unused bits, which after shuffling are located in specific random positions along the signal vector, are randomly set to zero or one for all signals, that is, they remain the same irrespective of meaning. The mixed language used to teach the first pupil is calculated as a look up table where for each meaning in  $\mathcal{M}$  the encoded signal from the  $\mathbb{A}$  language is chosen with probability  $p$ , otherwise the  $\mathbb{B}$  language is used.

### Quantifying language properties

A Hamming similarity is used to quantify the similarity between languages. The fraction of meanings on which they agree is calculated:

$$\tilde{f} = \frac{|\{m : e_1(m) = e_2(m) \text{ for } m \in \mathcal{M}\}|}{|\mathcal{M}|} \quad (1)$$

where  $e_1$  and  $e_2$  are the two encoder maps. A background value,  $f_0$ , is calculated for random pairs of languages, the similarity is then

$$f = \frac{\tilde{f} - f_0}{1 - f_0}, \quad (2)$$

Here  $a$  is used to denote the similarity from the  $\mathbb{A}$  language and  $b$  from the  $\mathbb{B}$  language. In addition, expressivity, compositionality and stability are calculated; of these, stability is the most straightforward, it is the similarity between a language and the language for the previous generation. Expressivity is based on the fraction of the signal space that is employed by the language

$$\tilde{x} = \frac{|\{e(m) \text{ for } m \in \mathcal{M}\}|}{|S|} \quad (3)$$

where  $x$  is calculated from  $\tilde{x}$  by removing the background value as before. It is more difficult to define a good measure of compositionality; roughly speaking the measure used here is based on taking each bit in the meaning vector and calculating the conditional entropy, over all values of the other bits in the meaning vector, of the best predicted bit in the signal. This is described in more detail in Bunyan et al. (2024), in summary though,  $c$ , is one for a perfectly compositional language and zero if the language is no more compositional than chance.

### Simulations

For all the simulations where the meaning vector has bit length ten  $|\mathcal{B}| = 80$  and  $|\mathcal{A}| = 240$ ; for the example where the meaning vector has bit length 20 a larger bottleneck is used:  $|\mathcal{B}| = 185$  and  $|\mathcal{A}| = 555$ . Results for 50 simulations are presented for all examples.

### 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, Compose v0.9.5 and Colors v0.12.10. Data was handled using DataFrames v1.5.0 and CSV v0.10.11. The code used to produce the figures and the corresponding simulated data is available at [github.com/IteratedLM/2024\\_04.mix](https://github.com/IteratedLM/2024_04.mix).

## Results

Figure 2 shows the result of language mixing for three different values of the mixing parameter  $p$ . The initial languages  $\mathbb{A}$  and  $\mathbb{B}$  are perfectly expressive and compositional. However, the mixed language is not and there is initially a low value for  $x$  and  $c$ . Nonetheless the language quickly recovers these properties through iterated learning: within ten generations the languages are expressive and compositional and become stable from generation to generation. This is true for all values of  $p$ , however, if  $p = 0.5$  the language rarely ends up matching either  $\mathbb{A}$  or  $\mathbb{B}$  while for  $p = 0.75$  it almost invariably matches  $\mathbb{A}$ .

This is further illustrated in Fig. 3 where the fraction of languages that have  $a > 0.9$  or  $b > 0.9$  is plotted against the value of  $p$  (see Fig. 3A) and the average value of  $a$  and  $b$  is plotted, again, against  $p$  (see Fig. 3B). It is apparent that there is a very rapid transition from the behaviour near  $p = 0.5$  where a new language is often formed, to the behaviour away from  $p = 0.5$  where one of the two original languages is quickly reconstructed from the elements that are present in the original mixture. What has not been assessed is the nature of the mixed languages at  $p = 0.5$ ; although typically they are not similar to either  $\mathbb{A}$  nor  $\mathbb{B}$  it is still possible that they contains some evidence of each as is the case with mixed languages like Maltese (ISO 639-3: mlt). However, Fig. 3B indicates that this is rare, the low average value of  $a$  is not the results of a large number of

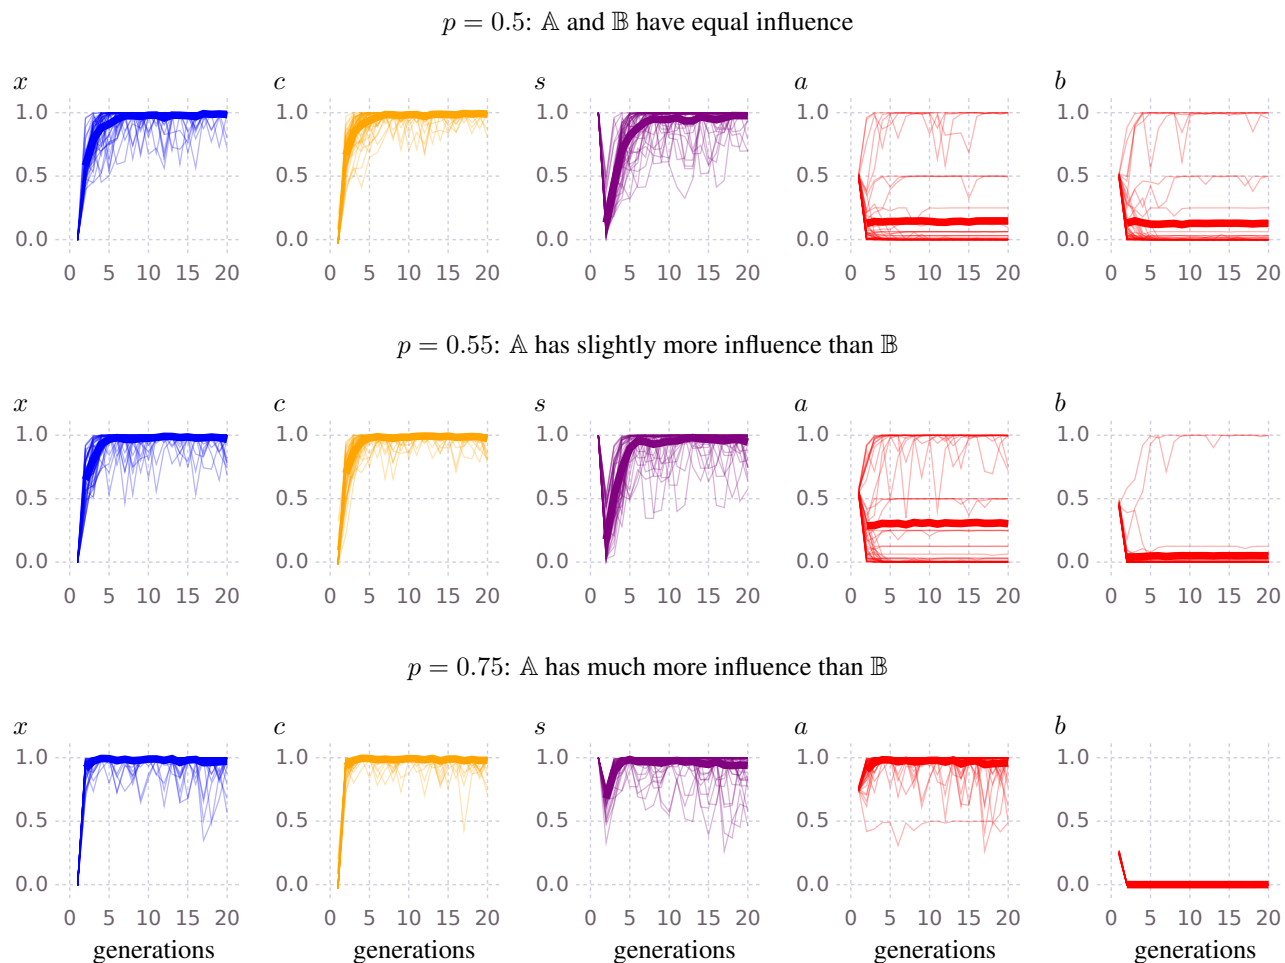


Figure 2: **Evolution of the mixed language.** Here the language is evolved for 20 generations and the results of 50 simulations are plotted for five different language attributes. In every graph the individual simulations are plotted with thin pale lines, the mean with a thicker, darker line. The first three columns plot attributes of the language itself, showing its expressivity ( $x$ ), compositionality ( $c$ ) and the stability ( $s$ ) of the iterated learning. The final two columns show how similar the languages are to the initial  $\mathbb{A}$  and  $\mathbb{B}$  languages, these similarities are  $a$  and  $b$ . The three rows present results for three different values of  $p$ , the fraction of the initial language taken from language  $\mathbb{A}$ , these are  $p = 0.5$  when the new language is almost always different from either  $\mathbb{A}$  or  $\mathbb{B}$ ,  $p = 0.55$  when the new language is typically similar to  $\mathbb{A}$ , sometimes to  $\mathbb{B}$  and sometimes to neither, and  $p = 0.75$  when the new language almost always resembles  $\mathbb{A}$ . In each case a  $10 \times 10 \times 10$  agent is used.




Figure 3: **Modeling language contact.** Language contact is simulated 50 times for  $p$  values from 0.5 to 1.0 in increments of 0.05. **A** plots the fraction of simulations which result in a value of  $a$ , upper line, or  $b$ , lower line, greater than 0.9. In **B**  $a$  is plotted vertically, although the simulations were only run for  $p \in [0.5, 1.0]$  there is no distinction between the  $\mathbb{A}$  language and the  $\mathbb{B}$  language, this means that the  $b$  values in  $p \in [0.5, 1.0]$  can be used to extend the range over which  $a$  is plotted to all of  $p \in [0.0, 1.0]$ . Each black dot corresponds to one of the 50 simulations at each  $p$  value. Horizontal jitter of  $\pm 0.025$  has been added to aid visualization. The red line is the average. Both **A** and **B** describe the same set of simulations using a  $10 \times 10 \times 10$  agent.

trials in which the eventual value of  $a$  is low, it is the result of a large number of trials giving an  $a$  value of zero and a smaller number where  $a = 1.0$  or  $a = 0.5$ . It will be interesting in future work to consider in more detail the language that results when  $p$  is near 0.5 and to consider more complicated multi-language mixing as happens in the formation of Creoles, though, in the case of many Creoles there are other important and challenging sociological factors involved in their formation.

In Fig. 4 the performance of the  $10 \times 10 \times 10$  Semi-Supervised ILM is compared to three other Semi-Supervised ILMs. The Semi-Supervised ILM appears to perform more robustly if the hidden layer has more nodes than the bit length of the signal and meaning spaces Bunyan et al. (2024). Here, the same set of simulations as above was performed using a Semi-Supervised ILM with a  $10 \times 12 \times 10$  agent. In addition, it might be suspected that the model returns to the  $\mathbb{A}$  language for  $p > 0.5$  because the signal space is restricted, in a perfectly expressive language all signals are used. With this in mind simulations, were performed using a Semi-Supervised ILM with a  $9 \times 11 \times 12$  agent and with a  $10 \times 15 \times 20$ . In fact, across all these examples, the behaviour of the model was largely unchanged.

A much larger language is considered in Fig. 5; here a  $20 \times 30 \times 20$  agent is used so the signal and meaning space both have bit length 20. The ILM in this case produces languages which are both expressive and compositional with values near one, as before; however the value of stability fluctuates a lot. Stability is a measure of how many meanings map to the same signal from one generation to the next,




Figure 4: **The behaviour does not change substantially when the size of hidden layer or signal space is changed.** Four different architectures are compared here. Here the average value of  $a$  is given for the  $10 \times 10 \times 10$ ,  $10 \times 12 \times 10$  and  $9 \times 11 \times 12$  and  $10 \times 15 \times 20$  agents. The four curves are similar.

in other words it operates at the level of meanings and signals, not facts and morphemes and so it is harsh test of a compositional language: if two bits are swapped, this would reduce stability by a half. In fact, the behaviour seems very similar to the 10-10-10 example, the language quickly becomes expressive and compositional, for the  $p = 0.5$  case it generally does not resemble either the  $\mathbb{A}$  or  $\mathbb{B}$  language, while for  $p = 0.75$  it largely resembles  $\mathbb{A}$ .

## Discussion

Language contact is complex, both in its internal dynamics and in the broader contexts in which language mixing occurs. The political and culture setting for language contact can be extremely consequential for the populations involved. Language loss can be traumatic and destructive and is often accompanied by other challenging forms of loss; conversely, culture exchange can be enriching and exciting and a useful mechanism under which languages become capable of more efficiently expressing new ideas. None of this complexity is included in the work presented here. The agent modeling also ignores bilingualism and code-switching. However, what has been demonstrated is that a mechanism which can explain the development of expressivity and compositionality in languages can also explain the stability of languages during language contact, even without invoking the stabilizing influence of a population of speakers not involved in contact.

Recently, population-based versions of the ILM have been considered (Brace et al., 2015; Brace and Bullock, 2016; Sains et al., 2023) and a population approach is implicit in the simulations described here in that the initial tutor, which provides examples to the first pupil agent, draws those examples from two different languages. However, this is an extremely restricted version of the richer population dynamics that studying the problem of language contact in more detail would require. In fact, because the Semi-Supervised ILM does not have a sharp distinction between its learning and teaching phases, it lends itself to population models and these will form the basis of further work.

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 Fedzechkina et al. (2016); Smith et al. (2017); Culbertson and Schuler (2019) for reviews. A full exploration of language contact should incorporate historical, often very qualitatively evidence, along with simulations and matched data from participant studies.

## Conclusion

In this paper we employed a new class of ILM to explore the outcome of contact between two mature languages that are already expressive, compositional and stable. The semi-supervised nature of this new iterated learning paradigm allows the model to dispense with the computationally expen-

sive and unrealistic oververber mechanism that has tended to limit the size and complexity of languages explored in previous ILM studies. Here, for the first time we are able to model binary languages with up to  $2^{20}$  possible meanings and signals.

The results here demonstrate that when the two pre-contact languages contribute equally to the language exposure of a naïve pupil, the resultant language often does not resemble either of the pre-contact languages, and that this is particularly true for larger meaning-signal spaces where recovery of either of the original languages is very rare. However, when one language has some degree of advantage over the other in terms of its contribution to the learning experiences of the naïve pupil, this language or something closely resembling it is often able to re-emerge after a period of language destabilization during language contact. This demonstrates that the same dynamics that drives languages to compositionality and expressivity can also stabilize them during language contact.


## Acknowledgements

SB is supported by UKRI grant EP/Y028392/1: AI for Collective Intelligence (AI4CI). Simulations were carried out using the computational facilities of the Advanced Computing Research Centre at the University of Bristol, [www.bristol.ac.uk/acrc/](http://www.bristol.ac.uk/acrc/). We are grateful to Dr Stewart whose philanthropy supported some of the compute resource used in this project.

## References

- Brace, L. and Bullock, S. (2016). Understanding language evolution in overlapping generations of reinforcement learning agents. In Gershenson, C., Froese, T., Siqueiros, J. M., Aguilar, W., Izquierdo, E. J., and Sayama, H., editors, *Artificial Life XV: Proceedings of The Fifteenth International Conference on the Synthesis and Simulation of Living Systems*, pages 492–499. MIT Press.
- Brace, L., Bullock, S., and Noble, J. (2015). Achieving compositional language in a population of iterated learners. In Andrews, P., Caves, L., Doursat, R., Hickinbotham, S., Polack, F., Stepney, S., Taylor, T., and Timmis, J., editors, *ECAL 2015: Proceedings of the Thirteenth European Conference on Artificial Life*, pages 349–356. MIT Press.
- Bunyan, J., Bullock, S., and Houghton, C. (2024). An Iterated Learning Model of language change that mixes supervised and unsupervised learning. *arXiv*, [arxiv.org/abs/2405.20818](https://arxiv.org/abs/2405.20818).
- Culbertson, J. and Schuler, K. (2019). Artificial language learning in children. *Annual Review of Linguistics*, 5:353–373.
- Fedzechkina, M., Newport, E. L., and Jaeger, T. F. (2016). Miniature artificial language learning as a complement to typological data. In Ortega, L., Tyler, A. E., Park, H. I., and Uno, M., editors, *The usage-based study of language learning and multilingualism*, pages 211–232. Georgetown University Press, Washington, DC.

$p = 0.5$ :  $\mathbb{A}$  and  $\mathbb{B}$  have equal influence



$p = 0.75$ :  $\mathbb{A}$  has much more influence than  $\mathbb{B}$

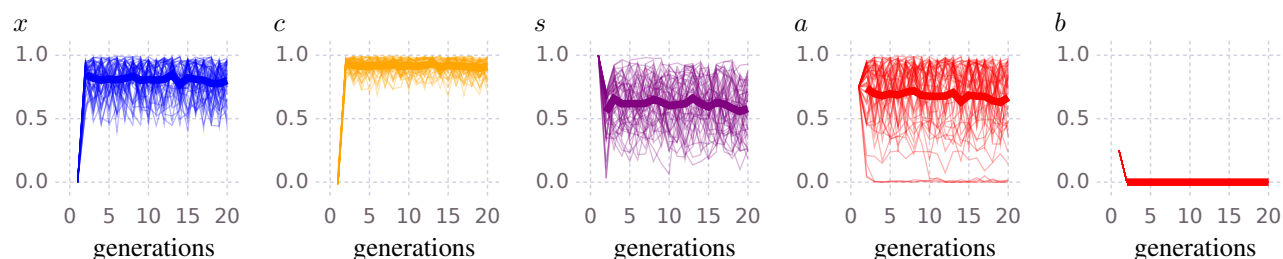


Figure 5: **Mixing larger languages.** Here the language is evolved for 20 generations and the results of 50 simulations are plotted for the usual five different language attributes:  $x$ ,  $c$ ,  $s$ ,  $a$  and  $b$ . The two rows give two different values of  $p$ . In each case the encoder and decoder maps have a 20-30-20 layout.

Filppula, M., Klemola, J., and Paulasto, H. (2008). *English and Celtic in Contact*. Routledge.

Kirby, S. (2001). Spontaneous evolution of linguistic structure—an iterated learning model of the emergence of regularity and irregularity. *IEEE Transactions on Evolutionary Computation*, 5(2):102–110.

Kirby, S. and Hurford, J. (2002). The emergence of linguistic structure: An overview of the iterated learning model. In Cangolosi, A. and Parisi, D., editors, *Simulating the Evolution of Language*, pages 121–148. Springer London.

Lodge, R. A. (2013). *French: From Dialect to Standard*. Routledge.

Matras, Y. and Sakel, J. (2007). *Grammatical Borrowing in Cross-Linguistic Perspective*, volume 38 of *Empirical Approaches to Language Typology*. De Gruyter Mouton.

Mufwene, S. S. (2001). *The Ecology of Language Evolution*. Cambridge University Press.

Oliphant, M. (1999). The learning barrier: Moving from innate to learned systems of communication. *Adaptive Behavior*, 7(3-4):371–383.

Oliphant, M. and Batali, J. (1997). Learning and the emergence of coordinated communication. *Center for Research on Language Newsletter*, 11(1):1–46.

Rickard, P. (2003). *A history of the French language*. Routledge.

Sains, G., Houghton, C., and Bullock, S. (2023). 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., Izquierdo, E., Suzuki, R., and Balteri, M., editors, *Ghost in the Machine: Proceedings of the 2023 Artificial Life Conference*, pages 377–385.

Sakel, J. and Matras, Y. (2008). *Modelling Contact-induced Change in Grammar*, volume 35 of *Empirical Approaches to Language Typology*. De Gruyter Mouton.

Smith, K., Perfors, A., Fehér, O., Samara, A., Swoboda, K., and Wonnacott, E. (2017). Language learning, language use and the evolution of linguistic variation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1711):20160051.

Thomason, S. G. and Kaufman, T. (1988). *Language Contact, Creolization, and Genetic Linguistics*. University of California Press.