Cognitively Inspired Developmental Trajectories Improve Explore-Exploit Dynamics in Neural Agent Emergent Communication

Jan Dziewoński, Flor Miriam Plaza del Arco and Tessa Verhoef
Leiden Institute of Advanced Computer Science
Leiden University
jandziewonski@google.com, f.m.plaza.del.arco@liacs.leidenuniv.nl
t.verhoef@liacs.leidenuniv.nl

Abstract

Emergent communication models support interaction-based language learning, benefiting both Natural Language Processing (NLP) applications and simulations of language evolution, but they are prone to destabilizing language drift. Inspired by developmental trajectories in human language acquisition, this paper investigates whether age-based plasticity, where younger agents learn quickly and older agents maintain stable representations, can reduce language drift. In our set-up, static populations first reliably develop shared languages, followed by a phase in which population turnover gradually replaces older agents with new learners. Age-based plasticity significantly reduces drift in this setting, maintaining high accuracy and language similarity. In contrast, in populations with uniformly low plasticity agents cannot adapt quickly enough to integrate newcomers and in those with uniformly high plasticity the language changes faster than stable conventions can form. These findings demonstrate that developmental trajectories in individual learners substantially reduce overall language drift in dynamic populations.

1 Introduction

Emergent communication models have gained prominence as a way to computationally capture the interactive and adaptive aspects of human language (Brandizzi, 2023). Within natural language processing (NLP), these models support the development of task-oriented languages grounded in communicative interaction rather than explicit supervision (Lazaridou et al., 2017; Foerster et al., 2016; Lazaridou and Baroni, 2020). In parallel, recent computational advances have sparked renewed interests in using such models to simulate processes of human language emergence and change (Chaabouni et al., 2021; Lian et al., 2023; Carlsson et al., 2024; Lian et al., 2025; Galke et al., 2022).

A central methodological tool in this line of research is the use of language games, which have a long tradition in the study of language evolution (Steels, 1997; Steels and McIntyre, 1998; De Boer, 2006; Beuls and Steels, 2013). A common scenario is the referential game, where a sender describes a target object and a receiver identifies it among distractors (Lazaridou and Baroni, 2020). Modern implementations of emergent communication rely on neural network–based agents, trained through reinforcement learning, where typically one agent always speaks and the other guesses the meaning. Learned protocols in this setting often diverge significantly from human language, due to the absence of crucial cognitive and communicative pressures (Galke et al., 2022). In response, researchers have started modeling multi-agent settings with group interaction (Tieleman et al., 2019; Graesser et al., 2019; Kim and Oh, 2021; Rita et al., 2022; Michel et al., 2023; Lian et al., 2024) and generational transmission (Chaabouni et al., 2019; Lian et al., 2021; Ren et al., 2020; Li and Bowling, 2019; Cogswell et al., 2020; Carlsson et al., 2024), leading to more realistic pressures from learning and interaction processes (Smith, 2022). As agents repeatedly interact with other individuals, their language naturally drifts: symbols shift in usage and grammatical patterns change. For simulating the evolution of language, some drift is welcome, and even needed to see the changes and potential emergence of patterns we intend to study , but too much of it can lead to an unnatural collapse of the language to a degenerate state where it becomes uninterpretable and optimized to a form that is far from human-like (Lazaridou and Baroni, 2020; Lazaridou et al., 2018; Lu et al., 2020). In addition, there are many other application areas where emergent communication set-ups can be useful (Boldt and Mortensen, 2024), such as synthetic data generation or multi-agent cooperation. In such cases drift could also be problematic as it decreases mutual intelligibility or causes the lan-

Proceedings of the 30th Conference on Computational Natural Language Learning, pages 104–120

guage to lose properties that made it useful in the first place. The problem perhaps most directly affects settings in which agents are first pretrained on human language and then finetuned through interaction to maximize task rewards (Lee et al., 2019; Lu et al., 2020; Lazaridou and Baroni, 2020; Lazaridou et al., 2020), as it may result in agent messages that humans can no longer understand. Especially in more dynamic populations with transmission to new learners, agents who entered at different times may struggle to preserve mutual intelligibility with humans and/or other agents.

Here, we propose cognitively inspired agent developmental trajectories as a solution to achieve language stability. The role of a developmental trajectory in human language learning is emphasized in the Critical Period hypothesis (Lenneberg, 1967), which proposes that humans can only acquire language fully and effortlessly during a constrained time window, which is usually assumed to end around puberty, or slightly later (Hartshorne et al., 2018). Evidence for this hypothesis comes from a combination of first language deprivation cases and studies on age-related decline in second language acquisition success (Curtiss et al., 1974; Oyama, 1976; Newport, 1988; Johnson and Newport, 1989). Gopnik (2020) proposes a theoretical framework that reframes the critical period as a solution to a well-known computational problem: the exploration-exploitation trade-off. In complex environments, there is no truly optimal way to balance exploration and exploitation, but one effective strategy is to begin with broad, high-temperature search that samples widely across the space of possibilities, then gradually cool to narrower, low-temperature search that converges on stable solutions (Kirkpatrick et al., 1983; Gopnik, 2020). Gopnik argues that human development demonstrates precisely this strategy. Children exhibit high plasticity, broad hypothesis search, and characteristics suited to exploration. Adults, by contrast, display more stable behavior, focused attention, and characteristics suited to exploitation. This division of labor allows a human to explore the space of possible languages, behaviors, and structures during a protected period of immaturity, then exploit that knowledge efficiently in adulthood (Gopnik, 2020). Likewise, in emergent communication settings older agents could provide stable learning targets that anchor the system against drift while younger agents could have the liberty to explore and learn the language quicker. Long before the deep learning era and the accompanying recent wave of emergent communication models, a simulation of vowel system emergence (Verhoef and de Boer, 2011), building on work by De Boer and Vogt (1999), demonstrated that age-structured populations preserved their vowel systems significantly better than either homogeneous high-plasticity (all children) or homogeneous low-plasticity (all adult) populations. Developmental trajectories in individual agents therefore limited drift and increased stability at the population level. Here we demonstrate that an analogous mechanism affects neural network agents playing referential games: introducing age-based plasticity into populations with turnover significantly increases language stability and cross-generational accuracy in a modern emergent communication model.

2 Related Work

2.1 Simulating Populations

Early neural emergent communication models typically simulated a single sender-receiver pair, even though it has been found that population characteristics affect human language structure, with larger communities developing more systematic and structured languages (Wray and Grace, 2007; Lupyan and Dale, 2010; Raviv et al., 2019). Thus, more recent set-ups have started to introduce multi-agent populations (Graesser et al., 2019; Tieleman et al., 2019; Kim and Oh, 2021; Michel et al., 2023; Rita et al., 2022; Lian et al., 2024). Scaling to larger populations did not immediately reproduce the effects observed in human experiments and real language though. Chaabouni et al. (2022) conducted large-scale experiments with populations of up to 100 agents on realistic image datasets and found no improvement in language generalization or robustness from population size alone. Michel et al. (2023) provided a theoretical explanation for the initial negative results. In standard training, receivers co-adapt to all senders they interact with, preventing them from developing one shared language and at equilibrium the population-level objective takes the same functional form regardless of population size. They proposed partitioning agents (connecting them into sender-receiver pairs) to prevent this co-adaptation, which improved population size effects on language structure and enabled successful communication with previously unseen partners. A similar solution with role-alternating agents was proposed in the NeLLCom-X framework (Lian et al., 2024), where agents first learn an artificial language and then communicate in groups.

2.2 Population Heterogeneity

Most group communication simulations employ homogeneous populations where all agents share identical architectures, learning rates, and optimization procedures. Rita et al. (2022) proposed that this homogeneity may also explain the initial null effect of group size. When they introduced asymmetries in learning speed across agents, larger populations began developing more stable and structured languages. Mahaut et al. (2025) studied populations of pre-trained visual networks with different architectures and training histories, finding that such heterogeneous agents could develop shared communication protocols through referential games. New agents learned existing community protocols more efficiently than communities could develop new protocols from scratch. Both lines of work suggest that agent diversity need not impede, and may facilitate, the emergence of shared conventions. Here, we introduce a developmental form of heterogeneity through age-based plasticity.

2.3 Transmission to New Learners

Real language populations are not static, but languages are constantly transmitted to new generations (Kirby et al., 2014; Smith, 2022). This process has been simulated in various emergent communication set-ups with single agent pairs (Chaabouni et al., 2019; Lian et al., 2021; Ren et al., 2020; Carlsson et al., 2024), but only a few have combined it with populations. Li and Bowling (2019) model one speaker communicating with multiple generations of listeners. This encourages the senders to form languages that are easier to teach to newcomers. Cogswell et al. (2020) introduced a framework that simulates generational transmission by resetting agents, turning them into newcomers with no knowledge of the existing language. This process improved compositional generalization compared to static populations. Prior work has mainly examined how population turnover affects compositionality and learnability, but less attention has been paid to mechanisms that might regulate drift and preserve cross-generational communication. Our setup is very similar to that of Cogswell et al. (2020) in the way agents are gradually replaced to form new generations, but here we use role-alternating agents and, importantly, agents are not static but develop in a cognitively inspired way when they age.

2.4 Mitigating Language Drift

A few other studies have addressed the problem of destabilizing language drift. Lee et al. (2019) studied a translation game where pre-trained agents communicate in an English pivot language. When fine-tuned for task completion, the intermediate English diverged from natural language. They mitigated drift by adding additional constraints: a language model rewarding grammatical plausibility, and a visual grounding model aligning captions with corresponding images. Combining both syntactic and semantic constraints best preserved language properties while also improving communication performance. Lu et al. (2020) explored the same scenario using Seeded Iterated Learning (SIL), which bears partial resemblance to our approach. Their method uses transmission bottlenecks to preserve structure: a teacher agent fine-tunes for task completion, then generates data for a student agent to imitate. The imitation step filters out irregular patterns that emerged during optimization. This alternation between teacher fine-tuning and student imitation resembles the exploration-exploitation dynamics central to our work. However, SIL implements this through temporal phases within one agent, while we distribute these roles across a population through age-based heterogeneity.

Both Lee et al. (2019) and Lu et al. (2020) prevent divergence from a pre-trained human language in two-agent settings. Our work addresses a different problem: drift within multi-agent populations where a language emerges from scratch, but is not preserved throughout generations. We aim to improve stability in this setting by introducing developmental trajectories.

3 Methods

3.1 Game Design

We use a referential game based on the `objects_game` in the existing EGG framework (Kharitonov et al., 2019; Lazaridou et al., 2018). Each round, a sender observes a target object and transmits a message to a receiver, who must identify the target among 5 distractors. Objects are discrete feature vectors with 4 attributes, each taking 5 values, yielding 625 possible objects. Messages are variable-length sequences with vocabulary size 20, including end-of-sequence (EOS) and a maximum length of 10. The sender generates messages autoregressively using Gumbel-Softmax relaxation (Jang et al., 2017) for differentiable discrete communication; the receiver scores candidates via dot product attention. Both minimize cross-entropy loss. Implementation details on the elements we borrowed from EGG are listed in Appendix A, while we focus below on our extensions.

The single-pair game extends to populations following prior work (Rita et al., 2022; Michel et al., 2023; Chaabouni et al., 2022). Each batch randomly pairs N agents as senders and receivers. Self-play (same agent is sampled for both roles) is permitted. The training loss averages across pairs. Evaluation tests all N^2 sender-receiver combinations to measure population-wide accuracy.

3.2 Agent Architecture

Each agent can act as both sender and receiver, inspired by the role-alternation in NeLLCom-X (Lian et al., 2024). An agent consists of two modules: a sender module that encodes target objects and generates messages via LSTM, and a receiver module that encodes candidates and scores them against the received message. The modules use separate parameters (hidden dimension 128, embedding dimension 64). See also Appendix A.2.

Each agent has its own Adam optimizer, enabling per-agent learning rate control necessary for age-based plasticity. Only agents sampled in a training batch receive gradient updates. Separate optimizers allow for properly modeling agent turnover - newcomers should not be influenced by a shared optimizer. When a new agent enters the population, its optimizer initializes with a fresh state.

3.3 Population Dynamics

To study language transmission across generations, we implement agent turnover following Cogswell et al. (2020): agents are periodically replaced by newcomers who must learn the existing language through interaction. Each agent has an *age* measured in epochs since birth. The parameter k controls the turnover rate: every k epochs, the oldest agent is replaced. There is an initial warmup phase (until the population reaches an accuracy threshold) before the turnover begins.

We pre-allocate slots for all agents that will ever exist. At all times only N agents are *alive* and only they participate in training. Initially, all other agents start as *dormant*. When the oldest (random in case of a tie) living agent leaves, it is *replaced*.

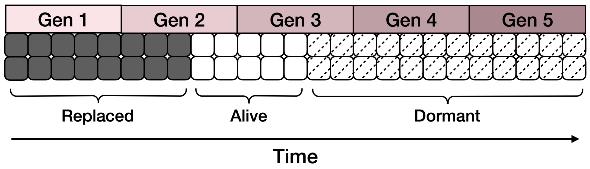

Figure 1: Example of Agent States and generational overlap. Here, four of the alive agents are from generation 2, and the other six are from generation 3.

It is excluded from training and its parameters are frozen at their values at replacement. At the same time, the next *dormant* agent activates with fresh, random initialization. This design makes cross-generational comparisons possible, as the *replaced* agents can be "resurrected" and be paired with their successors. Figure 1 illustrates these agent states and how they are placed in generations.

Generations group agents by birth order. Generation 1 contains the founding population (agents 0 through $N-1$) and each subsequent N deaths marks a new generation. For the experiments reported below we use $N = 10$ and $k = 10$. This yields agent lifespans of at most 100 epochs before replacement. Preliminary experiments confirmed that populations of 2 - 20 agents reliably develop shared languages, and that turnover causes cumulative language drift regardless of rate and population size (see Appendix B).

3.4 Age-Based Plasticity

We model developmental trajectories by decreasing agents' plasticity as they age. Young agents learn quickly and explore broadly, old agents learn slowly and do less exploration. A plasticity function $p(a) \in [0, 1]$ maps age to learning capacity using a sigmoid:

$$p(a) = \frac{1}{1 + \exp(\beta(\bar{a} - \mu))} \tag{1}$$

where $\bar{a} = a/(a_{\max} - 1)$ normalizes age to $[0, 1]$, β controls transition steepness, and μ sets the midpoint. Plasticity modulates two training parameters:

First, Gumbel-Softmax temperature τ controls how discrete the samples are. Notably, low values of τ characterize more discrete, sharp categories but also more difficult learning, which is more adult-like, while high values of τ result in more continuous representations, softer categories and easier learning, which is more child-like.

$$\tau(a) = \tau_{\min} + (\tau_{\max} - \tau_{\min}) \cdot p(a) \tag{2}$$

Changing τ from high to low in the development of an agent can be compared to the ability of infants to discriminate a wide range of speech sounds after birth, while losing this ability for many sounds later on as they become specialized in sound categories of their native language only (Werker and Tees, 1984). Similarly, the often observed phenomenon of overextension in object naming (Clark, 1978; Rescorla, 1980) indicates the fluidity of children's categorical understanding and active experimentation in word production. As such, decreasing τ implements the general explore-exploit developmental shift in humans (Gopnik, 2020).

Second, Learning Rate η controls adaptation speed, with adults being able to adapt more slowly than children, following Verhoef and de Boer (2011):

$$\eta(a) = \eta_{\min} + (\eta_{\max} - \eta_{\min}) \cdot p(a) \quad (3)$$

So, young agents (high p) produce diverse messages and adapt quickly; old agents (low p) produce consistent messages and resist change. At any given time, the population spans this spectrum, creating heterogeneity. Parameters update after each epoch based on current age.

3.5 Metrics and Evaluation

We evaluate populations using two main complementary measures: **accuracy** and **language similarity**. For language similarity, we follow Michel et al. (2023) and Rita et al. (2022) and compute pairwise normalized edit distances between messages from different agents averaged across all pairs. Values near 1 indicate similar messages for the same inputs. We first evaluate these within a single population snapshot (see Appendix D.4).

To quantify drift, we extend both metrics across time using reference points. **Cross-generational** metrics compare each generation against the founding population (Gen 1, captured after warmup). Cross-generational accuracy tests whether agents from different generations (that may have never coexisted) can still communicate (in both directions):

$$A_{\text{cross-gen}} = \frac{1}{2}\left(A_{G_1 \to G_t} + A_{G_t \to G_1}\right) \quad (4)$$

Similarly, cross-generational similarity measures whether they produce similar messages for the same inputs. Declining values in either indicate cumulative drift from the original language. We also measure **previous-generation** metrics that use the previous generation (Gen $g-1$) as the reference point to capture drift at each transition.

3.6 Experimental Conditions

To test whether developmental trajectories reduce language drift, we compare four plasticity conditions (Table 1). The baseline uses fixed moderate parameters without age modulation. The adult-only condition uses uniformly low plasticity, modeling a population where all agents learn slowly and resist change. The child-only condition uses uniformly high plasticity, modeling a population where all agents learn quickly and explore frequently. All conditions load from a shared pre-trained checkpoint (using baseline settings) where the founding population has converged to 90% accuracy. In the age-based condition, Gen 1 agents subsequently begin as adults since they have already established the language. Starting from a converged checkpoint provides a fixed reference language against which to measure drift in a comparable way across conditions. Appendix C shows that age-heterogeneous populations can, however, also reliably develop shared languages from random initialization. New agents enter with high plasticity and mature over their lifespan. Training runs for 400 epochs and spans 5 generations. Each configuration ran with 3 random seeds and the results below report mean $\pm$ standard deviation. To isolate the two plasticity components, we run two additional ablation conditions. For temperature-only, we lock η at the baseline value, while age-coupling τ. For learning-rate-only, we do the reverse. All other parameters match the baseline condition.

Condition	Temperature τ	Learning rate η
Baseline	1.0	3×10^{-3}
Adults	0.5	10^{-3}
Children	5.0	2.5×10^{-2}
Age-based	0.5–5.0	10^{-3}–2.5×10^{-2}

Table 1: Plasticity conditions. All use $N = 10$ and $k = 10$.

4 Results

4.1 Within-generation accuracy and similarity

Figure 2 shows test accuracy over training epochs. Firstly, in all conditions, we can observe a cycle caused by agent replacement, where accuracy drops as a newcomer enters and recovers when the newcomer learns the language. The baseline, adult-only, and age-based populations all maintain relatively similar accuracy levels, with adults slightly behind the other two. Their accuracies are in the range of 70-90% depending on the phase of the cycle. The child-only condition however, already starts from very low initial accuracy of around 25%. This suggests that the population is so chaotic that it immediately loses the language formed during the warmup. The population regains some of the accuracy but plateaus in the 60-70% range and does not show significant improvement past epoch 150.

Figure 3 plots language similarity and shows that all conditions maintain reasonably synchronized languages over training epochs, though they all show high variance across seeds. Adults are noticeably worse than other populations. Notably, the drop in accuracy for the child-only population does not occur in the language similarity curve.

4.2 Cross-generational accuracy and similarity

Cross-generational metrics reveal substantial differences between conditions. Figure 4 shows cross-generational accuracy and similarity measured against the founding generation (Gen 1). First, all curves decrease over time as expected, demonstrating that more distant generations have a harder time communicating successfully. The decrease is more significant at first, and diminishes over time, as languages adapt to be transmitted more easily. Crucially, the age-based condition maintains significantly higher cross-generational accuracy than all other conditions. Cross-generational similarity shows a similar pattern. The child-only population experiences a substantial drop in both metrics almost immediately after turnover is activated. The adult-only condition's curve resembles the shape of the baseline curve, only more pronounced, with

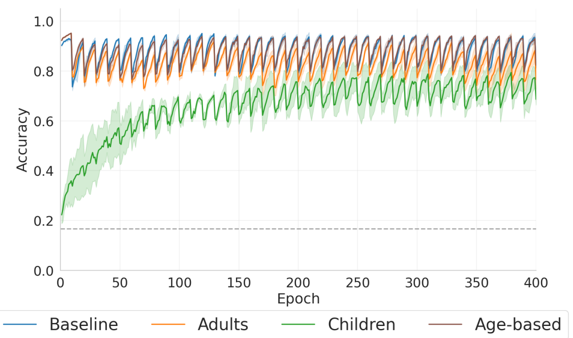

Figure 2: Test accuracy over epochs. The child-only condition drops right after the warmup and stays low. Lines show means across 3 seeds; shaded regions indicate ± 1 standard deviation. The dashed line in (a) indicates random baseline accuracy (16.7%).

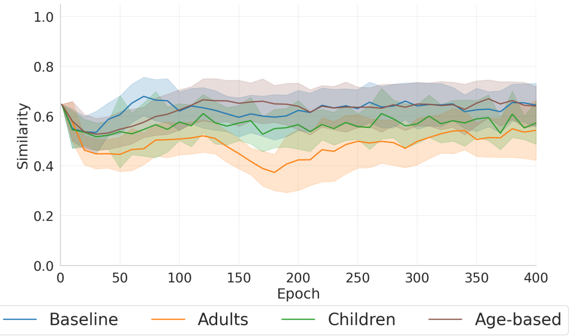

Figure 3: Language similarity over epochs. All conditions maintain relatively synchronized languages.

lower values (higher drift) the majority of the time. The exception is the very beginning of turnover: at first, adults preserve more of the original language (likely due to their superior stability), but this effect is quickly lost, as more newcomers are introduced who are not able to learn as effectively.

To capture the rate of language change between consecutive generations, Figure 5 shows cross-generational accuracy and similarity measured against the previous generation rather than Gen 1. The metrics confirm the trend visible in Figure 4: more drift in the beginning, less drift further in the simulation. The age-based condition maintains 80-90% accuracy with its previous generations at all times, indicating low drift. These metrics reveal that the age-based population preserves the language throughout generations not just cumulatively but at each generational transition.

Table 2 summarizes all metrics after 5 generations. The cross-generational metrics compare the final generation to the founding generation, while the previous-generation metrics are averaged across all generational transitions (generations 2 through 5) to capture the overall drift rate between consecutive generations.

The comparison to baseline reveals the central finding of this experiment. Both homogeneous conditions perform worse than baseline on cross-generational metrics: adults show 22% lower cross-generational accuracy and 49% lower similarity, while children show 28% lower similarity. However, the age-based condition—which combines the properties of both—substantially outperforms the baseline: 109% higher cross-generational accuracy (0.606 vs 0.290) and 131% higher cross-generational similarity (0.309 vs 0.134). The average previous-generation accuracy tells a similar story: children (0.608) and adults (0.657) both fall

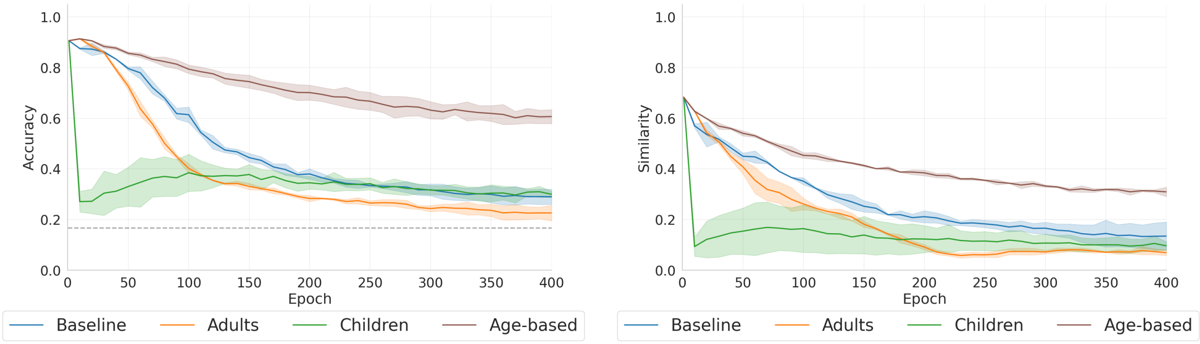

(a) Cross-generational accuracy by epoch. The age-based condition maintains substantially higher accuracy over generations, suggesting a better preservation of the language over time.

(b) Cross-generational similarity by epoch. The age-based condition preserves message similarity far better than uniform plasticity conditions.

Figure 4: Language drift across plasticity conditions, measured against the original generation (Gen 1).

Metric	Baseline	Adults	Children	Age-based
Test Accuracy (%)	82.2 ± 0.8	76.6 ± 3.0	68.6 ± 1.8	$\mathbf{82.3 \pm 0.9}$
Language Similarity	0.705 ± 0.02	0.601 ± 0.10	0.615 ± 0.08	$\mathbf{0.709 \pm 0.03}$
Cross-Gen Accuracy	0.290 ± 0.02	0.226 ± 0.03	0.299 ± 0.01	$\mathbf{0.606 \pm 0.02}$
Cross-Gen Similarity	0.134 ± 0.05	0.069 ± 0.01	0.096 ± 0.01	$\mathbf{0.309 \pm 0.02}$
Avg Prev Accuracy	0.789 ± 0.12	0.657 ± 0.16	0.608 ± 0.15	$\mathbf{0.883 \pm 0.05}$
Avg Prev Similarity	0.479 ± 0.11	0.348 ± 0.12	0.396 ± 0.18	$\mathbf{0.597 \pm 0.09}$

Table 2: Metrics after 5 generations (mean $\pm$ std across 3 seeds). Cross-Gen metrics compare Gen 5 to Gen 1. Prev metrics are averaged across all consecutive generational transitions (Gen g to $g-1$).

below baseline (0.789), but age-based (0.883) exceeds it by 12%.

This pattern supports the hypothesis that age-based heterogeneity matters. The adult-only condition fails because agents cannot adapt quickly enough: newcomers struggle to learn from rigid teachers, and the population fragments. The child-only condition fails because agents change too rapidly: the language drifts faster than any stable convention can form. Neither extreme works well in isolation. But when combined in the age-based condition, each compensates for the other's weakness. Older agents with low plasticity provide stable, consistent learning targets. Younger agents with high plasticity can efficiently acquire this stable language before their own plasticity decreases. The result is a population that maintains both high within-generation accuracy and low cross-generational drift.

4.3 Decomposing Plasticity: η versus τ

To test whether the cross-generational stability of the age-based condition can be attributed to either temperature or learning rate alone, we compare it to conditions that lock one of the parameters at the baseline value. Neither single parameter condition achieves the cross-gen accuracy gain of full plasticity (Figure 6a). The learning-rate-only condition is destabilizing: cross-generational similarity drops below baseline after some generations (Figure 6b). Temperature-only performs only slightly better than baseline on cross-generational accuracy. We interpret these results as follows. Modulating only η with age gives young agents fast adaptation, but a sharp message distribution, which makes them commit to whatever they sample in the beginning, potentially without exploring enough to align fully with the messages used in the population. Modulating only τ gives young agents broad exploration, but with everyone learning at the same rate, the only added effect is noisier message sampling from young agents, who update too slowly for high temperature to change what the population actually learns, while adults do not provide stable targets. Modulating both parameters lets newcomers explore broadly enough to find the existing language quickly, and adapt fast enough to fully (or mostly) commit to it, while learning from a stable input.

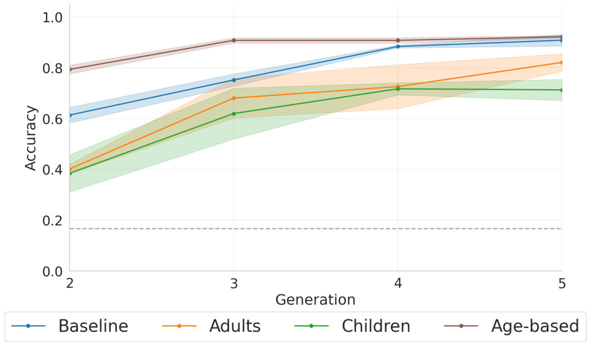

(a) Cross-generational accuracy by generation. The age-based condition maintains over 80% accuracy with the previous generation at all times.

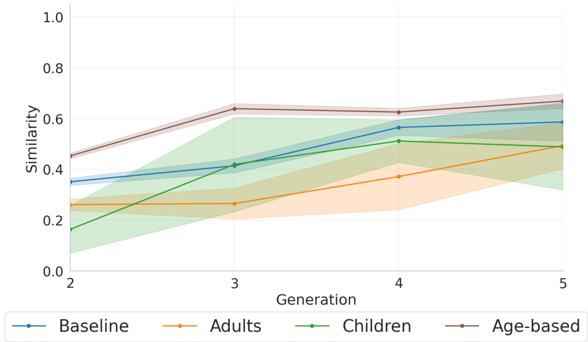

(b) Cross-generational similarity by generation. The age-based condition shows the highest and most consistent similarity with its predecessor.

Figure 5: Language drift across plasticity conditions, measured against the previous generation (Gen $g-1$).

5 Discussion

We investigated whether age-based plasticity improves language stability in dynamic agent populations. Our experiments strongly support this hypothesis. The developmental trajectory in individual agents substantially reduces language drift compared to homogeneous populations, with the heterogeneous (age-based) population achieving more than double the cross-generational accuracy.

Turnover creates two distinct challenges for language stability. First, newcomers disrupt communication immediately upon entering the population, with the severity scaling with turnover frequency: populations with frequent replacement never fully recover between events, while populations with infrequent replacement maintain high accuracy. Second, the language itself drifts over time regardless of turnover rate. Cross-generational accuracy and similarity both decline as generations accumulate, confirming that the founding language gradually gives way to successor variants.

Our results reveal that neither extreme of uniform plasticity works well in populations with agent replacement. The adult-only condition (slow learning, low exploration) fails because agents cannot adapt quickly enough. Newcomers are too slow to learn the language properly before being replaced and the language gradually decays. The child-only condition fails for the opposite reason: there are no stable teachers to learn from and even though children can learn quickly, the language drifts faster than stable conventions can form. Additionally, there is an immediate collapse in accuracy right after the warmup, indicating that children alone are simply not able to maintain the developed language. Both homogeneous conditions underperform when compared to the baseline on cross-generational metrics, establishing that the problem is not simply having too little or too much plasticity, but having the wrong distribution of it.

The age-based condition combines both extremes and substantially outperforms all alternatives. Cross-generational accuracy reaches 0.606, more than double the baseline (0.290). The previous-generation metrics confirm this is not merely a coincidence: age-based populations maintain 80-90% accuracy with their immediate predecessors throughout training, higher than homogeneous conditions. This pattern supports the hypothesis that the age-based structure drives the stabilizing effect. Older agents with low plasticity provide stable, consistent learning targets. Younger agents with high plasticity efficiently acquire this stable language before their own plasticity decreases. Neither component works in isolation, but together they create a population that balances adaptability with stability. When decomposing our plasticity function and testing the effect of η versus τ separately, we find that their combination is what robustly causes this benefit.

These findings provide important insights for the development of emergent communication set-ups. They demonstrate that population heterogeneity can substantially affect language dynamics, extending prior work by Rita et al. (2022) on learning rate asymmetries. In our case, heterogeneity is age-dependent. The failure of both homogeneous extremes, combined with the success of the heterogeneous condition, suggests that the interaction between stable and plastic agents is more important than the absolute level of either property, and can provide an effective solution against destabilizing

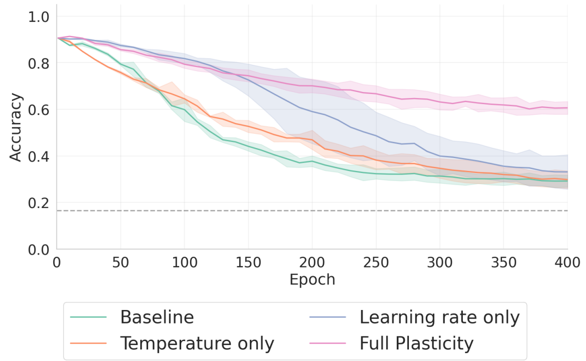

(a) Cross-generational accuracy by epoch. The learning-rate-only and temperature-only variants do not preserve accuracy as effectively as the full-plasticity condition.

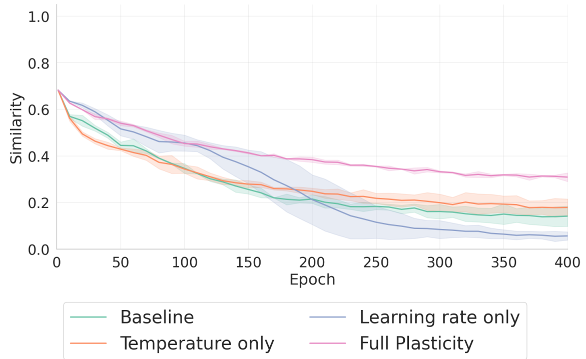

(b) Cross-generational similarity by epoch. Full-plasticity variant is best at preserving the similarity of the language. Learning-rate-only falls below baseline.

Figure 6: Language drift across different variants of plasticity measured against Gen 1. The learning-rate-only and temperature-only variants do not preserve language similarity as effectively as the full-plasticity condition.

language drift.

For theories of language evolution, the results provide computational evidence that human development contributes to language stability and cross-generational intelligibility in populations. The findings strengthen Gopnik's case that in human communities children's high plasticity and adults' stability may together be beneficial at the population level (Gopnik, 2020). They also extend the work by De Boer and Vogt (1999) and Verhoef and de Boer (2011) to current emergent communication settings with deep neural network agents and find similar patterns.

6 Conclusion

This paper provides computational evidence that age-based plasticity can substantially reduce language drift in dynamic populations. Heterogeneous populations, combining stable adults and highly plastic children, outperform homogeneous populations in maintaining cross-generational accuracy. This suggests that age-based plasticity could play a meaningful role in preserving intelligibility in human communities, supporting theories that children's high plasticity and adults' stability together benefit population-level language maintenance.

The implementation of plasticity as temperature and learning rate modulation provides only one possible approximation of biological processes. Human age-related changes in language learning have been connected to various different mechanisms including neural maturation or memory limitations. According to the "less is more" hypothesis (Newport, 1990), for example, children's limited cognitive resources, such as a smaller working memory capacity, actually help language learning. Mita et al. (2025) recently implemented this idea in a language model and showed that it improved grammatical learning efficiency in an individual learner, which creates a promising basis for exploring population-level effects.

The current framework investigates language drift mostly as degradation, but language change also serves an adaptive function. Human languages evolve to accommodate new concepts, technologies, and social structures. Future work could test whether age-based populations adapt better to changing environments by periodically introducing new input combinations that require learning new conventions. The question then becomes whether heterogeneous populations will be better at simultaneously maintaining high accuracy among currently alive agents and efficiently inventing new words for novel concepts. Our findings therefore open up novel research avenues connecting emergent communication research to theories of human language evolution and suggest that population structure deserves attention alongside the properties of individual agents.

7 Limitations

Several limitations constrain the interpretation of these results. The experimental design creates a puzzle between turnover rate and total training time. Varying the kill epoch parameter k simultaneously changes how much time newcomers have to learn and how many optimization steps occur per generation. This makes it difficult to isolate whether observed differences stem purely from learning dynamics or simply the number of interactions between agents, when these settings vary. In our main experiment we kept turnover rate and population size constant, but future work could aim to better decouple these factors through alternative experimental designs.

The current experiment tests one configuration of plasticity parameters. Systematic exploration of the temperature and learning rate ranges could identify optimal settings and clarify how extreme the difference between young and old agents needs to be for the stabilizing effect to emerge. The plasticity trajectory (controlled by the sigmoid's steepness and critical point) also calls for investigation: does a gradual decline work better than a sharp transition?

The metrics used to measure drift have inherent limitations. Cross-generational accuracy conflates true language preservation with the possibility that task constraints push independently evolved languages toward similar solutions. Cross-generational similarity measures surface form but not semantic content. The parallel decline of both metrics suggests genuine drift, but this interpretation rests on assumptions that cannot be fully verified.

The scale and complexity of these experiments differ substantially from human language communities. A population size of 10 agents is far smaller than human speech communities. The task, with 625 possible inputs and 6 candidate objects, may be too simple to require all the complex linguistic structures that age-based plasticity might help preserve. More complex tasks should be explored and could reveal larger effects or different dynamics entirely.

Acknowledgments

We thank the Leiden Institute of Advanced Computer Science for providing the computational resources that made this research possible, and the REL Compute team for maintaining the infrastructure on which all experiments were conducted.

References

Katrien Beuls and Luc Steels. 2013. Agent-based models of strategies for the emergence and evolution of grammatical agreement. *PloS one*, 8(3):e58960.

Brendon Boldt and David Mortensen. 2024. A review of the applications of deep learning-based emergent communication. *Preprint*, arXiv:2407.03302.

Nicolo' Brandizzi. 2023. Toward more human-like ai communication: A review of emergent communication research. *IEEE Access*, 11:142317–142340.

Emil Carlsson, Devdatt Dubhashi, and Terry Regier. 2024. Cultural evolution via iterated learning and communication explains efficient color naming systems. *Journal of Language Evolution*, 9(1-2):49–66.

Rahma Chaabouni, Eugene Kharitonov, Emmanuel Dupoux, and Marco Baroni. 2021. Communicating artificial neural networks develop efficient color-naming systems. *Proceedings of the National Academy of Sciences*, 118(12):e2016569118.

Rahma Chaabouni, Eugene Kharitonov, Alessandro Lazaric, Emmanuel Dupoux, and Marco Baroni. 2019. Word-order biases in deep-agent emergent communication. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5166–5175, Florence, Italy. Association for Computational Linguistics.

Rahma Chaabouni, Florian Strub, Florent Altché, Eugene Tarassov, Corentin Tallec, Elnaz Davoodi, Kory Wallace Mathewson, Olivier Tieleman, Angeliki Lazaridou, and Bilal Piot. 2022. Emergent communication at scale. In *International Conference on Learning Representations*.

Eve V Clark. 1978. Strategies for communicating. *Child Development*, pages 953–959.

Michael Cogswell, Jiasen Lu, Stefan Lee, Devi Parikh, and Dhruv Batra. 2020. Emergence of compositional language with deep generational transmission. In *International Conference on Learning Representations*.

Susan Curtiss, Victoria Fromkin, Stephen Krashen, David Rigler, and Marilyn Rigler. 1974. The linguistic development of genie. *Language*, pages 528–554.

Bart De Boer. 2006. Computer modelling as a tool for understanding language evolution. In *Evolutionary epistemology, language and culture: A non-adaptationist, systems theoretical approach*, pages 381–406. Springer.

Bart De Boer and Paul Vogt. 1999. Emergence of speech sounds in changing populations. In *European Conference on Artificial Life*, pages 664–673. Springer.

Jakob Foerster, Ioannis Alexandros Assael, Nando De Freitas, and Shimon Whiteson. 2016. Learning to communicate with deep multi-agent reinforcement learning. In *Advances in neural information processing systems*, volume 29.

Lukas Galke, Yoav Ram, and Limor Raviv. 2022. Emergent communication for understanding human language evolution: What's missing? In *Emergent Communication Workshop at ICLR 2022*.

Alison Gopnik. 2020. Childhood as a solution to explore–exploit tensions. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 375(1803):20190502.

Laura Harding Graesser, Kyunghyun Cho, and Douwe Kiela. 2019. Emergent linguistic phenomena in multi-agent communication games. In *Proceedings of the 2019 Conference on Empirical Eethods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3700–3710.

Joshua K. Hartshorne, Joshua B. Tenenbaum, and Steven Pinker. 2018. A critical period for second language acquisition: Evidence from 2/3 million english speakers. *Cognition*, 177:263–277.

Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical reparameterization with gumbel-softmax. In *International Conference on Learning Representations*.

Jacqueline S Johnson and Elissa L Newport. 1989. Critical period effects in second language learning: The influence of maturational state on the acquisition of english as a second language. *Cognitive psychology*, 21(1):60–99.

Eugene Kharitonov, Rahma Chaabouni, Diane Bouchacourt, and Marco Baroni. 2019. EGG: a toolkit for research on emergence of lanGuage in games. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP): System Demonstrations*, pages 55–60, Hong Kong, China. Association for Computational Linguistics.

Jooyeon Kim and Alice Oh. 2021. Emergent communication under varying sizes and connectivities. *Advances in Neural Information Processing Systems*, 34:17579–17591.

Simon Kirby, Tom Griffiths, and Kenny Smith. 2014. Iterated learning and the evolution of language. *Current Opinion in Neurobiology*, 28:108–114. SI: Communication and language.

Scott Kirkpatrick, C Daniel Gelatt Jr, and Mario P Vecchi. 1983. Optimization by simulated annealing. *science*, 220(4598):671–680.

Angeliki Lazaridou and Marco Baroni. 2020. Emergent multi-agent communication in the deep learning era. *Preprint*, arXiv:2006.02419.

Angeliki Lazaridou, Karl Moritz Hermann, Karl Tuyls, and Stephen Clark. 2018. Emergence of linguistic communication from referential games with symbolic and pixel input. In *International Conference on Learning Representations*.

Angeliki Lazaridou, Alexander Peysakhovich, and Marco Baroni. 2017. Multi-agent cooperation and the emergence of (natural) language. In *International Conference on Learning Representations*.

Angeliki Lazaridou, Anna Potapenko, and Olivier Tieleman. 2020. Multi-agent communication meets natural language: Synergies between functional and structural language learning. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7663–7674. Association for Computational Linguistics.

Jason Lee, Kyunghyun Cho, and Douwe Kiela. 2019. Countering language drift via visual grounding. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4385–4395. Association for Computational Linguistics.

Eric H Lenneberg. 1967. *The biological foundations of language*. New York : Wiley.

Fushan Li and Michael Bowling. 2019. Ease-of-teaching and language structure from emergent communication. *Advances in neural information processing systems*, 32.

Yuchen Lian, Arianna Bisazza, and Tessa Verhoef. 2021. The effect of efficient messaging and input variability on neural-agent iterated language learning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10121–10129. Association for Computational Linguistics.

Yuchen Lian, Arianna Bisazza, and Tessa Verhoef. 2023. Communication Drives the Emergence of Language Universals in Neural Agents: Evidence from the Word-order/Case-marking Trade-off. *Transactions of the Association for Computational Linguistics*, 11:1033–1047.

Yuchen Lian, Arianna Bisazza, and Tessa Verhoef. 2025. Simulating the emergence of differential case marking with communicating neural-network agents. In *Proceedings of the 47th Annual Conference of the Cognitive Science Society (CogSci)*, pages 508–515. Cognitive Science Society.

Yuchen Lian, Tessa Verhoef, and Arianna Bisazza. 2024. NeLLCom-X: A comprehensive neural-agent framework to simulate language learning and group communication. In *Proceedings of the 28th Conference on Computational Natural Language Learning*, pages 243–258, Miami, FL, USA. Association for Computational Linguistics.

Yuchen Lu, Soumye Singhal, Florian Strub, Aaron Courville, and Olivier Pietquin. 2020. Countering language drift with seeded iterated learning. In *International Conference on Machine Learning*, pages 6437–6447. PMLR.

Gary Lupyan and Rick Dale. 2010. Language structure is partly determined by social structure. *PloS one*, 5(1):e8559.

Matéo Mahaut, Roberto Dessì, Francesca Franzon, and Marco Baroni. 2025. Referential communication in heterogeneous communities of pre-trained visual deep networks. *Transactions on Machine Learning Research*.

Paul Michel, Mathieu Rita, Kory Wallace Mathewson, Olivier Tieleman, and Angeliki Lazaridou. 2023. Revisiting populations in multi-agent communication. In *The Eleventh International Conference on Learning Representations*.

Masato Mita, Ryo Yoshida, and Yohei Oseki. 2025. Developmentally-plausible working memory shapes a critical period for language acquisition. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics*, pages 9386–9399.

Elissa L Newport. 1988. Constraints on learning and their role in language acquisition: Studies of the acquisition of american sign language. *Language sciences*, 10(1):147–172.

Elissa L Newport. 1990. Maturational constraints on language learning. *Cognitive science*, 14(1):11–28.

Susan Oyama. 1976. A sensitive period for the acquisition of a nonnative phonological system. *Journal of psycholinguistic research*, 5(3):261–283.

Limor Raviv, Antje Meyer, and Shiri Lev-Ari. 2019. Larger communities create more systematic languages. *Proceedings of the Royal Society B*, 286(1907):20191262.

Yi Ren, Shangmin Guo, Matthieu Labeau, Shay B Cohen, and Simon Kirby. 2020. Compositional languages emerge in a neural iterated learning model. In *8th International Conference on Learning Representations*.

Leslie A Rescorla. 1980. Overextension in early language development. *Journal of child language*, 7(2):321–335.

Mathieu Rita, Florian Strub, Jean-Bastien Grill, Olivier Pietquin, and Emmanuel Dupoux. 2022. On the role of population heterogeneity in emergent communication. In *International Conference on Learning Representations*.

Kenny Smith. 2022. How language learning and language use create linguistic structure. *Current Directions in Psychological Science*, 31(2):177–186.

Luc Steels. 1997. The synthetic modeling of language origins. *Evolution of communication*, 1(1):1–34.

Luc Steels and Angus McIntyre. 1998. Spatially distributed naming games. *Advances in complex systems*, 1(04):301–323.

Olivier Tieleman, Angeliki Lazaridou, Shibl Mourad, Charles Blundell, and Doina Precup. 2019. Shaping representations through communication: community size effect in artificial learning systems. *arXiv preprint arXiv:1912.06208*.

Tessa Verhoef and Bart de Boer. 2011. Language acquisition age effects and their role in the preservation and change of communication systems. *Linguistics in Amsterdam*, 4:1–23.

Janet F Werker and Richard C Tees. 1984. Cross-language speech perception: Evidence for perceptual reorganization during the first year of life. *Infant behavior and development*, 7(1):49–63.

Alison Wray and George W Grace. 2007. The consequences of talking to strangers: Evolutionary corollaries of socio-cultural influences on linguistic form. *Lingua*, 117(3):543–578.

A Implementation details

A.1 Game dynamics

Input Space Objects are discrete feature vectors $x = (x_1, \ldots, x_F)$ with F attributes. Attribute f takes values in $\{1, \ldots, d_f\}$, and the dimension vector $D = (d_1, \ldots, d_F)$ specifies the number of values per attribute. The input space $\mathcal{X}$ contains all valid attribute combinations.

Game Round Each round proceeds as follows: (1) sample a target $x^* \in \mathcal{X}$ uniformly; (2) sample n_d distractors from $\mathcal{X} \setminus \{x^*\}$ without replacement to form candidate set C; (3) shuffle C to randomize target position; (4) the sender observes x^* and produces message m; (5) the receiver observes m and C, then outputs a probability distribution over candidates. Communication succeeds when the receiver assigns highest probability to the target.

Communication Channel Messages are variable-length sequences from vocabulary $V = \{0, 1, \ldots, |V| - 1\}$, where symbol 0 denotes end-of-sequence (EOS). With maximum length L:

$$m = (m_1, \ldots, m_T), \quad m_t \in V, \quad T \leq L \tag{5}$$

The sender generates symbols autoregressively via an LSTM, sampling from policy $\pi_\theta(m_t \mid x^*, m_{<t})$. Gumbel-Softmax relaxation (Jang et al., 2017) with temperature τ enables gradient-based training with discrete symbols.

Receiver Discrimination The receiver processes the message through an LSTM (via EGG's `RnnReceiverGS`) to obtain embedding h_m. For each candidate $c_j \in C$, it computes score $s_j = h_m^\top f_\phi(c_j)$, where f_ϕ is the receiver's learned encoding function. Scores convert to probabilities via softmax:

$$\rho_\phi(c_j \mid m, C) = \frac{\exp(s_j)}{\sum_{c \in C} \exp(s_c)} \tag{6}$$

Training Objective The receiver minimizes cross-entropy loss:

$$\mathcal{L} = -\log \rho_\phi(x^* \mid m, C) \tag{7}$$

Gumbel-Softmax relaxation allows gradients to flow through the entire sender-receiver pipeline, enabling joint optimization. Accuracy measures the proportion of rounds where the receiver's argmax matches the target:

$$\text{Acc} = \frac{1}{|\mathcal{B}|} \sum_{(x^*,C)\in\mathcal{B}} \mathbf{1}\left[\arg\max_{c\in C} \rho_\phi(c \mid m, C) = x^*\right] \tag{8}$$

where $\mathcal{B}$ denotes the evaluation batch.

A.2 Agent architecture

Role alternation Agent i consists of two modules: a sender module parameterized by θ_i and a receiver module parameterized by ϕ_i. Both modules encode feature vectors from the same input space $\mathcal{X}$, but serve different roles during communication. When agent i acts as sender, it uses θ_i to encode the target and initiate message generation. When acting as receiver, it uses ϕ_i to encode candidates and compute discrimination scores. The modules do not share parameters—each agent maintains independent sender and receiver capabilities.

Sender Module The sender module encodes the target object into a hidden representation:

$$h_s = \tanh(W_s x^* + b_s) \tag{9}$$

where $W_s \in \mathbb{R}^{H\times F}$ is a weight matrix, $b_s \in \mathbb{R}^H$ is a bias vector, and H is the hidden dimension. This encoding initializes EGG's `RnnSenderGS` wrapper, which generates messages autoregressively via LSTM. At each step t, the LSTM produces a hidden state from which logits over the vocabulary are computed and a symbol is sampled (using Gumbel-Softmax during training, greedy decoding during evaluation). The message generation process is described in Section 3.1.

Receiver Module The receiver module encodes each candidate object into the same hidden space:

$$h_r^{(j)} = \tanh(W_r c_j + b_r) \tag{10}$$

where c_j is the j-th candidate and $W_r \in \mathbb{R}^{H\times F}$, $b_r \in \mathbb{R}^H$ are learnable parameters. EGG's `RnnReceiverGS` wrapper processes the incoming message through an LSTM to produce embedding h_m. The receiver then computes discrimination scores via dot product between the message embedding and each candidate encoding:

$$s_j = h_m^\top h_r^{(j)} \tag{11}$$

These scores are normalized to probabilities via softmax, as described in Section 3.1.

B Agent Turnover and Population Size

This preliminary experiment characterizes how agent turnover affects communication success and language stability. It was used to choose parameter settings for the main experiment. Turnover introduces two challenges: newcomers must learn the existing language from experienced agents, and the language itself may drift as the population composition changes over time.

B.1 Design

Two sub-experiments were conducted. The **rate study** varies turnover frequency while holding population size constant, testing how quickly agents can be replaced before communication breaks down. The **population study** varies population size while holding turnover rate constant, testing whether larger or smaller populations better maintain language stability under turnover.

Table 3 summarizes the experimental configurations. The rate study uses a fixed population of $N = 10$ agents and varies the turnover rate $k \in \{1, 2, 5, 10, 20\}$, where k specifies the number of epochs between consecutive agent deaths. The population study uses a fixed turnover rate of $k = 10$ and varies population size $N \in \{2, 4, 8, 16\}$. Each configuration runs with three random seeds.

The relationship between epochs, deaths, and generations requires clarification. A death occurs every k epochs, so after E epochs the population has experienced E/k deaths. A generation represents one complete population replacement cycle of N deaths. Thus generation $g = \lfloor \text{deaths}/N \rfloor + 1$. Both sub-experiments run for 5 generations, but this corresponds to different total epoch counts. For the rate study with $N = 10$: 5 generations requires $5 \times N \times k = 50k$ epochs, ranging from 50 epochs at $k = 1$ to 1000 epochs at $k = 20$. For the population study with $k = 10$: 5 generations requires $5 \times N \times k = 50N$ epochs, ranging from 100 epochs at $N = 2$ to 800 epochs at $N = 16$.

The turnover rate k creates a confounding variable in the experimental design: varying k simultaneously alters the learning opportunity for newcomers and the total optimization time for the population. A population with frequent turnover (low k) faces high disruption but undergoes fewer total gradient updates per generation, whereas a stable population (high k) accumulates more parameter

Sub-experiment	Variable	Values	Fixed	Warmup	Seeds
Rate study	k	1, 2, 5, 10, 20	$N = 10$	Pre-trained	3
Population study	N	2, 4, 8, 16	$k = 10$	Until 95% acc.	3

Table 3: Appendix B experiment configurations. Both sub-experiments run for 5 complete generations.

updates - and potentially more drift - over the same generational span. This trade-off makes it difficult to attribute observed drift patterns to a single cause, and motivates examining multiple complementary metrics.

The rate study loads agents from a pre-trained checkpoint where the population has already converged, isolating turnover effects from initial learning dynamics. The population study includes integrated warmup: training proceeds normally until accuracy reaches 95%, at which point turnover begins and generation counting starts. All other parameters follow those used in the main paper (see Table 4).

B.2 Results

Larger populations require more training to establish an initial shared language before turnover can begin. Figure 7(a) shows warmup duration for the population study. Populations of $N = 2$ reach the 95% accuracy threshold in approximately 70 epochs, while populations of $N = 16$ require approximately 290 epochs.

Once turnover begins, its rate strongly affects communication success. Figure 7(b) shows test accuracy over epochs for the rate study. With $k = 20$ (one death every 20 epochs), accuracy remains stable above 90%. As turnover rate increases, accuracy degrades progressively: $k = 10$ maintains accuracy around 80%, $k = 5$ shows more disruption at 65–70%, $k = 2$ drops to 45–50%, and $k = 1$ causes severe degradation to 30–40%. The pattern reflects a repeating cycle: a newcomer enters with random weights, accuracy drops, the newcomer partially learns the language, accuracy recovers, then another newcomer enters. With very frequent turnover, the population never fully recovers between replacement events.

Beyond immediate accuracy disruption, turnover causes cumulative language drift. Figure 8 shows cross-generational metrics for the rate study. Cross-generational accuracy (a) measures whether current agents can communicate successfully with the founding generation. Starting near 95%, it declines steadily as agent replacements accumulate. All turnover rates eventually converge to values between 25–35%, approaching the random baseline of 16.7%. However, slower turnover ($k = 20$) maintains higher cross-generational accuracy for longer, while faster turnover shows more rapid initial decline. Cross-generational similarity (b) shows parallel decline, confirming that the message structure itself changes as the population turns over.

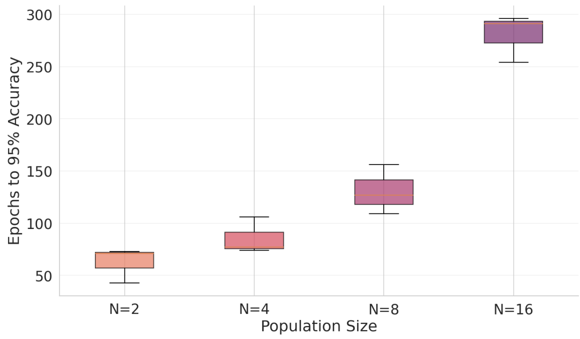

(a) Warmup duration for different population sizes with $k = 10$: epochs required to reach 95% accuracy before turnover starts. Larger populations need more epochs to converge.

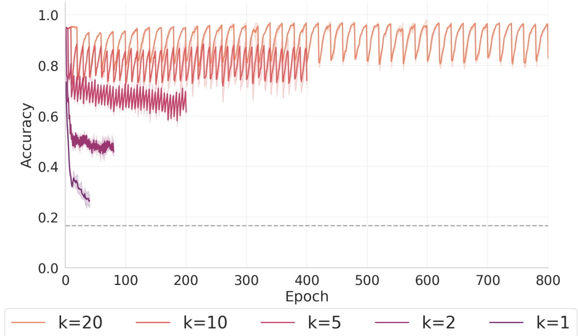

(b) Test accuracy for different turnover rates with population size $N = 10$ after warmup. Each replacement causes a drop, followed by (potentially partial) recovery. More frequent replacements lower the accuracy ceiling. Curves end earlier for smaller k because the same amount of replacements happen in less epochs.

Figure 7: Differences in warmup and training for different turnover rates k and population sizes N.

The results reveal how the k trade-off manifests in practice. Higher k values produce higher within-generation accuracy because newcomers have time to learn, but they also allow more total epochs for the language to shift through optimization. Lower k values prevent effective learning, causing immediate accuracy collapse, but span fewer total

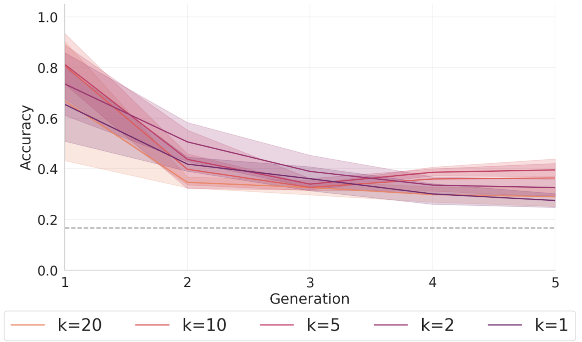

(a) Cross-generational accuracy for varying k and $N = 10$ measured for each generation against Gen 1. All populations lose intelligibility.

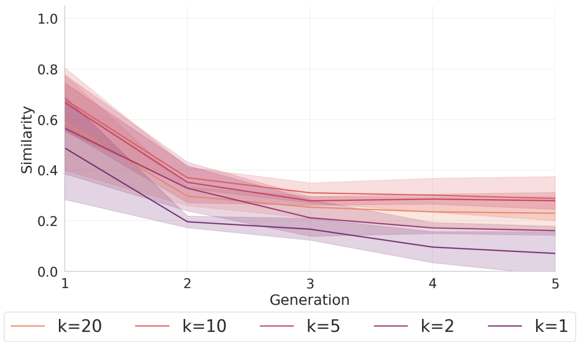

(b) Cross-generational language similarity for varying k and $N = 10$ measured for each generation against Gen 1. Similarity declines alongside accuracy indicating language drift.

Figure 8: Language drift for varying turnover rate k in population size $N = 10$.

epochs. When normalized by generation rather than epoch, all conditions show similar drift trajectories, suggesting that the number of replacement events matters more than the total training time.

The population study reveals how population size affects drift dynamics. Figure 9 plots cross-generational metrics by generation, normalizing the x-axis to enable comparison across population sizes. For cross-generational accuracy (a), smaller populations ($N = 2$, $N = 4$) maintain higher values through early generations, while larger populations ($N = 16$) show faster initial decline. In smaller populations, each agent represents a larger fraction of the shared language, so newcomers must conform more strictly to existing conventions. In larger populations, individual agents have less influence, allowing more variation to accumulate. By generation 5, all population sizes show substantial drift, though smaller populations retain modestly higher cross-generational accuracy. Cross-generational similarity (b) shows a similar pattern.

Interpreting cross-generational accuracy requires care. A value above the random baseline does not

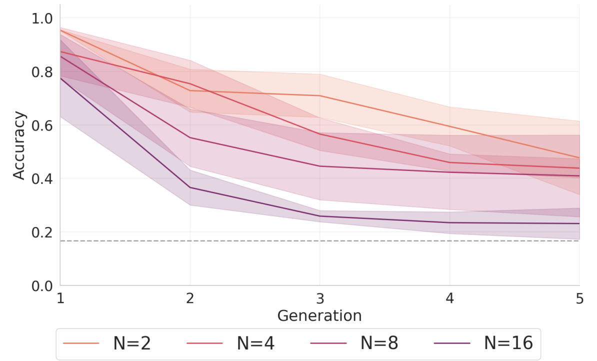

(a) Cross-generational accuracy for varying N and $k = 10$ measured for each generation against Gen 1. Smaller populations retain intelligibility better.

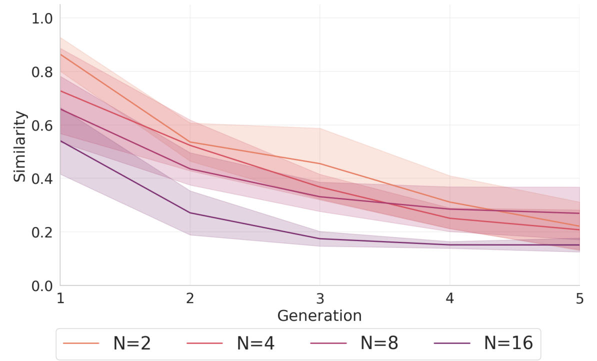

(b) Cross-generational language similarity for varying N and $k = 10$ measured for each generation against Gen 1. Similarity shows the same pattern as accuracy - smaller populations showcase less drift.

Figure 9: Language drift for varying population size N with turnover rate $k = 10$.

unambiguously indicate that the original language persists. Two alternative explanations exist: the language may have drifted while preserving some features of the original structure, or independently evolved languages may achieve above-chance accuracy due to convergent solutions imposed by the task constraints. Cross-generational similarity helps disambiguate these cases. When both metrics decline together, as observed here, the language form has genuinely changed. When accuracy remains high but similarity drops, agents may have found functionally equivalent but structurally different encodings. The parallel decline of both metrics in these experiments suggests true drift rather than convergent solutions.

These results establish that turnover creates fundamental challenges for language stability. Newcomers struggle to learn the existing language quickly enough, causing immediate accuracy drops that scale with turnover frequency. Over longer timescales, the language drifts regardless of turnover rate or population size. The founding generation's language gradually gives way to successor languages that each generation can use internally but that become increasingly incompatible with their predecessors.

C Language Emergence under Heterogeneous Initialization

The main experiments load from the same warmup checkpoint under baseline conditions so that language drift can be measured against the same language. This experiment shows that the age-based population can also reliably develop a language from random initialization (start "from scratch").

C.1 Design

In this experiment the age-based plasticity and agent turnover are enabled from the start. We run three conditions that differ only in the initial age distribution of the population. Young start initializes all agents at age 0, mature start at age 100, and mixed start staggers the ages at [0, 10, 20 ... 90] to maximize initial heterogeneity. All other parameters match the age-based condition from 3.6. Each condition runs for 500 epochs with 3 seeds.

C.2 Results

All three conditions converge to high-accuracy shared languages (Figure 10). The mature and mixed condition perform similarly, while the young condition is visibly slower to reach high accuracy. We do not report cross-generational metrics vs. Gen 1, because Gen 1 here is itself randomly initialized, so there is no converged language for later generations to be measured against. Cross-generational metrics vs. Gen $g-1$ indicate that the languages that emerge still exhibit some amount of language drift. The results demonstrate that a heterogeneous population of agents with age-based plasticity can lead to the emergence of a reliable language from scratch as well.

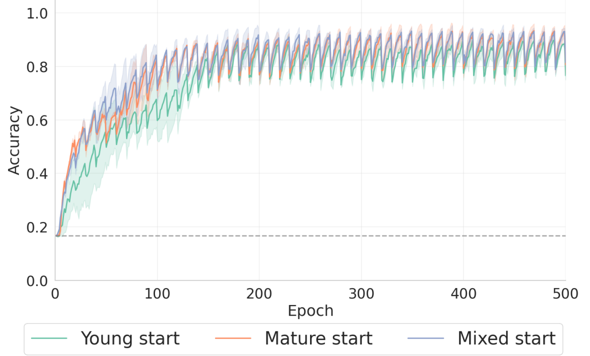

(a) Test accuracy for age-based populations. All conditions develop effective communication. Young-start is slower to converge because initially no agents have low plasticity that would anchor the language.

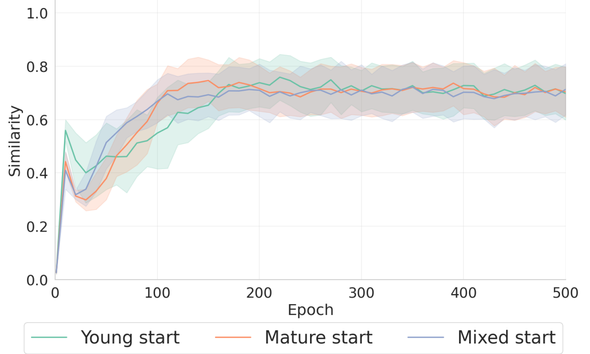

(b) Language similarity for age-based populations. All conditions converge to synchronized languages.

Figure 10: Populations with turnover and age-based plasticity enabled from the start develop a shared language successfully.

D Reproducibility

D.1 Code Availability

The source code is available at `https://github.com/Jano04/aging-game`. The implementation extends the EGG framework (Kharitonov et al., 2019) with modules for population dynamics and age-based plasticity in `egg/zoo/aging/`. The repository README provides installation instructions and usage details.

D.2 Hardware

All experiments ran on REL Compute infrastructure at the Leiden Institute of Advanced Computer Science (LIACS), Leiden University. Specifically, we used the `vibranium.liacs.nl` server:

- CPU: 24 Intel Xeon Silver 4214 cores @ 2.20GHz (48 threads)
- GPU: 2× NVIDIA GeForce RTX 3090 (24GB memory each)
- RAM: 256GB
- OS: Rocky Linux 9

D.3 Default Hyperparameters

Table 4 lists hyperparameters shared across all experiments. The main paper documents experiment-specific variations.

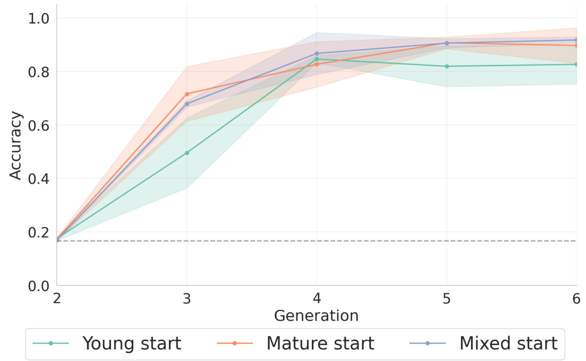

(a) Cross-generational accuracy for age-based populations.

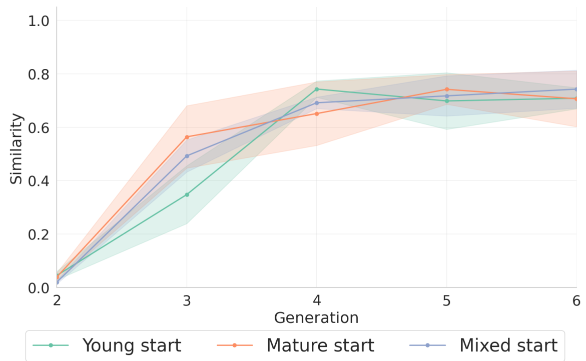

(b) Cross-generational similarity for age-based populations.

Figure 11: Cross-generational metrics measured for each generation against Gen $g-1$. Early generations are still forming a stable language, but by gen 3-4 the language is well established and next generations inherit it.

Parameter	Value
Batch size	1024
Learning rate	10^{-3}
Weight decay	10^{-5}
Hidden dimension	128
Embedding dimension	64
Vocabulary size	20
Max message length	10
Optimizer	Adam
Data seed	111
Training seeds	0,1,2

Table 4: Default hyperparameters.

D.4 Snapshots

A snapshot records a population state at a specific training epoch. Each snapshot stores: messages produced by all alive agents for the validation set $\mathcal{X}_{\text{val}}$, the corresponding inputs (needed for cross-generational evaluation), agent ages and birth epochs, and sender/receiver weights. Snapshots are captured when warmup completes (the generation 1 baseline) and immediately before each agent death.