Emergent Communication in Merging Artificial Agent Populations

Piotr M. Kosela[1[0000-0003-1380-7957]]

Institute of Computer Science, Maria Curie-Skłodowska University, Poland
piotr.kosela@mail.umcs.pl

Abstract. While emergent communication in artificial agents has been widely studied, interactions between previously separated populations remain underexplored, despite their real-world relevance. Our aim is to build a model of two pre-learned populations that meet and attempt to communicate. We develop an agent-based language evolution model, where agents are designed to resemble human internal development as closely as possible. These agents participate in 'language games'—atomic, scripted communication scenarios. When merging two pre-learned populations, we observe a significantly higher rate of successful communication compared to training all agents together from the beginning. This effect persists even after extended simulation of the merged population. Our findings suggest that merging pre-learned populations can enhance communication efficiency, offering practical insights for designing collaborative AI systems.

Keywords: Artificial Agents · Language Evolution · Language Games.

1 Introduction

Human language evolves on various timescales and through diverse mechanisms, often too gradually to be observed within a single generation or even through historical records. This makes capturing the evolution of communication a challenging task. A significant breakthrough in studying language dynamics has been made possible with advances in computational methods, enabling empirical-like linguistic simulations. Various approaches in computational linguistics have yielded insights into these processes [1]. In this study, we examine differences between language evolution in a single, uniform population and communication developed by two distinct groups with pre-existing internal lexicons. We demonstrate that the latter scenario fosters more efficient communication.

Section 2 begins with an overview of agent-based modeling, followed by a review of computational studies examining the impact of social structures on language development. Section 3 details our model, while Section 4 presents the results obtained in the experiment. In Section 5, we analyze these findings, leading to the conclusions in Section 6. Finally, Section 7 outlines directions for future research.

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97557-8_19

2 Related works

Performing experiments *in silico* became feasible with the rise of high-performance computing, allowing for the simulation of communication down to the level of individual utterances. Unlike mathematical formulations based on mean-field approximations, agent-based modeling enables the inspection of individual agents, making it a powerful tool for studying language development. This study adopts the *language games paradigm*, an approach pioneered by Steels in the well-known "talking heads" experiment [15, 14], which has significantly shaped computational research on language evolution.

The "talking heads" experiment was originally designed to address the symbol grounding problem [14, 2]. Its success led to a wave of research employing agent-based models to investigate linguistic dynamics. Spike et al. [13] provide an extensive, though not exhaustive, survey of such models, analyzing the minimal agent capabilities required for communication consensus. While such foundational models have theoretical and philosophical significance, their ability to explain large-scale linguistic phenomena is limited. Increasing agent complexity is necessary for studying population-wide linguistic behavior.

A crucial question in agent-based modeling is its alignment with real-world language emergence. While such cases are rare, empirical observations exist. Richie et al. [11] study homesigners—individuals who invent lexicons to communicate within deaf families—and compare their lexicon development to the slower progress observed in Nicaraguan Sign Language. Their computational model supports claims that social structure plays a crucial role in language evolution. In small home environments, interaction is centralized around a single deaf individual, whereas Nicaraguan Sign Language users engage in broader, diffused communication networks, mirroring findings from computational models.

Agent-based models have been instrumental in exploring the impact of social topology on language development. Labov's Harlem study [7] empirically demonstrated how social structure influences lexical evolution, inspiring computational models such as those by Fagyal et al. [4]. Their results corroborate Labov's findings, showing that communication frequency strongly affects linguistic consensus. Effective information processing within a group requires both leaders and loners—agents with high and low communication frequencies, respectively.

Further research investigates how network topology affects language evolution. Zubek [20] models interacting agent populations using language games, testing information flow efficiency under various social network configurations. Fully connected networks are optimal for information transfer, but star-like structures with designated leaders demonstrate greater adaptability to environmental change and higher overall communicative success. Gong et al. [5] identify key factors shaping linguistic development, including average vertex degree, shortcut connections between clusters, and network centrality. Their findings indicate that higher centrality reduces the number of linguistic categories required for communication.

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97557-8_19

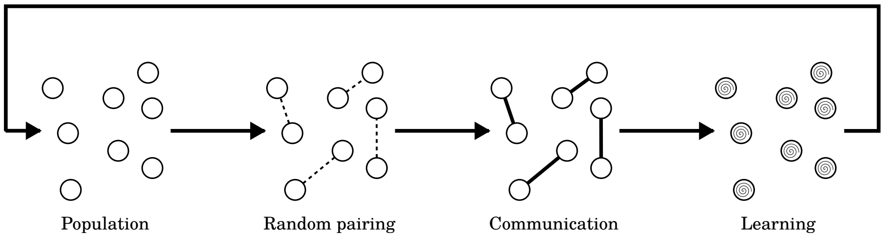

Fig. 1: Steps of simulation repeated in every epoch.

While prior studies focus on language evolution leading to eventual consensus, less attention has been given to how internal population diversity shapes linguistic equilibrium. Josserand et al. [6] employ Bayesian models to study the spread of sociolinguistic variants, demonstrating that even a small subset of biased agents can significantly influence the system's development and final linguistic landscape.

Existing research predominantly examines single populations reaching consensus from scratch. In contrast, real-world scenarios often involve the interaction of groups with pre-established lexicons [11]. In this study, we construct an agent-based computational model comparing two cases: one where a single population develops language from the ground up and another where two distinct populations interact. Our results suggest that merging populations leads to greater communicative success than learning from scratch.

3 Methods

The core objective of our model is to simulate the lexical dynamics of a population of agents engaged in simple linguistic interactions. This process is iterative, with each iteration referred to as an *epoch*. In every epoch, agents are paired uniformly at random[1], engage in scripted communication, and adjust their internal states based on the interaction outcome. Figure 1 illustrates this general process. The remainder of this section details the model's components, beginning with the internal architecture of agents, followed by communication mechanisms, and concluding with the parameter values and structures used in the experiments, distinguishing between control and research trials.

3.1 Agent Architecture

Embodied agents form the fundamental units of our model. Their capabilities are divided into three components [15]:

– Body
– Categorization system

[1] Or *almost random*, as explained in Subsection 3.3.

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97557-8_19

– Lexicon

The body is the only publicly accessible component, facilitating communication. The categorization system and lexicon remain private to each agent.

Agents possess several specific abilities:

– Perceiving objects
– Perceiving (hearing) words
– Uttering words
– Pointing to objects
– Inferring pointing targets

While embodiment enables communication, further exploration of bodily aspects lies beyond this study's scope.

The core capability of agents is categorization, requiring an evolving internal category structure. Since categorization is dynamic, no predefined categories exist; they emerge through agent interactions. This is implemented using a modified SUSTAIN algorithm [9], which models human categorization. Agents categorize objects represented as d-dimensional vectors $v \in \mathbb{R}^d$, where features correspond to real-valued attributes.

A category abstracts similar objects. For instance, a blue pen and a green pen may form the *pen* category, while a yellow pencil belongs to *pencils*. Higher-level categories, such as *writing utilities*, may also emerge. The model accounts for prototypical examples, as observed in human cognition [12].

Each category consists of a set of prototype-object pairs, with weights denoting importance:

$$C = \big\{(p_0, w_0), (p_1, w_1), \dots, (p_k, w_k)\big\},$$

where $p_i \in \mathbb{R}^d$ are prototypes, and $w_i \in \mathbb{R}+$ are their corresponding weights. Categorization of an object $x = (x_0, x_1, \dots, x_d) \in \mathbb{R}^d$ is determined by computing the activation function for each category:

$$h_C(x) = \sum_{j=0}^{k} w_k \exp\left(-\frac{1}{2}\sum_{i=0}^{d}(x_i - p_{j,i})^2\right).$$

The category with the highest activation is selected. However, modifications to categorization occur based on language game outcomes (Subsection 3.2), not categorization alone. Each category also maintains a *communicative success* score, representing the ratio of successful interactions.

To introduce system-wide dynamics, we define a *degeneration rate* δ. After each epoch, prototype weights decay:

$$w(t+1) = w(t)(1-\delta).$$

This mechanism weakens outdated categories unless actively maintained, preventing static structures and ensuring continual adaptation.

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97557-8_19

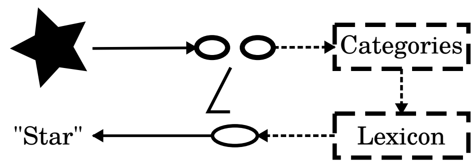

Fig. 2: Speaker perceiving an object and producing an utterance during a guessing game. Dashed lines denote internal agent processes.

Since categories are internal constructs, a lexicon is required to link them with words. Each category may be associated with multiple words, with selection favoring previously successful terms. The lexicon is a many-to-many mapping between categories and words, with weights determining preference. Words are atomic and drawn from a global pool; no two agents independently invent the same word. Weights evolve dynamically, with lexicon links removed when weights fall below $\varepsilon = 10^{-7}$.

3.2 Language Games

Linguistic interactions are modeled using language games, commonly employed in similar simulations [16]. Specifically, we utilize the *discrimination game* and *guessing game*. Both require a small object set, termed the *environment*, containing five elements. Each object has three features drawn from a normal distribution:[2] One element is designated as the *topic*. Both games conclude with either `SUCCESS` or `FAILURE`, affecting agent adaptation.

The *discrimination game* ensures that an agent can uniquely categorize the topic within its environment. The process is as follows:

1. The agent perceives a random environment.
2. The agent assigns categories to objects.
3. If the topic shares a category with another object, the game results in `FAILURE`.
4. If the topic has a unique category, the game results in `SUCCESS`.

The *guessing game* involves two agents: a *speaker* and a *hearer*. Intuitively, the speaker perceives the topic, categorizes it, and utters a word, while the hearer attempts to identify the corresponding object. Communication proceeds as follows:

1. Participants are presented with a random, shared environment.
2. Speaker performs discrimination game.
3. If the discrimination game's result is `FAILURE`:
   – If the topic's category communicative success is greater than 0.95, the topic is added as a new prototype to this category in the speaker's internal memory. Otherwise, a new category centered on the topic is created.

[2] Means = (66.97, 18.65, 38.36), standard deviations = (20.73, 35.11, 39.77)[20].

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97557-8_19

  - Weights of all prototype in the speaker's categories structure are decreased.
  - The guessing game ends with `FAILURE`.
4. If there is no word associated with the topic's category, the speaker invents one and immediately associates it with the category.
5. Speaker utters the word with the strongest association with the topic's category.
6. Hearer hears the spoken word.
7. If the hearer does not know the word:
  - Hearer performs the discrimination game. In case of `FAILURE`, if the topic's category communicative success is greater than 0.95, the topic is added as a new prototype to this category in the hearer's internal memory. Otherwise, a new category centered on the topic is created.
  - Topic's category is associated with the spoken word.
  - Weights of all prototypes in the hearer's categories structure are decreased.
  - The guessing game ends with `FAILURE`.
8. Hearer chooses the category with the strongest association with the spoken word.
9. Hearer points to the object, which suits the chosen category best.
10. If the hearer does not point to the topic meant by the speaker:
  - Speaker decreases the strength of association between the chosen category and spoken word.
  - Hearer decreases the strength of association between the chosen category and spoken word.
  - If the topic's category communicative success is greater than 0.95, the topic is added as a new prototype to this category in the hearer's internal memory. Otherwise, a new category centered on the topic is created.
  - Weights of all prototypes in the hearer's categories structure are decreased.
  - The guessing game ends with `FAILURE`.
11. The guessing game ends with `SUCCESS`.

Prototype weight reduction following an unsuccessful guessing game mirrors the degeneration process, but instead of the degeneration rate δ, we apply a distinct parameter, the *forgetting rate* ϕ, according to the equation:

$$w(t+1) = w(t)(1-\phi).$$

This ensures that less effective prototypes gradually lose influence, facilitating adaptation to new linguistic patterns. The key difference is that degeneration follows every language game, whereas forgetting is applied only after failed games, serving as a penalty for incorrect classification.

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97557-8_19

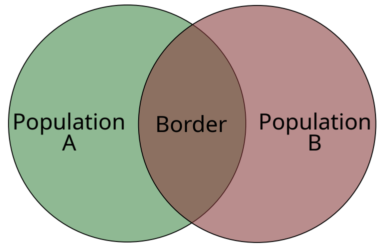

Fig. 3: Diagram of the merged populations, highlighting the border area.

3.3 Experimental Procedure

We detail the experimental procedure below. The control sample is generated using the method described at the beginning of Section 3, simulating 1,000,000 epochs with 200 agents.

The research trials consist of two stages. The first stage aims to develop basic categorization abilities and small lexicons in agents. Two separate populations of 100 agents each are simulated for 200,000 steps. This duration, chosen heuristically, ensures sufficient communicative success while maintaining agent flexibility to adapt to changes in population and environment.

The second stage combines the two populations into a single set of 200 agents, simulating an additional 1,000,000 epochs. The pairing rules are modified: 20 agents from each population are designated as *border agents*, with a probability $\beta = 0.8$ of interacting with border agents from the other population. One could think of them as interpreters or translators — individuals who engage with members of the other population, acquire their linguistic conventions, and transmit this knowledge back to their own group. Non-border agents interact only within their original populations. A schematic of this setup is shown in Figure 3.

The research procedure was repeated independently 20 times to allow for result averaging. Key learning parameters were held constant: $\phi = 0.0005$ and $\delta = 0.0000625$. These values, chosen heuristically, balance communicative success, computational complexity, and population flexibility. Increasing ϕ and δ tends to destabilize lexicons, lowering the model's overall success, while decreasing them makes populations resistant to new variants introduced by border agents. Investigating these effects in large populations is challenging due to the substantial computational resources required by the simulation.

After each epoch, we compute several metrics:

– *Communicative Success (CS)*: The ratio of successful guessing games to total interactions.
– *Average Number of Categories (AVGC)*: The mean number of categories retained by agents.
– *Category Operations*: Total numbers of *created (C)*, *modified (M)*, and *deleted (D)* categories. Modifications refer to changes in category prototypes.

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97557-8_19

4 Results

In this section, we present the experimental results, comparing the development of the research sample (first and second stages) with the control sample. Each plot is derived by pointwise averaging data from 20 independent trials, followed by a moving average with a block size of 10,000. We begin with a comparison of the first $200k$ steps for both samples.

The control sample consists of 200 agents communicating uniformly. Figure 4 compares three populations: A, B, and *control*. Figure 4a shows communicative success, with the control sample stabilizing at 0.25, while the research sample achieves 0.32 due to its smaller, independent populations. This aligns with findings in [20].

Figure 4b reveals linear growth in the average number of categories per agent across all populations. Figure 4c shows minimal deletions and stable modifications (~ 0.2) after $15k$ epochs. The total number of created categories remains constant, except for initial generations where agents compensate for the lack of categories. Notably, the slowing of category creation does not trigger deletions.

After $200k$ steps, populations A and B are connected via border agents. Figure 5 presents post-merge results, with epoch 0 marking the merge.

Communicative success in the research sample drops to 0.3 post-merge (Figure 5a), consistent with border agents. The control sample shows no change. Figure 5b shows the research sample stabilizing at $11k$ categories per agent by $300k$ epochs, while the control sample reaches $12k$ by $575k$ epochs.

Figure 5c shows higher category creation in the research sample post-merge (4 to 4.5). Modifications increase as category growth stabilizes, with deletions starting concurrently. The research sample exhibits two peaks in deletions, while the control sample stabilizes quickly. Figure 5d shows border agents with fewer creations and delayed deletions.

Finally, Figure 6 presents heatmaps of word usage. Figure 6c shows the control sample, with early words reused successfully. Figure 6a depicts post-merge word usage, with new signs dominating early words. Figure 6b shows border agents following population trends, with successful interactions around generation $500k$.

5 Discussion

In this section, we analyze the findings of our model, beginning with an interpretation and explanation of the results. We then take a broader perspective on language evolution models, relating our observations to real-world situations.

Let us begin by focusing on the behavior of communicative success. It drops immediately after the merge. This decrease is likely caused by the onset of trans-population interactions, while the groups do not share any words. However, the metric does not recover [3]. Thus, we conjecture that the drop in the research

[3] We performed many more simulation steps to confirm this.

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97557-8_19

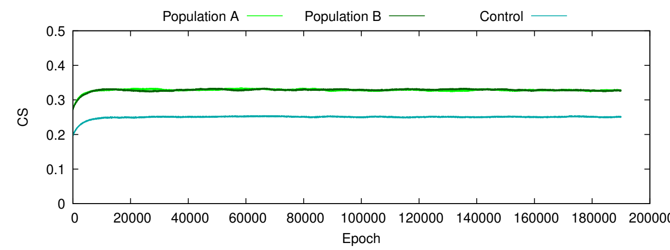

(a) Communicative success across research and control samples.

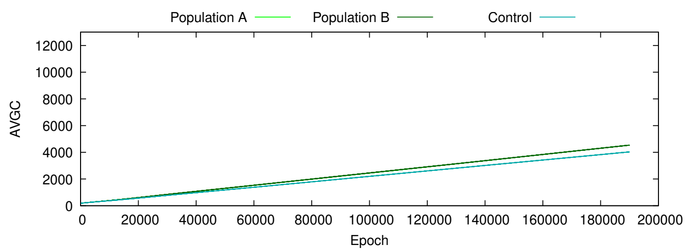

(b) Average number of categories per agent in research and control populations.

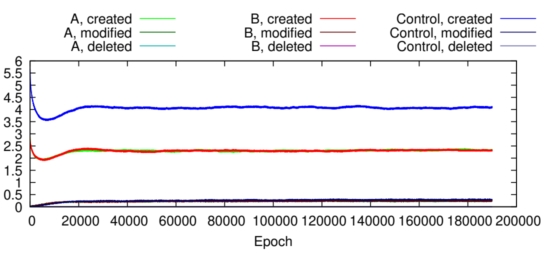

(c) Average total number of created, modified, and deleted categories in research and control samples.

Fig. 4: Comparison of communicative success, agent categories, and category operations between the control sample and research sample populations.

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97557-8_19

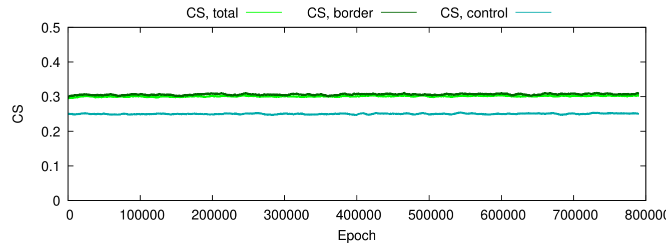

(a) Communicative success for control, research, and border agents.

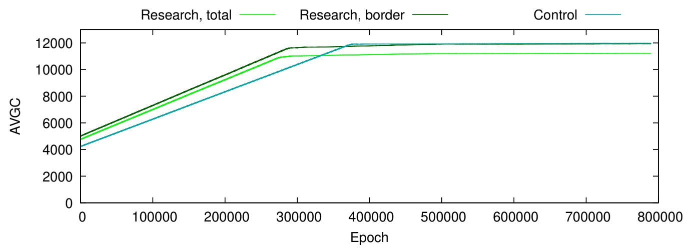

(b) Average number of categories per agent for control, research, and border agents.

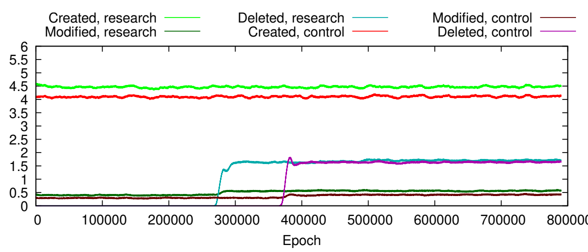

(c) Average total number of created, modified, and deleted categories for research and control samples.

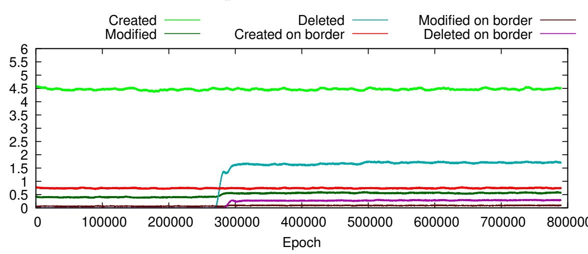

(d) Average total number of created, modified, and deleted categories for research sample and border agents.

Fig. 5: Communicative success and category dynamics in the second stage for control, research, and border agents.

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97557-8_19

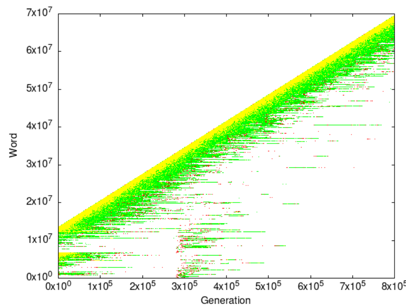

(a) Word usage in merged populations.

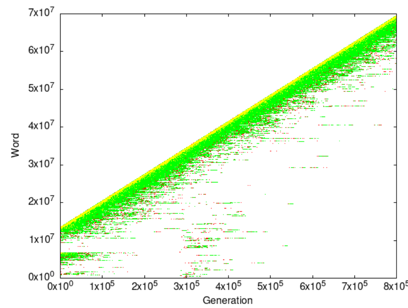

(b) Word usage in the border of merged populations.

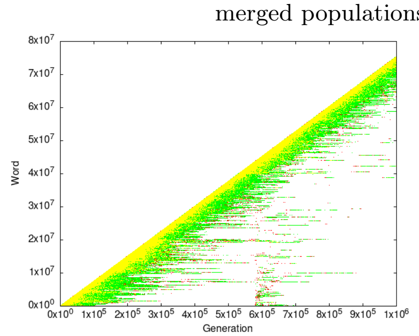

(c) Word usage in control sample.

Fig. 6: Heatmaps showing word usage (Y-axis) across generations (X-axis). Colors indicate interaction results: green (success), red (failure), yellow (mixed).

sample is an artifact resulting from model scaling issues [20]. CS at the border does not change over time. It is not significantly smaller, so it cannot affect the measure for the entire population. Nevertheless, the final population achieves better results than the control sample, even though it contains the same number of agents. The loss of CS from scaling is much smaller than in the case of agents learning together from the very beginning.

The CS of the research sample oscillated around 0.3 with fewer than $11k$ categories per agent. In contrast, the control sample had almost $12k$ categories on average, with CS at the level of 0.25. Overall, it appears that the population created from the two smaller ones performs better, surpassing the CS barrier established by the control sample. This suggests that building a population from smaller groups leads to easier adaptation and improved overall model performance. Despite having fewer categories, the research sample demonstrates better object recognition.

It was recently shown that, in the case of human categorization development, exposure to foreign cultures does not affect mature categorization networks [10]. We propose that an analogous effect may be present in our model; note that the weights in our categorization systems are not bounded, and the categories from

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97557-8_19

the preliminary stage of the simulation persist with border agents for an extended period. The average number of categories developed by the border agents is significantly greater than the population average. These rich categorization systems operate with more dynamic vocabularies. Notably, the horizontal lines, starting from generation $300k$ in Figure 6a, are longer than those in Figure 6b. Words do not survive long on the border but are exported to the rest of the population, where their usage is prolonged; this is consistent with [4]. These agents influence the main population and create trends. We believe this is the factor enabling the higher communicative success. The border does not, in fact, consist of leaders in the sense of [4]; Figure 3 may be slightly misleading in this context. These agents interact with the rest of the population only 20% of the time. This subpopulation may function as a small, organized subsystem, exporting this well-designed innovation externally. We might better consider them a collective, more complex outsider, balancing the system.

6 Conclusion

The results of this study demonstrate that populations formed by merging smaller groups achieve higher communicative success compared to populations that learn from scratch. This enhanced performance can be attributed to the richer categorization systems developed by the border agents, who, despite interacting less frequently with the main population, export valuable innovations. These agents facilitate more efficient adaptation, resulting in a significant increase in overall model performance. In contrast, populations that evolve from scratch face greater challenges in reaching similar levels of communicative success, highlighting the advantage of merging pre-existing lexicons.

7 Further research

While our study provides insights into language evolution in merged populations, several open questions remain. Here, we outline key areas for future research that could deepen our understanding of population interactions and model scalability.

One key difference between the control and research samples is the presence of additional peaks in category modifications and deletions in the latter. We believe this is not an inherent phenomenon in the control sample but rather a consequence of slower transition. While such peaks appear in both samples, their patterns differ. Investigating this variation could provide deeper insights into population interactions.

We find it worthwhile to incorporate spatial distribution into the model. While social network topology partially addresses this, no agent-based model we are aware of considers gradual environmental shifts based on agent locations. For instance, the need for specialized snow descriptions in the Arctic Circle compared to the Congo Valley is not well represented in existing models [19]. Social

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97557-8_19

topology influences interaction probability alongside spatial factors. Spatial arrangement shapes populations into smaller societies, while uniform environments foster more sophisticated communication.

Simulating agents in a diverse world is feasible only if communities are large enough. Otherwise, the model would merely reflect a social topology framework with distinct environments for individual agents. At some scale, resource constraints hinder computational feasibility [17]. While [17] highlights model complexity, it underestimates categorization scaling, the most computationally demanding aspect. As [18] shows, similar tasks cannot be achieved in polynomial time. Efficient categorization remains crucial for advancing agent-based language models. Currently, no comprehensive mathematical frameworks exist for explanatory sociolinguistic models, aside from highly simplified cases [8]. While real-world language evolution models exist, their reliance on mathematical apparatus often limits explanatory power [3].

Our results show that merging two independent agent populations leads to higher communicative success, indicating that linguistic convergence benefits from increased diversity in lexical and categorical structures. This finding underscores the importance of cross-population dynamics in language evolution. Future research should extend this framework by incorporating spatially distributed environments, allowing for a more nuanced exploration of regional linguistic variation. Additionally, addressing the computational complexity of our models will be crucial for scaling simulations and capturing more intricate dynamics of category and lexicon formation. Exploring these aspects will deepen our understanding of emergent communication and its sensitivity to environmental constraints.

Acknowledgments

I would like to thank Jarosław Bylina and Piotr Giza for their valuable insights and thoughtful discussions, which greatly supported the development of this work.

References

1. Bailes, R., Cuskley, C.: The cultural evolution of language (2022), publisher: PsyArXiv
2. Bougsty-Marshall, S.: The Dynamics of Growing Symbols: A Ludics Approach to Language Design by Autonomous Agents. In: Goertzel, B., Panov, A.I., Potapov, A., Yampolskiy, R. (eds.) Artificial General Intelligence, vol. 12177, pp. 44–53. Springer International Publishing, Cham (2020). https://doi.org/10.1007/978-3-030-52152-3_5, `http://link.springer.com/10.1007/978-3-030-52152-3_5`, series Title: Lecture Notes in Computer Science
3. Dębowski, L.: Information Theory Meets Power Laws: Stochastic Processes and Language Models. Wiley, 1 edn. (Sep 2020). https://doi.org/10.1002/9781119625384, `https://onlinelibrary.wiley.com/doi/book/10.1002/9781119625384`

DOI: 10.1007/978-3-031-97557-8_19

4. Fagyal, Z., Swarup, S., Escobar, A.M., Gasser, L., Lakkaraju, K.: Centers and peripheries: Network roles in language change. Lingua **120**(8), 2061–2079 (Aug 2010). https://doi.org/10.1016/j.lingua.2010.02.001, https://linkinghub.elsevier.com/retrieve/pii/S0024384110000203
5. Gong, T., Baronchelli, A., Puglisi, A., Loreto, V.: Exploring the Roles of Complex Networks in Linguistic Categorization. Artificial Life **18**(1), 107–121 (Dec 2011), https://direct.mit.edu/artl/article/18/1/107-121/2715
6. Josserand, M., Allassonnière-Tang, M., Pellegrino, F., Dediu, D.: Interindividual Variation Refuses to Go Away: A Bayesian Computer Model of Language Change in Communicative Networks. Frontiers in Psychology **12**, 626118 (Jun 2021). https://doi.org/10.3389/fpsyg.2021.626118, https://www.frontiersin.org/articles/10.3389/fpsyg.2021.626118/full
7. Labov, W.: Language in the inner city: Studies in the Black English vernacular. University of Pennsylvania Press (1972), issue: 3
8. Loreto, V., Baronchelli, A., Puglisi, A.: Mathematical Modeling of Language Games. In: Nolfi, S., Mirolli, M. (eds.) Evolution of Communication and Language in Embodied Agents, pp. 263–281. Springer Berlin Heidelberg, Berlin, Heidelberg (2010). https://doi.org/10.1007/978-3-642-01250-1_15, https://link.springer.com/10.1007/978-3-642-01250-1_15
9. Love, B.C., Medin, D.L., Gureckis, T.M.: SUSTAIN: A Network Model of Category Learning. Psychological Review **111**(2), 309–332 (2004). https://doi.org/10.1037/0033-295X.111.2.309, http://doi.apa.org/getdoi.cfm?doi=10.1037/0033-295X.111.2.309
10. Naji, A.: Semantic representations of abstract and concrete categories in the mental lexicon of mono- and bilingual Jordanians: A prototype analysis. Alkalmazott nyelvtudomány **vol. 21. issue 2. ISSN 1587-1061**, eISSN 24984442 (2021). https://doi.org/10.18460/ANY.2021.2.008, http://alkalmazottnyelvtudomany.hu/wordpress/wp-content/uploads/naji_tan.docx.pdf, publisher: Pannon Egyetem
11. Richie, R., Yang, C., Coppola, M.: Modeling the Emergence of Lexicons in Homesign Systems. Topics in Cognitive Science **6**(1), 183–195 (Jan 2014). https://doi.org/10.1111/tops.12076, https://onlinelibrary.wiley.com/doi/10.1111/tops.12076
12. Rosch, E.: Cognitive representations of semantic categories. Journal of Experimental Psychology: General **104**(3), 192–233 (Sep 1975). https://doi.org/10.1037/0096-3445.104.3.192, http://doi.apa.org/getdoi.cfm?doi=10.1037/0096-3445.104.3.192
13. Spike, M., Stadler, K., Kirby, S., Smith, K.: Minimal Requirements for the Emergence of Learned Signaling. Cognitive Science **41**(3), 623–658 (Apr 2017). https://doi.org/10.1111/cogs.12351, https://onlinelibrary.wiley.com/doi/10.1111/cogs.12351
14. Steels, L.: The Origins of Ontologies and Communication Conventions in Multi-Agent Systems. Autonomous Agents and Multi-Agent Systems **1**, 169–194 (1998)
15. Steels, L.: Evolving grounded communication for robots. Trends in Cognitive Sciences **7**(7), 308–312 (Jul 2003). https://doi.org/10.1016/S1364-6613(03)00129-3, https://linkinghub.elsevier.com/retrieve/pii/S1364661303001293
16. Steels, L., Belpaeme, T.: Coordinating perceptually grounded categories through language: A case study for colour. Behavioral and Brain Sciences **28**(4), 469–489 (Aug 2005). https://doi.org/10.1017/S0140525X05000087, https://www.cambridge.org/core/product/identifier/S0140525X05000087/type/journal_article

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97557-8_19

17. Vogt, P., Divina, F.: Language evolution in large populations of autonomous agents: issues in scaling. In: Proceedings of AISB. pp. 80–87. Citeseer (2005)
18. Woensdregt, M., Cummins, C., Smith, K.: A computational model of the cultural co-evolution of language and mindreading. Synthese **199**(1-2), 1347–1385 (Dec 2021). https://doi.org/10.1007/s11229-020-02798-7, https://link.springer.com/10.1007/s11229-020-02798-7
19. Yule, G.: The Study of Language. Cambridge University Press, 7 edn. (Nov 2019). https://doi.org/10.1017/9781108582889
20. Zubek, J., Denkiewicz, M., Barański, J., Wróblewski, P., Rączaszek-Leonardi, J., Plewczynski, D.: Social adaptation in multi-agent model of linguistic categorization is affected by network information flow. PLOS ONE **12**(8), e0182490 (Aug 2017). https://doi.org/10.1371/journal.pone.0182490, https://dx.plos.org/10.1371/journal.pone.0182490

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97557-8_19