

EMERGENCE OF GROUNDED, OPTIMALLY COMPOSITIONAL SPATIAL LANGUAGE AMONG HOMOGENEOUS AGENTS

Anonymous authors

Paper under double-blind review

ABSTRACT

A mechanism of effective communication is integral to human existence. An essential aspect of a functional communication scheme among a rational human population involves an efficient, adaptive, and coherent apparatus to convey one’s goal to others. Such an effective macro characteristic can emerge in a finite population through adaptive learning via trial and error at the individual (micro) level, with nearly consistent individual learning faculty and experience across the population. In this paper, we study and hypothesize pertinent aspects of glossogenetics, specifically primal human communication mechanisms, through computational modeling. In particular, we model the process as a language “guessing” game within the fabric of a decentralized, multi-agent deep reinforcement learning setting, where the agents with local learning and neural cognitive faculties interact through a series of dialogues. Our homogeneous agents seek to achieve the principle of least effort and overcome the poverty of stimulus through efficient concept selection, guided feedback and mirror learning. In our examinations, we observe the emergence of successful and structured communication among static and dynamic agent populations through guided, reinforced and continual learning.

1 INTRODUCTION

Effective communication via signals is the key to success in a cooperative world, where the goal is to complete the desired tasks by efficiently coordinating among themselves. A functional communication language should be nearly unambiguous, efficient, easily acquirable (culturally transmitted), and rooted in the environment. Language Chomsky (2006); Montague et al. (1970); de Saussure (2011) is an autonomous, culturally transmitted, complex adaptive system realised through multiple modalities—either vocal or gestures—that translate mental representations, which are internal structures, to utterances that represent the surface structure. In cooperative AI and cognitive science, language games Wittgenstein (1954); David (1969); Arrington (1954); Steels (1997; 2003); Wagner et al. (2003) which was motivated by the picture theory of language Wittgenstein (1954) and operant conditioning theory Skinner (1986) are empirical computational models developed to study the origin, evolution, and acquisition of human languages. The game setting involves a bottom-up simulation model which usually consists of multiple artificial agents (neural or statistical) equipped with sufficient cognitive/inductive reasoning abilities and sometimes sensory-motor systems interacting in a shared environment through vocal or non-vocal means and subsequently learning from the outcomes of the interactions. The language structures that emerge in these settings never entirely mirror human languages, since human languages are refined through millions of years of cultural evolution. However, language games can provide deep insights into the emergence of various aspects of human language, such as syntactic structures, compositionality, word order, generalization, brevity, stability, statistical regularity, complexity, coherence, and linguistic divergence.

With the recent advancement in the field of deep learning Mnih et al. (2013) with respect to computational tractability, one could observe rigorous applications of deep learning and deep reinforcement learning involving multiple agents in language games Lazaridou & Baroni (2020); Dafoe et al. (2020), especially Lewis signalling/referential/discrimination games Lazaridou et al. (2017); Havrylov & Titov (2017), reconstruction games Kharitonov et al. (2020), navigation/action games Kajić et al. (2020); Mordatch & Abbeel (2018) and visual communication games Qiu et al. (2022) without customized tweaking. The central feature of all the existing work is that it belongs to the

054 "Naming game" category, where the semantic aspect used to distinguish the topic from other con-
 055 text elements is clear and free from ambiguity. A few of these focus on the emergence of coherent
 056 communication protocols from scratch (tabula rasa) in a multi-deep-agent setup Sukhbaatar et al.
 057 (2016); Foerster et al. (2016); Havrylov & Titov (2017); Lazaridou et al. (2017; 2018). A few others
 058 target the pertinent linguistic universals of natural languages, such as the symbolic grounding Mor-
 059 datch & Abbeel (2018); Kottur et al. (2017); Lin et al. (2021), compositionality Mordatch & Abbeel
 060 (2018); Kottur et al. (2017); Li & Bowling (2019); Ren et al. (2020); Wang et al. (2016); Andreas
 061 (2018), generalization Baroni (2020); Chaabouni et al. (2020), brevity regularity Rita et al. (2020);
 062 Kharitonov et al. (2020), the cultural and architectural transmission Dagan et al. (2020); Ren et al.
 063 (2020), language structures through ease-of-teaching pressure Li & Bowling (2019) and networked
 064 communication Gupta et al. (2020). Some of the recent works also provide deeper analysis pertain-
 065 ing to the nature and factors affecting the semiotic dynamics underlying the emergence of language
 066 and language constructs. Kottur et al. (2017); Resnick et al. (2020); Tucker et al. (2022) delve into
 067 the factors and constraints such as selectionist criteria, utility, informativeness, memory capacity,
 068 and learning capabilities that contribute to the development of compositionality and Graesser et al.
 069 (2019); Eccles et al. (2019); Gaya et al. (2016) analyze conditions, inductive biases and intrinsic
 070 motivation required for the emergence of a coherent language. Another direction in which language
 071 emergence is being evaluated is along the dimension of scale Chaabouni et al. (2022); Rita et al.
 072 (2021), where the correlation between language characteristics and system complexity and popula-
 073 tion dynamics is examined, while Lazaridou et al. (2020) incorporates pre-trained general language
 074 models to develop task-specific language models. Choi et al. (2018) explores the obverter technique
 075 Batali (1998) which explores emergent communication among pseudo-homogeneous agents, where
 076 the speaker is assumed to be always true, while the listener calibrates its parameters to align itself
 077 with the speaker. Chaabouni et al. (2019) studies the existence of inverse correlation of word length
 and input frequency that exists in natural human language.

078 **Our contribution:** In this paper, we develop a decentralized, uncorrelated, multi-agent "Guessing
 079 game" with significant combinatorial complexity compared to the Naming games, primarily due to
 080 the ambiguity in inferring the topic from the context, which, in turn, significantly hinders vocabu-
 081 lary alignment and learning. Our unique approach includes the integration of an explicit agent
 082 ontology to distinguish between referents and concepts, as well as the provision of tangible feed-
 083 back, which differs from existing works that share feedback directly, mimicking "telepathy". In
 084 this unique guessing game framework, we investigate the emergence of certain coherent properties
 085 of languages such as interchangeability, and compositionality. The interchangeability property of
 086 languages enables the agents to act as either listener or speaker (contrary to dedicated roles in the
 087 existing works), where they synchronise through mirror learning. This ensures continual learning,
 088 which resembles human language acquisition. Additionally, we integrate the principle of least effort
 089 which encourages agents to convey information in a way that minimizes cognitive effort. The goal
 090 of the paper is to explore the dynamics of these components within the guessing game setup and to
 analyze and hypothesize the emergent behaviors that arise from their interaction.

091 2 PROBLEM FORMULATION

092 In this paper, our objective is to enable the emergence of coherent symbolic structures among a popu-
 093 lation of deep neural agents through decentralized learning and self-organization in guessing games.
 094 Our setting consists of N deep neural agents populated on a graph world $G = (V, E)$ which is em-
 095 bedded on a bounded 2D plane (flat earth), where V is the set of vertices and E is the set of edges.
 096 All the nodes are similar in shape. However, they possess two relevant features, location and color
 097 which distinguish them from each other. The location is unique for a node, although they can have
 098 the same color. Our agents are homogeneous in nature, where they can perform both comprehension
 099 and reproduction of language. This property is referred to as *homogeneity/interchangeability* which
 100 is one of the core properties of human language. Our language game (Figure 3) is as follows: At
 101 each instant in the game, one agent is paired against another to initiate a semiotic cycle of dialogue
 102 consisting of D conversations. In dialogue, one random agent takes the role of speaker, while the
 103 other agent is the listener (step ① in Figure 3). In a conversation, the speaker randomly selects a
 104 topic node unknown to the listener and communicates it using a suitable conceptualization (using its
 105 ontology) and utterance via a noise-free, discrete channel (steps ② to ⑤). The listener attempts to
 106 decipher the meaning of the temporal sequence of utterances by correctly identifying the topic node.
 107 The interaction is subsequently rewarded according to the interpretation outcome which is shared
 among the participants (steps ⑥ to ⑧). In case of failed communication, the speaker discloses the

topic node to the listener (step ⑨). Learning occurs through the induction of hypotheses (the innate linguistic structure that is characterized by neural networks) based on payoffs, and disclosures.

We formulate our setting using a multi-agent Markov game framework Littman (1994); Puterman (2014) since we aim for the emergence of symbolic structures through interactions among agents who possess the cognitive ability to extract and reinforce commonalities across multiple experiences. Here, we assume that the agents only have a partial observation of the environment, which aligns with real human scenarios where one can only be aware of his local surroundings and perceive the world in a coarse form. The state of the environment at time step t is denoted by $s^{(t)} \in \mathcal{S}$, where \mathcal{S} is the set of all the environment states. We let $o_i^{(t)} \in \mathcal{O}$ be the partial observation of agent i , which is characterized by the function $f_i : \mathcal{S} \mapsto \mathcal{O}$, where \mathcal{O} is the set of all possible observations. At time instant t , agent i chooses a random action $a_i^{(t)}$ which is dependent on the current observation according to a parameterized stochastic policy $\pi_{\theta_i}(\cdot | o_i^{(t)})$ which is a conditional probability mass function over \mathcal{A} conditioned on the observation $o_i^{(t)}$. For agent i , each state transition yields a random reward $r_i^{(t)}$ according to the function $\mathcal{R} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto \mathbb{R}$. The system evolution is stochastic in nature and characterized by the probability transition function $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto [0, 1]$, where $\mathcal{P}(s, a, s') = \mathbb{P}r(s^{(t+1)} = s' | s^{(t)} = s, a^{(t)} = a)$ which is the conditional probability of next state is s' conditioned on the current state and action being s and a respectively. The collective goal of the agent population is to collaboratively seek a policy $\pi_{\theta_*} = [\pi_{\theta_1^*}, \pi_{\theta_2^*}, \dots, \pi_{\theta_N^*}]$ that maximizes the globally averaged long-term return over the network based solely on local information, *i.e.*,

$$\theta_i^* = \arg \max_{\theta \in \Theta} J_i(\theta), \text{ with } J_i(\theta) = \mathbb{E}_{\pi_{\theta}, \mu} \left[\sum_{t=0}^{T-1} \mathbf{r}_i^{(t)} \right], \quad (1)$$

where $\mathbb{E}_{\pi_{\theta}, \mu}[\cdot]$ is the expectation with respect to all T length trajectories generated using the stochastic policy π_{θ} with initial distribution μ and $\Theta \subset \mathbb{R}^s$ is a compact and convex set.

3 DOMAIN ONTOLOGY

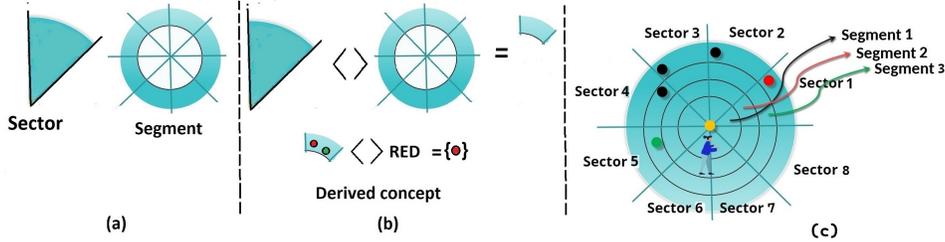


Figure 1: Agent Ontology

The ontology Guarino & Giaretta (1995); Mark et al. (2003) \mathcal{C} of the agent is the concept space through which an agent perceives the space. In this paper, $\mathcal{C} = \mathcal{H} \cup \mathcal{W} \cup \mathcal{B} \cup \{\perp\}$ which consists of a finite collection of segments \mathcal{H} , sectors \mathcal{W} , colors \mathcal{B} and the NULL concept \perp . We let $\bar{\mathcal{H}} = \mathcal{H} \cup \{\perp\}$, $\bar{\mathcal{W}} = \mathcal{W} \cup \{\perp\}$ and $\bar{\mathcal{B}} = \mathcal{B} \cup \{\perp\}$. A segment is a strip of region in the 2D plane encompassed by outer and inner concentric circles (Figure 1 (a)) centered at a certain point. A sector is defined to be a part of a disc made of the arc of the disc centered at a certain point along with its two radii extending to the boundary of the world. These are spatial deixis which are generally perceived relative to the location of the central point. The segments and sectors provide a conceptualization of space that is grounded in the sensory and physical interactions of the agents with the world and one can relate it to the concepts of cardinal directions in the human discourse. Each node possesses the intensive property of color and we assume that the agents possess the sensory mechanism to capture the hue range of colors. Hence, we also consider colors as concepts. Roughly, \mathcal{C} represents the hierarchical deep structure of concepts (semantic entities) where one can be either specific (finer) or general (coarser), or disjoint than the other. The concept space \mathcal{C} is equipped with an operation $\langle \cdot, \cdot \rangle : \mathcal{C} \times \mathcal{C} \rightarrow \mathcal{C}'$, where \mathcal{C}' is the set of derived concepts, which are concepts which can be derived from the basis concepts Andreas (2018); Montague et al. (1970). In our setting, the operation $\langle \cdot, \cdot \rangle$ is set intersection since our concept space consists of regions and colors. An illustration is provided in Figure 1 (b). We also maintain a pre-defined encoder $\Gamma : \mathcal{C} \rightarrow \mathbb{Z}$ which maps the

abstract basis concepts in \mathcal{C} to discrete scalars. Given any topic node, the agent can conceptualize the vertex in terms of the tuple $\langle \text{segment}, \text{sector}, \text{color} \rangle \in \mathcal{H} \times \mathcal{W} \times \mathcal{B}$ relative to the current location of the agent. Here, we abuse the notation, as $\langle \rangle$ typically denotes a binary operation, and in this case, it should be interpreted as $\langle \text{segment}, \langle \text{sector}, \text{color} \rangle \rangle$.

Complexity of Guessing Game: For a given vertex $u \in V$, we consider the function $\mathcal{C}_u : V \rightarrow 2^{\mathcal{H} \times \mathcal{W} \times \mathcal{B}}$ which maps vertices to their corresponding conceptualizations relative to the source vertex u . For a given (source, topic) vertex pair (u, x) , one can have more than one conceptualization possible, *i.e.*, $\mathcal{C}_u(x) \subseteq \mathcal{H} \times \mathcal{W} \times \mathcal{B}$. Hence our setting can be categorized as “guessing game” (Section 1.3.2 of Steels (2012)). The complexity of the guessing game is substantially high due to the inherent meaning ambiguity arising from the existence of more than one possible distinct concept for a given unknown message. Meaning uncertainty arises because multiple concepts can possibly be associated to a novel word and the listener cannot, with only one exposure, determine which meaning is intended by the speaker. The available information for learning or understanding is limited or insufficient for the listener to comprehend. This complicates the construction of a shared vocabulary, as aligning meanings through communication becomes more challenging. This is Quine’s “Gavagai” problem also referred to as *Poverty of stimulus* Quine (1960).

Assumption: In this paper, we assume that the ontology possessed by all the agents is identical and they all conform to the same ontological framework to avoid inconsistent perspectives and thus evade the Tower of Babel situation Iliadis (2019); Mark et al. (2003).

4 GROUNDED VOCABULARY LEARNING

The lexis (Ψ) of a language is a finite catalog of all q -letter words available a priori to an agent. A vocabulary bidirectionally maps lexis (phonological entities) to meanings (semantic entities), where one is able to evoke the other Ren et al. (2020). This symbolic association is referred to as the property of groundedness. For a population of agents to successfully communicate, there should exist a shared, coherent vocabulary among the population. This implies that the vocabulary possessed by the agents should hold the same meaning for everyone to successfully communicate verbally among themselves. Ideally, the mapping should be isomorphic. Apparently, in every realistic scenario, this is not the case, which transpires into various language characteristics like homonyms and synonyms. Initially, there is no ex-ante meaning associated with the words, and hence no coherence among the agents exists and we aim to foster common grounding among agents incrementally, which is fully shaped by past linguistic experience. This is referred to as the symbol grounding problem Steels (2012). We achieve this through verbal interactions between them, where they extract and reinforce similarities across multiple episodes incrementally through evidence of understanding which can be either positive or negative. This trial and error based calibration process shapes, reshapes, and enforces the mental mapping, where the phonological expressions become more efficient and established through repeated use Bisk et al. (2020); Arrington (1954), and eventually drives the system to a dissipative structure Prigogine (1987) which enables common ground for expressing concepts. The speaker exhibits an inherent categorical bias, to structure messages in a specific order ($\mathcal{M}(\text{segment}), \mathcal{M}(\text{sector}), \mathcal{M}(\text{color})$) reflecting the agent’s conceptualization process. However, the listener module does not share this bias and instead treats each word as potentially belonging to any category. This discrepancy in how the speaker and listener process information makes it more challenging to establish a shared language, as the listener does not rely on the same categorical ordering as the speaker.

Assumption: We assume that all the agents possess the same lexis.

5 POLICY ARCHITECTURE

The policy architecture of the agent is modeled using stochastic neural networks. Each agent consists of two modules: speaking (concept-selection and utterance) and listening. All the modules are implemented using RNN (recurrent neural networks) to allow for continuous and sequential communication. Here θ , ψ , and ϕ represent the parameters of the utterance network, concept-selection network, and listening network respectively. All the modules of listener and speaker have to synchronize through trial and error for a successful communication language to emerge. In our setting, we perform decentralized learning with decentralized execution Foerster et al. (2016). Our agents are independent learners Tan (1993) and the channel between speaker and listener is non-differentiable, which implies that the back-propagation of the listener does not transmit the gradient backward to the speaker. In our 2D environment, there are N agents and M vertices. The state \mathcal{S} of the game set

consists of all relevant details that define the environment. The state of the environment at time t is given by $\mathbf{s}_t = \left[\mathbf{x}^{(1), \dots, (N)}, \mathbf{z}_t^{(1), \dots, (N)}, \mathbf{q}^{(1), \dots, (N)}, \mathbf{u}_t^{(1), \dots, (N)} \right]^\top \in \mathcal{S}$, where $\mathbf{x}^{(i)} \in \mathbb{R}^2$ is the location of the i^{th} vertex in the world, $\mathbf{z}^{(i)} \in \{1, 2, \dots, N\}$ is the current location of agent i , $\mathbf{q}^{(i)} \in \mathbb{R}$ is the color of vertex i and $\mathbf{u}_t^{(i)}$ is the utterance in the conversation involving agent i . The agent i locally perceives the environment which characterizes the observation vector of the speaker agent $\mathbf{o}_t^{(i)} \stackrel{f_i}{=} \left[\mathbf{z}_t^{(i)}, \mathbf{g}_t^{(i)}, \mathbf{u}_t^{(i)}, \mathbf{q}^{(\mathbf{g}_t^{(i)})}, \text{DNN}(\mathbf{d}^{(1), \dots, (M)} + \epsilon_d, \mathbf{w}^{(1), \dots, (M)} + \epsilon_w) \right]^\top$, where $\epsilon_d \sim \mathcal{N}(0, 1)$ and $\epsilon_w \sim \mathcal{N}(0, 1)$ are white Gaussian noises, $\mathbf{g}_t^{(i)} \in \{1, 2, \dots, M\}$ is the topic vertex, and \mathbf{d}, \mathbf{w} represent the distance and the angle of vertices from the speaker’s current vertex respectively. Here DNN represents a deep neural network which embeds the graph relative to the source vertex. The interaction pathway consists of multiple networks across the speaker and listener agents operating sequentially. The concept-selection network π_{ψ} operates in a one-to-many mode, where the initial hidden vector is obtained through a linear transformation of the observation vector \mathbf{o}_t , and the output is fed back as input. This network outputs the conception-selection bit-vector \mathbf{b}_t which is then passed through a differentiable channel to the speaking module π_{θ} (many-many mode) along with the spatial description \mathbf{c}_t of the topic vertex as $\mathbf{c}_t \odot \mathbf{b}_t$, where \odot is co-ordinate-wise vector product. The network utters the message \mathbf{m}_t which is transmitted to the listener through a non-differentiable (naive categorical sampling) noise-free channel. We use Gumbel-Softmax Jang et al. (2016); Maddison et al. (2016) based sampling to enable differentiability of concept-selection to utterance channel allowing gradients to flow through the sampling process. The listening module π_{ϕ} in the listener agent operates in a many-to-many mode, which means it processes the words in the generated message \mathbf{m}_t sequentially and generates a probability distribution $\pi_{\phi}(\cdot | \mathbf{m}_t)$ over the entire concept space \mathcal{C} . This distribution represents the agent’s interpretation of the message in terms of different concepts within the concept space. This distribution is further used to generate the listener interpretation \mathbf{c}'_t through categorical sampling. The complete architecture of the agents is depicted in Figure 4.

6 PERFORMANCE MEASURE

The objective function of our language game consists of three components: Regularized communication feedback, description length loss, and mirror loss. Each component plays a specific role in shaping the communication outcome and the characteristics of the emergent language.

6.1 REGULARIZED, GUIDED COMMUNICATION FEEDBACK

We consider the finite horizon cumulative reward with an entropy regularization term as given in Equation (4). The regularizer offers a few advantages that are conducive to language games. First, entropy regularization encourages exploration and helps prevent early convergence to sub-optimal policies. Second, the resulting policies can serve as a good initialization for fine-tuning to a more specific behavior. Third, the maximum entropy framework provides a better exploration mechanism for seeking out the best mode in a multimodal reward landscape. In the language game, we follow a stochastic, guided feedback mechanism. During a failed interaction, the speaker plausibly guides the listener by pointing out the topic vertex to the listener with a probability $\lambda \in [0, 1]$. This implies that the speaker may or may not provide effective guidance with a certain probability. The speaker and listener subsequently reinforce with respect to the spatial concept \mathbf{c}'_t corresponding to the plausibly communicated topic vertex. This implies that during the interaction between the speaker A and listener B , the interpreted concept \mathbf{c}'_t is taken as

$$\mathbf{c}'_t \sim \lambda \pi_{\phi_B}(\cdot | \mathbf{m}_t) + (1 - \lambda) \delta_{\mathbf{c}_t}, \text{ where } \lambda \in [0, 1] \text{ and for } E \subseteq \mathbb{R}^k, \delta_x(E) = \begin{cases} 1 & \text{if } x \in E, \\ 0 & \text{otherwise.} \end{cases}$$

Here δ_x is the Dirac measure at x which is a singular measure that places all its probability mass at the single point x . In the case of effective guidance, a full reward is associated with the interaction.

6.2 PRINCIPLE OF LEAST EFFORT

According to the principle of least effort Zipf (2016); Cancho & Solé (2003), language evolves as speakers of the language tend to simplify their speech in various ways to achieve a trade-off between understanding and effort. When choosing how to express themselves in a language, speakers take into account both their current and future communication requirements. This drives the speakers to consider linguistic constructs that are effective in meeting their communication goals and efficient in optimizing their labour. A similar hypothesis connecting the overarching fairness between cognitive load and language exposition is the principle of the economy of thought Mach (1898). It suggests

that the human mind, with its limited cognitive resources, seeks to represent the infinite complexities of the world in a way that is efficient and economical. Hence, we believe that languages tend to evolve in ways that promote the economy of least thought and linguistic effort, where the language users communicate using messages that are relatively easy to produce and comprehend. Hence, in the post-transient phase of language evolution, message length tends to decrease Futrell et al. (2015).

6.3 MIRROR NETWORKS

To enable interchangeability of the language, we consider mirror networks. A mirror neuron Di Pellegrino et al. (1992); Rizzolatti et al. (1996), strictly defined, is a type of neuron that is fired both when the individual executes certain actions and when it observes a strictly or broadly congruent set of actions. In our setting, we want the speaking, listening, and concept selection networks of an agent to be consistent with each other so that the information gained through the comprehension of the language is used for its reproduction and vice versa. Since our networks represent stochastic policies, by consistency, we mean in the Bayesian probabilistic sense (posterior distribution of the parameters conditioned on the message or concept). By ensuring consistency between these policies, you're seeking a coherent, bidirectional relationship between how the agent generates its responses (speaking) and how it interprets and understands incoming information (listening). This implies that the calibration pathway has to update and synchronize all the relevant networks in the direction of consistency. Hence, we consider the following mirror loss:

$$\mathbb{E} \left[\underbrace{\alpha_1 \mathcal{D}_{KL} \left(\pi_{\theta^A}(\cdot|\mathbf{m}) \parallel \pi_{\phi^A}(\cdot|\mathbf{m}) \right)}_{\text{speaker congruence}} + \underbrace{\alpha_2 \mathcal{D}_{KL} \left(\pi_{\phi^B}(\cdot|\mathbf{c}') \parallel \pi_{\theta^B}(\cdot|\mathbf{c}') \right)}_{\text{listener congruence}} + \alpha_3 \underbrace{\mathcal{D}_{KL} \left(\pi_{\phi^B}(\cdot|\mathbf{m}) \parallel \pi_{\psi^B}(\cdot|\mathbf{o}) \right)}_{\text{concept selection congruence}} \right], \text{ where } \mathcal{D}_{KL}(p_1 \parallel p_2) = \sum_x \log p_1(x) \frac{\log p_1(x)}{\log p_2(x)}$$

is the Kullback-Leibler divergence and $\alpha_{1...3} \geq 0$. The above loss quantifies the dissimilarity between the conditional probabilities of the mirror networks and the corresponding active networks. Minimizing the mirror loss during calibration implies making the mirror networks as similar as possible to the active networks resulting in bidirectional language use.

6.4 POVERTY OF STIMULUS

The poverty of stimulus appears in guided feedback scenarios, where the speaker reveals the topic vertex \mathbf{g} (which is the tangible component) to the listener at the end of a failed conversation. However, the conceptualization $\mathcal{C}_z(\mathbf{g})$ of the topic vertex g consists of more than one element, which makes the novel message \mathbf{m} of the conversation potentially ambiguous and uncertain. This uncertainty poses a challenge for the listener. To address the meaning uncertainty inherent in the poverty of stimulus situation due to insufficient information, the listener relies on contextual cues, where it distributes the message \mathbf{m} across all the possible conceptualizations $\mathcal{C}_z(\mathbf{g})$ of the topic vertex \mathbf{g} with respect to the source vertex \mathbf{z} and assign different normalized weights or probabilities w_b to each interpretation b based on some predispositions (the factors could be prior knowledge and context):

$$\log \pi_{\phi^B}(\mathbf{c}'|\mathbf{m}) = \frac{\sum_{b \in \mathcal{C}_z(\mathbf{g})} w_b \log \pi_{\phi^B}(b|\mathbf{m})}{\sum_{b \in \mathcal{C}_z(\mathbf{g})} w_b}, \text{ where } w_b \geq 0. \quad (2)$$

Over time, as the listener gains more exposure to the word and its usage in various contexts, the uncertainty decreases, and the listener becomes more adept at determining the intended meaning based on the context of its usage.

6.5 OBJECTIVE FUNCTION

The performance measure $J(\theta, \psi, \phi)$ of the language game is defined as follows: Let $\mathbb{E}_{\mathcal{I}}[\cdot]$ be the expectation induced by the *r.v.s.* $\mathbf{m} \sim \pi_{\theta^A}(\cdot|c)$, $\mathbf{c}' \sim \pi_{\phi^B}(\cdot|m)$, $\mathbf{s} \sim \mu$, $\mathbf{o} = f_A(\mathbf{s})$, $\mathbf{o} \rightarrow \mathbf{c}$ and $\mathbb{E}_{\mathcal{I}_t}[\cdot]$ be the expectation induced by the *r.v.s.* $\mathbf{m}_t \sim \pi_{\theta^A}(\cdot|\mathbf{c}_t)$, $\mathbf{b}_t \sim \pi_{\psi^A}(\cdot|\mathbf{o}_t)$, $\mathbf{c}'_t \sim \lambda \pi_{\phi^B}(\cdot|\mathbf{m}_t) + (1 - \lambda)\delta_{\mathbf{c}_t}$, $\mathbf{s}_t \sim \mu$, $\mathbf{o}_t = f_A(\mathbf{s}_t)$, $\mathbf{o}_t \rightarrow \mathbf{c}_t$.

Then $J(\theta, \phi, \psi) = \kappa_1 \mathcal{L}_1(\theta, \phi, \psi) + \kappa_2 \mathcal{L}_2(\theta, \phi, \psi) + \kappa_3 \mathcal{L}_3(\theta, \phi, \psi)$, where $\kappa_1, \kappa_2, \kappa_3 \geq 0$, (3)

$$\mathcal{L}_1(\theta, \phi, \psi) = \mathbb{E}_{\mathcal{I}_t} \left[\underbrace{\sum_{t=0}^{T-1} \mathbf{r}_t + \beta \mathcal{H}(\pi_{\theta^A}(\cdot|\mathbf{c}_t)) + \beta \mathcal{H}(\pi_{\phi^B}(\cdot|\mathbf{o}_t))}_{\text{Regularized cumulative reward}} \right], \beta \geq 0, \quad (4)$$

$$\mathcal{L}_2(\theta, \phi, \psi) = \underbrace{-\mathbb{E}_{\mathcal{I}} [\|\mathbf{b}\|_2^2 + \beta' \mathcal{H}(\pi_{\psi^A}(\cdot|\mathbf{s}))]}_{\text{Description length loss (Principle of least effort)}}, \beta' \geq 0 \quad (5)$$

$$\mathcal{L}_3(\theta, \phi, \psi) = \mathbb{E}_{\mathcal{I}} \left[\alpha_1 \mathcal{D}_{KL} \left(\pi_{\theta^A}(\cdot|\mathbf{m}) \parallel \pi_{\phi^A}(\cdot|\mathbf{m}) \right) + \alpha_2 \mathcal{D}_{KL} \left(\pi_{\phi^B}(\cdot|\mathbf{c}') \parallel \pi_{\theta^B}(\cdot|\mathbf{c}') \right) + \alpha_3 \mathcal{D}_{KL} \left(\pi_{\phi^B}(\cdot|\mathbf{m}) \parallel \pi_{\psi^B}(\cdot|\mathbf{o}) \right) \right] \text{ (Mirror loss)}, \quad (6)$$

with $\mathcal{H}(\pi(\cdot|s)) = -\sum_a \pi(a|s) \log \pi(a|s)$ is the entropy regularizer, where $\mathcal{H}(\pi(\cdot|s))$ represents the entropy of a policy π conditioned on state s . The entropy measures the uncertainty or randomness associated with the actions chosen by that policy when in a particular state. Further, we obtain the gradient of J as follows:

$$\nabla J(\theta, \phi, \psi) = \kappa_1 \nabla \mathcal{L}_1(\theta, \phi, \psi) + \kappa_2 \nabla \mathcal{L}_2(\theta, \phi, \psi) + \kappa_3 \nabla \mathcal{L}_3(\theta, \phi, \psi), \text{ where} \quad (7)$$

$$\nabla \mathcal{L}_1(\theta, \phi, \psi) = \mathbb{E}_{\mathcal{I}} \left[(Q_{\mathcal{I}}(\mathbf{s}, \mathbf{m}, \mathbf{b}, \mathbf{c}') - \beta \log \pi_{\theta^A}(\mathbf{m}|\mathbf{c}) - \beta \log \pi_{\phi^B}(\mathbf{c}'|\mathbf{m}) - \beta) (\nabla_{\theta^A} \log \pi_{\theta^A}(\mathbf{m}|\mathbf{c}) + \nabla_{\phi^B} \log \pi_{\phi^B}(\mathbf{c}'|\mathbf{m})) \right], \text{ where } Q_{\mathcal{I}}(s, m, b, c') = \mathbb{E}_{\mathcal{I}} \left[\sum_{t=0}^{T-1} \mathbf{r}_i^{(t)} | s, m, b, c' \right]. \quad (8)$$

$$\begin{aligned} \nabla \mathcal{L}_2(\theta, \phi, \psi) &= -\mathbb{E}_{\mathcal{I}} \left[(-\beta' \log \pi_{\psi^A}(\mathbf{b}|\mathbf{o}) - \beta') \nabla_{\psi^A} \log \pi_{\psi^A}(\mathbf{b}|\mathbf{o}) \right] \\ &\quad - \mathbb{E}_{\substack{s \sim \mu, \\ \mathbf{o} = f_A(s)}} \left[\nabla_{\psi^A} \mathbb{E}_{\mathbf{b} \sim \pi_{\psi^A}(\cdot|\mathbf{o})} [\|\mathbf{b}\|_2^2] \right] \text{ and} \\ \nabla \mathcal{L}_3(\theta, \phi, \psi) &= -\mathbb{E}_{\substack{\mathcal{I}, \\ \mathbf{c}' \rightarrow \mathbf{b}'}} \left[-\alpha_1 \frac{\mathbb{P}(\mathbf{c})}{\mathbb{P}(\mathbf{m})} \nabla_{\phi^A} \log \pi_{\phi^A}(\mathbf{c}|\mathbf{m}) - \alpha_2 \frac{\mathbb{P}(\mathbf{m})}{\mathbb{P}(\mathbf{c}')} \nabla_{\theta^B} \log \pi_{\theta^B}(\mathbf{m}|\mathbf{c}') \right. \\ &\quad \left. - \alpha_3 \frac{\mathbb{P}(\mathbf{b}')}{\mathbb{P}(\mathbf{o})} \nabla_{\psi^B} \log \pi_{\psi^B}(\mathbf{b}'|\mathbf{o}) \right], \quad (9) \end{aligned}$$

Here, Equation (8) is obtained by appealing to soft policy gradient theorem Shi et al. (2019) and multi-agent policy gradient theorem Zhang et al. (2018).

7 EXPERIMENTS & DISCUSSION

In all the experiments, model parameters are initialised with random values. The hyper-parameters (step size, batch size, and regularization strength) are fine-tuned through iterative experimentation. The feasible reward values for various scenarios are obtained through an exhaustive yet rational search. We consider a 2D world consisting of a complete graph with 5 vertices whose positions are randomly chosen. There are four agents in this world leading to a total of 12 unique speaker-listener pairs which likely allows for a rich variety of interactions and communication scenarios. Each dialogue consists of 100 conversations. The role switching (speaker to listener and vice versa) occurs after every 500 iterations through random selection. We consider two-timescale networks Chung et al. (2018) to obtain synchronized convergence, where the utterance network is calibrated along a faster timescale compared to the concept selection network. In this approach, the concept selection network can be considered to be pseudo-stationary, while the utterance network converges with respect to the stationary values of the concept selection network and this cycle repeats itself in the long run. To achieve this, we employ the vanilla stochastic gradient algorithm with learning rates of the respective networks differing by order of magnitude. This can be formalized as follows: Let $\{e_t\}$ and $\{e'_t\}$ be the learning rates of the concept selection network and utterance network respectively. Then $\{e_t\}$ and $\{e'_t\}$ satisfy the following:

$$e_t, e'_t \in (0, 1), \sum_{t \geq 0} e_t = \sum_{t \geq 0} e'_t = \infty, \sum_{t \geq 0} e_t^2 + e'^2_t < \infty, \lim_{t \rightarrow \infty} \frac{e_t}{e'_t} = 0. \quad (10)$$

The concept space \mathcal{C} consists of 4 sectors, 3 segments and 4 colors. The concept space \mathcal{C} is illustrated in Figure 6. Since there are overlapping sectors (Sectors 1, 3, and 4) we have a poverty of stimulus situation. The lexis size ($|\Psi|$) is 25. The speaking and listening module within the agent's architecture utilizes an LSTM cell. The observation vector o_t by the speaker agent is transformed into a feature vector ($\in \mathbb{R}^{25}$) by passing it through a fully connected neural network. This feature

vector forms the hidden input of the concept selection module (ψ) whose hidden size is also taken as 25. The speaking and listening modules are implemented as a single-layer LSTM cell with a hidden size 250. The LSTM networks output the sequence of words or concepts with a maximum length of 3. In the Gumbel softmax of the the concept selection module, the temperature parameter τ is taken as 0.5. Also, successful communication rewards both the agents with +100, while the misinterpretation incurs a penalty of -50.

We hypothesize the following based on the observations emerging from our guessing game:

Hypothesis I: *A measure of tangible guided feedback supplemented with reinforcement learning is required for a shared language to emerge among a population through decentralized learning.*

In our experiments, we observe that our decentralized guessing games are converging very often. The vocabulary-concept mappings developed by the individual agents during the transient phase are random which enables sufficient exploration to drive the evolution towards coherence in a finite number of dialogues. Coherence implies a more consistent and meaningful use of language, where words and expressions convey clear, grounded and shared meanings. This is corroborated by the convergence of loss functions and the maximization of average reward (average of the rewards of the conversations in a dialogue) as illustrated in Figures 13b and 16b. Positive rewards, attributed to successful communication events, are considered rare. When such events occur, positive rewards are reinforced by policy gradient-based calibration operations. This is further boosted by the guided feedback mechanism, where the listener uses the topic vertex information conveyed by the speaker to calibrate his parameters to match the conceptualization of the topic vertex (possibly ambiguous) with the utterance. These mechanisms help reinforce successful communication schemes. The success ratio is defined as the frequency of conversations in a dialogue, where the listener is able to identify the topic vertex. The evolution of this success ratio is depicted in Figure 16a over the course of dialogues across the population. This signifies that all the dialogue interactions in the conversation among the agent population are successful after a finite number of steps which suggests that the participants are able to achieve their communication objectives (identifying topic vertex) effectively through the medium of language, and this is independent of the nature of the agent executing the role of listener and speaker. As the process unfolds, it becomes increasingly apparent that all agents converge towards possessing highly similar vocabularies (Figure 2). The convergence towards near-identical vocabularies likely stems from the agents’ interactions and the need for effective communication. Through repeated interactions, agents gradually align their linguistic representations, leading to a shared, grounded lexicon.

Hypothesis II: *The locutors of a language exhibits an inclination towards minimizing cognitive effort, reflected in the emergent word order and Zipfian characteristic of the language.*

The agents involved in the language game have demonstrated a remarkable ability to grasp the essential concepts required for effective communication through the application of the principle of least effort. We observe that the order 101 which denotes $\langle \text{segment}, \perp, \text{color} \rangle$, is the emergent word order, suggesting a common consensus among the agents that for every topic vertex, communicating two concepts ($\text{segment}, \text{color}$) is adequate (Figure 14). The agents tend to communicate in a way that minimizes cognitive effort. The language tend to adapt to be as efficient as possible, with speakers and listeners preferring forms of expression that require the least amount of cognitive resources. This adaptation is driven by the principle of least effort component of the learning dynamics. To corroborate this claim further, we consider cases where the effect of principle of least effort is reduced significantly (Figure 8). For this purpose, consider the scenarios where κ_2 , representing the weight associated with loss \mathcal{L}_2 that quantifies the principle of least effort, takes on values of 1.1, 0.01, and 0.001 (with $\kappa_1 = 1$ and $\kappa_3 = 1$). For the case when $\kappa_2 = 1.1$, we have minimal conceptualization ($\langle \text{segment}, \perp, \text{color} \rangle$), of the topic vertex, however for the remaining cases, *i.e.*, $\kappa_2 \in \{0.01, 0.001\}$, we have maximal conceptualization ($\langle \text{segment}, \text{sector}, \text{color} \rangle$), of the topic vertex which signifies that the principle of least effort is fundamentally attributed to the cognitive economy exhibited by the agents, with higher weights for principle of least effort leading to simpler conceptualizations and lower values prompting more detailed ones. To understand the nature of principle of least effort with respect to the complexity of the word, we consider a setting with four agents having the same concept space as before, residing in a graph world of 50 vertices. This world is notably complex, presenting a challenge for the agents to fully comprehend using their predefined concepts. When it comes to conceptualizing the topic vertices, there is an intriguing trend. About 60% of the agents prefer to describe vertices using a combination of segment, sector, and color. On the other hand, the remaining 40% opt for a simpler approach, using segment,

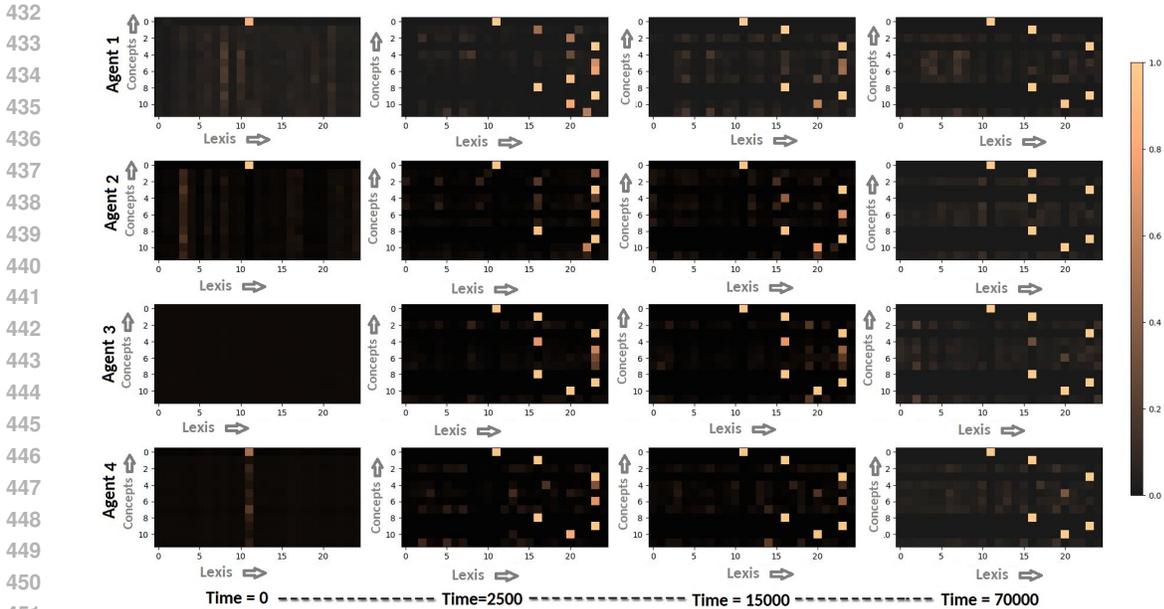


Figure 2: Convergence of shared vocabularies among the population is depicted here. Each rectangle represents the vocabulary of an agent, with each row (k , where $1 \leq k \leq 4$) showing the evolution of agent k 's vocabulary. It is notable that in the later stages of the process, all agents exhibit vocabularies that are nearly identical.

a placeholder for silence \perp , and color. This diversity in conceptualization approaches highlights the complexity of the world these agents are trying to understand. Despite their efforts to adapt and communicate effectively, the intricate nature of the graph world poses a significant challenge for the agents in establishing a shared understanding.

Hypothesis III *The language reflects the complexity of the environment it describes.*

Since the topic of conversation is a vertex that is conceptualized using sectors, segments, and colors, a few sectors and segments rarely appear in the conversations due to the distribution of vertices in the 2D world. They remain mostly dormant and hence no dominating concepts emerge for these less-discussed concepts. Contrary to the dormant concepts are the active words which are consistently used by the agents to express the topic vertices and hence are alive in the population. This phenomenon suggests that certain concepts may become dominant in communication due to their frequent usage, while others remain less prominent. The above bifurcation of the concept space is highly sensitive to the nature of the graph world and the distribution bias of the source-topic vertex pairs used during interactions. This is illustrated in Figure 16d, where we allude to concepts 4, 5, 6 and 7 which are unused and concept 1, which sees minimal usage ($\approx 1.0\%$). Since each color appears once in the world, except for yellow ($\Gamma(\text{yellow}) = 9$) which occurs twice and with the word order emerging as 101 ($\langle \text{segment}, \perp, \text{color} \rangle$), all the color concepts are active with yellow being more salient in communication. A shared word for expressing silence (NULL/ \perp) emerges and it is the most popular word. The dynamics of the game settle down to employing a single word to represent \perp concept irrespective of the position of \perp in the message, a phenomenon vividly illustrated in Figure 16c. As the communication game evolves, a consensus emerges among the agents to use a single word to represent the \perp concept. This transformation is additionally accompanied by a noteworthy reduction in the probabilities associated with other words, underscoring the agents' adaptability and the efficiency of their evolving communication system. There is a clear and consistent way of representing the absence of a specific concept in a message. Since 101 ($\langle \text{segment}, \perp, \text{color} \rangle$) is the globally accepted word order in the population, a \perp concept is always present in every conversation which makes its word the most popular word.

Hypothesis IV: *Mirror learning fosters continuity in learning, allowing participants to apply the knowledge they acquired in a specific role while engaging in the mirrored role.*

In our experiments, we observe near-complete synchronization of bidirectional mappings in the agent population through mirror networks. During the interactions, the individual agents calibrate

486 their corresponding mirror networks (listener network for the speaker and utterance network for the
 487 listener) *w.r.t.* their corresponding active networks (utterance network for the speaker and listener
 488 network for the listener) which ensures continuity during role switching. This is achieved by min-
 489 imizing the mirror loss which decays to 0 as illustrated in Figure 13b. The continuity in learning
 490 (comprehension and reproduction) is exemplified in Figure 16a, where one can observe the dips
 491 in success-ratio (due to role switching) eventually vanish ensuring continuity. The term "continual
 492 learning" suggests that, over time, the system improves at maintaining smooth flow of commu-
 493 nication as agents switch between speaker and listener roles, applying knowledge gained during
 494 comprehension to reproduction and vice versa, leading to effective bidirectional language use.

495 7.1 EMERGENT LANGUAGE CHARACTERISTICS

496 **Emergence of Compositionality:** The economy of thought gradually transpires into composition-
 497 ality in the emergent language since language encapsulates thought. Hence one could observe the
 498 emergence of efficient communication strategies, where shorter phrases are preferred over longer
 499 ones. In our setting, the language that emerged is two word compositional, where it follows a word
 500 order which involves `segment` followed by `color`. More complex scenarios can also arise. To
 501 illustrate compositionality more vividly, we consider a smaller graph with $N = 3$ vertices (Figure
 502 10) and two homogeneous agents with the concept space as before (Figure 6). The outcomes are
 503 depicted in Figures 10 and 11. Notably, no specific word order emerges in this scenario; instead, a
 504 blend of word combinations dominates the discourse. It is noteworthy that, in this setup, the number
 505 of available colors in the agents' concept space is 4, which is more than the number of vertices. This
 506 implies that the agents can communicate the topic vertices by just referring to the color alone. The
 507 same holds true for sectors. However, the observed trend reveals that in the majority of conversations
 508 (50%), the agents opt for the color alone ($\langle \perp, \perp, \text{color} \rangle$), and in 14% of instances, it relies on
 509 the sector alone ($\langle \text{segment}, \perp, \perp \rangle$). Interestingly, in (20%) of conversations, the agents prefer
 510 a combination of segment and color ($\langle \text{segment}, \perp, \text{color} \rangle$). The above mixture (dominated by
 511 an optimal concept selection) is a sub-optimal limiting behaviour with respect to concept selection
 512 and this scenario arises due to the non-convex nature of the objective function operating over a de-
 513 centralized setting, as shown by the convergence to sub-optimal values in Figure 10). However, the
 514 success ratio (Figure 10) Despite this, the success ratio is nearly 100% (Figure 10), and all agents
 515 have nearly identical vocabularies (Figure 12), indicating the emergence of an effective semi non-
 516 compositional language with 65% of the conversations being non-compositional and 24% two-word
 517 compositional. Similar sub-optimal behaviors can be expected in human scenarios. Nevertheless,
 518 what stands out is that the agents employ one, sometimes two, and rarely three words or stay silent
 during the dialogue, mirroring patterns observed in human interactions.

519 **Adherence to Zipf's Law:** Zipf's law is a linguistic phenomenon that describes the distribution of
 520 word frequencies within a language. It states that the frequency of any word is inversely proportional
 521 to its rank. In other words, the most common word occurs approximately twice as often as the second
 522 most common word, three times as often as the third most common word, and so on. We see in the
 523 Figure 15 (world setting from Figure 10), the occurrence of the word decreases exponentially for
 524 the lower rank words in the language emerged among agents. This aligns with the characteristic
 525 pattern described by Zipf's law. To further understand this relationship, we studied the impact
 526 of the principle of least effort on the Zipfian characteristic of emerged languages. Specifically,
 527 we looked at scenarios where the weight (represented by κ_2) associated with the loss function that
 528 quantifies the principle of least effort varied (Figure 8). When κ_2 had higher values (1.1 in one case),
 529 the discrepancy between the observed word usage frequency and the expected Zipf distribution is
 530 minimal. This indicates that the sensitivity of the Zipfian characteristic to the weight of the principle
 531 of least effort. In other words, higher weights lead to word frequencies that more closely follow an
 inverse relationship to their ranks, as predicted by Zipf's law.

532 8 CONCLUSION

533 In this paper, we develop a computational language game framework to model the factors influenc-
 534 ing language dynamics involving a finite number of homogeneous deep neural agents in a guessing
 535 game setting. We integrate the following: explicit ontology, guided feedback mechanism to over-
 536 come poverty of stimulus, mirror learning for interchangeability and factoring silence as a symbol.
 537 From our observations, we conclude the following: "*The emergence of a shared, grounded, struc-
 538 tured, cognitively local optimal and Zipfian language among a population arising from ambiguous
 539 interactions is attributed to diverse phenomena weaving together in the intricate tapestry of decen-
 tralized, guided and reinforced learning*".

REFERENCES

- 540
541
542 Jacob Andreas. Measuring compositionality in representation learning. In *International Conference*
543 *on Learning Representations*, 2018.
- 544 Robert L Arrington. Ludwig wittgenstein: Philosophical investigations. In *Central Works of Philos-*
545 *ophy v4*, pp. 257–279. Routledge, 1954.
- 546 Marco Baroni. Linguistic generalization and compositionality in modern artificial neural networks.
547 *Philosophical Transactions of the Royal Society B*, 375(1791):20190307, 2020.
- 548 John Batali. Computational simulations of the emergence of grammar. *Approaches to the Evolution*
549 *of Language-Social and Cognitive Bases-*, 1998.
- 550
551 Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella
552 Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, et al. Experience grounds lan-
553 guage. *arXiv preprint arXiv:2004.10151*, 2020.
- 554 Ramon Ferrer I Cancho and Ricard V Solé. Least effort and the origins of scaling in human language.
555 *Proceedings of the National Academy of Sciences*, 100(3):788–791, 2003.
- 556
557 Rahma Chaabouni, Eugene Kharitonov, Emmanuel Dupoux, and Marco Baroni. Anti-efficient en-
558 coding in emergent communication. *Advances in Neural Information Processing Systems*, 32,
559 2019.
- 560 Rahma Chaabouni, Eugene Kharitonov, Diane Bouchacourt, Emmanuel Dupoux, and Marco Baroni.
561 Compositionality and generalization in emergent languages. In *Proceedings of the 58th Annual*
562 *Meeting of the Association for Computational Linguistics*, pp. 4427–4442, Online, July 2020.
563 Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.407. URL <https://aclanthology.org/2020.acl-main.407>.
- 564
565 Rahma Chaabouni, Florian Strub, Florent Alché, Eugene Tarassov, Corentin Tallec, Elnaz Davoodi,
566 Kory Wallace Mathewson, Olivier Tieleman, Angeliki Lazaridou, and Bilal Piot. Emergent com-
567 munication at scale. In *International Conference on Learning Representations*, 2022.
- 568
569 E Choi, A Lazaridou, and N de Freitas. Compositional obverter communication learning from raw
570 visual input. *iclr 2018. arXiv preprint arXiv:1804.02341*, 2018.
- 571
572 Noam Chomsky. *Language and mind*. Cambridge University Press, 2006.
- 573
574 Wesley Chung, Somjit Nath, Ajin Joseph, and Martha White. Two-timescale networks for nonlinear
575 value function approximation. In *International conference on learning representations*, 2018.
- 576
577 Allan Dafoe, Edward Hughes, Yoram Bachrach, Tantum Collins, Kevin R McKee, Joel Z
578 Leibo, Kate Larson, and Thore Graepel. Open problems in cooperative ai. *arXiv preprint*
arXiv:2012.08630, 2020.
- 579
580 Gautier Dagan, Dieuwke Hupkes, and Elia Bruni. Co-evolution of language and agents in referential
581 games. *arXiv preprint arXiv:2001.03361*, 2020.
- 582 LEWIS David. *Convention: a philosophical study*, 1969.
- 583
584 Ferdinand de Saussure. *Course in general linguistics*. Columbia University Press, 2011.
- 585
586 Giuseppe Di Pellegrino, Luciano Fadiga, Leonardo Fogassi, Vittorio Gallese, and Giacomo Rizzo-
587 latti. Understanding motor events: a neurophysiological study. *Experimental brain research*, 91:
176–180, 1992.
- 588
589 Tom Eccles, Yoram Bachrach, Guy Lever, Angeliki Lazaridou, and Thore Graepel. Biases for
590 emergent communication in multi-agent reinforcement learning. *Advances in neural information*
591 *processing systems*, 32, 2019.
- 592
593 Jakob Foerster, Ioannis Alexandros Assael, Nando De Freitas, and Shimon Whiteson. Learning
to communicate with deep multi-agent reinforcement learning. *Advances in neural information*
processing systems, 29, 2016.

- 594 Richard Futrell, Kyle Mahowald, and Edward Gibson. Large-scale evidence of dependency length
595 minimization in 37 languages. *Proceedings of the National Academy of Sciences*, 112(33):10336–
596 10341, 2015.
- 597
- 598 Miquel Cornudella Gaya, Thierry Poibeau, and Remi van Trijp. The role of intrinsic motivation
599 in artificial language emergence: a case study on colour. In *26th International Conference on*
600 *Computational Linguistics (COLING 2016)*, pp. 1646–1656, 2016.
- 601
- 602 Laura Graesser, Kyunghyun Cho, and Douwe Kiela. Emergent linguistic phenomena in multi-agent
603 communication games. In *2019 Conference on Empirical Methods in Natural Language Process-*
604 *ing and 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP*
605 *2019*, pp. 3700–3710. Association for Computational Linguistics, 2019.
- 606
- 607 Nicola Guarino and Pierdaniele Giaretta. Ontologies and knowledge bases. *Towards very large*
608 *knowledge bases*, pp. 1–2, 1995.
- 609
- 610 Shubham Gupta, Rishi Hazra, and Ambedkar Dukkipati. Networked multi-agent reinforcement
611 learning with emergent communication. In *Proceedings of the 19th International Confer-*
612 *ence on Autonomous Agents and MultiAgent Systems, AAMAS '20*, pp. 1858–1860, Richland,
613 SC, 2020. International Foundation for Autonomous Agents and Multiagent Systems. ISBN
614 9781450375184.
- 615
- 616 Serhii Havrylov and Ivan Titov. Emergence of language with multi-agent games: Learning to com-
617 municate with sequences of symbols. *Advances in neural information processing systems*, 30,
618 2017.
- 619
- 620 Andrew Iliadis. The tower of babel problem: making data make sense with basic formal ontology.
621 *Online Information Review*, 43(6):1021–1045, 2019.
- 622
- 623 Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv*
624 *preprint arXiv:1611.01144*, 2016.
- 625
- 626 Ivana Kajić, Eser Aygün, and Doina Precup. Learning to cooperate: Emergent communication in
627 multi-agent navigation. *arXiv preprint arXiv:2004.01097*, 2020.
- 628
- 629 Eugene Kharitonov, Rahma Chaabouni, Diane Bouchacourt, and Marco Baroni. Entropy minimiza-
630 tion in emergent languages. In *International Conference on Machine Learning*, pp. 5220–5230.
631 PMLR, 2020.
- 632
- 633 Satwik Kottur, José Moura, Stefan Lee, and Dhruv Batra. Natural language does not emerge ‘nat-
634 urally’ in multi-agent dialog. In *Proceedings of the 2017 Conference on Empirical Methods in*
635 *Natural Language Processing*, pp. 2962–2967, 2017.
- 636
- 637 Angeliki Lazaridou and Marco Baroni. Emergent multi-agent communication in the deep learning
638 era. *arXiv preprint arXiv:2006.02419*, 2020.
- 639
- 640 Angeliki Lazaridou, Alexander Peysakhovich, and Marco Baroni. Multi-agent cooperation and the
641 emergence of (natural) language. In *5th International Conference on Learning Representations,*
642 *ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, 2017.
- 643
- 644 Angeliki Lazaridou, Karl Moritz Hermann, Karl Tuyls, and Stephen Clark. Emergence of linguistic
645 communication from referential games with symbolic and pixel input. In *International Confer-*
646 *ence on Learning Representations*, 2018.
- 647
- 648 Angeliki Lazaridou, Anna Potapenko, and Olivier Tieleman. Multi-agent communication meets
649 natural language: Synergies between functional and structural language learning. In *Proceedings*
650 *of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 7663–7674,
651 Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.
652 685. URL <https://aclanthology.org/2020.acl-main.685>.
- 653
- 654 Fushan Li and Michael Bowling. Ease-of-teaching and language structure from emergent commu-
655 nication. *Advances in neural information processing systems*, 32, 2019.

- 648 Toru Lin, Jacob Huh, Christopher Stauffer, Ser Nam Lim, and Phillip Isola. Learning to ground
649 multi-agent communication with autoencoders. *Advances in Neural Information Processing Sys-*
650 *tems*, 34:15230–15242, 2021.
- 651 Michael L Littman. Markov games as a framework for multi-agent reinforcement learning. In
652 *Machine learning proceedings 1994*, pp. 157–163. Elsevier, 1994.
- 653 Ernst Mach. *Popular scientific lectures*. Number 21. Open Court Publishing Company, 1898.
- 654 Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous
655 relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.
- 656 David M Mark, Werner Kuhn, Barry Smith, and Andrew G Turk. Ontology, natural language and
657 information systems: Implications of cross-linguistic studies of geographic terms. In *6th AGILE*
658 *conference on Geographic Information Science*, 2003.
- 659 Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wier-
660 stra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint*
661 *arXiv:1312.5602*, 2013.
- 662 Richard Montague et al. Universal grammar. 1974, pp. 222–46, 1970.
- 663 Igor Mordatch and Pieter Abbeel. Emergence of grounded compositional language in multi-agent
664 populations. In *Thirty-second AAAI conference on artificial intelligence*, 2018.
- 665 Ilya Prigogine. Exploring complexity. *European journal of operational research*, 30(2):97–103,
666 1987.
- 667 Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John
668 Wiley & Sons, 2014.
- 669 Shuwen Qiu, Sirui Xie, Lifeng Fan, Tao Gao, Jungseock Joo, Song-Chun Zhu, and Yixin Zhu.
670 Emergent graphical conventions in a visual communication game. *Advances in Neural Informa-*
671 *tion Processing Systems*, 35:13119–13131, 2022.
- 672 Willard Van Orman Quine. Word and object mit press. *Cambridge MA*, 1960.
- 673 Yi Ren, Shangmin Guo, Matthieu Labeau, Shay B Cohen, and Simon Kirby. Compositional lan-
674 guages emerge in a neural iterated learning model. *arXiv preprint arXiv:2002.01365*, 2020.
- 675 Cinjon Resnick, Abhinav Gupta, Jakob Foerster, Andrew M Dai, and Kyunghyun Cho. Capac-
676 ity, bandwidth, and compositionality in emergent language learning. In *Proceedings of the 19th*
677 *International Conference on Autonomous Agents and MultiAgent Systems*, pp. 1125–1133, 2020.
- 678 Mathieu Rita, Rahma Chaabouni, and Emmanuel Dupoux. “lazimpa”: Lazy and impatient neural
679 agents learn to communicate efficiently. In *Proceedings of the 24th Conference on Computational*
680 *Natural Language Learning*, pp. 335–343, 2020.
- 681 Mathieu Rita, Florian Strub, Jean-Bastien Grill, Olivier Pietquin, and Emmanuel Dupoux. On the
682 role of population heterogeneity in emergent communication. In *International Conference on*
683 *Learning Representations*, 2021.
- 684 Giacomo Rizzolatti, Luciano Fadiga, Vittorio Gallese, and Leonardo Fogassi. Premotor cortex and
685 the recognition of motor actions. *Cognitive brain research*, 3(2):131–141, 1996.
- 686 Wenjie Shi, Shiji Song, and Cheng Wu. Soft policy gradient method for maximum entropy deep
687 reinforcement learning. In *Proceedings of the 28th International Joint Conference on Artificial*
688 *Intelligence*, pp. 3425–3431, 2019.
- 689 BF Skinner. The evolution of verbal behavior. *Journal of the Experimental analysis of Behavior*, 45
690 (1):115, 1986.
- 691 Luc Steels. The synthetic modeling of language origins. *Evolution of communication*, 1(1):1–34,
692 1997.

702 Luc Steels. Evolving grounded communication for robots. *Trends in cognitive sciences*, 7(7):308–
703 312, 2003.

704

705 Luc Steels. Grounding language through evolutionary language games. In *Language grounding in*
706 *robots*, pp. 1–22. Springer, 2012.

707 Sainbayar Sukhbaatar, Rob Fergus, et al. Learning multiagent communication with backpropaga-
708 tion. *Advances in neural information processing systems*, 29, 2016.

709

710 Ming Tan. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings*
711 *of the tenth international conference on machine learning*, pp. 330–337, 1993.

712 Mycal Tucker, Roger Levy, Julie A Shah, and Noga Zaslavsky. Trading off utility, informativeness,
713 and complexity in emergent communication. *Advances in neural information processing systems*,
714 35:22214–22228, 2022.

715

716 Kyle Wagner, James A Reggia, Juan Uriagereka, and Gerald S Wilkinson. Progress in the simulation
717 of emergent communication and language. *Adaptive Behavior*, 11(1):37–69, 2003.

718 Sida I. Wang, Percy Liang, and Christopher D. Manning. Learning language games through interac-
719 tion. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*
720 *(Volume 1: Long Papers)*, pp. 2368–2378, Berlin, Germany, August 2016. Association for Com-
721 putational Linguistics. doi: 10.18653/v1/P16-1224. URL [https://aclanthology.org/
722 P16-1224](https://aclanthology.org/P16-1224).

723

724 Ludwig Wittgenstein. *Philosophical investigations*. John Wiley & Sons, 1954.

725 Kaiqing Zhang, Zhuoran Yang, Han Liu, Tong Zhang, and Tamer Basar. Fully decentralized multi-
726 agent reinforcement learning with networked agents. In *International Conference on Machine*
727 *Learning*, pp. 5872–5881. PMLR, 2018.

728

729 George Kingsley Zipf. *Human behavior and the principle of least effort: An introduction to human*
730 *ecology*. Ravenio Books, 2016.

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

A APPENDIX

A.1 DEFINITIONS

Emergent vocabulary: An emergent vocabulary \mathcal{M} is a shared mapping (not necessarily bijective) function between lexis Ψ and the concept space \mathcal{C} , *i.e.*, $\mathcal{M} : \Psi \leftrightarrow \mathcal{C}$ collectively agreed upon by all the agents in the population de Saussure (2011). Note that there are $|\mathcal{C}|^{|\Psi|}$ possible vocabularies for all the agents to agree upon, which makes it unlikely for all agents to converge on the same vocabulary without some mechanism for coordination and consensus.

Compositionality: A language is compositional Andreas (2018); Montague et al. (1970) if the utterance of each derived concept is determined by the utterances of its basis concepts. Formally, for the derived concept $c = \langle g, h, q \rangle$, we have $\mathcal{M}(c) = \mathcal{M}(g)\mathcal{M}(h)\mathcal{M}(q)$, with the implicit operation of concatenation connecting them.

A.2 GUESSING GAME SETUP

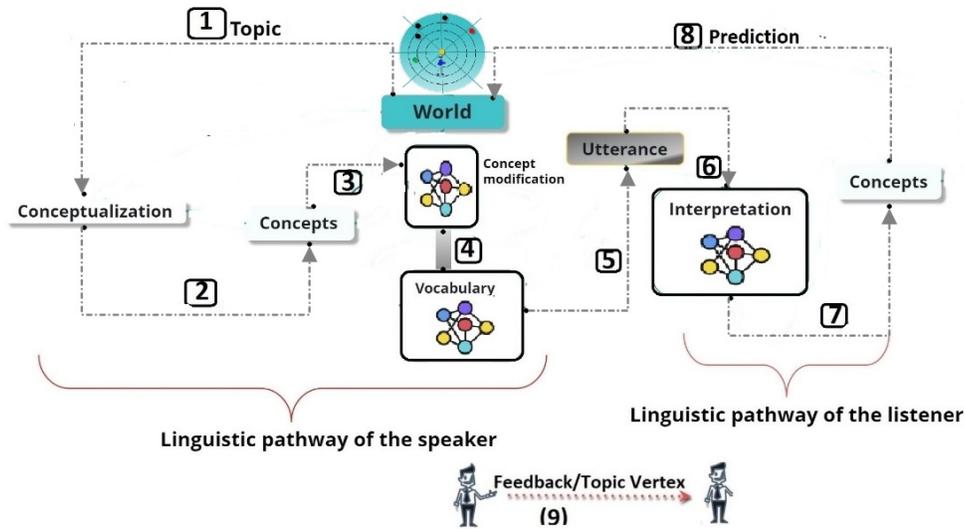


Figure 3: Semiotic pathway illustration of the guessing game

encourages agents to refine their communication strategies over time.

$$\mathbf{r}_t = \begin{cases} \zeta_1 (\in \mathbb{R}), & \text{if } g_t = g'_t, \\ \zeta_2 (\in \mathbb{R} \wedge \zeta_2 < \zeta_1), & \text{if } \mathcal{C}_{z_t}(g_t) \cap \mathcal{C}_{z_t}(g'_t) \neq \emptyset, \\ \zeta_3 (\zeta_3 \leq 0 \wedge \zeta_3 < \zeta_2), & \text{otherwise.} \end{cases} \quad (11)$$

The concept-selection module of the speaker seeks to select the optional spatial description to refer to the topic vertex by deactivating redundant concepts. The mechanism aims to ensure that the sentence corresponding to the generated spatial description is of optimal length to convey the intended meaning effectively. To support optimal word-order selection, we penalize the speaker for choosing a sub-optimal sequence of concepts. In cases where a concept is de-activated, the agent chooses to remain silent at that particular instant of the corresponding generated message. To enable this, the utterance module chooses a “NULL” utterance $\mathcal{C}(\perp)$ to indicate silence. The concept of \perp utterance is significant since we do not explicitly impose it a priori, rather it is learned through interactions. In order to promote consistency and coherence in the use of the $\mathcal{C}(\perp)$ utterance across different word categories in a sentence, we employ a strategy to positively reward \mathbf{r}' the speaker for the reuse of the same word for the \perp irrespective of its temporal position in a sentence. This reward system encourages the emergence of a common word for the \perp across different contexts, regardless of its temporal position in the message \mathbf{u}_t .

$$\mathbf{r}'_t = \begin{cases} \zeta'_1 (\in \mathbb{R}), & \text{if } |\{\mathcal{M}(a) | a \in \mathbf{u}_t \wedge a = \perp\}| = 1, \\ \zeta'_2 (\zeta'_2 < \zeta'_1), & \text{otherwise.} \end{cases} \quad (12)$$

B.2 EXPERIMENTS AND FURTHER DISCUSSION

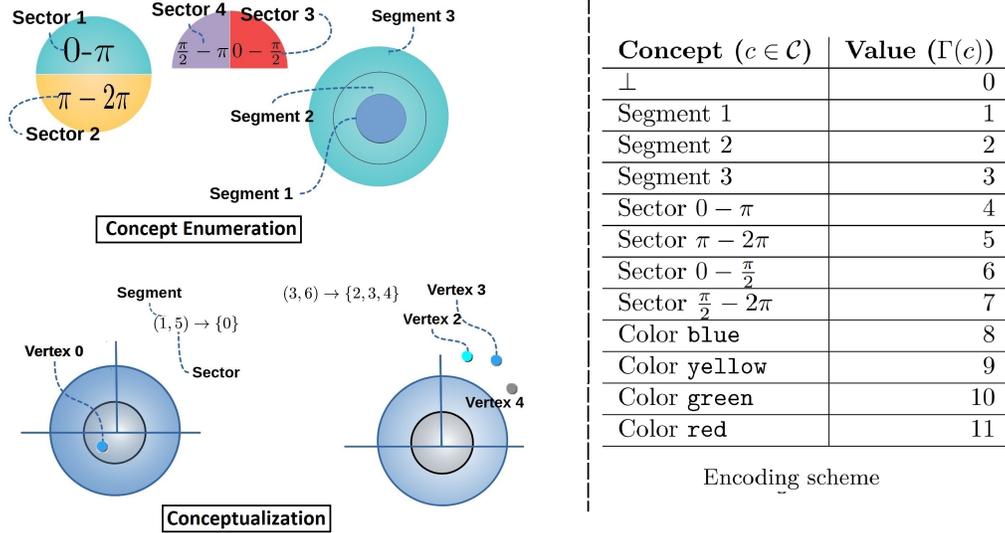
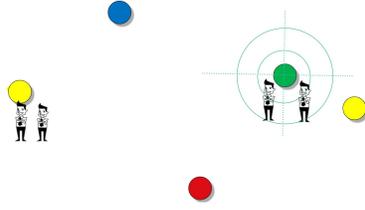


Figure 6: Concept space and the corresponding conceptualization of vertices

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971



Source	Conceptualization
0	$(3, 5) \rightarrow \{1, 2, 3\}, (3, \perp) \rightarrow \{1, 2, 3\}, (\perp, 5) \rightarrow \{1, 2, 3\}, (2, 5) \rightarrow \{4\}, (2, \perp) \rightarrow \{4\}$
2	$(3, 7) \rightarrow \{0, 3, 4\}, (3, \perp) \rightarrow \{0, 3, 4\}, (\perp, 7) \rightarrow \{1, 3, 4, 0\}, (3, 4) \rightarrow \{0, 3, 4\}, (\perp, 4) \rightarrow \{1, 3, 4, 0\}, (2, 7) \rightarrow \{1\}, (2, \perp) \rightarrow \{1\}, (2, 4) \rightarrow \{1\}$
1	$(3, 4) \rightarrow \{0\}, (3, \perp) \rightarrow \{0\}, (\perp, 4) \rightarrow \{3, 4, 0\}, (3, 7) \rightarrow \{0\}, (\perp, 7) \rightarrow \{4, 0\}, (2, 5) \rightarrow \{2\}, (2, \perp) \rightarrow \{4, 2, 3\}, (\perp, 5) \rightarrow \{2\}, (2, 4) \rightarrow \{3, 4\}, (2, 6) \rightarrow \{3\}, (0, 6) \rightarrow \{3\}, (2, 7) \rightarrow \{4\}$
3	$(3, 7) \rightarrow \{0\}, (3, \perp) \rightarrow \{2, 0\}, (\perp, 7) \rightarrow \{0\}, (3, 4) \rightarrow \{0\}, (\perp, 4) \rightarrow \{0\}, (2, 5) \rightarrow \{1, 4\}, (2, \perp) \rightarrow \{1, 4\}, (\perp, 5) \rightarrow \{2, 4, 1\}, (3, 5) \rightarrow \{2\}$
4	$(2, 7) \rightarrow \{0\}, (2, \perp) \rightarrow \{3, 0, 1\}, (\perp, 7) \rightarrow \{0\}, (2, 4) \rightarrow \{0, 3\}, (\perp, 4) \rightarrow \{0, 3\}, (2, 5) \rightarrow \{1\}, (\perp, 5) \rightarrow \{2, 1\}, (3, 5) \rightarrow \{2\}, (3, \perp) \rightarrow \{2\}, (2, 6) \rightarrow \{3\}, (0, 6) \rightarrow \{3\}$

Figure 7: The conceptualization of vertices with respect to each source vertex (referred to as source in the table). The syntax followed in the table is as follows: For the source vertex $v \in V$ (source column), the conceptualization column contains $(a, b) \in \Gamma(\mathcal{H}) \times \Gamma(\mathcal{W}) \rightarrow \{u \in V | (a, b) \in \mathcal{C}_v(u)\}$

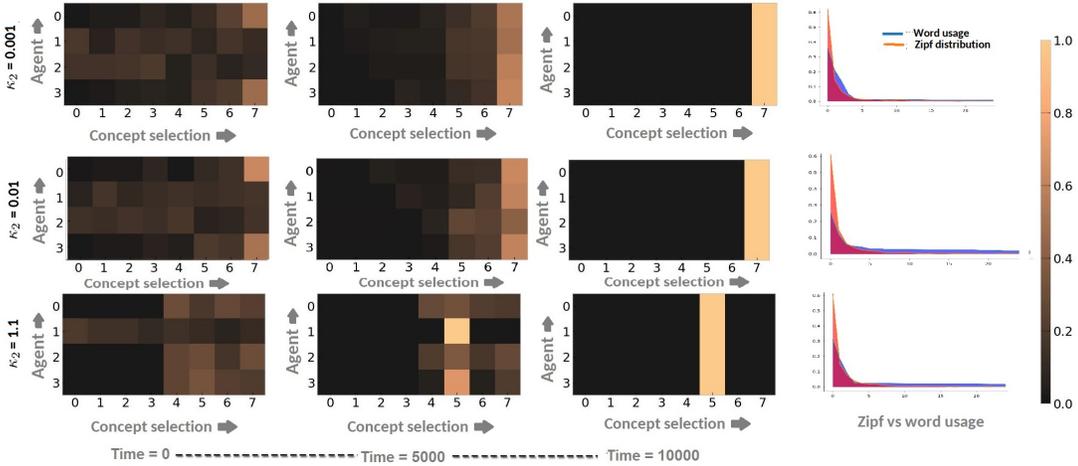


Figure 8: Sensitivity of principle of least effort on the cognitive economy

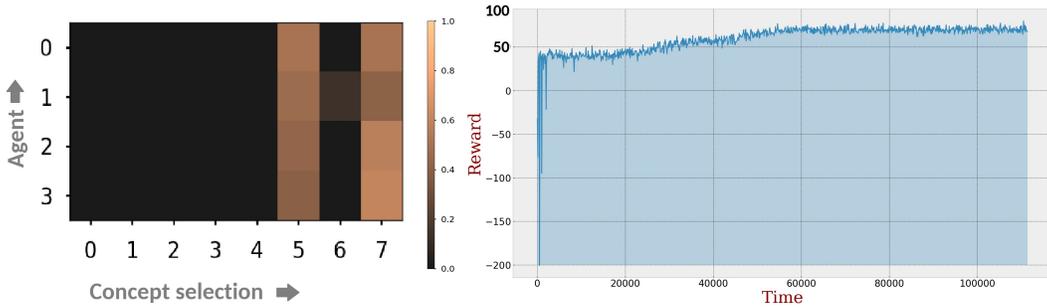


Figure 9: In this setting, we have 4 agents with the same static concept space as before residing a graph world with 50 vertices. Note that the world is too complex for the agents to perceive uniquely using their pre-defined concept space. The conceptualization of the topic vertices seems to use as much concepts as possible (60% preferring $\langle \text{segment}, \text{sector}, \text{color} \rangle$, while 40% preferring $\langle \text{segment}, \perp, \text{color} \rangle$)

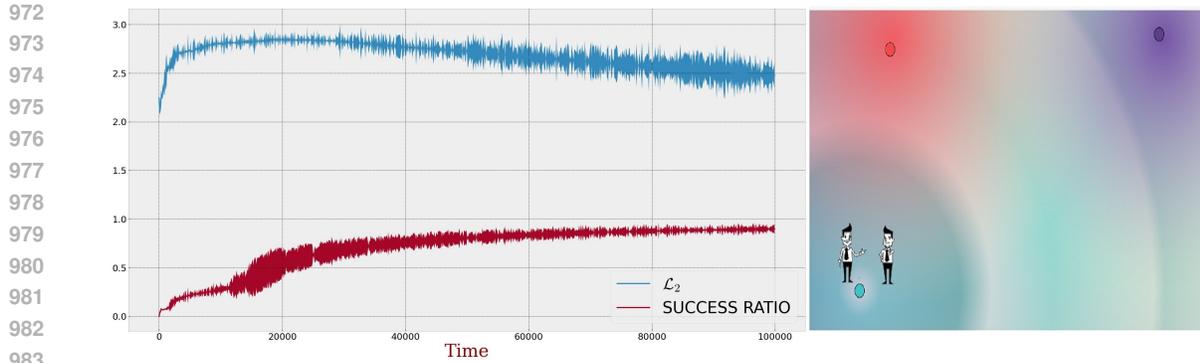


Figure 10: Success ratio for the setting with $N = 3$ vertices and 2 homogeneous agents and the same concept space as in the previous setting

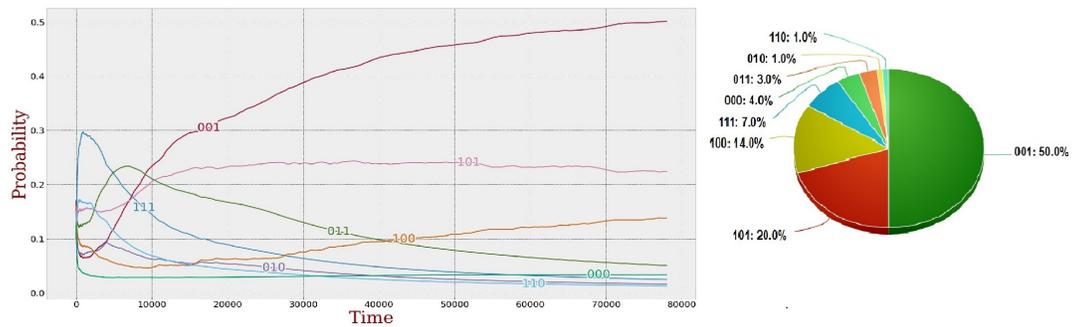


Figure 11: Emergence of concept selection pattern for the setting with $N = 3$ vertices and 2 homogeneous agents with the same concept space as in the previous setting. The plot illustrates the frequency ratio of the chosen concept selection during the dialogue utterances. The continuity in the trajectories indicates that the distribution over the concept selection space remains consistent across the agents.

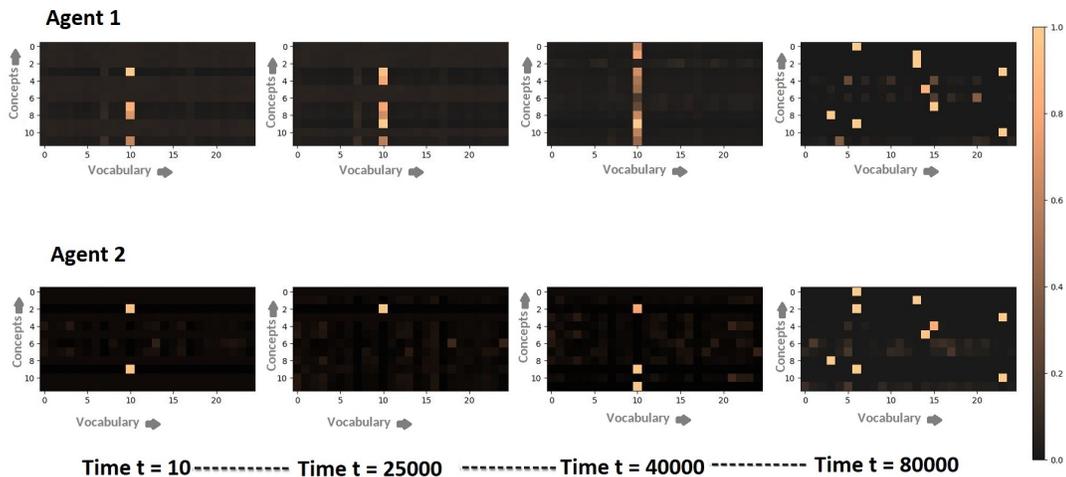


Figure 12: Success ratio for the setting with $N = 3$ vertices and 2 homogeneous agents and the same concept space as before

B.2.1 EMERGENCE OF DOMINATING WORDS

From Figure 2, one can observe that for each active concept ($\{0, 1, 2, 3, 8, 9, 10, 11\}$ from Figure 16d) there exists a dominating word (with probability one) except concept 11 (color red). In other

1026
1027
1028
1029
1030
1031
1032
1033
1034
1035

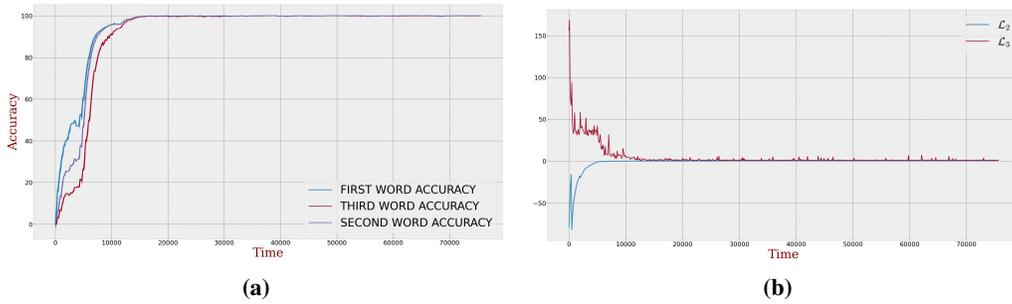


Figure 13: (a) First, second and third word accuracy: These metrics measure the frequency of interactions in a dialogue where the listener correctly interprets the first, second, and third words of a sentence, respectively across the population. (b) The evolution of the mirror loss and description length loss across the event population with respect to time

1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057

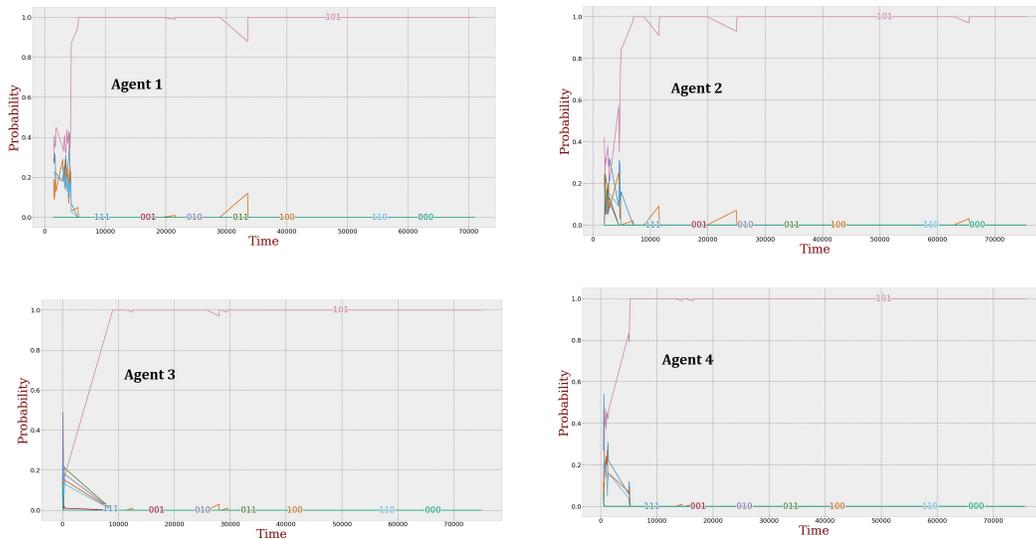


Figure 14: Cognitive economy: Evolution of the common word order among the population of agents. Here y -axis is the probability of choosing the specific word order. All the agents converge to a common word order 101 which represents $\langle \text{segment}, \perp, \text{color} \rangle$

1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072

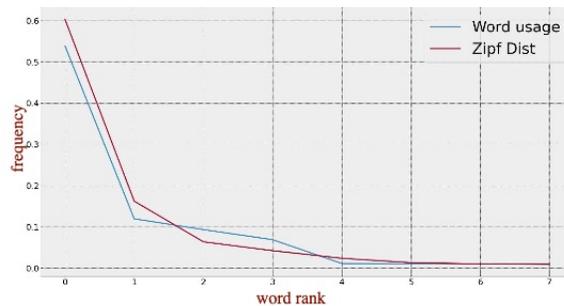


Figure 15: The usage of a word reduces in accordance with their rank supporting the Zipf law.

1073
1074
1075
1076
1077
1078
1079

words, for each active concept being discussed (except red), there is a specific word that is always used when referring to that concept. Further exploring the emergent language, one can observe certain intriguing patterns. For example, the concepts 9 (yellow) and 3 (segment 3) are represented by the same word 23 (See Figure 2). Since yellow is categorized as color and segment 3 is categorized as segment, the two concepts take different positions in the conceptualization, and

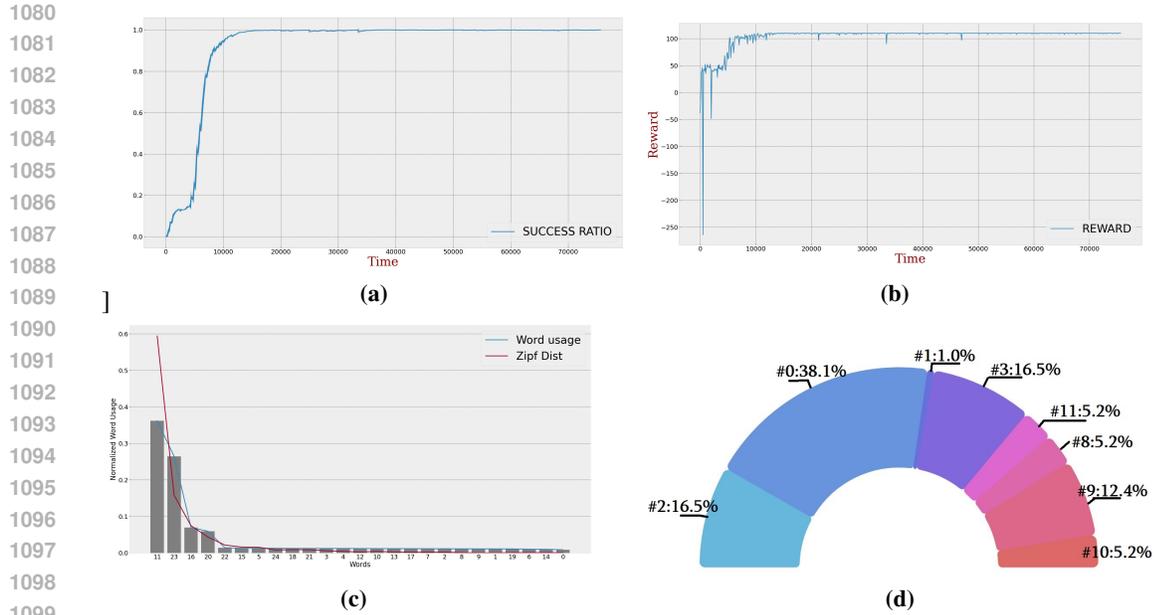


Figure 16: (a) Evolution of communication success ratio during dialogues over time (b) Trend of average reward over interactions within a dialogue over time (c) Word usage across all the dialogues over time among the agent population (d) Concept usage across all the dialogues over time among the agent population.

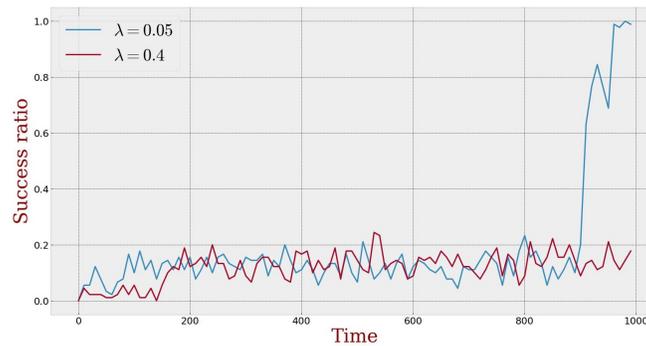


Figure 17: The plot illustrates how the language evolves and stabilizes in relation to the guidance feedback weight parameter λ . For low values of λ , which represents a greater emphasis on feedback during the learning process, we observe a convergence in the language’s structure and usage among the agents. This stabilization indicates that a higher feedback weight facilitates more effective communication and understanding.

hence the presence of the same word at different positions of the message generates different concepts making them unambiguous. Thus the word 23 is a homonym. A similar observation applies to another homonym, word 16, which is associated with concepts 1 (segment 1) and 8 (blue). It seems that for the red vertex, segment information is enough to deduce the vertex and hence color red is not learnt. Hence one can observe (Figure 16c), the emergence of four dominating words from the lexis ($|\Psi| = 25$) is mapped to the concept space ($|\mathcal{C}| = 12$). The convergence towards an unambiguous dominating word for each concept could be beneficial in fostering effective communication, as it reduces ambiguity and promotes a shared understanding of the intended meanings behind specific terms.

B.2.2 LANGUAGE EMERGENCE AT SCALE

The agent population is upscaled to observe the emergent behaviour among large populations. We consider complex settings with the number of different agent pairs equal to 12, 30 and 42. The emergence of a shared language among a larger population is cumbersome which requires a large

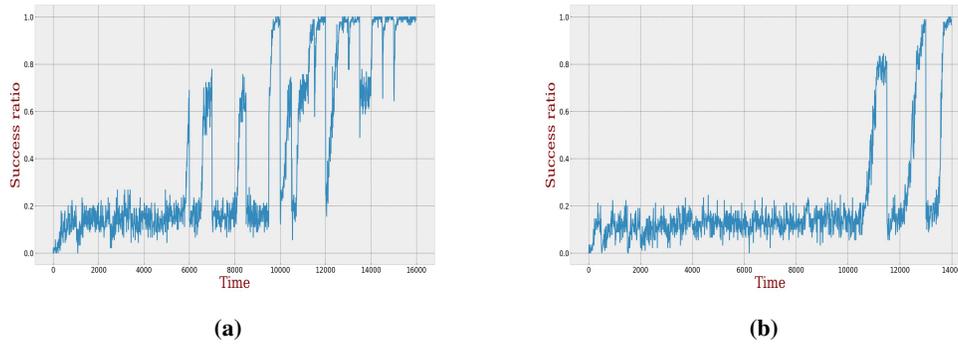


Figure 18: (a) Plot of success ratio with 30 pairs of agents and 5 vertices. (b) Plot of success ratio with 42 pairs of agents and 5 vertices

number of iterations. The underlying policy gradient algorithm develops coherence by reinforcing successful interactions. However, in the case of more agent pairs the probability of propagation of mappings involving successful interaction among the population is minimal. This behaviour is illustrated in Figures 18. Hence at scale, language emergence among larger populations proves challenging, necessitating numerous iterations to propagate successful interactions and generate shared understanding.