"What do you mean by that?"

Idiolects, Casual Miscommunication, and the Evolutionary Fitness of Languages

Alexander (Sasha) Boguraev
Class of 2024

Advisors:
Dr. Morten Christiansen, Cornell University, Department of Psychology
Dr. Eugene Vinitsky, NYU Tandon, Department of Civil & Urban Engineering
Dr. Veit Elser, Cornell University, Department of Physics

submitted in partial fulfillment
of the requirements for a
degree of B.A. with honors in
College Scholar, *Computer Science*, and *Mathematics*

Cornell University
Ithaca, New York
May 1, 2024

Abstract

Traditionally, human languages are perceived as a single cohesive set of words — a vocabulary — and grammatical rules. In reality, every individual who speaks a given language has their own personal understanding of it — their idiolect. In this work, we use Multi-Agent Reinforcement Learning to train four populations of agents, each possessing different combinations of idiolectic traits. We then compare their performance on a variety of tasks, including in-distribution performance, zero-shot generalization to new environments, zero-shot generalization to new agents, and robustness to external noise. Our findings indicate that the incomplete communication generated by idiolects may aid languages in not over-fitting to situations encountered during evolution; they further suggest that the existence of unique language representations for each individual within a population, their idiolects, likely aids in providing robustness to external communication failures.

Contents

List of Figures

List of Tables

1 Introduction

Language is an extraordinarily complex and uniquely human trait. Due to the inherent entanglement of language and cognition, many fields including psychology and cognitive science see it as a window to study various human processes. However, there exist certain aspects of human cognition which can only be readily observed when studies consider broader timescales, such as cognitive development (Chater and Christiansen, 2010). In these cases, one benefits not only from studying language at any given point in time, but its evolution, and how this evolution has been shaped by human cognition. It is this perspective with which we present our work, gaining deeper understandings of human cognition by studying what exactly drives a specific aspect of language evolution.

1.1 Idiolects and 'Good-Enough' Communication

When considering language, one usually thinks in terms of a single cohesive set of words — a vocabulary — and a single cohesive set of grammatical rules. These rules enable the construction of more cohesive, compositional elements that possess a depth of meaning larger than the words themselves. However, in thinking of language this way, we lose sight of a critical part of language: that every individual's understanding of it is uniquely different.

No two individuals who speak the same language have the exact same understanding of it. From each person's unique vocabulary to their grammar and even their own pronunciation, all who employ a given language have a distinct way of using it – their personal idiolect. This term commonly denotes the unique speech variety, or linguistic system, used by a particular individual. Idiolects are shaped by the personal experiences and cultural influences within the life of a language user. It follows that instead of viewing language merely as a set vocabulary and the grammatical rules that allow compositionality from that vocabulary, perhaps there should be a paradigm shift to view language from a more idiolectic perspective: as an amalgamation of the idiolects of all its users.

It may seem counter-intuitive that every speaker in a given population possesses their own individual understanding of a language, as this may potentially lead to difficulties in communication. This logic follows a common assumption: that both the production and comprehension of language between communicating agents is complete. By complete, we mean that a) the message that the speaker intends to convey is completely captured in their utterance and b) this is then completely deciphered and understood by the receiver. However, this assumption is faulty. Within human languages, inter-agent communication operates on a framework of incomplete production and comprehension. Despite this lack of completion, language is largely able to convey a 'good-enough' approximation of the intended meaning between individuals. In these cases, where the conveyed meaning is not total, but not completely lost, the incomplete production and comprehension are referred to as 'good-enough' production (Goldberg and Ferreira, 2022) and 'good-enough' comprehension (Ferreira and Patson, 2007).

"Good-enough" communication is driven in part by the 'Now-or-Never Bottleneck' (Christiansen and Chater, 2016). This phenomena consists of fundamental limitations in the sensory and cognitive memory function of the brain, and forces the immediate processing of linguistic input. If a language user does not immediately process linguistic input, it will not be recalled, and thus the information contained therin is lost forever. Immediate comprehension is driven by a process termed as "chunk and pass", in which the brain uses predictive mechanisms to ensure that any ambiguities in a received message are resolved during the first pass. Once a first pass is completed, the original input is lost and cannot be recovered. As such, this bottleneck "implies that the processing system will build the most abstract and complete representation that is justified, given the linguistic input", but it does not guarantee perfectly complete communication (Christiansen and Chater, 2016). The bottleneck does, however, select for the linguistic structures which are easiest to process, whilst fitting within its constraints.

The combination of a lack of complete communication with the selection for linguistic structures which are easiest to process results in 'good enough' communication, i.e. communication which is not necessarily complete, but through the process of "chunk and pass", good enough to be understood.

There further exist many roadblocks to complete communication between agents on both the speaker and receiver's end. For example, a speaker may not have the necessary vocabulary to optimally express their message, or a receiver may misinterpret the phrases they receive – whether through an inability to parse, a predisposition to different interpretations, or another cause. There are an endless number of ways in which communication with human language as a medium may not be complete. It is this lack of complete encoding and decoding of messages that is encapsulated by these aforementioned concepts of good-enough production and comprehension.

More roadblocks in perfect communication stem from language being more than a system consisting of words, phrases and sentences. These are just the tip of the 'communication iceberg', with aspects such as factual knowledge, unspoken rules, conventions, empathy, and more, all existing below the surface, but underpinning language (Christiansen and Chater, 2022). Language is not a system with which the meaning of a thought can be neatly packaged and delivered from the sender to the receiver; instead, language is a charade-like method in which there are a vast many cues which can help signal meaning between communicating agents.

So clearly, in language, universal understanding is by no means necessary, or even plausible, for successful communication to occur. Merely all that is needed is good-enough production and comprehension.

1.2 Language Evolution through the Eyes of Biological Evolution

There is a wealth of literature attempting to understand language evolution through a comparison to its biological equivalent. In this manner, cognitive scientists have viewed language itself as akin to a biological organism – a self-organizing adaptive system – which evolves to adapt to the brain, whilst subjected to the pressures of human learning and processing mechanisms (Christiansen and Chater, 2008). This way of thinking leads directly to applying evolutionary theories, such as natural selection, to explain the emergence and subsequent persistence of many core linguistic traits. It follows that many linguistic traits confer some sort of evolutionary advantage upon a language that allows for its survival, and subsequent reproduction (in a language's case, new users).

With this perspective in mind, the ubiquity of idiolectic languages suggests an evolutionary benefit in a language being non-universal, possessing variance in the understanding between its many users. However, there has been little to no work done to understand what evolutionary advantage idiolectic languages possess. This study attempts to fill that gap and duly provide insights into what advantages idiolectic languages hold over their universal counterparts.

One such theory we present is best viewed via the aforementioned link between biological and linguistic evolution. In biological natural selection, traits that give an evolutionary advantage to those who possess them get artificially selected for reproduction, and show up with higher prevalence in the next generation of individuals. However, there must be a fuel for these new traits to arise. In biological evolution, this fuel is random genetic mutation, leading to unforeseen variations and the emergence of new traits. It follows that perhaps idiolectic diversity is exactly the fuel that provides new linguistic traits which can be selected for during language evolution. If this were the case, we would expect that while a universal language may perform better at tasks in environments similar to the ones trained upon, a language with broad idiolectic diversity will likely perform better on, or at least be better equipped to adapt to, tasks involving novel environments. Similarly, we would expect language systems with this diversity between its speakers to be able to learn new concepts much more efficiently, due to the exploration inherent to the use of these given language systems. In contrast, a universal language, with no inherent exploration, would likely find it much more difficult to learn new concepts.

Note that, in this study, we explicitly stray away from studying exactly how idiolectic languages developed, leaving this for future work, because while a fascinating question, it is much more difficult to model. Further, while we have discussed theoretical benefits idiolects may lend to a language's evolution, we do not measure idiolect's affects on a language's speed of adaptation to new settings. This decision was dictated by our chosen methods, and their associated computational limitations – which we discuss at length in Section 4.4. Of course, none of what is discussed here precludes the prospect that this preliminary study may provide results or methods that inform such future studies.

1.3 Reinforcement Learning and Language

There exist a host of papers in the field of reinforcement learning (RL) which apply the field's frameworks towards the development of models with capacities to understand and use language in a humanoid manner. At its core, reinforcement learning utilizes an agent housed in an environment. This agent is then provided tasks and goals, duly being rewarded for completing a given task, and penalized for failing to complete said task. While agents' behaviors begin as completely random in nature, slowly their good behavior (in relation to the tasks at hand) are biased for (i.e. rewarded), and their bad behavior biased against (i.e. punished). Models trained using RL to aid language development are typically trained in population-based settings, with multiple agents. This sub-field of RL is colloquially referred to as multi-agent reinforcement learning (MARL), in which multiple agents, or models, interact to accomplish their goals.

Studies to date have simulated the emergence of languages between agents which have been trained on cooperative tasks, such as image identification (Lazaridou et al., 2016), prisoner's dilemma style problems (Foerster, 2018), basic directing of movement for other agents (Mordatch and Abbeel, 2018), and others. While it must be noted that many of the resulting emergent languages are either extremely rudimentary, completely symbolic in nature, or both, there have been a surprising number of studies in which the language that emerges is compositional in nature (Mordatch and Abbeel, 2018), or where agents manage to develop seemingly natural language (Lazaridou et al., 2016).

In this study, we train a host of populations which simulate populations of agents with, and without, various traits we construct to simulate idiolects. We then proceed to evaluate the performance of these populations on four different tasks. We find no performance difference when evaluating populations in environments which necessitate zero-shot generalization to new colors, and paradigms which necessitate communication with new agents from a separately trained population. Notably, we do find performance differences between populations with and without idiolects when performing in environments drawn from the training distribution, and on tasks which inject variable amounts of external noise. We view these results as evidence that idiolects may help with the development of robust languages in human populations, which are well-fit to the environments and tasks at hand.

2 Emergent Communication Development

Inspired by work from Mordatch and Abbeel 2018, and subsequently Lowe et al. (2017), our evolutionary environments are formulated as a partially observable Markov game (Littman, 1994). By definition of a Markov game, this consists of a set, $\mathcal{S}$, of all configurations of agents and landmarks in our game, a set of actions, $\mathcal{A}_1 \dots \mathcal{A}_N$, for each agent, and a set of observations, $\mathcal{O}_1 \dots \mathcal{O}_N$, for each agent. Furthermore, state transitions are defined by a function $\mathcal{T}$: $\mathcal{S} \times \mathcal{A}_1 \times \dots \times \mathcal{A}_N \rightarrow \mathcal{S}$. Each agent, i, in a given population makes an action, physical and communicative, as dictated by a learned function, $\mu_{\theta_i} : \mathcal{S} \rightarrow \mathcal{A}$. Further, at every time-step, each agent makes observations with function $o_i : \mathcal{S} \rightarrow \mathcal{O}$, and receives reward $r_i : \mathcal{S} \times \mathcal{A}_i \rightarrow \mathcal{R}$.

2.1 Environmental Setup

Our environment, as seen in Figure 1, consists of two agents and three landmarks. Each agent and landmark, i, is randomly initialized with position, $p_i = [p_{i,x}, p_{i,y}]$, where $p_{i,x}, p_{i,y} \in [-1, 1]$. During training, one landmark is assigned to be red, another green, and the other one blue. Our agents also possess a goal vector, g_i, consisting of the landmark to which they would like the other agent to move. These goals are private and not observed by the other agent. Associated with each agent is a private "reference frame", consisting of a randomly initialized rotation matrix, by which all of their positional observations are rotated. This rotation ensures that developed communications focus on the traits of the environment, and not merely relative positions of agents and landmarks in the environments.

At every time point, t, an agent takes a step in a certain direction, with their action represented by u_{i_t}, and their new position being calculated by $p_{i_{t+1}} = p_{i_t} + u_{i_t}$. Further, at each time t, both agents will utter a communication vector, c_{i_t}, which is a 10-dimensional vector, observable by the other agent. Over time, our agents develop the ability to convey to their partner agent how to complete their goal, by utilizing this communication vector.

The state of a given environment at any time, t, can be represented by

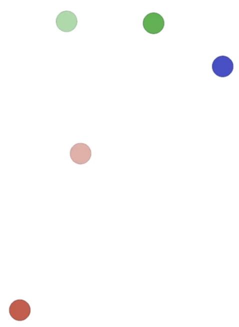

Figure 1: **Example environment.** Each agent and landmark has their x and y coordinates randomly initialized within [-1, 1]. Of the three landmarks, one is red, one blue, and one green. In the rendered environment, pictured above, the three landmarks have solid colors, whereas the two agents are colored lightly, with their colors corresponding to the landmark they should move to – this is equivalent to the landmark corresponding to the other agent's goal. In the pictured environment, the red agent's goal is for the green agent to move to the green landmark, and vice-versa.

$\mathcal{S}_t = [p_{1_t} \cdots p_{5_t} \quad c_{1_t} \quad c_{2_t} \quad g_1 \quad g_2]$, and a given agent's observation by the vector $o_{i_t} = [{}_ip_{1_t} \cdots {}_ip_{5_t} \quad c_{1_t} \quad c_{2_t} \quad g_i]$, where ${}_ip_{j_t}$ is the j^{th} entity's position seen through agent i's reference frame at time t.

Throughout training and evaluation, we roll-out environments for 100 frames before terminating. Rewards are cooperative, calculated by summing the $\mathcal{L}_2$ distances between each agent's goal landmark and the position of their partner agent (i.e. the agent they want to move to their goal landmark). This sum is then negated, ensuring that when we minimize our loss (i.e. the negative of the reward) we directly minimize the distance our agents are from where the other agent would like them to move.

In an effort to aid development of our environments, Bettini et al. (2022)'s Vectorized Multi-Agent Simulators, also known as VMAS, are used as a starting point from which we develop our environments. The environments already present within VMAS include the multi-particle environments from Lowe et al. (2017), with the additional bonus that they are developed in a vectorized form[1]. Vectorization allows for parallel execution of multiple environments at once, speeding up training time by many orders of magnitudes.

2.2 Algorithm Setup

Agents learn policies via Lowe et al. (2017)'s off-policy Multi-Agent Deep Deterministic Policy Gradients (MADDPG), with a decentralized critic. The decision to use decentralized critics was intended to allow each individual in a population to develop their own "world view", in an attempt to instill real-life credibility to our simulations. In order to train our populations of two agents, Bettini et al. (2023)'s BenchMARL package was used, chosen due to its built in implementation of MARL algorithms, such as MADDPG, native integration with VMAS, and modularity.

As suggested by the VMAS fine-tuned hyper-parameters in Bettini et al. (2023) for off-policy algorithms, our populations are trained with 60 environments per batch. Populations are trained on a total of 54,000 environments, with intermediate checkpoints repeatedly saved after exposure to 3,000 environments. Throughout the rest of this paper, we label checkpoints by the amount of frames they have seen[2]. Full hyper-parameters, which we source from Bettini et al. (2023)'s fine-tuned VMAS hyper-parameters, are specified in Appendix A.

2.3 Idiolect Development

In order to simulate idiolects, we develop two training regimens: one inherent to the environment, and one inherent to the agent's policies.

2.3.1 Environmental Idiolects

One such method to develop idiolects simulates their effects, i.e. good-enough production (Goldberg and Ferreira, 2022), and good-enough comprehension (Ferreira and Patson, 2007). To achieve this, we distort agents' communications to each other. For each agent in a population the following procedure is performed. Two values are selected at random from a Gaussian distribution and used to parameterize a Beta-distribution. During every time-step in a given episode, each agent will add a 10-dimensional vector, sampled at random from their unique Beta-distribution and scaled down by

[1] In the context of this study, vectorized environments and batched learning are synonymous.

[2] For example the first checkpoint, which has been trained on 3,000 environments – and subsequently 300,000 frames – is labelled Checkpoint 300,000

2, to their observed and outgoing communications. In doing this, each agent has small, unique, losses of completion in their "production" – sending of a communication – and "comprehension" – receiving of a communication.

2.3.2 Policy-Based Idiolects

The other method utilized to simulate idiolects in an environment is by directly modifying the characteristics of a population's learned policy. Recollecting that idiolects are merely an individual's own unique understanding of a language, in order to simulate idiolects, policies must be modified to be unique to each agent in a population. To accomplish this, populations trained to develop this representation of idiolects are trained with decentralized actor networks. As such, agents do not share the parameters of their actor network, letting them each develop a unique understanding of their emergent langauges.

2.3.3 Trained Populations

Given the two different methods to simulate idiolects, we train four different types of populations:

1. A population with shared actor-network parameters and no noise in the communication channel.
2. A population with shared actor-network parameters and noise in the communication channel.
3. A population with unshared actor-network parameters and no noise in the communication channel.
4. A population with unshared actor-network parameters and noise in the communication channel.

By training with these four paradigms, we allow for like-to-like comparisons between different representations of idiolects in multi-agent populations. We further train three populations[3], with distinct seeds, for each group, enabling us to evaluate the statistical soundness of any observed phenomena or results.

3 Experiments

All populations are evaluated on four different metrics: performance on in-distribution environments (Section 3.1), performance during zero-shot generalization (Section 3.2), zero-shot agent integration (Section 3.3), and performance in variable noise conditions (Section 3.4). These metrics were chosen to test the population's abilities to perform when faced with real-world challenges.

Throughout this section, we refer to agent performance using their total rewards[4] throughout an episode. To frame this discussion, we find it necessary to mention that in our environments, agents whose total rewards are in the range of -70 or higher typically have developed the ability to complete our given task. In contrast, the optimal non-verbal strategy – which corresponds to agents jointly moving to the geometric average of the triangle formed by the landmarks – scores nearer to -120.

3.1 Performance on In-Distribution Environments

Setup

Populations are first measured on their abilities to complete tasks in environments selected from the same distribution on which they were trained. For each population, total reward-per-episode is evaluated at every checkpoint on 5,000 environments (1,000 each for 5 different seeds). Performances are subsequently averaged per group (i.e. each different training paradigm) and then compared between groups.

Results

Figure 2 shows all four population groups evaluated on environments in the training distribution across all saved checkpoints. At first glance, it seems that across all checkpoints evaluated there are no consistent, significant, differences between any populations[5]. However, when we fit a linear mixed effects model to our data after convergence (i.e. performance at checkpoints 3,600,000 to 5,400,000), as seen in Table 1, we find that the coefficient for the noisy training indicator variable is both positive and significant, with no other variables significant. This indicates that when our populations are trained with a noisy channel, they possess clear performance benefits when compared to populations trained with a clean communication channel, and further, that this is the only factor in this setup which provides significant benefit.

[3] This number was chosen due to computational limitations in our study.

[4] In a given episode, total reward is the sum of each agents' rewards at each time-step.

[5] Numerical values can be seen in Appendix B

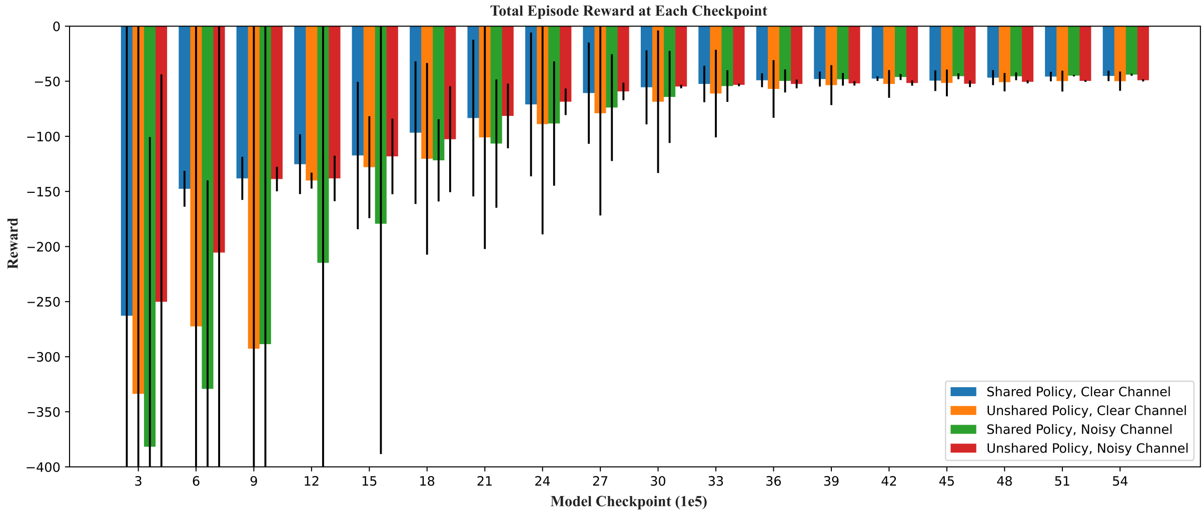

Figure 2: **In-Distribution Environment Total Reward.** Plotted for all populations when evaluated at every saved checkpoint. For each population, the performance is evaluated at every checkpoint on 5,000 environments – 1,000 each for 5 different seeds, with mean reward in a given environment averaged across all evaluated environments.

Variable	Coefficient	Standard Error	Z-Score	P > |Z|
Noisy Channel	5.363	1.190	4.509	<.00001
Shared Network	1.838	2.291	0.802	0.422
(1|Checkpoint)	0.000	0.000	4.945	0.000
(1|Population)	0.551	0.364	1.513	0.130

Table 1: **Linear Mixed Effects Model for In-Distribution Environments**. Coefficients and statistics fit for population performance on in-distribution environments. The model is fitted with the formula: Reward = 1 + (Noisy Channel) + (Shared Network) + (1|Checkpoint) + (1|Population). In this case, noisy channels and shared parameters are both indicator variables and treated as fixed effects. We control for the checkpoint and population by enforcing them as random effects.

3.2 Zero-Shot Generalization to New Colors

Setup

Populations are next measured on their performance in environments sampled from outside the training distribution. To create novel environments, we randomize the color of every landmark, instead of having one red, one green and one blue. This task then uses the same metrics as presented in Section 3.1.

Results

Figure 3 shows all four population groups evaluated on environments in the training distribution across all saved checkpoints. We observe that, similarly to the in-distribution environment task, there are no consistent, significant, performance differences between groups. However, unlike the in-distribution environments, our populations struggle to even complete the task when we use completely random colors. This inability can be seen in their total reward plateauing in-between -120 and -160, indicating a regression to the optimal non-verbal strategy[6]. When we fit the same linear mixed effects model (Table 2) as in Section 3.1, we now see there are no significant variables. These results indicate that none of our populations possess traits allowing them to better generalize to new environments.

3.3 Zero-Shot Agent Integration

Setup

Inspired by the phenomena of language contact, we design a task to evaluate the performance of our agents when interacting with agents trained in a different population. In this setting, two agents are isolated from different populations and introduced to each other. Total reward-per-episode is then measured across 500 environments (100 environments across 5 different seeds). Similar to our evaluation in Section 3.1, performance is then averaged per group and compared between groups.

This task, known in RL-settings as Cross-Play (Hu et al., 2021), is performed between every checkpoint after convergence in every population, with convergence occurring in our populations around exposure to 3,600,000 environments.

Results

Figure 4 demonstrates that agents struggle to perform in cross-play settings. We see, along the diagonals of our heat maps, that when populations are measured performing cross-play with their own population, i.e. a normal roll out, they successfully complete a given task. In comparison, when we look at the off diagonals, we see that agents' scores are consistent with the optimal non-verbal strategy. To us, this clearly indicates that agents struggle to communicate when introduced to agents that were trained separately, and subsequently speak a different language.

3.4 Variable Noise Conditions

Setup

The final set of tasks is inspired by work from Logan et al. (1991), McMillan and Saffran (2016), and Austin et al. (2022) among others, which suggests that noise and variability help drive more robust language learning in individuals. To investigate this idea through our simulations, we evaluate the performance of our populations, post-convergence, in three specific environments: environments with no noise at all, environments where populations are evaluated in the same noise conditions as they were trained in, but with additional Gaussian noise added to the communication channel, and environments with no idiolectic noise (i.e. noise unique to an agent and simulating "good enough communication") but possessing Gaussian noise. This Gaussian noise is intended to corrupt the communication channel, acting similarly to external noise in daily communication – we refer to this Gaussian noise as "external noise" going forward. All the above tasks feature environments sampled from our training distribution.

Each of these tasks are designed to specifically capture different ways in which idiolects may provide performance advantages to a population of speakers. The first task attempts to compare the raw language which develops in a given population, stripping away any communication failures that occur as a result of idiolects. In comparison, the second task attempts to probe how robust performance is to external noise when it is present in the setting in which a language developed. Finally, the third task focuses specifically on this external noise, attempting to discern how robust languages are to purely external noise, removing any noise present due to idiolects.

Similarly to our cross-play evaluations, we measure total reward across 500 environments (100 environments each for 5 different seeds). Further,

[6] Numerical values can be seen in Appendix C

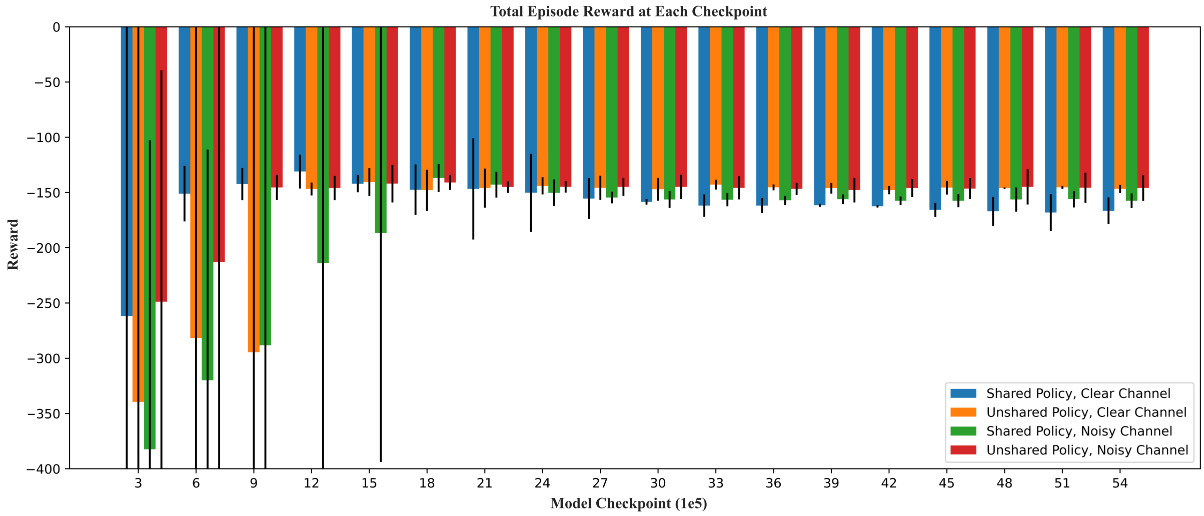

Figure 3: **Novel Environment Total Reward.** Plotted for all populations when evaluated at every saved checkpoint. For each population, the performance is evaluated at every checkpoint on 5,000 environments – 1,000 each for 5 different seeds, with mean reward in a given environment averaged across all evaluated environments.

Variable	Coefficient	Standard Error	Z-Score	P > |Z|
Noisy Channel	-4.284	3.224	-1.329	0.184
Shared Network	-6.013	3.823	-1.573	0.116
(1|Checkpoint)	-0.000	0.000	-0.980	0.327
(1|Population)	-0.591	0.382	-1.545	0.122

Table 2: **Linear Mixed Effects Model for Novel Environments**. Coefficients and statistics fit for population performance on novel environments. The model is fitted with the formula: Reward = 1 + (Noisy Channel) + (Shared Network) + (1|Checkpoint) + (1|Population). In this case, noisy channels and shared parameters are both indicator variables and treated as fixed effects. We control for the checkpoint and population by enforcing them as random effects.

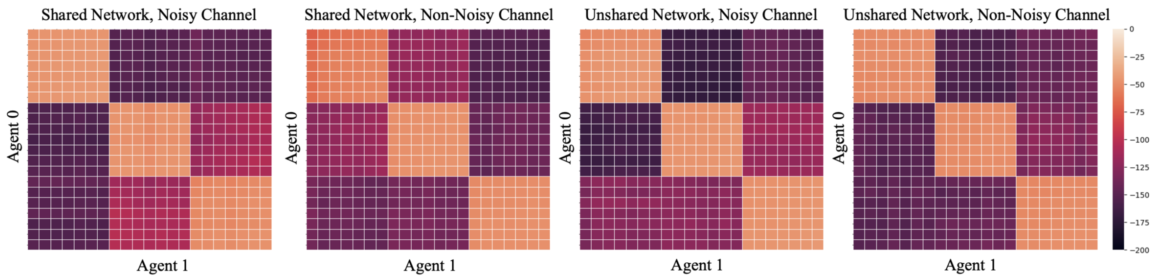

Figure 4: **Cross-Play Matrix**. Visualization of agent performance (total reward) when paired with other checkpoints (within trained population and out of trained population). y-axis represents Agent 1, and the x-axis Agent 0. **Agents 0-8:** Checkpoints 3600,000 to 5,400,000 of Population 0, spaced by intervals of 300,000. **Agents 9-17:** Checkpoints 3,600,000 to 5,400,000 of Population 1, spaced by intervals of 300,000. **Agents 18-26:** Checkpoints 3600,000 to 5,400,000 of Population 2, spaced by intervals of 300,000. Each block in the grid is obtained by evaluating the pair on 500 randomly generated environments, 100 environments across 5 seeds.

for each task that involves adding Gaussian noise, we evaluate with Gaussian noise at 20%, 40%, 60%, 80% and 100% of the full noise level, where "full noise" is described as selecting from a Gaussian distribution with $\mu = 0$ and $\sigma = 1$.

Results

We fit a linear mixed effects model to determine the effect of our various training paradigms on final performance, where the "*shared_params*" and "*noisy_training*" variables are indicator variables for whether actor network parameters are shared and training includes noisy channels respectively. We report our results in Tables 3, 4 and 5. In all tasks, we can see that our noisy indicator variable has a positive, significant, coefficient, indicating that when we train with noise, it leads to significant increases in performance. Furthermore, in both of our tasks which add external Gaussian noise, we see that our shared parameters indicator variable has a negative, marginally significant, coefficient. These results indicate that when we train without shared parameters, we also find evidence of increases in performance.

4 Discussion

In this work, we built upon a reinforcement-learning framework for emergent communication to develop simulations of real-world populations with varied linguistic traits. Our experiments on four different tasks demonstrate that idiolects likely aid the development of languages better fit to their general environment and task. Furthermore, our results also indicate idiolects may help drive the evolution of languages more robust to external noise. In contrast, we see no benefits to idiolects in zero-shot generalization, or zero-shot cross-play. In this section, we will walk through our various tasks, and discuss our results further.

4.1 In-Distribution and Novel Environment Performance

We demonstrate that populations of agents possessing idiolects perform better than populations of agents without idiolects when evaluated in environments drawn from their training distribution. Specifically, we find that populations of agents trained with noisy channels simulating "good-enough" production and comprehension, show performance benefits when evaluated on in-distribution environments. In contrast, we find no significant performance benefit derived from training with or without shared actor networks.

Our injection of idiolectic noise to the communication channel can be viewed in a similar vein to adding noise during training in classical deep learning. This has been reported by Srivastava et al. (2014), among countless others, to provide performance advantages by smoothing the data distribution, which in turn helps prevent overfitting[7]. Applying these ideas to our results, in which performance is significantly positively affected by noise during training, would suggest that the presence of idiolects – and the noisy communication generated by them – may help in disrupting any overly precise patterns or structures present in the encountered environments. This data smoothing likely pushes our languages to develop more generalized communicative systems, better fit to the task and overall distribution of environments seen during training. In contrast, universal languages might over-fit to the specific environmental configurations seen during their evolution.

We also see that when performing our evaluations in novel environments, with landmarks possessing colors not seen during training, neither noisy training nor shared networks lend performance advantages to our populations. We perform these evaluations in a stop-gap attempt to measure the capacity for adaptation, recognizing that there are fundamental limitations in our methodology – discussed in Section 4.4 – that prevent us from measuring adaptation speed, duly limiting us to zero-shot tasks. As such, this failure in generalization is not particularly surprising; our originally presented theory presents idiolects as a source for rapid linguistic mutation for natural selection, which should manifest in quicker adaptation to new environments. However, zero-shot generalization tasks, such as the one we use for evaluation, are largely impossible for our populations when not provided the chance to adapt to these new environments.

4.2 Zero-Shot Agent Integration

Our results for the zero-shot agent integration task suggest that agents with idiolects cannot communicate with agents from other populations any better

[7] It is pertinent to note that in deep learning one classically removes noise regularization during test time to realize performance advantages, however the evaluations reported here occur (and duly performance advantages are realized) with noise present for populations trained with noisy channels.

Variable	Coefficient	Standard Error	Z-Score	P > |Z|
Noisy Channel	5.927	1.175	5.043	<0.0001
Shared Network	0.950	2.263	0.420	0.675
(1|Checkpoint)	0.000	0.000	5.037	0.000
(1|Population)	0.431	0.360	1.198	0.231

Table 3: **Linear Mixed Effects Model for Non-Noisy Environments**. Coefficients and statistics fit for population performance on non-noisy environments. The model is fitted with the formula: Reward = 1 + (Noisy Channel) + (Shared Network) + (1|Checkpoint) + (1|Population). In this case, noisy channels and shared parameters are both indicator variables and treated as fixed effects. We control for the checkpoint and population by enforcing them as random effects.

Variable	Coefficient	Standard Error	Z-Score	P > |Z|
Noisy Channel	5.126	0.496	10.324	<0.0001
Shared Network	-1.835	0.971	-1.889	0.059
(1|Checkpoint)	0.000	0.000	11.640	0.000
(1|Population)	-0.171	0.155	-1.102	0.271
(1|Noise)	-0.027	0.049	-0.544	0.587

Table 4: **Linear Mixed Effects Model for Environments with Training Noise Conditions and Added External Noise**. Coefficients and statistics fit for population performance on their training environments with various amounts of Gaussian noise in the communication channel. The model is fitted with the formula: Reward = 1 + (Noisy Channel) + (Shared Network) + (1|Checkpoint) + (1|Population) + (1|Noise). In this case, noisy channels and shared parameters are both indicator variables and treated as fixed effects. We control for the checkpoint, population, and external noise level by enforcing them as random effects.

Variable	Coefficient	Standard Error	Z-Score	P > |Z|
Noisy Channel	5.199	0.495	10.507	<0.0001
Shared Network	-2.029	0.968	-2.096	0.036
(1|Checkpoint)	0.000	0.000	11.692	0.000
(1|Population)	-0.199	0.155	-1.284	0.199
(1|Noise)	-0.051	0.049	-1.046	0.296

Table 5: **Linear Mixed Effects Model for Externally Noisy Environments**. Coefficients and statistics fit for population performance on non-idiolectically noisy environments with various amounts of Gaussian noise in the communication channel. The model is fitted with the formula: Reward = 1 + (Noisy Channel) + (Shared Network) + (1|Checkpoint) + (1|Population) + (1|Noise). In this case, noisy channels and shared parameters are both indicator variables and treated as fixed effects. We control for the checkpoint, population, and external noise level by enforcing them as random effects.

than agents without idiolects can, when evaluated in a zero-shot setting. This similar performance is also not surprising to us, as for there to be notable difference, one of two things would likely need to be true: either the presence of idiolects would provide the ability for generalization of communication to new languages, or idiolects would need to act as "regularizers", driving populations to the same communication methods for similar environments and tasks. However, neither of these phenomena are observed in the real world.

We do observe that, in this absence of communicative ability, agents adopt the optimal non-verbal strategy. This is promising to observe, as it is what we would expect to see in human settings, and lends more real-world credence to our simulations.

4.3 Variable Noise Settings

One task where we observe significant differences is when we consider variable noise settings. As a reminder, in these evaluations, we evaluate our populations on three tasks: the first task comparing purely the language developed in different training regimens, with no noise of any type present, the second task probing how robust in-distribution performance is to external noise, and the third task focusing specifically on how robust each language is to external noise, possessing only environmental, but no idiolectical noise.

Our first point of note is that the presence of a idiolectically noisy channel during training drives significant performance advantages for a given population in all of our evaluatory tasks. As mentioned in our discussion in Section 4.1, we believe this is an indicator of the noise resulting from idiolects "smoothing" the distribution of environments encountered, driving the development of languages not over-fit to the environments seen during training. Given that idiolectic noise in our simulations mirrors good-enough production (Goldberg and Ferreira, 2022) and comprehension (Ferreira and Patson, 2007), we conclude that "good-enough communication" – in part driven by idiolects – helps guide language evolution toward languages better-fit to the tasks and environments encountered during said evolution than those which evolve with no idiolectic noise. Of course, these results do not preclude the possibility that other aspects of "good-enough communication" are equally, or more, important in driving effective communication; however these are not questions we can answer given the structure of our work, instead necessitating future work.

When we evaluate our populations in settings with external noise, we also observe the emergence of shared parameters in our agent's actor networks as a detriment to performance. In comparison to noisy training, which seems to universally drive the performance of languages better fit to the environment and tasks at hand, learning non-ubiquitous representations of language only lends performance benefits when we test for their robustness to communication corruption. This performance drop associated with shared parameters in our agent's actor networks indicates that a non-universal representation of language drives the development of robust communicative signaling. This robustness provides performance advantages in the most "real-world" settings we consider, in which noise occurs not only from individual communicative points-of-failure, but also from the external environment. While our earlier results regarding noise are notable, they are not purely beholden to idiolects, with other sources of noise likely to give similar results. However, non-ubiquity in language understanding is a direct representation of idiolects, and as such we find these results extremely encouraging, suggesting this decentralized view of language lends a robustness not seen in universal languages; an extremely important trait to have.

4.4 Limitations and Future Work

One avenue of immediate work which could further this project could comprise training more populations. Our sample size of three populations per group is atypically small, due to our limited computational resources. If we include more populations in our statistical analyses, we expect to see a large reduction in our p-values, allowing us to be more confident in our results and conclusions. This is particularly relevant to our results presented in Section 3.4, but also broadly applicable.

Another line of immediate inquiry would be to reproduce our results on other MARL tasks necessitating communication. There exist a host of other MARL tasks, some of which are already integrated into VMAS, which necessitate the emergence of communication between agents for completion. Reproducing our observed results on these tasks and environments would also allow us to draw stronger conclusions supporting the findings that we have thus far demonstrated with a single task.

Perhaps the largest limitation to our study was our use of batched simulators, necessitated by our lack of adequate computational resources. In a batched setting, such as the VMAS environments (Bettini et al., 2022) we used in this study, algorithms process "batches" of data at a time, with subsequent weight updates performed utilizing averages across all the rewards seen in the various environments in the batch. In this manner, our algorithm is able to process a large amount of environments at a time (in our case 60), thus speeding up training by an order of magnitudes.

However, while computationally efficient, batched learning is quite dissimilar to human learning. In settings similar to the ones we present to our agents, humans will only process, and learn from, one environment and scenario at a time. It is a large part due to this sequential cognitive processing that an un-batched RL algorithm is much more cognitively sound of a model. Further, as far as we are aware, there is no concrete evidence that relative learning speeds between various populations in a batched setting will hold in an un-batched setting. As such, we refrained from making any cognitive conclusions from our batched learning speeds. Regrettably, due to computational constraints, training all our populations in an un-batched setting would take months to years longer than was possible for us to train. Consequently, we were fundamentally unable to test a core idea of our hypothesis, which was that idiolects provide natural "mutations" in language, allowing for faster evolution and adaptation to new environments. Instead, we were left only able to evaluate our agents on the stop-gap zero-shot tasks, which, as expected and previously addressed in furthur detail, did not turn up any significant differences between populations.

Future work on this topic could address this, developing simulations in which valid cognitive conclusions could be made regarding population learning speed, whether merely training in an un-batched manner, or devising some other method to study this factor.

In hindsight, another limitation of this study was the use of a communication vector which both allowed continuous values in each channel, and was high-dimensional. This high-dimensional, continuous signalling allowed our agents to develop communication strategies which did not explicitly map specific communications to meanings, as human language does. Instead our populations mapped

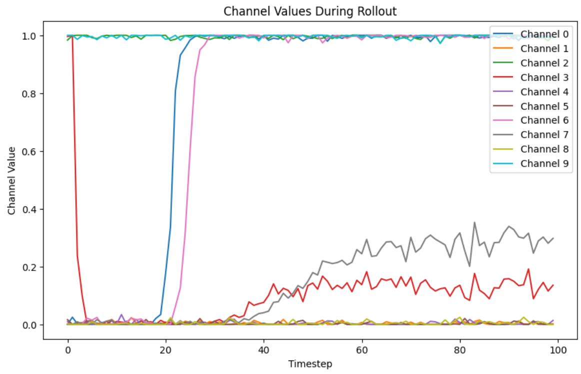

Figure 5: **Agent Communication Channels.** Values in the channels of a given agent's communication vector throughout an example roll-out.

continuous changes in specific channels to meaning. This can be better viewed in Figure 5 in which we plot the values in each channel of a given agent's communication over the course of a roll out. Here, it is clear that the signalling that is occurring is not via discrete communications, as words and sentences are in human language, but by continuous changes in a given channel. This suggests that the communication aspect of our modeling is not completely representative of human language. One way future work may address this problem by lessening the dimensionality of the communication vector, and enforcing discrete values in each channel. This could perhaps lead to the development of more human traits in the language, as Mordatch and Abbeel (2018)'s work suggests that limiting the "vocabulary" for agents leads to enhanced compositionality in their communications. Further ideas for developing more "natural" language is discussed in work such as Kottur et al. (2017), and is worth integrating into future simulations.

Another avenue of exploration, informed by work from Piantadosi et al. (2012), might approach this problem from an opposite direction. Instead of directly simulating idiolects in the training paradigm, one could train multiple populations from similar starting points, with varying costs on their communication function. One could then try to evaluate whether idiolects and ambiguity arise in some due to a pressure for greater linguistic efficiency. This methodology does, of course, come with its own issues chief among them being the ideas that "words are cheap". Different words do not necessitate the same physical capacity as traits in other species which could be modeled utilizing costs – as can be done with mating rituals or parental care of young. This phenomena makes it particularly hard to design a meaningful cost function with which we can drive idiolectic development. Another issue is that it is not immediately clear how to measure linguistic efficiency, with work from Coupé et al. (2019) suggesting that many languages encode information at similar information rates, despite their clear linguistic differences. As such, it would be particularly difficult to design metrics which measure the linguistic efficiencies of these potential final languages.

A final avenue we see for potential future work concerns treating the amount of noise in the channel as a hyper-parameter, or even a random variable. We select the scaling value of $\frac{1}{2}$ for our "good-enough" communication noise somewhat arbitrarily, wanting to select a value which was large enough to make a difference to the communication, but small enough to not completely corrupt the channel. Future work may try varying this completely to gain better intuition into how this value effects our various populations and tasks.

5 Conclusion

In sum, our simulations provide evidence that idiolects likely provide multiple benefits to evolving languages. In particular, we find they contribute to the development of languages not over-fit to the environments witnessed during evolution, and further that they aid in driving the evolution of languages robust to external communication failures.

References

Alison C Austin, Kathryn D Schuler, Sarah Furlong, and Elissa L Newport. 2022. Learning a language from inconsistent input: Regularization in child and adult learners. *Language Learning and Development*, 18(3):249–277.

Matteo Bettini, Ryan Kortvelesy, Jan Blumenkamp, and Amanda Prorok. 2022. Vmas: A vectorized multi-agent simulator for collective robot learning.

Matteo Bettini, Amanda Prorok, and Vincent Moens. 2023. Benchmarl: Benchmarking multi-agent reinforcement learning. *arXiv preprint arXiv:2312.01472*.

Nick Chater and Morten H Christiansen. 2010. Language acquisition meets language evolution. *Cognitive science*, 34(7):1131–1157.

Morten H Christiansen and Nick Chater. 2008. Language as shaped by the brain. *Behavioral and brain sciences*, 31(5):489–509.

Morten H Christiansen and Nick Chater. 2016. The now-or-never bottleneck: A fundamental constraint on language. *Behavioral and brain sciences*, 39:e62.

Morten H Christiansen and Nick Chater. 2022. *The language game: How improvisation created language and changed the world*. Random House.

Christophe Coupé, Yoon Mi Oh, Dan Dediu, and François Pellegrino. 2019. Different languages, similar encoding efficiency: Comparable information rates across the human communicative niche. *Science advances*, 5(9):eaaw2594.

Fernanda Ferreira and Nikole D Patson. 2007. The 'good enough' approach to language comprehension. *Language and linguistics compass*, 1(1-2):71–83.

J Foerster. 2018. *Deep multi-agent reinforcement learning*. Ph.D. thesis, University of Oxford.

Adele E Goldberg and Fernanda Ferreira. 2022. Good-enough language production. *Trends in Cognitive Sciences*, 26(4):300–311.

Hengyuan Hu, Adam Lerer, Alex Peysakhovich, and Jakob Foerster. 2021. "other-play" for zero-shot coordination.

Satwik Kottur, José MF Moura, Stefan Lee, and Dhruv Batra. 2017. Natural language does not emerge'naturally'in multi-agent dialog. *arXiv preprint arXiv:1706.08502*.

Angeliki Lazaridou, Alexander Peysakhovich, and Marco Baroni. 2016. Multi-agent cooperation and the emergence of (natural) language. *arXiv preprint arXiv:1612.07182*.

Michael L Littman. 1994. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*, pages 157–163. Elsevier.

John S Logan, Scott E Lively, and David B Pisoni. 1991. Training japanese listeners to identify english/r/and/l: A first report. *The Journal of the Acoustical Society of America*, 89(2):874–886.

Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. *Neural Information Processing Systems (NIPS)*.

Brianna TM McMillan and Jenny R Saffran. 2016. Learning in complex environments: The effects of background speech on early word learning. *Child Development*, 87(6):1841–1855.

Igor Mordatch and Pieter Abbeel. 2018. Emergence of grounded compositional language in multi-agent populations.

Steven T Piantadosi, Harry Tily, and Edward Gibson. 2012. The communicative function of ambiguity in language. *Cognition*, 122(3):280–291.

Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958.

A Appendix - Training Hyper-Parameters

```
gamma: 0.9
lr: 0.00005
adam_eps: 0.000001
clip_grad_norm: True
clip_grad_val: 5
soft_target_update: True
polyak_tau: 0.005
hard_target_update_frequency: 5
exploration_eps_init: 0.8
exploration_eps_end: 0.01
exploration_anneal_frames: 1_000_000
max_n_frames: 10_000_000
off_policy_collected_frames_per_batch:
    6000
off_policy_n_envs_per_worker: 60
off_policy_n_optimizer_steps: 1000
off_policy_train_batch_size: 128
off_policy_memory_size: 1_000_000
checkpoint_interval: 300_000
```

B Appendix - Numerical Values for Evaluations of Populations on In-Distribution Environments

Checkpoint	Shared Noise	Shared Noiseless	Unshared Noise	Unshared Noiseless
300,000	-262.78 ± 454.276	-333.714 ± 448.376	-381.738 ± 281.134	-250.189 ± 206.438
600,000	-147.626 ± 16.288	-272.513 ± 535.208	-329.255 ± 189.334	-205.563 ± 279.538
900,000	-138.172 ± 19.519	-292.733 ± 598.28	-288.663 ± 327.983	-138.792 ± 11.094
1,200,000	-125.307 ± 27.138	-140.134 ± 7.247	-214.815 ± 321.981	-138.186 ± 20.629
1,500,000	-117.412 ± 66.899	-127.968 ± 46.261	-179.387 ± 208.939	-118.199 ± 34.311
1,800,000	-96.684 ± 64.702	-120.381 ± 86.994	-121.796 ± 37.3	-102.629 ± 48.065
2,100,000	-83.442 ± 71.088	-100.948 ± 101.303	-106.602 ± 58.202	-81.474 ± 29.354
2,400,000	-71.031 ± 65.205	-88.963 ± 99.997	-88.386 ± 56.406	-68.622 ± 12.041
2,700,000	-60.776 ± 45.935	-79.083 ± 92.719	-73.915 ± 48.479	-59.232 ± 7.944
3,000,000	-55.497 ± 33.594	-68.59 ± 64.698	-64.232 ± 41.795	-54.78 ± 1.589
3,300,000	-52.508 ± 16.538	-61.191 ± 39.744	-54.361 ± 14.354	-53.327 ± 1.132
3,600,000	-49.134 ± 6.275	-56.989 ± 26.146	-49.682 ± 10.468	-52.449 ± 3.976
3,900,000	-48.009 ± 6.776	-53.516 ± 18.138	-48.231 ± 5.72	-51.93 ± 1.936
4,200,000	-47.621 ± 2.07	-52.431 ± 12.555	-46.159 ± 2.723	-51.648 ± 2.38
4,500,000	-49.456 ± 9.406	-51.541 ± 12.162	-45.478 ± 2.605	-52.306 ± 3.008
4,800,000	-46.745 ± 6.805	-50.877 ± 8.321	-45.528 ± 3.531	-50.583 ± 1.289
5,100,000	-45.769 ± 4.303	-49.869 ± 9.447	-44.945 ± 0.702	-49.721 ± 0.797
5,400,000	-45.168 ± 4.714	-49.93 ± 8.76	-44.037 ± 0.963	-49.16 ± 1.061

C Appendix - Numerical Values for Evaluations of Zero-Shot Generalization to New Colors

Checkpoint	Shared Noise	Shared Noiseless	Unshared Noise	Unshared Noiseless
300,000	-261.812 ± 460.153	-339.496 ± 425.395	-382.414 ± 279.709	-248.817 ± 209.382
600,000	-151.037 ± 25.128	-281.636 ± 514.422	-319.989 ± 208.89	-212.994 ± 268.554
900,000	-142.477 ± 14.547	-294.705 ± 595.838	-288.403 ± 316.584	-145.45 ± 11.337
1,200,000	-131.035 ± 15.353	-146.802 ± 5.867	-214.006 ± 318.971	-146.018 ± 10.979
1,500,000	-142.044 ± 7.749	-140.604 ± 12.632	-186.769 ± 207.135	-141.995 ± 17.021
1,800,000	-147.457 ± 22.9	-147.927 ± 18.624	-136.926 ± 12.613	-140.929 ± 6.818
2,100,000	-146.756 ± 45.838	-145.98 ± 17.692	-142.888 ± 11.729	-145.073 ± 5.13
2,400,000	-150.196 ± 35.314	-144.012 ± 7.735	-150.201 ± 12.032	-144.826 ± 5.236
2,700,000	-155.504 ± 18.466	-145.725 ± 10.986	-154.559 ± 5.296	-144.891 ± 8.233
3,000,000	-158.476 ± 2.476	-147.12 ± 10.249	-156.347 ± 7.527	-144.885 ± 11.032
3,300,000	-161.849 ± 10.07	-142.852 ± 4.449	-156.492 ± 6.048	-145.806 ± 10.433
3,600,000	-161.845 ± 6.825	-145.374 ± 2.657	-157.131 ± 4.223	-146.667 ± 5.539
3,900,000	-161.589 ± 1.41	-146.143 ± 4.829	-156.16 ± 4.429	-147.929 ± 11.053
4,200,000	-162.544 ± 1.047	-147.926 ± 3.695	-157.45 ± 3.901	-146.03 ± 8.204
4,500,000	-165.6 ± 6.353	-145.598 ± 6.214	-157.417 ± 5.929	-146.442 ± 9.504
4,800,000	-167.063 ± 13.164	-146.01 ± 0.629	-156.379 ± 10.924	-144.911 ± 15.931
5,100,000	-168.136 ± 16.526	-145.357 ± 1.332	-156.02 ± 7.536	-145.633 ± 13.604
5,400,000	-166.519 ± 12.136	-146.755 ± 3.674	-157.427 ± 6.602	-145.971 ± 11.602

Acknowledgements

I am often told my writing can be long and flowery. As an aspiring academic, it is something that I usually have to work to beat out of my writing, to instead write up clear, concise science. However, every once in a while a piece of writing comes around, in which I feel I need all the words and expression I can possibly use. This is one of those. What follows is my completely unabridged thanks to so many of you who have played a large role in my academic journey. Thank you all, from the bottom of my heart.

Thank you, first and foremost, to my direct advisors. Dr. Morten Christiansen, thank you for all your assistance in conceiving, shaping, and executing this project. I first met with you as a student with a mere interest in language evolution, and am leaving as someone with both a deep fascination and a thesis in the field. You saw every iteration of my ideas for this project, from my first, extremely disjointed proposal, all the way to my final thesis, and have been immeasurably helpful and supportive the whole way. I couldn't have done this without your help. Dr. Eugene Vinitsky, I owe a massive debt of gratitude to your willingness to respond to a cold email from a student at a different university professing interest in completing a grandiose project with your help. Without your continued guidance and assistance, I would likely still be trying to re-implement all my Multi-Agent RL algorithms and environments from scratch (and losing much hair along the way). Thank you for all the time you took to help me, you have made me an immeasurably better computational researcher. Dr. Veit Elser, you watched my project change from its initial idea of Quantum Physics and Philosophy all the way to Language Evolution and Multi-Agent RL. Thank you for all your help in finding committee members, being a soundboard for my all-to-large ideas, and general academic guidance throughout the College Scholar Program.

Speaking of the College Scholar Program, it would be remiss to mention just how much this program has shaped my time at Cornell. Without the freedom to truly explore all my academic interests, I would likely still be studying Chemistry, Statistics, Biology, or another of the many subjects I had forays into for my first two years, and much less academically satisfied with the work I was doing. To that end, Dr. Michael Goldstein, Mike, you have been one of the firmest points of unadulterated support and advice throughout my time at Cornell, and for that I thank you. Thank you, Aliou Kamau Gambrel, for being a role model and friend throughout my time in this program. Without meeting you, I would likely never have found this program, and even if I did, likely never have pursued the questions I ended up pursuing. Even though you graduated two years ago, I know I can still text you for random advice – about the program or life – and expect a helpful, honest, and thoughtful answer back. Finally, thank you to Dr. Malcolm Bare for writing fulfilling the *very late* recommendation letter request that that accompanied my application into this program. Your freshman year writing seminar taught me that my time at Cornell need not be only stress and working, but also fun?

To everyone else who contributed to my Academic and Scientific journey, even if not directly related to my thesis, you have shaped my academic trajectory and my thanks to you are endless. Dr. Klaas van Wijk, you let me work in your lab as an extremely green college freshman, and while I no longer work in the plant sciences, many of the lessons I learnt from you and your lab made me into the researcher I am today. For that, I thank you. Thank you to Pat Dubon, for giving me the opportunity to work at such a place as NASA's Jet Propulsion Laboratories. To you, David Lane, and Asitang Mishra, I owe a debt of gratitude for your support and guidance this summer; I hope to work with all of you again in the not-so-distance future. I am also so lucky to have received the exceptional scientific education I did during my time in grade-school. To Mr. Neumann, Mr. Ludwig, Dr. Dowling, Ms. Nichol, Mr. Ceconi, Ms. Peborde, and Mr. (Coach) Dossena, you encouraged my academic curiosity whilst providing me the necessary background to come to an institution like Cornell and succeed, I thank you immensely. Last but not least in this category, thank you to my many amazing professors at Cornell, chief amongst you (in no particular order) Dr. Lillian Lee, Dr. Marten van Schijndel, Dr. Rachee Singh, Dr. Kilian Weinberger, Dr. Steven Strogatz, Dr. Adrian Lewis, Dr. Alice Nadeu, Dr. Alexander Vladimirsky, Dr. Irena Peeva, Dr. Karola Mészáros, and Dr. Phil Milner. While I am merely a drop in the water in a large body of water which is a Cornell classroom, and some of you may not even remember me, your teaching and time have shaped who I am over these past four years, and I thank you that.

To the many friends I have made during the past four years, you've made Cornell all the amazing things it has been for me. I would've struggled immensely to do any of this – my majors, my thesis, my four years – without you. To Clare Della Valle, I don't know how we got to where we are, but my am I happy we did. You have believed in me time and time again when I didn't myself. Your constant, undying support has propelled me through some of my lowest moments at this school, and to where I am today. You and Cornell are synonymous in my mind, and I wouldn't change that if I could. And to Jake Koneval, I am both so extremely unsurprised and so extremely grateful that our friendship has lasted as long as it has. You mean so much to me, and I can't imagine the past 7 years without you. I know you'll visit me in Austin, and I'll definitely be in Salt Lake again in the future. Thank you to both of you for your constant presence in my life, and being unadulteratedly yourselves.

To my roommates, Leo Davies, Noah Rice, Olaf de Rohan Wilner, thank you for choosing to live with me and sharing so many memories in this short year. Thank you also for tolerating my incessant complaints about the state of our living areas, hopefully I'll become a bit less insufferable during these final weeks. Leo, who would've thought that a kid from a remote area of Northwest Colorado and a kid from Westchester would gel so well together? Not me, that's for sure. But I can't overstate how glad I am that we did. You've made the past three years of my life so much more enjoyable and truly fun than it would be without. To Olaf, thank you for being a mainstay of British culture in my life. I feel like you understand me in a way that very few others don't sometimes; our wine nights and late night FIFA games kept me sane throughout the past year, and I will cherish those memories for a long time. And Noah, now that our theses are completed, perhaps we can both relax and let loose a little... god knows we both deserve it. Max Carothers, thank you for being an endless source of positivity, reason, and laughs. I'm constantly in awe of your drive, and can't wait to see where you go to Med School. To everyone who graduated in the past two years, and I don't get to see as frequently anymore – Gabe Schiffer, Bertie Hoal, Alvaro Perez Triay, Fernando Celaya, Isaac Oluwakuyide, Brady Sites, David Thuman, Quimey Moure – just know how much I miss all of you every day. And to all of you who will be here after I leave – Jin Bae, Victor Lançon, Ben Terhaar, Reid Papafloratous, Pedro Moreira, Max Whitton – I'm so happy that I met you, even if it was only for 3 years. I hope all of you will miss me as much as I will miss you. And of course, Izzy Goodrow, thank you for your constant positivity, especially in the face of the hardest math problem either of us has ever seen. Whilst I may not be a chemistry major anymore, I hope you keep me up to date on whatever niche part of the field you end up mastering during your PhD.

And last, but not least, thank you to my family. To Anna-Sophia thank you for proving me wrong and putting your English degree to good use editing every important piece of writing I've put together (including this one). We've grown so much from the elementary school days of you kicking me off a chair and me tossing a butter-knife at you (now I'd only throw a spoon). Instead of traumatising our mother with our fights, we seem to have joined forces and instead decided to trigger her by strapping our 40L bags to our backs and traipsing off into the woods together in some of the worst weather conditions man has witnessed. But in all seriousness, you've provided me so much support and guidance over the last eight years, for which I am so grateful. While we may still argue every so often, I know our bond will only continue to grow over the coming years. To Father, your influence on my academic interests is self explanatory, and on my soon-to-be-started PhD in Computational Linguistics even more so. And to Mother, you have spent so much of your time listening to me ranting and raving about this project, and the preceding one, and really whatever is on my mind at any given point. You have read every draft of every grad school application, research proposal, thesis and more. To both of you jointly, my Computer Scientist Parents, thank you for supporting me while I explored my many interests, even if I eventually ended up right back in the realm where you both have spent so many years. You have done a fantastic job in raising two children to the points of success where Anna and I currently stand, and I don't think I can thank you enough for that.

To each and every one of you named here, you have made my academic career what it is so far. As Dr. Lillian Lee would say, my debt to you is unbounded.