THESIS FOR THE DEGREE OF LICENTIATE OF ENGINEERING

Efficiency, learnability, and structure in recursive systems of communication

ANDREA SILVI

Department of Computer Science and Engineering
CHALMERS UNIVERSITY OF TECHNOLOGY | UNIVERSITY OF GOTHENBURG
Gothenburg, Sweden, 2026

Efficiency, learnability, and structure in recursive systems of communication

Andrea Silvi

ISSN 1652-876X

Department of Computer Science and Engineering
Division of Data Science and AI
Chalmers University of Technology | University of Gothenburg
SE-412 96 Göteborg,
Sweden
Phone: +46(0)31 772 1000

Printed by Chalmers Digitaltryck,
Gothenburg, Sweden 2026.

Efficiency, learnability, and structure in recursive systems of communication

Andrea Silvi

Department of Computer Science and Engineering
Chalmers University of Technology | University of Gothenburg

Abstract

Language is fundamentally shaped by pressures for efficient communication, typically characterised as a trade-off between simplicity and informativeness. While both information-theoretic models and multi-agent reinforcement learning have successfully captured these dynamics, they predominantly treat linguistic terms as atomic labels. Consequently, these existing frameworks struggle to account for the systematicity and recursiveness that are prevalent across languages.

This thesis extends the study of efficient communication to compositional domains with productive morphosyntax, focusing primarily on recursive numeral systems. Across four contributions, we employ computational modeling to explore how functional and learning pressures shape structured languages. First, we show through multi-agent reinforcement learning that artificial agents optimised solely for communicative success tend to prefer more efficient recursive numeral systems. Second, we argue that previous efficiency measures cannot account for regularity and introduce a different trade-off that can separate human systems from artificial ones that were previously considered optimal but were lacking human-likeness. Third, we connect regularity to learnability, using reinforcement learning to show that human numeral systems exhibit high regularity because they are inherently easier to learn. Finally, we expand this framework to an open-ended collaborative building task, showing that agents utilising procedural abstractions develop languages that minimise similar efficiency trade-offs.

Overall, this work attempts to bridge the gap between efficient communication models and the compositional reality of language, demonstrating how structure is consistently preferred because of communicative and cognitive constraints.

Keywords

Efficient communication, Reinforcement Learning, Regularity, Learnability, Recursive numeral systems.

List of Publications

Appended publications

This thesis is based on the following publications:

[**Paper I**] **A. Silvi**, J. D. Thomas, E. Carlsson, D. Dubhashi, M. Johansson, *Learning Efficient Recursive Numeral Systems via Reinforcement Learning.*
Proceedings of the Annual Meeting of the Cognitive Science Society 47, 2025.

[**Paper II**] P. Prasertsom, **A. Silvi**, J. Culbertson, D. Dubhashi, M. Johansson, K. Smith, *Recursive numeral systems are highly regular and easy to process.*
Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers, pages 4873–4885), 2026.

[**Paper III**] **A. Silvi**, P. Prasertsom, J. Culbertson, D. Dubhashi, M. Johansson, K. Smith, *Evaluating the relationship between regularity and learnability in recursive numeral systems using Reinforcement Learning. Preprint.*

[**Paper IV**] J. D. Thomas, **A. Silvi**, D. Dubhashi, M. Johansson, *PACE: Procedural Abstractions for Communicating Efficiently.*
Proceedings of the Annual Meeting of the Cognitive Science Society 47, 2025.

Acknowledgment

I am extremely grateful to my supervisor Moa Johansson for giving me the opportunity to pursue fascinating cross-disciplinary research and both coming to my aid whenever I got stuck on various issues and trusting me when I wanted to explore a direction we had not considered yet. A special thanks also goes to my co-supervisor Devdatt Dubhashi, who also introduced me to many of these topics. Both of them have allowed me to pursue fun research and have created the bases for great collaborations. I am also thankful for said collaborators: Ponrawee, Jonathan and Emil. All three brought different perspectives to the table and I have learnt a lot from all of them. I am also very grateful to Kenny Smith and Jennifer Culbertson at Edinburgh, who together with Ponrawee make the linguistic side of much of the research in this work much more solid and interesting that I could make it out to be. I am also grateful to my examiner Nir Piterman, who has shown genuine interest in my research and has been at the same time understanding and demanding in a fair way, and has driven me to work even harder.

To all my friends that either support me closeby here in Göteborg or far away across the world, thanks for listening to my complaints and thanks for the time spent together! While it might not always show, I am always grateful for your company and always wish to spend more time with you. To my family in Italy, thanks for always being there for me, even just for a phone call, and while sometimes it feels like we are too much far away from each other, every time we meet I can truly cherish those moments and I would not have it any other way! To my extended family in Sweden, you have become such an important part of my life so quickly and I cannot thank you enough for welcoming me to Värnamo so warmly every time and for making me explore new parts of Sweden every now and then. To my Swedish home family — Pippi, Tomo and Klara — every moment spent together is such a gift. Thank you for putting up with me.

This project was funded by Vetenskapsrådet (registration number 2022-03486). The computational resources provided by Chalmers e-Commons at Chalmers and by the National Academic Infrastructure for Supercomputing in Sweden (NAISS) have enabled some of the research presented in this thesis.

Contents

Part I

Introductory Chapters

Chapter 1

Introduction

Language is one of the most fascinating aspects of human evolution. It has evolved concurrently with human societies over the centuries and millennia thanks to the introduction of new terms for emerging concepts, the obsolescence of ones that might have fallen out of fashion or gotten incorporated into more general terms. Language evolution in general has been driven by a multitude of factors, from cross-cultural contacts, to environmental changes, to cultural and technological innovations, that necessitate new concepts while rendering others obsolete.

When studying language evolution at the individual speaker's level, a prominent line of research suggests that languages are most likely shaped by pressures for *efficient communication* (Gibson et al., 2019; Kemp & Regier, 2012): as individuals use language to coordinate between each other, they might utter new, more efficient terms, which after many more of these micro-level communicative processes, will drive linguistic change at a macro-level. This efficiency has been mostly characterised as a trade-off between *simplicity* and *informativeness*, meaning that languages are "stretched" in opposite directions by competing pressures to both be as simple as possible, to minimise cognitive load, and at the same time be as informative as possible, to maximise the ability to refer accurately to a vast array of concepts. More recently, similar questions have been explored with computational models through the lens of machine learning. A promising strand of research utilises multi-agent reinforcement learning to demonstrate how functional pressures for efficiency can indirectly manifest when dyads of artificial agents learn to coordinate and successfully communicate in specific linguistic domains (Carlsson, 2024). Another related line of research studies how certain cross-linguistic, domain-general biases like *convexity* correlate with the *ease of learning* of linguistic systems, typically operationalised as sample complexity (Douven & Verheyen, 2024; Steinert-Threlkeld & Szymanik, 2020).

While efficient communication and these related frameworks offer interesting information-theoretic explanations to constrained linguistic variation, they mostly analyse terms as **atomic** labels. This assumption often stems either from the nature of the semantic domains which are being studied and the cross-linguistic human data gathered for them, or as a simplification of the problem. However, as a consequence, the existing efficiency models struggle to account for the systematicity of forms across whole lexicons — a highly prevalent feature in many domains that allows for compositionality.

This thesis extends the aforementioned approaches in the context of more complex linguistic domains characterised by compositional semantics and productive morphosyntax. The seminal work of (Denić & Szymanik, 2024) provided a crucial first step in applying efficient communication models to said domains, specifically recursive numeral systems. Taking their work as our starting point, we expand on their line of research in several directions:

Paper 1 expands on the framework of Carlsson et al. (2021), who focused on approximate and exact-restricted numeral systems, to the context of recursive ones, showing how a simple two-agent framework, optimised solely for communicative success, consistently prefers more efficient systems;

Paper 2 argues that the pressures identified by (Denić & Szymanik, 2024) cannot account for the regularity of recursive numeral systems and introduces a different trade-off based on regularity and processing complexity. We analyse recursive numeral systems under these measures through a formal model based on Deterministic Finite Automata (DFA), finding that human systems are significantly more efficient for our measures than previously considered optimal systems;

Paper 3 studies via reinforcement learning the connection between regularity and learnability in the context of recursive numeral systems and argues that most human systems are highly regular because they are inherently easier to learn;

Paper 4 applies similar concepts of Paper 1 to a more open-ended building blocks domain, where two agents are tasked to collaborate to build block towers. Here, we show that through the introductions of procedural abstractions, the agents' languages tend to minimise the same trade-off as Denić and Szymanik (2024).

Chapter 2

Background

2.1 Constrained variation in languages and efficient communication

While languages are widely different, a significant amount of research in cognitive science and linguistics has explored the shared features of languages across all levels of language (Evans & Levinson, 2009). Cross-linguistic features like Zipf's law of abbreviation (Zipf, 1949), typological universals like the "subject-first" phrase structure (Croft, 2002), and countless others, can help us better understand *constrained variation*: as different languages have emerged and evolved with commonalities, these shared features can help us better understand the cognitive and linguistic constraints that shape languages.

A general framework to understand how constrained variation came to be is offered by *efficient communication*: languages share similar features not because of the features themselves, but because languages are shaped by the same *pressures*, specifically communicative ones (Gibson et al., 2019; Kemp & Regier, 2012). In particular, *efficient communication* argues that languages are simultaneously shaped by a pressure to be *informative*, to maximise communicative accuracy, and by a pressure to be *simple*, to minimise cognitive load. These pressures clearly trade off with each other: a simpler language will have fewer terms or rules, but at the same time, it will be harder to refer accurately to many concepts, and some might have to fall under the same umbrella term. This notion of communicative efficiency is grounded in the information-theoretic communication model of Shannon (1948): a speaker samples a concept and chooses a message to communicate to the listener, who tries to decode the message into the corresponding concept. In these terms, simplicity is the number of terms chosen to carve the concept space, whereas informativeness is the difference between the listener's reconstruction and the concept referred to.

Efficient communication has been successfully applied in many semantic domains, from colour-naming terms (Regier et al., 2015; Zaslavsky et al., 2019), kinship categories (Kemp & Regier, 2012), household containers' names (Xu

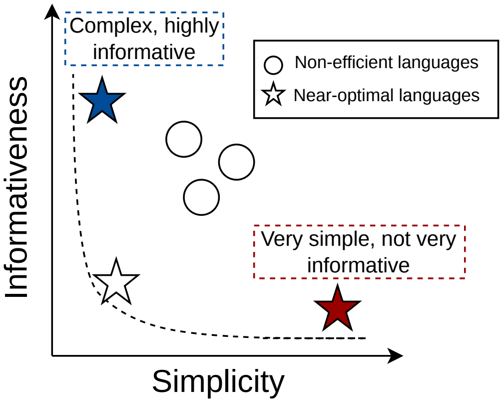

Figure 2.1: An illustrative example of the simplicity-informativeness trade-off in efficient communication: human languages are on average much more efficient than other unattested but possible ones, as they usually lie closer to the Pareto frontier (the dashed curve).

et al., 2016), person systems (Zaslavsky et al., 2021), grammatical markers (Mollica et al., 2021), quantifiers (Steinert-Threlkeld, 2021), indefinite pronouns (Denić et al., 2022), boolean connectives (Uegaki, 2022), spatial demonstratives (Chen et al., 2023), numeral systems (Xu et al., 2020), and even sign language (Yin et al., 2024). Most research first operationalises the two pressures for simplicity and informativeness into two domain-specific measures which can then be used in conjunction with human data from the same domain to compare efficiency across languages. A more general framework is the IB model from Zaslavsky et al. (2018), which argues that languages optimise the information bottleneck (IB) trade-off between accuracy and complexity of the lexicon.

With pressures quantified into measures, most of this research then generates unattested artificial languages to compare the overall space of possible languages to that of human systems. Human languages then usually appear on average much closer to optimality than other unattested systems. Oftentimes a *Pareto frontier* is also either mathematically derived or approximated, which represents the hypothetical line over which languages cannot exist, as improving for one measure would mean automatically worsening for the other. We illustrate the simplicity-informativeness trade-off with an example in Figure 2.1: in all domains referred to above, human languages usually lie close to Pareto frontier as the stars of Figure 2.1, making them near-optimal. Other unattested but plausible alternative languages instead usually are mostly non-efficient, with the closest ones to the Pareto frontier usually sharing some human-like characteristics.

Finally, a very influential element in efficient communication research is the *communicative prior*: this is effectively defined as a probability distribution over the possible concepts of the interested domain. It signifies, for each concept, the need we have to communicate it. Gibson et al. (2017), for example, shows how across languages, warm colours are communicated more efficiently than cool colours, because objects, which we typically are interested in referring to, are usually warm-coloured. This in turn influences their communicative need, effectively making it more likely to adopt more informative terms for warm colours. Another example is that of numerical quantities, which are in general believed to be governed by a left-skewed prior (Dehaene & Mehler, 1992): smaller numbers have a higher need to be communicated than larger numbers. Then, it is not a coincidence that in approximate and exact restricted numeral systems - which are systems that have either only inexact terms or a few exact ones followed by a non-recursive, catch-all term - the most informative terms refer for the overwhelming majority to lower quantities (Xu et al., 2020). This is of course also connected to Zipf's principle of least effort (Zipf, 1949), as languages tend to dedicate shorter, distinct and crucially more informative forms to concepts they need to express most often.

2.1.1 Efficient communication and productive morphosyntax

While efficient communication offers an attractive theoretical framework to understand constrained variation in language, most work based on this theory assumes monomorphemic utterances in their communication model. This means that there can be no overlap between forms, and these ought to be analysed as holistic. This can cause issues in both how simplicity and informativeness are operationalised. For simplicity, there might be different degrees of cognitive load to learn and store in memory a holistic term like the Swedish *mor* for mother versus a compositional one that reuses previously stored terms like the Swedish *mormor* for grandmother on the mother's side (whose literal translation is *mother's mother*). Something similar can be said for informativeness, as a language with morphosyntactically complex terms can allow speakers to be more informative about a variety of different meanings without having to introduce new ones (as the Swedish *mormor* exemplifies well).

A crucial contribution to this issue and for our thesis is that of Denić and Szymanik (2024), who reformulate this trade-off in the semantic domain of *recursive* numeral systems. In their work, they notice that some linguistic systems exhibit maximal informativeness, even though the number of lexicalised terms is low. Within these linguistic systems, talking about pressures for simplicity and informativeness trading-off loses meaning, because informativeness always stays maximal. In fact, by definition, with recursive numerals we can think of any numerical quantity, no matter how large, and express an utterance to exactly represent it, hence their maximal informativeness. Xu et al. (2020) originally argued that recursive numeral systems are just a more informative and more complex category of numeral systems, whereas other categories like *approximate* and *exact restricted* numeral systems are simpler and less informative. Also, crucially, all the terms in non-recursive numeral systems are monomorphemic, which makes the analysis more in line with previously discussed efficient communication research. But Denić and Szymanik (2024) argue how in reality natural recursive numeral systems lie much further away from the Pareto frontier than their other natural, non-recursive counterparts.

They instead introduce *average morphosyntactic complexity* as the competing pressure that shapes recursive numeral systems, operationalising simplicity as lexicon size. More specifically, the underlying morphosyntactic forms of recursive numerals are composed of a combination of rules and monomorphemic terms, which are the ones that must be learned by heart in a language (e.g. *one*, *two*, ..., *ten* in English). With these terms, we can then compose any numeral that refers with perfect accuracy to any given number (e.g. *two-hundred forty-six*, whose underlying morphosyntactic form is $2 * 100 + 4 * 10 + 6$). Given this formalisation, operationalising the two competing pressures is then very easy. As for lexicon size, it is simply calculated as the number of monomorphemic forms in a specific numeral systems. For example, Mandarin and many other systems lexicalise the numbers from 1 (一 , *yī*) to 10 (十, *shí*), combining them through multiplication and addition to form the remaining numerals[1] (e.g. 11, 十一, *shíyī* and so on). At the same time, the average morphosyntactic complexity of a system, which is interchangeable for average utterance length, is defined as such:

$$\overline{UL}(L) = \sum_{n \in [1,\ldots,99]} P(n) \cdot UL(n, L) \tag{2.1}$$

where $UL(n, L)$ is the utterance length (or morphosyntactic complexity) of number n under language L, and $P(\cdot)$ is the prior over numbers, an analogue of the communicative prior, usually defined as $P(n) \propto n^{-2}$ (Dehaene & Mehler, 1992).

The authors then demonstrate how natural recursive numeral systems are near-optimal in terms of this new trade-off of lexicon size and average utterance length. To do so, they first annotate 128 numeral systems collected in (Comrie, 2022) following Hurford's grammar (J. Hurford, 2007; J. R. Hurford, 1975):

$$\begin{aligned} &\textbf{Num} \rightarrow D \mid Phrase \mid Phrase + Num \mid Phrase - Num \\ &\textbf{Phrase} \rightarrow M \mid Num * M \end{aligned} \tag{2.2}$$

Here, morphemes of category M are multipliers, corresponding to bases, and D are digits. Once natural numeral systems are annotated following these rules, lexicon size is just the cardinality of the union set of Digits and Multipliers $|D \cup M|$, whereas the morphosyntactic complexity of a numeral is just the sum of the monomorphemic morphemes (of type D or M) and the combinator morphemes that compose it. For example, $4 * 10 + 3$'s utterance length is 7. They then generate artificial systems, again according to the grammar of

---

[1] Note that both Xu et al. (2020) and Denić and Szymanik (2024) restrict the range of numbers to 1-99 for tractability and also because the influence of higher terms on either work's complexity measures is close to zero because of the left-skewed need distribution considered (Dehaene & Mehler, 1992).

Equation 2.2, and use an evolutionary algorithm to approximate the Pareto frontier (Denić et al., 2022; Steinert-Threlkeld, 2021), and finally show how human systems are all near-optimal in terms of the new trade-off, while other unattested artificial systems on average lie far away from the Pareto frontier.

An issue with this work is that several optimal artificial systems that compose the Pareto frontier are not very human-like, as they often exhibit alternating bases or a very low level of systematicity. Yang and Regier (2025) attempt to resolve this issue by constraining Hurford's grammar to generate only human-like recursive numeral systems. Another detail is that because Hurford's grammar overgenerates potential forms for a single number, Denić and Szymanik (2024) always select the shortest construction for each numeral. While this choice often aligns with Hurford's packing strategy(J. Hurford, 2007) — a cross-linguistic constraint over numerals construction — the mechanisms differ. Ideally, one would hope that communicative and cognitive pressures could explain the rise of the packing strategy as an optimal byproduct rather than having to enforce it as a given universal. In general, however, this work is an important step towards extending the framework of efficient communication to semantic domains that have productive and recursive morphosyntax, and more generally, to natural language, which is itself highly recursive.

2.1.2 Simplicity as systematicity

A part of this thesis explores the concept of simplicity in the context of semantic domains like recursive numeral systems. As previously stated, linguistic systems in these domains make use of productive morphosyntax to produce different-length utterances that combine smaller lexicalised terms to refer to non-lexicalised concepts, i.e. larger numbers. In this sense, if simplicity is to be considered solely as the number of lexicalised terms, then a missing factor that makes these linguistic systems human-like would be the *systematicity* with which these sub-forms are combined together.

An influential work for this thesis is that of Brighton (2003), which explores simplicity as a driving force for linguistic evolution and the emergence of systematicity within language. In their work, Brighton (2003) draws inspiration from the shortest description length principle (Li, Vitányi et al., 2008), which essentially connects simplicity to the theory of *Kolmogorov complexity*: the simplicity of a language is related to the **shortest** program capable of generating all strings in it. Unfortunately, while Kolmogorov complexity theory is elegant and sound, it is also impractical, so (Brighton, 2003) approximated some of these intuitions by utilising the *Minimum Description Length* (MDL) principle (Rissanen, 1978). The MDL principle essentially states that for some observed data D, the best hypothesis is the one that minimises the sum of the encoding length of the hypothesis itself and of the data encoded in terms of the hypothesis. van de Pol et al. (2023), for example, demonstrates the explanatory power of the MDL principle in the domain of logical quantifiers, showing that semantic universals exhibit in general the shortest description length.

In order to adapt this theory in our work, we employ partial Deterministic Finite Automata (DFA) to represent the underlying grammar of the linguistic

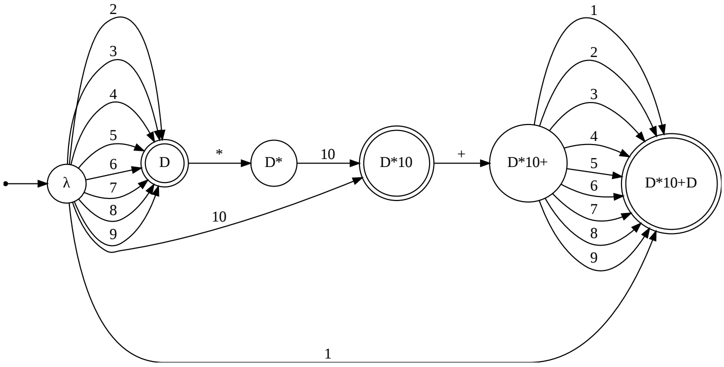

Figure 2.2: The minimal partial DFA of Mandarin numerals.

systems considered that generates all and only valid strings in it. We are specifically interested in the *minimal* partial DFA, as that relates to the best hypothesis indicated by the MDL principle. An example of a DFA for the numeral systems analysed in our work is shown in Figure 2.2: from the entry state λ, one can generate all and only the valid terms that form the first 99 Mandarin numerals by going through consecutive pairs of transitions (the lines) and states (the circles), halting only at one of the final states (the double circles). Then simplicity can be easily characterised as how *compressible* this description language is: a smaller DFA will reuse most of the same structures to encode different data points, making in turn the encoded data have a higher degree of systematicity. We quantify the opposite of this measure (irregularity) into the number of bits required to represent this minimal DFA:

$$\underbrace{|Z|\big(2 \cdot log_2|S| + log_2|\Sigma|\big)}_{\text{(A)}} + \underbrace{log_2|S|}_{\text{(B)}} + \underbrace{|S|}_{\text{(C)}} \tag{2.3}$$

where (A) is the number of bits required to encode the DFA's transition table, with Z being the set of transitions, S the set of states and Σ the set of symbols (unary-length morphemes and combinators), while (B) and (C) are the numbers of bits required to encode the identity of initial state and accepting states.

2.1.3 Simplicity as learnability

Another related theory that attempts to explain constrained variation and which is relevant to this thesis is that of *learnability*: some properties of languages are more prevalent than others because they are easier to learn. This hypothesis is very attractive because a main drawback of efficient communication is the domain-specificity of the complexity measures usually employed. Instead, framing simplicity in terms of learnability could condense all the domain-specific measures into a more general one, possibly rooted in learning theory. The issue arises when quantifying learnability, as there is no obvious way to do so. Some experimental research with humans has connected systematicity itself

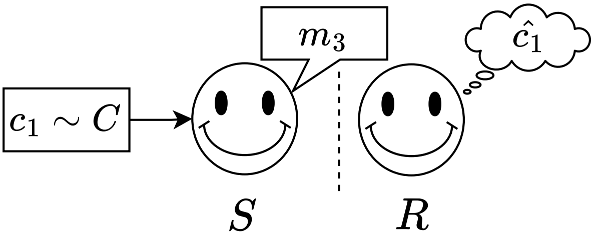

Figure 2.3: A round of the Lewis signalling game (Lewis, 1969): a speaker S observes a concept c and encodes it into a message m. The receiver R observes the message and selects the concept $\hat{c}$ that they think m refers to.

to learnability (Raviv et al., 2021; Shepard et al., 1961; Wang et al., 2025), but then we fall back to the problem of having to define simplicity for any domain in which we might want to study the learnability of languages. Recent research has shown how deep learning models can come to our aid when quantifying learnability (Galke et al., 2024; Osmelak et al., 2026; Steinert-Threlkeld, Szymanik et al., 2019; Steinert-Threlkeld & Szymanik, 2020). These methods usually utilise off-the-shelf deep learning models like recurrent neural networks (RNNs) to learn different languages, and quantify learnability as empirical learning metrics of these networks. Similarly to efficient communication results, human and human-like languages are often more learnable than non-human ones. As for what human-likeness means, this is again domain-specific: as an example, for colour terms *convexity* is cross-linguistically prevalent (Gärdenfors, 2000; Jäger, 2010), and can be explained by ease of learning (Carlsson et al., 2024; Steinert-Threlkeld & Szymanik, 2020); the same can be said for semantic universals (Steinert-Threlkeld, Szymanik et al., 2019) and systematicity between forms (Osmelak et al., 2026).

2.2 Language evolution, communication and Reinforcement Learning

While efficient communication and learnability offer powerful theories to explain the cross-linguistic prevalence of some linguistic features, a question may arise: how do these pressures actually shape linguistic systems? Languages are not holistic entities but a byproduct of the populations that speak them, so it is not obvious how these pressures manifest. While an idea might be to study how efficiency has evolved in historical corpora, diachronic data is scarce and usually not domain-specific but more general, although some exceptions exist (Zaslavsky et al., 2022). An alternative is to *simulate* language evolution in a vacuum, in order to easily isolate the effects of specific pressures on language.

A plausible model to use is suggested by the communication model of Shannon (1948), which is the framework of efficient communication: in this, language is framed as a two-player game where a speaker and a listener learn how to coordinate over a set of concepts. This functional view on language can be further formalised as a *Lewis signalling game* (Lewis, 1969) (cf. Figure 2.3): in this game, a sender S and a receiver R learn how to coordinate by developing a common convention to refer to concepts. A round goes as follows: a concept c is observed by S, which then chooses a message m from the set of messages M to encode the meaning. The receiver R observes m and acts accordingly by choosing an action a from the set of actions A, which usually corresponds to choosing the concept $\hat{c}$ from C that S thinks m refers to.

A plenitude of studies have simulated the evolution of artificial languages in a vacuum, either through experimental research with humans (Fay et al., 2010; Kirby et al., 2008) or computational research with artificial agents (Jäger et al., 2011; Kirby, 2002; Kirby et al., 2015; Smith et al., 2003). A fundamental body of work for this thesis is called *emergent communication*, which employs multi-agent deep reinforcement learning (RL) (Sutton & Barto, 1998) agents to study language emergence and/or evolution (Carlsson, 2024; Chaabouni et al., 2020; Foerster et al., 2016; Jorge et al., 2016; Lazaridou & Baroni, 2020; Lazaridou et al., 2017).

In RL, an artificial agent learns how to solve a task, framed in terms of maximising a reward, by interacting with an environment $\mathcal{E}$. Generally, at each timestep t an observation o_t is sampled from $\mathcal{E}$ and observed by the agent. At their disposal, they also have a set of actions A to sample from, as a response to o_t. The sampled action a_t then yields a reward r_t, and could potentially change the environment state s_{t+1} (which can correspond to o_{t+1} or could be a superset of it). An extension of this framework which we adopt in **Papers 1 and 4** is multi-agent RL, which simply involves multiple agents jointly solving a collaborative task. It is simple to note how this general framework can be easily adapted to the aforementioned models stemming from Shannon (1948)'s: the sender S will now sample a message m conditioned on the observation o representing a concept sampled from a distribution; the receiver R will then receive m and in turn sample the concept $\hat{c}$ conditioned on m. Both agents S and R are rewarded if and only if c and $\hat{c}$ coincide (or they might receive partial reward if they are "close enough", depending on whether the concept space is ordered). The agents will each learn a policy, respectively π_S that maps from concepts to messages and π_R that maps from messages back to concepts, that they will optimise through interactions with the environment.

Throughout the thesis we make use of the term *contextual bandit* to refer to this learning setup. In fact, throughout this thesis our reinforcement learners are always playing a contextual bandit: each one observes a context, plays an action (or an arm of the bandit) according to their policy, and obtains a reward based on it, which they in turn will use to update their policy. The difference with the general RL setup is that the environment state does not transition based on the agent's actions; rather, a new independent context is drawn at each step.

As our learning algorithm we adopt REINFORCE (Williams, 1992) throughout the thesis. It is a policy gradient RL algorithm that utilises reward to update a policy, usually parametrised by a neural network, following the rule:

$$\theta_{t+1} = \theta_t + \alpha \nabla \log \pi_{\theta_t}(a_t|o_t) r_t \tag{2.4}$$

where α is the learning rate. In **Paper 4** we adopt the Gumbel-Softmax estimator (Jang et al., 2017) as an alternative algorithm to train discrete communication. Gumbel-Softmax provides a continuous differentiable approximation of a categorical distribution. The primary advantage of this approach is that it still conserves the discrete communication channel while allowing for backpropagation of the gradients through the communication bottleneck, which can reduce high variance gradient estimates that can hinder training in REINFORCE when the task is sufficiently hard.

Carlsson (2024) motivates the use of RL to explore the mechanisms that lead to efficient languages in several ways. Firstly, the functional view of language envisioned by efficient communication pairs well with the reward-maximising framework of reinforcement learning: language is used as the means to solve a task (i.e. refer to specific concepts) and thus users are pushed to use it efficiently, and eventually modify it according to these pressures. Also, Carlsson (2024) argues that in the multi-agent setting of emergent communication agents have both an implicit bias for simpler languages in order to more easily coordinate between each other, and also to maximise informativeness as a sub-goal of maximising reward, as a more informative message will more likely lead to a higher reward. Finally, utilising RL to estimate learnability can be connected to learning theory on sample complexity (Kakade, 2003), which essentially studies the number of samples required for different RL algorithms to learn across different learning setups.

2.3 Efficient communication and procedural abstractions in collaborative games

A second aim of this thesis is to extend the emergent communication framework to other complex domains. One of these is the collaborative framework of McCarthy et al. (2021), where two users collaborate through natural language to learn how to compose towers with building blocks. It is an asymmetric framework, as only one user, the *architect*, can visualise the tower that ought to be built. The architect then has to use natural language to communicate to the second user, the *builder*, how to procedurally recreate the tower from scratch using only two different building blocks. McCarthy et al. (2021) notes how humans when repeatedly collaborating on this task tend to increasingly shorten their language over time. They do so by introducing **procedural abstractions**: from using atomic instructions that refer directly to the smallest building blocks, humans extend their vocabulary to include larger abstractions that refer to combinations of blocks that appear frequently in their interactions, thus introducing terms that carry more information and allow for shorter utterances. This phenomenon directly connects to the average utterance length-lexicon size trade-off introduced by Denić and Szymanik (2024) and discussed in Section 2.1.1.

A computational approach we adopt in this work to study the phenomenon of procedural abstraction is *library learning*. It extends the classic approach of program synthesis by extracting common, repeated subprograms into new library terms that can be subsequently used when sampling programs to solve new tasks. A deep learning-based library learning approach relevant to our thesis is the DreamCoder framework of Ellis et al. (2021), which learns to efficiently solve increasingly harder and more complex tasks by learning procedural abstractions that allow the learning algorithm to more efficiently explore the space of solutions and sample concise programs that can solve these increasingly harder tasks.

Chapter 3

Summary of Papers

3.1 Learning Efficient Recursive Numeral Systems via Reinforcement Learning

In this paper we extend the framework of Carlsson et al. (2021) to the domain of recursive numeral systems. In their work, Carlsson et al. (2021) employed artificial agents using deep RL to communicate about numerical quantities. They showed that the sole pressure for successful communication made the agents develop artificial protocols that resembled approximate and exact restricted systems. Given the findings of Denić and Szymanik (2024), we aim to provide a computational model based on RL that demonstrates how efficient recursive numeral systems are consistently preferred over inefficient ones when used as a basis for communication.

We utilise a slightly modified version of Hurford's grammar (J. Hurford, 2007; J. R. Hurford, 1975) to produce recursive numeral systems that are then used to sample different numerals, to be used as input for the artificial speaker. The speaker then encodes them neurally and sends the message to an artificial listener, which then predicts a numerosity (cf. Figure 3.1). The speaker can actually use at any point two alternative grammars, which will correspond to two different sets of numerals, to communicate. We model this as a simple *bandit* algorithm, where communicative success with each grammar determines the policy of the speaker that governs which grammar it should use. After each communication round, the most successful grammar is kept and the alternative one is substituted with another one, which we obtain by performing a slight modification to the (D, M) pair that represents the grammar.

Across our results (Figure 3.2), we show how more efficient recursive numeral systems are consistently preferred by our RL agents. This work both aligns with and complements the theoretical findings of Denić and Szymanik (2024) in this domain, providing an example of how the effects of the pressures predicted by them can be reproduced with a simple computational multi-agent model. Our results also connect to work on ease of learning (Steinert-Threlkeld, Szymanik et al., 2019; Steinert-Threlkeld & Szymanik, 2020), as more efficient recursive

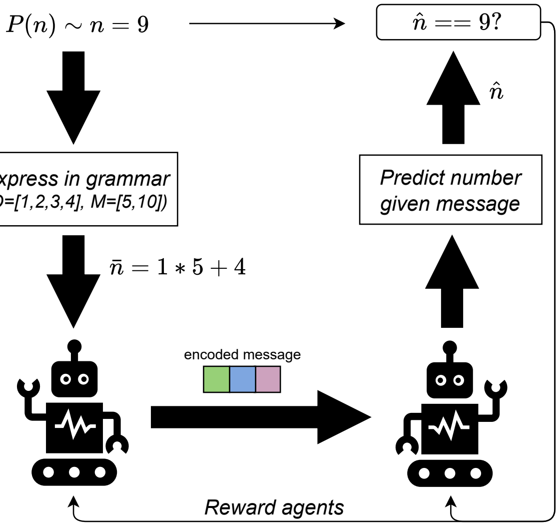

Figure 3.1: An example of a communication round: a number, e.g. $n = 9$, is sampled from the need distribution. The corresponding representation in terms of the current (D, M) is passed to the speaker, which encodes it neurally. The listener receives this message and outputs the number it thinks it refers to. Both agents are rewarded if the guess is correct. Taken from Silvi et al. (2025).

numeral systems are easier to learn for our two-agent RL model.

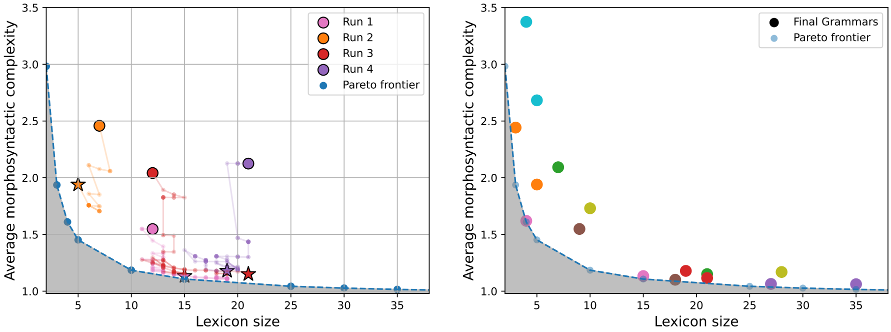

(a) Evolutionary trajectories.

(b) Final grammar convergence.

Figure 3.2: Combined results showing language evolution and final convergence toward the Pareto frontier. The grey area represents unobtainable configurations. (a) Some trajectories showing the evolution of languages (not all trajectories are included for visual clarity). (b) Final languages plotted in terms of lexicon size and average morphosyntactic complexity, colour coded based on their starting (D, M) pair. Taken from Silvi et al. (2025).

3.2 Recursive numeral systems are highly regular and easy to process

In this paper we address the issue raised by Denić and Szymanik (2024), where recursive numeral systems that are optimal for the pressures they introduce (which we analyse in Section 2.1.1) are not necessarily human-like, indicating that their pressures, while accounting for some aspects that are shared across human recursive numeral systems (namely their lexicon size-average utterance length efficiency), cannot account for other shared features. Specifically, we argue that **regularity** is the key ingredient missing, as human systems often reuse the same sub-forms to compose larger numbers. We apply a model based on the Minimum Description Length framework (Brighton, 2003) and Deterministic Finite Automata (DFA) to quantify regularity, and posit that it might be in contention with *processing complexity*, which is roughly the number of possible paths in the DFA to parse a numeral.

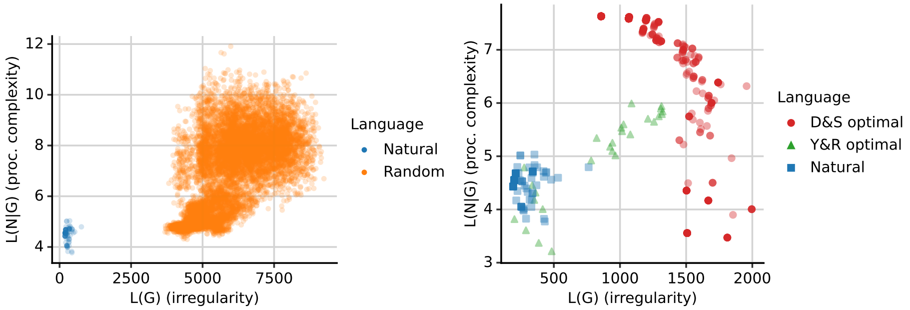

(a) Natural vs. random languages.

(b) Natural vs. optimal systems.

Figure 3.3: Comparison of irregularity (x-axis) and processing complexity (y-axis) across different language models. (a) Natural languages (blue) plotted alongside 10,000 randomly generated baseline artificial languages (orange). (b) Natural systems (blue, square) compared with optimal systems in Denić and Szymanik (2024) (red, circle) and Yang and Regier, 2025 (green, triangle). Taken from Prasertsom et al. (2026).

We use this model to analyse different sets of human and artificial recursive numeral systems. Our results suggest that human systems are much more efficient than both randomly sampled ones, Figure 3.3(a), and artificial languages that would be optimal in terms of the lexicon size-average utterance length trade-off, Figure 3.3(b). This suggests that our measures can both account for the efficiency of human recursive numeral systems and can differentiate the previously considered recursive numeral systems which lacked human-likeness (i.e., systematicity) from human ones. More generally, it suggests that regularity is a key aspect of linguistic systems and that efficient communication should account for it by not analysing forms atomically.

3.3 Evaluating the relationship between regularity and learnability in recursive numeral systems using Reinforcement Learning

In this paper we explore the connection between regularity and ease of learning in the domain of recursive numeral systems. We posit that regularity should ease learning, and expect that, when quantifying learnability with machine learning models, we should see a correlation with systematicity as measured with the model introduced in **Paper 2**.

We introduce a single-agent RL framework to estimate learnability of different types of recursive numeral systems. We operationalise learnability as the Area Under the Curve of the test accuracy, and let our agents learn a system over a fixed number of iterations by sampling from the need distribution of Dehaene and Mehler (1992), and guessing a numerical quantity based on the numeral of the corresponding numeral system we are attempting to quantify learnability for. After each training epoch, the agent is tested on numerical quantities sampled from the test distribution, so we can quantify learnability as the AUC of this statistic.

Our results show that when considering the same skewed distribution of Dehaene and Mehler (1992), no correlation between regularity and learnability seems to emerge. Instead, introducing a uniform test distribution (while still training on the inverse power-law one from the literature) produces results more in line with our expectations, although regularity seems to influence learnability more than other measures only in the high regularity range (Figure 3.4(a)). Instead, when regularity is low, other measures like average utterance length seem more important for learnability (Figure 3.4(b)). This makes sense, as regularity can be exploited for learning only if there is enough systematicity between forms; otherwise, other features might be more exploitable for more efficient learning, for example how short the forms are.

This duality between the inverse exponential need distribution used for training and uniform generalising distribution for testing is an effective setup to formalise the dual pressures that probably shape the learning of these recursive numeral systems: a need to communicate about numerical quantities that most likely follows Dehaene and Mehler (1992) and a need to be able to generalise to any and every numerical quantity in an exact manner.

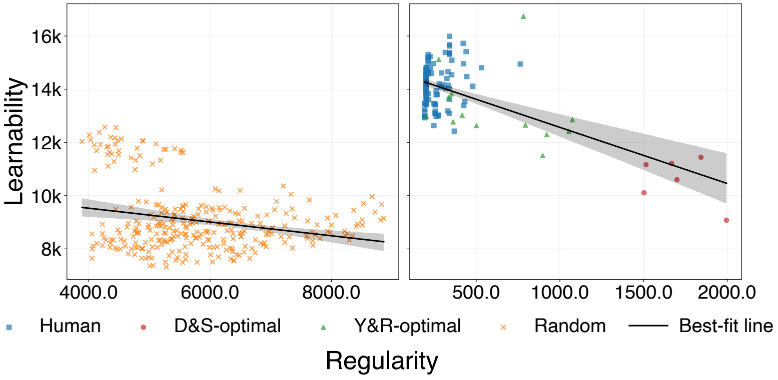

(a) Regularity.

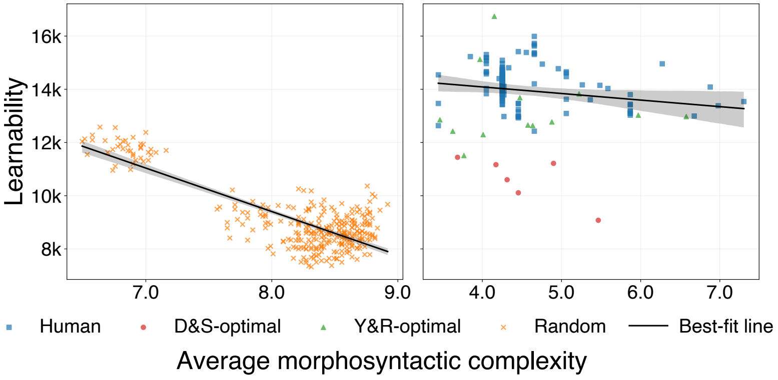

(b) Morphosyntactic complexity.

Figure 3.4: Average learnability for random baseline (left within panels) and highly regular systems (right within panels), under the inverse power law training distribution and a uniform test distribution. (a) Evaluated against regularity. (b) Evaluated against morphosyntactic complexity.

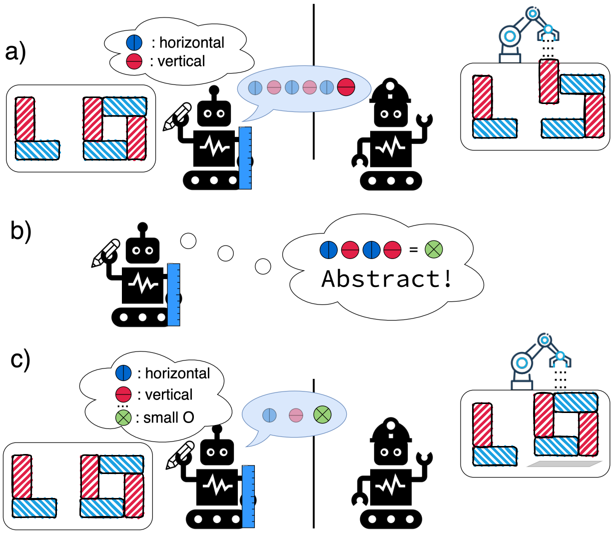

Figure 3.5: Two artificial agents playing the architect-builder game, starting from a small artificial language. Initially, the architect messages refer to horizontal or vertical blocks (a). After multiple interactions, the architect tries to introduce an abstraction (b), which after a learning period allows for shorter communication to solve the task (c). Taken from Thomas et al. (2025).

3.4 PACE: Procedural Abstractions for Communicating Efficiently

In this paper we explore the role of procedural abstractions in shaping collaboration in a procedural building block task. We extend the two-player framework of McCarthy et al. (2021) to artificial agents. In this setting, an architect visualises a tower made of building blocks and has to communicate a procedure to the builder to construct it. Instructions can be atomic and refer to single building blocks, or *abstractions* that refer to chunks of blocks, reducing the number of instructions needed to complete a tower but introducing uncertainty, as the listener might not know how to interpret said abstractions (cf. Figure 3.5).

The base model of communication is very similar to the one of **Paper 1**, but with supervised learning paired with the Gumbel-Softmax relaxation instead of REINFORCE to train our agents. The architect and builder first learn how to coordinate on the base language consisting of terms that refer only to single building blocks, but then are allowed to introduce new terms to their vocabulary. This is done explicitly by modelling the task of decomposing the target tower into a *program*, which is then communicated instruction by instruction. Then, via library learning (Ellis et al., 2021), existing programs are analysed to discover common subprograms that are consistently communicated correctly. They are then explicitly abstracted into new terms and new programs are introduced to the architect to describe towers containing them. The architect's choice of program is then modelled as a simple contextual bandit with contextual arms, where the architect learns to pick programs that are composed of instructions that yield most communicative success.

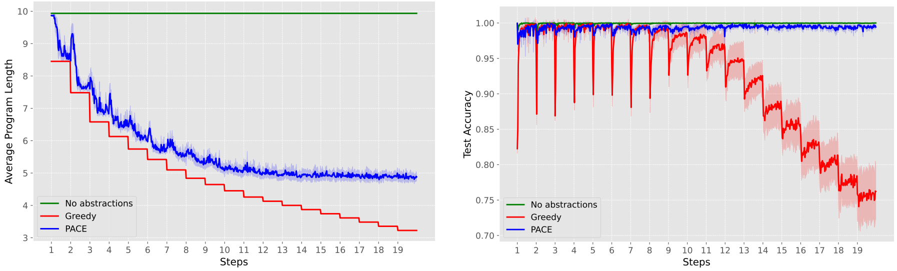

Figure 3.6: Comparison between PACE, *Greedy* and *No abstractions* in terms of program length and test reward over time. Line indicates mean value and shaded regions indicate the 95% confidence interval. Taken from Thomas et al. (2025).

Our results show that, compared to baselines that either do not introduce abstractions or only use programs that contain abstractions, our computational model's results are the most in line with the ones of McCarthy et al. (2021) when testing humans on this task: over successive interactions, pairs of architects and builders learn to communicate using increasingly more succinct language by introducing abstractions (Figure 3.6). Moreover, as with humans our agents' languages converge after introducing some abstractions, as there is a law of diminishing returns that makes it inconvenient and inefficient to further abstract terms that might be too infrequent. While new abstractions might carry more informativeness, they will also carry uncertainty that frequently cannot be overcome by learning their meaning via interactions, as their frequency might just be too low.

In the paper we also show that in general the abstractions that are most likely to be introduced in our agents' shared protocols are the ones that are most efficient in terms of the same trade-off between average utterance length (here measured as average program length) and lexicon size, where the lexicon is the number of different instructions the architect can use to build programs (Figure 3.7). These results show how pressures for efficient communication can explain what procedural abstractions are most likely to be introduced to facilitate collaboration.

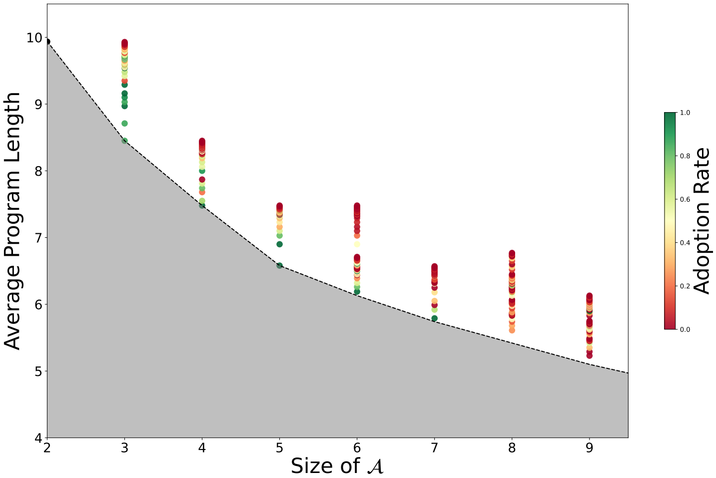

Figure 3.7: Adoption rate of possible abstractions by language size. As the language grows, fewer new abstractions are adopted. The dashed line represents the interpolation of the (discrete) Pareto Frontier calculated as the trade-off between average program length and size of the language $\mathcal{A}$. The grey area represents unachievable languages. Taken from Thomas et al. (2025).

Chapter 4

Discussion and Future Avenues

Throughout our papers we consistently show how domains like recursive numeral systems or the architect-builder framework of McCarthy et al. (2021) are highly promising avenues to explore efficient communication theories and adjacent concepts, such as the connection between ease of learning and regularity in **Paper 3**. **Paper 2** and **Paper 3** demonstrate how systematicity has a key role to play in explaining cross-linguistic features in domains with productive linguistic systems: without it, models of efficient communication cannot fully account for human-likeness. At the same time, systematicity seems to ease learning, indicating why it might be consistently high throughout all human recursive numeral systems.

Future work could aim to provide a holistic model for the emergence and adoption of different numeral systems, from approximate and exact-restricted to recursive ones. Studying how other modalities besides communication and learning shape these systems, like the process of **counting** (Holt & Barner, 2025), could be an exciting avenue and might be easily modelled with artificial RL agents, extending the model of **Paper 1**. As Hurford argues that counting might give rise to the packing strategy (J. Hurford, 2007), a related question one could explore is if we could reproduce the emergence of this strategy with computational models just from communicative and learning pressures instead of having to impose it as a constraint. At the same time, a holistic model of efficient communication that incorporates all three types of numeral systems into one would be an exciting extension of **Paper 2**. To achieve this, the full model of Brighton (2003) could be employed, as it allows for a form to refer to multiple meanings (such as *muthaa*, which refers to numerical quantities larger than 4 in Kayardild, an exact-restricted system).

The results of **Paper 2** could be extended to find the Pareto frontier of the optimal systems in terms of regularity and processing complexity, but extending the approach of Denić and Szymanik (2024) is non-trivial, as exploring differing levels of regularity would mean abandoning the high-level (D, M) representation derived from Hurford's grammar. One could employ available frontier language models to propose modifications to whole systems instead, leveraging approaches similar to AlphaEvolve (Novikov et al., 2025).

Finally, the current rise of large language models presents an exciting research prospect if pursued correctly. Recent research has studied the evolution of LLMs' language and how they can be meaningful models of language evolution if framed correctly (Imel & Zaslavsky, 2026; Ren et al., 2024). An obvious extension of **Paper 2** and **Paper 3** in this direction would be to study the bias for systematicity that these models possess, and compare it to human systems.

A similar argument can be made for **Paper 4**. As LLMs collaborate with humans via repeated interactions, an interesting question would be to discern if they employ procedural abstractions in a progressively more efficient manner like humans over repeated interactions.

Another timely question within the study of efficient communication that could benefit from deep learning approaches is quantifying the need distribution that governs different semantic domains. These are often assumed to be universal, and oftentimes uniformly distributed for simplicity's sake. However, geographical factors, cultural and historical differences and many other factors (including randomness) are most certainly key ingredients to the *variation* side of constrained variation, making languages unique despite their similarities, some of which we have explored in this thesis. Analysing how variation in these factors may influence the linguistic systems that emerge at the end of an evolutionary process would be a very interesting future contribution. At the same time, frameworks like the one of **Paper 4** could also be extended to systematically analyse the degree of difficulty to build conventions between partners that have developed specialised languages from environments with differing priors over the same concepts. The same prior over numbers from Dehaene and Mehler (1992) has been shown to present regular peaks across languages over numbers like 10, 12, 15, 20, 100 and so on, because of a *roundness effect* (Jansen & Pollmann, 2001); these should be studied in conjunction with RL agents if we desire the linguistic system they converge upon to be human-like. Finally, a framework of language evolution we have yet to employ is the one of *iterated learning* (Kirby et al., 2008; Smith et al., 2003), which would complement our results by adding the crucial dimension of vertical, inter-generational cultural transmission. Previous results have demonstrated that both within humans (Kirby et al., 2015) and deep RL agents Carlsson et al., 2024 the combined pressures of communication and cultural transmission are necessary to obtain a linguistic system at the end of an evolutionary process that is simultaneously information-theoretically optimal and structurally human-like.

Bibliography

Brighton, H. (2003). *Simplicity as a driving force in linguistic evolution* [PhD thesis]. The University of Edinburgh. http://hdl.handle.net/1842/23810 (cit. on pp. 9, 18, 25).

Carlsson, E. (2024). *Reinforcement learning: Efficient communication and sample efficient learning* [Doctoral dissertation]. Chalmers University of Technology. (Cit. on pp. 3, 12, 13).

Carlsson, E., Dubhashi, D., & Johansson, F. D. (2021). Learning approximate and exact numeral systems via reinforcement learning. *Proceedings of the Annual Meeting of the Cognitive Science Society, 43*. https://doi.org/10.48550/arXiv.2105.13857 (cit. on pp. 4, 15).

Carlsson, E., Dubhashi, D., & Regier, T. (2024). Cultural evolution via iterated learning and communication explains efficient color naming systems. *Journal of Language Evolution*, *9*(1-2), 49–66. https://doi.org/10.1093/jole/lzae010 (cit. on pp. 11, 26).

Chaabouni, R., Kharitonov, E., Bouchacourt, D., Dupoux, E., & Baroni, M. (2020). Compositionality and generalization in emergent languages. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, 4427–4442 (cit. on p. 12).

Chen, S., Futrell, R., & Mahowald, K. (2023). An information-theoretic approach to the typology of spatial demonstratives. *Cognition*, *240*, 105505 (cit. on p. 6).

Comrie, B. (2022). The arithmetic of natural language: Toward a typology of numeral systems. *Macrolinguistics*, *10*(1), 1–35 (cit. on p. 8).

Croft, W. (2002). *Typology and universals*. Cambridge university press. (Cit. on p. 5).

Dehaene, S., & Mehler, J. (1992). Cross-linguistic regularities in the frequency of number words. *Cognition*, *43*(1), 1–29. https://doi.org/10.1016/0010-0277(92)90030-L (cit. on pp. 7, 8, 19, 26).

Denić, M., Steinert-Threlkeld, S., & Szymanik, J. (2022). Indefinite pronouns optimize the simplicity/informativeness trade-off. *Cognitive Science*, *46*(5), e13142 (cit. on pp. 6, 9).

Denić, M., & Szymanik, J. (2024). Recursive numeral systems optimize the trade-off between lexicon size and average morphosyntactic complexity. *Cognitive Science*, *48*(3), e13424. https://doi.org/https://doi.org/10.1111/cogs.13424 (cit. on pp. 4, 7–9, 13, 15, 18, 25).

Douven, I., & Verheyen, S. (2024). Concept learning: Convexity versus connectedness. i. douven, s. verheyen. *Erkenntnis*, *91*, 445–462. https://doi.org/10.1007/s10670-024-00909-1 (cit. on p. 3).

Ellis, K., Wong, C., Nye, M., Sablé-Meyer, M., Morales, L., Hewitt, L., Cary, L., Solar-Lezama, A., & Tenenbaum, J. B. (2021). Dreamcoder: Bootstrapping inductive program synthesis with wake-sleep library learning. *Proceedings of the 42nd ACpM sigplan international conference on programming language design and implementation*, 835–850 (cit. on pp. 14, 21).

Evans, N., & Levinson, S. C. (2009). The myth of language universals: Language diversity and its importance for cognitive science. *Behavioral and brain sciences*, *32*(5), 429–448 (cit. on p. 5).

Fay, N., Garrod, S., Roberts, L., & Swoboda, N. (2010). The interactive evolution of human communication systems. *Cognitive Science*, *34*(3), 351–386. https://doi.org/https://doi.org/10.1111/j.1551-6709.2009.01090.x (cit. on p. 12).

Foerster, J. N., Assael, Y. M., de Freitas, N., & Whiteson, S. (2016). Learning to communicate with deep multi-agent reinforcement learning. *Proceedings of the 30th International Conference on Neural Information Processing Systems*, 2145–2153 (cit. on p. 12).

Galke, L., Ram, Y., & Raviv, L. (2024). Deep neural networks and humans both benefit from compositional language structure. *Nature Communications*, *15*(1). https://doi.org/10.1038/s41467-024-55158-1 (cit. on p. 11).

Gärdenfors, P. (2000). Conceptual spaces: The geometry of thought. *Tijdschrift Voor Filosofie*, *64*(1), 180–181 (cit. on p. 11).

Gibson, E., Futrell, R., Jara-Ettinger, J., Mahowald, K., Bergen, L., Ratnasingam, S., Gibson, M., Piantadosi, S. T., & Conway, B. R. (2017). Color naming across languages reflects color use. *Proceedings of the National Academy of Sciences* (cit. on p. 7).

Gibson, E., Futrell, R., Piantadosi, S. P., Dautriche, I., Mahowald, K., Bergen, L., & Levy, R. (2019). How Efficiency Shapes Human Language. *Trends in Cognitive Sciences*, *23*(5), 389–407. https://doi.org/10.1016/j.tics.2019.02.003 (cit. on pp. 3, 5).

Holt, S., & Barner, D. (2025). Learning a novel number system: The role of compositional rules and counting procedures. *Cognitive Science*, *49*(6), e70071. https://doi.org/10.1111/cogs.70071 (cit. on p. 25).

Hurford, J. (2007). A performed practice explains a linguistic universal: Counting gives the packing strategy [Linguistic perspectives on numerical expressions]. *Lingua*, *117*(5), 773–783. https://doi.org/10.1016/j.lingua.2006.03.002 (cit. on pp. 8, 9, 15, 25).

Hurford, J. R. (1975). *The linguistic theory of numerals*. Cambridge University Press. (Cit. on pp. 8, 15).

Imel, N., & Zaslavsky, N. (2026). Evolution and compression in LLMs: On the emergence of human-aligned categorization. *The Fourteenth International Conference on Learning Representations*. https://openreview.net/forum?id=s7gSTR2AqA (cit. on p. 26).

Jäger, G. (2010). Natural color categories are convex sets. In M. Aloni, H. Bastiaanse, T. de Jager & K. Schulz (Eds.), *Logic, language and meaning* (pp. 11–20). Springer Berlin Heidelberg. (Cit. on p. 11).

Jäger, G., Metzger, L. P., & Riedel, F. (2011). Voronoi languages: Equilibria in cheap-talk games with high-dimensional types and few signals. *Games and Economic Behavior*, *73*(2), 517–537. https://doi.org/https://doi.org/10.1016/j.geb.2011.03.008 (cit. on p. 12).

Jang, E., Gu, S., & Poole, B. (2017). Categorical reparameterization with gumbel-softmax. https://arxiv.org/abs/1611.01144 (cit. on p. 13).

Jansen, C., & Pollmann, M. (2001). On round numbers: Pragmatic aspects of numerical expressions. *Journal of Quantitative Linguistics*, *8*(3), 187–201. https://doi.org/10.1076/jqul.8.3.187.4095 (cit. on p. 26).

Jorge, E., Kågebäck, M., Johansson, F. D., & Gustavsson, E. (2016). Learning to Play Guess Who? and Inventing a Grounded Language as a Consequence (cit. on p. 12).

Kakade, S. M. (2003). *On the sample complexity of reinforcement learning* [Doctoral dissertation]. University College London. (Cit. on p. 13).

Kemp, C., & Regier, T. (2012). Kinship categories across languages reflect general communicative principles. *Science*, *336*(6084), 1049–1054. https://doi.org/10.1126/science.1218811 (cit. on pp. 3, 5).

Kirby, S. (2002). Natural language from artificial life. *Artificial Life*, *8*(2), 185–215. https://doi.org/10.1162/106454602320184248 (cit. on p. 12).

Kirby, S., Cornish, H., & Smith, K. (2008). Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences*, *105*(31), 10681–10686. https://doi.org/10.1073/pnas.0707835105 (cit. on pp. 12, 26).

Kirby, S., Tamariz, M., Cornish, H., & Smith, K. (2015). Compression and communication in the cultural evolution of linguistic structure. *Cognition*, *141*, 87–102 (cit. on pp. 12, 26).

Lazaridou, A., & Baroni, M. (2020). Emergent multi-agent communication in the deep learning era. (Cit. on p. 12).

Lazaridou, A., Peysakhovich, A., & Baroni, M. (2017). Multi-agent cooperation and the emergence of (natural) language. *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*, 1–11 (cit. on p. 12).

Lewis, D. K. (1969). *Convention: A philosophical study*. Wiley-Blackwell. (Cit. on pp. 11, 12).

Li, M., Vitányi, P., et al. (2008). *An introduction to kolmogorov complexity and its applications* (Vol. 3). Springer. (Cit. on p. 9).

McCarthy, W., Hawkins, R., Holdaway, C., Wang, H., & Fan, J. (2021). Learning to communicate about shared procedural abstractions. *Proceedings of the 43rd Annual Conference of the Cognitive Science Society* (cit. on pp. 13, 21, 22, 25).

Mollica, F., Bacon, G., Zaslavsky, N., Xu, Y., Regier, T., & Kemp, C. (2021). The forms and meanings of grammatical markers support efficient communication. *Proceedings of the National Academy of Sciences*, *118*(49), e2025993118. https://doi.org/10.1073/pnas.2025993118 (cit. on p. 6).

Novikov, A., Vũ, N., Eisenberger, M., Dupont, E., Huang, P.-S., Wagner, A. Z., Shirobokov, S., Kozlovskii, B., Ruiz, F. J., Mehrabian, A., et al. (2025). Alphaevolve: A coding agent for scientific and algorithmic discovery. *arXiv preprint arXiv:2506.13131* (cit. on p. 26).

Osmelak, D., Xu, Y., Hahn, M., & McCurdy, K. (2026). Systematicity between forms and meanings across languages supports efficient communication. https://arxiv.org/abs/2601.17181 (cit. on p. 11).

Prasertsom, P., Silvi, A., Culbertson, J., Dubhashi, D., Johansson, M., & Smith, K. (2026, March). Recursive numeral systems are highly regular and easy to process. In V. Demberg, K. Inui & L. Marquez (Eds.), *Proceedings of the 19th conference of the European chapter of the Association for Computational Linguistics (volume 1: Long papers)* (pp. 4873–4885). Association for Computational Linguistics. https://doi.org/10.18653/v1/2026.eacl-long.226 (cit. on p. 18).

Raviv, L., de Heer Kloots, M., & Meyer, A. (2021). What makes a language easy to learn? a preregistered study on how systematic structure and community size affect language learnability. *Cognition*, *210*, 104620. https://doi.org/https://doi.org/10.1016/j.cognition.2021.104620 (cit. on p. 11).

Regier, T., Kemp, C., & Kay, P. (2015). Word meanings across languages support efficient communication. *The handbook of language emergence*, 237–263. https://doi.org/10.1002/9781118346136.ch11 (cit. on p. 5).

Ren, Y., Guo, S., Qiu, L., Wang, B., & Sutherland, D. J. (2024). Bias amplification in language model evolution: An iterated learning perspective. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak & C. Zhang (Eds.), *Advances in neural information processing systems* (pp. 38629–38664, Vol. 37). Curran Associates, Inc. https://doi.org/10.52202/079017-1220 (cit. on p. 26).

Rissanen, J. (1978). Modeling by shortest data description. *Automatica*, *14*(5), 465–471. https://doi.org/https://doi.org/10.1016/0005-1098(78)90005-5 (cit. on p. 9).

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, *27*(3), 379–423. https://doi.org/https://doi.org/10.1002/j.1538-7305.1948.tb01338.x (cit. on pp. 5, 11, 12).

Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological monographs: General and applied*, *75*(13), 1 (cit. on p. 11).

Silvi, A., Thomas, J., Carlsson, E., Dubhashi, D., & Johansson, M. (2025). Learning efficient recursive numeral systems via reinforcement learning. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *47*. https://escholarship.org/uc/item/3cc5053z (cit. on pp. 16, 17).

Smith, K., Kirby, S., & Brighton, H. (2003). Iterated learning: A framework for the emergence of language. *Artificial life*, *9*(4), 371–386 (cit. on pp. 12, 26).

Steinert-Threlkeld, S. (2021). Quantifiers in natural language: Efficient communication and degrees of semantic universals. *Entropy*, *23*(10), 1335 (cit. on pp. 6, 9).

Steinert-Threlkeld, S., Szymanik, J., et al. (2019). Learnability and semantic universals. *Semantics and Pragmatics*, *12*, 401–435. https://doi.org/10.3765/sp.12.4 (cit. on pp. 11, 15).

Steinert-Threlkeld, S., & Szymanik, J. (2020). Ease of learning explains semantic universals. *Cognition*, *195*, 104076. https://doi.org/10.1016/j.cognition.2019.104076 (cit. on pp. 3, 11, 15).

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction* (Second). The MIT Press. (Cit. on p. 12).

Thomas, J., Silvi, A., Dubhashi, D., & Johansson, M. (2025). Pace: Procedural abstractions for communicating efficiently. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *47*(0). https://escholarship.org/uc/item/2mf0m2c3 (cit. on pp. 21–23).

Uegaki, W. (2022). The informativeness/complexity trade-off in the domain of boolean connectives. *Linguistic Inquiry*, 1–23. https://doi.org/10.1162/ling_a_00461 (cit. on p. 6).

van de Pol, I., Lodder, P., van Maanen, L., Steinert-Threlkeld, S., & Szymanik, J. (2023). Quantifiers satisfying semantic universals have shorter minimal description length. *Cognition*, *232*, 105150 (cit. on p. 9).

Wang, F., Kirby, S., & Culbertson, J. (2025). The learning bias for cross-category harmony is sensitive to semantic similarity: Evidence from artificial language learning experiments. *Language*, *101*(1), 109–150 (cit. on p. 11).

Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Mach. Learn.*, *8*(3–4), 229–256. https://doi.org/10.1007/BF00992696 (cit. on p. 12).

Xu, Y., Liu, E., & Regier, T. (2020). Numeral systems across languages support efficient communication: From approximate numerosity to recursion. *Open Mind*, *4*, 57–70. https://doi.org/10.1162/opmi_a_00034 (cit. on pp. 6–8).

Xu, Y., Regier, T., & Malt, B. C. (2016). Historical semantic chaining and efficient communication: The case of container names. *Cognitive science*, *40*(8), 2081–2094 (cit. on p. 5).

Yang, D., & Regier, T. (2025). Re-examining the tradeoff between lexicon size and average morphosyntactic complexity in recursive numeral systems. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *47*. https://escholarship.org/uc/item/5k8646v4 (cit. on pp. 9, 18).

Yin, K., Regier, T., & Klein, D. (2024). Pressures for communicative efficiency in american sign language. *Annual Conference of the Association for Computational Linguistics (ACL)* (cit. on p. 6).

Zaslavsky, N., Garvin, K., Kemp, C., Tishby, N., & Regier, T. (2022). The evolution of color naming reflects pressure for efficiency: Evidence from the recent past. *Journal of Language Evolution*, *7*(2), 184–199. https://doi.org/10.1093/jole/lzac001 (cit. on p. 11).

Zaslavsky, N., Kemp, C., Regier, T., & Tishby, N. (2018). Efficient compression in color naming and its evolution. *Proceedings of the National Academy of Sciences*, *115*(31), 7937–7942 (cit. on p. 6).

Zaslavsky, N., Kemp, C., Tishby, N., & Regier, T. (2019). Color naming reflects both perceptual structure and communicative need. *Topics in Cognitive Science*, *11*(1), 207–219. https://doi.org/https://doi.org/10.1111/tops.12395 (cit. on p. 5).

Zaslavsky, N., Maldonado, M., & Culbertson, J. (2021). Let's talk (efficiently) about us: Person systems achieve near-optimal compression. *Proceedings of the annual meeting of the cognitive science society*, *43*(43) (cit. on p. 6).

Zipf, G. K. (1949). *Human behaviour and the principle of least effort*. Addison-Wesley. (Cit. on pp. 5, 7).