

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

Andrea Silvi¹ Ponrawee Prasertsom² Jennifer Culbertson²
Devdatt Dubhashi¹ Moa Johansson¹ Kenny Smith²

{silvi, moa.johansson, dubhashi}@chalmers.se

{ponrawee.prasertsom, jennifer.culbertson, kenny.smith}@ed.ac.uk

¹Chalmers University of Technology and Gothenburg University ²University of Edinburgh

Abstract

Human recursive numeral systems (i.e., counting systems such as English base-10 numerals), like many other grammatical systems, are highly regular. Following prior work that relates cross-linguistic tendencies to biases in learning, we ask whether regular systems are common because regularity facilitates learning. Adopting methods from the Reinforcement Learning literature, we confirm that highly regular human(-like) systems are easier to learn than unattested but possible irregular systems. This asymmetry emerges under the natural assumption that recursive numeral systems are designed for generalisation from limited data to represent all integers exactly. We also find that the influence of regularity on learnability is absent for unnatural, highly irregular systems, whose learnability is influenced instead by signal length, suggesting that different pressures may influence learnability differently in different parts of the space of possible numeral systems. Our results contribute to the body of work linking learnability to cross-linguistic prevalence.

Keywords: regularity; reinforcement learning; learnability; typology; recursive numerals; numerical cognition

Introduction

Linguistic systems are highly diverse, nevertheless they also exhibit similarities (Croft, 2002; Evans & Levinson, 2009). Research has shown that cross-linguistically prevalent features can be explained by a pressure for efficiency (Gibson et al., 2019; Kemp et al., 2018). In this paper, we focus on recursive numeral systems, i.e., systems of numerals that can express exact quantities (e.g., English *one, two, three, . . . , ten, eleven, twelve, . . . , twenty one, twenty two, twenty three, . . .*). Prasertsom et al. (2026) show that natural recursive numeral systems, as opposed to theoretically possible but unattested systems, are characterised by their efficiency; while such systems are maximally informative (i.e., any integer can be precisely conveyed), they exhibit a trade-off between processing complexity and regularity. Here, we focus on the latter, and explore its connection to learnability.

Previous research has connected learnability to cross-linguistic prevalence across a number of linguistic domains using artificial language learning experiments with humans, and computational models (e.g., Culbertson et al., 2012; Steinert-Threlkeld & Szymanik, 2020). We investigate whether regularity in recursive numeral systems is driven by a pressure for learnability, using Reinforcement Learning (RL) as a tool to measure learnability. We ask whether computational agents can learn human-like, regular recursive numeral systems bet-

ter than irregular ones. We find that regular recursive numeral systems are indeed more learnable than unnatural irregular ones in an RL setting. Importantly this result depends on the assumption that ability to generalise from limited training data to express a larger range of numbers is an important aspect of learnability. Interestingly, this holds in particular for systems that are highly regular, like human recursive numeral systems; regularity has less of an effect on the learnability of highly irregular systems.

We begin by briefly reviewing three relevant strands of research on language learning, cross-linguistic tendencies, and RL, which provide a basis for our study.

Efficiency of recursive numeral systems

Recent research seeks to explain why linguistic systems are the way they are by appealing to the fact that they are optimised for efficient communication (Gibson et al., 2019; Kemp et al., 2018), i.e., they minimise complexity while maximising communicative accuracy. A recent line of work has looked at numeral systems under this lens. Denić and Szymanik (2024) and Yang and Regier (2025) argue that recursive numeral systems optimise the trade-off between lexicon size (the number of unique morphemes that are used in the numeral system) and average morphosyntactic complexity (average number of morphemes per numeral). Intuitively, a system with a large lexicon consisting of many unique monomorphemic *numerals* for specific *numbers* can get by with shorter numerals (e.g. in a system with unique numerals for numbers in the range 1–20, all numbers in that range can be expressed with a single monomorphemic form of morphosyntactic complexity 1); a system with a smaller lexicon (e.g. dedicated morphemes only for numbers 1–10) needs to use longer numerals to express larger numbers (e.g. 11–20 might be expressed by numerals consisting of “ten plus one”, “ten plus two” etc).

While this trade-off has some explanatory power, Prasertsom et al. (2026) argue that it fails to take into account the regularity of numeral forms. Natural numeral systems tend to be highly regular (e.g. Mandarin in Table 1), but the measures used by Denić and Szymanik (2024) assign the same level of efficiency to highly irregular, unnatural systems (see example hypothetical system in Table 1), as long as they have the same lexicon size and average length.

Drawing on the Minimum Description Length approach from Brighton (2003), Prasertsom et al. (2026) propose that these systems are thus better thought of as being subject to

Number	Numeral (Mandarin)	Numeral (unnatural)
18	10 + 8	2 * 9
19	10 + 9	10 + 9
20	2 * 10	2 * 10
21	2 * 10 + 1	2 * 9 + 3
22	2 * 10 + 2	2 * 10 + 2
23	2 * 10 + 3	2 * 9 + 5
24	2 * 10 + 4	2 * 10 + 4
25	2 * 10 + 5	2 * 9 + 7

Table 1: Numerals in Mandarin (first column) and in a hypothetical system (second column) that is matched for lexicon size and average morphosyntactic complexity for the number range 18 – 25. Notice that Mandarin is consistently base-10, whereas the unnatural system has alternating bases (9, 10).

competing pressures to maximise *regularity* (how much form reuse there is) and minimise *processing complexity* (how difficult it is to parse or generate numerals). This approach is able to delineate natural recursive numeral systems from theoretically possible but unattested ones, with the natural ones exhibiting higher regularity and lower processing complexity.

Simplicity, learnability and cross-linguistic prevalence

Prasertsom et al. (2026) also suggest that *learning* may be the mechanism through which recursive numeral systems are optimised for regularity: regular systems have short description lengths, which make them easier to learn, more stable in intergenerational transmission, and therefore more widespread. There is a wealth of experimental evidence showing that simpler systems (i.e. those than can be described in a shorter or more compressed way in a given description language) are easier for humans to learn (e.g. Moreton et al., 2017; Raviv et al., 2021; Shepard et al., 1961; Wang et al., 2025; see Chater and Vitányi, 2003 for a broad review of simplicity principles in cognition). Experiments have also shown that learners introduce regularity into irregular linguistic systems (e.g., Keogh et al., 2024; Smith & Wonnacott, 2010).

Similar results also often arise in machine learning, where more human-like systems are more machine-learnable (Galke et al., 2024; Osmelak et al., 2026; Steinert-Threlkeld, Szymanik, et al., 2019; Steinert-Threlkeld & Szymanik, 2020). This suggests that machine learning models may suitably approximate human learning in this respect. We therefore hypothesise that human(-like), simpler, regular recursive numeral systems should also be easier for simulated agents (and thus plausibly also humans) to learn. Next, we turn to how learnability can be measured in this domain.

Measuring learnability via reinforcement learning

RL has been used as a tool to study the emergence of communication (Foerster et al., 2016; Lazaridou et al., 2017; Mordatch & Abbeel, 2018), often modelling the evolution of language as a cooperative process entailing the maximisation of future rewards. In this area, a growing body of work has shown

that the most communicatively efficient languages are also the ones that are easiest for deep RL agents to learn and converge on (Carlsson et al., 2021, 2024; Kågeback et al., 2020; Silvi et al., 2025). In other words, one could use RL to study the learnability of different theoretical types of linguistic systems.

Specifically, the learnability of a system can be quantified by how rapidly a single agent learns to minimise errors given a fixed number of learning exposures or budget of interactions with the environment. To measure this, theoretical work often adopts sample complexity (i.e., the number of samples required for convergence) to derive asymptotic bounds (Kakade, 2003). Similarly, empirical studies quantify efficiency by computing the Area Under the Curve (AUC) of learning statistics (Agarwal et al., 2021; Taylor & Stone, 2009). We adopt the latter approach to quantify the learnability of recursive numeral systems.

Experiment 1: Learnability of recursive numerals with varying regularity

First, we investigate the learnability of recursive numeral systems with different degrees of regularity. We expect that regularity and learnability under RL should correlate.

Methods

Across all experiments, we use single-agent RL (Sutton & Barto, 1998), where an agent tries to maximise cumulative reward over iterations (epochs). The goal is to learn to map each numeral (i.e., the form) to their number (i.e., the numerical quantity) in a given recursive numeral system. At each iteration, a batch of numerals are sampled with replacement from a training distribution. A single RL agent processes the numerals morpheme by morpheme and predicts numbers between 1–99,¹ and is rewarded depending on how close the predicted number is to the target number.

More formally, we characterise the task of learning recursive numerals as a contextual bandit problem with continuous action space (Agrawal, 1995; Lattimore & Szepesvári, 2020). Our agent is a single-layer LSTM network (Hochreiter, 1997). Each morpheme in a numeral is represented by a learned size-5 embedding, and the resulting sequence of embeddings is processed one-by-one by the LSTM (hidden size of 10). We use the last hidden state to predict a scalar mean $\mu \in [0.0, 1.0]$, which is interpreted as the mean of a Gaussian policy $\mathcal{N}(\mu, \sigma)$. We decrease the standard deviation σ as the training proceeds to encourage exploration early while encouraging exploitation later. The distribution is then used to sample a prediction, which is then multiplied by 100 to obtain a number between 1–99. Our reward function $\mathcal{R}$ is bounded, exponentially-decaying, Gaussian-like:

$$\mathcal{R}(\hat{n}, n) = \mathbb{1}[0 < \text{int}(\hat{n}) < 100] \exp(-\alpha|\hat{n} - n|) \quad (1)$$

where $\text{int}(\hat{n})$ approximates $\hat{n}$ to its nearest integer, $\alpha = 0.5$ determines the sharpness of the reward function. The first

¹For tractability and comparability with previous work, we only consider the numerals from 1 to 99.

part of the reward function (the indicator function) ensures predictions outside the range 1 – 99 are not rewarded. The second part penalises predictions that are far away from the target, with the reward rapidly approaching and becoming effectively 0 if the distance is 10 or more.

The RL agent is trained via the classical policy-gradient objective REINFORCE (Williams, 1992) over 30,000 epochs, with 5 batches of 32 numerals per epoch, sampled with replacement from the *training distribution* $p(n)$. The agent is tested on a sample of 99 numerals from the *test distribution* $q(n)$. The *learnability* of a given recursive numeral system is the AUC of accuracy at test against epochs, i.e. the approximation of the integral under the curve obtained by plotting epoch against test accuracy, where higher AUC corresponds to more rapid learning. We adopt this measure as it is more robust to training noise than other common ones such as final accuracy (Steinert-Threlkeld & Szymanik, 2020).

This setup effectively frames the problem of learning numeral systems as a regression problem, rather than a classification one used in prior literature (Carlsson et al., 2021; Guo et al., 2019; Silvi et al., 2025). As previous classification approaches treat numbers as discrete orthogonal labels, the agent there would need to learn that some numbers are closer to each other than others via the reward function. Under our regression approach, the agent does not have to learn this, as the output space already has the required ordering in place (i.e., it is already a number line). For example, a classification task would require the agent to learn that the class for 15 is close to the class for 16 via a distance-based reward function, whereas in our regression task, there is no need for the agent to learn that 15 ($\mu = 0.15$) is close to 16 ($\mu = 0.16$), so the agent simply has to learn the form-meaning mappings.

To investigate the effect of regularity on learnability, we follow the same formalisation of *irregularity* in Prasertsom et al. (2026). The measure rests on the intuition that regular systems can be represented with smaller sets of grammar rules, which can be represented by a finite-state machine that generates all and only valid numeral strings. Figure 1 gives an example of an inferred minimal partial DFA that generates all and only valid Mandarin numerals, a highly regular system. The *irregularity* is the size of the numeral grammar, measured by the number of bits required to encode the DFA:

$$\underbrace{|Z|(2 \cdot \log_2|S| + \log_2|\Sigma|)}_{(A)} + \underbrace{\log_2|S|}_{(B)} + \underbrace{|S|}_{(C)} \quad (2)$$

where (A) is the number of bits required to encode the transition table, where Z is the set of transitions, S the set of states and Σ the set of symbols (morphemes). (B) and (C) are the numbers of bits for encoding the identity of initial state and accepting states. We use the same algorithm (Mihov & Schulz, 2019) and implementation (`automata-lib` in Python) for minimal DFA inference as Prasertsom et al.

Experiment setup

We quantify the relationship between learnability and regularity in the sample of recursive numeral systems investigated




Figure 1: Inferred minimal partial DFA representing the Mandarin recursive numeral grammar. Transitions (arrows) represent output morphemes. States (circles) are labelled roughly according to the relevant grammar rules they capture. Double borders indicate accepting states. Generating the numeral for 20 ($2 * 10$) would involve going from the initial state (λ) to D (emitting 2), from D to D^* (emitting $*$), and from D^* to $D * 10$ (emitting 10) and terminate the machine there.

by Prasertsom et al. (2026). This includes 1) 128 human recursive numeral systems from Denić and Szymanik (2024), 2) a sample of artificially generated systems under the lexicon size/average morphosyntactic complexity trade-off from Denić and Szymanik (data from 2024, henceforth, *D&S optimal systems*), 3) a sample of optimal artificial systems under the same trade-off with post-hoc constraints imposed to enforce human-likeness from Yang and Regier (2025) (henceforth, *Y&R optimal systems*), and 4) a sample of 300 random baseline unattested systems that make use of the symbols (digits and multipliers) found in human systems. Note that all the artificial languages are either sampled or selected to have a lexicon size which is comparable to the ones attested in human systems (i.e., lexicon size of 5–13). To ensure robustness, we compute the learnability as described above 20 times for each system, taking the mean.

In this initial setup, our training and test distributions are both the power law distribution $p(n) \propto n^{-2}$, $q(n) \propto n^{-2}$, following previous work (Carlsson et al., 2021; Denić & Szymanik, 2024; Silvi et al., 2025; Xu et al., 2020), which drew on the estimated communicative need of each number n (Dehaene & Mehler, 1992). Under this highly-skewed distribution, about 75% of training and test samples will consist of numbers 1 and 2 alone, and about 90% will consist of numbers in the range 1-6.

Results and discussion

Figure 2 (left) plots irregularity against learnability, along with a least-square regression fit to this data. Surprisingly, we found that there is little correlation (as indicated by the almost horizontal regression line of slope $\beta = -0.020$): there is no difference between regular, human(-like) systems and irregular systems in terms of learnability.

What could explain this result? We assumed that the pri-




Figure 2: Irregularity and learnability for human (blue), D&S optimal (red), Y&R optimal (green), and random baseline (orange) recursive numeral systems under power law (left panel) and the uniform (right panel) test distributions $q(n)$. The training distribution $p(n)$ is always the power law distribution.

mary distribution that shapes recursive numeral systems is the power law distribution derived from frequency of use. This means that, at test, the agent trained on the system only has to perform well on the lower ranges of numerals in order for that system to score high on learnability. However, as Prasertsom et al. (2026) point out, it is likely that recursive numeral systems are sensitive to the pressure to precisely communicate about *any* integer, including very large numbers (cf. also Frank et al., 2008; Pitt et al., 2022). This is at odds with using the power law distribution to evaluate learnability, under which larger numbers are virtually never sampled and which therefore contribute negligibly to measured learnability.

Experiment 2: Learnability as accuracy of generalisation

We therefore separate out two relevant distributions that influence our measure of the learnability of recursive numeral systems. One is the *learning distribution*, which reflects the statistics of the input to our learners. The other is the *generalising distribution*, which we argue should reflect the need to precisely communicate all numbers.

In Experiment 2, we assume that the training distribution $p(n)$ is based on the frequency of use, and thus we use the power law distribution suggested in prior work. By contrast, $q(n)$ will be flat (an extreme case of a less-skewed distribution), reflecting the idea that numerals in the long tail also require high precision. We thus explore the relationship between regularity and learnability in terms of the capacity to generalise to uniform $q(n)$.

Experiment setup

We adopt the same methods as the previous experiment, but use a uniform distribution over $1 - 99$, i.e., $p(n) \propto n^{-2}$ and $q(n) \propto \frac{1}{99}$.

Results and discussion

Figure 2 (right) plots the learnability of the systems under a uniform $q(n)$. As expected, highly regular human systems now have on average the highest learnability and there is a negative correlation between irregularity and learnability ($\beta = -0.828$). This is likely because the agents can learn subparts of numerals and then reuse them to more efficiently learn the rarer, larger numbers. In contrast, the sampled random systems are harder to learn, likely because they tend to jump between bases, reusing the same sub-forms less frequently and less systematically.

Figure 3 (left) provides a closer look at human and highly regular languages. D&S systems, despite being efficient in terms of lexicon size and length, have relatively high irregularity and hence low learnability, whereas Y&R systems were constrained to be human-like and regular, and therefore have higher learnability.

There are however a few clear outliers, such as the highly learnable Y&R system (the green triangle at the top of Figure 3, left), which is more irregular than most human systems yet more learnable. This could be because the current measure of regularity does not take into account the structure of the output space, but only the degree of overall regularity (i.e., how much form is reused across *all* numerals). However, it is plausible that learnability is largely affected by *local* regularity (i.e., how much form is reused between consecutive numbers). French, for instance, uses inconsistent bases, but uses a consistent base for a given range of consecutive numbers: base 10 in the 30s up to the 60s (e.g., *quar-ante cinq* $4 * 10 + 5$), base 20 in the 80s (e.g., *quatre-vingt-un* $4 * 20 + 1$), and both in the 70s and 90s (e.g., *quatre-vingt-dix-sept* $4 * 20 + 10 + 7$), and so on.

We therefore compute *average local regularity*, adapted from the previous regularity measure. We first extract recursive numeral subsystems that correspond to 10 consecutive numerals from a given system (i.e., numerals in the ranges 1-10, 2-11, 3-12, ...). We compute the irregularity for each of these (local) subsystems, then take the unweighted aver-




Figure 3: Irregularity (left) and average local irregularity (right) against learnability for human (blue), D&S optimal (red) and Y&R optimal (green) systems, under the power law training distribution and a uniform test distribution.

age to obtain the average local irregularity. Figure 3 (right) shows that once local regularity is taken into account, the outliers appear more in line with the trend, suggesting that local regularity may indeed play key role in the learnability of a recursive numeral system for our RL agents.

While it seems that regularity correlates well with learnability overall (Figure 2, right, $\beta = -0.828$) and for highly regular systems (Figure 3, left, $\beta = -2.110$), random artificial systems do not follow the same trend (see the orange cluster in Figure 2, right), where the negative trend is weaker ($\beta = -0.259$ among this cluster). For these highly irregular systems, we instead see an influence of average morphosyntactic complexity (i.e., average length): higher morphosyntactic complexity corresponds to lower learnability. This trend is almost absent for highly regular systems (Figure 4). This result suggests that in different areas of the space of possible recursive numeral systems, different measures can have a higher or lower impact on learnability. Specifically, the degree of regularity (i.e., the complexity of the rules involved) plays a large role when the system exhibits some regularity, whereas for highly irregular systems, it is so difficult to find any pattern that changes in regularity matter much less (cf. Raviv et al., 2021). In this case, length becomes the main determinant for learnability since longer forms are harder to memorise.

In the next section, we test this idea more explicitly, and examine the regularity-learnability link in families of highly similar, human-like, regular recursive numeral systems with fixed average morphosyntactic complexity and lexicon size.

Experiment 3: Learnability among highly regular, human-like systems

Experiment setup

We adopt the same methods as the previous two experiments, but our samples now come from the most and least efficient systems in the *local neighbourhoods* of human recursive numeral systems. These are families of artificially generated systems in Prasertsom et al. (2026) that have the same digits, multipliers, arithmetic combinators (+, -, *) and numeral lengths as human recursive systems. For example, Basque uses 10




Figure 4: Average morphosyntactic complexity and learnability for random baseline (left) and highly regular systems (right), under the power law training distribution and a uniform test distribution.

and 20 as multipliers, and 91 is expressed as $4 * 20 + 10 + 1$ (length = 7). Thus, its alternative numerals of the same length are $\{10 + 4 * 20 + 1, 10 + 8 * 10 + 1, 20 + 7 * 10 + 1, 4 * 20 + 10 + 1, 7 * 10 + 20 + 1, 8 * 10 + 10 + 1\}$. Basque’s local neighbourhoods consist of systems comprised of combinations of these alternative numerals. This effectively also means that systems in each local neighbourhood will all have the same lexicon size and average morphosyntactic complexity. The samples only consist of the most and least regular systems obtained from a greedy algorithm. Since some neighbourhoods feature very small variation, we only include 30 out of 37 neighbourhoods that have more than 10 systems.

Results and discussion

As predicted, we again found a negative correlation between irregularity and learnability in most of these neighbourhoods. We fit a least-square regression predicting learnability from irregularity for each local neighbourhood, and found a negative trend in 21 out of 30 neighbourhoods (Figure 5, left; Mean $\beta_{\text{irregularity}} = -0.249$; S.E. = 0.451). Figure 5 (right) gives illustrative examples from four neighbourhoods (see the full plot that includes every neighbourhood see [our anonymised GitHub repository](#)).




Figure 5: Left: best-fit regression lines between irregularity and learnability for the higher-populated local neighbourhoods (one colour per neighbourhood) from Prasertsom et al. (2026). Right: Irregularity and learnability of example local neighbourhoods (Arawak-like, Diola-like, Nahuatl-like and Yakut-like systems).

General discussion

In this work, we explored whether learnability pressures can explain the cross-linguistic prevalence of regularity in recursive numeral systems. In particular, using RL, we measured the learnability of attested and hypothetical recursive numeral systems with varying degrees of regularity, and tested whether regularity correlates with learnability. We found that this is born out under the assumption that the relevant measure of learnability reflects the need to generalise from limited data to precisely express all integers. By contrast, it does not hold under the assumption that both the training and testing distributions are the power law communicative need distribution (Dehaene & Mehler, 1992). We additionally demonstrated that regularity has less influence on the learnability of highly irregular systems. Our overall results thus corroborate the idea that learnability explains the cross-linguistic prevalence of regular recursive numeral systems. They also point to the particular importance of generalisation in the learning of numeral systems, even if frequency of use is highly skewed.

We have shown the learnability advantage of regularity with RL. A natural prediction is that regularity should similarly confer a learnability advantage for humans in this domain. However, in two artificial numeral learning experiments, Holt and Barner (2025) found little evidence of English-speaking adults’ learning performance with different amounts of exposure to generalisable, regular rules. While this may seem at odds with our prediction, it is worth noting that the numerals used in this study were in the lower ranges especially during training (1–10 at train, 11–20 at test), and there was only one round of learning and generalisation. In other words, unlike humans in real-life situations who have a real need to generalise to large numerals, the participants were not presented with any direct pressure to perform well at generalisation. Future designs might implement this pressures artificially, e.g., by iteratively alternating between learning and generalisation.

Another issue worth exploring concerns the role of input

orders. Our model treats the input as random samples from the distribution of use. However, recursive numerals are typically learned through counting routines (Carey, 2000; Carey & Barner, 2019; Fuson, 2012), perhaps even necessarily so (Pitt et al., 2022). One way counting routines could facilitate learning is by artificially boosting the frequency of numerals over 10 (e.g. in the range 10–30), which are not frequently used for communication according to the power law, but are often part of counting routines. This would expose learners to a range of numerals from which highly generalisable rules can be extracted. For instance, in Mandarin, numerals in the range of 1–10 will not allow the learner to extract any general pattern because they all have dedicated forms, whereas the range 1–30 contains all the patterns required for all numerals up to 99 (e.g., you can reuse the rules generating $3 * 10 + 1$ for any numeral of the same structure like $4 * 10 + 5$).² It would be possible to investigate this with the current model, or a Markov Decision Process (Sutton & Barto, 1998) where an agent is exposed to consecutive numbers in a single episode.

Our results also suggest that the relationship between regularity and learnability is likely non-linear, and other aspects of complexity (like length) might influence learnability more in some parts of the space. Going beyond regularity, an interesting future direction would be to identify the influence on learnability not only of morphosyntactic complexity/length, but also other pressures such as processing complexity, hypothesised by Prasertsom et al. (2026) to also trade off with regularity among highly efficient numeral systems. The fact that learnability is multiply determined also means that a learnability-based measure of complexity (e.g., Osmelak et al., 2026) may be of limited use as a substitute for a more transparent measure of a specific aspect of complexity (like our regularity measures), as learnability may depend on different features of systems in different regions of the space.

Finally and more broadly, our study may have implications for the evolution of regularity in general. We find that regularity is most useful when the learning metric requires generalisation beyond the learning distribution (i.e., when the train and test distributions are different). This shows that regular systems in general enjoy the benefits of being robust to different need distributions. We suggest that, besides the learning vs. generalising distribution distinction, this distinction could also be framed in terms of *changing* communicative needs. In addition to being robust to limited learning samples (cf. the *transmission bottleneck*; Smith et al., 2003), then, regular systems may also be selected in language evolution because they are a robust solution to changes in communicative needs over individual and historical times (Karjus et al., 2020).

²This is also reflected in the local regularity of Mandarin numerals in these ranges. It increases as more numbers are covered but no longer increase after 30 (56.219 for 1-10, 201.192 for 1-20, 241.039 for 1-30, 1-40, 1-50 up to 1-99).

Acknowledgments

P.P. is supported by Anandamahidol Foundation, Thailand. A.S. is supported by the Swedish Research Council. The computations were enabled by resources provided by Chalmers e-Commons at Chalmers.

References

- Agarwal, R., Schwarzer, M., Castro, P. S., Courville, A., & Bellemare, M. G. (2021). Deep reinforcement learning at the edge of the statistical precipice. *Proceedings of the 35th International Conference on Neural Information Processing Systems*.
- Agrawal, R. (1995). The continuum-armed bandit problem. *SIAM J. Control Optim.*, 33(6), 1926–1951. <https://doi.org/10.1137/S0363012992237273>
- Brighton, H. (2003). *Simplicity as a driving force in linguistic evolution* [PhD thesis]. The University of Edinburgh. <http://hdl.handle.net/1842/23810>
- Carey, S. (2000). The origin of concepts. *Journal of Cognition and Development*, 1(1), 37–41.
- Carey, S., & Barner, D. (2019). Ontogenetic origins of human integer representations. *Trends in cognitive sciences*, 23(10), 823–835.
- 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.
- 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>
- Chater, N., & Vitányi, P. (2003). Simplicity: A unifying principle in cognitive science? *Trends in Cognitive Sciences*, 7(1), 19–22. [https://doi.org/https://doi.org/10.1016/S1364-6613\(02\)00005-0](https://doi.org/https://doi.org/10.1016/S1364-6613(02)00005-0)
- Croft, W. (2002). *Typology and universals*. Cambridge university press.
- Culbertson, J., Smolensky, P., & Legendre, G. (2012). Learning biases predict a word order universal. *Cognition*, 122(3), 306–329.
- 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](https://doi.org/10.1016/0010-0277(92)90030-L)
- 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>
- 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.
- 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.
- Frank, M. C., Everett, D. L., Fedorenko, E., & Gibson, E. (2008). Number as a cognitive technology: Evidence from pirahã language and cognition. *Cognition*, 108(3), 819–824. <https://doi.org/https://doi.org/10.1016/j.cognition.2008.04.007>
- Fuson, K. C. (2012). *Children's counting and concepts of number*. Springer Science & Business Media.
- 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>
- 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>
- Guo, S., Ren, Y., Havrylov, S., Frank, S., Titov, I., & Smith, K. (2019). The emergence of compositional languages for numeric concepts through iterated learning in neural agents. *CoRR, abs/1910.05291*. <http://arxiv.org/abs/1910.05291>
- Hochreiter, S. (1997). Long short-term memory. *Neural Computation MIT-Press*.
- 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>
- Kågeback, M., Carlsson, E., Dubhashi, D., & Sayeed, A. (2020). A reinforcement-learning approach to efficient communication. *PLOS ONE*, 15(7), 1–26. <https://doi.org/10.1371/journal.pone.0234894>
- Kakade, S. M. (2003). *On the sample complexity of reinforcement learning* [Doctoral dissertation]. University College London.
- Karjus, A., Blythe, R. A., Kirby, S., & Smith, K. (2020). Quantifying the dynamics of topical fluctuations in language. *Language Dynamics and Change*, 10(1), 86–125.
- Kemp, C., Xu, Y., & Regier, T. (2018). Semantic typology and efficient communication. *Annual Review of Linguistics*, 4(1), 109–128. <https://doi.org/10.1146/annurev-linguistics-011817-045406>
- Keogh, A., Kirby, S., & Culbertson, J. (2024). Predictability and variation in language are differentially affected by learning and production. *Cognitive Science*, 48(4), e13435.
- Lattimore, T., & Szepesvári, C. (2020). *Bandit algorithms*. Cambridge University Press.
- Lazaridou, A., Peysakhovich, A., & Baroni, M. (2017). Multi-agent cooperation and the emergence of (natural) language. *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. <https://openreview.net/forum?id=Hk8N3ScIsg>
- Mihov, S., & Schulz, K. U. (2019, August 1). *Finite-State Techniques: Automata, Transducers and Bimachines*

- (1st ed.). Cambridge University Press. <https://doi.org/10.1017/9781108756945>
- Mordatch, I., & Abbeel, P. (2018). Emergence of grounded compositional language in multi-agent populations. *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence*.
- Moreton, E., Pater, J., & Pertsova, K. (2017). Phonological concept learning. *Cognitive science*, 41(1), 4–69.
- 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>
- Pitt, B., Gibson, E., & Piantadosi, S. T. (2022). Exact number concepts are limited to the verbal count range. *Psychological Science*, 33(3), 371–381.
- Prasertsom, P., Silvi, A., Culbertson, J., Johansson, M., Dubhashi, D., & Smith, K. (2026). Recursive numeral systems are highly regular and easy to process [arXiv:2501.01234. To appear in Proceedings of the EACL]. <https://arxiv.org/abs/2510.27049>
- 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>
- Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological monographs: General and applied*, 75(13), 1.
- 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>
- Smith, K., Kirby, S., & Brighton, H. (2003). Iterated learning: A framework for the emergence of language. *Artificial life*, 9(4), 371–386.
- Smith, K., & Wonnacott, E. (2010). Eliminating unpredictable variation through iterated learning. *Cognition*, 116(3), 444–449.
- 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>
- Steinert-Threlkeld, S., & Szymanik, J. (2020). Ease of learning explains semantic universals. *Cognition*, 195, 104076. <https://doi.org/10.1016/j.cognition.2019.104076>
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction* (Second). The MIT Press.
- Taylor, M. E., & Stone, P. (2009). Transfer learning for reinforcement learning domains: A survey. *J. Mach. Learn. Res.*, 10, 1633–1685.
- 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.
- 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>
- 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
- 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>