

# Visual Theory of Mind Enables the Invention of Writing Systems

Benjamin A. Spiegel (bspiegel@cs.brown.edu); Lucas Gelfond; George Konidaris

Department of Computer Science, Brown University

## Abstract

Abstract symbolic writing systems are *semiotic codes* that are ubiquitous in modern society but are otherwise absent in the animal kingdom. Anthropological evidence suggests that the earliest forms of some writing systems originally consisted of *iconic pictographs*, which signify their referent via visual resemblance. While previous studies have examined the emergence and, separately, the evolution of pictographic writing systems through a computational lens, most employ non-naturalistic methodologies that make it difficult to draw clear analogies to human and animal cognition. We develop a multi-agent reinforcement learning testbed for emergent communication called a *Signification Game*, and formulate a model of inferential communication that enables agents to leverage *visual theory of mind* to communicate actions using pictographs. Our model, which is situated within a broader formalism for animal communication, sheds light on the cognitive and cultural processes that led to the development of early writing systems.

**Keywords:** cognitive semiotics; reinforcement learning; language learning; sketch understanding; bayesian modeling

## Introduction

The ability to spontaneously invent, use, and understand natural language is unique to humans and separates us from the rest of the animal kingdom. However, the capacity for *signification*, i.e. communicating a mental referent using a *sign*, is shared with animals to some extent. For example, sophisticated acoustic signaling has been observed in many species, including Sperm Whales (Sharma et al., 2024), Chickadees (Ficken et al., 1994) and the Japanese great tit bird (Toshitaka N. Suzuki, 2016); and gesture-based visual signaling has been observed in chimpanzees, bonobos, and bees (Gillespie-Lynch et al., 2014; Halina et al., 2013; Von Frisch, 1962).

One variety of signification systems not found elsewhere in nature are *symbolic writing systems*, ubiquitous tools in modern society that practically every human can use. While diachronic linguistics has long studied the evolution of spoken language, the early origins of written language are primarily a matter of prehistoric record. Substantial archaeological evidence indicates that the earliest writing systems were comprised of *iconic pictographs*, which signify their referents via a visual resemblance relationship (Taylor, 1883; Trigger, 1998). These stand in contrast to the *abstract symbolic* signs used in modern writing systems, which instead rely on codified convention to convey meaning. The evolution of writing from *icons* to *symbols* raises fundamental questions about the cognitive mechanisms that enabled the invention of the first

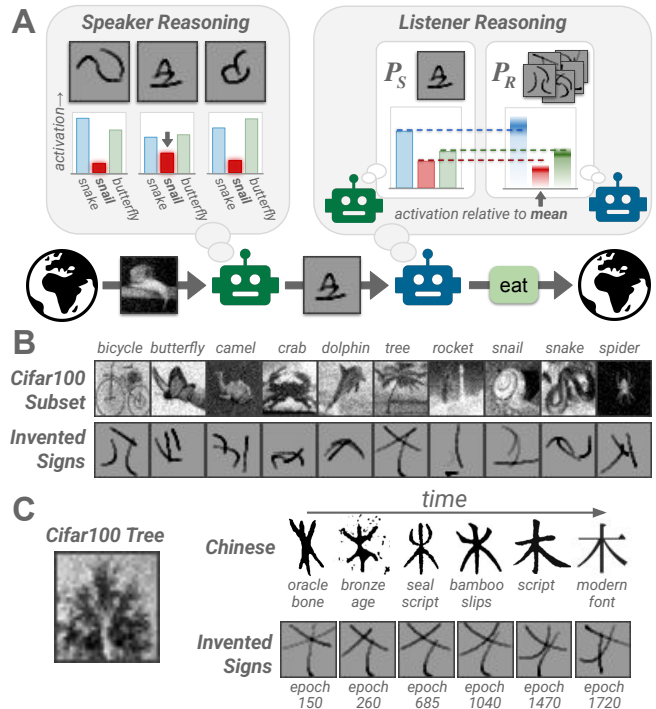


Figure 1: **A)** Our model of inferential communication based on *visual theory of mind* has agents consider background recognition likelihoods for computing a referent sensitivity term  $P_R(r_i)$  for each referent, sensitizing them to referents that are hard to generate signs for. **B)** Signs invented by agents to mimic classes from the Cifar100 dataset. Agents create pictographs given discrete states that serve as class labels. **C)** We simulate communication over long time horizons, finding that initially iconic signs grow more abstract over time. We compare the evolution of a sign for Cifar100 palm trees for a Chinese character for “tree.”

pictographs and the cultural processes driving their transition to abstract symbolic signs.

While studies using human subjects have shed light on the short-term evolution of ad-hoc writing systems (Fan et al., 2023; Fan et al., 2018; Fay et al., 2014; Garrod et al., 2007; Hawkins et al., 2023), current computational simulations treat the invention of pictographs and the evolution of writing systems as separate phenomena and rely on technical assump-

tions that lack faithfulness to a naturalistic setting for communication. For example, Mihai and Hare (2021) present a two-agent model for sketching that is end-to-end differentiable, enabling gradients to pass between agents that represent separate entities; Fan et al. (2019) present a sketching model that is not generative, and instead communicates by selecting a sketch from a dataset of human-drawn sketches; Qiu et al. (2022) pre-train a sender model on an existing sketch dataset and use a hand-designed architecture that extracts salient edges for easier sketching; and Jiang et al. (2024) model the evolution of Chinese using a domain-specific-language with primitives that are crafted specifically for Chinese, and also do not consider how the first Chinese characters emerged. These methodological limitations often make it difficult to draw clear analogies to human cognition, and fall short of situating *pictographic signification* within a broader formalism for animal communication, which could shed light on why even our closest relatives in the clade of Great Apes did not develop writing systems and cannot understand iconic pictographs as well as humans can (Persson, 2008).

We present a unified computational approach to simulating *both* the invention of iconic pictographs and their transition to abstract symbols over time. In Part I of this paper we introduce the *Signification Game*, a naturalistic communication testbed inspired by referential and signaling games (Lewis, 1969; Skyrms, 2010) in which agents engaged in a decision process communicate by painting splines on a canvas, an action space designed to mimic the limitations of early human writing tools. In contrast to referential games, agents lack specialized hardware for processing communicative signals, and must instead learn to interpret signals using the same perception system used for processing state observations. We show that while agents with simple policy architectures can learn to communicate from reward-maximization alone, such a means of language learning is severely limited. Namely, the constraints imposed by a writing-like signal induces a *signification gap*, making some referents impossible to communicate without substantial learning by listeners. In Part II, we formalize an inferential model of communication that leverages *visual theory of mind*, and show that agents can wield it to surmount the signification gap, enabling them to signify most referents on much shorter time scales.

## Part I: A Model of Primitive Communication

We describe a setting for simulating the emergence of visual communication called a *Signification Game*. We start by motivating our experimental design by drawing from theories on the emergence of primitive communication in nature.

### The Origins of Communication in Nature

One plausible origin story for early animal communication is that primitive communicative acts constituted the use of signals to induce behavior in others that conferred an advantage to speakers. For example, an animal that has associated a screeching noise with the presence of a predator may have learned to reliably take a fleeing action, a behavior that could




Figure 2: Signaling can be used for manipulative or cooperative ends. Actions yielding positive and negative rewards are denoted with green and red bubbles, respectively.

be exploited by another animal who screeches so that it may selfishly engorge itself on a scarce food source without competition. Krebs and Dawkins (1978) cite many similar examples in which animals learn to “exploit other animals as tools” because they respond in “robot-like fashion to key stimuli.” Furthermore, Krebs and Dawkins note that the payoff structure resulting from the induced behavior leads to substantially different kinds of signals, with mutually-advantageous cooperative communication leading to cost-minimizing signals and non-cooperative signaling resulting in conspicuous, “typical ritualized” signals. We explore both of these settings.

Formally, an animal who forms a stimulus-response (S-R) behavior is conditioned to associate a stimulus  $S$  (e.g., a screeching sound) with a response  $R$  (e.g., a fleeing action) using a classifier  $f_{S-R} : X \rightarrow [0, 1]$ , where  $X$  is the space of possible stimuli and  $f_{S-R}(S)$  outputs the probability that  $S$  is recognized as the stimulus and triggers the associated response behavior (Skinner, 1938). Such an S-R behavior can be exploited by another animal that can generate  $\hat{S}$ , a convincing replica of stimulus  $S$  that fools  $f_{S-R}$  into inducing response  $R$ . Under this formulation, a primitive language can emerge comprised of utterances  $\{\hat{S}_1, \hat{S}_2, \dots\}$  with meanings co-opted from precursor S-R behaviors  $\{R_1, R_2, \dots\}$ .<sup>1</sup> The communication-as-manipulation hypothesis is attractive because it suggests that communication can emerge without any special-purpose cognition—as soon as biological hardware that can be used to generate signals emerges, the same cognitive tools for maximizing reward in the environment can be applied to such hardware to learn primitive languages.

## Methods

**Signification Games** A Signification Game is an iterated multi-agent reinforcement learning (Littman, 1994) task for simulating the emergence of communication under naturalistic assumptions. In its basic form, a population of agents begins by learning S-R behaviors—implemented as policy functions from observations to actions,  $\pi : o \rightarrow a$ —that maximize reward in an underlying decision process. Over time, agents are gradually given the opportunity to interact, allowing them to communicate using a continuous signal—in our case, greyscale images—that is interpreted by the same be-

<sup>1</sup>A similar action-oriented account of language is described in Wittgenstein (1953).

havioral policy used for interacting with the environment. In its most basic formulation, agents cannot distinguish between their observations of the world and the signals generated by other agents; agents must learn to exploit the behavioral policies of others to induce desired behaviors, which may be self-ish or mutually-advantageous. Signification Games can be played over millions of rounds to give us a picture of how communication protocols emerge and evolve over time.

Formally, a Signification Game is a decentralized partially-observable Markov Decision Process (Oliehoek & Amato, 2016) denoted as a tuple  $\langle S, \{A_i\}_{i=1}^N, \{\Omega_i\}_{i=1}^N, T, R \rangle$ , where  $S$  is a set of states,  $A_i$  is the action space for agent  $i$ ,  $\Omega_i(o_i|s)$  is the observation function for agent  $i$ ,  $T(s'|s, \{a_i\})$  represents the transition function given the current state  $s$  and the joint actions  $\{a_i\}$ , and  $R(s, \{a_i\})$  is the reward function. From each state  $s \in S$ , a collection of contextual bandits<sup>2</sup>  $\{c_1, \dots, c_n\}$  (Langford & Zhang, 2007) with winning arms  $\{a_1, \dots, a_n\}$  is created for a subset of agents designated as *listener* agents. The contexts  $\{x_1, \dots, x_n\}$  for each bandit are greyscale images drawn from two sources: 1) a multi-class dataset whose classes correspond to the winning arms, or 2) from *speaker* agents who observe the winning arms before generating a signal. This gives agents the opportunity to induce actions in others based on knowledge of the underlying decision process.

**Agents** Agents in a Signification Game have a high-level policy  $\pi$  comprised of two sub-policies: a behavioral policy  $\pi_L$  invoked when agents play a *listener* or *observer* role and a speaker policy  $\pi_S$  invoked when agents play a *speaker* role. The behavioral policy acts directly on the environment, taking images as input—either observed from the environment or generated by other agents—and outputting discrete actions to the underlying decision process. The speaker policy generates an image conditioned on the discrete state of the underlying decision process. Formally, each agent  $i$  seeks to maximize cumulative returns using its high-level policy  $\pi_i$  over the duration of the signification game:

$$\mathbb{E} \left[ \sum_{t=0}^{\infty} \gamma^t R(s_t, \pi_i(\Omega_i(s_t))) \right],$$

where sub-policies are conditionally deployed based on the observation function  $\Omega_i : s_t \rightarrow o_i$ , which returns an image if the agent is tasked with *observing* or *listening*, or a one-hot state vector if the agent is tasked with *speaking*. Behavioral policies are implemented as Convolutional Neural Networks (Krizhevsky et al., 2012) with an attached neural network, whose logits are akin to the activations for a set of S-R classifiers  $\{f_a|a \in A\}$  from which actions are sampled according to a categorical distribution obtained by applying a softmax:  $P(a|o) = \text{softmax}(f_a(o))$ . Speaker policies are implemented with neural networks that output a multivariate Gaus-

sian distribution over spline parameters, which are sampled and painted on a canvas to create an image. Speaker policies are akin to a set of generators  $\{g_a|a \in A\}$  for producing stimuli that induce actions in listeners. Policies are trained independently with PPO (Schulman et al., 2017) using separate rollouts  $\{(s_t, \pi_S(o_t^i), r_t), \dots\}$  and  $\{(s_t, \pi_L(o_t^i), r_t), \dots\}$ .

## Experiments

Signification Games afford a number of tunable parameters that can drastically affect experiment outcomes, including the payoff structure for communication (e.g., cooperative or manipulative), the population size, the signal-rendering system for generating images, the frequency of agent-agent interaction over time, and the multi-class image dataset used in the underlying decision process. Unless otherwise stated, experiments use the following parameters.

- Each experiment begins with a solipsistic phase in which agents interact solely with the underlying decision process until performance plateaus, at which point agents are gradually exposed to signals generated by other agents. We imagine a setting in which the hardware for communication emerges in a population of agents who are already proficient at maximizing reward in their environment.
- Observations for the underlying decision process are drawn from the MNIST dataset (LeCun et al., 1998). When observing images from the environment, agents are rewarded for taking a discrete action that corresponds to the true class of the observed image. While agents are essentially performing a classification task, this experimental setup is roughly analogous to animals learning to take advantageous actions in their environment, e.g. fleeing or pursuing, based on observed stimuli, e.g. from predators, manipulators, or allies. We use the MNIST dataset due to its easily-recognizable classes, though a more natural dataset is used in Part II.
- Speakers output parameters for two Bézier curves that are rendered on a canvas and peppered with Gaussian noise to simulate an imperfect communication medium. Speakers also receive rewards inversely proportional to the entropy of listener policies if successful communication occurs, reflecting agents with higher confidence in their observations being more likely to act on them. Effectively, this encourages speakers to push away from the decision boundaries of listeners to generate more prototypical stimuli.
- Each experiment is run with a population of 5 agents. Experiments with up to 20 agents yield similar results, but using a small number of agents enables longer simulations. In general, increasing the number of agents acts as a *regularizer*, as agents must convince a larger population with their signals.

We additionally train an *iconicity probe*: a CNN-based classifier trained on MNIST images to assess the iconicity of agent signals. The probe returns a distribution over MNIST classes for a given image. High entropy of this distribution indicates that the probe struggles to recognize the class of the image. We find that *iconic* images yield lower entropy

<sup>2</sup>Using a contextual bandit simplifies credit assignment, though all methods explored here can be extended to temporally-extended decision processes.

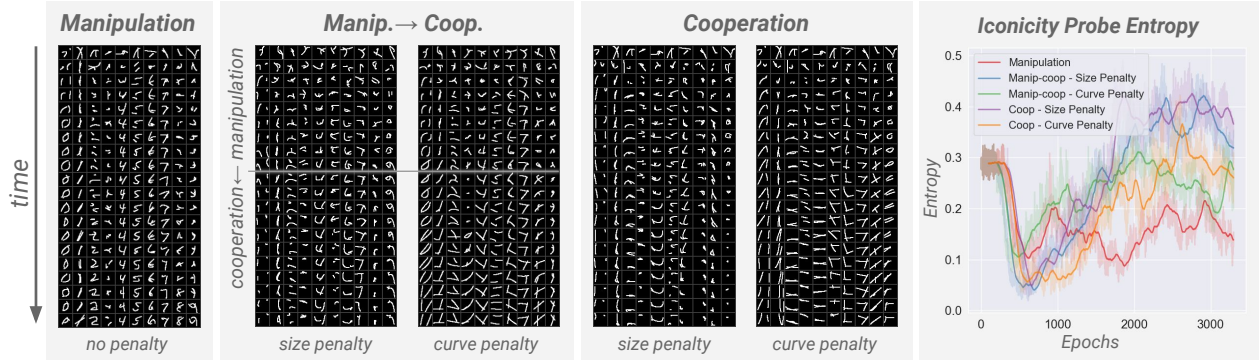


Figure 3: **Visualization of signaling over time.** Agent signals appear random at first (top rows) and grow increasingly ordered over time. In a manipulation setting (left), speakers learn to generate convincing replicas of state observations to trick increasingly picky listeners. In a manipulation-cooperation setting (second from left), agents switch to a cooperative setting after 300 epochs of manipulation, resulting in a drift from iconic signals into abstract symbols. In a fully-cooperative setting (second from right) agents generate vaguely iconic signals that quickly shift into symbols. Environmental pressures on signal production and interpretation—simulated by size and curve penalties—greatly affect the appearance of signals over time. Qualitative analysis of the *iconic* and *symbolic* nature of generated signals is supported by an *iconicity probe* (right). We find that agent-generated signals achieve low probe entropy (i.e. higher confidence of recognition) shortly after communication begins before increasing, reflecting a transition from icons to abstract symbols.

scores, reflecting higher confidence that the image belongs to a specific class, while more *abstract symbolic* images yield higher entropy scores, reflecting more class ambiguity.

**Deceptive Signaling Leads to Iconic Replicas.** We first assess the emergence of communication in an adversarial context. In a natural setting with a scarcity of resources, e.g., actions induced by stimuli from self-interested animals are more likely to benefit the speaker than the listener. In this experiment, speakers are rewarded for successfully tricking listeners into taking an action, while listeners are penalized for reacting. As agents balance their capacity to classify environment observations while interpreting deceptive signals, an arms race unfolds in which agents generate increasingly convincing stimuli.<sup>3</sup> As a result, agents learn to generate a form of iconic signs called *replicas* (see Figure 3, left), which confer meaning by virtue of being perfect facsimiles of existing meaningful stimuli (Eco, 1976). Crucially, iconic signification using replicas is possible only when the signaling space of speaker policies directly overlaps with relevant environment stimuli. In Part II we eliminate this possibility.

**Cooperative Communication Leads to Abstract Symbols.** In contrast to adversarial signaling, listeners in a cooperative setting have an incentive to communicate and should share the effort of signal transmission by assuming an interpretive burden. While initial signals are likely to be iconic—so that their referents may be recognized by listeners—the eventual result of cooperative communication should be cost-minimizing signals. In two experiments, we simulate the evolution of cooperative communication and find that it results in

muted but effective signals (see Figure 3).

We run experiments under two conditions: one in which agents begin cooperating after a period of adversarial communication—providing them with a set of strongly-iconic signals—and another in which their first use of signals is cooperative. In both conditions we simulate a scenario where communicating with others is more lucrative than interacting with the environment directly, reflecting an emerging population of agents with specialized roles that relies on communication to survive. As the proportion of agents who interact directly with the environment decreases, agents become the primary sources of observed stimuli, alleviating pressures on speakers to create signals resembling environment observations. The result is the collapse of initially iconic signals into *abstract symbols* that drift while continuing to serve as reliable signifiers for referent behaviors. Furthermore, as agents grow proficient at communicating, the initial pressures to merely cross target listener decision-boundaries cease to dominate the reward landscape, and environmental constraints that influence the perception and creation of signals become exaggerated, e.g. as constraints imposed by writing implements and mediums are hypothesized to have impacted the look of many insular Asian scripts (Miller, 2014). We simulate these pressures using two signal penalties applied to speakers: a *curvature* penalty, which rewards agents for drawing straighter lines, and a *size* penalty, which penalizes agents for generating large signals.

## Part II: A Model of Pictorial Signification

In Part I, we demonstrated that primitive communication can spontaneously emerge once agents become capable of emitting stimuli, even when the method of processing observed

<sup>3</sup>We note a similarity to a GAN (Goodfellow et al., 2014).

stimuli remains the same. However, languages learned by such means are significantly limited by what they can reference. For a sign-referent pair to be added to the lexicon, speakers must induce referent behaviors by subjecting listeners to convincing stimuli. If the stimuli for existing S-R behaviors is complex relative to the signaling medium—as was the case for early human writing—agents may be unable to generate convincing replicas, leaving them unable to communicate. We refer to this barrier to communication as a *signification gap*, and suggest that new methods for generating and processing stimuli are necessary to bridge this gap.

## Models of Inferential Communication

Over the past decade, significant advancements have been made in formal accounts of *inferential* communication, wherein listeners glean the meaning of utterances by inferring the intentions of speakers based on context and shared knowledge (Degen, 2023). Most notably, Goodman and Frank (2016) introduced the Rational Speech Acts (RSA) framework, which formalized Gricean maxims (Grice, 1975; Sperber & Wilson, 1986) by modeling communication as a recursive process where agents infer intentions based on models of their interlocutor. Similar Bayesian inference models have previously been introduced to infer the goals of agents based on their behavior (Baker et al., 2009; Ullman et al., 2009) and more recent “inverse-inverse” planning methods (Chandra et al., 2023) have leveraged such models to guide the production of new behavior in a manner akin to *acting*. The core mathematics underpinning these methods are virtually identical: listeners infer meaning (e.g., as worlds  $w$  or goals  $g$ ) from signals (e.g., utterances  $u$  or behavior  $s \rightarrow a$ ) using Bayes’ rule and a “likelihood” model of speaker behavior:

RSA	Inverse Planning
$P_L(w u) \propto P_S(u w)P(w)$	$P(g s \rightarrow a) \propto P(s \rightarrow a g)P(g)$
$P_S(u w) \propto \exp(\alpha U(u; w))$	$P(s \rightarrow a g) \propto \exp(Q^g(s, a))$

At the core of these inferential models of communication lies an established relationship between the actions of an agent and its intended goal.<sup>4</sup> Before the invention of early writing systems, however, there was no conventional measure for relating drawn signs to their referents. Humans instead relied on a shared visual understanding to signify visually-complex referents using signs produced from simple writing implements.

## Pictographic Signification with Visual Theory of Mind

We seek to formalize the uniquely human capacity for *pictographic signification*. Recall that pictographs are effective not because viewers are genuinely convinced they are observing

<sup>4</sup>This is formalized in RSA by a utility function,  $U(u; w)$ , based on a denotational semantics providing the truth-value of a sentence in a given world. In inverse- and inverse-inverse planning it is formalized by a set of optimal Q-functions,  $\{Q^g(s, a)\}$ , capturing the degree to which behaviors are in service of a set of goals.

the referent,<sup>5</sup> but because they consider the space of possible referents when interpreting them. Interpreters of pictographs ask, “What does this pictograph best resemble, given what it could refer to and how it otherwise could have looked?” while creators of pictographs ask, “How can I best show likeness to my target referent while showing distinctness from other referents, given what I can draw?” By leveraging such *visual theory of mind* to reason about the drawing abilities and perception systems of others, humans can communicate referents using even crudely-drawn iconic pictographs. We formalize this model of communication as:

$$P_S(s|r_i) \propto \frac{f_{r_i}(s)}{\sum_{j \neq i} f_{r_j}(s)},$$

$$P_L(r_i|s) \propto P_S(s|r_i)P(r_i)P_R(r_i),$$

$$P_R(r_i) = \frac{1}{\int_S f_{r_i}(s)p(s) ds}.$$

Speakers select pictographs that maximally trigger the S-R classifier for a target referent and minimally trigger other classifiers,<sup>6</sup> while listeners apply Bayes’ rule to infer the likely referents of pictographs. Listeners are additionally tempered by a referent sensitivity term,  $P_R(r_i)$ , for capturing the overall likelihood of recognizing any pictograph as a given referent. This term is larger for referents with stimuli that are harder to replicate using pictographs, sensitizing listeners to referents with a greater *signification gap*.

## Methods

### Using Mirror Cognition to Implement Theory of Mind.

The inferential communication model described above relies on two critical assumptions: 1) that agents have shared access to the same S-R classifiers,  $\{f_{r_i}, \dots\}$ , when computing  $P_S(s|r_i)$  and  $P_R(r_i)$ ; and 2) that agents can explore the entire space of pictographs to compute the integral in  $P_R(r_i)$ . Neither are reasonable in a naturalistic setting where communication occurs between agents with siloed cognitive systems. Instead, we implement agents that leverage models of their own cognition as if they were models of their interlocutor (Gallese & Goldman, 1998), prompting agents to generate and interpret their own pictographs to estimate  $P_S(s|r_i)$  and  $P_R(r_i)$ . When speaking, agents generate a small number of pictographs and select the one that maximizes  $P_S(s|r_i)$ .

## Experiments

Experiments in this section use the following parameters:

- Observations from the underlying decision process are drawn from 5000 samples across 10 classes from the Cifar100 dataset (Krizhevsky et al., 2009) and converted to grayscale (see Figure 1). A more naturalistic dataset is chosen to demonstrate how communication emerges when there are large signification gaps. An additional iconicity probe is trained on this dataset.

<sup>5</sup>A pictographic horse is unlikely to convince you it is a real horse.

<sup>6</sup>This notion is similar to a *contrastive loss* (Hadsell et al., 2006).






Figure 4: **Comparison of Behaviorist and Inferential Signaling.** Agents wielding pictographic signification quickly succeed at communicating by generating iconic stimuli, achieving high rates of communicative success after only 300 epochs (second from right). As in Part I, we find that environmental pressures affect the appearance of pictographs over time (second group from left). We also report one agent’s estimate of  $P_R(r_i)$  for five referents, finding that palm trees are initially difficult to signify (evidenced by a large  $P_R$ ) while spiders are very easy to signify (evidenced by a small  $P_R$ ). As agents grow more proficient at signaling, most  $P_R$  values tend strictly toward the mean, while a small number are more volatile (e.g. dolphins and rockets) before drifting back toward the mean. We hypothesize this is due to agents mastering some referents before others, likely due to inherent difficulties in creating iconic signs for some referents. Analysis of the Cifar100 dataset finds that there is greater visual variance in the classes for which  $P_R$  is more volatile (e.g. rockets are sometimes obscured by plumes of exhaust smoke).

- The number of splines agents can draw is increased from 2 to 3 and the speaker canvas background color is adjusted to be closer to that of the average Cifar100 image.
- Each experiment is run with a population of 10 agents, which can now differentiate between environment observations and agent-generated pictographs, and use their inferential communication model only when interacting with other agents. Half of all observations processed are generated by agents, and half from the environment.
- Speakers generate and choose from 10 images at first, to a single image at epoch 750, reflecting agents with increased confidence in signaling capability over time.

**Visual Theory of Mind Bridges Signification Gaps Induced by Writing.** In a final set of experiments we examine the inferential capacities of agents wielding pictographic signification. Specifically, an analysis of one agent’s referent sensitivities  $P_R(r_i)$  suggests that the constraints of writing make some referents significantly harder to signify than others (everything appears to be a spider at first; see Figure 4). While this poses a problem for behaviorist agents without any meta-cognition, agents with theory of mind are unaffected because they can adjust their policies of sign interpretation based recognition likelihood, enabling the swift communication of most referents using iconic pictographs.

## Discussion

We set out to shed light on the mystery of why symbolic writing systems are unique to humans despite the ubiquity of symbolic communication in nature (Rosenthal, 2007) and the ability for some animals to learn abstract sign systems in structured learning settings (Patterson, 1978; Rumbaugh et al., 1974). These phenomena suggest that the capacity for symbolic communication alone is likely not sufficient for

the invention of writing systems, which differ in form from the signaling systems already used by many communicative animals. Furthermore, the presence of visually-recognizable features in the earliest forms of human writing suggest that abstract visual resemblance played a critical bootstrapping role in the adoption and invention of writing systems. We therefore suggest that this capacity—to leverage competence about one’s own visual system and a space of possible referents when interpreting the intended referent of an iconic pictograph—could be a unique feature of human cognition that enabled early humans to invent the first writing systems.

We assess this hypothesis by analyzing emergent languages invented by groups of reward-maximizing agents engaged in *signification games*, multi-agent decision-processes with partial observability that can be bridged with communication. We first implement a behaviorist model of communication akin to animal signaling and identify a critical shortcoming: agents struggle to communicate most referents due to a *signification gap* between their own capacity to signal and the stimuli needed to evoke those referents. We then formalize a model of *pictographic signification* inspired by Bayesian inference that leverages *visual theory of mind* to surmount signification gaps, enabling the communication of most referents with crude pictographs that resemble early forms of human writing. We additionally analyze the evolution of such pictographs and find that they grow visually-abstract over time, taking on a similar trajectory to ideographic writing systems invented by humans. Altogether, our results suggest that behaviorist models of communication are insufficient at explaining the emergence of writing systems, and that visual theory of mind may have played a crucial role in the invention of the first pictographic sign systems. We believe these results warrant further investigation into the extent to which other animals have the theory of mind necessary for performing signification.

## References

- Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning [Reinforcement learning and higher cognition]. *Cognition*, *113*(3), 329–349.
- Chandra, K., Li, T.-M., Tenenbaum, J., & Ragan-Kelley, J. (2023). Acting as inverse inverse planning. *ACM SIG-GRAPH 2023 Conference Proceedings*.
- Degen, J. (2023). The rational speech act framework. *Annual Review of Linguistics*, *9*(Volume 9, 2023), 519–540.
- Eco, U. (1976). *A theory of semiotics*. Indiana University Press.
- Fan, J. E., Bainbridge, W. A., Chamberlain, R., et al. (2023). Drawing as a versatile cognitive tool. *Nature Reviews Psychology*, *2*, 556–568.
- Fan, J. E., Hawkins, R. D., Wu, M., & Goodman, N. D. (2019). Pragmatic inference and visual abstraction enable contextual flexibility during visual communication. *Computational Brain & Behavior*, *3*(1), 86–101.
- Fan, J. E., Yamins, D. L. K., & Turk-Browne, N. B. (2018). Common object representations for visual production and recognition. *Cognitive Science*, *42*(8), 2670–2698.
- Fay, N., Ellison, M., & Garrod, S. (2014). Iconicity: From sign to system in human communication and language. *Pragmatics & Cognition*, *22*(2), 244–263.
- Ficken, M. S., Hailman, E. D., & Hailman, J. P. (1994). The chick-a-dee call system of the mexican chickadee. *The Condor*, *96*(1), 70–82.
- Gallese, V., & Goldman, A. (1998). Mirror neurons and the simulation theory of mind-reading. *Trends in Cognitive Sciences*, *2*(12), 493–501.
- Garrod, S., Fay, N., Lee, J., Oberlander, J., & MacLeod, T. (2007). Foundations of representation: Where might graphical symbol systems come from? *Cognitive Science*, *31*(6), 961–987.
- Gillespie-Lynch, K., Greenfield, P. M., Lyn, H., & Savage-Rumbaugh, S. (2014). Gestural and symbolic development among apes and humans: Support for a multimodal theory of language evolution. *Frontiers in Psychology*, *5*.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, & K. Weinberger (Eds.), *Advances in neural information processing systems* (Vol. 27). Curran Associates, Inc.
- Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, *20*(11), 818–829.
- Grice, H. P. (1975). Logic and conversation. In D. Davidson (Ed.), *The logic of grammar* (pp. 64–75). Dickenson Pub. Co.
- Hadsell, R., Chopra, S., & LeCun, Y. (2006). Dimensionality reduction by learning an invariant mapping. *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, *2*, 1735–1742.
- Halina, M., Rossano, F., & Tomasello, M. (2013). The ontogenetic ritualization of bonobo gestures. *Animal cognition*, *16*, 653–666.
- Hawkins, R. D., Sano, M., Goodman, N. D., et al. (2023). Visual resemblance and interaction history jointly constrain pictorial meaning. *Nature Communications*, *14*, 2199.
- Jiang, G., Hofer, M., Mao, J., Wong, L., Tenenbaum, J. B., & Levy, R. P. (2024). Finding structure in logographic writing with library learning.
- Krebs, J. R., & Dawkins, R. (1978). Animal signals: Mind-reading and manipulation. In *Behavioural ecology: An evolutionary approach*. Blackwell Scientific Publications.
- Krizhevsky, A., et al. (2009). Learning multiple layers of features from tiny images.
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. Burges, L. Bottou, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems* (Vol. 25). Curran Associates, Inc.
- Langford, J., & Zhang, T. (2007). The epoch-greedy algorithm for multi-armed bandits with side information. In J. Platt, D. Koller, Y. Singer, & S. Roweis (Eds.), *Advances in neural information processing systems* (Vol. 20). Curran Associates, Inc.
- LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proc. IEEE*, *86*, 2278–2324.
- Lewis, D. (1969). *Convention: A philosophical study*. Harvard University Press.
- Littman, M. L. (1994). Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994* (pp. 157–163). Elsevier.
- Mihai, D., & Hare, J. (2021). Learning to draw: Emergent communication through sketching. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, & J. W. Vaughan (Eds.), *Advances in neural information processing systems* (pp. 7153–7166, Vol. 34). Curran Associates, Inc.
- Miller, C. (2014). A survey of indigenous scripts of indonesia and the philippines. *Proceedings of the international workshop on endangered scripts of Island Southeast Asia*, 1–49.
- Oliehoek, F. A., & Amato, C. (2016). *A concise introduction to decentralized pomdps*. Springer International Publishing.
- Patterson, F. G. (1978). The gestures of a gorilla: Language acquisition in another pongid. *Brain and Language*, *5*(1), 72–97.
- Persson, T. (2008). *Pictorial primates: A search for iconic abilities in great apes* [Doctoral Thesis (monograph)]. Lund University Cognitive Studies. Cognitive Science.
- Qiu, S., Xie, S., Fan, L., Gao, T., Joo, J., Zhu, S.-C., & Zhu, Y. (2022). Emergent graphical conventions in a visual communication game. *Proceedings of the 36th International Conference on Neural Information Processing Systems*.
- Rosenthal, G. G. (2007). Spatiotemporal dimensions of visual signals in animal communication. *Annual Review of*

- Ecology, Evolution, and Systematics*, 38(Volume 38, 2007), 155–178.
- Rumbaugh, D. M., Von Glasersfeld, E., Warner, H., Pisani, P., & Gill, T. V. (1974). Lana (chimpanzee) learning language: A progress report. *Brain and Language*, 1(2), 205–212.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms.
- Sharma, P., Gero, S., Payne, R., Gruber, D. F., Rus, D., Torralba, A., & Andreas, J. (2024). Contextual and combinatorial structure in sperm whale vocalisations. *Nature Communications*, 15(1), 3617.
- Skinner, B. F. (1938). *The behavior of organisms: An experimental analysis*. Appleton-Century.
- Skyrms, B. (2010). *Signals: Evolution, learning, and information*. OUP Oxford.
- Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition*. Harvard University Press.
- Taylor, I. (1883). *The alphabet: An account of the origin and development of letters*. Kegan Paul, Trench.
- Toshitaka N. Suzuki, M. G., David Wheatcroft. (2016). Experimental evidence for compositional syntax in bird calls. *Nature Communications*, 7.
- Trigger, B. G. (1998). Writing systems: A case study in cultural evolution. *Norwegian Archaeological Review*, 31(1), 39–62.
- Ullman, T., Baker, C., Macindoe, O., Evans, O., Goodman, N., & Tenenbaum, J. (2009). Help or hinder: Bayesian models of social goal inference. *Advances in neural information processing systems*, 22.
- Von Frisch, K. (1962). Dialects in the language of the bees. *Scientific American*, 207(2), 78–89.
- Wittgenstein, L. (1953). *Philosophical investigations*. Macmillan.