PREDICTABILITY SHAPES ADAPTATION: AN EVOLUTIONARY PERSPECTIVE ON MODES OF LEARNING IN TRANSFORMERS

Anonymous authorsPaper under double-blind review

000

001

002

004

006

008 009 010

011 012

013

014

015

016

017

018

019

021

023

025

026

027

028

029

031

033

036

038

040

041

042

043

044

046

047

048

051

052

ABSTRACT

Transformer models learn in two distinct modes: in-weights learning (IWL), encoding knowledge into model weights, and in-context learning (ICL), adapting flexibly to context without weight modification. To better understand the interplay between these learning modes, we draw inspiration from evolutionary biology's analogous adaptive strategies: genetic encoding (akin to IWL, adapting over generations and fixed within an individual's lifetime) and phenotypic plasticity (akin to ICL, enabling flexible behavioral responses to environmental cues). In evolutionary biology, environmental predictability dictates the balance between these strategies: stability favors genetic encoding, while reliable predictive cues promote phenotypic plasticity. We experimentally operationalize these dimensions of predictability and systematically investigate their influence on the ICL/IWL balance in Transformers. Using regression and classification tasks, we show that high environmental stability decisively favors IWL, as predicted, with a sharp transition at maximal stability. Conversely, high cue reliability enhances ICL efficacy, particularly when stability is low. Furthermore, learning dynamics reveal taskcontingent temporal evolution: while a canonical ICL-to-IWL shift occurs in some settings (e.g., classification with many classes), we demonstrate that scenarios with easier IWL (e.g., fewer classes) or slower ICL acquisition (e.g., regression) can exhibit an initial IWL phase later yielding to ICL dominance. These findings support a relative-cost hypothesis for explaining these learning mode transitions, establishing predictability as a critical factor governing adaptive strategies in Transformers, and offering novel insights for understanding ICL and guiding training methodologies.

1 Introduction

Transformer architectures (Vaswani et al., 2017) underpin the success of modern large language models (LLMs), demonstrating impressive performance across diverse tasks (Brown et al., 2020). A key emergent capability in these models is in-context learning (ICL), where models adapt to novel tasks specified by examples within the input prompt, without requiring gradient-based parameter updates (Brown et al., 2020; Dong et al., 2022; Liu et al., 2023). This rapid, flexible adaptation contrasts with standard in-weights learning (IWL), where knowledge is gradually encoded into model parameters during training (Chan et al., 2022; Singh et al., 2023). The mechanisms enabling ICL are an active area of research, with theories encompassing induction heads that copy and complete patterns (Olsson et al., 2022; Elhage et al., 2021), function vector heads that compute task representations (Todd et al., 2023; Hendel et al., 2023), and interpretations of ICL as implicit optimization or Bayesian inference within the forward pass (Von Oswald et al., 2023; Dai et al., 2022; Xie et al., 2021).

While IWL traditionally yields stable knowledge from consistent task exposure, the interplay and developmental trajectory between IWL and ICL remain incompletely understood. Notably, ICL capabilities can be transient, emerging early in training but sometimes diminishing or being superseded by IWL as tasks become more predictable or frequently encountered (Singh et al., 2023; Anand et al., 2024). Characteristics of the data distribution, such as input noise and dynamic input-label mappings – which influence task predictability – have also been shown to affect the balance

between ICL and IWL (Chan et al., 2022; Reddy, 2023). Despite these advances, a comprehensive framework for why and when Transformers shift their reliance between these learning modes is still developing.

In this paper, we argue that evolutionary biology offers one explanatory lens for understanding this dynamic interplay. Organisms adapt to variable environments via analogous mechanisms: phenotypic plasticity and genetic evolution. Phenotypic plasticity – the ability of a genotype to produce different phenotypes (e.g., changes in morphology, physiology, or behavior like learning) in response to environmental cues (West-Eberhard, 2003; Scheiner, 1993) – allows rapid, within-lifetime adaptation, akin to a Transformer's ICL adapting to a prompt. This contrasts with genetic evolution – the slower, generational modification of the genotype – which encodes more permanent traits, analogous to IWL.

Evolutionary biology has established that environmental predictability critically determines the favored adaptive strategy (Stearns, 1989). Phenotypic plasticity is generally favored when the environment fluctuates and reliable cues are available to guide an appropriate adaptive response (Stearns, 1989; Stephens, 1989). However, plasticity may be selected against if the environment is extremely stable (negating the need for flexibility and favoring genetically hardcoded traits), if cues are unreliable, or if the inherent costs of plasticity are too high (Stearns, 1989; Stephens, 1989; DeWitt et al., 1998). Furthermore, a key long-term dynamic is genetic assimilation: if an environment that previously favored a plastic response becomes stable, the consistently advantageous phenotype, initially induced plastically by cues, can become genetically encoded, shifting the trait from flexible to fixed (Waddington, 1953; Pigliucci et al., 2006). This offers a compelling parallel to the observed transience of ICL.

Building on this evolutionary framework, we hypothesize that analogous dimensions of task predictability govern which modes of learning Transformers adopt and how their balance shifts over time. We operationalize two key dimensions, inspired by these biological concepts: *cue reliability*, defined as the clarity and informativeness of information within a single prompt for specifying the required task, and *environmental stability*, defined as the consistency or predictability of the task structure across different prompts encountered over the course of training. We predict that high environmental stability will favor IWL, while high cue reliability will enhance ICL efficacy, particularly when stability is low. Moreover, we anticipate that conditions of high stability for tasks initially learned via ICL may lead to an "assimilation" of these solutions into IWL.

To test these hypotheses, we employ a controlled experimental design that systematically investigates the independent and interactive effects of cue reliability and environmental stability on the use and development of ICL and IWL. We deploy this approach in two distinct domains: a parametric sinusoid regression task and a few-shot binary classification task on Omniglot images. Our findings aim to establish predictability as a critical factor shaping adaptive learning in Transformers, offering insights for understanding ICL and potentially guiding future training methodologies.

2 LEARNING DYNAMICS THROUGH AN EVOLUTIONARY LENS

Our approach leverages analogies between biological adaptation and Transformer learning to investigate the influence of environmental predictability. The core idea is that the emergence, persistence, and transitions between ICL and IWL in Transformers can be understood as parallels to how organisms evolve to use phenotypic plasticity versus genetically determined traits, and how these strategies themselves can shift (e.g., via genetic assimilation).

To make these parallels concrete:

ICL as **Phenotypic Plasticity:** Consider the water flea (Daphnia), which develops protective helmets and spines (phenotype) only when detecting chemical cues from predators (environmental cue) (Tollrian & Harvell, 1999). Similarly, many animals adjust coat thickness and color based on changing daylight hours – a cue predictive of seasonal temperature and snowfall (Underwood & Reynolds, 1980; Otte et al., 2024). This rapid, reversible, cue-driven adaptation mirrors how a Transformer uses in-context examples (cue) to tailor its output (phenotype) to a novel task without altering its weights (genotype).

IWL as Genetically Determined/Canalized Traits: The innate, species-specific song of certain songbirds develops reliably, even in acoustic isolation. This exemplifies a genetically underpinned trait that is invariant to environmental variation, a phenomenon known as environmental canalization (Marler, 2014). Analogously, IWL captures fundamental, consistently useful knowledge (e.g., grammar for an LLM) that is "hardcoded" into model weights through sustained exposure to stable patterns.

ICL Transience as Genetic Assimilation: Waddington's classic experiments showed that a heat-shock-induced crossveinless wing phenotype in Drosophila (a plastic response), after generations of selection, became genetically assimilated, appearing even without the heat shock (Waddington, 1953). Analogously, if a Transformer initially uses ICL for a task type that later becomes a stable and frequent part of training, the model might "assimilate" this solution into its weights (IWL), potentially reducing reliance on ICL for that specific pattern.

Guided by this framework, particularly the established role of environmental stability and cue reliability in shaping biological adaptation, our experimental design systematically manipulates these two factors within the training regime of Transformers. For cue reliability, we vary the clarity or consistency of the in-context examples provided within each prompt. For environmental stability, we manipulate the consistency of the underlying task rule across successive training batches.

3 EXPERIMENTAL DESIGN

3.1 MODEL AND TRAINING REGIME

Drawing on the previous literature, we use both a sinusoid regression task (Finn et al., 2017) and an Omniglot classification task (Lake et al., 2015). For both experimental tasks, we employ a decoder-only Transformer architecture with 4 layers, 4 attention heads per layer, an embedding dimension (d_{model}) of 128, and learned positional encodings.

Input processing varies by task. For the sinusoid regression task, scalar inputs and outputs are linearly projected to $d_{\rm model}$. For Omniglot classification, character images are processed by a shallow ResNet featuring three stages, each with 2 residual blocks, and outputting 16, 32, and 64 channels per stage, respectively. This ResNet is trained jointly with the Transformer and projects the extracted image features to $d_{\rm model}$.

Models are optimized to minimize either the mean squared error for sinusoid regression or the binary cross-entropy loss for Omniglot classification, calculated on their predictions for the query item (y_q) . We use the AdamW optimizer (Loshchilov & Hutter, 2017) with a peak learning rate of 1×10^{-4} , subject to a cosine decay schedule following 1,000 warmup steps. All models are trained for a total of 50,000 training steps, using a batch size of 128. Each training sequence presented to the model consists of N prompt example pairs followed by a single query item. The specific value of N differs per task, as detailed below.

3.2 TASK 1: SINUSOID REGRESSION

The first task requires the model to perform regression on points sampled from a sinusoidal function (see Figure 1), a setup inspired by prior meta-learning studies (Finn et al., 2017).

Each training batch corresponds to a single environmental timestep t. Within a batch, all training sequences are generated from the same true underlying sinusoid $f_t(x) = A_t \sin(x + \phi_t)$. This setup mirrors a population adapting to a common, evolving environment.

For each sequence, the model is provided with N=10 prompt examples $\{(x_i,y_i)\}_{i=1}^N$. Inputs x_i are sampled uniformly from $[-\pi,\pi]$, and corresponding outputs are $y_i=f_t(x_i)+\delta_i$, where δ_i represents Gaussian observation noise specific to each prompt example. The model's objective is to predict the noise-free value $y_q=f_t(x_q)$ for a given query input x_q .

Cue Reliability: This is manipulated by adjusting the variance, $\sigma_{\text{reliability}}^2$, of the Gaussian observation noise $\delta_i \sim \mathcal{N}(0, \sigma_{\text{reliability}}^2)$ applied to the y_i values in the prompt. Lower $\sigma_{\text{reliability}}^2$ corresponds to higher cue reliability.

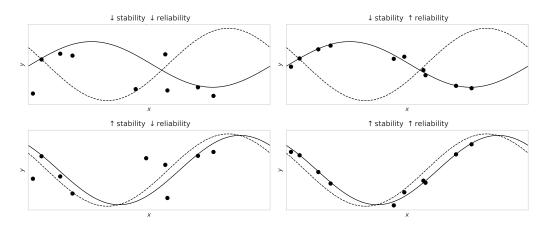


Figure 1: Illustration of the sinusoid regression task under varying predictability conditions. Each panel shows prompt examples (black dots) for a given underlying sinusoid at timestep t (solid line), and the sinusoid at the preceding timestep t-1 (dashed line). Top row: Low environmental stability, meaning the underlying sinusoid (solid line) changes significantly at each step. Bottom row: High environmental stability, where the underlying sinusoid is more stable across time. Left column: Low cue reliability, where prompt examples are noisy. Right column: High cue reliability, where prompt examples are less noisy and closer to the true underlying sinusoid.

Environmental Stability: This is manipulated by controlling the parameter $\alpha_{\text{stability}}$ in an AR(1) process governing the evolution of the sinusoid parameters $\theta_t = [A_t, \phi_t]^T$ across successive environmental timesteps t. A higher $\alpha_{\text{stability}}$ indicates greater stability. The parameter update rule is:

$$\theta_t = \alpha_{\text{stability}} \theta_{t-1} + (1 - \alpha_{\text{stability}}) \theta_{\text{innovation}}^{(t)} \tag{1}$$

where θ_0 is initially sampled from prior distributions. For each subsequent timestep t, $\theta_{\rm innovation}^{(t)}$ is a vector of innovations drawn afresh from these same priors: amplitude $A \sim U[0.5, 1.5]$ and phase $\phi \sim U[0, 2\pi]$. This formulation ensures mean-reversion towards the center of the prior distributions, with $\alpha_{\rm stability}$ dictating the rate of change.

3.3 TASK 2: OMNIGLOT FEW-SHOT BINARY CLASSIFICATION

The second task is few-shot binary classification of Omniglot characters (Lake et al., 2015) (see Figure 2). Similar to the sinusoid task, each training batch corresponds to a single environmental timestep t. At each timestep t, an underlying global mapping, $M_t: \mathcal{C} \to \{0,1\}$, assigns a binary label to each character c from a set \mathcal{C} of 1623 unique Omniglot character classes.

A prompt sequence comprises N=2 example image-label pairs $\{(x_i,y_i)\}_{i=1}^N$, followed by a query image x_q (an exemplar of character class c_q). The model predicts the true binary label $M_t(c_q)$. Critically, one example character in the prompt, c_k , matches the query's class $c_k=c_q$. This allows the model to predict the query label by relying on the in-context label y_k (ICL) or by accessing an internalized representation of $M_t(c_q)$ (IWL).



Figure 2: Sequences drawn from the Omniglot binary classification task. (**Left**) A sequence illustrating a reliable in-context cue (prompt label y_1 for character $c_1 = c_q$ matches $M_t(c_q)$). (**Right**) A sequence illustrating an unreliable in-context cue (prompt label y_2 for $c_2 = c_q$ is flipped relative to $M_t(c_q)$).

Cue Reliability: This is manipulated by controlling the probability, $p_{\text{reliability}}$, that the label y_i for a prompt character c_i matches its true label $M_t(c_i)$. Specifically:

$$y_i = \begin{cases} M_t(c_i) & \text{with probability } p_{\text{reliability}} \\ 1 - M_t(c_i) & \text{with probability } 1 - p_{\text{reliability}} \end{cases}$$
 (2)

Higher $p_{\text{reliability}}$ indicates higher cue reliability.

Environmental Stability: This is manipulated by varying the probability, $p_{\text{stability}}$, that the true mapping $M_t(c)$ for each character c persists from the previous timestep $M_{t-1}(c)$. The initial mapping $M_0(c)$ is sampled from Bernoulli(0.5) for all characters. For t > 0:

$$M_t(c) = \begin{cases} M_{t-1}(c) & \text{with probability } p_{\text{stability}} \\ 1 - M_{t-1}(c) & \text{with probability } 1 - p_{\text{stability}} \end{cases}$$
 (3)

Higher $p_{\text{stability}}$ signifies greater environmental stability. The first-order temporal autocorrelation for each character's label in this process is $\alpha_{\text{stability}} = 2p_{\text{stability}} - 1$.

3.4 EVALUATION: ASSESSING PREFERENCE BETWEEN ICL AND IWL

To assess a trained model's preference for ICL versus IWL solutions, we employ a targeted evaluation protocol inspired by prior work (Chan et al., 2022). This evaluation is performed periodically during training. It uses novel evaluator tasks sampled from the same prior distributions as the training tasks but unseen during training. This allows us to contrast mappings learned in weights against those learnable from context during evaluation.

Evaluator Task Generation: For sinusoid regression, an evaluator task is a specific function $f_e(x) = A_e \sin(x + \phi_e)$ with (A_e, ϕ_e) newly sampled from their priors (Section 3.2). For Omniglot classification, an evaluator task is a complete binary mapping $M_e: \mathcal{C} \to \{0,1\}$ where each $M_e(c)$ is independently sampled from Bernoulli(0.5).

Prompting and Prediction: For each evaluator task $(f_e \text{ or } M_e)$ and a query input x_q , we construct a noiseless evaluation prompt $P_e = \{(x_j, y_j)\}_{j=1}^N$, with examples derived directly from f_e or M_e . For Omniglot, one example in P_e matches c_q . The model then predicts \hat{y}_q (sinusoid) or $\hat{p}_q = P(y_q = 1|P_e, x_q)$ (Omniglot).

Target Definitions for Preference Assessment: At the moment of periodic evaluation, we define two distinct targets for the model's prediction:

- ICL Target $(y_{q,ICL})$: The ground-truth for x_q from the current evaluator task: $f_e(x_q)$ or $M_e(c_q)$. This target is independent of the training environment's state.
- IWL Target $(y_{q,\text{IWL}})$: The ground-truth for x_q from the current state of the training environment generator (i.e., $f_t(x_q)$ or $M_t(c_q)$ at the time of evaluation). This reflects what a model relying purely on an internalized representation of the latest training environment state would predict.

Preference Quantification: We calculate the model's prediction error relative to both targets ($E_{\rm ICL}$ and $E_{\rm IWL}$) using mean squared error (sinusoid) or binary cross-entropy (Omniglot). A model prefers ICL if $E_{\rm ICL} > E_{\rm IWL}$. We quantify this preference using the score:

$$S_{\rm ICL} = \frac{E_{\rm IWL}}{E_{\rm ICL} + E_{\rm IWL} + \epsilon_{\rm stab}} \tag{4}$$

where ϵ_{stab} (e.g., 1×10^{-8}) ensures numerical stability. $S_{\text{ICL}} > 0.5$ indicates IWL preference. This metric captures strategy preference, not absolute task performance.

4 ANALYSIS AND RESULTS

Our experiments reveal that the learning strategies adopted by Transformers are systematically shaped by the predictability of their training environment, in ways that parallel adaptive mechanisms in evolutionary biology. We first examine how environmental stability and cue reliability

dictate the asymptotic preference for ICL vs. IWL. We then explore the dynamic learning trajectories, or transience, of these strategies. Finally, we introduce and provide evidence for a *relative-cost hypothesis*, which posits that the computational cost of acquiring ICL versus IWL solutions governs these dynamic preferences and reconciles observed behaviors with biological principles.

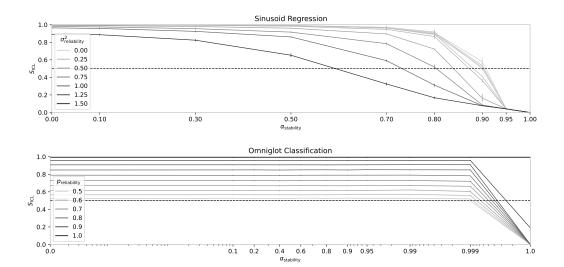


Figure 3: ICL preference score ($S_{\rm ICL}$) as a function of environmental stability and cue reliability. Higher S $S_{\rm ICL}$ (y-axis) indicates greater ICL reliance. Environmental stability (x-axis) increases from left to right. Different lines represent varying levels of cue reliability. The dashed line at $S_{\rm ICL}=0.5$ marks balanced preference between ICL and IWL. (**Top**) Sinusoid regression task: Environmental stability is parameterized by the AR(1) autocorrelation $\alpha_{\rm stability}$. Cue reliability increases with decreasing prompt noise variance $\sigma^2_{\rm reliability}$ (lighter lines denote higher reliability). (**Bottom**) Omniglot binary classification task: Environmental stability is parameterized by the task mapping autocorrelation ($\alpha_{\rm stability}=2p_{\rm stability}-1$, plotted on a logit scale). Cue reliability increases with prompt label correctness probability $p_{\rm reliability}$ (darker lines denote higher reliability).

4.1 Environmental Predictability Governs Asymptotic Learning Preferences

Figure 3 visualizes how dimensions of environmental predictability influence the asymptotic preference between ICL and IWL, supporting our core hypothesis.

In the sinusoid regression task (Figure 3, Top), increasing environmental stability (higher $\alpha_{\rm stability}$) consistently leads to a marked reduction in ICL preference ($S_{\rm ICL}$), with scores declining sharply as $\alpha_{\rm stability}$ approaches 1. This mirrors evolutionary principles where stable environments favor fixed, genetically encoded traits over costly phenotypic plasticity (Stearns, 1989; DeWitt et al., 1998). Furthermore, higher cue reliability (lower $\sigma^2_{\rm reliability}$, lighter lines) generally fosters greater ICL preference, especially at low-to-moderate environmental stability. This is analogous to biological systems where reliable cues are crucial for adaptive plasticity in variable environments (Stephens, 1989; Scheiner, 1993). However, cue reliability's impact diminishes under very high environmental stability ($\alpha_{\rm stability} \approx 1$), where IWL becomes dominant regardless of prompt clarity, consistent with the biological view that if an environment seldom changes, plasticity (and cue sensitivity) offers little benefit.

The Omniglot binary classification task (Figure 3, Bottom) shows a generally sharper response but reflects similar principles. Across a broad range of environmental stability (up to $\alpha_{\text{stability}} \approx 0.95$), cue reliability (higher $p_{\text{reliability}}$, lighter lines) primarily governs ICL preference, underscoring the tenet that reliable cues are essential for plasticity in non-static environments (Stephens, 1989). As the environment approaches perfect stability ($\alpha_{\text{stability}} = 1$), ICL preference plummets, indicating a robust shift towards IWL, aligning with evolutionary concepts like genetic assimilation or trait fixation (Waddington, 1953; Pigliucci et al., 2006).

Notably, in the Omniglot task with perfectly reliable cues ($p_{\text{reliability}} = 1$), models can maintain strong ICL preference even in a perfectly stable environment. This contrasts with typical biological assumptions that maintaining plasticity machinery incurs costs (DeWitt et al., 1998), suggesting that when ICL and IWL strategies are equally effective at task performance, their relative computational cost to acquire or implement might determine preference. This intuition is formalized as the relative-cost hypothesis in Section 4.3.

Overall, these asymptotic results demonstrate that environmental predictability systematically shapes Transformer learning strategies, extending prior work showing that factors like noise and dynamic mappings affect ICL (Chan et al., 2022), and aligning broadly with predictions from evolutionary biology.

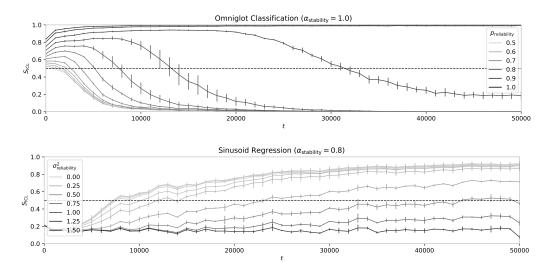


Figure 4: Contrasting learning dynamics of ICL preference ($S_{\rm ICL}$) over training steps (t) for different levels of cue reliability, at a fixed level of environmental stability. (**Top**) Omniglot classification (perfect environmental stability, $p_{\rm stability} = 1.0$): Illustrates typical ICL transience, where high initial $S_{\rm ICL}$ decays over training. Decay rates vary with cue reliability ($p_{\rm reliability}$, different lines). (**Bottom**) Sinusoid regression (environmental stability $\alpha_{\rm stability} = 0.8$): Demonstrates IWL transience, where $S_{\rm ICL}$ starts lower and gradually increases, especially with higher cue reliability (lower $\sigma^2_{\rm reliability}$, lighter lines). The dashed line indicates $S_{\rm ICL} = 0.5$.

4.2 LEARNING DYNAMICS: TASK-DEPENDENT TRANSIENCE BETWEEN ICL AND IWL

Beyond asymptotic states, learning trajectories reveal dynamic shifts in strategy preference – a phenomenon known as *transience* (Singh et al., 2023; Anand et al., 2024; Singh et al., 2025). As shown in Figure 4, this changing balance between ICL versus IWL is task-dependent, reminiscent of developmental or evolutionary shifts in organisms adapting to sustained environmental pressures.

Specifically, for the Omniglot task under high environmental stability (Figure 4, Top), we typically observe "ICL transience" (Singh et al., 2023). The model exhibits strong preference for ICL early in training but then gradually shift towards an IWL solution, as the model internalizes the stable character-to-label mapping. This progression strongly parallels *genetic assimilation*, where an initially plastic response becomes genetically encoded under prolonged, stable selection (Waddington, 1953; Pigliucci et al., 2006).

In contrast, the sinusoid regression task (Figure 4, Bottom) can exhibit the opposite phenomenon of "IWL transience" under high stability. Here, the model initially exhibits a lower preference for ICL, or a preference that increases more slowly, with a notable increase often occurring later in training, particularly with high cue reliability. This suggests an initial phase where a simpler global IWL solution is more readily acquired, with the subsequent rise in ICL preference reflecting the slower development of more complex ICL circuitry. This pattern is analogous to an organism first establishing a basic viable trait before refining more sophisticated plastic capabilities, and it also relates to

phenomena like "grokking," where generalization emerges after initial memorization (Power et al., 2022).

These contrasting transience patterns indicate that the particular learning paths depend on the specific computational challenges posed within a given task – a notion explored next.

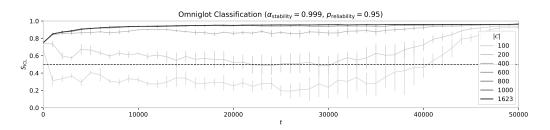


Figure 5: Reversal of transience in the Omniglot task by varying IWL difficulty (total character set size $|\mathcal{C}|$). Plots show ICL preference ($S_{\rm ICL}$) across training steps (t) under high environmental stability ($\alpha_{\rm stability}=0.999$) and high cue reliability ($p_{\rm reliability}=0.95$). Reducing $|\mathcal{C}|$ (e.g., to $|\mathcal{C}|=100$, lightest gray line) makes IWL easier, leading to an initial phase of lower $S_{\rm ICL}$ (IWL preference) followed by a gradual rise, characteristic of IWL transience. The dashed line is $S_{\rm ICL}=0.5$.

4.3 THE RELATIVE-COST HYPOTHESIS

We propose a *relative-cost hypothesis* to explain both the persistence of ICL under perfect conditions in Omniglot (Section 4.1) and the observed task-dependent transience patterns (Section 4.2). This hypothesis posits that the choice between ICL and IWL, and the direction of any transience, is governed by the relative computational cost – in terms of ease of learning, representational complexity, or inference efficiency – of implementing a solution via each mode for a specific task configuration.

For instance, our Omniglot involves a large character set ($|\mathcal{C}|=1623$), such that ICL (e.g., pattern matching from prompt examples) might be computationally simpler to discover initially. In contrast, IWL (memorizing the entire $M_t:\mathcal{C}\to\{0,1\}$ mapping) is a high-capacity memorization task, becoming viable only with extended training if M_t is stable. Conversely, for sinusoid regression, an IWL approach (learning a single global sinusoid, especially with moderate stability) could be simpler to parameterize quickly. Developing sophisticated ICL for arbitrary few-shot regression might represent a more complex circuit to learn, thus emerging later.

To directly test this, we manipulated IWL difficulty in the Omniglot task by varying the total character set size, $|\mathcal{C}|$. As Figure 5 shows, significantly reducing $|\mathcal{C}|$ (e.g., to 100, lightest gray line), thereby making IWL substantially easier, reverses the learning dynamics under high environmental stability and cue reliability. $S_{\text{ICL}}(t)$ starts low, indicating an initial preference for the easier IWL strategy, and then gradually increases. This induced IWL transience—in a task that typically shows ICL transience (Singh et al., 2023; 2025)—strongly supports our hypothesis: making IWL relatively easier to acquire leads to its initial preference.

These results align with previous findings that increasing classifier class count can affect ICL transience (Singh et al., 2023), and with models relating strategy acquisition speed to transience (Nguyen & Reddy, 2024; Singh et al., 2025). They also resonate with studies on "circuit efficiency" and regularization in grokking (Thilak et al., 2022; Nanda et al., 2023; Varma et al., 2023).

Together, these findings suggest that Transformers, like biological systems, navigate a trade-off. Instead of metabolic or developmental costs, Transformers appear to balance computational costs and learning efficiencies of different internal mechanisms. An easier strategy may be adopted first, while a more complex but potentially more effective or general strategy develops over a longer timescale, echoing how biological systems often evolve solutions favoring immediate viability while allowing subsequent refinement. The framework of *resource rational analysis* (Lieder & Griffiths, 2020), introduced in cognitive science but inspired by earlier work in AI (Horvitz, 1987; Russell & Subramanian, 1994), may provide a set of tools for formalizing this tradeoff in a way that can offer further insight into the behavior of AI models.

5 DISCUSSION

Our investigation, drawing analogies from evolutionary biology, has demonstrated that core principles of environmental predictability systematically govern the balance and dynamics between incontext learning (ICL) and in-weights learning (IWL) in Transformer models. We found that high environmental stability strongly promotes IWL, mirroring the selection for genetically fixed traits in stable biological environments, while high cue reliability enhances ICL efficacy, akin to phenotypic plasticity's reliance on dependable environmental signals. Furthermore, our findings revealed distinct, task-dependent transience patterns: models sometimes shift from ICL to IWL (resembling genetic assimilation) or, conversely, from IWL to ICL (resembling the evolution of plasticity). We proposed and provided evidence for a relative-cost hypothesis, suggesting these dynamics are driven by the comparative ease of learning and implementing each strategy within a given model and task context. These results affirm the utility of an evolutionary lens for understanding adaptive learning in artificial neural networks.

5.1 LIMITATIONS AND FUTURE DIRECTIONS

While our controlled experiments are aimed at providing foundational insights, certain limitations define the scope of the current work. We also describe exciting areas for future work.

Task Simplification: The sinusoid regression and Omniglot classification tasks are, by design, simplifications of the complex, high-dimensional environments (e.g., natural language) where large language models typically operate. Such controlled settings are, however, essential for isolating variables and establishing clearer causal links between predictability and learning dynamics. They provide a foundational understanding that can inform future investigations in more complex domains.

Computational versus Mechanistic Analogy: This paper offers a computational-level analysis in the tradition of Marr (Marr, 1982; Ku et al., 2025). We characterize the adaptive challenges a system faces (coping with environmental predictability) and the functional solutions it employs (ICL as flexible plasticity, IWL as assimilated knowledge). The analogy to evolutionary biology operates at this functional level of adaptation to environmental statistics; we do not claim direct equivalence between the learning mechanisms in Transformers and the specific genetic or developmental pathways in biology.

The Baldwin Effect and Learning Speed: A key argument for biological plasticity, beyond managing uncertainty, is its potential to accelerate evolution on complex fitness landscapes (e.g., via the Baldwin Effect). Plastic individuals can more effectively explore and identify advantageous phenotypic regions, which can then be refined and genetically assimilated (Hinton et al., 1987; Fernando et al., 2018). An analogous question for Transformers is whether ICL not only provides flexibility but also speeds up the acquisition of robust IWL solutions for complex tasks, perhaps by IWL reusing or refining circuits initially developed for ICL (Singh et al., 2025).

Exploring Diverse Predictability Regimes: Our use of noise for cue reliability and AR(1) processes for environmental stability represents specific implementations of predictability. Future work should explore other forms of environmental structure and change. This includes systematically varying the nature of aleatory uncertainty beyond our current manipulations – for example, by introducing non-stationary predictability or more complex temporal structures (e.g., environments with punctuated equilibria (Gould & Eldredge, 1977)).

5.2 Conclusion

This research underscores the significant value of drawing on established principles from evolutionary theory to illuminate complex learning behaviors in artificial systems. The insights gained not only contribute to a more fundamental understanding of how Transformers adapt to their training environments but also pave the way for designing more effective training methodologies. Ultimately, this line of inquiry can help guide the development of artificial systems capable of robustly and flexibly navigating the diverse and dynamic environments they increasingly encounter.

ETHICS AND REPRODUCIBILITY STATEMENTS

Ethics Statement: This research aims to provide a foundational understanding of the learning dynamics in Transformer models by drawing an analogy to evolutionary biology. Our experiments were conducted exclusively on synthetic data and the publicly available Omniglot dataset, as detailed in Section 3. We do not foresee any direct negative societal consequences stemming from this work.

Reproducibility Statement: To ensure the reproducibility of our findings, we provide a detailed description of our model architecture, training regime, and task design in Section 3. Appendix A contains a comprehensive list of all hyperparameters for our experimental sweeps, as well as the compute resources used. The complete code for data generation, model training, and analysis will be released upon publication.

REFERENCES

- Suraj Anand, Michael A Lepori, Jack Merullo, and Ellie Pavlick. Dual process learning: Controlling use of in-context vs. in-weights strategies with weight forgetting. *arXiv preprint arXiv:2406.00053*, 2024.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901, 2020.
- Stephanie Chan, Adam Santoro, Andrew Lampinen, Jane Wang, Aaditya Singh, Pierre Richemond, James McClelland, and Felix Hill. Data distributional properties drive emergent in-context learning in transformers. *Advances in Neural Information Processing Systems*, 35:18878–18891, 2022.
- Damai Dai, Yutao Sun, Li Dong, Yaru Hao, Shuming Ma, Zhifang Sui, and Furu Wei. Why can gpt learn in-context? language models implicitly perform gradient descent as meta-optimizers. *arXiv* preprint arXiv:2212.10559, 2022.
- Thomas J DeWitt, Andrew Sih, and David Sloan Wilson. Costs and limits of phenotypic plasticity. *Trends in Ecology & Evolution*, 13(2):77–81, 1998.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Jingyuan Ma, Rui Li, Heming Xia, Jingjing Xu, Zhiyong Wu, Tianyu Liu, et al. A survey on in-context learning. *arXiv preprint arXiv:2301.00234*, 2022.
- Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, et al. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 1(1):12, 2021.
- Chrisantha Fernando, Jakub Sygnowski, Simon Osindero, Jane Wang, Tom Schaul, Denis Teplyashin, Pablo Sprechmann, Alexander Pritzel, and Andrei Rusu. Meta-learning by the baldwin effect. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pp. 1313–1320, 2018.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pp. 1126–1135. PMLR, 2017.
- Stephen Jay Gould and Niles Eldredge. Punctuated equilibria: the tempo and mode of evolution reconsidered. *Paleobiology*, 3(2):115–151, 1977.
- Roee Hendel, Mor Geva, and Amir Globerson. In-context learning creates task vectors. *arXiv* preprint arXiv:2310.15916, 2023.
 - Geoffrey E Hinton, Steven J Nowlan, et al. How learning can guide evolution. *Complex systems*, 1 (3):495–502, 1987.
 - Eric J Horvitz. Reasoning about beliefs and actions under computational resource constraints. *Proceedings of the third conference on uncertainty in artificial intelligence*, pp. 429–44, 1987.

- Alexander Ku, Declan Campbell, Xuechunzi Bai, Jiayi Geng, Ryan Liu, Raja Marjieh, R Thomas McCoy, Andrew Nam, Ilia Sucholutsky, Veniamin Veselovsky, et al. Using the tools of cognitive science to understand large language models at different levels of analysis. *arXiv preprint arXiv:2503.13401*, 2025.
 - Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
 - Falk Lieder and Thomas L Griffiths. Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, 43:e1, 2020.
 - Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pretrain, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM computing surveys*, 55(9):1–35, 2023.
 - Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint* arXiv:1711.05101, 2017.
 - Peter Marler. The instinct to learn. In *The epigenesis of mind*, pp. 37–66. Psychology Press, 2014.
 - David Marr. Vision: A Computational Investigation into the Human Representation and Processing of Visual Information. Henry Holt and Co., Inc., New York, NY, USA, 1982.
 - Neel Nanda, Lawrence Chan, Tom Lieberum, Jess Smith, and Jacob Steinhardt. Progress measures for grokking via mechanistic interpretability. *arXiv preprint arXiv:2301.05217*, 2023.
 - Alex Nguyen and Gautam Reddy. Differential learning kinetics govern the transition from memorization to generalization during in-context learning. *arXiv* preprint arXiv:2412.00104, 2024.
 - Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, et al. In-context learning and induction heads. *arXiv preprint arXiv:2209.11895*, 2022.
 - Pieter J Otte, Joris PGM Cromsigt, Christian Smit, and Tim R Hofmeester. Snow cover-related camouflage mismatch increases detection by predators. *Journal of Experimental Zoology Part A: Ecological and Integrative Physiology*, 341(4):327–337, 2024.
 - Massimo Pigliucci, Courtney J Murren, and Carl D Schlichting. Phenotypic plasticity and evolution by genetic assimilation. *Journal of Experimental Biology*, 209(12):2362–2367, 2006.
 - Alethea Power, Yuri Burda, Harri Edwards, Igor Babuschkin, and Vedant Misra. Grokking: Generalization beyond overfitting on small algorithmic datasets. *arXiv preprint arXiv:2201.02177*, 2022.
 - Gautam Reddy. The mechanistic basis of data dependence and abrupt learning in an in-context classification task. *arXiv* preprint arXiv:2312.03002, 2023.
 - Stuart J Russell and Devika Subramanian. Provably bounded-optimal agents. *Journal of Artificial Intelligence Research*, 2:575–609, 1994.
 - Samuel M Scheiner. Genetics and evolution of phenotypic plasticity. *Annual review of Ecology and Systematics*, 24(1):35–68, 1993.
 - Aaditya Singh, Stephanie Chan, Ted Moskovitz, Erin Grant, Andrew Saxe, and Felix Hill. The transient nature of emergent in-context learning in transformers. *Advances in Neural Information Processing Systems*, 36:27801–27819, 2023.
- Aaditya K Singh, Ted Moskovitz, Sara Dragutinovic, Felix Hill, Stephanie CY Chan, and Andrew M Saxe. Strategy coopetition explains the emergence and transience of in-context learning. *arXiv* preprint arXiv:2503.05631, 2025.
 - Stephen C Stearns. The evolutionary significance of phenotypic plasticity. *Bioscience*, 39(7):436–445, 1989.

- DW Stephens. Variance and the value of information. *The American Naturalist*, 134(1):128–140, 1989.
- Vimal Thilak, Etai Littwin, Shuangfei Zhai, Omid Saremi, Roni Paiss, and Joshua Susskind. The slingshot mechanism: An empirical study of adaptive optimizers and the grokking phenomenon. *arXiv preprint arXiv:2206.04817*, 2022.
- Eric Todd, Millicent L Li, Arnab Sen Sharma, Aaron Mueller, Byron C Wallace, and David Bau. Function vectors in large language models. *arXiv preprint arXiv:2310.15213*, 2023.
- Ralph Tollrian and C Drew Harvell. *The ecology and evolution of inducible defenses*. Princeton University Press, 1999.
- L. S. Underwood and Patricia Reynolds. Photoperiod and fur lengths in the arctic fox (Alopex lagopus L.). *International Journal of Biometeorology*, 24(1):39–48, March 1980. ISSN 1432-1254. doi: 10.1007/BF02245540. URL https://doi.org/10.1007/BF02245540.
- Vikrant Varma, Rohin Shah, Zachary Kenton, János Kramár, and Ramana Kumar. Explaining grokking through circuit efficiency. *arXiv preprint arXiv:2309.02390*, 2023.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 2017.
- Johannes Von Oswald, Eyvind Niklasson, Ettore Randazzo, João Sacramento, Alexander Mordvintsev, Andrey Zhmoginov, and Max Vladymyrov. Transformers learn in-context by gradient descent. In *International Conference on Machine Learning*, pp. 35151–35174, 2023.
- Conrad H Waddington. Genetic assimilation of an acquired character. *Evolution*, 7(2):118–126, 1953.
- Mary Jane West-Eberhard. Developmental plasticity and evolution. Oxford University Press, 2003.
- Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. An explanation of in-context learning as implicit Bayesian inference. *arXiv* preprint arXiv:2111.02080, 2021.

A TECHNICAL APPENDICES AND SUPPLEMENTARY MATERIAL

A.1 EXPERIMENTAL DETAILS: PARAMETER SWEEPS AND REPLICATION

To investigate the effects of environmental predictability, we performed dense parameter sweeps for the Sinusoid regression and Omniglot classification tasks. Unless otherwise specified, all results reported in figures are the mean across 3 runs using different random seeds, and plotted error bars represent ± 1 standard error of the mean (SEM).

Sweep 1: Sinusoid Regression

The following grid of parameter values was used:

- $\alpha_{\text{stability}} \in \{0.0, 0.1, 0.3, 0.5, 0.7, 0.8, 0.9, 0.95, 0.99, 0.999, 1.0\}$
- $\sigma_{\text{reliability}} \in \{0.0, 0.005, 0.01, 0.02, 0.05, 0.1, 0.2, 0.3, 0.5, 0.75, 1.0, 1.5\}$

Sweep 2: Omniglot Classification

The following grid was used, where $\alpha_{\text{stability}}$ is determined from $p_{\text{stability}}$ via $\alpha_{\text{stability}} = 2p_{\text{stability}} - 1$, as described in the main text.

- $\alpha_{\text{stability}} \in \{0., 0.1, 0.2, 0.4, 0.6, 0.8, 0.9, 0.95, 0.99, 0.999, 1.\}$
- $p_{\text{reliability}} \in \{0.5, 0.55, 0.6, 0.65, 0.7, 0.75, 0.8, 0.85, 0.9, 0.95, 0.99, 1.\}$

Sweep 3: Omniglot Relative-Cost Experiment

To test the relative-cost hypothesis (Figure 5), environmental stability and cue reliability were held at high values ($\alpha_{\text{stability}} = 0.999$, $p_{\text{reliability}} = 0.95$), while varying the IWL difficulty via the character set size:

• $|\mathcal{C}| \in \{100, 200, 400, 600, 800, 1000, 1200, 1400, 1623\}$

These sweeps comprise $(11 \times 12) + (11 \times 12) + 9 = 273$ unique experimental configurations. Including the 3 random seeds per configuration, this resulted in a total of 819 model training runs.

A.2 COMPUTE RESOURCES

All models were trained on single TPUv3 cores. Each of the 819 training runs took approximately 30 minutes to complete the 50,000 training steps (detailed in Section 3). The total compute usage was approximately 410 TPU-hours.

B USE OF LARGE LANGUAGE MODELS

During the preparation of this manuscript, a large language model was used to assist with editing the prose for clarity and to aid in literature discovery. This latter function was particularly valuable for exploring the evolutionary biology literature that provides the conceptual framework for this study. All scientific claims and the final content of the paper were written and verified by the authors.