UNIVERSITY OF TARTU
Institute of Computer Science
Computer Science Curriculum

Johan Kirikal

Modelling Creativity Using Artificial Neural Networks

Bachelor's Thesis (9 ECTS)

Supervisor:
Jaan Aru, PhD

Tartu 2025

Modelling Creativity Using Artificial Neural Networks

Abstract:

The mental processes and algorithms that lead to creativity are still largely unknown. Psychological theories suggest that creativity arises from a balance between associative memory structure and the executive processes that retrieve and recombine information. In this thesis, I explore whether artificial neural networks can exhibit creative-like behaviour through controlled entropy modulation. Using the Continuous Generative Flow Network (C-GFN), a biologically inspired architecture, I simulate the impact of increased stochasticity on creative output. By varying the standard deviation in the model's inner dynamics, I test the hypothesis that higher internal entropy leads to more creative outputs, as the entropy modulation theory of creativity predicted. As a result, models with elevated entropy not only generated more diverse symbolic representations but also did so more quickly and covered a greater distance in their latent space, emulating the faster-and-further phenomenon observed in human creativity research. These findings provide computational support for the associative and executive theories of creativity and highlight the role of entropy modulation in creative behaviour.

Keywords: Artificial intelligence, modelling, psychology, creativity

CERCS: P176 Artificial intelligence

Loovuse mudeldamine tehisnärvivõrkudega

Lühikokkuvõte:

Loovusega seotud protsesse ja algoritme mõistetakse siiamaani üsna kehvasti. Psühholoogias levinud teooriate kohaselt sünnib loovus assotsiatiivse mälu struktuuri ja sellel toimivate algoritmide tulemusel. Selles bakalaureusetöös uurin, kas tehisnärvivõrkudel põhinevat mudelit saab panna loovamalt käituma, kui selle protsesse mõjutada entroopiat ehk suvalisuse määra muutes, nagu väidab entroopia modulatsiooni teooria. Selle hüpoteesi testimiseks kasutan *Continuous Generative Flow Network* (C-GFN) mudelit, mille arhitektuur on inimajust inspireeritud. Et hüpoteesi uurida, vaatlen suurenenud entroopia mõju mudeli väljundite loomingusesele. Selleks mõjutan katsetes mudeli standardhälvet, mida see kasutab oma sisemistes protsessides otsuste langetamiseks. Katsete tulemusel selgus, et suurema entroopiaga mudelid tootsid mitmekesisemaid väljundeid, need asetsesid mudeli semantilises ruumis algsest sisendist kaugemal ja see kõik toimus kiiremini, kui madala-entroopia mudelites. Selline käitumine viitab *faster-and-further* ehk kiiremini ja kaugemale fenomenile, mida on loova mõtlemisega seostatud. Need tulemused toetavad loovuse assotsiatiivset ja täidesaatvat teooriat ning rõhutavad entroopia dünaamilise muutmise tähtsust loovas käitumises.

Võtmesõnad: Tehisintellekt, mudeldamine, psühholoogia, loovus

CERCS: P176 Tehisintellekt

Table of Contents

1. Introduction

Studying the human brain is one of the most significant challenges in modern science. The complexity of this subject stems from the fact that the brain is a black box whose inner workings and architecture are complex to measure and observe. While it is possible to measure brain activity on different scales, the underlying algorithms and processing language of the mind are still a mystery. Solving these algorithms is a guessing game, and multiple fields try to solve this problem from different angles. Neuroscience is the umbrella term under which different disciplines operate to solve this problem. Some subfields include psychology, which attempts to answer these questions by observing human behaviour; medicine, where the physical body and systems are being observed; biology, where the processes are described on the cellular and molecular level; and so on.

In this thesis, I apply the framework of **computational neuroscience** to study brain function. Computational neuroscience aims to model aspects of human cognition and behaviour through computer-based simulations, including artificial neural networks, neuromorphic computing systems, or algorithmic frameworks.

The purpose of this bachelor's thesis is to offer a modest contribution to the broader goal of understanding the human brain. Specifically, I focus on the phenomenon of **creativity**. To investigate this topic, it is first necessary to establish a working definition of creativity. The Cambridge Dictionary[1] defines creativity as *"the ability to produce or use original and unusual ideas."* Similarly, Wikipedia[2] describes creativity as *"the ability to form novel and valuable ideas or works using one's imagination."* However, these are only two of many definitions: social geographer Peter Meusburger estimated that over a hundred distinct definitions exist in the literature (Meusburger, 2009). In this work, I adopt the definition by Runco and Jaeger (2012) who characterize creativity as **an associative process over memory that involves the generation of new and effective ideas**.

From this foundation, the central hypothesis of the thesis is formulated: **raising entropy in the behaviour of an artificial neural network model increases its creativity**. This hypothesis explores the idea that destabilized systems may exhibit higher creative capabilities, potentially offering insights into the mechanisms underlying creativity in the human brain.

---

[1] https://dictionary.cambridge.org/dictionary/english/creativity
[2] https://en.wikipedia.org/wiki/Creativity

To present the research findings and results, this thesis is organized into four main sections:

1. **Background**, where I review various aspects of creativity and previous related work.
2. **Methods**, where I describe the approach taken in modelling and experimentation.
3. **Results**, where I present the outcomes of the experiments.
4. **Discussion**, where I discuss results, reflect on the limitations of this work and suggest directions for future research.

Appendix I: Glossary provides a comprehensive overview of the field-specific terminology used in this thesis. The results, code, and instructions for reproducibility can be found under Appendix III: Data and Code Availability.

This work was edited using ChatGPT[3] and Grammarly[4].

[3] OpenAI (2025). ChatGPT (version 4o): https://chatgpt.com/
[4] https://app.grammarly.com/

2. Background

Successful modelling of creative behaviour requires a comprehensive understanding of the phenomenon itself—how creativity is defined, its core characteristics, and the tools used to measure it. Additionally, familiarity with prior attempts to model creativity is essential.

2.1 What is Creativity

Associative thinking, the process of connecting and combining different concepts to form ideas, art, or music, has long been considered the driving mechanism behind creativity (Beaty & Kenett, 2023). When applied to the generation of **novel** ideas and combinations, this process manifests as **creativity**.

This thesis is built on the definition that creativity is an associative process over **memory** that involves the generation of **new** and **effective** ideas (Runco & Jaeger, 2012). The level of creativity is a function of the number of associations produced per cue item. Individuals with lower levels of creativity tend to produce fewer but stronger associations, whereas highly creative individuals generate a greater number of weaker associations (Beaty & Kenett, 2023).

Two primary factors that influence creativity are **memory structure** and the retrieval algorithm or **process** that operates over it (Beaty et al., 2014; Benedek et al., 2017).

2.1.1 The Associative Theory

The associative theory of creativity suggests that the memory structure of an individual influences how quickly and effectively creative solutions can be reached. As Mednick (1962, p. 222) states, “The organization of an individual’s associations will influence the probability and speed of attainment of a creative solution.”

A widely used metaphor for this concept is that of a rugged terrain (Hills & Kenett, 2022). In this mental landscape, ideas are positioned on peaks, and the challenge of moving between these peaks accounts for the differences in creativity levels among individuals. Mednick (1962) argued that more creative people have “flat” associative landscapes that are easier to navigate. In contrast, those with lower creativity tend to have “steep” associative landscapes where it is difficult to move between different peaks. Figure 1 illustrates the metaphor between steep and flat mental landscapes.

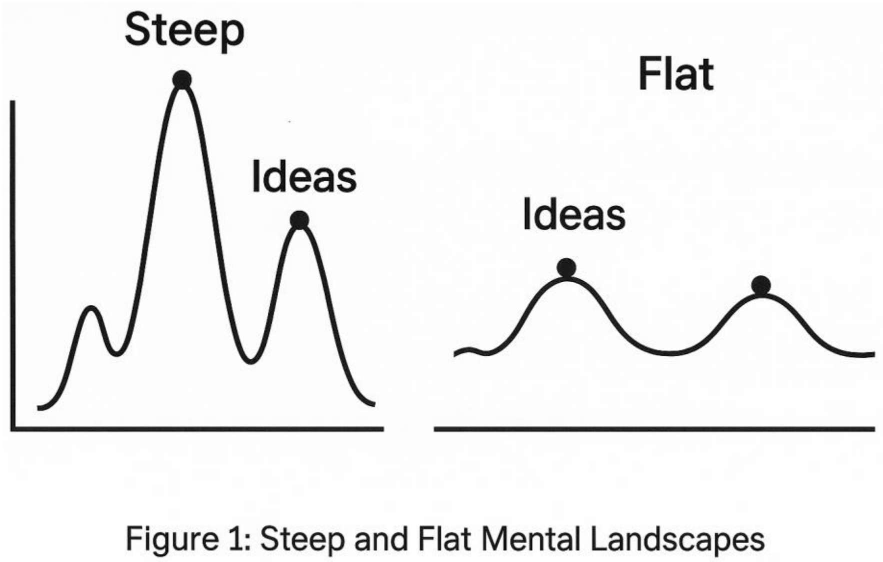

Figure 1. ChatGPT[5] image generation, April 2025. Flat associative landscape compared to a steep associative landscape.

Therefore, individuals with strong associations are limited in how far they can traverse their associative landscape, often resulting in fewer and more stereotypical responses (e.g., *table–chair*). In contrast, a broader and weaker associative network allows for greater cognitive flexibility, enabling access to more distant concepts and producing more varied and creative solutions (e.g., *table–clouds*).

Expanding on this landscape metaphor, recent evidence shows that highly creative individuals form **equally strong links** to both **common and unusual** concepts**,** enabling them to reach beyond conventional pairings and access more remote ideas (Rossmann & Fink, 2010). Regarding the terrain analogy, their landscape is flatter and **more navigable in every direction**, which affords greater cognitive flexibility and, ultimately, more innovative outcomes. Because the probabilities of reaching different associations are distributed more evenly, their associations are less predictable—mathematically, they exhibit **higher entropy** (see Section 2.2) (Hills & Kenett, 2025).

Recent advances in computational modelling—particularly **distributional semantic models**—have enabled researchers to simulate aspects of semantic memory (Jackson et al., 2021; Mandera et al., 2017). These models allow for clustering concepts and quantifying semantic distances between them, which is valuable for traditional psychological studies and newer

[5] OpenAI (2025). ChatGPT (version o4-mini): https://chatgpt.com/

computational approaches. Researchers can better understand the associative mechanisms underlying creative thought by modelling how closely concepts are related in meaning. An illustrative example of noun clustering based on semantic similarity is presented in Figure 2.

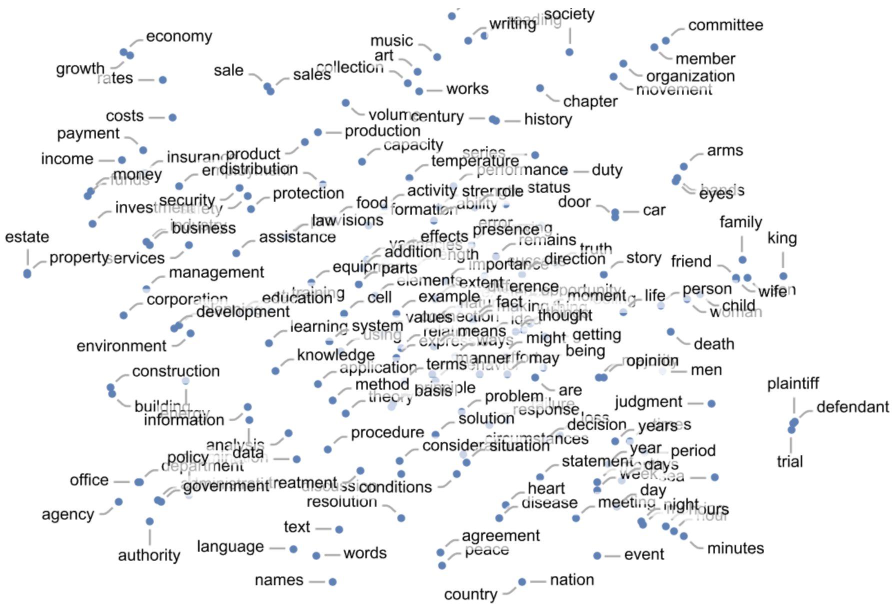

Figure 2. A semantic feature space where semantically similar words are placed nearby. Projected to 2D space (Wolfram, 2023).

Therefore, semantic memory structure forms the foundation for creative thinking, providing the framework upon which new combinations are formed through associations (Abraham & Bubic, 2015). The second key factor contributing to creativity, discussed next, is the retrieval process—or the algorithm—that operates over memory (Beaty et al., 2014).

2.1.2 The Executive Theory

Searching and retrieving information from memory has been conceptualized as an internal form of foraging—similar to how animals search for food in the wild—where individuals navigate through the landscape of mental concepts (exploration) and focus on particular categories (exploitation), depending on the demands of the task (Todd & Hills, 2020). The executive theory of creativity builds on this metaphor, proposing that creative thought emerges from top-down control processes that guide and constrain which memories are accessed and how they are combined (Beaty et al., 2014; Benedek et al., 2023; Chrysikou, 2019; Silvia, 2015).

In modern psychology, multiple creative "algorithms" have been defined upon which the brain acts. Each retrieval process serves a specific purpose, depending on the situation and need.

Creative processes fall under **free-association** tasks or **goal-directed** tasks (Beaty & Kenett, 2023).

Free association is mainly based on stochastic processes and memory structure and lacks any control mechanism, as the name suggests (Richie et al., 2023). Free association involves tasks where an individual is presented with a word, and they must present the first (random) word that comes to mind. An example would be “dog–leash” and “table–chair”. These tasks do not have any restrictions and are useful in measuring a person’s creativity level. In this case, highly creative associations would be “dog–church” and “table–moon” so that the response to a stimulus is farther away from the stimulus itself in the **semantic space** of a person. Therefore, an individual’s creativity in free-association tasks is measured by the distance of produced associations in a representational space (Gray et al., 2019).

In contrast, goal-directed creativity uses specific constraints to outline the goal and uses different strategies to reach this goal, such as:

- Switching - alternating between semantic categories (e.g., animals; “birds, snakes, fish”) (Troyer et al., 1998)
- Clustering - grouping related concepts together (e.g., small animals; “mouse, rabbit, hamster”) (Troyer et al., 1998)
- Convergent creativity - a targeted search for a single destination (Malaie et al., 2024)
- Divergent creativity - a wide exploration of space to find many candidate ideas (Malaie et al., 2024)
- Inhibition - the ability to suppress dominant or automatic responses (like repeating sequences) (Benedek et al., 2012)
- Entropy modulation - heightened neural variability that may increase exploratory behaviours (Hills & Kenett, 2025)

These algorithms are used by a **cognitive control** mechanism, which decides how much and when to use any of those strategies (Avery & Krichmar, 2017; Pezzulo et al., 2015).

2.2 Entropy in Neural Systems

Noise, irregularity, and entropy describe the same underlying concept: random variation in neural activity that injects stochasticity into information processing. Recent research suggests that such noise is not merely background clutter but an important driver of brain dynamics (Rolls & Deco, 2010; Touboul et al., 2012). Crucially, adding an optimal amount of noise tunes association retrieval distance so that more distant associates have a greater chance of being retrieved, thereby fostering originality (Heinen & Johnson, 2018; Rastelli et al., 2022) and "more noise" does not mean uniformly random retrieval.

In modelling studies, this stochastic influence is typically parameterized by **σ**, which can reflect variability generated within single neurons (ion-channel noise, synaptic release variability) or across populations through their connectivity patterns (Deco et al., 2011). The amplitude of σ appears to be a determining factor in whether an individual shows higher or lower creative capability: everyone can adjust their noise level, yet highly creative people seem able to modulate σ over a broader range or to a larger amplitude (Hills & Kenett, 2025).

2.2.1 The Entropy Modulation Theory of Creativity

The entropy modulation theory is based on the findings that cognitive processes are fundamentally **probabilistic**. Probabilistic sampling is a core component of many models of cognitive function (Anderson et al., 1997; Bhatia & Loomes, 2017; Faisal et al., 2008; Hills et al., 2012; Luce, 2005; Zhu et al., 2020), which model human variability remarkably well. These models suggest that variance, or entropy, plays a key role in the degree of exploration in individuals (e.g., Yechiam et al., 2005), which brings us to the topic of entropy modulation.

The theory of entropy modulation is based on diverse behavioural and neural evidence indicating that both humans and animals can strategically enhance the randomness and unpredictability of their responses (Hills, 2019; Wilson et al., 2014, 2021). For instance, they may increase exploratory behaviours in uncertain environments through heightened neural variability. This kind of strategic exploratory behaviour is facilitated by specific prefrontal regions, possibly by disrupting stable neural attractors based on reward-related cues (Cools et al., 2007; Daw et al., 2006; Rolls et al., 2024). Therefore, cognitive control can adaptively tune the regularity of background neural firing in different brain areas to tune behavioural flexibility (Aston-Jones & Cohen, 2005; Avery & Krichmar, 2017; Cools & D'Esposito, 2011; Doya, 2008; Eppinger et al., 2021; Gilbertson & Steele, 2021; Lin & Vartanian, 2018; Pezzulo et al., 2015). Additionally, variations in this modulatory capability are linked to individual differences in executive functioning (Omidvarnia et al., 2021; Shine et al., 2019) and creativity (Boot et al., 2017).

Ultimately, according to the entropy modulation theory, the link between entropy control and creativity resides in the altering of the memory retrieval process. Higher creative individuals are more creative because they produce more **variance in the activation of associations**, which means that associations further away from the cue are more likely to get activated. This increase in activation is in addition to any other differences in activation thresholds or structural differences of memory (Benedek et al., 2023).

2.2.2 Faster and Further

Creativity research consistently highlights two main findings we aim to explore: compared to individuals with lower creativity, those with higher creativity tend to (a) generate associations more quickly and (b) produce less typical, more original (creative) associations. This pattern appears widely across studies, indicating a strong link between creativity and total production amount (Jung et al., 2015; Ovando-Tellez et al., 2022; Simonton, 1997; Weiss et al., 2024). Hills and Kenett (2025) refer to the pattern as **faster-and-further**.

As discussed earlier, differences in creative performance have often been interpreted in two ways: one supports an **associative theory**, which is based on differences in **cognitive representations**, while the other aligns with an executive theory, suggesting that more creative individuals employ **cognitive control** to approach their cognitive representations in a distinct manner (see Section 2.1). Although the role of semantic memory structure in creativity has been the subject of many previous computational models (Benedek et al., 2023; Kenett, 2025; Kenett & Faust, 2019; Volle, 2018), the executive processes—specifically, the **executive control of entropy** in memory retrieval—underlying the faster-and-further phenomenon had not been formally modelled until the work of Hills and Kenett (2025).

To address this gap, Hills and Kenett (2025) proposed the entropy modulation theory, integrating both structural and process-based perspectives on creativity. They hypothesized that the brain modulates the level of entropy in neural firing patterns and that this capacity is linked to individual differences in creativity. Their central aim was to mechanistically demonstrate how executive control processes could account for the faster-and-further pattern.

They continued to construct an experiment to validate their hypothesis. For their experiment, they created a model, where the process was formalized as follows. When presenting the model a retrieval cue (e.g., a concept), it starts sampling associates from its memory. Each associate is sampled with a probability proportional to their similarity to the cue. Sampling an associate awards the sampled concept an increment to its activation. The sampling and activation process is continued until one associate exceeds a threshold and it gets retrieved and reported. This kind of process is a version of evidence accumulation or race model, where various potential targets compete to reach a threshold level of activation (Busemeyer et al., 2019; Heathcote & Matzke, 2022). In this framework, response time is measured by the number of iterations required for an associate to reach the threshold (Kvam et al., 2015).

Next, Hills and Kenett (2025) ran simulations with high- and low-noise (i.e., entropy) variants of the model. The high-entropy model retrieved more distant associations and required fewer activation steps, thereby demonstrating both speed and originality in retrieval. These findings support the idea that entropy modulation can produce the faster-and-further phenomenon.

While Hills and Kenett (2025) employed a framework based on evidence accumulation or race models, they did not test the entropy modulation theory using an artificial neural network that simulates human cognition. My thesis addresses this gap by validating the entropy modulation theory within a model grounded in artificial neural networks.

2.3 A Brain-Inspired Model Architecture

To test the hypothesis that creativity can be increased through entropy modulation, it is essential to ground the investigation in a neural architecture capable of reflecting the key characteristics of human cognition. While entropy modulation—adding and tuning the noise in a system—can, in theory, be applied to a wide range of artificial neural networks (ANNs), most common architectures lack mechanisms that adequately capture the structure, dynamics, and probabilistic reasoning observed in the human brain.

2.3.1 The Core Characteristics of Cognition

Human cognition is marked by a **hybrid of sub-symbolic and symbolic processing** (McClelland et al., 2010). Our brains operate through vast networks of neurons interacting via continuous-valued signals—this is sub-symbolic computation. Yet, we routinely engage in structured, symbolic reasoning: combining concepts, manipulating language, solving equations. Classical digital computers are symbolic by design (all computations reduce to manipulations of discrete binary states), while ANNs have typically belonged to the sub-symbolic paradigm. Bridging these two levels—how discrete symbols can emerge from continuous dynamics—has been a core challenge in cognitive modelling (Fodor & Pylyshyn, 1988; Rumelhart et al., 1986).

Although recent successes with large language models (LLMs) have shown that ANNs can, in practice, exhibit symbol-like behaviours (e.g., language generation), these systems often still struggle with tasks requiring **systematic symbolic reasoning**—such as mathematical deduction or multi-step logical inference (Achiam et al., 2023) (although the cutting-edge reasoning models are getting better, see LLM Leaderboard[6]). This highlights an ongoing need for models that can **naturally give rise to symbolic structure** from neural processes, rather than relying on it being explicitly programmed in.

One promising direction for achieving this synthesis of sub-symbolic and symbolic processing is through **attractor dynamics**. Attractors offer a biologically plausible way for continuous systems (like neurons or neural networks) to settle into stable, discrete states, which computationally might correspond to symbols. In a dynamical system, an attractor is a stable configuration that the system evolves toward, even when starting from different initial conditions or under small perturbations. This mechanism aligns well with neuroscientific findings: for instance, when a monkey has to decide between two options in a binary choice task, neurons in its cortex evolve into one of two stable firing patterns—each pattern corresponding to one of the decisions (Inagaki et al., 2019; Shadlen & Newsome, 2001).

In parallel, human reasoning exhibits a **chain-of-thought** structure (L. Hu et al., 2024). We rarely arrive at conclusions in a single step; rather, ideas evolve over time through internal iteration, hypothesis testing, and symbolic recombination. Traditional feedforward neural networks do not capture this process well, as they compute outputs in a single pass. A more realistic model would involve recurrent or iterative dynamics, where intermediate representations evolve across multiple time steps—mirroring the process of thinking or reasoning (Wang, 2008).

Another essential aspect of cognition is its **probabilistic nature** (Knill & Pouget, 2004; Tenenbaum et al., 2011). The brain does not operate deterministically—it constantly makes predictions, updates beliefs based on uncertainty and evaluates multiple possible interpretations. In contrast, many simple ANNs are fully deterministic (Gawlikowski et al., 2023), lacking a mechanism to model uncertainty or variability. A suitable model for studying creativity should therefore support **Bayesian inference** or similar probabilistic reasoning.

While creativity has been explored in LLMs (Zhao et al., 2025), there is ongoing debate about the extent to which transformer architectures reflect biologically plausible neural behaviour (Kozachkov et al., 2023). In the context of LLMs, entropy modulation typically refers to adjusting the "temperature" parameter. This key hyperparameter influences the randomness or unpredictability of the model's responses by altering the probability distribution over output

[6] https://www.vellum.ai/llm-leaderboard

logits. However, Zhao and colleagues' (2025) experiments found that varying temperature did not have a noticeable impact on creative performance, which appeared largely unpredictable. As such, LLMs may not provide an ideal foundation for studying entropy modulation.

2.3.2 The Architecture for Modelling Creativity

The model introduced by Nam and colleagues (2023), a **Continuous Generative Flow Network (C-GFN)**, provides an ideal foundation for this thesis. It is explicitly designed to combine:

- Attractor dynamics - to support discrete symbolic states emerging from continuous latent space.
- Probabilistic inference - the model samples from posterior distributions over symbolic representations.
- A chain-of-thought process - via iterative latent-state transitions.
- And the seamless integration of symbolic and sub-symbolic processing.

The C-GFN model is a **neural stochastic dynamical system** trained to form stable attractor states in its latent space, each corresponding to a symbolic representation (e.g., a token sequence). Crucially, these attractors and the symbols associated with them are **not pre-programmed**—they emerge naturally from unsupervised learning under the pressure to generate meaningful, compact, and compositional outputs. This makes the model a biologically inspired, cognitively plausible platform for exploring the relationship between **stochasticity (entropy)** and **creative behaviour**.

The core mechanism of the C-GFN network, the **Generative Flow Network** (GFlowNet) was developed by Bengio and colleagues (2021). The GFlowNet is a generative model or stochastic policy that generates an output x through a sequence of steps. The intermediate steps can be thought of as internal actions or thoughts that result in an idea. The model samples x with a probability proportional to a reward function R(x). Therefore, the range of different outputs x generated will represent a distribution of R(x). In contrast, typical generative models usually generate only the most likely outcome.

A good example description of the sampling process of the GFlowNet is of building a Lego ship, where the sequence of steps resembles the sequence of adding different pieces together. One specific state during the process of ship building can be a result of many different paths to that state, where there may be slight variations in the order of piecing together the Lego bricks.

The reward of the resulting ship describes how well the output resembles a ship. Since there can theoretically exist an innumerable amount of different Lego ships, the reward function acts like a criterion that a good ship must satisfy.

Therefore, the original GFlowNet operates with discrete steps. This model was further developed (Lahlou et al., 2023) to enable the model to work in continuous spaces and ultimately learn attractor basins.

The third important part of this model is the Expectation-Maximization (EM) algorithm (Dempster et al., 1977), which was developed to estimate the values of latent variables that are involved in observation generation. This EM algorithm was applied to the GFlowNet architecture (E. J. Hu et al., 2023) and the basic intuition behind the EM-loop looks like this:

E-Step:

The model asks: "Given each input, what are the likely symbolic sequences (and latent attractor states) that could describe it?"—this is a posterior inference over sequences, which the GFlowNet learns to sample from.

M-Step:

The model asks: "Given these sampled descriptions for each input, how should I update the representation of those descriptions (and the encoding of inputs) so that they better match each other?"—this updates the parameters of the generative model—encoding and decoding of items between the latent space. Therefore, this step updates the locations of attractor states, defining what each symbol means.

By iterating these steps, the system gradually improves both its inference policy (E-step) and its generative parameters (M-step), so that they converge to consistent symbolic encodings.

The C-GFN model consists of multiple neural modules working together, outlined in Table 1 below:

Table 1. Modules in the C-GFN model. The modules optimized in the E-step and M-step are parameterized by θ and ϕ, respectively.

Module	Description	Expression
Input Encoder (M-Step)	Encodes an object x in the input space as an information-rich state z_0 in the latent representation space.	$P_\phi(z_0 \mid x)$
Input Decoder (M-Step)	Decodes a state z in the latent representation space to an object x in the input space.	$P_\phi(x \mid z)$
Forward Dynamics (E-Step)	Defines the forward stochastic transition dynamics from z_t to z_{t+1} in the latent representation space.	$P_\theta(z_{t+1} \mid z_t)$
Backward Dynamics (E-Step)	Defines the backward stochastic transition dynamics from z_{t+1} to z_t in the latent representation space. The backward dynamics are dependent on time t and input x.	$P_\theta(z_t \mid z_{t+1}, t+1, x)$
Sentence Encoder (M-Step)	Encodes a token sequence s as a state $\widehat{z_s}$ in the latent representation space.	$\widehat{z_s} \leftarrow e_\phi(s)$
Sentence Decoder (M-Step)	Decodes a state z in the latent representation space to a token sequence s autoregressively by sampling a token s_i conditioned on z and the previous tokens $s_{:i}$. We refer to the repeated application of the decoder to sample a full sequence $s \sim P_\theta(s \mid z)$ as the discretizer.	$P_\theta(s_i \mid s_{:i}, z)$
Flow Correction (E-Step)	Measures the flow for a state z_t using a correction quantity g.	$g \leftarrow g_\theta(z_t, t, x)$

The workflow of the C-GFN model consists of three main steps:

1. **Input Encoder** takes an input x and projects it onto a multi-dimensional latent space Z, initializing the trajectory at z_0.
2. The **Forward Dynamics** module constructs a trajectory $z_0 \rightarrow z_T = z_0, z_1, \ldots, z_T$ in the latent space Z over T discrete timesteps. At each intermediate timestep t, the model produces a **mean** and **standard deviation** upon which the next step is chosen using a normal distribution.
3. The **Sentence Decoder** samples a sentence s based on the terminal latent state z_T. The sentence s is a token sequence and the sampling process adds one token at a time to the generated sentence until a termination token is sampled.

For every "forward-action" module there exists a module for performing the reverse action, enabling training through bidirectional consistency. This bidirectional consistency is supported by a measure called "flow", unique to the Generative Flow Networks, which estimates the expected reward from a state. This measure indicates the probabilities of transitions to the next state. The training process aims to match the transition probabilities between the forward sampling and backward sampling processes.

Figure 3 demonstrates an example of the whole sampling process, starting from a stimulus and ending with a sampled sentence s.

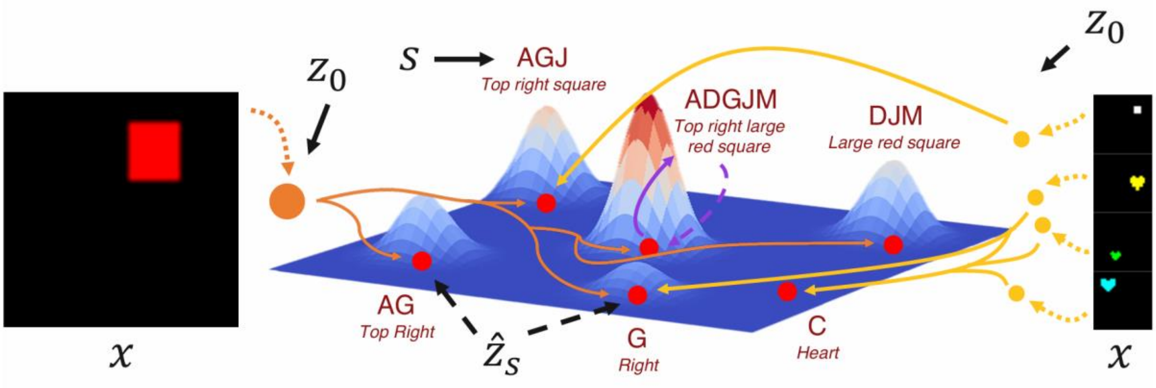

Figure 3. An illustration of the sentence sampling process. First, the model is presented with a stimulus, such as a picture with a red square. Next, the Input Encoder module encodes this stimulus onto the latent space Z to location z_0. Now starts the trajectory sampling process from z_0 to z_T. The terminal z_T is located at an attractor basin, which resembles the best description for the original input. Lastly, the Sentence Decoder module samples tokens based on the attractor (Nam et al., 2023).

Ultimately, the model learns to sample token sequences, which are composed into a kind of sentence. This symbolic language is learned by the model itself through unsupervised learning. In addition, these sentences are generated to resemble the structural and compositional features of a given input. The pressure for the model to learn how to describe the input using symbols effectively comes from the pressure that the model should be able to regenerate the original stimulus from the inferred token sequence. In addition, the reward function (see Appendix II: Equations) can be supplied with a prior $P(s)$ that favours shorter symbol sequences and therefore pressures the model to supply each token with an actual meaning. Therefore, the pressure for the model to learn a structural and compositional symbolic language comes from the fact that the sampled symbolic sequences (low informational fidelity) should be sufficient for regenerating the original input x (high informational fidelity).

In modelling creativity within this framework, several abstract parallels can be drawn with human cognition. The latent space Z, where the model generates its trajectories, can be compared to a person's semantic space. The sentences produced during the sampling process depend on the structure of this space, specifically, the locations of learned attractor basins. Furthermore, the forward-stepping policy reflects an abstract form of executive control in human brains. As such, the model encapsulates both associative and executive theories of creativity within its operations.

3. Methods

The goal of this thesis is to evaluate whether increasing entropy in an artificial neural network enhances the creativity of its outputs. As defined in the previous chapter, creativity is a process by which an entity produces ideas that are both original and useful. Additionally, creative thinking is characterized by the faster-and-further phenomenon—creative individuals generate more distant associations more rapidly compared to less creative counterparts. To computationally test this hypothesis, I selected the C-GFN model (see Section 2.3.2) and modulated its internal entropy, specifically adjusting the standard deviation within the forward-stepping module. An illustration of this process is depicted in Figure 4.

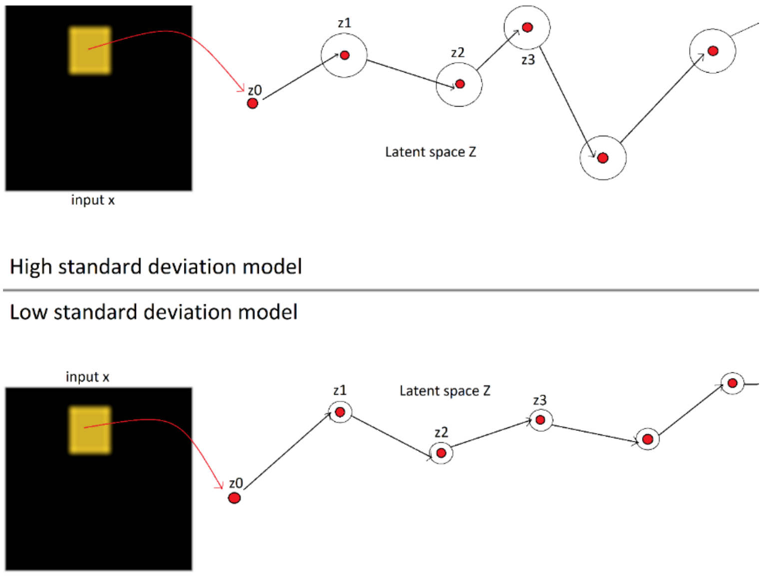

Figure 4. Example comparison of different C-GFN model variants under high- and low-entropy settings. An input x is first processed by the Input Encoder, projecting it into the latent space Z at location z_0. From z_0, the Forward Dynamics module determines the next step z_1 using a predicted mean and standard deviation. Both high- and low-entropy models produce similar mean predictions (shown as red circles), but the high-entropy model samples from a broader distribution, illustrated by the larger black circle around the mean.

The decision to use the forward-stepping module for experimenting with different entropy settings is based on its role as the primary driver of the model's exploratory behaviour and its utility in simulating search strategies.

3.1 Model Setup

The source code for the C-GFN model, available on GitHub[7], is written in Python[8] and utilizes the PyTorch[9] framework for neural network architecture. Due to computational constraints on my personal resources, I trained the model using the University of Tartu's high-performance computing (HPC) Rocket[10] server. Preliminary tests confirmed the anticipated result that GPU training significantly outperformed CPU training due to superior parallel processing capabilities (Buber & Diri, 2018).

For model configuration, I employed the existing configuration file "diffusion_gfn_dynamics_em.yaml" from the GitHub repository, designed specifically to utilize the full capabilities of the C-GFN model by integrating all available modules into the training process.

3.2 Datasets

Nam and colleagues (2023) employed two structured datasets (HBV and dSprites, detailed in Sections 3.2.1 and 3.2.2) characterized by either hierarchical or compositional structures. Figure 5 illustrates these datasets.

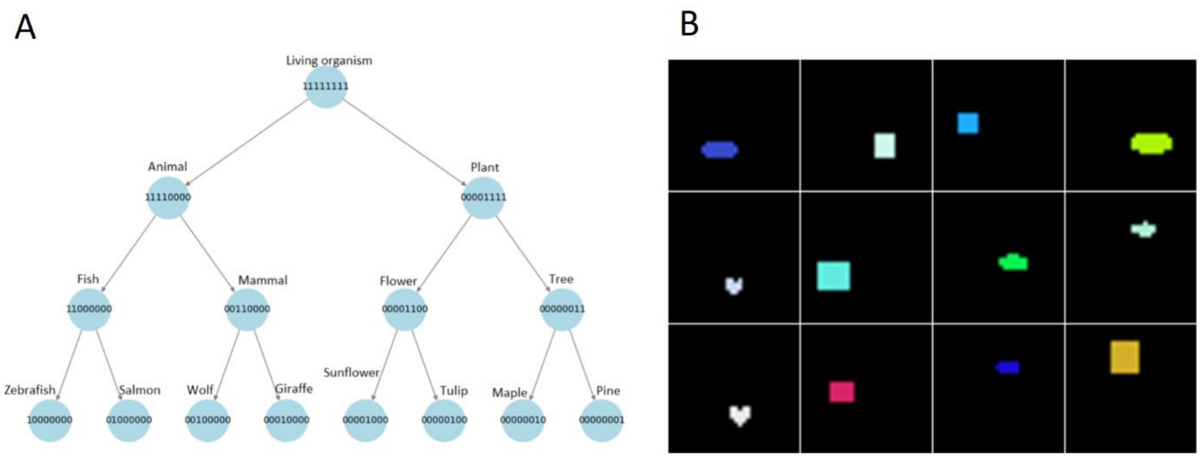

Figure 5. (A) A simplified representation of the HBV dataset with 8-bit vectors and a depth-3 hierarchy tree. (B) Examples from the dSprites dataset.

---

7 https://github.com/andrewnam/gfn_attractors
8 https://www.python.org/
9 https://pytorch.org/
10 https://hpc.ut.ee/services/HPC-services/Rocket

Structured datasets are critical, as symbolic sequences require an underlying ruleset to meaningfully describe complex inputs. Without structural underpinning, symbolic sequences lose fidelity, failing to adequately represent complex information.

3.2.1 Hierarchical Binary Vectors (HBV)

The original purpose of the C-GFN model was to evaluate how well their model can learn and represent hierarchical structures. For that, Nam and colleagues (2023) used a dataset modelled after real-world classification systems—for instance, first distinguishing between broad groups like animals and plants, then further dividing them into subcategories such as birds, fish, trees, and flowers (see Figure 5A). This hierarchy is abstracted using N-bit binary vectors that simulate an idealized phylogenetic tree (McClelland & Rogers, 2003; Saxe et al., 2019), structured by depth (D), where each node corresponds to a class prototype. The tree begins at a root node with all bits set to 1. At each level, a node splits into two children, each inheriting non-overlapping halves (left and right) of the parent's 1s. This process continues recursively until the tree reaches the specified depth D, with leaf nodes still retaining 1s.

I used 128-bit vectors and a tree with a depth 6 in my experiments. To avoid the model simply memorizing specific branch identifiers (e.g., using "A, B, C, D" instead of combinations like "AA, AB, BA, BB" for second-generation branches), I constrained the model with 12 tokens arranged into 6 exclusive pairs—only one token from each pair can appear in a given sequence. For instance, a sequence may include either token "A" or "B", or neither, but not both. This setup results in 6**3 = 216 possible token sequences, each up to a maximum length of 6.

3.2.2 dSprites

To test the model with a more complex input modality, I used a modified version of the dSprites (see Figure 5B) dataset from Higgins and colleagues (2017), which contains images of simple shapes that vary in characteristics like size and position. Nam and colleagues (2023) developed the modified version of this dataset. They extended the dataset by adding colour, assigning each shape one of seven distinct RGB colour combinations with minor variations—excluding black, which was reserved for the background. Additionally, they adjusted the placement of the shapes to lie on a 4x4 grid, introducing slight random shifts to better mimic natural variability. They removed the original orientation feature, as it can be ambiguous for symmetrical shapes, and instead focused on four key attributes: shape, colour, size, and position. To provide the model with enough flexibility to represent overlapping semantic meanings, I allowed it to use up to 32 distinct tokens, with each input sequence containing a maximum of 7 tokens.

3.3 Developing Visualization Tools

Effective visualization was critical for validating model performance and outcomes. While the original code by Nam and colleagues (2023) provided basic visualization functions, additional intuitive tools were developed:

- **HBV graph visualization** - the original code had a method for constructing the hierarchical trees into a variable. On top of that, I modified the graph node labels and created functions for visualizing the graphs.
- **Trajectory visualization** - the original code used a method for constructing a GIF of the different time steps in the latent space Z during the trajectory-building phase, which uses the forward-stepping module. This GIF only showed a dot for a specific timestep, but the trajectory leading up to that moment was lost. I left the dots from each timestep on the plot and connected each dot with a line so that it appears to form a trajectory up to the present moment.
- **Metrics visualization** - after training a C-GFN model, the program creates a metrics .csv file, which holds information about the model's performance during each training step and epoch. To understand how the model's performance improved over time for each of the metrics, I created a separate Python[11] program to visualize these results.

The plots shown in Figure 7, Figure 8, and Figure 9 were generated with assistance from ChatGPT[12] (see Appendix IV: ChatGPT Prompts)

3.4 Determining the Optimal Standard Deviation

The C-GFN model generates internal trajectories within multidimensional space, transitioning between states via sampling from a normal distribution characterized by a trained mean and standard deviation (σ). These parameters are predicted by a simple 3-layer multilayer perceptron (MLP)—the forward-stepping module—taking the current state as input and outputting distribution parameters. Increasing entropy involves amplifying the σ parameter.

[11] https://www.python.org/

[12] OpenAI (2025). ChatGPT (version o4-mini-high): https://chatgpt.com/

Two approaches were considered:

1. Manually select a fixed standard deviation at each step while allowing the forward-stepping module to predict the mean. For example, one model version uses a σ of 0.05 at each forward- and backward-step while an enhanced entropy model version uses a σ of 0.10 at each step.
2. Use a separate, pre-trained forward-stepping MLP with frozen weights to predict the standard deviation at each timestep. Therefore, the σ remains within a certain fixed boundary but varies a little depending on the situation and is not influenced by further training.

Method (1) is simpler since the standard deviation can be hard-coded into a variable. The draw-back of this strategy is that the optimal standard deviation might be slightly different for every step depending on the previous state. This difference could only be reflected using a neural network for σ predictions.

Method (2) is more dynamic in predicting a σ according to the situation. However, this strategy might not be necessary for validating my hypothesis, since optimal model performance is not mandatory for observing differences in creative-like behaviour.

Regardless of the chosen method, I needed to find a reference σ where the C-GFN model performs optimally. For both of the tasks (HBV and dSprites), I trained the C-GFN model for 500 epochs so that the model performance plateaus to a stabilized level.

Next, I ran this model to print out the σ that the forward-stepping module predicts at each timestep for later usage in method (1). In parallel, I saved the forward-stepping (and backward-stepping) module into a PyTorch[13] model (.pt) file for later usage, which was important for conducting my experiments according to the method (2).

3.5 Training the Model on Different Standard Deviations

The central research question of this thesis—whether the level of creativity can be modulated by regulating the amount of entropy in a neural system—was addressed through the following experiments.

[13] https://pytorch.org/

1. For method (1), using a fixed (hard-coded) standard deviation, I modified the configuration file (*diffusion_gfn_dynamics_em.yaml*) to incorporate the optimal standard deviations identified in Section 3.4. These values were 0.055 for the forward- and backward-stepping models in the HBV task, and 0.028 for the dSprites task. To explore the effects of varying entropy levels, I created three variations of the configuration file, each differing by a standard deviation multiplier of 0.5, 1, and 2, respectively.
2. For the second approach, I introduced additional *fixed_mlp* model variables within the codebase for both the forward- and backward-stepping models. I then implemented functionality to load pretrained models for these dynamics and modified the C-GFN workflow to utilize the *fixed_mlp* model for standard deviation predictions. To ensure that further training would not alter the pretrained model, I froze its weights during training. As with the first method, I prepared three configuration files, each corresponding to a standard deviation multiplier of 0.5, 1, or 2.

Both C-GFN model variants—the hard-coded standard deviation model and the learned standard deviation model—were trained for 400 epochs under each multiplier setting. Consequently, for both the HBV and dSprites tasks, a total of 12 distinct model configurations were trained to validate the central hypothesis.

4. Results

To demonstrate that entropy modulation is associated with creative-like behaviour, three criteria from Section 2.2.2 must be met:

1. The high-entropy model moves further in the latent space.
2. The high-entropy model moves faster in the latent space.
3. The high-entropy model produces more associations.

An initial overview of model behaviour under different standard deviation conditions is presented in Figure 6, illustrating intuitively how standard deviation affects trajectory paths and learning of attractor points.

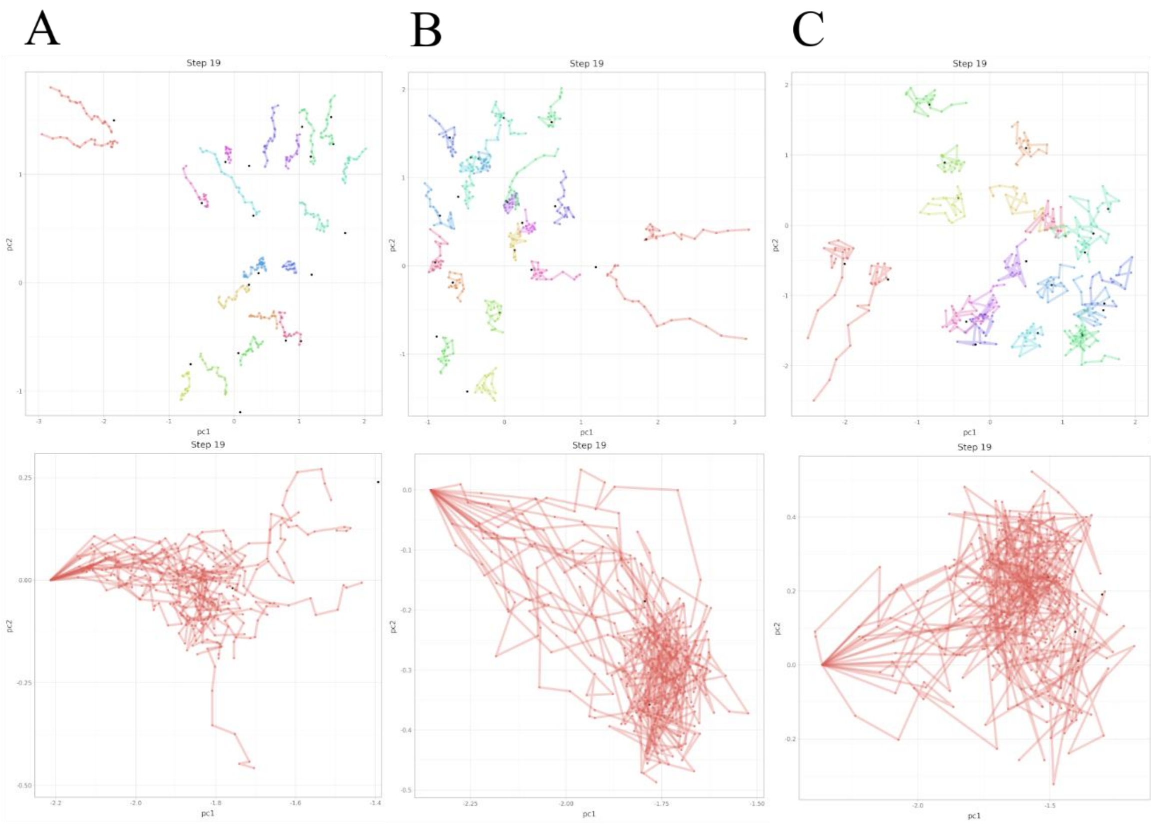

Figure 6. Model trajectories in latent space Z, projected onto a 2D space. Columns A, B, and C correspond to half, equal, and double standard deviation models, respectively. The first row shows trajectories originating from various starting points, while the trajectories in the second row originate from a single starting point. Colours indicate the input x, and black dots represent trained attractor locations.

Figure 6 reveals that models with lower standard deviation learn attractors closer to the original stimulus, while higher standard deviation models reach attractors in fewer steps and subsequently circle these attractors. The trajectories are projected onto the 2D space using principal component analysis.

4.1 Distance Metric (Further)

The "further" component was quantified using the sum of Euclidean distances between consecutive points along each trajectory. Results are depicted in Figure 7.

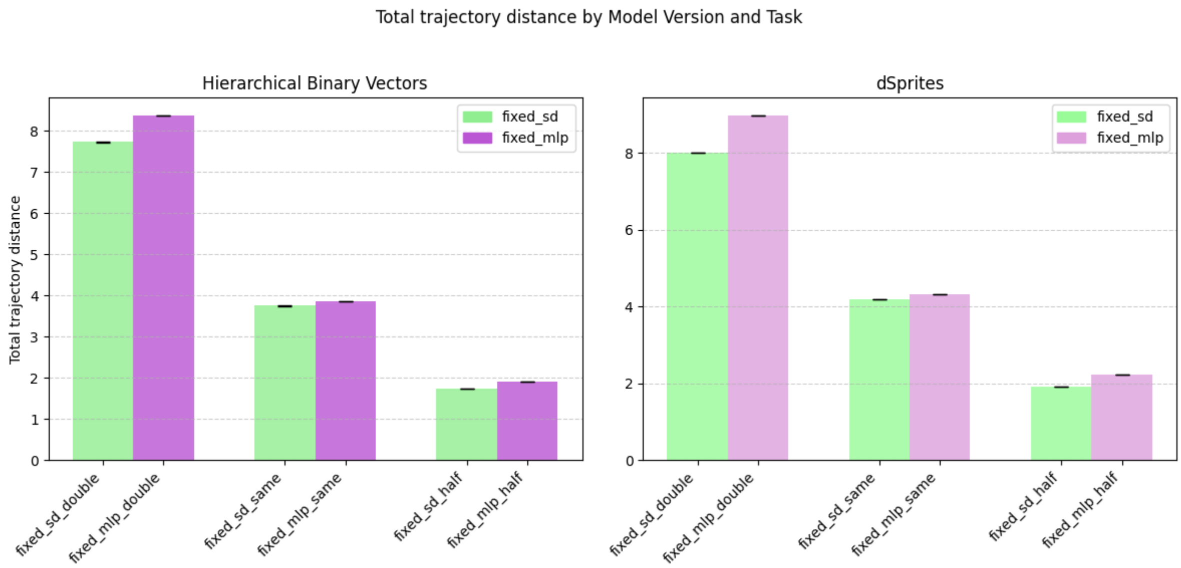

Figure 7. Total trajectory distances for three standard-deviation conditions (×0.5, ×1, ×2). Bars denote mean ± 95 % CI across 10 000 trajectories per model; colours distinguish training methods (1) fixed-sd and (2) fixed-mlp.

These plots show a consistent decrease in total trajectory distance with decreasing standard deviation settings. In addition, we can see that method (2) results in slightly longer trajectory distances.

4.2 Speed Metric (Faster)

The second criterion—speed—is measured by the number of steps required to reach an attractor. Given that trajectories are consistently 20 steps, determining the step at which models begin orbiting an attractor is crucial. Visually, Figure 6 suggests that higher standard deviation models initiate orbiting sooner. However, visual assessment alone is insufficient. Therefore, the speed was calculated using a relative metric: the ratio of total trajectory length to the distance between the attractor and starting point. This is formalized in Equation 1:

$$Relative\ speed = \frac{total\ trajectory\ length}{attractor\ distance} \tag{1}$$

This approach assumes trajectories are relatively direct paths toward attractors and effectively reflects the leftover trajectory orbiting an attractor. Results are presented in Figure 8.

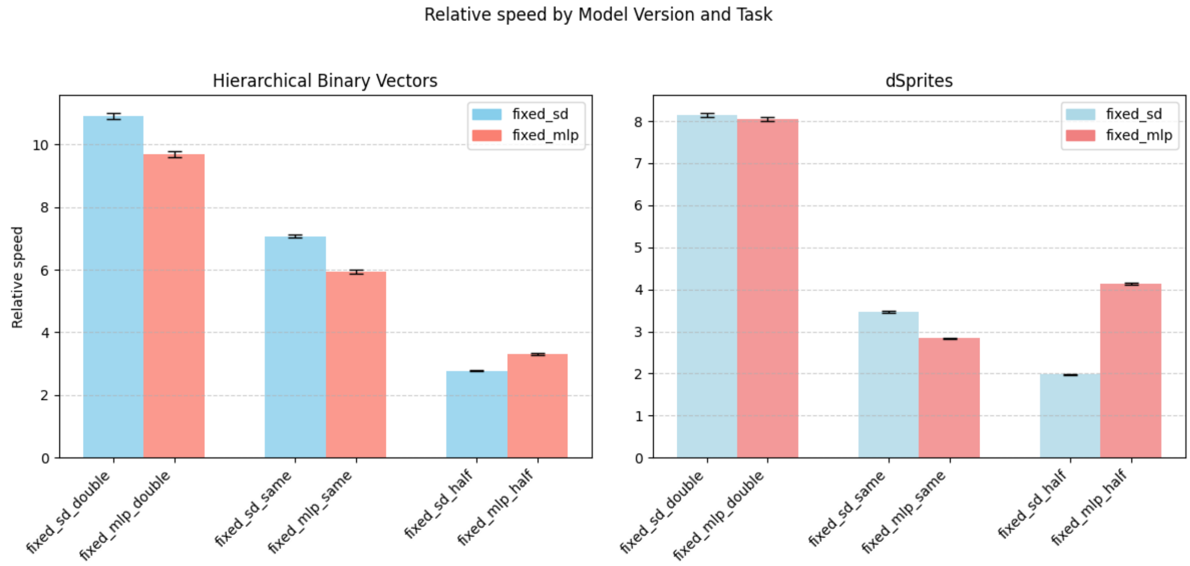

Figure 8. Relative speed to reach attractors for three standard-deviation conditions (x0.5, x1, x2). Bars denote mean ± 95 % CI across 10 000 trajectories per model; colours distinguish training methods (1) fixed-sd and (2) fixed-mlp.

The bars display a pattern similar to that in Figure 7, with one exception: the "fixed_mlp_half" model shows a higher speed than its x1 standard deviation counterpart. All other models exhibit a consistent decrease in speed as the standard deviation setting is reduced.

4.3 Produced Associations

Creativity is frequently measured by the quantity of associations produced per cue item (Beaty & Kenett, 2023). In the context of the C-GFN model, creative-like behaviour can be quantified by the variety of unique descriptions generated for each input. Figure 9 shows median numbers of distinct token sequences from a single input.

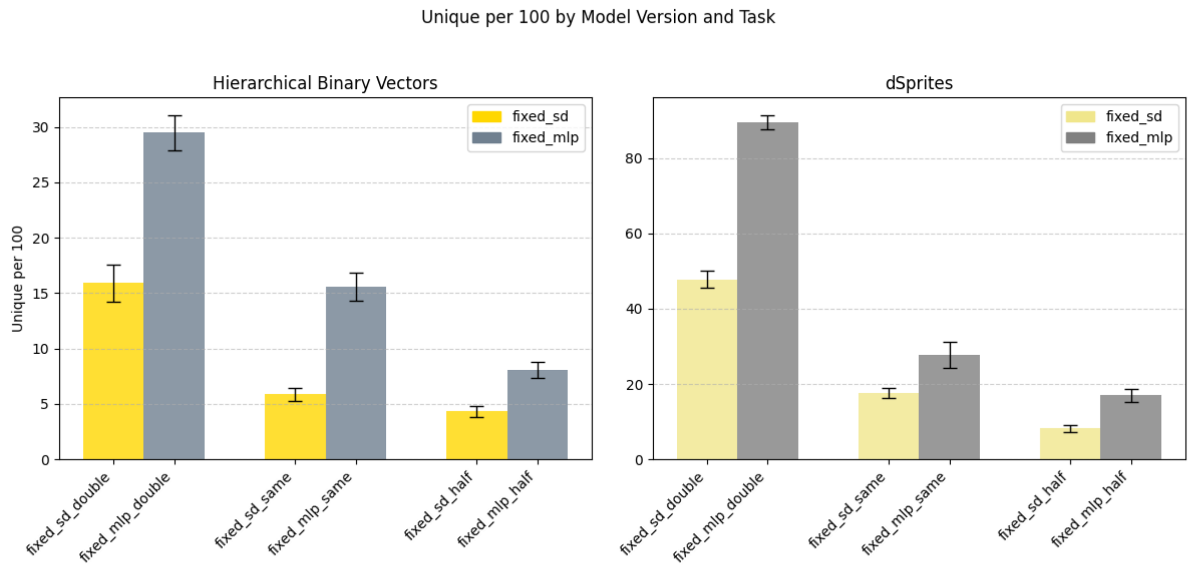

Figure 9. Number of unique token sequences for three standard-deviation conditions (×0.5, ×1, ×2). Bars denote mean ± 95 % CI across 100 batches of 100 trajectories per model; colours distinguish training methods (1) fixed-sd and (2) fixed-mlp.

To account for variability between different input items, the metric was calculated from 100 different starting points, where from each starting point, 100 trajectories were sampled and a token sequence generated (*e.g.* from starting point *3*, 100 trajectories were generated and the model produced 66 unique token sequences), from which the unique sequences were counted. Based on these 100 starting points, a mean value of this unique token number was taken.

These plots reveal a clear correlation between the standard deviation setting and the number of unique token sequences generated: models with higher standard deviations produce more unique sequences.

5. Discussion

In this thesis, I investigated how modulating entropy within artificial neural networks influences creative behaviour. Prominent theories of creativity suggest that humans can modulate the level of randomness in cognitive processes, and this modulation plays a key role in facilitating creative thought. Additionally, we explored the faster-and-further phenomenon, wherein creative individuals form associations more quickly and across a wider semantic space than less creative individuals. While this concept is well established in psychological literature, it lacked computational validation until Hills and Kenett (2025) introduced a model demonstrating entropy modulation. However, their model was based on a race model rather than artificial neural networks, leaving a gap in cognitively plausible modelling.

To address this, I sought a neural architecture that more closely emulates human cognitive reasoning. The Continuous Generative Flow Network (C-GFN) architecture was selected for its resemblance to several cognitive characteristics, including statistical inference, symbolic and sub-symbolic processing, and the emergence of attractor states. Within this framework, I focused on the forward-stepping module—a key component responsible for navigating toward attractor states in the latent space. By varying the standard deviation in this process, I created three versions of the model with differing levels of entropy.

There was a deliberate reason for choosing to replicate the faster-and-further phenomenon rather than relying on conventional creativity metrics such as the Torrance Tests of Creative Thinking (Torrance, 1966) or the Divergent Association Task (Olson et al., 2021). The rationale is straightforward: traditional creativity tests are based on evaluating an entity's textual output. Since the C-GFN model is not a language model, it is not designed to generate or process text. Therefore, a different evaluation approach was necessary, and the faster-and-further metric was well-suited for this purpose.

Experimental results across tasks such as HBV and dSprites demonstrated a consistent pattern: models with higher entropy reached attractor states more quickly and settled in regions further from their starting points in the latent space. These higher-entropy models also exhibited greater variation in generating candidate responses to a cue. Together, these findings align with the faster-and-further phenomenon and suggest that entropy modulation within artificial neural networks can indeed influence creative-like behaviour.

It is worth noting that all results were generated using the same random seed. This decision was based on tests showing that varying the seed had no significant effect on the outcomes. To verify this, I conducted speed, distance, and trajectory evaluations using two approaches: (1) a single seed with a batch size of 10,000, and (2) 100 different seeds with a batch size of 100 each, ensuring the total number of trajectories remained the same in both cases. The mean values obtained from both methods were nearly identical, and the 95% confidence intervals fell within a similar range. I conducted the final experiments using a single seed based on these findings.

Although theories linking creativity to entropy and associative exploration have existed for decades, they lacked computational support. The combined work of Hills and Kenett (2025) and this thesis helps bridge that gap, offering computational evidence that supports the idea that humans can modulate entropy to enhance creative behaviour.

It is important to note that the C-GFN model is an abstract representation rather than a complete simulation of cognitive processes. While it provides valuable insights, it cannot be used to draw definitive conclusions about creativity. Nonetheless, abstract models like the C-GFN can serve as valuable tools for hypothesis generation and guiding future empirical investigations.

In Section 2.3.1, I briefly addressed the topic of large language models (LLMs). One of my primary arguments against incorporating LLMs into this thesis was their lack of mastery in symbolic reasoning. Although the LLM Leaderboard[14] (as of May 2025) shows state-of-the-art models achieving impressive results in mathematics, these results can be misleading. This conclusion is based on my experimentation with these models. For instance, I found that the commonly used ChatGPT[15] model, version 4o (May 2025), struggled with performing accurate multiplications involving large numbers. I then tested the best-performing model at the time, ChatGPT version o3, which consistently provided correct answers to math-related queries. However, upon inspecting the model's chain-of-reasoning output, it became clear that it was executing Python[16] code to arrive at its answers. This suggests that the model did not engage in symbolic reasoning internally but instead relied on an external computational tool. Consequently, the model's high performance in mathematics can be seen as deceptive in assessing genuine reasoning capabilities. Given the rapid pace at which LLMs are evolving, there is

[14] https://www.vellum.ai/llm-leaderboard
[15] OpenAI (2025). ChatGPT (version o4): https://chatgpt.com/
[16] https://www.python.org/

currently no robust scientific methodology to reliably capture their latest capabilities, making most research on their utility quickly outdated.

In addition, I briefly mentioned the temperature scaling parameter in LLMs, which effectively adjusts the entropy of the model's outputs. However, as demonstrated in Zhao and colleagues' (2025) experiments, changing this parameter had no noticeable impact on the model's creative performance.

Several limitations emerged over the course of this work. The metric used to quantify the "faster" component of the phenomenon requires refinement. Current methods assume that trajectories approach an attractor in a relatively straight line and then orbit around it, a behaviour inferred from trajectory plots but not directly verified. Future work could involve stopping trajectories when they reach proximity to an attractor, allowing for variable-length sequences that better capture differences in convergence speed.

Further research should also evaluate the usefulness of generated associations, acknowledging that creativity involves more than randomness. This might be achieved by imposing task constraints and assessing whether the model can generate meaningful, constraint-satisfying outputs. Although the C-GFN model produces outputs that symbolically describe objects, this study did not measure its capacity for symbolic compositionality under different entropy settings. A high symbolic compositionality score would indicate useful responses in that case. The original work by Nam and colleagues (2023) included such analyses, but their code did not support replicating this component. I have contacted the authors, who indicated that the code will be released alongside the journal publication of their work sometime in the future.

Finally, it is worth emphasizing that while the C-GFN model produces variant expressions of input items (e.g., generating "scarlet" as a synonym for "red"), it does not create genuine conceptual associations such as "table–chair" or "table–tennis." Future efforts should aim to develop models that go beyond synonym generation and generate truly associative outputs, either by extending the current GFlowNet framework or by exploring alternative architectures.

6. Conclusion

This thesis proposes that modulating entropy in artificial neural networks may influence their capacity for creative behaviour. To evaluate this claim, the faster-and-further phenomenon was investigated in models with varying entropy levels. Empirical results from experiments using the C-GFN model indicate that higher-entropy systems exhibit both accelerated exploration and broader generative reach, suggesting an enhancement of creative-like behaviour. Future research should examine the practical utility of outputs produced under different entropy regimes and extend the analysis to additional reasoning algorithms and datasets.

References

Abraham, A., & Bubic, A. (2015). Semantic memory as the root of imagination. *Frontiers in Psychology*, *6*. https://doi.org/10.3389/fpsyg.2015.00325

Achiam, J., OpenAI, Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., Avila, R., Babuschkin, I., Balaji, S., Balcom, V., Baltescu, P., Bao, H., Bavarian, M., Belgum, J., … Zoph, B. (2023). *GPT-4 Technical Report* (Version 6). arXiv. https://doi.org/10.48550/ARXIV.2303.08774

Anderson, J. R., Matessa, M., & Lebiere, C. (1997). ACT-R: A Theory of Higher Level Cognition and Its Relation to Visual Attention. *Human–Computer Interaction*, *12*(4), 439–462. https://doi.org/10.1207/s15327051hci1204_5

Aston-Jones, G., & Cohen, J. D. (2005). AN INTEGRATIVE THEORY OF LOCUS COERULEUS-NOREPINEPHRINE FUNCTION: Adaptive Gain and Optimal Performance. *Annual Review of Neuroscience*, *28*(1), 403–450. https://doi.org/10.1146/annurev.neuro.28.061604.135709

Avery, M. C., & Krichmar, J. L. (2017). Neuromodulatory Systems and Their Interactions: A Review of Models, Theories, and Experiments. *Frontiers in Neural Circuits*, *11*. https://doi.org/10.3389/fncir.2017.00108

Beaty, R. E., & Kenett, Y. N. (2023). Associative thinking at the core of creativity. *Trends in Cognitive Sciences*, *27*(7), 671–683. https://doi.org/10.1016/j.tics.2023.04.004

Beaty, R. E., Silvia, P. J., Nusbaum, E. C., Jauk, E., & Benedek, M. (2014). The roles of associative and executive processes in creative cognition. *Memory & Cognition*, *42*(7), 1186–1197. https://doi.org/10.3758/s13421-014-0428-8

Benedek, M., Beaty, R. E., Schacter, D. L., & Kenett, Y. N. (2023). The role of memory in creative ideation. *Nature Reviews Psychology*, *2*(4), 246–257. https://doi.org/10.1038/s44159-023-00158-z

Benedek, M., Franz, F., Heene, M., & Neubauer, A. C. (2012). Differential effects of cognitive inhibition and intelligence on creativity. *Personality and Individual Differences*, *53*(4), 480–485. https://doi.org/10.1016/j.paid.2012.04.014

Benedek, M., Kenett, Y. N., Umdasch, K., Anaki, D., Faust, M., & Neubauer, A. C. (2017). How semantic memory structure and intelligence contribute to creative thought: A network science approach. *Thinking & Reasoning*, *23*(2), 158–183. https://doi.org/10.1080/13546783.2016.1278034

Bengio, Y., Lahlou, S., Deleu, T., Hu, E. J., Tiwari, M., & Bengio, E. (2021). *GFlowNet Foundations* (Version 4). arXiv. https://doi.org/10.48550/ARXIV.2111.09266

Bhatia, S., & Loomes, G. (2017). Noisy preferences in risky choice: A cautionary note. *Psychological Review*, *124*(5), 678–687. https://doi.org/10.1037/rev0000073

Boot, N., Baas, M., Van Gaal, S., Cools, R., & De Dreu, C. K. W. (2017). Creative cognition and dopaminergic modulation of fronto-striatal networks: Integrative review and research agenda. *Neuroscience & Biobehavioral Reviews*, *78*, 13–23. https://doi.org/10.1016/j.neubiorev.2017.04.007

Buber, E., & Diri, B. (2018). Performance Analysis and CPU vs GPU Comparison for Deep Learning. *2018 6th International Conference on Control Engineering & Information Technology (CEIT)*, 1–6. https://doi.org/10.1109/CEIT.2018.8751930

Busemeyer, J. R., Gluth, S., Rieskamp, J., & Turner, B. M. (2019). Cognitive and Neural Bases of Multi-Attribute, Multi-Alternative, Value-based Decisions. *Trends in Cognitive Sciences*, *23*(3), 251–263. https://doi.org/10.1016/j.tics.2018.12.003

Chrysikou, E. G. (2019). Creativity in and out of (cognitive) control. *Current Opinion in Behavioral Sciences*, *27*, 94–99. https://doi.org/10.1016/j.cobeha.2018.09.014

Cools, R., & D'Esposito, M. (2011). Inverted-U–Shaped Dopamine Actions on Human Working Memory and Cognitive Control. *Biological Psychiatry*, *69*(12), e113–e125. https://doi.org/10.1016/j.biopsych.2011.03.028

Cools, R., Sheridan, M., Jacobs, E., & D'Esposito, M. (2007). Impulsive Personality Predicts Dopamine-Dependent Changes in Frontostriatal Activity during Component Processes of Working Memory. *The Journal of Neuroscience*, *27*(20), 5506–5514. https://doi.org/10.1523/JNEUROSCI.0601-07.2007

Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, *441*(7095), 876–879. https://doi.org/10.1038/nature04766

Deco, G., Jirsa, V. K., & McIntosh, A. R. (2011). Emerging concepts for the dynamical organization of resting-state activity in the brain. *Nature Reviews Neuroscience*, *12*(1), 43–56. https://doi.org/10.1038/nrn2961

Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum Likelihood from Incomplete Data Via the *EM* Algorithm. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, *39*(1), 1–22. https://doi.org/10.1111/j.2517-6161.1977.tb01600.x

Doya, K. (2008). Modulators of decision making. *Nature Neuroscience*, *11*(4), 410–416. https://doi.org/10.1038/nn2077

Eppinger, B., Goschke, T., & Musslick, S. (2021). Meta-control: From psychology to computational neuroscience. *Cognitive, Affective, & Behavioral Neuroscience*, *21*(3), 447–452. https://doi.org/10.3758/s13415-021-00919-4

Faisal, A. A., Selen, L. P. J., & Wolpert, D. M. (2008). Noise in the nervous system. *Nature Reviews Neuroscience*, *9*(4), 292–303. https://doi.org/10.1038/nrn2258

Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, *28*(1–2), 3–71. https://doi.org/10.1016/0010-0277(88)90031-5

Gawlikowski, J., Tassi, C. R. N., Ali, M., Lee, J., Humt, M., Feng, J., Kruspe, A., Triebel, R., Jung, P., Roscher, R., Shahzad, M., Yang, W., Bamler, R., & Zhu, X. X. (2023). A survey of uncertainty in deep neural networks. *Artificial Intelligence Review*, *56*(S1), 1513–1589. https://doi.org/10.1007/s10462-023-10562-9

Gilbertson, T., & Steele, D. (2021). Tonic dopamine, uncertainty and basal ganglia action selection. *Neuroscience*, *466*, 109–124. https://doi.org/10.1016/j.neuroscience.2021.05.010

Gray, K., Anderson, S., Chen, E. E., Kelly, J. M., Christian, M. S., Patrick, J., Huang, L., Kenett, Y. N., & Lewis, K. (2019). "Forward flow": A new measure to quantify free thought and predict creativity. *American Psychologist*, *74*(5), 539–554. https://doi.org/10.1037/amp0000391

Heathcote, A., & Matzke, D. (2022). Winner Takes All! What Are Race Models, and Why and How Should Psychologists Use Them? *Current Directions in Psychological Science*, *31*(5), 383–394. https://doi.org/10.1177/09637214221095852

Heinen, D. J. P., & Johnson, D. R. (2018). Semantic distance: An automated measure of creativity that is novel and appropriate. *Psychology of Aesthetics, Creativity, and the Arts*, *12*(2), 144–156. https://doi.org/10.1037/aca0000125

Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., & Lerchner, A. (2017). Beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. *International Conference on Learning Representations*. https://openreview.net/forum?id=Sy2fzU9gl

Hills, T. T. (2019). Neurocognitive free will. *Proceedings of the Royal Society B: Biological Sciences*, *286*(1908), 20190510. https://doi.org/10.1098/rspb.2019.0510

Hills, T. T., Jones, M. N., & Todd, P. M. (2012). Optimal foraging in semantic memory. *Psychological Review*, *119*(2), 431–440. https://doi.org/10.1037/a0027373

Hills, T. T., & Kenett, Y. N. (2022). Is the Mind a Network? Maps, Vehicles, and Skyhooks in Cognitive Network Science. *Topics in Cognitive Science*, *14*(1), 189–208. https://doi.org/10.1111/tops.12570

Hills, T. T., & Kenett, Y. N. (2025). An entropy modulation theory of creative exploration. *Psychological Review*, *132*(1), 239–251. https://doi.org/10.1037/rev0000511

Hu, E. J., Malkin, N., Jain, M., Everett, K., Graikos, A., & Bengio, Y. (2023). *GFlowNet-EM for learning compositional latent variable models* (Version 2). arXiv. https://doi.org/10.48550/ARXIV.2302.06576

Hu, L., Liu, L., Yang, S., Chen, X., Tan, Z., Ali, M. A., Li, M., & Wang, D. (2024). *Understanding Reasoning in Chain-of-Thought from the Hopfieldian View* (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2410.03595

Inagaki, H. K., Fontolan, L., Romani, S., & Svoboda, K. (2019). Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature*, *566*(7743), 212–217. https://doi.org/10.1038/s41586-019-0919-7

Jackson, J. C., Watts, J., List, J.-M., Puryear, C., Drabble, R., & Lindquist, K. A. (2021). From Text to Thought: How Analyzing Language Can Advance Psychological Science. *Perspectives on Psychological Science*, *17*(3), 805–826. https://doi.org/10.1177/17456916211004899

Jung, R. E., Wertz, C. J., Meadows, C. A., Ryman, S. G., Vakhtin, A. A., & Flores, R. A. (2015). Quantity yields quality when it comes to creativity: A brain and behavioral test of the equal-odds rule. *Frontiers in Psychology*, *6*. https://doi.org/10.3389/fpsyg.2015.00864

Kenett, Y. N. (2025). The Role of Knowledge in Creative Thinking. *Creativity Research Journal*, *37*(2), 242–249. https://doi.org/10.1080/10400419.2024.2322858

Kenett, Y. N., & Faust, M. (2019). A Semantic Network Cartography of the Creative Mind. *Trends in Cognitive Sciences*, *23*(4), 271–274. https://doi.org/10.1016/j.tics.2019.01.007

Knill, D. C., & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neurosciences*, *27*(12), 712–719. https://doi.org/10.1016/j.tins.2004.10.007

Kvam, P. D., Pleskac, T. J., Yu, S., & Busemeyer, J. R. (2015). Interference effects of choice on confidence: Quantum characteristics of evidence accumulation. *Proceedings of the National Academy of Sciences*, *112*(34), 10645–10650. https://doi.org/10.1073/pnas.1500688112

Lahlou, S., Deleu, T., Lemos, P., Zhang, D., Volokhova, A., Hernández-García, A., Ezzine, L. N., Bengio, Y., & Malkin, N. (2023). *A theory of continuous generative flow networks* (Version 2). arXiv. https://doi.org/10.48550/ARXIV.2301.12594

Lin, H., & Vartanian, O. (2018). A Neuroeconomic Framework for Creative Cognition. *Perspectives on Psychological Science*, *13*(6), 655–677. https://doi.org/10.1177/1745691618794945

Luce, R. D. (2005). *Individual choice behavior: A theoretical analysis.* Dover Publications. https://doi.org/10.1037/14396-000

Malaie, S., Spivey, M. J., & Marghetis, T. (2024). Divergent and Convergent Creativity Are Different Kinds of Foraging. *Psychological Science*, *35*(7), 749–759. https://doi.org/10.1177/09567976241245695

Mandera, P., Keuleers, E., & Brysbaert, M. (2017). Explaining human performance in psycholinguistic tasks with models of semantic similarity based on prediction and counting: A review and empirical validation. *Journal of Memory and Language*, *92*, 57–78. https://doi.org/10.1016/j.jml.2016.04.001

McClelland, J. L., Botvinick, M. M., Noelle, D. C., Plaut, D. C., Rogers, T. T., Seidenberg, M. S., & Smith, L. B. (2010). Letting structure emerge: Connectionist and dynamical systems approaches to cognition. *Trends in Cognitive Sciences*, *14*(8), 348–356. https://doi.org/10.1016/j.tics.2010.06.002

McClelland, J. L., & Rogers, T. T. (2003). The parallel distributed processing approach to semantic cognition. *Nature Reviews Neuroscience*, *4*(4), 310–322. https://doi.org/10.1038/nrn1076

Mednick, S. (1962). The associative basis of the creative process. *Psychological Review*, *69*(3), 220–232. https://doi.org/10.1037/h0048850

Meusburger, P. (2009). Milieus of Creativity: The Role of Places, Environments, and Spatial Contexts. In P. Meusburger, J. Funke, & E. Wunder (Eds.), *Milieus of Creativity* (Vol. 2, pp. 97–153). Springer Netherlands. https://doi.org/10.1007/978-1-4020-9877-2_7

Nam, A., Elmoznino, E., Malkin, N., McClelland, J., Bengio, Y., & Lajoie, G. (2023). *Discrete, compositional, and symbolic representations through attractor dynamics* (Version 2). arXiv. https://doi.org/10.48550/arXiv.2310.01807

Olson, J. A., Nahas, J., Chmoulevitch, D., Cropper, S. J., & Webb, M. E. (2021). Naming unrelated words predicts creativity. *Proceedings of the National Academy of Sciences*, *118*(25), e2022340118. https://doi.org/10.1073/pnas.2022340118

Omidvarnia, A., Zalesky, A., Mansour L, S., Van De Ville, D., Jackson, G. D., & Pedersen, M. (2021). Temporal complexity of fMRI is reproducible and correlates with higher order cognition. *NeuroImage*, *230*, 117760. https://doi.org/10.1016/j.neuroimage.2021.117760

Ovando-Tellez, M., Benedek, M., Kenett, Y. N., Hills, T., Bouanane, S., Bernard, M., Belo, J., Bieth, T., & Volle, E. (2022). An investigation of the cognitive and neural correlates of semantic memory search related to creative ability. *Communications Biology*, *5*(1), 604. https://doi.org/10.1038/s42003-022-03547-x

Pezzulo, G., Rigoli, F., & Friston, K. (2015). Active Inference, homeostatic regulation and adaptive behavioural control. *Progress in Neurobiology*, *134*, 17–35. https://doi.org/10.1016/j.pneurobio.2015.09.001

Rastelli, C., Greco, A., De Pisapia, N., & Finocchiaro, C. (2022). Balancing novelty and appropriateness leads to creative associations in children. *PNAS Nexus*, *1*(5), pgac273. https://doi.org/10.1093/pnasnexus/pgac273

Richie, R., Aka, A., & Bhatia, S. (2023). Free association in a neural network. *Psychological Review*, *130*(5), 1360–1382. https://doi.org/10.1037/rev0000396

Rolls, E. T., & Deco, G. (2010). *The Noisy BrainStochastic Dynamics as a Principle of Brain Function*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199587865.001.0001

Rolls, E. T., Deco, G., Huang, C.-C., & Feng, J. (2024). The connectivity of the human frontal pole cortex, and a theory of its involvement in exploit versus explore. *Cerebral Cortex*, *34*(1), bhad416. https://doi.org/10.1093/cercor/bhad416

Rossmann, E., & Fink, A. (2010). Do creative people use shorter associative pathways? *Personality and Individual Differences*, *49*(8), 891–895. https://doi.org/10.1016/j.paid.2010.07.025

Rumelhart, D. E., McClelland, J. L., & AU. (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition: Foundations*. The MIT Press. https://doi.org/10.7551/mitpress/5236.001.0001

Runco, M. A., & Jaeger, G. J. (2012). The Standard Definition of Creativity. *Creativity Research Journal*, *24*(1), 92–96. https://doi.org/10.1080/10400419.2012.650092

Saxe, A. M., McClelland, J. L., & Ganguli, S. (2019). A mathematical theory of semantic development in deep neural networks. *Proceedings of the National Academy of Sciences*, *116*(23), 11537–11546. https://doi.org/10.1073/pnas.1820226116

Shadlen, M. N., & Newsome, W. T. (2001). Neural Basis of a Perceptual Decision in the Parietal Cortex (Area LIP) of the Rhesus Monkey. *Journal of Neurophysiology*, *86*(4), 1916–1936. https://doi.org/10.1152/jn.2001.86.4.1916

Shine, J. M., Breakspear, M., Bell, P. T., Ehgoetz Martens, K. A., Shine, R., Koyejo, O., Sporns, O., & Poldrack, R. A. (2019). Human cognition involves the dynamic integration of neural activity and neuromodulatory systems. *Nature Neuroscience*, *22*(2), 289–296. https://doi.org/10.1038/s41593-018-0312-0

Silvia, P. J. (2015). Intelligence and Creativity Are Pretty Similar After All. *Educational Psychology Review*, *27*(4), 599–606. https://doi.org/10.1007/s10648-015-9299-1

Simonton, D. K. (1997). Creative productivity: A predictive and explanatory model of career trajectories and landmarks. *Psychological Review*, *104*(1), 66–89. https://doi.org/10.1037/0033-295X.104.1.66

Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to Grow a Mind: Statistics, Structure, and Abstraction. *Science*, *331*(6022), 1279–1285. https://doi.org/10.1126/science.1192788

Todd, P. M., & Hills, T. T. (2020). Foraging in Mind. *Current Directions in Psychological Science*, *29*(3), 309–315. https://doi.org/10.1177/0963721420915861

Torrance, E. P. (1966). *Torrance Test of Creative Thinking: Directions Manual and Scoring Guide*. Personnel Press.

Touboul, J., Hermann, G., & Faugeras, O. (2012). Noise-Induced Behaviors in Neural Mean Field Dynamics. *SIAM Journal on Applied Dynamical Systems*, *11*(1), 49–81. https://doi.org/10.1137/110832392

Troyer, A. K., Moscovitch, M., Winocur, G., Alexander, M. P., & Stuss, D. (1998). Clustering and switching on verbal fluency: The effects of focal frontal- and temporal-lobe lesions. *Neuropsychologia*, *36*(6), 499–504. https://doi.org/10.1016/S0028-3932(97)00152-8

Volle, E. (2018). Associative and Controlled Cognition in Divergent Thinking: Theoretical, Experimental, Neuroimaging Evidence, and New Directions. In R. E. Jung & O. Vartanian (Eds.), *The Cambridge Handbook of the Neuroscience of Creativity* (1st ed., pp. 333–360). Cambridge University Press. https://doi.org/10.1017/9781316556238.020

Wang, X.-J. (2008). Decision Making in Recurrent Neuronal Circuits. *Neuron*, *60*(2), 215–234. https://doi.org/10.1016/j.neuron.2008.09.034

Weiss, S., Goecke, B., & Wilhelm, O. (2024). How Much Retrieval Ability Is in Originality? *The Journal of Creative Behavior*, *58*(3), 370–387. https://doi.org/10.1002/jocb.659

Wilson, R. C., Bonawitz, E., Costa, V. D., & Ebitz, R. B. (2021). Balancing exploration and exploitation with information and randomization. *Current Opinion in Behavioral Sciences*, *38*, 49–56. https://doi.org/10.1016/j.cobeha.2020.10.001

Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology: General*, *143*(6), 2074–2081. https://doi.org/10.1037/a0038199

Yechiam, E., Busemeyer, J. R., Stout, J. C., & Bechara, A. (2005). Using Cognitive Models to Map Relations Between Neuropsychological Disorders and Human Decision-Making Deficits. *Psychological Science*, *16*(12), 973–978. https://doi.org/10.1111/j.1467-9280.2005.01646.x

Zhao, Y., Zhang, R., Li, W., Huang, D., Guo, J., Peng, S., Hao, Y., Wen, Y., Hu, X., Du, Z., Guo, Q., Li, L., & Chen, Y. (2025). Assessing and Understanding Creativity in Large Language Models. *Machine Intelligence Research*. https://doi.org/10.1007/s11633-025-1546-4

Zhu, J.-Q., Sanborn, A. N., & Chater, N. (2020). The Bayesian sampler: Generic Bayesian inference causes incoherence in human probability judgments. *Psychological Review*, *127*(5), 719–748. https://doi.org/10.1037/rev0000190

Appendices

Appendix I: Glossary

Artificial neural network: computational model inspired by biological neural networks, consisting of interconnected processing units (neurons) that learn patterns from data through adaptive connections.

Associative thinking: a cognitive process involving the connection and linkage of ideas or memories based on their relationships or similarities, often spontaneously or creatively.

Attractor: a stable state or pattern toward which a dynamical system evolves, often used metaphorically in cognitive and neural models to represent persistent or recurrent states.

Bayesian inference: a statistical method of updating beliefs or probability estimates based on prior knowledge and newly observed evidence using Bayes' theorem.

Divergent thinking: the ability to generate many different ideas in response to open-ended prompts (e.g., unusual object uses).

Distributional semantic modelling: a natural language processing modelling approach used to analyse and represent the meaning of words in a given context by studying their distribution patterns in a large corpus of text.

Dynamical system: a mathematical system characterized by variables that evolve over time according to specific rules or equations, describing how states change dynamically.

Entropy: a measure of randomness, disorder, or unpredictability within a system, commonly used in physics, information theory, and statistics.

Evidence accumulation (race) model: a decision-making model describing the cognitive process by which evidence supporting multiple alternatives accumulates over time, until a decision threshold is reached.

Free association: a thought process in which a person freely associates ideas, words, or memories without deliberate censorship or cognitive control.

Generative model: a computational model that learns to generate new data samples resembling a given dataset by capturing underlying data distributions.

Goal-directed association: a thought process in which a person associates ideas, words, or memories with a specific goal or purpose in mind.

Large language model: a computational model trained on vast amounts of textual data, capable of generating human-like text and performing language-related tasks.

Latent space: an abstract multidimensional space where hidden or underlying variables represent complex data or phenomena in a compressed or meaningful way.

Latent variables: hidden or unobserved variables inferred from observed data that explain relationships or variations in a dataset or model.

Multilayer perceptron (MLP): a neural network with multiple layers of connected neurons that learn complex input-output mappings using nonlinear activation and backpropagation.

Neural stochastic dynamical system: a computational model combining neural network architectures with stochastic dynamics to represent systems whose behaviour includes both neural interactions and random processes.

Neuromorphic computing: a method of computing modelled after the architecture and operations of the human brain.

Normal distribution: a symmetric, bell-shaped probability distribution characterized by a mean and standard deviation, widely used in statistics and data analysis.

Phylogenetic tree: a branching diagram representing evolutionary relationships among biological species or entities based on genetic, morphological, or other data.

Principal component analysis: a linear technique for reducing dimensionality, commonly used for data exploration, visualization, and preprocessing.

Prior P(s): a probability distribution expressing initial beliefs or assumptions about a variable before observing new evidence, central to Bayesian reasoning.

Random walkers: computational entities or models that move randomly through a network or space, commonly used to model diffusion processes or information spread.

Reward function: a mathematical expression defining the goals or objectives of an agent within reinforcement learning frameworks, guiding behaviour through positive or negative reinforcement.

Semantic distance: a measure of the inverse of similarity (or relatedness) between two words or concepts, often used in natural language processing and computational linguistics.

Semantic memory: a form of long-term memory involving the storage and retrieval of factual knowledge, concepts, and meanings independent of personal experiences.

Spreading activation: a cognitive model describing how activation of one concept in a memory network spreads to related concepts, influencing retrieval and associations.

Sub-symbolic reasoning: cognitive or computational processes that operate directly on patterns or representations at a numerical or distributed level, without explicit symbols.

Symbolic reasoning: cognitive or computational processes involving explicit manipulation of symbols and symbolic expressions according to defined rules and logic.

Unsupervised learning: a type of machine learning in which models discover patterns or structures within unlabelled datasets without explicit guidance or predefined labels.

Appendix II: Equations

The model is trained by considering that sampling from P_θ should be drawn in proportion to an unnormalized learned density, which we refer to as the reward function $R_\phi(z_T, s; x)$ as defined in Equation 2:

$$P_\theta(z_T, s \mid x) \propto R_\phi(z_T, s; x) = \exp\left(-D_{\mathrm{KL}}\left(P_\phi(z_0 \mid x) \parallel P_\phi(z_0 \mid \widehat{z_s})\right)\right) \cdot P_\phi(s \mid z_T) \cdot P(s). \quad (2)$$

This reward function integrates several components. The first is a KL-divergence term, $D_{\mathrm{KL}}\left(P_\phi(z_0 \mid x) \parallel P_\phi(z_0 \mid \widehat{z_s})\right)$, which captures how much information is lost when the input is converted into a discrete form—thus, its negative reflects how well the original information is preserved after tokenization. The second term, $P_\phi(s \mid z_T)$, evaluates how likely different sentences are given the same latent representation, regardless of the specific input x This accounts for the fact that semantically similar sentences (like "red" and "scarlet") may correspond to the same underlying concept. The assumption here is that if someone perceives an apple and thinks "red", they could just as plausibly have thought "scarlet", assuming both words lead to similar latent representations—disregarding the fact that "red" is more commonly used.

Appendix III: Data and Code Availability

The code and results generated during this thesis are available at: https://github.com/jki-rikal/gfn_attractors_entropy. This repository also includes the instructions for reproducing these results.

Appendix IV: ChatGPT Prompts

ChatGPT[17] was used for creating the plots in Figure 7, Figure 8 and Figure 9. The resulting code from these prompts can be found under the provided code repository (see Appendix III: Data and Code Availability) with the file name "plot_final_results.py". The following prompts were used for achieving these results:

1. "Create 3 separate bar plots for this data using python: [copy-paste data from Microsoft Excel[18]]"
2. "Now differentiate the "mlp" and "sd" versions from each other using color"
3. "Now replace the legend of "SD" and "MLP" with "fixed sd" and "fixed MLP""
4. "Now set different color themes for the different plots (relative speed plot has different theme than total trajectory distance plot etc.)"
5. "Create the same kinds of plots for this data, that is a little bit different now: Data combined from HBV and dsprites tasks. Each metric is a mean from 4 different models (2 for each task) [copy-paste data from Excel]"
6. "The plots should have a grid and the bars a little bit transparent, so that the grids behind can be seen. In addition, keep the different coloring themes on the different plots"
7. "also make the plots 4 units high and 6 units wide"
8. "How would I modify this code: [copied code from previous "plot_final_results.py"] to include the 95% CI higher and lower bounds? [copied data from Excel spreadsheet]"
9. "Now construct two plots side-by-side to also include the results from the dsprites task on every plot: [copied additional data from spreadsheet]"
10. "I want one subplot per metrics, but on the subplot, on the left part of the plot should be hierarchical binary vectors task results and on the right dsprites results"
11. "Right now, you constructed 1 figure. Can you do 3 separate figures?"
12. "now in each figure, separate the tasks by different plots"
13. "Now use different scalings for the different tasks, also do so that the "double" bars are touching, "same" bars are touching and "half" bars are touching"
14. "Now add legends explaining the colors. Therefore there should be fixed_sd and fixed_mlp"

[17] OpenAI (2025). ChatGPT (version o4-mini-high): https://chatgpt.com/
[18] https://www.microsoft.com/en-us/microsoft-365/excel

License