Bachelor Computer Science & Datascience and Artificial Intelligence

Luhmannian Society Digitized:

A study on structured creativity as an autopoeitic process within the DIFI framework

Victor Michael de Munck

First supervisor : Rob Saunders
Second supervisor: Tessa Verhoef

BACHELOR THESIS

Leiden Institute of Advanced Computer Science (LIACS)
www.liacs.leidenuniv.nl 07/12/2026

1 Introduction

Creativity is often thought of as a thought. A "Eureka!" moment where suddenly everything connects. Yet, while the fire of creativity is often imagined as growing from that "spark" where everything seems to suddenly connect, there must be more to this process. Where do these thoughts emerge from, what drives them towards something that is inherently "creative," and how do we know what this is? Creativity, like fire, is a continuous interaction between elements and their environment[Iba10]. Understanding how to model the system will help us understand how thoughts become creative discoveries. Understanding how the elements and the environment interact within the system helps us understand how something is deemed creative, and another thing is not.

Creativity, as it is studied here, is seen as a social system the model of which was presented by Mihaly Csikszentmihaly[Csi99]. The individual exists in this system as a producer of artifacts, and as artifacts flow through society, they may be deemed "creative," at which point they join the domain. Creativity is therefore a cultural phenomenon[Csi99, SG05, Iba10]. The personal mechanisms by which individuals define creativity are not the central object of study; what is essential is how these mechanisms work together to constitute a creative culture.

In the middle of the twentieth century, the creativity field was dominated by psychological perspectives[Iba10, Csi99, SL99]. Sternberg and Lubart trace a succession of such paradigms—among them the psychodynamic, psychometric, cognitive, and social-personality—each locating creativity in the intrapsychic and tying it to concepts such as personality and cognition[SL99]. Individualistic paths of inquiry sought to identify the creative individual's traits and to understand how these traits influence creativity. Problems emerged when focusing solely on these traits as explanatory variables. Traits alone could not explain why some people with apparently creative dispositions stopped producing, nor why others with few such traits succeeded [Csi99]. Recognizing this limitation, one is forced to consider many interacting forces, such as intellectual abilities, knowledge, thinking styles, personality, motivation, and environment [SL99].

Margaret Boden differentiates personal creativity (P-creativity) from historical creativity (H-creativity). P-creativity refers to one's intrinsic threshold for novelty, when an individual considers something they have made to be new and creative to themselves[Bod04]. This is not necessarily socially important, for even if you derive Newton's law of gravitation on your own, society will not bat an eye, as Newton already did so centuries ago. Although for H-creativity to occur there has to be first P-creativity. Our study focuses on the social mechanisms that turn P-creativity into a socially recognized H-creativity. We focus on the forces in the individual's environment.

Hence, the study of creativity evolved into the systems-property perspective mentioned earlier. Creativity is typically thought to require both novelty and usefulness[RJ12, SL99, Bod04, TTJ19]. Something must be new, must bring discovery, but it must also be deemed valuable. Otherwise it is quickly forgotten, a splash that momentarily distracts the continuous stream of creativity[Csi99, Ass10]. To model creativity as a systems property, then, we must consider how to build a system that deems something to be new and valuable historically speaking(H-creativity)[Bod04].

Simulations have become increasingly important to this direction of research[SG05, Sau19]. By simulating the interactions of many "producer" and "evaluator" agents we model a society in which the sharing of ideas and opinions determines the creativity of an artifact[SGJ09, Sau19]. If each agent retains a memory of the artifacts they have seen or created, they can determine whether something is new, and by evaluating an artifact against their preferences they determine its value[LH18]. As these social interactions are scaled up, the society defines its own creative values and builds its own history[Mot25, SG05]. Through many shared evaluations, the personal meanings of creativity merge into a creative culture. Every contextual variable, the agents' preferences for novelty, their memories, their own abilities, becomes generalized in the domain of the society[Csi99, Iba10, SG05]. The domain thus serves as cultural history, and one can understand the rules and meanings that define creativity in a given period by studying the artifacts of that time that survive in the domain[Csi99]. Researchers have examined how the communication and sharing of artifacts builds this domain and how contextual variables affect those communications[SG05, SGJ09, SB15, Mot25]. Thus, research into creativity has veered towards interactionalist models.

The point of this paper is to extend the systems perspective to the domain itself. How do contextual variables influence what is remembered as culture or defined as meaningful? How does the cultural domain influence the future of culture? How does recognizing some structure in culture feed back into the interactions themselves?[Luh95, Iba10]. This paper's research uses another social systems theory to address the creative social system that Csikszentmihayli has presented, which treats social systems as self-reproducing through their differentiation from their environment[Luh95]. By building on the perspective of creativity as a social system, this paper attempts to apply such a theory to creativity. This theory directly addresses a gap in current research, where the cultural domain has been treated as passive storage, a list to which "successful" artifacts are added and nothing more[Mot25, Sau19]. The point of culture is the shared meaning and rules it defines for creativity. The cultural values that determined past successful artifacts ought to influence what defines the value and novelty of newly created artifacts on a societal level[Luh95, Ass10]. What can be mined from the field of exploration, and what makes certain discoveries successful, is observed socially. Knowledge that is retained by society is often done so under the presumption of its utility [TTJ19]; ignoring how the perception of utility evolves, and the population with it, is limiting current models [Iba10]. From this idea three commitments follow. First, if the domain is a memory rather than a store, its structure must lie in how it is retrieved. A potter does not walk into a warehouse and pick a pot at random to base their next piece on; in real creative fields — scientific, artistic, technical — influence is traceable, and it is that traceability that lets one work critique, extend, or answer another [Csi99, Ass10, BO12]. Representing culture as a pile discards exactly the property that makes it culture. Second, a structure is inert unless something moves through it, so agents require a navigation strategy — and one answerable both to their own interest and to the field's acceptance, since it is the conversion of P-creativity into H-creativity that concerns us. Third, prior models assume that the agents who build a culture are the same ones who use it forever; under that assumption an agent's private memory can outgrow the domain and cultural memory has nothing left to do. Only when agents die and the domain persists does the navigation of an inherited culture become necessary [Ass10, BO12].

2 Research Question

To measure how modeling an active domain can refine the research and understanding of creative societies, this paper poses the following research questions:

How does a structured cultural domain, through which creative agents can navigate, affect the growth and richness (in terms of differing levels of artifact complexity) of the domain?

How does the navigation of a structured domain reveal itself in the shaping of the creative society's interaction network?

Sub-questions

- **RQ1a:** Which domain-association rule produces the greatest stylistic diversity as measured by average difference in feature vectors?
- **RQ1b:** Does 'higher level' (lineage-based) domain association increase an association based path-dependence (lock-in), or does it enable sustained novel production by providing a more stable "rough ground" for recursive discovery?
- **RQ2a:** How does generational turnover influence the field's structure and individual agent strategies for acquiring domain knowledge?
- **RQ2b:** How does a heterogeneous vs homogeneous mixing of the initiated strategy affect whether/what navigation strategies emerge?

3 Theoretical Background

3.1 The DIFI Framework

The DIFI framework by Csikszentmihalyi has given researchers a foundation for modeling social creativity[Csi99]. Creativity emerges under the lattice of influence between the individual, the field, and the domain as seen in figure 1[Csi99].

The individual does not just randomly produce variation, they draw from the domain and field for exposure to novel and valuable works[Csi99, Sau19]. The different knowledge, memory, preferences, and intrinsic motivations of the individual are the variation property. What is considered novel and valuable by other agents is added to the domain. How other agents evaluate artifacts maintains the same kinds of idiosyncrasies as the creator[Sau19]. The point is that culture, as the active total memory of the past, is what allows agents to communicate meaningfully under all these variables[Ass10]. The other agents make up the field. The field and its gatekeepers select the knowledge that is worth preserving. The field consists of agents that have knowledge over the domain and maintain influence over its contents[Csi99]. The cultural memory of the domain is essential for the creative process, for one to evaluate the novelty of a creative artifact there must be an ability to reference to what has already been done. A similar argument applies to the value of an artifact. Csikszentmihaly refers to the domain, "a set of already existing objects, rules, representations, or notations"[Csi99, p. 315]. The individual and the field are inherited unchanged from Motz [Mot25]. What this thesis addresses is the domain, which in the inherited model is

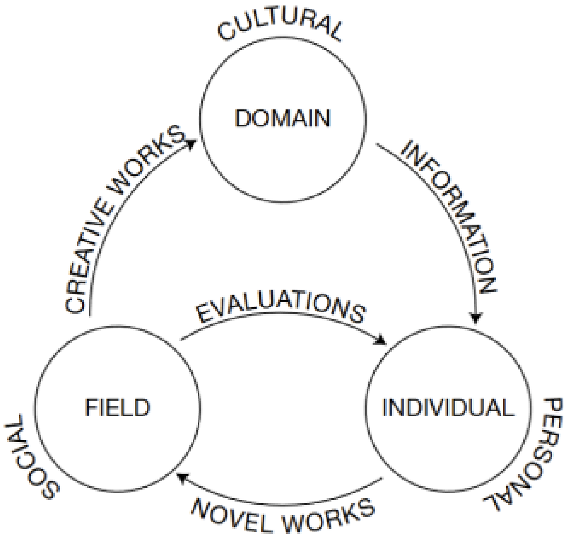

Figure 1: Csikszentmihalyi's Systems View of Creativity (DIFI Framework), showing the interaction between the Cultural Domain, the Personal Individual, and the Social Field. Adapted from [Csi99] by [Sau19].

an undifferentiated list — every stored artifact equally reachable by every agent — so that the domain-to-individual transmission DIFI describes as selective is, in implementation, uniform. Giving that transmission an internal structure is the work of section 5.

3.2 Autopoietic Creative Social Systems

Takashi Iba identified bridges between creative social systems and autopoietic social systems[Iba10]. Iba thus makes the claim that creativity is an autopoietic system which runs on discoveries. This paper uses this claim as a design premise, building and analyzing our model as if creativity is autopoietic. Autopoietic systems are self-organizing and self-producing, consisting of elements and functions that operate only between those elements[Moe12, Luh95]. Elements are "momentary events that have no duration"[Iba10, p. 6616]. For our model, the share–accept event is the element in our system. Though the event vanishes it leaves behind a trace, or a structure in the domain which contains the next element. To maintain itself the system must continuously reproduce itself[Moe12]. Luhmann applies this biological term to social systems, arguing that society, rather than being made up of individuals, is made up of communication[Luh95, Moe12]. Communication filters out all that is not communication: thoughts, gestures, and words that are not interpreted[Moe12]. In our model, once an artifact is shared it can affect the system; therefore it is the sharing event which filters what is creative for society.

These creative autopoietic systems are said to reproduce discoveries, or differentiate discoveries from non-discoveries, or P-creativity from H-creativity[Bod04]. The reproduction of communication in Iba's theory follows a track of information, utterance, and understanding[Iba10]. These three steps mirror the DIFI structure in that information is acquired in the individual's memory (generation and evaluation phases), then uttered under the pretense of a shared meaning created through culture/domain (sharing phase), and the field attempts its own understanding of the value (interaction phase). This value is not objectively true, but rather represents the recognition of something creative by a subset of the field[SG05, Jen10]. This is an aspect of the functional closure of autopoietic systems, in that the thoughts of an individual are mangled in their expression and interpretation; everything that is not communication exists in the environment of the system, only able to perturb it[Moe12]. Even the code and its functions exist outside the system, as changes to the code only change the system from the outside rather than being an operation of the system. This places a rule on every comparison reported here: the code is the base and the retrieval structure is the contingency, so only what varies in communication, once the substrate is held fixed, may be read as cultural.

Discoveries emerge from a translated synthesis, drawn from Luhmann, consisting of an idea, an association, and a consequence[Iba10]. This idea of contingency originates with Luhmann; that every discovery is contingent on a previous discovery, yet is as likely to occur as any other discovery, lending contingency to the discovery itself[Moe12, Iba10]. The domain artifacts remain as media within this pool of idiosyncrasy, contingency becoming the association edge between past and present[Iba10, Ass10]. Each accepted artifact is bred from a retrieved prior artifact so a domain entry only becomes reachable through earlier artifacts. Association can be read in at least two ways: by resemblance between works, and by descent from one work to the next. These need not coincide, and that difference motivates the two retrieval structures compared in section 5.

Consider two ends of a spectrum: creativity that is more logical, like mathematics, and creativity that is more emotional, like painting. Each can be expressed in historically creative ways and each discovery can be traced[Bod04]. Einstein's theory of general relativity extended Newton's law of universal gravitation, and Manet's *Le Déjeuner sur l'herbe* was essentially a reworking of Marcantonio Raimondi's *The Judgement of Paris*[Axe05, Læs05]. It is the recognizable structure in both types of creative fields that allows one work to critique another, another to adjust, and discoveries to emerge. From the theoretical anchors of Niklas Luhmann and Takashi Iba, the case for an active domain emerges. The self-differentiation of an autopoietic system from its environment becomes a creative domain whose structured past filters what counts as creative in the future. Iba's contingent discovery inspires the associations made between each artifact.

3.3 Cumulative Cultural Evolution

Evolution and creativity have gone hand in hand for humans. As we became social creatures, communication allowed us to share good ideas, habits, and techniques[BO12, SB15]. This allowed groups to exploit the most renowned practices. As is the case with all exploitation, there is an exploratory trade-off. Groups that fostered creativity and the variability of their individuals sometimes found techniques so beneficial that they spread throughout society[BO12, SG05]. This is only helpful if culture is cumulative, that is, if it is stored in some sense, allowing new members and generations to adopt and reinterpret history. Through culture, society builds an atemporal framework, transcending the individual's time and relating past, present, and future[Ass10]. This concept shares the atemporal and structural features of Iba's autopoietic creative system. Yet for an artifact or communication to exist in this framework it must be capable of passing through time.

In its biological form, evolution runs on three tightly defined mechanisms: *mutation*, *natural selection*, and *transmission*. Cumulative culturalists borrow this three-part logic and apply it to ideas rather than genes[BO12]. The borrowing is an analogy, but it is useful in understanding the transmission and context of culture. In our model, the analogue of variation is the artifact an individual produces, the idea; the analogue of selection is the field's evaluation of which artifacts are deemed creative; and the analogue of transmission is the domain returning accepted artifacts to individuals. Read this way, the cultural system, like the biological one, both produces its own reproduction and the conditions for that reproduction. As culture accumulates, whatever is considered novel or valuable must confront that accumulated knowledge[Csi99, Iba10]. The domain is what facilitates this confrontation by making general relations between artifacts. Culture has been called "a memory of society... that transcends external symbols"[Ass10, p. 97]. To capture culture we need to look outside the individual symbols and look for the patterns which may explain symbols in terms of some culture.

This framework allows culture to accumulate beyond any individual's lifespan[BO12, Ass10]. That is why to understand how culture can accumulate the model must let agents die while the domain persists. Then culture persists: it maintains the chain of discoveries, where individuals are simply there at the right time[Iba10]. As ideas communicate, some dying out with the generations, new space for ideas and norms emerges. This is the cultural evolution of creativity[BO12]. It can be measured by an accumulating complexity[Csi99, Mot25]. As the environment (new agents, preferences, memories) differentiates, so does culture[Csi99]. The use of the domain within these subcultures also becomes increasingly integrated. As culture becomes this evolving, reactive process, agents with a more up-to-date and encompassing understanding of culture would be revered, and it is something we attempt to check with a correlation test of age and communication [SB15]. Those agents that communicate most successfully may become gatekeepers of creativity. Lock-in can occur if a field becomes normative and stagnant in its creative variety[SG05, BS99]. On the other hand, dispersed communication can annihilate the hierarchy of gatekeepers and culture may become too fragmented for meaningful communication[BO12, Mot25].

4 Related Computational Work

The twentieth century saw a change in simulations as the use of it as a research tool infiltrated a plethora of domains (psychology, sociology, physics, etc.) [Axe05]. While it has many uses, simulations thrive at modeling societies and interactions: in games, in theory testing, in prediction[Axe05]. As the paradigm in creativity research shifted from cognitive models to social and distributed perspectives, there arose an interest in the use of simulation for modelling these creative societies[Mot25, SG01]. Models such as Gabora's *Memes and Variations*, Galanter's complexity-based artifact generators, and others emerged to examine the interactions between individuals that drive social creativity and its evolution through imitation and mutation of ideas[Gab95, Gal03]. Rob Saunders and his Digital Clockwork Muse were pivotal in showing that creativity itself can emerge from a model of the DIFI framework. Each agent is equipped with novelty-seeking mechanisms, which allows the population to produce novelty (mutation), evaluate others' novelty (selection), and update their own novelty-seeking mechanism (adaptation)[SG01]. Saunders created the architecture, but his research, while revealing, was constrained in its scale and in the interactions between agents[Mot25]. The scale was also limited in the works of Sosa and Gero who examined how roles such as "gatekeepers" emerge naturally from such artificial societies[SG05].

Motz filled this gap in order to better understand whether the basic essence of the DIFI framework was enough for all that Saunders and others revealed. Motz scaled up the agents, the interactions, and the society. Creative societies had never been able to reveal so much of their emergent complexity. Motz tracked these complexities in the form of clique formation, stylistic complexity, and knowledge transfers. These emergent features revealed the bridge between small-scale experimental setups and real-world creative communities.[Mot25]

As the foundational works, this paper uses many of the same features and mechanisms. By adding the theories of evolution and autopoiesis we are able to directly compare any changes. The main mechanisms inherited are the Wundt curve for novelty driven interest, kNN-based novelty detection over a per-agent memory matrix, and ResNet-based feature extraction[SG01, Mot25]; each is characterized in where it bears on this paper's additions (Section 5.2).

5 System Design

Section 1 argued for three changes to the inherited model; this section describes how each is implemented. Everything else is inherited, largely unchanged, from Motz's architecture[Mot25] — chosen because it was built for the large, emergent experiments this paper also needs — so that any effect we observe can be attributed to the domain and not to the rest of the system.

5.1 The Simulation Loop

A simulation runs n agents over t discrete steps, batched on the GPU through Motz's ParallelScheduler so per-step operations vectorise across the population[Mot25]. Each step is a fixed sequence of phases modelling one turn of the DIFI cycle. Inherited phases are listed plainly; those this paper rewrites or adds are marked.

1. **Update interest:** each agent refreshes its interest against its novelty buffer.
2. **Generation:** each agent renders a new artifact (the individual's variation); this is where an agent may draw inspiration from the domain.
3. **Evaluation:** features are extracted and novelty computed
4. **Individual evaluation:** the agent updates its interest given its own new artifact.
5. **Sharing:** if interest clears the agent's threshold, the artifact is shared with a random set of peers (the individual meeting the field).
6. **Interaction** *(modified):* recipients evaluate the shared artifact, accept or reject it for the domain, and return a reward signal — which also trains agent strategy (Section [5.4]).

7. **Boredom / retrieval** *(modified):* agents whose interest has fallen retrieve from the domain; retrieval is structured and strategy-driven rather than random (Sections [5.3]–[5.4]).
8. **Threshold update:** adaptive thresholds (boredom, novelty normalization, domain acceptance) are updated.
9. **Lifespan** *(new):* each agent ages, and agents reaching their lifespan die and are replaced (Section [5.5]).

The three marked phases are where this paper departs from the inherited model.

5.2 Inherited Substrate

See Motz [Mot25] for full detail.

Artifacts. Agents produce 32×32 RGB images rendered from a quaternion expression tree whose nodes are functions such as sin. cos, *swirl*, and *blend*; deeper trees yield more complex artifacts (Figure 2)[SG01, Mot25]. The trees are "genetic": agents breed them by depth-bounded mutation and crossover. Each artifact also carries metadata – creator_id, generation step, genetic parents, root_creator_id, lineage_depth, and a domain_parent_id set when it is bred from a domain retrieval. These fields are inert in the baseline model but become the substrate on which this paper's domain structure is built (Section [5.3]).

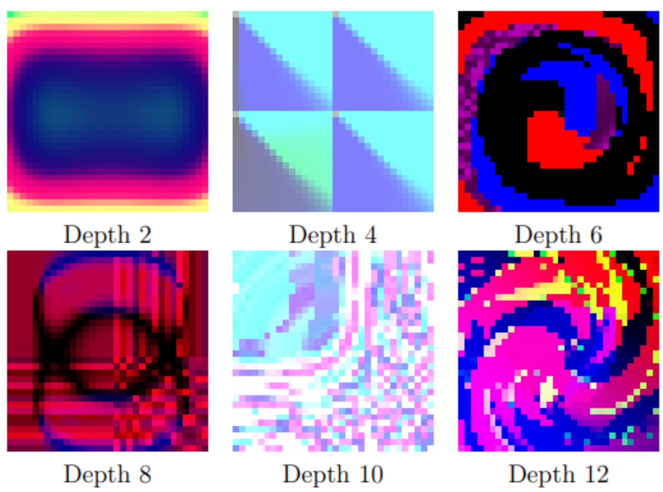

Figure 2: A example grid of artifacts (32x32 pixels) generated via quaternion-based expression trees at depths 2, 4, 6, 8, 10, and 12. Source: [Mot25].

Features. Artifacts are compared through a lightly modified ResNet-18 – a 3×3 first-layer kernel and a layer-2 output, tuned by Motz[Mot25, MAP13]. The result is a 128-dimensional L2-normalized vector. Features are projected to 16 dimensions by PCA before storage; all reported runs use this setting.

Subjective novelty. Novelty is per-agent. Each agent keeps a circular buffer (size 1000) of recently seen vectors and scores a query artifact by its mean similarity to the k=15 nearest, scaled by their spread and normalized across the population over a rolling window[Mot25, MAP13].

Sharing. The sharing threshold is now per-agent: an artifact is shared when its interest exceeds the 80th percentile of that agent's own recent self-interest scores, rather than of a population-wide pool. This locates the sharing decision in the individual as according to the 'I' in DIFI.

Motivation. Interest follows a Wundt curve: the difference of two cumulative Gaussians about a per-agent `preferred novelty` N(0.5,0.155)), rewarding artifacts that are novel but not too novel[SG01]. The computation is batched in GPU. One inherited detail is load-bearing for our additions: when an agent's memory is small, its aversion to high novelty is scaled down, so inexperienced agents tolerate more[Mot25]. This does little in a fixed population, but becomes essential once agents are born and die (Section [5.5]).

Additional Performance Plumbing. Several non-behavioural performance changes from Motz's ongoing work were merged such as rcompilation of expression trees into Reverse Polish Notation for vectorised rendering, and bilinear interpolation of the pretrained ResNet first-layer kernel[Bok24].

5.3 The Active Domain

The domain is implemented as a `Domain` class that supports three retrieval modes. These structures do not explicitly exist in the storage of artifacts, rather in their retrieval. The artifacts themselves are stored identically across all modes in a list. It is in the retrieval phase that artifacts are returned based on the mode-specific scoring system. An artifact must pass a threshold once shared to another agent to be deemed creative socially. This threshold requires the peer-evaluated interest to exceed a rolling window of interest generated by the 80th percentile of recently received artifacts across the field[Mot25]. This is adaptive, as the population's standards for creativity evolve. The percentile is strict enough that the domain can be considered selective (roughly the top fifth) and stable enough to prevent runaway acceptance or rejection[Jen10]. This affects only domain entry. The lowest structure represents the Motz model: a baseline domain with no structure, a pile of creative artifacts. Therefore, changes recognized in higher-structured modes are a result of how the domain is navigated not necessarily the content.

5.3.1 Similarity-Based Retrieval

Similarity mode is the first step in structure. It models how creators seek inspiration in what resembles the familiar, capturing the subjective, style-based associations of creative fields, like art [Csi99, Læs05]. The mechanism reuses machinery already in the system: candidates are ranked by the cosine similarity of their ResNet feature vectors – the very metric agents use for novelty – so an agent's sense of *stylistically close* when drawing from the domain stays consistent with its sense of *not novel* when evaluating. The relation is thus intrinsic (it translates the agent's own evaluative machinery) yet built on an extrinsic, observable property of the artifact: objective features, subjective interpretation of the features.

5.3.2 Lineage-Based Retrieval

Lineage mode grants a concrete, domain-controlled relation: an artifact is related to another by descent, requiring no subjective interpretation. Descent is traced through the domain_parent_id metadata rather than through agents' perceptions of lineage; the latter would leave holes in the tree and cost more to compute, and tracing through the domain keeps the relation under the domain's control. Relations are scored on a weighted ladder, figure 3: direct links score highest, and weight decays with distance.

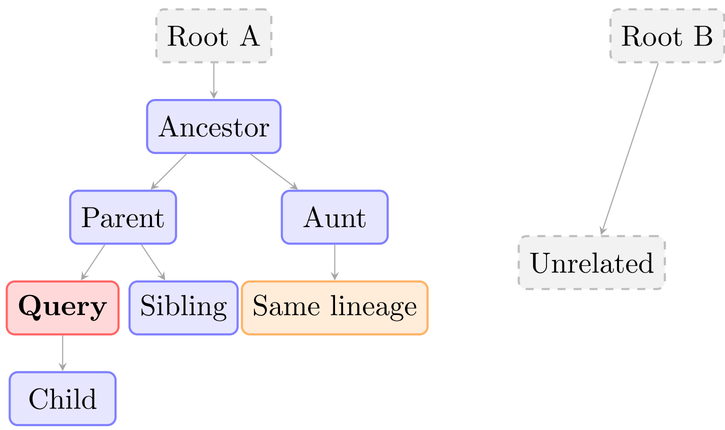

Figure 3: Lineage relations relative to query artifact. Directly lateral connections (parent, child) receive the highest scores and the other relationship types diminish in weight as they move further away. Arrows trace `domain_parent_id` relationships.

5.4 Agent Strategy and the Learning Rule

A structured domain is only useful if agents can navigate it; the retrieval modes return a ranked list, and where an agent looks in that list is its strategy. Each agent holds a `strategy_pref` setting that position — low for refinement (close artifacts), high for exploration — learned from two reward signals whose pairing is deliberate. A self-reward pulls the preference toward regions that satisfy the agent's own interest; a peer-feedback reward, returned when a shared artifact is accepted, pulls it toward what the field values. The two are kept inseparable, so navigation answers to both the self and the field.

5.5 Generational Turnover

Each agent is assigned a lifespan drawn from a uniform distribution and an `age` that increments every step; when age reaches lifespan the agent dies and is replaced, holding the population constant. The replacement is a genuine newcomer – a fresh globally-unique id (ids are never reused, preserving lineage integrity), an empty kNN memory, a newly sampled Wundt curve, and a reset strategy_pref. Because its memory is empty, the inherited experience annealing (Section [5.1]) lets it tolerate high novelty at first. With the flag off, the phase is inert and the model reduces exactly to the fixed-population baseline.

6 Methodology

6.1 The Domain Class

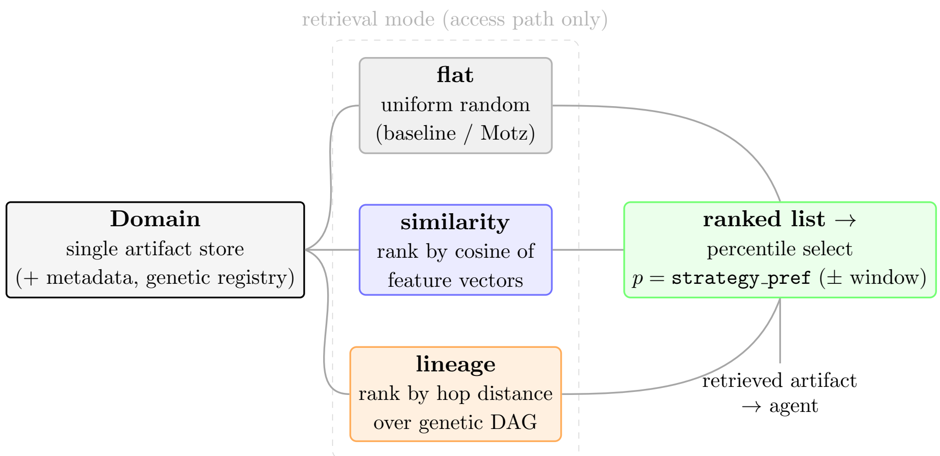

Figure 4: The three retrieval modes share one backing store and one selection policy; they differ *only* in how the candidate list is ranked. `flat` returns a uniform random member (the Motz baseline), `similarity` ranks by cosine similarity of ResNet feature vectors, and `lineage` ranks by hop distance over the genetic DAG.

The very relationships drawn from artifacts' metadata make up the structure, and these relationships are built in the class. In lineage mode, the ancestral links are added into the newly born domain artifacts' metadata, describing the genetic history of each artifact, and in similarity mode, the feature data is updated. The modes/structures are a property of the access, not the storage. These modes only differ in their path of retrieval. This feature is what makes them comparable in the forthcoming experiments.

6.2 The Math Behind Similarity Mode

Similarity mode represents a relative style between artifacts. The meaning in an artifact that the mode draws on is its feature vectors. Specifically, this mode compares the L2-normalized (unit) feature vectors between two artifacts. If the direction encodes the style, then cosine similarity measures stylistic relatedness. Mathematically, this is done by considering two feature vectors c (candidate) and p (query). Cosine similarity is then

$$\cos = c \cdot p \in [-1,\, 1].$$

This value is then converted into a positional coordinate via

$$percentile = \mathrm{clip}\left(\frac{1-\cos}{2},\ 0,\ 1\right).$$

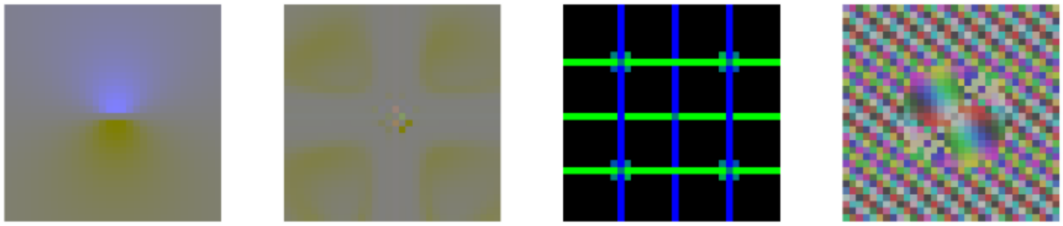

Figure 5: An example of pairs of artifacts with high vs. low cosine similarity. The left pair have a cosine similarity of 0.995 while the right pair have a similarity of 0.494.

The positional coordinate is used in building the weighted list that agents retrieve from. Interpreting the math as a position on the list, one can say a candidate artifact with a position near 0 represents a stylistically identical artifact (refine) while a candidate near position 1 represents a stylistically orthogonal vector (novel).

6.3 The Math Behind Lineage Mode

The source of lineage mode is the parent-child relationship. Every familial label can be defined this way; for example your uncle is just your parents' parent's child. Although, that 'child' has to be a different person than your parent, and a simple check of an agent's ID can capture this. Building on this idea, there is a global registry

$$\texttt{Artifact.genetic\_parents} : id \to (parent1,\ parent2)$$

that maps each artifact's ID to its parents' IDs. This is stored and grows throughout the experiment, and each domain-added artifact uses this global registry to build its meaning in the context of culture i.e. Aunt, Grandpa, etc. When an artifact is added, it finds its nearest domain relatives by walking up this registry. The result is 3 mappings; a mapping up, a mapping down, and a depth measure:

$$\begin{aligned} \texttt{domain\_ancestors[id]} &\to \text{frozenset of nearest in-domain ancestor ids (up-edges)} \\ \texttt{domain\_children[id]} &\to \text{ids that name id as an ancestor (down-edges)} \\ \texttt{domain\_depth[id]} &\to 1 + \max(\text{ancestor depth}), \quad 0 \text{ for roots} \end{aligned}$$

This adjacency is calculated once at artifact insertion and never recomputed. Furthermore, the number of ancestor nodes is 2^8 or 256, limiting any computation cost for genealogically rich artifacts. The registry is global rather than per-agent because relative position must be measured over the full genealogy: without it, an artifact one breeding step removed and one a hundred steps removed would score identically whenever both share the same nearest domain relative.

These nearest domain artifacts give the relative positioning of the query artifact. One can assume a newly bred artifact is not yet in the domain, and so it has no node in the domain ancestry graph. The domain can also be assumed to store many artifacts, in a run with 100 agents by step 100 there can be thousands of domain artifacts so figuring out where a lone artifact sits and then how it relates has to be simplified to save time. Each artifact also carries its domain parents IDs, and from these, an artifact can be pinpointed in the structure. Each domain parent is seeded at a distance of 1. The distance is important to keep track of as this is our relational measure, and the radius of search is set to 8. Any node exceeding this distance is considered unrelated. Each level of BFS corresponds to an increase in distance, and as the BFS assigns these distances, we develop relationships. The retrieval mode grabs from a weighted list of artifacts that share a greater connection under the structure; for lineage, this weighted list is made up of a distance score,

$$score = 1 - \frac{\min(d,\, unreached)}{unreached},$$

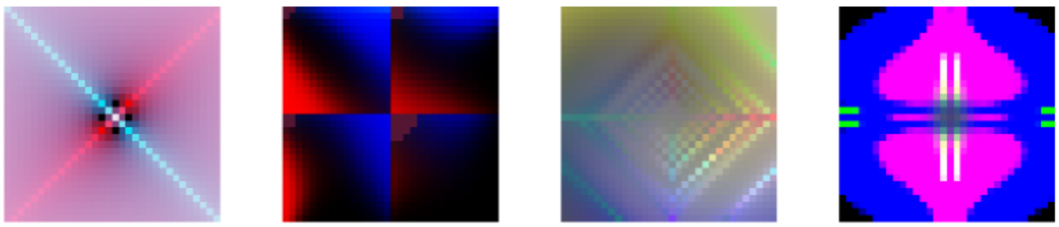

Figure 6: An example of pairs of artifacts with high vs. low lineage hops. The left pair are a difference of 1 lineage hop, while the right pair is a difference of four. Although they seem in stark contrast across one can see the traces on the left pair. The first image shares the colors and 4 quadrants of the second. It seems as though the first image is pulling the second image to the center. The second pair have very little in common other than the horizontal streak across the middle and a general diamond-like shape both of which have morphed.

where 1 is adjacent and 0 is unrelated. For logging purposes, these distances have been converted into the recognizable terms for family where: distance = 0 is self, = 1 is parent or child/adjacent, and = 2 is sibling or grandkin. These are considered close relations and are logged as such, as are mid relations and far. This way we can track types of relationships, in kinship terms, that are most prominent across modes.

6.4 Structure Revelation

All three modes converge on this selection policy, so its process must remain the same but robust. The ranked pool/list of size n is addressed by a target percentile, $p \in [0, 1]$, as seen below.

```
target_idx = round( p * (n - 1) )
window     = max(1, round(W * n))    # W = STRATEGY_WINDOW = 0.05  (LINEAGE_WINDOW =
    0.10)
selected   = uniform_random_int(target_idx - window, target_idx + window)
```

This list shares the same $[0, 1]$ geometry as the novelty score that each agent evaluates creativity with. The idea is that a position in the list is isomorphic to novelty in some form. An artifact can be novel in terms of similarity and also genetically, and it is likely these are often quite similar measures. Furthermore, the reason that the preference is turned into a window around that percentile is twofold. First, this injects some controlled exploration noise into the experiment so the policy is

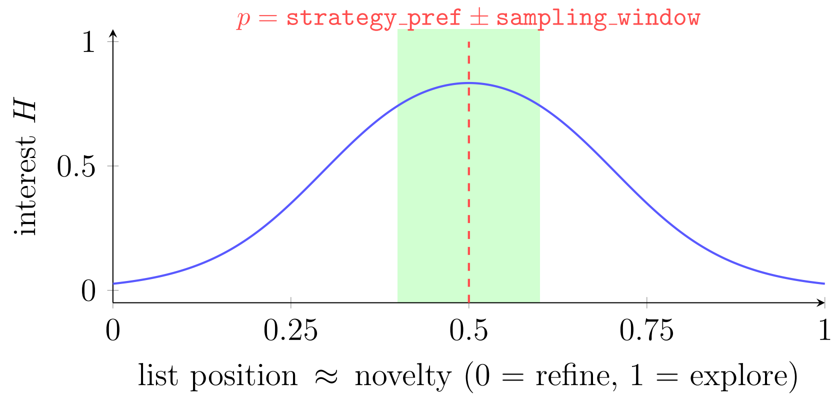

Figure 7: An agent's learned `strategy_pref` sets a target percentile p, and an artifact is drawn uniformly from a window $p \pm W$ (with $W =$ `STRATEGY_WINDOW` $= 0.05$, or 0.10 in lineage mode).

not so deterministic. Considering that the policy is learned, any noisy exploration is beneficial in learning robustly. Second, this serves as a fallback in case the candidate pool is small. The window represents a fraction of the pool, so any steering is proportional as the domain grows. In lineage mode, the window is increased to 0.10. This is because all equal-distant artifacts are considered equal so to capture the sampling from a relationship such as 'adjacent' a larger window is better. Also, it takes a while for a rich lineage to develop, so a wide enough fallback window is essential for early simulation stages.

6.5 The Math Behind Strategy

Agents remember the last position from which they sampled, giving the ability to adapt their preference based on past results. The strategy preference also maintains a learning rate for the update rule. The learning update exists as a single method on the Agent class:

`update_strategy_pref(outcome_interest, reference_interest)`.

This method works as follows: when an artifact produces an interest above the agent's average interest, the preference value is nudged towards the position from which it most recently sampled. If the generated interest is below the average, the preference is nudged away. Mathematically [HL18]:

$$\begin{aligned} \Delta &= \text{strategy\_lr} \times \text{reward} \times (\text{sampled\_position} - \text{current\_pref}) \\ \text{new\_pref} &= \text{clip}(\text{current\_pref} + \Delta,\ 0,\ 1) \end{aligned}$$

The sampled-position-minus-preferred-position term is what moves the update in a certain direction. This term also self-stabilizes: the closer the preference is to the agent's sampled position, the smaller the update, preventing oscillations. Extreme cases are prevented by clipping the updated preference value, preventing runtime errors.

Because sampling is drawn from a narrow window around `strategy_pref`, the update term $(\textit{sampled\_position} - \textit{current\_pref})$ is both small and randomly signed. An ϵ-greedy exploration term and mixed preference initialization are added for the preference to move measurably. The $\epsilon = 0.2$, where the retrieval is drawn uniformly. Higher values were avoided as they begin to swamp the signal the preference is trying to settle on. The mixing is done by drawing the initial preference (default 0.5) from a uniform distribution, which can be manipulated to start with more refining agents or more exploratory agents. While the strategy preferences may be vastly different, the update function remains the same. This gives us a constant baseline in a homogeneous start, with exploration in the effects of heterogeneous and unequal (between exploring and refining agents) starts

The reward signal of this update fires in two locations during a single iteration: in the self-evaluation phase and the peer-feedback phase. The reward in terms of interest it generated has to be relative. Some artifact may return a low interest, but if it is still more interesting than what usually comes through, this should be seen in the formula. For the reward signals in the personal evaluation phase, these reveal how interesting an artifact is to the agent. This is made relative to the agent's running baseline which is their average interest levels:

$$\textit{reward} = \mathrm{clip}(\textit{outcome\_interest} - \textit{reference\_interest},\ -1,\ 1)\,.$$

This reward signal also fires when an agent breeds artifacts. The code calculates the structural position of the breeding pair within the agent's memory. The structural position is found using the same mechanics as the structure mode that the experiment is running in. In lineage mode, the position is found from the hop distances over the domain DAG divided by the search radius. Similarity mode returns distances from the cosine similarity of feature vectors. Flat mode simply returns None. None here implies a structural position that cannot be measured. This could also be an unanchorable partner in the lineage or missing features in similarity.

The reward signal for a shared artifact depends on whether the artifact shared got accepted into the domain. This formula uses-rather than an agent's own average interest- an `acceptance_rate_ema` value to set a baseline for acceptance into the domain. See below:

$$\begin{aligned} \textit{advantage} &= (1 \text{ if accepted else } 0) - \textit{accept\_rate\_ema} \\ \textit{accept\_rate\_ema} &\leftarrow \beta \cdot \textit{accept\_rate\_ema} + (1-\beta)\cdot \textit{accepted} \qquad (\beta = \texttt{accept\_rate\_beta} = 0.9) \end{aligned}$$

The beta value is an exponential moving decay factor of the acceptance rate. These reward signals are made comparative because, if acceptance is common, a binary reward is constant regardless of which region was sampled. This scenario would carry no information about whether the current search region is more novel and the signal would just drift. Subtracting the agent's own moving acceptance rate means the reward is only positive if the retrieval zone beats the agent's usual success in sharing. This way, the measured performance of a sample region in any mode is carried over into the sharing reward signal. Also, this setup allows the simulation to track the population's acceptance rate over time and incorporate it into the measurement. When a strategy becomes common, its acceptance rate normalizes; the advantage decays, discouraging the collapse of agents on a certain region. This is a natural frequency-dependent pressure that allows navigation to become a genuinely adapting search rather than a constant drift. This consideration is noticeable in the baselines used for each of the reward signals described. That they are then fed into the same strategy update formula is what makes this method adaptive and meaningful to the design of our system.

6.6 Generational Turnover

The mechanism is its own experimentable condition (toggled via flag `--agent_lifespan`, which can also accept an integer for some specific lifespan. Our experiments uses `min_age` 40 and `max_age` 100 for a 500-iteration experiment. This makes about 7 generations, staggered around the boundaries of lifespan, giving rise to enough information necessary for our analysis. The heterogeneity of these lifespans is deliberate, as a uniform age at birth and death would mean each cohort building the culture from scratch. The staggered lifespans ensure a mixing of 'newborns' and 'elders' at all times, allowing us to explore age difference dynamics. The lifespan phase itself runs after threshold updates so a death never corrupts the batched tensors (for parallel processing of many agents' phases at a time) midstep. The agent-index map used to create batches is refreshed after every replacement. When lifespan is turned off there are no operations and the baseline is as it was.

7 Measurements

7.1 Richness

A both complex, diverse and large cultural memory is interpreted to be the result of a successful memory structure. These metrics are calculated per step; a snapshot of the domain artifacts `tree_depth` and node count are all we need for complexity while a simple counting of domain artifacts reveals its growth. As more nodes are used in combination at greater depths it becomes less and less obvious at a glance the functions that went into making the artifact. Using the node counts' standard deviation (std) is how richness is measured. This way we not only see how complex the domain artifacts become, but the spread of that complexity. A domain with many similar artifacts, whether high or low complexity, reveals a less rich domain through std. On the other hand, a small domain with many different complexities returns a high std. The std is exactly the "differing levels of artifact complexity" metric.

7.2 Stylistic Diversity

Considering that the computation cost would be $O(sample^2)$ we limit the sample to 128 artifacts. As the domain grows into the thousands and exceeds this value, the 128 artifacts are randomly selected, potentially leading to some volatility in our measurements. The unit vectors of these artifacts' features make up the rows of a matrix (M) which is then used to form the similarity matrix

$$S = M\,M^{\top}.$$

Then from the similarity matrix we take the means of the off diagonal entries via

$$\frac{S.\mathrm{sum}() - \mathrm{diagonal.sum}()}{m \cdot (m-1)}.$$

This captures the average style similarity between distinct pairs, ignoring the similarity between a vector and itself. The dispersion is calculated to be 1 minus this mean of off-diagonal similarity. The higher this value the more feature vectors are pointing in different directions, and it is interpreted to imply many styles. The lower the more the styles are clustered. This is calculated and averaged over seeds of experiments. Tracking this value as it rises or falls may also reveal whether the system is able to sustain novelty or if novelty converges. This is implied because novelty measure also uses cosine similarity, so more diversity would mean more styles of novelty explored (more sustainable).

7.3 Lock-in

Rather than the sustainability of styles, lock-in refers to the ability for other paths of creation to exist. Where the styles may be quite similar, measuring path dependence allows one to decide whether there is a single type in that style, or whether there are many paths for inspiration within a style. Path dependence is measured over the lineage graph, whose nodes are artifacts and whose edges are breeding relations. During the run, snapshots of the roots `roots[root_creator_id]` $\rightarrow$ *count* create a histogram of distinct roots over time. Then

$$root\_concentration = \frac{\max(counts)}{n}$$

to show us what fraction of the domain has come from the biggest lineage. 1.0 would mean that one lineage has taken over whereas the closer to 0.0, the more lineages there are that coexist. As was mentioned lock-in can also be revealed in a slowing down of what's accepted, growth stalling under the closed-mindedness of culture. Hence, we track every share event as either accepted (1) or not (0) and average the acceptance over 20-step windows. From this we can generate a curve showing the acceptance rate over time. A sharp dive implies the domain's growth is closing off and a rise might imply a creative breakthrough. Lock-in can then be diagnosed by either a rising `root_concentration` and/or falling acceptance rate.

7.4 Navigation and Network

To understand what navigation is actually occurring our metrics look at the strategy preference value. Pooling the values over the agents and tracking the mean, std, drift, and range allows us to understand how strategy is evolving. The key derived number is drift as it shows us how much the population changes its strategy from the starting point. The spread (std) and drift together are evidence for agents actually differentiating their navigation rather than staying uniform. To figure out if these strategies are actually drifting towards anything meaningful, we track their acceptance. Every shared artifact keeps (in the metadata) the position its domain parent was sampled in. Then we can compute the acceptance rate (same method as seen before) for each position bucket. The spread is the max - min of these bucket's acceptance rates, and this reveals whether there is actually some dependence between where an agent samples and their success. If the spread is greater than 0.10 it can be said that there is some learnable gradient for a strategy learner to climb. We use the Wilson confidence interval to calculate the proper confidence bounds for bucket proportions, as this is better than trusting a raw sample/total measurement, especially for smaller buckets(low sample).

The method we use is the same that Motz uses to provide a good image for comparison on what has changed. Another CSV is made to track the network of interaction, logging every share (edge) between two agents (nodes). By putting the threshold of tracking interest-generating shares (only shares that generate above a certain level of interest) to be the same as the domain threshold, we can observe the network of shares that actually built the domain. The network is generated over time too, every 100 steps, with a stride of 50 (50% overlap of shares); a frame is captured to visualize a fluid change. This data is put into the same Louvain algorithm Motz uses to record the communities and then modularity Q to measure group strength. Motz set a goal of a Q-value greater than 0.3 to assert clique formation, and we adopt the same goal. Understanding under what contexts- structure mode, strategies, lifespans- give higher amounts of group formation is important to understanding the social aspect of emergent creativity. The Q value is output as a curve over time to show the growth of communities. Usually, by the end of a simulation every agent has interacted, so this is essential for revealing the stages leading up.

Because agents of different generations cannot communicate, whole-run modularity would separate cohorts as clusters by construction, so network evaluations are computed over overlapping windows instead. In addition, we have become interested in age-difference communication, or more specifically the communication across generations. This way we can observe how communication adapts with lifespan, and whether inter-generational agents communicate more as the torchbearers of cultural memory. We use the same communication graph as before, where interest generated is the weight of an edge. Separately, we calculate the betweeness centrality for every agent. Using a Pearson r calculation for linear regression averaged over the seeds between the betweeness centrality and agent age, we get our metric. A positive r value; the higher it is, the greater the correlation between older agents and being a generational medium.

7.4.1 Domain Network

Community detection is still a topological measure, and the topology of these experiments runs on random wiring. The domain, by contrast, is not randomly wired: its retrieval relationships already constitute a graph. Where previous experiments had the domain as a random pile, our domain is full of relationships. These can build a new network graph, from a perspective within the domain. The domain is practically already a network: in lineage mode the ancestor DAG and in similarity, the nearest neighbor graph over features. Our domain retrieval already logs the non-random edges at insertion (nearest-neighbour in similarity, nearest in-domain ancestors in lineage, remembered domain parent in flat), where our domain artifacts are the nodes. Connected components, largest-component fraction, isolated-node count, and Louvain Q are computed on this graph. Because single-nearest attachment yields tree-like graphs, and trees score near 1 on modularity by construction, each Q is compared against a density-matched random null; only the observed-minus-null gap is interpretable.

7.5 Homogeneous vs. Heterogeneous Start

In regards to the experiments for a homogeneous vs. heterogeneous start, all these metrics are observed as their resulting richness, complexity, and field dynamics are to be compared. The baseline experiment, if not specified, is using a heterogeneous spread, as this seems to capture the greater portion of reality and give the agents a deeper idiosyncrasy.

We look only at the lineage runs for our skewed starts and include the homo start in the metrics too. This is because lineage is to represent our highest structured mode, using flat mode would make no sense in terms of our exploration of strategy. We are also more interested here in the strategy's influence as opposed to the structures so we keep structure constant.

7.6 Success Criteria

The metrics above are tied to the research questions through the explicit criterion referenced in Table 1. Each is read against the "flat baseline" structure which uses random domain artifact retrieval.

Question	Metric	Criterion for success	Compared to	Outcome
RQ1 — Growth	Domain size over time	Sustained, non-stalling growth in the count of accepted artifacts	Flat baseline	**Met** in every mode; the criterion does not discriminate between them
RQ1 — Richness	Node-count std.	Higher complexity spread than baseline; rising or stable, not collapsing	Flat baseline	**Not met**. The spread between modes (0.417) is smaller than the largest between-seed deviation (0.881)
RQ1a — Stylistic diversity	Feature dispersion	Higher and/or sustained dispersion (does not converge toward 0)	Flat baseline	**Partly met**. Dispersion falls from ≈0.25 to ≈0.22 and then holds; similarity retains the most (0.2279)
RQ1b — Lock-in	root_concentration; binned acceptance rate	Concentration stays well below 1.0 and acceptance does not sharply collapse	Flat baseline	**Met**. Concentration plateaus at 0.10–0.16 and acceptance holds near 0.20
RQ2a — Generational bridging	Pearson r(age, betweenness centrality)	$r > 0.2$ (older agents act as cross-generational bridges)	Lifespan runs	**Met** in all modes (0.355–0.395)
RQ2b — Navigation differentiation	strategy_pref mean, spread, within-lifetime drift	Populations with different starts diverge in preference rather than retaining their initialisation	Uniform start	**Not met**. Skewed populations move −0.001 and +0.010 from their own initial means and do not converge
RQ2 — Learnable gradient	Acceptance spread across position buckets (Wilson-bounded)	Spread > 0.10, indicating a gradient to climb	Within mode	**Not met**. 0.024 (similarity) and 0.023 (lineage)
RQ2 — Community structure	Louvain modularity Q over time	$Q > 0.3$ (clique formation, per Motz)	Flat baseline / Motz	**Not met**. $Q \approx 0.247$ in every mode

Table 1: Success criteria linking each research question to its metric, the threshold that counts as a positive result, the comparison condition, and the outcome. All curves are averaged over five seeds; structured modes are judged against the flat baseline except where a within-mode or lifespan-only comparison is the natural control.

7.7 Reproducibility Statement

All runs are reproducible: seeds are propagated through Python, NumPy, and PyTorch, fixed per configuration, with exact CLI invocations and runner scripts in the repository. Every inherited parameter (Wundt curve, kNN settings, mutation rate, share count, thresholds) is held constant unless noted, so observed differences trace to this paper's modifications rather than to retuning .Unless stated otherwise, all below reported values are means over five seeds.

[1]

8 Results

8.1 Richness Results

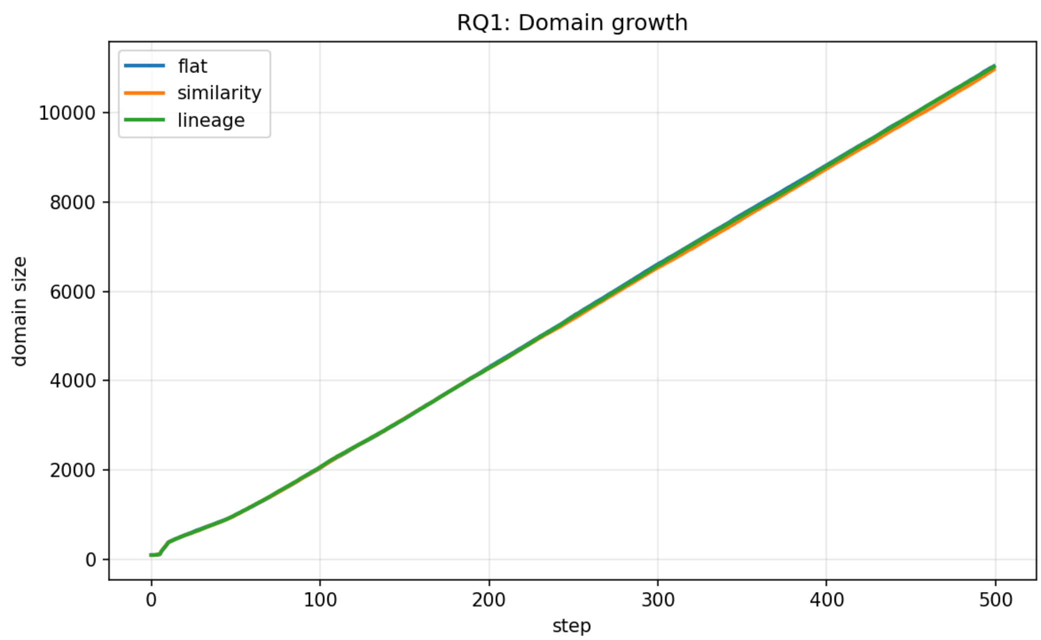

Figure 8: Domain growth, as measured by the number of artifacts over time steps. The purpose is to visualize the domain's substance across different structures.

Fig. 8 provides a visual of domain growth for each mode. Each mode/domain structure reveals a near identical pattern of growth. Each ends up with a domain size of $\approx 10,500$ artifacts over 500 steps (averaged across 5 seeds). This implies that per step the 100 agents are adding ≈ 21 artifacts to the domain with each step. There also appears to be a relatively linear growth with hardly any interference in the inflow of the domain from beginning to end.

Figure 9 captures our measures for artifact complexity, as seen by the artifact's expression tree depth and node count. Combined these give us an idea of the true magnitude of the expression tree

[1]Repository link is supplied with the submission

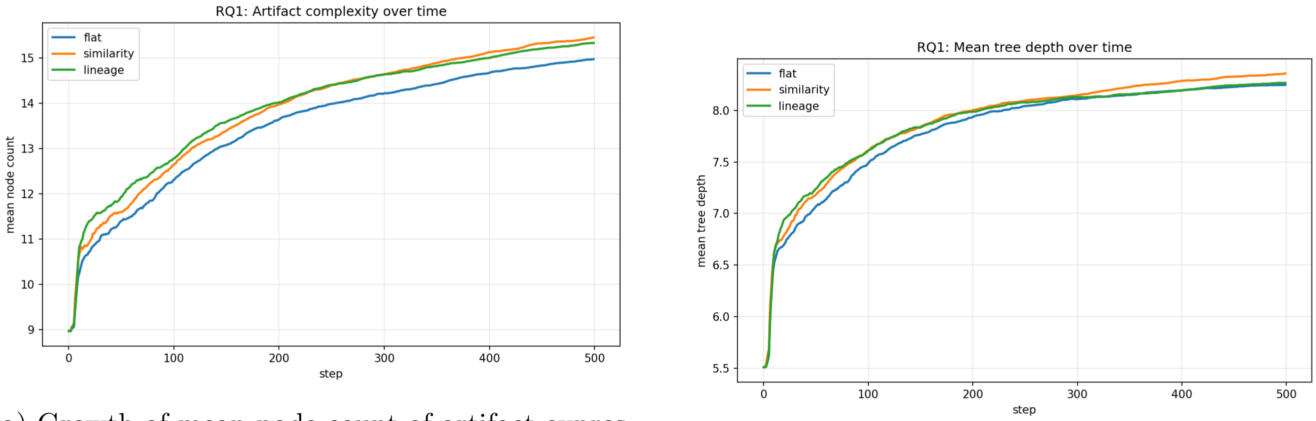

(a) Growth of mean node count of artifact expression tree over time spread across the three domain structures.

(b) Mean tree depth of artifacts' expression tree over timestep contrasting the three domain structures.

Figure 9: Together, these measures track the complexity of the artifact over time. Specifically, the complexity of the artifact's expression tree considering the magnitude of depth and width.

over time. Beginning with the node count in Fig. 9a all structures grow relatively in tandem with one another. There is a noticeable gap between the flat mode and 'more structured' modes that widens over time, capping at the end with flat having an average node count of 15 and similarity and lineage with a node count of $15.3 - 15.4$. Amongst the structured modes, it seems that lineage has the strongest start, but then is overtaken by similarity at around step 200. This occurs in Fig. 9b as well, where lineage mode is headed out by similarity by the end of the runs. Lineage mode finishes on par with flat mode (with an average tree depth of $8.2 - 8.3$) while similarity finishes noticeably higher (depth of 8.4). Furthermore, both artifact complexity measures show a stark increase in the beginning stages of the simulation. All curves are monotonically increasing; tree depth flattens after $\approx$ step 300 while node count continues to rise

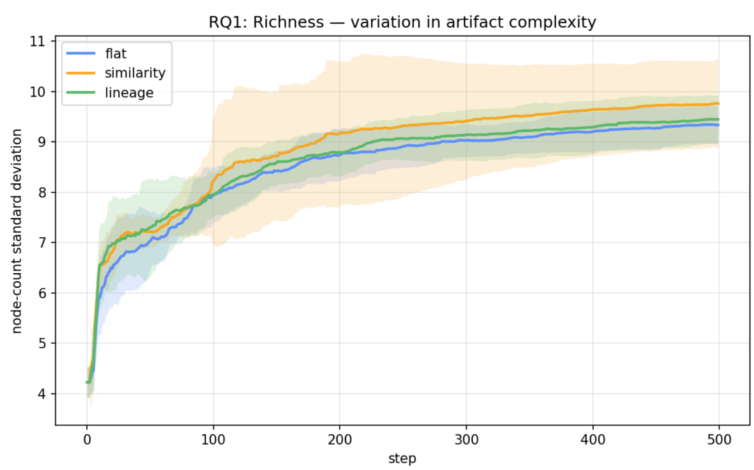

Figure 10: Standard deviation of node counts across all domain artifacts, recorded each step and averaged over five seeds. The line is the spread of complexity within a run; the shaded band is ± 1 standard deviation across seeds. The three modes end within one another's seed bands.

Richness, the standard deviation of node counts across the domain, is given in Figure 10. All modes begin at 4.22 and rise to settle over the final fifty steps at 9.32 for flat, 9.42 for lineage, and 9.74 for similarity. The spread between modes is 0.417 while the largest standard deviation between seeds is 0.881, so the modes cannot be separated on this metric. Similarity carries the widest seed variation of the three at 0.881 , against 0.37 for flat and 0.47 for lineage.

8.2 Stylistic Richness Results

This metric provides a further glance at the richness of creativity in these domains by measuring stylistic richness.

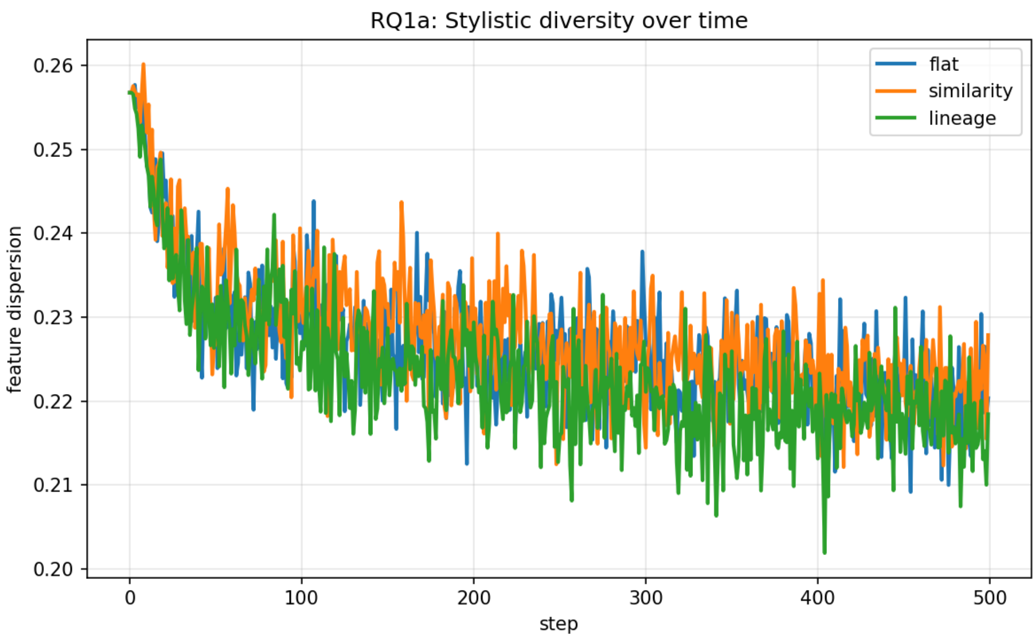

Figure 11: Domain stylistic diversity, measured by the cosine similarity of artifacts sampled from the domain. The purpose is to visualize the evolution of styles during creative evolution.

Figure 11 reveals a noisy graph of our feature diversity metric. Across all structures, there is a noticeable decrease in feature diversity, moving from the 0.25 range to the 0.22 range. Numerically, similarity diminishes the least with a final average diversity score of 0.2279, followed by flat mode with 0.2204 and lineage with 0.2184.

8.3 Sustained Novelty Results

To further explore the richness of creativity in this domain, we also look at the sustained creative lineages throughout the simulations. We track the number of distinct root nodes (artifacts with no domain parents, the first of their kind) with the different structured domain modes.

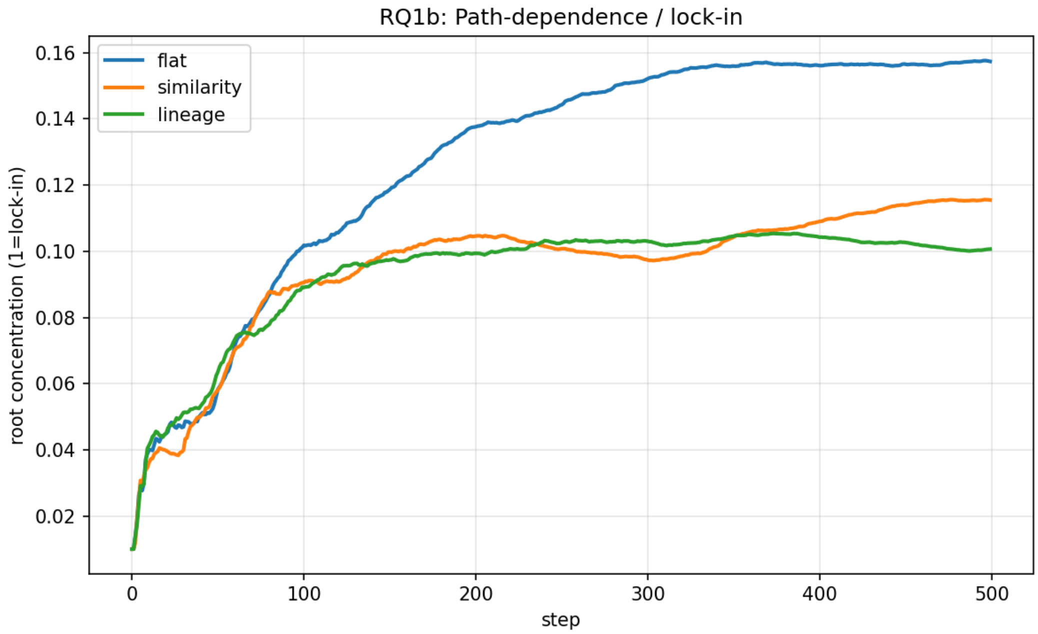

Figure 12: Domain root concentration. This tracks what percentage of the domain is owned or produced within the largest lineage. 1.0 implies 100% is owned by a single lineage of domain artifacts.

Figure 12 reveals the history of lineage. Flat remains noticeably higher than the structured modes, with a concentration of 0.16. Similarity is the second lowest with a concentration of 0.12, and Lineage mode, with the most dispersed roots, has a concentration of 0.10. Root concentration with the number of distinct root nodes (all modes starting with 100 distinct roots per agent): flat has 234 distinct roots by the end, similarity comes in with the lowest of 230, and lineage with the highest of 241.

8.4 Generational Turnover Results

We compare Figure 13 to what we have seen of the artifact complexity and the growth for simulations with a lifespan. The domain growth of lifespan runs ($10,500$ artifacts) far exceeds the no-lifespan runs ($7,000$ artifacts). On the other hand, the node count and tree depth remain about the same. With lifespan off, flat mode becomes the new leader of these complexity measures, whereas in lifespan runs the structured modes lead. In addition to these measures, we also use Figure 14 to observe the evolution of style. All our modes seem to decrease uniformly from the feature dispersion of lifespan runs (≈ 0.22) to a new low of ≈ 0.20 in no lifespan runs. Also, the flat mode becomes the worst (by a margin) of the modes in terms of stylistic diversity. Similarity remains, having the best dispersion of vectors in the feature space.

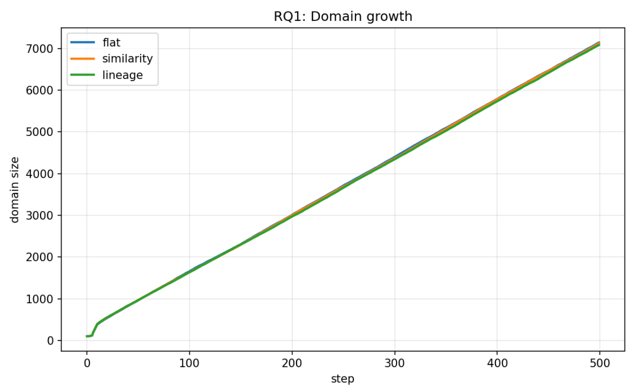

(a) Growth of domain size as the number of artifacts over time.

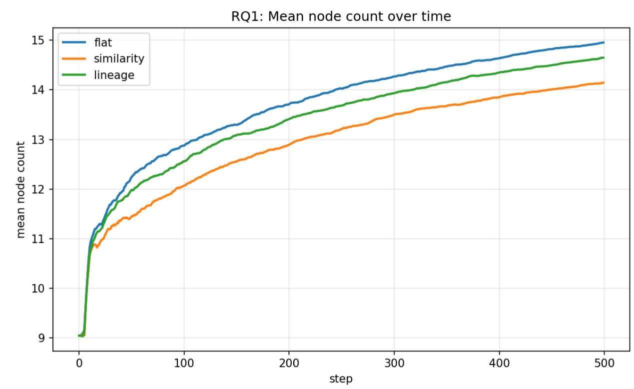

(b) Mean node count of artifacts in domain over time.

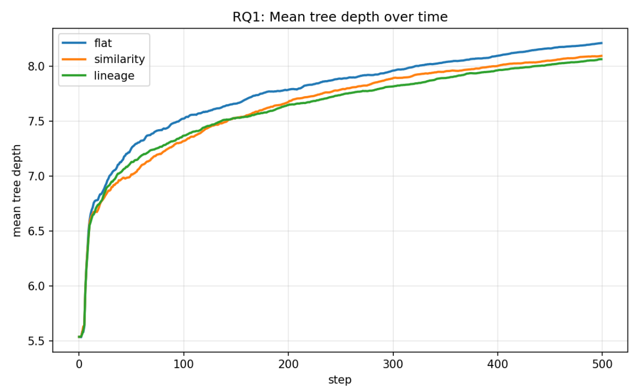

(c) Mean tree depth of artifacts' expression tree over timestep.

Figure 13: The internal artifact richness and general growth of the domain when there is no lifespan.

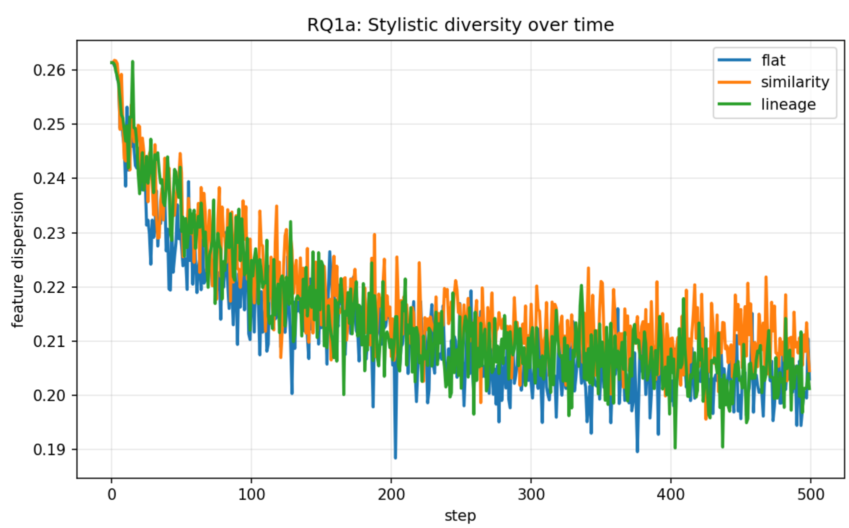

Figure 14: Feature dispersion of feature vectors sampled from all domain modes.

Finally, we observe the richness of heritage in simulations with no lifespan, as shown in Figure 15.

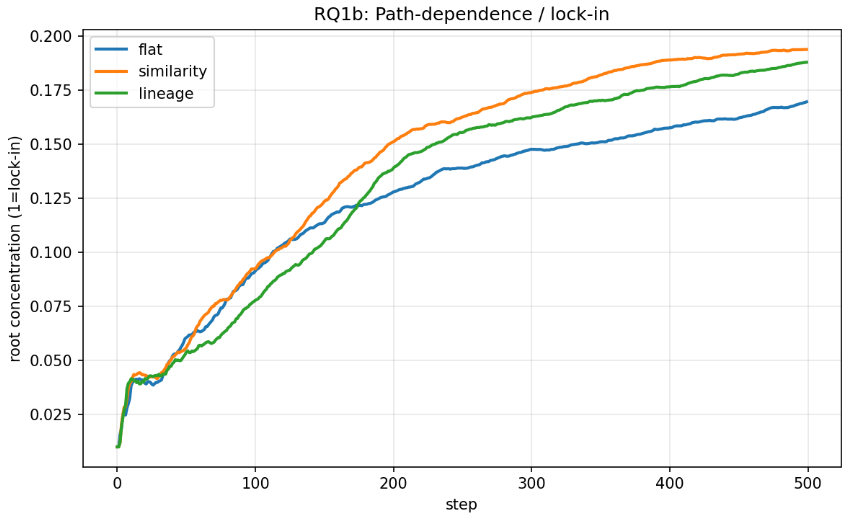

Figure 15: Root concentration of domain artifacts over time.

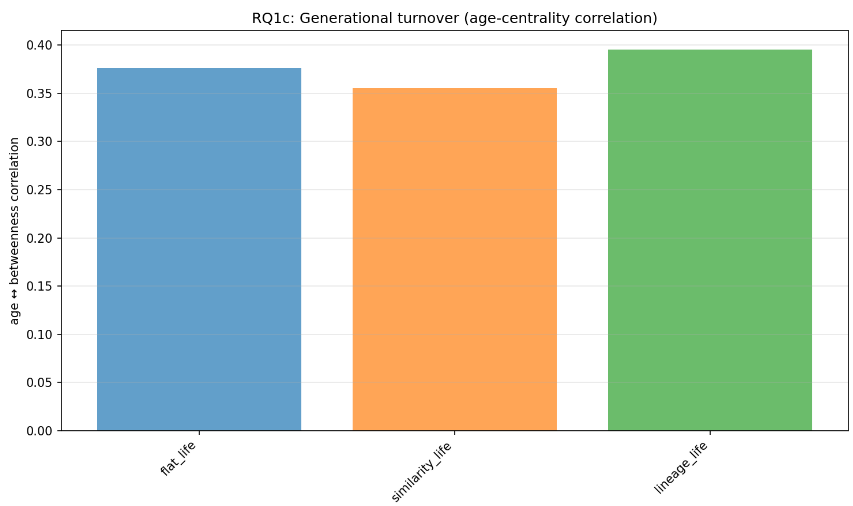

Figure 16: Correlation figures between age and centrality in sharing. Each structure is represented by a bar, and lifespan is turned on for all.

All structures have been shown to share a positive correlation between age and centrality. They pass the test of correlation (0.2), with lineage mode at the forefront (score of 0.395). There does seem to be a slight effect of the structure of this centrality factor: flat mode in the middle (score of 0.376) and similarity at the bottom (score of 0.355).

8.5 Navigation and the Acceptance Landscape

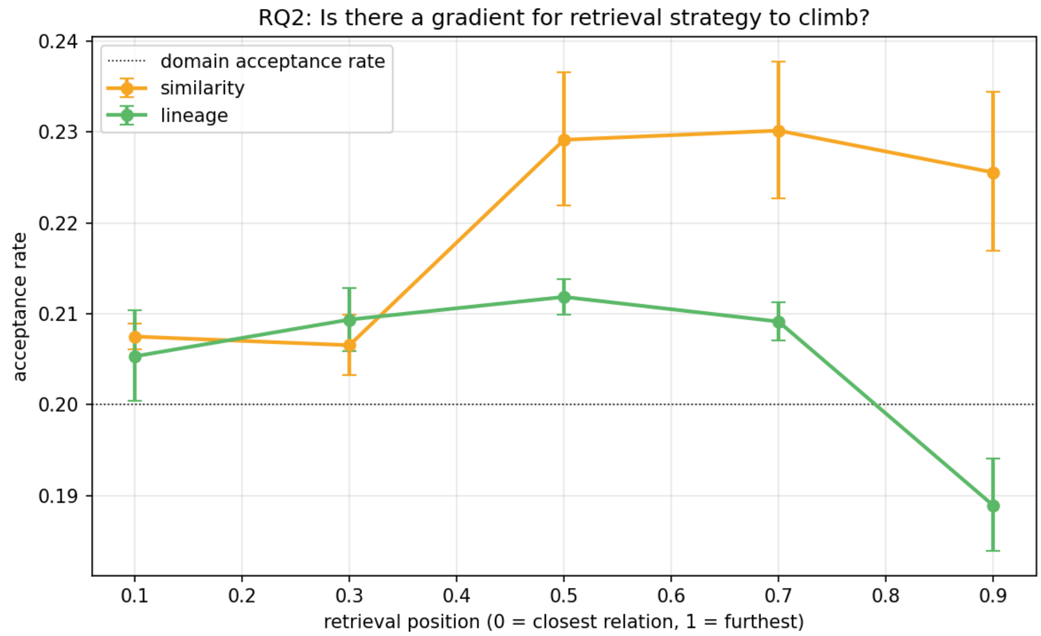

Figure 17: Proportion of shared artifacts accepted into the domain, grouped by the position in the ranked candidate list from which their parent was retrieved, where 0 is the closest relation and 1 the furthest. Pooled over five seeds; error bars are 95% Wilson intervals; the dotted line marks the domain acceptance rate of 0.20. Flat mode is absent because it records no retrieval position. N = 415,665 (similarity) and 409,635 (lineage).

Figure 17 gives the acceptance rate of shared artifacts against the position their parent was retrieved from. In similarity the rate rises across the buckets from 0.207 at the closest to 0.226 at the furthest. In lineage it peaks at 0.212 in the middle bucket and falls to 0.189 at the furthest. The spread is 0.024 for similarity and 0.023 for lineage, both below the 0.10 threshold of Table 1, so the criterion is not met in either mode. With upwards of four hundred thousand linked shares per mode the Wilson intervals are narrow enough that these differences sit outside them, but they remain small in absolute terms. The retrievals are not evenly spread across the positions: in similarity 78% fall in the closest bucket, whereas in lineage the bulk sits in the middle two.

Strategy preference is reported in two ways. Measured as the displacement of the population mean from its own initial value over 500 steps, flat moves -0.017, similarity +0.037, and lineage -0.022. Measured within a lifetime (Figure 19a), the mean preference of the oldest agents differs from that of the youngest by +0.046 in flat, +0.167 in similarity, and +0.018 in lineage, with correlations between age and preference of +0.018, +0.163, and -0.012 respectively. The spread of the population widens slightly over the run in the two structured modes, from 0.248 to 0.259 in similarity and from 0.247 to 0.265 in lineage, while flat holds at 0.230 to 0.228.

8.6 Hetero vs. Homo Initialization Results

From Figure 18, we can see that both initializations retain a similar shape to their initial starting positions. From both, there is a trickling spread towards the edges. The homo and hetero start (0.493 and 0.492 respectively) shift ever so slightly to the conservative side. The standard deviation for both is relative to where they started (homo STD = 0.018, while hetero STD = 0.240), and both reflect the increase in spread over time. In terms of actual change, the drift measured for both is tiny, with 0.007 for homo and 0.008 for hetero.

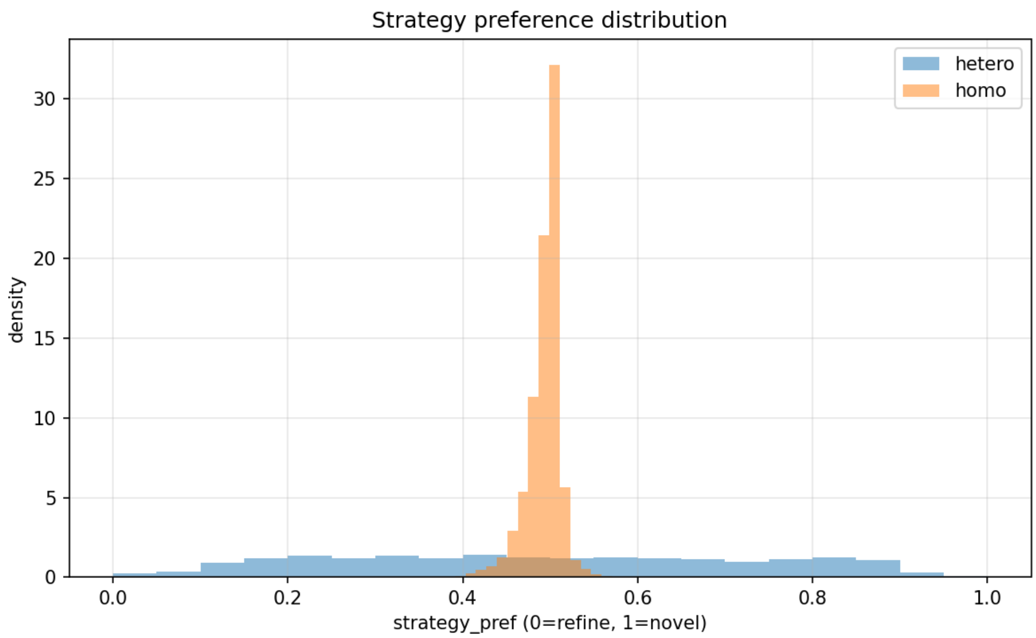

Figure 18: Strategy preference distribution averaged over modes and compared between a homogeneous vs. heterogeneous distribution of strategy values to begin the experiment. This graph shows the final spread by the end of the experiment.

For the skewed starts (Figure 19b) the right-skewed population moves −0.001 from its own initial mean of 0.051, and the left-skewed population +0.010 from 0.933. Within a lifetime the right-skewed agents move −0.028 between the youngest and the oldest at a correlation of −0.112, and the left-skewed agents +0.022 at +0.098. The two conditions do not converge on a shared preference; each retains the character of its initialisation.

Table 2: RQ1 Growth Metrics Across Different Distributions

Configuration	Final Domain	Mean Node Count	Tree Depth
RQ1 growth homo	≈ 11008	$9.0 \rightarrow 15.6$	$5.5 \rightarrow 8.3$
RQ1 growth skewedLeft	≈ 11040	$9.0 \rightarrow 15.8$	$5.5 \rightarrow 8.4$
RQ1 growth skewedRight	≈ 11017	$9.0 \rightarrow 15.3$	$5.5 \rightarrow 8.3$

In addition with the heterogeneous start the values seen in table 2 are all quite similar. It seems that no matter the starting criteria or structure agents in this simulation find creativity in artifact

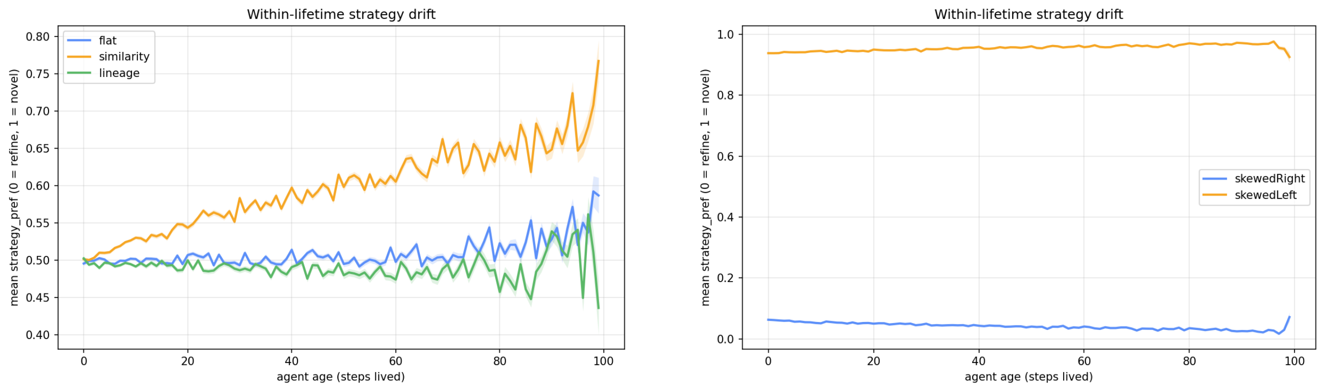

(a) Three retrieval modes from a centred start.

(b) Two skewed starts under lineage mode.

Figure 19: Mean strategy preference against the number of steps an agent has lived, pooled over five seeds and restricted to ages with at least thirty recorded shares. Shading is ±1 standard error. Lifespan is drawn at birth independently of behaviour, so age carries no selection on performance.

complexity and add artifacts to the domain at a relative rate (11000 per 500 steps). The values for stylistic diversity are similar in this regard as seen in Table 3.

Table 3: Feature Dispersion Changes by Distribution Type

Distribution	Initial Value	Final Value	Trend
Homo	0.2533	0.2217	falling/flat
Skewed Left	0.2498	0.2190	falling/flat
Skewed Right	0.2498	0.2237	falling/flat

Although these values are slightly lower than the heterogeneous start values seen in figure 11 (which is to be expected as a heterogeneous start implies a spread of agent types / artifact styles) they follow the same downward sloping curve of feature dispersion. Right skew has the highest feature distribution.Now our lock-in test as seen in Table 4.

Table 4: Lock-in Metrics Evolution by Distribution Type

	Root Concentration		Distinct Roots	
Distribution	**Initial**	**Final**	**Initial**	**Final**
Homo	0.01	0.20	100	232
Skewed Left	0.01	0.12	100	235
Skewed Right	0.01	0.15	100	238

The two measures disagree: a right skew ends with the most distinct roots (238), while a left skew ends with the lowest concentration (0.12) and so is the most evenly distributed lineage influence.

8.7 Interaction Networks

Our network analysis is conducted on experiments with a share value of 5 and a population of 100. The field metrics are seen in Table 5.

Table 5: Average network statistics across five runs for each domain retrieval mode with lifespan on.

Mode	Accepted Edges	Modularity Q	Range Min	Range Max
Flat	16,195	0.246	0.176	0.278
Lineage	16,213	0.248	0.176	0.284
Similarity	16,209	0.247	0.182	0.280

Once again there is a slight change in the structured modes as we see higher modularity levels reached, yet the average modularity remains near-identical across all runs. The modularity score does not diminish as in Motz's but rather stays around a certain score, as seen in Fig 20. Q remains below the 0.3 threshold adopted from Motz in all modes hence failing to prove a sustained clique formation.

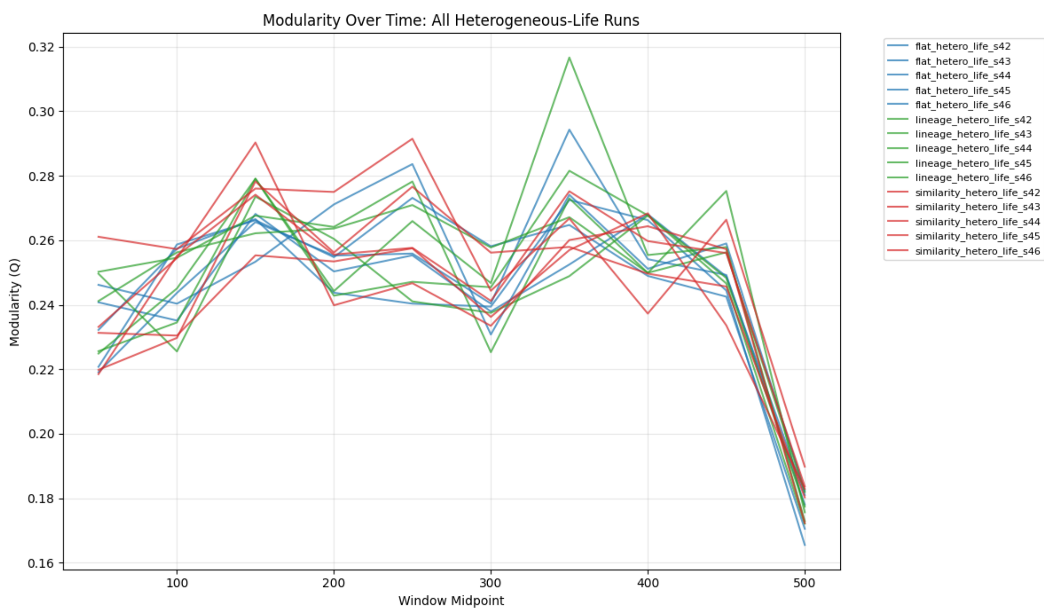

Figure 20: Modularity scores at the time of rolling window for all hetero, lifespan runs .

8.8 The Autopoietic Loop Made Measurable

In the following experiments, we explore the resulting domain network across three structures of association between artifacts. To define a clique we can look at connected components in this network. These are groups of artifacts that are linked together via edge walk; artifacts in different components cannot be reached by the associations we have drawn above. This metric reveals that in flat there are 352 components, 1 in similarity, and 5 in lineage on average (the image is a single run and so it contains different values). We can also observe the weights of these components or what percentage of domain artifacts live on each island, or more concisely, the largest island. Flat returns 0.157 meaning that the largest component in flat mode consists of ≈ 16% of domain artifacts. Both similarity and lineage have a largest island weight of 1.

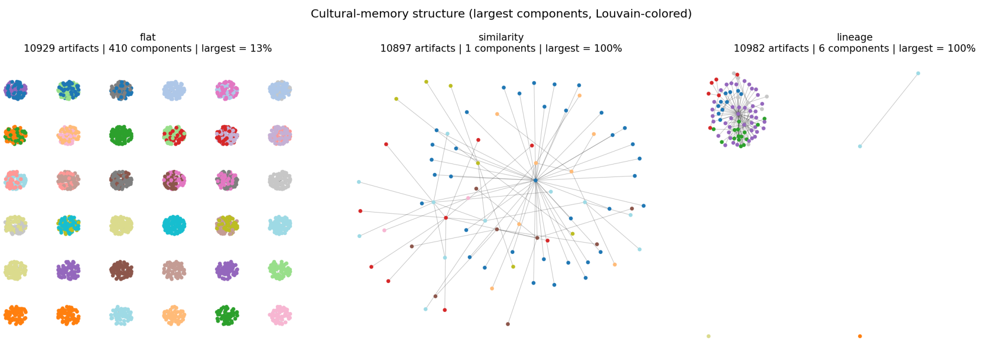

Figure 21: The three domain structure networks visualized using a spring layout to separate nodes cohesively. Nodes are Louvain colored, where the same clique nodes (as measured using Louvain) are colored the same. .

Figure 21 gives us a fraction of the domain structure that reflects the larger pattern of fragmentation in flat, and the general connectivity of the structured modes. By looking at the number of isolated nodes in each graph we see flat has 227, similarity has 0, and lineage has 4. Just as in Motz's paper, we use the Louvain algorithm, but on the domain structures. We find that the Q-values for similarity (0.979) and flat (0.947) are high while lineage is lower (0.608).

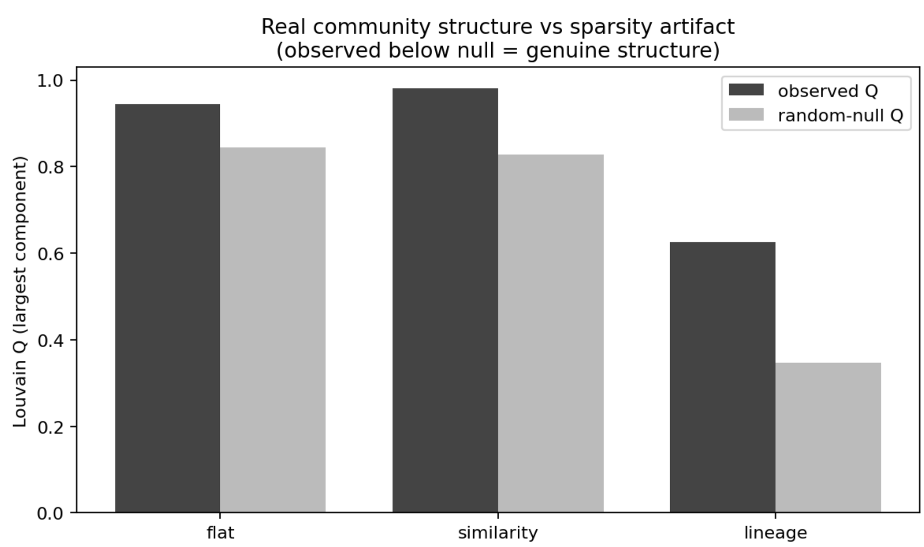

Figure 22: Bar graph of observed Q values vs. a density matched random null. The difference between these values can be interpreted as a degree of clique formation(if positive).

We run an analysis on a network with the same node and edge count randomly structured. Any increase in Q value from this random graph to our observed graph could be a result of clique formation. Such is the case for similarity and flat domain networks seen in Figure 22. Flat and similarity exceed the null graph values ($\approx .83$) by about 0.1 (slightly more for similarity) while lineage exceeds its null graph by ≈ 0.3.

Before, we looked at the age centrality relationship. We are still interested in the field, so we explore that relationship by taking the same agent centrality measure from before, and how anchored their artifacts are in the domain. We correlate each agent's eigenvector centrality in the accepted-share network with the mean cultural in-strength of the artifacts they contributed (Pearson r, averaged over five seeds). Our measures return that flat mode has an r value of -0.01, similarity has a value of 0.178, and lineage with 0.058.

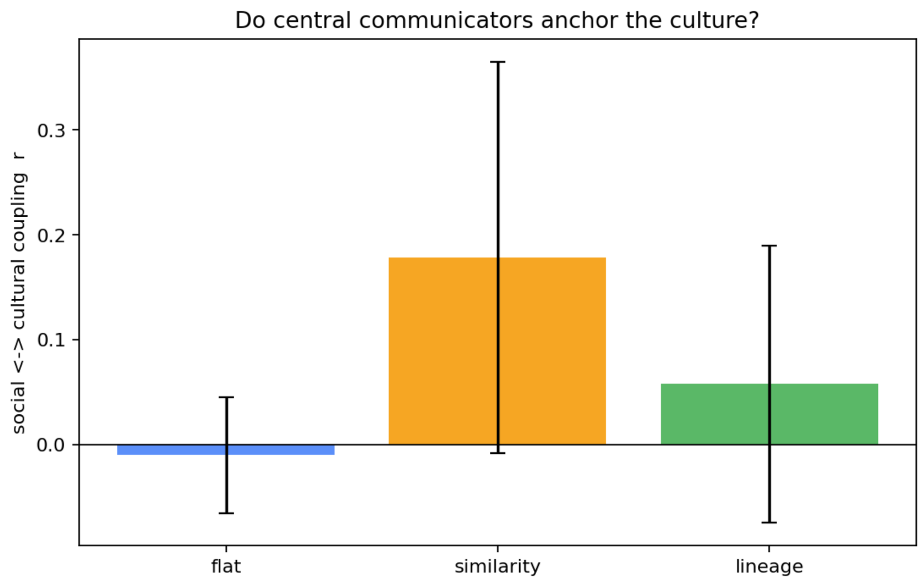

Figure 23: Graph of the r values measured over 5 seeds where bars represent the means and error bars are the standard deviation for each structured mode.

To get a central overview of the evolution of our domain structures over time one needs only to look at the graphs of Figure 24.

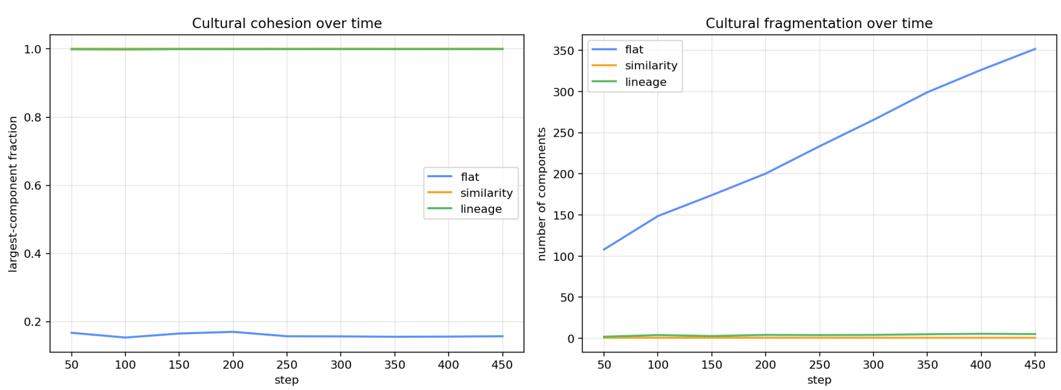

Figure 24: Graph tracking the largest component fraction and number of components over time, one line per domain structure.

9 Discussion

Throughout this paper we have been implying that culture is a kind of constraint on what is possible. It is the contingency of Luhmann, in that past representations of what was possible shape what becomes possible, and it is the line of association between past and present that Takashi Iba describes [Luh95, Iba10]. Read this way, the pre-notion of the domain as *wholly* representing culture was wrong. Culture is not the pile, nor even the structured pile; it is whenever contingency narrows the possibilities of the present, and in a simulation such contingencies are everywhere. There is the constraint of the whole system's design, the constraint of the tools that go into cultural production, the associations drawn through communication and through the memory of the agents themselves. We must try to separate a useful concept of culture from this vast definition.

Consider how an anthropologist works. They do not consider what norms do not happen. Culture becomes easiest to see in comparison, norms versus norms. What the anthropologist never asks is why gravity holds people to the ground, or why sound is carried between particles, because every culture they might compare exists under the same set of physical capabilities. Those constraints are not worth consideration for an anthropologist; they are simply the base. In simulation this base is a design choice. We can change what an agent is composed of, how it is capable of communicating, what it is capable of producing, and how it evaluates an artifact. This is a genuine advantage, but it comes with an obligation. If everything in the system is a constraint, and every constraint may be culture, then the claim "the active domain adds culture" cannot be shown false by any measurement we could make. The obligation, then, is to declare for each comparison what is being held still, what is the base (Section 3.2).

There is one further clarification, and Luhmann himself has raised the question. He was openly hostile to the very word we have been using, calling culture one of the most detrimental concepts ever invented [Luh00, p. 247]. His objection is worth taking seriously: "culture" invites us to treat shared meanings and values as substances floating above the system, when for him a social system consists of communication and nothing else, and everything else—bodies, minds, and in our case the code—belongs to the environment, able only to perturb the system from outside [Moe12, Luh95]. We do not use culture in the sense he rejected. We use it as he uses contingency: the constraint that a system's own past operations impose on its next operation. This gives a clean rule for what *is not* culture. The code does not evolve, and it reveals no norm; it simply is. When we change the code—suppose we cut the pixel budget of an artifact—the change itself is not culture. What may follow from it is: the new regularities of production, of sharing, of what agents come to treat as successful. The code change is the *cause*; the culture is the emergent behaviour that settles around it.

With the substrate fixed we can consider the arisen contingencies. There is the artifact as a *form*: what its expression tree looks like, how many nodes it carries, which direction its feature vector points. And there is the artifact as a *descendant*: where it came from, what it was bred out of, which domain artifact it was anchored to. Complexity and style are two readings of the first; the domain network is a reading of the second. They are like two observations of the same system. Two distinctions drawn across the same stream of events, each marking one side and leaving the other unmarked. Complexity marks what an artifact is and says nothing of its descent; lineage marks where it came from and says nothing of its form. The point is that this is not a difference in how we choose to look as the two come apart in the data we have so far seen. Figure 6 shows a pair of artifacts one breeding step apart that visually do not seem related, while artifacts with no common ancestor can be stylistically near-identical (Figure 5). Descent is a historical fact about an artifact; form is a fact about it in the present. That they can be dissociated is precisely why similarity and lineage are two different modes and not two views of the same structure. This difference raises new contingencies in how artifacts are retrieved from the domain, this directly influences creativity and the loop with everything else as the base.

Beginning with creative growth, our early figures—domain growth (Fig. 8) and artifact complexity (Fig. 9)—show that structure has almost no effect on the artifact as form. To sustain constant domain growth, artifacts increase their complexity: they branch out early and then build on the branched foundation with finer detail and more interesting combinations. This is a characteristic trajectory of creative evolution, and its strength is that it is invariant. It holds with lifespan and without it, across all three structures, and across homogeneous, heterogeneous, and skewed starts. Part of this trajectory is enforced by the code base itself: the refining phase coincides with the hardcoded ceiling on tree depth, so once agents can no longer go deeper they can only go finer. The reason agents go deeper first may also be due to the hard-code, in that novelty is the goal and novelty is more possible if you go deeper first. This way of perceiving the artifact evolution gives us an illustration, but the substrate (code) directly constrains the shape of the trajectory both explicitly and more implicitly, as our new framing predicts. Structure is not important here, because structure is a reading of the relational, and the growth of complexity and novelty detection is a fact about form. The information contained in some past artifact is enough to carry complexity forward; it does not matter through which relation the agent reached it. This is why complexity accumulates even through death and the extinction of individual memory: the domain holds the form, whatever the path.

Richness asks a different question of the same artifacts than complexity does: not how elaborate the domain became, but how varied that elaboration is. A domain of uniformly complex artifacts and one split between the simple and the elaborate return the same mean and very different standard deviations. Figure 10 shows the modes ending inside one another's seed bands, so structure leaves this untouched as well. Whatever the association rule does to cultural memory, it does not reach the form of the individual artifact, and this holds for the variety of forms as much as for their average.

The same trajectory appears in our measure of style, but it is worth being careful about what style is. Both style and complexity are readings of the same artifacts, and our complexity measure is in some sense housed inside the feature vector that defines style. They are two observations of the same form. A style is a distinct level of complexity, especially here where both are separated (high vs. low complexity and one style vs. another) based on the feature vector/form. It follows suit that similarity mode, which is based on feature vectors for a measure of novelty access, would outperform the others. For lineage to be slightly lower than the baseline flat highlights the difference of lineage mode and its potential overlooking of any features of the artifacts themselves. Feature dispersion begins high (Fig. 11), a reflection of the early branching of expression trees where complexity increases and more styles can be staked, then convergence to a lower value. This is the same branch-then-refine shape read in a coarser grain. We recognise styles in art because we can perceive edges, colours, and symbols; we do not sort artistic styles by wavelengths outside our visual field, nor distinguish art made of undetectable edges. The constraints of perception fix the grain at which style can be observed at all. Our feature extractor and expression tree bounds plays that role here, and so the stylistic measure inherits the constraints and then the trajectory of the forms it reads.

The acceptance rate is invariant across structures (Fig. 17), and this invariance is architectural rather than empirical. Acceptance fires when an artifact clears a rolling 80th percentile of recent interest. The problem is the percentile threshold is self-normalising by construction: as the population's interest distribution drifts upward with rising complexity, the threshold drifts with it, because it is a quantile of that same distribution. Roughly one fifth of shares clear it forever, no matter how novelty, diversity, or complexity evolve underneath. The flat curve is therefore a fixed point of an adaptive threshold rather than a property of creativity, and it is evidence that the threshold is doing its job—tracking a moving distribution instead of saturating or collapsing. What follows is a methodological consequence, acceptance rate is one metric in this study that *cannot* carry a cultural signal. Just as in anthropology, the focus does not become the extent of culture and the rate of its growth, but rather what patterns constitute the culture, and how do they emerge.

Acceptance is nearly flat across retrieval position (Fig. 17), and this follows from the acceptance rule rather than from structure. The rolling percentile re-normalises whatever the field sends it, so an artifact bred from a close relation and one bred from a distant relation meet the same bar. What survives that normalisation is a residue of two to three percentage points, and the direction of the residue differs by mode. In similarity the rate rises with distance, giving a consistent direction to move in. In lineage it peaks in the middle and falls at the far end, giving none.

The within-lifetime measurements follow this exactly. Similarity, the one mode whose landscape has a consistent direction, is also the one mode where agents move with age, gaining 0.167 in preference between the youngest and oldest agents at a correlation of +0.163. Lineage, whose landscape offers no single direction, shows no net movement, and flat, where retrieval position carries no meaning at all, shows the same. Since lifespan is drawn at birth and independently of anything an agent does, age carries no selection on performance, so this is not an effect of which agents survive.

We read this as evidence that the update rule works and that what limits it is the landscape rather than the mechanism. Agents climb the gradient that is there; it is simply a very shallow one. The population mean does not record this because a replaced agent draws a fresh preference from the distribution its predecessor started from, so each death returns a learned value to the initial spread. What the population does record is its variance, which widens over the run in both structured modes, as one would expect of a walk with almost no slope to bias it.

This also sets a limit on what our navigation question could have shown. A preference over retrieval position is learnable only where position predicts acceptance, and under a threshold defined on relative rather than absolute interest it barely does. The flatness is a property of the acceptance rule we inherited, not of structured domains as such, and a rule that judged artifacts against a fixed bar would leave the gradient standing.

Which brings us to the axis where structure is loud. Structure did not touch style, but it moved root concentration decisively, and one need only look at Figure 21 to see why. We can say that under flat cultural memory is 352 fragments connected by certain domain parent id's, under similarity it is one connected whole based on similarity, and in lineage it is a giant connected genealogy with a few stragglers. Those disconnected nodes are artifacts that entered the domain and connected to nothing, with no parent ID. This number is reminiscent of the number of agents that last in a lifespan run;

so these artifacts may be the roots of each lineage. Unlike in lineage mode, these roots cannot be related to the past/other lineages; there is no cultural memory, and it then becomes obvious the lack of 'cultural memory' in flat mode. Random retrieval makes for a more individualistic culture of creativity: the possibilities open to an agent depend on its interactions outside the domain, with the field and with its own memory, and the domain cannot recognise where an artifact has come from other than from its creator. The structured domains are connected wholes; each lineage stems from some past lineage and flows into the next. As each is woven into another, it becomes hard for any single one to dominate. As one lineage gains ground, so should its relatives, for the system cannot label an artifact creative and a structurally near-identical one not. And because every artifact is connected, one can trace a path from any artifact to its nearest neighbour and onward across the whole domain, so that tugging on one strand tensions the rest of the web. This is the essence of structure: there is a place for every artifact in the structured domain, because the domain is built around each artifact. Within this structure we can observe the patterns of communication and what communicating groups may form. The near perfect modularity scores in figure 22 are not because we have discovered the formula for pure group formation; rather, similarity and flat are both thin tree-like graphs. This is because they both look at the single-nearest attachment, creating a distinct directed path, like a tree. The one informative Q value is lineage (0.608) precisely because it is lower. The lineage graph is denser; it allows crossover, as artifacts can share multiple edges (in domain ancestors). This creates genuine sub-families stemming from common roots, rather than a thin thread of associations.

None of the network results, however, would occur without navigation. Without agents actively steering through the structure, the artifacts would carry no meaningful relationships and root concentration would collapse to the flat-mode case; it is navigation that distributes the influence of each lineage. It may be that in terms of feature dispersion as every agents refines its own artifact their initial style becomes more clearly their own. This would be supported by the lowest stylistic diversity in novel seeking agents as this search is more likely to lead many to the same place. The lineage tracking does not pick up on this cconvergence of style so these metrics are quite seperate. Rather it is possible that the novelty seeking agents constant adoption of new things makes it hard for any one lineage to claim monopoly on creativity, as seen in Table 4. This is reinforced by Table 3 where refining agents have a greater style dispersion as they maintain distinct styles without converging so easily on any. Consider that because every structured agent carries a strategy and a learning update, agents can pivot away from unrewarding zones, whereas unstructured agents can only wait for a more interesting artifact or to retrieve a new one from memory. Since creativity in our system is defined in terms of novelty, and our structures each reflect a particular type of novelty, one might picture the domain network as a creative topology in which agents occupy novel-producing zones. Two agents working the same domain artifact are likely to produce something relatively similar, and so they compete: whoever produces first is the more novel. An exploring agent may then be stepping on the toes of a peer already refining that region. This would explain both the equitability of the structured domains and the failure of the skewed starts to converge on one another. We did not measure this overlap directly—whether outlier-strategy agents suffer lower acceptance when they sample regions their peers already occupy—and we mark it as a conjecture for future work.

In addition to structure we have also explored the role of lifespan, a feature rarely made explicit in creative simulations. We explored how cultural memory may be used by agents who are born into a culture, and compared with agents who have built their own cultural environment. Our atemporal runs show that structure has even less influence with the latter. With generational turnover, agents were motivated to depend on the domain for representations of past creativity; the domain offers successful artifacts to agents who possess no understanding of what success is. For the agents who built the domain themselves, such inspiration is unnecessary, since they evolved in tandem with the domain. The domain is only a highlighted subset of the far greater accumulation of creative possibility circulating in the field, and an immortal agent's memory can outgrow that subset entirely. Newborn agents have no such memory and need the domain. In lifespanless runs the domain becomes a trophy case for past artifacts, and it is doubtful that removing it entirely would change much at all in such a world.

Turnover also introduces something like wisdom. Older agents are naturally the more successful communicators; they have developed an understanding of creativity that the young can only absorb as it leaks into the field. Crucially, this age–centrality effect is largely independent of the domain structure—it holds, with only slight variation, across all three modes (Fig. 16). Agents who have learned the whole system, and who grow their expressions through field and domain interaction alike, become successful communicators with or without a navigable structure. Even though we have made culture navigable and concrete in the domain we can imagine that cultural memory is expressed across the whole system in its constraints. Wisdom is a visible case: new agents learn matured representations of creativity from old ones, and those representations implicitly link a past as lived by the elder to a present as interpreted by the newcomer [Ass10, BO12]. Culture dictates every interaction, not only domain retrievals, and its memory persists where it is never made explicit. In a system as constrained as ours, the architecture itself carries shared meaning to every agent: each artifact is limited in its expressions, each style limited in number, each agent limited in ability, and that limitation may already suffice as a normative constraint. The elders have simply made the best of these limitations, as the whole system strains against them. This also explains the lifespanless case, where every agent accumulates understanding from the very first step, so that no interaction—with an elder, or with the domain—is more essential than any other. Although these agents communicate culture it is not necessarily true that their artifacts are central to society. The stark contrast between lineage and similarity, as seen in figure 23 may imply the ability of similarity to structure domains around the more communicable concept. This means that an agent whose communication infiltrates the most discussions leaves a trace of their work that is detectable in similarity mode more so than in lineage. Yet lineage is disastrously low, for under the same logic an agent communicating would still leave its DNA trace in the immediately added artifacts. It is worth mentioning this metric is quite noisy, and the values varied for all modes. They seem far noisier in the structured mode, as seen by the error bars. This is likely a result of random sharing, and more tests would be needed to draw any conclusive conclusions on the matter; yet there may still be some pattern. What is robust is weaker but cleaner: across all five seeds, similarity is the only mode whose coupling stays positive, whereas flat and lineage straddle zero. We read this as suggestive, not conclusive, evidence that a similarity-structured domain — built on the same feature space agents use to judge value — is best able to reflect successful communicators in the artifacts that persist. Although the culture of the system can accumulate outside the structure, the structure allows itself to be used and traced. Consider the root concentration with lifespan off. As with lifespan on, there is a general decrease in sustainability as modes seem to converge on higher levels of concentration, as seen in Figure 14. Looking more closely, we see that the flat mode has taken the role of most dispersed lineage power, and yet it is the same value as when lifespan is on. What has really happened is that without lifespan, the structured modes are far worse at dispersing power and nearly 20% of the domain artifacts are of a single lineage. This is because structure, with strategy, allows agents to exploit more successful lineages. The role of structure is not necessarily as a distributor of culture, but as a relational tool which agents may use.

Read across our two observations, the shape of the results is clear. The product axis and the relational axis are near-independent: one may reorganise the entire relational stack, from a fragmented pile into a connected genealogy, and the forms produced along the way barely shift. This is not a disappointment but an architectural fact, and it is the reason structure is silent on form and loud on descent. It also explains the small exceptions/perturbations we did observe, where the structured modes carry a slightly higher node count or complexity than flat: a faint upward leak from the relational axis into the product axis, which is exactly what one would expect if the two are near-independent rather than perfectly so.

Our system, then, is like a sieve, and to do meaningful research we must search deeper into what comes through. This is why we have focused on norms—the patterns that survive the sifting of our hardcoded constraints. We take the growing complexity of artifacts and look deeper to find that cultural memory can be carried through older agents. We take the interconnected structure and look deeper to find agents manoeuvring a topology of novelty. We see the weakness of a structured memory when lifespan is switched off, and understand that the association between past and present already lives inside an immortal agent. Every one of these results is a statement about the behavioural norms that emerge once the substrate is fixed, and it is on those norms that cultural memory acts.

It remains to say what this means for the autopoietic loop. Autopoiesis in Luhmann is a strong claim: the system produces the very elements that produce it. Iba supplies the element we need, the discovery, each one contingent upon a prior discovery [Iba10, Luh95]. In our system a domain entry is that discovery: an artifact bred from a retrieval, shared, evaluated, and admitted, which then becomes the ground from which the next retrieval is made. Structure does not micromanage these operations—they remain constant, necessary, and largely random on their face—it shapes only what is remembered and what may be expected of a successful discovery. And here is the unifying lesson of our results: making the domain active did not manufacture culture, because culture was already latent in the constraints and operations of the system. What structure did was give that culture an explicit frame within which it could be observed. The passive list did not hold less culture than our structured domain; it collapsed the relational axis into invisibility, so that descent could not be marked at all. To structure the domain is not to add a floor to the tower of creativity. It is to draw the distinction that lets us watch a culture inherit itself.

10 Limitations of the System

There are many constraints to the system in all areas of the coded DIFI framework. The most obvious may be random sharing. If the strategy is to utilize the cultural memory, then that strategy ought to extend to whom it shares with. The navigable domain gives agents real freedom in their hunt for a rewarding retrieval, but the second reward signal—generated when a shared artifact is accepted—is delivered over an essentially random topology. Then there is the relatively random selection of bred artifacts too. Were agents able to navigate the cultural memory, the field, and their own memory alike, contingent culture might finally become manifest in all of its currently dormant forms. Global thresholds (boredom, domain acceptance) remove the individuality that would otherwise further perturb the system in unexpected ways. The unity in every phase- that each agent gets bored once per iteration and shares with X many agents, and so on all constrain the natural perturbations that would otherwise reach the system. We deliberately kept the sharing threshold personal, the one place we judged this individuality indispensable. Such uniformity is otherwise the price of measurability, since an entirely unconstrained system would be far harder to read and compute.

Some limitations of the system would still do well to address. The problem of a small learning update was addressed, but it would be natural for agents to be able to access the domain a few times before selecting an artifact. This would give greater opportunity for agents to match their interests within the domain structure. The reason this wasn't implemented was because of the potential slowdown that this would bring to the simulation, as hardware is another limitation. There are also limitations in our experiments. A similarity run in our skewed starts would be worth comparing and analyzing. Also, the culture network graph shouldn't use single-nearest, causing modularity (Q) to be inflated by sparsity. A denser k-nearest or thresholded graph, built from stored feature vectors, would be more faithful. Several measures would benefit from more seeds (for noise reduction), like the coupling metric and stylistic diversity. It is also noticeable that many graphs begin to plateau over time (9,11, 12, 18) and it would be worthwhile to run longer and larger experiments to see how the system stays creative as discovery is pushed to its limits, if it can be.

11 Contributions and Conclusions

We have made the domain something real and measurable. Furthermore, we have shown a tractable way to bring Luhmann's theory of autopoietic systems into computational creativity research. We see the system build a cultural memory (that feeds back into the system as seen by our experiments) out of the same random social churn that Motz's system piled into the domain list. The general contribution has not been that structure makes the system louder, but that it defines the domain as a memory in the first place; the next move would be to let the memory reach back into the field through communication, closing the loop in both operation and observation. On a larger scale, we have also presented a problem of working with these hierarchical systems; that changes to a system on one level, the domain in this case, are more obviously revealed in analysis on that level. We can probe and assess the changes translated to other levels of the system, but these are noisy and arduous to describe in process. Yet, the probes have been insinuating, and it is only with more research that the process clears.

If the fire is the culture of a creative system, then our structured modes build a single raging flame, each spark catching the next. The structureless baseline is like a scattered fire, sparks jumping but never catching enough energy to become one flame(like the 227 disconnected orphan nodes of flat structure). The system nonetheless sparks away, ceaselessly, giving one the illusion that a flame rages. This is the illusion of a domain as creativity produces while cultural memory is not yet present.

12 Future Work

The experiments conducted have thrown the door wide open for future research. There is the prospect of grander runs that push creativity to crisis—how do structure, generational turnover, and strategy respond when creative possibility grows scarce, as it may sometimes feel in our own reality? It would also be interesting to see how a stabilized culture(where agents have settled into their general preferences) would work in this system that runs on novelty. Culture is rarely so burgeoning; exploring the stability of culture with longer runs and also a forgetful nature to the domain would be worthwhile. It was considered to build a structure based on social parameters (views, likes, shares etc.), and this could be used with a forgetful aspect where artifacts that lack attention become forgotten. This would deepen the understanding of cultural memory and how the flame of creativity remains roaring.

There is potential for roles that emerge in the skewed societies between outlier agents and their conformist brethren. We saw outlier agents moving against a conformist strategy and exploring why and whether it is of benefit to the individual or society is a gap in our research. The broader question is how different starting distributions shape what emerges. Above all, there is implementing a biased share parameter: much work, including the foundations this paper builds upon, has accepted a randomly organized field, and the time has come to change that. Experimenting with these parameters that give the agent further control (share, boredom, breeding) and seeing their interactions with structure would be a sensible extension of this paper's research; as would exploring different structures and their resulting domain networks.

In this paper there may be furthered research in regards to clique formation in creative societies. Motz focused his research on this and this paper even devoted a section (8.6) to see whether structure and lifespan would have an affect. We have found older agents to be more central communicators and it is possible that cliques may be forming around them. In our figure 20 we see this trend of rising and falling stages of clique formation. It would be worth exploring the age distribution of these peaks and doing runs with consistent modularity measurements to see exactly what events take place around the peaks. This would give us a better idea of how age and communication/clique formation are related.

A single parameterisation models a single culture; the framework assumes cultures differ only in their contents, not in their mechanisms of transmission and evaluation. Varying those mechanisms is a natural extension. In the same spirit, applying the framework to other creative substrates—language, logic, sound—would test how far these findings generalise. A further possibility may be interacting with many systems of creativity, coupling them in a single environment and allowing translation between them, observing how cultures perturb one another. We already have a base for this as our very artifacts are mathematical expressions translated into color. This also deepens our understanding of Luhmannian theory as we can measure the reaction between two coexisting systems.

References

[Ass10] Aleida Assmann. Canon and archive. In Astrid Erll and Ansgar Nünning, editors, *Cultural Memory Studies: An International and Interdisciplinary Handbook*, pages 97–107. De Gruyter, 2010.

[Axe05] Robert Axelrod. Advancing the art of simulation in the social sciences. In Jean-Philippe Rennard, editor, *Handbook of Research on Nature Inspired Computing for Economy and Management*. Idea Group, 2005.

[BO12] R. Alexander Bentley and Michael J. O'Brien. Cultural evolutionary tipping points in the storage and transmission of information. *Frontiers in Psychology*, 3(569), 2012.

[Bod04] Margaret A. Boden. *The Creative Mind: Myths and Mechanisms*. Routledge, 2nd edition, 2004.

[Bok24] Jakub Boksansky. Sampling from a normal (gaussian) distribution on gpus, October 2024.

[BS99] Geoffrey C. Bowker and Susan Leigh Star. *Sorting Things Out: Classification and Its Consequences*. The MIT Press, 1999.

[Csi99] Mihaly Csikszentmihalyi. Implications of a systems perspective for the study of creativity. In Robert J. Sternberg, editor, *Handbook of Creativity*, pages 313–335. Cambridge University Press, 1999.

[Gab95] Liane Gabora. Meme and variations: A computational model of cultural evolution. In L. Nadel and D. Stein, editors, *1993 Lectures in Complex Systems*, pages 471–485. Addison-Wesley, 1995.

[Gal03] Philip Galanter. What is generative art? complexity theory as a context for art theory. In *Proceedings of the 6th Generative Art Conference*, 2003.

[HL18] Otto Hantula and Simo Linkola. Towards goal-aware collaboration in artistic agent societies. pages 248–255, 2018.

[Iba10] Takashi Iba. An autopoietic systems theory for creativity. *Procedia Social and Behavioral Sciences*, 2:6610–6625, 2010.

[Jen10] Kyle E. Jennings. Developing creativity: Artificial barriers in artificial intelligence. *Minds and Machines*, 20(4):489–501, 2010.

[Læs05] Rolf Læssøe. Édouard Manet's "le Déjeuner sur l'herbe" as a veiled allegory of painting. *Artibus et Historiae*, 26:195, 2005.

[LH18] S. Linkola and O. Hantula. On collaborator selection in creative agent societies: An evolutionary art case study. In *Proceedings of the Ninth International Conference on Computational Creativity (ICCC)*, 2018.

[Luh95] Niklas Luhmann. *Social Systems*. Stanford University Press, 1995. Trans. John Bednarz Jr. with Dirk Baecker.

[Luh00] Niklas Luhmann. *Art as a Social System*. Stanford University Press, Stanford, CA, 2000. Trans. Eva M. Knodt.

[MAP13] Lei Clifton Lionel Tarassenko Marco A.F. Pimentel, David A. Clifton. A review of novelty detection. *Elsevier*, 2013.

[Moe12] Hans-Georg Moeller. *The Radical Luhmann*. Columbia University Press, 2012.

[Mot25] Luuk Motz. Modeling social creativity with large-scale multi-agent systems and emergent properties of collective creative processes, 2025.

[RJ12] Mark A. Runco and Garrett J. Jaeger. The standard definition of creativity. *Creativity Research Journal*, 24(1):92–96, 2012.

[Sau19] Rob Saunders. Multi-agent-based models of social creativity. In T. R. Besold, M. Schorlemmer, and A. Smaill, editors, *Computational Creativity: The Philosophy and Engineering of Autonomously Creative Systems*, pages 305–326. 2019.

[SB15] Rob Saunders and Oliver Bown. Computational social creativity. *Artificial Life*, 21(3):366–378, 2015.

[SG01] Rob Saunders and John S. Gero. The digital clockwork muse: A computational model of aesthetic evolution. In *Proceedings of the AISB'01 Symposium on AI and Creativity in Arts and Science*, pages 12–21. SSAISB, 2001.

[SG05] Ricardo Sosa and John S. Gero. A computational study of creativity in design: the role of society. *AIEDAM*, 19(4):229–244, 2005.

[SGJ09] Ricardo Sosa, John S. Gero, and Kyle E. Jennings. Growing and destroying the worth of ideas. In *CC '09: Proceedings of the seventh ACM conference on Creativity and cognition*, pages 149–158. ACM, 2009.

[SL99] Robert J. Sternberg and Todd I. Lubart. The concept of creativity: Prospects and paradigms. In Robert J. Sternberg, editor, *Handbook of Creativity*, pages 3–15. Cambridge University Press, Cambridge, 1999.

[TTJ19] J. Y. Tsao, C. L. Ting, and C. M. Johnson. Creative outcome as implausible utility. *Review of General Psychology*, 2019.