

Creative Artificial Intelligence within the Artificial Life Installation “Infranet”

Graham Wakefield
York University
Toronto, Canada
grrrwaaa@yorku.ca

Haru Hyunkyung Ji
OCAD University
Toronto, Canada
hji@ocadu.ca

The authors explore how current mainstream data-driven AI approaches can be questioned critically from a perspective of computational creativity and ecosystemic art. This centres on a critique of the future as being over-determined by the past; both from the data used, and in the questions or objectives assumed by training. The main contributions of this paper are to apply alternative creative approaches to nature-inspired artificial intelligence, and to detail some of these through their embodiment in the authors’ artwork “Infranet”. Infranet is a neuro-evolutionary art installation that exhibited at three international locations over 2018-2019. It uses geospatial data of the host city not as a training material but as a habitat for artificial life. In contrast to training-based AI systems, in Infranet there is no objective or fitness function and very little evolutionary pressure or competition. Moreover, it eschews the trend of a large and pre-specified neural network structure in favour of a population of thousands of small interacting neural networks, each with distinct structure, in a “liquid” process of continuous reorganization; resonating with some contemporary theories and models of non-conscious cognition in biological and ecological systems.

Artificial intelligence. Artificial life. Data art. Computational creativity.

1. INTRODUCTION

The rapid growth of Artificial Intelligence (AI) in recent years, driven significantly by the remarkable successes of machine learning (ML) methods training deep networks upon vast amounts of data, has not only placed AI in the centre of public attention, but is also increasingly embedding it into the everyday systems within which we live our lives. Artists have responded to this growth in many ways, such as using AI as a creative or analytical tool or critiquing by revealing the vast scales of data used and the dangers of bias embedded in them.

Our concerns as artists regarding machine learning here became twofold. Both regard the past (the “known”) over-defining the future (the “unknown”).

First, an AI trained on a store of data propagates the strengths and the weaknesses of that data; and the subsequent application of such an AI effectively bottlenecks the future into the terms of its past (which is why the problem of bias becomes so important). A once-trusted dataset that turns out to be unrepresentative will promulgate biases as the networks trained upon it continue to be used. And

even an AI trained on a dataset that is believed to be unbiased may develop and propagate blind spots as it is applied in a complex and changing world (Roselli et al. 2019). More generally, the proliferation of data-trained AIs broadcasts a tendency to focus on whatever has been most easily measured, at the expense of what isn’t easily measured, as well as the risk of mistaking a map for the territory.

Second, training typically optimizes for utility performance for a clearly defined objective question, such as, “does this image contain a cat” or “how to most efficiently navigate a maze”. The objective function is by definition biased toward a specific utility (and against the implicit contrast space). Moreover, since the objective function is defined in advance, it also casts this shadow into the future.

The main contributions of this paper are to explore alternative creative approaches to artificial intelligence and artificial life, through the lens of the artwork “Infranet”. Infranet uses geospatial data as a habitat for a large population of artificially intelligent virtual creatures. The design of this system eschews mainstream methods of artificial

intelligence, instead taking inspiration from natural creativity and cognition. The paper will elaborate an analysis of the system's design and its evolutions, and we will share new insights in order to enrich broader discussion about creative artificial intelligence and living data in art.

2. SEARCHING FOR CREATIVE INTELLIGENCE

Machine learning with neural networks is an originally biologically inspired field of computer engineering coming from a historical desire for generalized prediction. A neural network model produced by ML is effective if it will often make accurate predictions. It thus follows a statistical genealogy—trained networks are essentially multidimensional variants of line graphs fitting experimental data—and represents *induction* rather than *deduction* (and has very little to say so far about *abduction*). Even in unsupervised learning, where the data is unlabelled, the goal is still to create a predictive model of a "latent space" the data suggests. The method is an approximation, carried out numerically, which means making lots of heuristically informed, somewhat randomized guesses (*exploration*), and empirically filtering down to the most satisfying results (*exploitation*). As effective as such approximations can be, when turned to questions of creative intelligence, it leads to several problems as outlined in this section.

2.1 The problem of algorithmic specificity

First, most algorithm structures are effective at certain kinds of problem but are ineffective at others. For example, perceptrons (neural networks with no hidden layers) can only learn linear functions and are incapable of awareness of nonlinear spaces, whereas networks with hidden layers can approximate more complex solution spaces. For example, convolutional neural network structures are well-suited for image recognition tasks, while recurrent neural network structures work better for speech recognition sequences, and so on. More generally, networks and algorithms have *structural bias*; their structure is an image of the problems they can fit well.

Second, building and training networks requires the choice of hyperparameters, such as the topologies of hidden layers, the learning rate, batch sizes, regularizer parameters, etc., which have to be tuned carefully so that a machine learning model can solve a desired problem.

But there's "no free lunch". Any algorithm's elevated performance over one class of problems is offset by poorer performance over another class, and no optimization algorithm outperforms any other when averaged over all problems (Wolpert & Macready 1995; Adam et al. 2019). Moreover,

there's no way to know in advance what neural topologies and what hyperparameters best fit a *new* problem domain, beyond a few heuristics and varieties of population-based trial-and-error (Burkov 2019).

Furthermore, the value of creativity is characteristic of problems that are often multiplicities that are difficult to define precisely or cannot be stated in advance, since the terms and structures of its questions may themselves need to change by context and in time. As such, there is no ideal best candidate structure to be found (though there may be a diversity of good candidates for any particular moment). Therefore, we suggest that an open-ended world requires a system that a) can rewrite its own structure to adapt to the world, and b) can propose and evaluate many possible structures at the same time. Genes and neural cortices are natural examples of such structural plasticity and collective evolution.

2.2 The problem of objective functions

There is an assumption often made in training an AI that the problem being optimized (the "search space") will produce optimal solutions under the pressure of some measure of success (validity, efficiency, effectiveness, reward, etc.) The pressure of this measure embodies an "objective function" that drives training of the AI through the landscape of the search space toward the optimal points that maximize this measure. The assumption is that not only such a landscape exists, but also that is amenable to hill-climbing/gradient descent methods. Any aspects or features that do not present such a landscape, or for which measures are confounding, are likely to be excluded.

Evolutionary approaches have been suggested as alternative methods for problems in which the "search space" is high-dimensional or not "well-behaved", as they can explore multiple parallel objectives to avoid getting stuck or being "deceived" by local optima (Miikkulainen 2020). However, the externally imposed and pre-defined objective function remains, now cast as an evolutionary fitness function. To the extent that such measures are defined externally and remain static, they represent selective breeding to a frozen standard rather than the creative capacities of natural evolution as such. There is a risk that an important outlier state, one which might in fact be key to a creative paradigm shift equivalent to natural speciation would not be included at all in the selection. Natural evolution cannot be understood in this way. There is no teleology, no fitness function as such, rather a fuzzy and changing set of conditions of *viability*. Furthermore, since ecologies form densely linked networks and histories of internal feedback, its viability

landscapes are actively constructed by the ecology itself.

The more we examined AIs pre-trained on pre-defined "objective functions", the more the "No-Face" character of *Spirited Away* came to mind (Sen to Chihiro no Kamikakushi 2001). No-Face is an initially small and voiceless god, who nevertheless grows monstrous as it gives more and more and more of whatever it is that you show that you like, seemingly unable to understand anything else, until you are consumed by it.

External (extrinsic) pressures, like objective and fitness functions, do not seem to lead to *creative* responses in natural, human, or artificial cases. In studies of human creativity, external pressures play much smaller roles in creative discovery than previously thought (Baldassarre 2011). Fitness functions alone are insufficient to explain evolutionary innovations such as sight and flight. And artificial systems driven by extrinsic, designer-given objectives can still get stuck (Lehman & Stanley 2008).

In contrast, we looked at algorithms that use intrinsic motivations or otherwise avoid the use of fitness or objective functions. For example, the Novelty Search method has been shown to outperform fitness-based measures in a variety of complex problem domains. In this method, evolutionary candidates are selected not by scores, costs, or similar measures of effectiveness, but simply for having done something that hasn't been done before (Lehman & Stanley 2011; Woolley & Stanley 2011). Even for objective-driven problems, a combination of novelty search with regular optimization has outperformed objective optimization alone (Mouret 2011; Cuccu & Gomez 2011; Nguyen et al. 2015b).

Removing extrinsic objectives might also alleviate problems of longer-term stagnation. It has been demonstrated that *evolvability*, that is, the tendency of living organisms to develop increasing diversity and evolutionary potential over time, results without any selective pressure at all, but rather from neutral drift through genotypic space coupled with evolution's passive tendency to accumulate niches (Lehman & Stanley 2013). This is reminiscent of a similar argument put forward by Kauffman with regard to the inevitable proliferation of organic polymers due to natural reaction gradients toward new compounds of existing molecules (Kauffman 2002). That is, life's creativity may have more to do with its "strongly constructive" microstructures (Fontana 2006) than selective pressures of evolution; and that beyond sufficient complexity, this process may be self-sustaining.

2.3 The problem of anthropocentrism

What AIs can already achieve is often quite astounding. For example, a Deep Neural Network trained on ImageNet to caption images may appear to have remarkable capacities approaching human-level common understanding of the content of new images (Krizhevsky et al. 2012). However, we must beware of anthropomorphic biases: the same networks can evoke the same responses when presented with inputs that appear very wrong or even random to humans (Szegedy et al. 2013; Nguyen et al. 2015a), revealing how alien "what it is like to be" the AI really is (Brooks 2017). We can take this as a warning to be wary of human-centric (and consciousness-centric) conceptions of cognition as unnecessarily limiting.

Similarly, we are wary of limiting investigations of *creativity* to human standards. A working definition supported by many researchers in computational creativity describes it as "machines producing artefacts that would be considered creative if produced by a human" (Jordanous 2012), but this human-centric definition engenders blind-spots toward non-human creativity and collaborative human-machine creativity (Arriagada 2020; Ragot et al. 2020). We do not even have a coherent grasp of what "intelligence" and "creativity" mean in the human case (for example, Cardoso et al. 2009). Our concern is that if we focus on a division between human and non-human, in creativity and in general, we will end up enforcing the division by ourselves (for example the circular argument in Hertzmann 2020).

Instead, we look to the adaptive creativity found in natural systems. Here we find resonance with a model of creativity found in Henri Bergson's distinction between mixtures of quantitative and qualitative difference (Bergson 1911, and elucidated in Deleuze 1990). The tendency of uncreative systems is an entropic "relaxation" toward equilibrium distributions whose variations are quantitative, and as such are generally predictable and can be numerically approximated by few parameters. The echoes of "gradient descent" methods of statistical modeling in AI here should be clear. In contrast, creative living systems tend to complexify toward heterogenous mixtures of non-equivalent qualities that are more resistant to reductive approximation (e.g. speciation and emergence). In short, natural creativity is a continual differentiation into new variations of *kind*. Here we find resonance with the production of new structures for adaptive change, and the seeking of difference as found in novelty search.

3. INFRANET

Infranet began as a commissioned work for the Gwangju Media Arts Festival. In responding to the theme of "Algorithm Society: Birth of the Machine God" we started with the notion that intelligence is not a human privilege but a phenomenon of cognition in general. As an embedded condition for all living systems, and now also machinic systems, it forms a "planetary cognitive ecology" (Hayles 2020).




Figure 1. *Infranet: Gwangju. Gwangju Media Art Festival 2018 "Algorithm Society: Birth of The Machine-God", Asia Culture Center (ACC), Gwangju, Korea. November-December 2018.*





Figure 2. *Infranet: Vancouver. IEEE VIS Arts Exhibition, Vancouver, Canada. October 2019.*



Figures 3a,b. *(Middle, Below images) Infranet: NYC. Korean Media Arts Festival: "Techno-imagination: Living Data." Sylvia Wald & Po Kim Art Gallery, Manhattan, New York, USA. August-December 2019.*

Infranet has exhibited at three locations during 2018 and 2019 (see Figures 1-3). In each case the landscape of the simulation consisted of a variety of geospatial data gathered on the metropolitan region surrounding the exhibition location; these have included Gwangju (South Korea), New York (USA), and Vancouver (Canada). However, Infranet is not a map or visualization of the city data; rather, the data is treated as a habitat for a population of AIs as artificial life creatures. Moreover, these AIs are not employed to discover or visualize specified aspects of interest to us from the data, rather the projected images are traces of the processes, experiences, and collective interests of the creatures themselves.

In ecosystemic art, it is well recognized that depth of an ecosystem's behaviour is deeply dependent on the qualities of the environment it inhabits (Antunes et al. 2014). A vibrant and diverse ecosystem needs an environment that is rich with niches in a variety of interestingly non-uniform and non-random distributions to support it. Geospatial city data can be very rich in this way. But to take advantage of a complex environment, life also needs to match its requisite variety, that is, to be able to generate a repertoire of responses which is (at least) as nuanced. For this purpose, we adopted neural networks with dynamic topologies within an evolutionary simulation that is unsupervised and without objective, and highly liquid through social exchanges inspired by non-human forms of collective and adaptive cognition.

3.1 A dynamic diversity of tastes

Infranet supports around 4000 to 8000 mobile creatures, each with its own neural network. They have sensors that can read the local underlying data, but these senses are biased by a personal "taste" or "scent" that filters most of the data away. A creature is most sensitive to features with very similar (positive correlation) or very dissimilar (negative correlation) taste, but blind to features that are unrelated to their taste. It is as if, for any

particular creature, the majority of the data simply doesn't exist.

As creatures move through data space they leave fogs and trails of their scent behind. These fogs and trails, and neighbour interactions, are the only element visualized in Infranet (see Figure 4). The colour depends on the creatures' scent, revealing shifting tastes and kinship relations through the population. Fogs disperse widely and quickly, revealing how they pulsate and exchange between each other. Thinner trails remain in place for a long time, a kind of long-term memory that gradually desaturates and then evaporates, revealing how creatures have experienced the city.

Creatures are also aware of these trails and fogs, as filtered by their own taste. Together, sensations of city data and creature trails feed into the neural network that is unique to each agent, which in turn feeds back to their motion, spiking, and changing taste (see Figure 5).




Figure 4. (Right to left, top to bottom) Nine detail frames in an evolution of Infranet: Gwangju, taken at an interval of two seconds between each image. Communications between neighbours—sharing ‘scent’ and potentially also neural networks—appear as transitory lines. Agents also inscribe fogs and trails of their passage that are comparatively more persistent in space, coloured according to their scent.

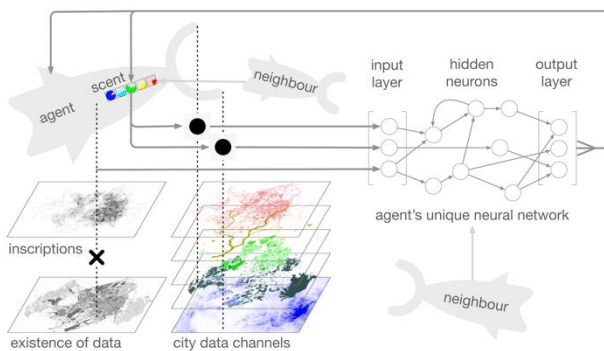


Figure 5. The principal dynamics of a creature in Infranet. The creature (top-left) senses the city data (bottom-left) as well as the existence of other creatures' inscriptions in space, the product of which feeds back to

the creature's intrinsic measure of well-being. The senses are filtered by the creature's internal bias (“scent”), and then fed into the input layer of the creature's neural network (top-right). The neural network's hidden structure is unique to the creature and changes during the creature's lifetime by horizontal transfer with other creatures (bottom-right). The output of the neural network directs the creature's motion, and changes of bias, which is also influenced by that of other creatures nearby (top-centre).

3.2 Viability conditions

In Infranet there is no “fitness function” and very little evolutionary pressure or competition in the system. Instead, there is only a viability condition expressed as two intrinsic motivations: 1) to have recently encountered data of the kind the agent has taste for, and 2) to have recently encountered traces of other agents. In plain terms: a primal need to verify the world still exists, and that one is not alone in it. In the absence of these experiences, a creature's internal sense of well-being decays, and if this reaches a critical threshold, the creature is removed and replaced with a newly created creature at a random location.

3.3 Dynamic micro-structures

Unlike many neural network approaches, in Infranet each creature's network is different in topological structure. To enable each agent to have a different structure of neural network, and still apply evolutionary dynamics, we utilized the NEAT methodology (Stanley & Miikkulainen 2002) as applied in (Waagenar 2017).

However, unlike typical evolutionary systems, network variation is not tied to reproductive birth, nor are generations run in synchronous batches. Instead, these structures can change within a creature's lifetime as creatures exchange networks between each other, overlapping asynchronously at high frequencies (sometimes multiple times per second). With each pulse, or chirp, they ‘sing’ their neural structure; and a neighbor may pick up this song, with mutations or by mixing with their own (via crossover operations), to produce a new neural structure in themselves. This highly dynamic process is not unlike the horizontal gene transfer that occurs in many microbial quasispecies.

Adoption of neighbour's networks is driven by relative measures of difference: creatures with lower wellbeing are more likely to adopt networks of their neighbours, especially if the neighbor's taste is very *different*.

3.4 Dynamic macro-structure

The creatures move through space in pulsations akin to neural spike trains. Spikes determine speed

(so over-excited creatures can 'overshoot' features, lower activations accumulate deeper inscriptions).

They also try to synchronize with their neighbours through entrainment, like flashing fireflies. This creates a second "social network" through which the more slowly evolving tastes of creatures can diffuse. A creature at a lower point of the pulse intensity is more likely to shift its taste to align with a neighbour whose intensity is currently higher. Visually, brighter creatures reach out with thin lines like feelers in proximity to their dimmer neighbours.

As creatures' movements cause neighbour relations to continually shift, this macro-structural network resembles a "liquid neural" network (Pinero & Solé 2019) of a second-order superorganism comprising around 40,000 neurons and 150,000 connections.

Movement also prevents the entrainment from stabilizing and the resulting intensity gradients lead to diffusions of taste through the population, clearly visible in the projected visualization as linear, circular and spiral waves of colour shift through the population.

4. ANALYSIS

We recorded all simulation activity, including agent and neural-network states over a 10-minute period, comprising over 7,000 simulation steps, in which each of 4096 creatures' networks activates (30 million activations in total).

Lifespan data verified our strategy to minimize selection pressure. More than half the population survived the entire period, and likely would have persisted far longer. Only around 3,000 times (once per 10,000 steps) did any creature's well-being fall below the viability threshold, causing them to be replaced by new creatures; and the majority of these events occurred within the first half-minute of the simulation, likely due to a lack of trails to discover.

Exchange of networks was far more fecund: over 10 minutes, around 244,000 events occurred in which a creature copied the network of a neighbour, with a 10% chance of applying a mutation. Examples of initialized and evolved networks in Infranet are shown in Figure 6.




Figure 6. Diagrams of creature neural networks recorded in Infranet. Weights and biases are omitted to focus on structural features. Left: All creatures have three inputs (bottom), four outputs (top), and begin with a fully connected network with no hidden nodes and logistic transfer functions. Right: An example neural network after living in the simulation for a period of time. This creature has developed two hidden neurons with different transfer functions (indicated by colour), and several neurons have developed lateral connections and feedback self-connections.

The average lifespan of a specific neural network instance is around 120 activations between each mutation. However, this rate appeared to be very bursty for each agent, with often a series of rapid adoptions followed by longer periods of stability; we presume this arises due to changes of neighbour communities and their relative well-being and pulse phase alignment.

As the 4096 creatures adopted and mutated neighbour's networks, over 10 minutes a total of 25,000 distinct mutated network structures were generated. We traced the lineages of inheritance of these adoptions, whose histogram revealed a very sharply exponential distribution in which most networks are rarely copied, while a very small number of networks lead very long lineages.

We ran another simulation for comparison in which adoptions occurred independently of creature dynamics, by selecting creatures and distributing adoption events randomly among the entire population. In general, the distribution in Infranet skews toward more prolific inheritance than in the randomized model. In Infranet, the most prolific candidate was adopted around 2,000 times (forming 1% of all recorded adoptions), compared to 1,000 in the random model. Ten networks in Infranet were half as prolific, compared to none in the random model. At the other extreme, around half of the networks in the random model were never copied at all, compared to 35% in Infranet. These figures identify that behaviour is certainly distinguished from well-mixed random history, and has a greater tendency toward inheritance, while at the same time not producing a monoculture.

We also ran a variant of the system in which adoption was implemented by crossover followed by mutation. The most striking difference found was divergence of network complexity: without

crossover, networks diverged to a range of up to 7-10 neurons and 7-17 connections. With crossover, networks diverged to 7-13 neurons and 11-100 connections. Since mutations are evenly balanced in probability of adding or removing nodes and connections, and selection pressure is low, this suggests an implicit neutral drift toward network complexity inherent to the crossover operation itself. Visually, the simulation with crossover also showed smaller neighbourhoods of similarity, but local variations lasted longer.

5. DISCUSSION

The use of machine learning to solve pre-stated problems articulated on pre-given data results in predictive models at risk of creating echo chambers.

Our observation is that an artificial intelligence mustn't work alone but must be open to and embedded within a world as a balanced complex adaptive system. Moreover, we want to see AIs that can do more than satisfy our demands by developing intrinsically, just as biological intelligence develops itself in the context of a living world. This makes artists (and developers, policymakers, etc.) into gardeners, an idea characteristic of the nascent days of artificial life research.

Our general approach to computational creativity is thus oriented to population-based systems that can rewrite their structures in coordination with their lived experiences, according to motivations that are intrinsic to the dynamics of the system. In *Infranet*, we allowed creatures' neural networks to change in structure over their lifetimes. They do this continuously, exchanging networks every few chirps by mutating their own neural structures according to the structures they hear from others. This is like the continuous genetic exchange of bacterial quasi-species, but perhaps also the contagious transfer of ideas. Moreover, since the creatures are moving and changing neighbours, the shifting population as a whole creates a much larger, "liquid" neural network of ever-changing structure, where new "ideas" can move in waves through the whole.

It seems that AI will become ubiquitous, smearing like new media to all aspects of life, into what Katherine Hayles described as a planetary cognitive ecology (Hayles 2020). We are therefore cautious about static performance measures and assumptions, as seen in supervised machine learning, that can only predict what is already known. We are cautious of the risk in working with 'low hanging fruit' of available data, readily measurable properties, known structures etc., as "exploiting too soon", and thus over-determining available paths of action and limiting futures. We

don't want to live in a future already paved by the past. Instead, we see an imperative for creativity oriented to an open future.

Now that 96% of the Earth's biomass is in service to us (Bar-On et al. 2018) we have reduced the wild unknown to rare pockets and islands. This seems characteristic of a more general tendency: directly or indirectly, we displace anything we can't measure or make use of, even without knowing what it might be, like dominating a market or controlling the gateways, and despite the fact that optimization breeds in fallibility. This points to the value of creativity beyond art; for example, again in thinking of the planetary cognitive ecology, and thinking beyond robustness toward antifragility (Equihua et al. 2020). We hope that a better understanding of creativity, beyond the human-centric, may point again to qualities of the wild. Perhaps including artificial creativity within the planetary cognitive ecology can bring us to the wild too.

6. REFERENCES

- Adam, S.P., Alexandropoulos, S.A.N., Pardalos, P.M. and Vrahatis, M.N. (2019) No free lunch theorem: A review. *Approximation and optimization*, pp.57-82.
- Antunes, R.F., Leymarie, F.F. and Latham, W. (2014) Two decades of evolutionary art using computational ecosystems and its potential for virtual worlds. *Journal of Virtual Worlds Research*, 7(3).
- Arriagada, L. (2020) CG-Art: demystifying the anthropocentric bias of artistic creativity. *Connection Science*, 32(4), pp.398-405.
- Bar-On, Y.M., Phillips, R. and Milo, R. (2018) The biomass distribution on Earth. *Proceedings of the National Academy of Sciences*, 115(25), pp.6506-6511.
- Bergson, H. (1998) *Creative evolution*. 1911. Trans. Arthur Mitchell. New York: Dover.
- Brooks, R. (2017) *What is it like to be a robot?* [Website]. <https://rodnebrooks.com/what-is-it-like-to-be-a-robot> (retrieved March 2021)
- Burkov, A. (2019) *The hundred-page machine learning book* (Vol. 1). Vancouver: Andriy Burkov.
- Cardoso, A., Veale, T. and Wiggins, G.A. (2009) Converging on the divergent: The history (and future) of the international joint workshops in computational creativity. *AI magazine*, 30(3), pp.15-15.
- Cuccu, G. and Gomez, F. (2011) April. When novelty is not enough. In *European Conference on the Applications of Evolutionary Computation*, pp. 234-243. Springer, Berlin, Heidelberg.

- Deleuze, G. (1990) *Bergsonism*. Trans. Tomlinson, H. & Habberjam, B. Zone Books.
- Equihua, M., Aldama, M.E., Gershenson, C., López-Corona, O., Munguía, M., Pérez-Maqueo, O. and Ramírez-Carrillo, E. (2020) Ecosystem antifragility: beyond integrity and resilience. *PeerJ*, 8, p.e8533.
- Fontana, W., 2006. The topology of the possible. In *Understanding Change*, pp. 67-84. Palgrave Macmillan, London.
- Hayles, N.K. (2020) *Unthought*. University of Chicago Press.
- Hertzmann, A. (2020) Computers do not make art, people do. *Communications of the ACM*, 63(5), pp.45-48.
- Jordanous, A. (2012) A standardised procedure for evaluating creative systems: Computational creativity evaluation based on what it is to be creative. *Cognitive Computation*, 4(3), pp.246-279.
- Kauffman, S. (2002) *Investigations*. Oxford University Press, USA.
- Krizhevsky, A., Sutskever, I. and Hinton, G.E. (2012) Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, pp.1097-1105.
- Lehman, J. and Stanley, K.O. (2011) Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, 19(2), pp.189-223.
- Lehman, J. and Stanley, K.O. (2013) Evolvability is inevitable: Increasing evolvability without the pressure to adapt. *PLoS one*, 8(4), p.e62186.
- Miikkulainen, R. (2020) Creative ai through evolutionary computation. In *Evolution in Action: Past, Present and Future*, pp. 265-269. Springer, Cham.
- Miyazaki, H. (2001) *Sen to Chihiro no Kamikakushi (Spirited Away)*, Directed by Hayao Miyazaki [Film], Japan, Studio Ghibli.
- Mouret, J.B. (2011) Novelty-based multiobjectivization. In *New horizons in evolutionary robotics*, pp. 139-154. Springer, Berlin, Heidelberg.
- Nguyen, A., Yosinski, J. and Clune, J. (2015) Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 427-436.
- Nguyen, A.M., Yosinski, J. and Clune, J. (2015) Innovation engines: Automated creativity and improved stochastic optimization via deep learning. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, pp. 959-966.
- Pinero, J. and Solé, R. (2019) Statistical physics of liquid brains. *Philosophical Transactions of the Royal Society B*, 374(1774), p.20180376.
- Ragot, M., Martin, N. and Cojean, S. (2020) AI-generated vs. Human Artworks. A Perception Bias Towards Artificial Intelligence? In *Extended abstracts of the 2020 CHI conference on human factors in computing systems*, pp. 1-10.
- Roselli, D., Matthews, J. and Talagala, N. (2019) Managing bias in AI. In *Companion Proceedings of The 2019 World Wide Web Conference*, pp. 539-544.
- Stanley, K.O. and Miikkulainen, R. (2002) Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2), pp.99-127.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I. and Fergus, R. (2013) Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*.
- Wagenaar, T (2017). *Neataptic*. <https://github.com/wagenaartje/neataptic> (retrieved March 2021)
- Wolpert, D.H. and Macready, W.G. (1995) *No free lunch theorems for search*. Technical Report SFI-TR-95-02-010, Santa Fe Institute.
- Woolley, B.G. and Stanley, K.O. (2011) On the deleterious effects of a priori objectives on evolution and representation. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pp. 957-964.