Video Games as a Testbed for Open-Ended Phenomena

Sam Earle
Tandon School of Engineering
New York University
Brooklyn, USA
sam.earle@nyu.edu

Julian Togelius
Tandon School of Engineering
New York University
Brooklyn, USA
julian@togelius.com

L. B. Soros
Cross Labs
Cross Compass, Ltd.
Tokyo, Japan
lisa.soros@cross-compass.com

***Abstract*—Understanding and engineering open-endedness, or the indefinite generation of novelty and complexity at arbitrary scales, has long been studied by implementing nature-inspired simulations specifically designed for artificial life studies. This paper argues that video games serve as a complementary domain for research on open-endedness. In support of this claim, experiments in this paper evaluate the effects of age-based and spatial destructive events in two game domains: an interactive Game of Life and the city-building game SimCity. These games are played by a neural-network-controlled gameplay agent trying to maximize reward. Results indicate that experiments with SimCity are more likely to identify statistically significant differences in complexity as a result of applied destructive events, highlighting the utility of this game domain for studying artificial life phenomena.**

I. Introduction

Identifying mechanisms that enable open-endedness is a fundamental question for the artificial life (ALife) community [1]. In fact, this question has implications reaching beyond ALife, as [2] hypothesize it may be critical not only for achieving ever-complexifying digital evolution, but also for creating a path to artificial general intelligence. The basic goal of open-endedness research is to make complexity evolve indefinitely and at arbitrary scales.[1] Towards this goal, it is instructive to understand how complexity and robustness change over time in response to phenomena observed in nature. This paper focuses on catastrophic destructive events in particular, which have proven capable of both spurring evolutionary change and reversing evolutionary expansions of biological diversity [4], but are not well-studied in the context of artificial life.

The history of ALife is rich with systems designed to achieve open-endedness, including virtual worlds such as Tierra [5], Avida [6], Polyworld [7], Geb [8], Division Blocks [9], Evosphere [10], and Chromaria [11]. All of these virtual worlds can be thought of as zero-player video games. Though such simulations explicitly designed to recapitulate features of the natural world have historically served as experimental domains for open-endedness, other types of domains could also serve as viable testbeds for empirical studies of the phenomenon. In particular, this paper argues that certain video games incorporate some form of open-ended dynamics, and that such domains could complement the existing catalogue of ALife worlds.

The goal of the experiments presented in this paper is to use open-ended video game domains to understand how periodic destructive events of various types affect the complexity of evolving spatial systems. Note that *catastrophe* and *destructive event* will be used synonymously in this paper. First, prior work on destructive artificial evolution is reviewed, followed by an introduction to open-endedness and emergence in games. The two domains used for empirical study are then introduced along with the gameplay agent and experimental parameters. The effects of age-based and spatial destructive events on map complexity are then analyzed for three different game sizes and three different destructive event frequencies.

II. Background

This section first introduces open-endedness research in artificial life, then studies of destructive events in artificial evolution, which have historically been studied in the context of evolutionary algorithms rather than in ALife simulations. It then provides a theoretical grounding for the video game domains used for empirical study.

A. Artificial life and open-ended evolution

The field of Artificial Life, or ALife, in general aims to simulate and synthesize living systems at all scales, from physics to chemistry to biology (with many possible levels of description in between). Instead of being tied explicitly to Earth-like systems, however, artificial life researchers seek to find viable abstractions of the processes that lead to systems *like* those observed on earth. Understanding open-ended evolutionary processes, which can be seen as a specific type of abstract creative process, is a particularly important major challenge for the field of artificial life with a rich research history. Framed imprecisely, the goal of open-ended evolution is to construct a system that could produce interesting things forever. Of course, one of the difficulties here is that what is "interesting" is hard to define precisely and objectively;

This work was supported by the National Science Foundation (Award number 1717324 - RI: Small: General Intelligence through Algorithm Invention and Selection).

[1]For a more thorough and nuanced discussion of OEE definitions, see [3], Chapter 2.

what counts as interesting varies from domain to domain and observer to observer.

The term *open-ended evolution* (OEE) was coined by the ALife community to describe processes in the spirit of natural evolution. Though the idea has been actively explored since its conception, its precise definition is still a contentious topic without a universally satisfying resolution. For example, OEE has been variously described as:

- "a process in which there is the possibility for an indefinite increase in complexity" [12],
- "a system in which components continue to evolve new forms continuously, rather than grinding to a halt when some sort of 'optimal' or stable position is reached" [13],
- a system in which the "number of possible types by far exceeds the number of individuals (copies, sequences, etc.) in a plausible (realistic) population" [14], or
- the "on-going and indefinitely creative production of significantly new kinds of adaptive responses to significantly new kinds of adaptive challenges and opportunities" [15].

This list illustrates the extent to which there is disagreement about the salient features of evolution. The characterization of evolution as continuously producing novel forms is one of the more popular definitions [16–18]. It should be noted that these definitions are not necessarily mutually exclusive.

B. Destructive events in artificial evolution

Results from a few recent studies of evolutionary algorithms suggest that intentionally removing individuals from the population can improve algorithm performance by increasing evolvability, where *evolvability* can mean either response to certain kinds of adaptive challenges [19–21] or potential for creative innovation [22–24]. Lehman and Miikkulainen [25] test the hypothesis that extinction events indirectly select for the ability to proliferate through vacated niches, thereby speeding up evolution. This work stems from Palmer and Feldman's earlier studies on indiscriminate extinctions in spatially segregated populations [26]. Beyond extinction events, Veenstra et al. [27] investigate the impact of *mortality* on evolution's ability to traverse search spaces while avoiding non-functional regions. In this paradigm, inspired by biological death, individuals are removed from the population after a fixed amount of time regardless of reward.

C. Open-ended games

Taking inspiration from the concept of emergence as used in complexity science and artificial life, and particularly as dichotomized into strong and weak forms by [28], Juul [29] formulated a theory of *open* versus *closed* games, also framed as *emergent* versus *progressive*. (Note that *games* in this context refers generally to playable board and video games.) In an attempt to clarify what is meant by games of *emergence* using Juul's term, Soler-Adillon [30] presents case studies of two games each capable of different degrees of emergence. Conway's Game of Life is selected as representative of the notion of emergence from the complexity science perspective, and SimCity is chosen as representative of an emergent game. Empirically evaluating the functional difference between these two kinds of systems of emergence thus has implications for both game studies and complexity studies.

III. Methodology

The experiments in this paper investigate the impacts of destructive events on reward and complexity of game states achieved by a reinforcement-learning agent playing adaptations of the representative emergent systems studied by [30].

A. Domain 1: Interactive Game of Life

The Game of Life (GoL) is a classic cellular automaton composed of cells occupying a discrete 2D grid, where each cell can either be "dead" or "alive." At each timestep, each cell's state is recalculated based on how many of its neighbors are dead or alive. This canonical implementation can also be thought of as a zero-player game [31], as can many ALife worlds. The experimental domain in this paper is an interactive adapted version that allows an automated game-playing agent to choose a single cell of the automaton and flip it. Pursuant to GoL cell update rules, the player's choices can have cascading effects on the rest of the evolving system.

B. Domain 2: Gym-city, a SimCity1 reinforcement learning environment

SimCity (Figure 1) is a city-building game designed by Will Wright and loosely based on system dynamicist Jay Wright Forrester's theory of urban dynamics [32]. In this single-player game, players spend money (which is a finite resource in the real game, but not in the test domain) to build cities composed of residential, commercial, and industrial zones on a fixed 2D grid. These zones are then procedurally filled with buildings such as houses (which increase the maximum population possible in the city), offices, and factories. However, zones will only be filled in if there is a need for development; there is no need to build more houses than the number of people in the city, for instance. In this way, SimCity is akin to a many-dimensional cellular automaton with an added interactive component (i.e. the human player, who is able to add and remove city components). In addition to zones, the player can construct growth-enabling components such as power lines and transportation infrastructure. The goal of the test environment is eventually to reach targets in certain city-wide metrics (including population, traffic, and mayor rating) after a fixed amount of time.

As noted by [33], "by coupling cellular automata and system dynamics style simulations, behavior [in SimCity] unfolds at both micro and macro scales, giving the microworld both finely detailed texture as well as an overarching organic coherence." Cities (both real and simulated) can also be thought of in terms of self-organization, exhibiting emergent features such as congestion and segregation[2] [34]. This game thus provides an excellent and novel domain for empirically investigating phenomena observed in real-world complex systems,

[2]The authors would like to clarify that not all emergent features are necessarily *good* features.

Fig. 1: **SimCity, an open-ended city-building simulation game.** The (human) player must place industrial, commercial, and residential zones in configurations that encourage population growth.

complementing traditional artificial life worlds and software platforms.

Following the API of OpenAI Gym[3], which provides portable environments for comparing reinforcement algorithms, Earle [35] recently re-packaged SimCity1 (whose code is freely available online) as a reinforcement learning (RL) environment nicknamed gym-city. This domain is used for all city-building experiments in this paper. Note that in the RL domain, money is not a finite resource (though money could be limited for future experiments).

C. Game-playing agent

Importantly, there is no evolutionary algorithm implemented for the purpose of the experiments reported in this paper. Instead, an artificial neural network (ANN) is trained to build, assuming the role that a human player normally would. The ANN is then responsible for planning the macro-scale city structure, with micro-level population dynamics governed by the aforementioned CA rules.

The ANN used in this study is a special type of convolutional neural network called a *fractal network*, which was specially designed to take advantage of spatial relationships in 2D spaces at multiple scales [35, 36]. A controllable RL agent is trained, as in [37], to achieve target values in certain high-level metrics in two CA-based domains: Conway's Game of Life (GoL) and SimCity. In each domain, the agent observes a one-hot encoded image of the map (observing tile-states rather than pixels in the case of SimCity) and builds or deletes one cell or structure on the map per step, followed by one or several ticks of the CA or city-simulation, respectively. In GoL, the sole target metric is net population in terms of the number of living cells on the board. In SimCity, there are multiple target metrics: the city's net residential, commercial, industrial, and traffic populations; the mayor rating; and the number of power plants on the map. We use an algorithm modeling absolute learning progress with gaussian mixture models (ALP-GMM) [38] to vary these target metrics and their weights over the course of training such that the agent's learning progress is maximised. These targets and weights are embedded in the agent's observation, allowing it to associate diverse sub-policies with various targets.

During inference in SimCity, high targets are provided for all 4 population types (roughly the maximum possible for each on a 16×16–tile map), a target of 1 power plant, and a 100% approval rating. Traffic has a low weight, making it less likely to be advantageous for the agent to exploit the game's traffic engine and produce excessive traffic by covering the map with road while providing only a few distant zones. The weight for residential population is similarly small, making the agent less prone to produce a large unemployed population via a repetitive, dense pattern of residential zones. The agent is thus encouraged to build with a variety of zones, leading to inhibitory adjacency effects between zone types, which can only be addressed by more complex city layouts; and global population dynamics, which make potential solutions more fragile and less flexible (i.e. when scaled to larger maps, they are prone to violent global population in- and out-fluxes).

The agent plays on maps produced by the game's built-in procedural terrain generator, and randomly scattered with a few additional structures, some of which cannot be deleted by the agent.

In Game of Life, the target population (number of "alive" cells corresponds to a densely-filled small map. In random initial layouts, each cell begins alive with 20% probability. In both domains, population targets do not scale with map size, so that on larger maps, it is in the agent's best interest to fill out the map only partially, providing a novel generalization challenge.

To produce the data in Figures 2–3, the best-performing column of the trained fractal network is sampled to produce actions, and the agent plays on a land-only map for 20 episodes of 1,000 steps, after which reward and compressibility tend to converge. On the largest maps, however, SimCity layouts nevertheless continue to evolve for longer, so Figure 5 renders snapshots of gameplay from 3,000–step episodes.

D. Experimental setup

In both domains[4], the gameplay agent plays on three different map scales: 16×16, 32×32, and 64×64. Each run lasts 2000 timesteps. In control runs, the agent is allowed to play without interference. In all other runs, one of the following types of destructive events occurs at fixed intervals (with only one type of destructive event occurring in a single run):

- **Random**: structures are selected for deletion at random, without replacement
- **Spatial**: as a metaphor for catastrophic geological events, a tile is selected at random from the map, and, spiralling

[3]https://gym.openai.com/

[4]Source code is available at https://github.com/smearle/gym-city

out from this central tile, any structure occupying a tile along this path is deleted

- **Age**: as a metaphor for biological mortality, structures are deleted in decreasing order of their age (oldest first)

Every step, the raw pixel map is converted to a tile map, and the type of structure occupying that tile determines its color according to the following functional categories: residential, commercial, industrial, infrastructure (transport, services), power, and other (natural structures, disaster-related objects like fire or radioactive waste). JPEG compression is then applied, giving an approximation of the complexity of the constructed cities. The general approach of using compressibility as complexity metric has been previously found to be comparable to more widely-used metrics, such as Shannon entropy, for detecting major evolutionary changes in cellular automata [39]. As noted by [40], compression-based proxies for complexity are not perfect metrics because extreme values correspond to complete order and complete disorder; it is unclear what intermediate value would be an indicator of "ideal" complexity. Nonetheless, JPEG compression is able to capture useful information about structured, spatial relationships.

IV. Results

A. Interactive Game of Life

Figure 2 illustrates the effects over time of destructive events at various frequencies and at various map sizes on an agent's gameplay in interactive Game of Life. On a small map at low frequency, facing the agent with destructive events leads to consistently greater average cumulative reward over the course of the episode. On these small maps where movement toward a denser map-state is rewarded, intermittent mass cell-death allows the prompts the agent to repeatedly repopulate the map, allowing it to climb out of local optima. Destructive events result in empty patches on the map which are reflected in steep drops in complexity (in terms of jpeg compressibility), though agents are able to return quickly to states with complexity matching that of states on an undisturbed map.

As the frequency of these events increases, average reward decreases. With excessive destruction, the agent has difficulty maintaining any living cells, as those remaining are likely to die from isolation. In the extreme case, this can lead to an irreversibly empty map. This tendency is reflected by downward trends in measures of both reward and complexity.

On a medium-sized map, the agent is rewarded for moving toward a population just less than the average on initialization. But, as on the small map, agents respond to catastrophe by fostering growth, this time at a detriment to reward. This tendency is repeated less pronouncedly on the large map, where destructive events have less impact overall. In both cases, increasing destructive frequency results in a loss of complexity and slight increase in reward as destructive events become more difficult to counteract with growth.

B. Gym-city

Figure 3 illustrates the effects of destructive interference on average attained reward and inverse JPEG compressibility over

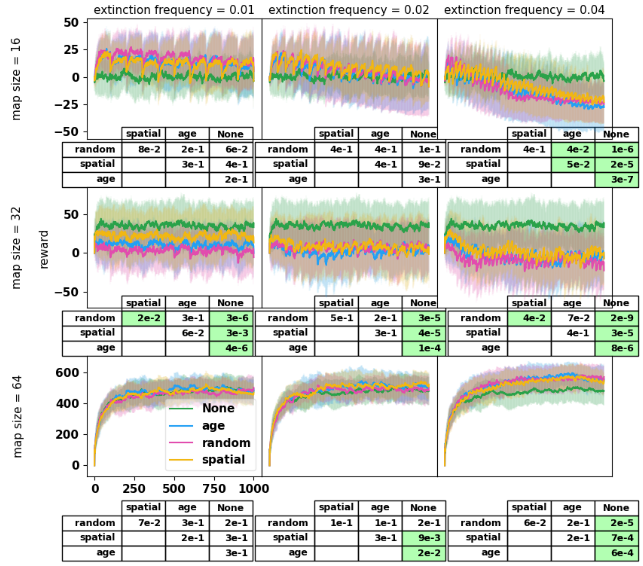

(a) Average reward

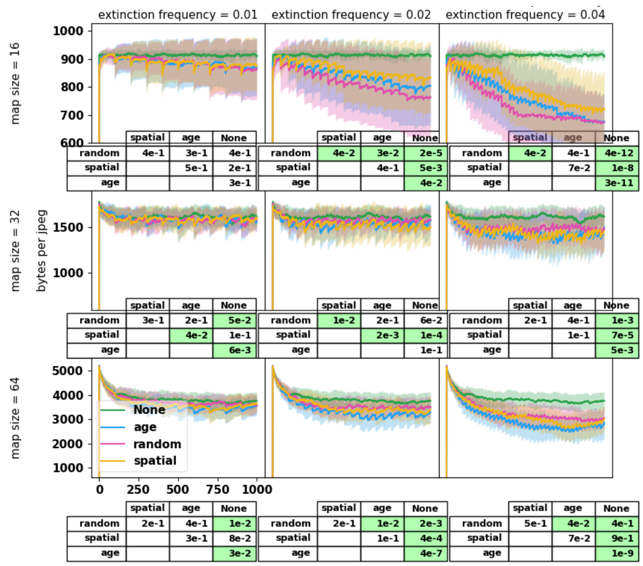

(b) Average inverse JPEG compressibility

Fig. 2: **Results for Interactive Game of Life.** Averages are calculated over 40 runs. The table below each subfigure gives p-value from a Mann-Whitney test. Green boxes indicate statistically significant differences ($p < 0.05$).

40 trials. Overall, catastrophic events tend to decrease, rather than increase, reward. This result is likely due to factors such as (intended) sudden destruction of populated zones, resulting in large dips in population, instability in neighboring zones (via disrupted desirability effects) and potential holes in the power grid. More frequent destruction events leads to more such destruction, and thus a greater negative affect on reward. In any case, in most cases, particularly on small and medium maps, age-based and spatial destruction are not significantly different from random destruction.

On large maps with infrequent distinction, however, spatial and random destruction lead to greater reward than on undisturbed maps. The city-building agent, having been trained on small maps, is unsuccessful on larger maps, leaving otherwise feasible residential neighborhoods disconnected from power (figure 4). Persistent (but non-overwhelming) destructive events disrupt this flawed stable state, and in forcing the agent to repair its broken map, lead to various other, more successful states, with dynamics that continue to evolve for several thousand time-steps. Spatial destruction (figure 5a) produces a constantly-migrating city center (comprising a power plant, a residential center, and an industrial periphery), while age-based destruction (figure **??**) produces a map that is constantly rippling or scanning upward, from left to right; destructive events, echoing past builds, leave horizontal bands of empty space on the map, which are repeatedly filled-in to produce various artefacts. In both cases, the irregular patterns resulting from destruction lead to greater and more robust connectivity (in terms of power) relative to the undisturbed map.

Complexity results in this domain are more informative with respect to evolutionary dynamics (and have a higher potential for cross-domain generalizability). Here, there are significant differences among most types of destructive events, though the large map is less able to discriminate between different destructive event types. On other map sizes, introducing destructive events tends to lead to increased file size after compression. Bearing in mind that compression size is not a perfect method for measuring complexity, it is not possible to say whether the achieved complexity is "ideal" or not, only that there is positively an increase. It should also be noted that there is a significantly higher variance between runs on large maps with lower frequencies of destructive events, suggesting that other parameter combinations are more predictable.

Of course, for the purpose of both validating the chosen metrics and gaining a better understanding of what's actually happening in the game, it is helpful to examine actual maps constructed by the game-playing agent. Figure 4 shows the final game state achieved by agent in the control scenario, where no catastrophes are introduced. Figure 5 depict representative runs subject to age-based and spatial catastrophes, respectively.

V. Discussion

In the Interactive Game of Life domain, age-based interventions are most likely to lead to all-out extinction, likely because older cells are more likely to remain alive than younger ones. An old cell is often part of some stable or partially oscillating multi-cellular configuration, while a new one is more likely to disappear on the next tick (or within a few) as it may be part of some travelling or oscillating configuration. The deletion of older cells would thus have a larger impact on future population counts (if only for the absence of the deleted cell itself from future time-steps). Spatial events are least likely to lead to all-out extinction, likely because, even when only a small number of cells are left after an event, these cells are

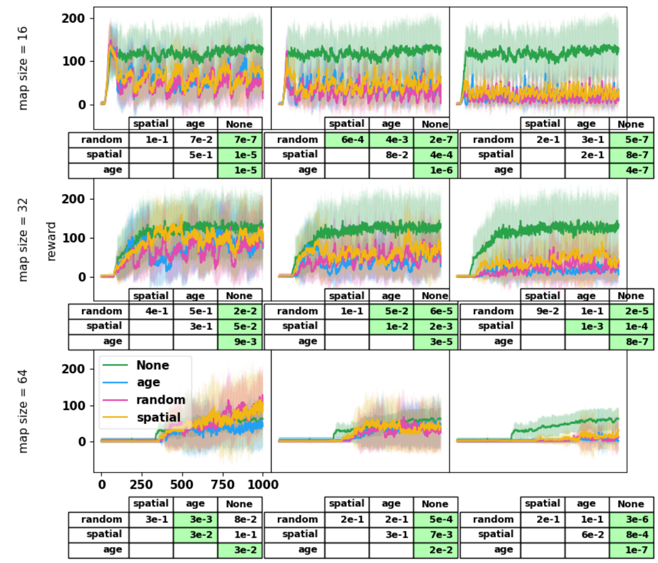

Map size 16:

	spatial	age	None
random	1e-1	7e-2	7e-7
spatial		5e-1	1e-5
age			1e-5

	spatial	age	None
random	6e-4	4e-3	2e-7
spatial		8e-2	4e-4
age			1e-6

	spatial	age	None
random	2e-1	3e-1	5e-7
spatial		2e-1	8e-7
age			4e-7

Map size 32:

	spatial	age	None
random	4e-1	5e-1	2e-2
spatial		3e-1	5e-2
age			9e-3

	spatial	age	None
random	1e-1	5e-2	6e-5
spatial		1e-2	2e-3
age			3e-5

	spatial	age	None
random	9e-2	1e-1	2e-5
spatial		1e-3	1e-4
age			8e-7

Map size 64:

	spatial	age	None
random	3e-1	3e-3	8e-2
spatial		3e-2	1e-1
age			3e-2

	spatial	age	None
random	2e-1	2e-1	5e-4
spatial		3e-1	7e-3
age			2e-2

	spatial	age	None
random	2e-1	1e-1	3e-6
spatial		6e-2	8e-4
age			1e-7

(a) Average reward

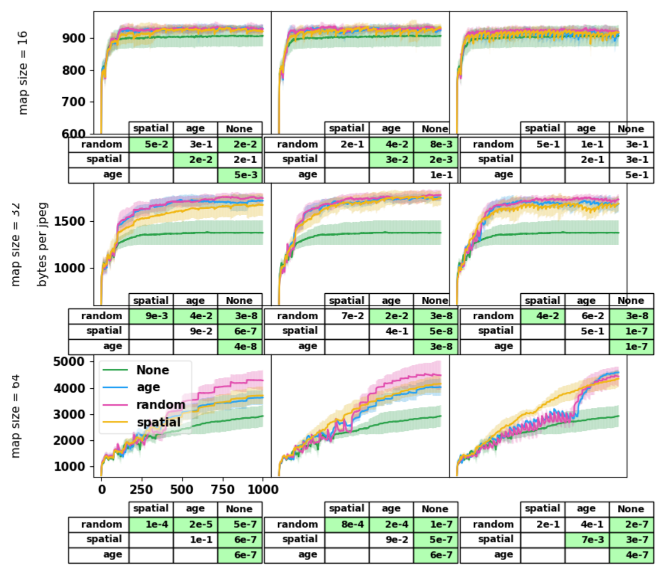

Map size 16:

	spatial	age	None
random	5e-2	3e-1	2e-2
spatial		2e-2	2e-1
age			5e-3

	spatial	age	None
random	2e-1	4e-2	8e-3
spatial		3e-2	2e-3
age			1e-1

	spatial	age	None
random	5e-1	1e-1	3e-1
spatial		2e-1	3e-1
age			5e-1

Map size 32:

	spatial	age	None
random	9e-3	4e-2	3e-8
spatial		9e-2	6e-7
age			4e-8

	spatial	age	None
random	7e-2	2e-2	3e-8
spatial		4e-1	5e-8
age			3e-8

	spatial	age	None
random	4e-2	6e-2	3e-8
spatial		5e-1	1e-7
age			1e-7

Map size 64:

	spatial	age	None
random	1e-4	2e-5	5e-7
spatial		1e-1	6e-7
age			6e-7

	spatial	age	None
random	8e-4	2e-4	1e-7
spatial		9e-2	5e-7
age			6e-7

	spatial	age	None
random	2e-1	4e-1	2e-7
spatial		7e-3	3e-7
age			4e-7

(b) Average inverse JPEG compressibility

Fig. 3: **Results for SimCity.** Averages are calculated over 40 runs. The table below each subfigure gives p-value from a Mann-Whitney test. Green boxes indicate statistically significant differences ($p < 0.05$). Destructive events generally lead to decreased reward, but more complexity (less compressibility) after 1000 timesteps. On large maps with infrequent destruction, they lead to increased reward.

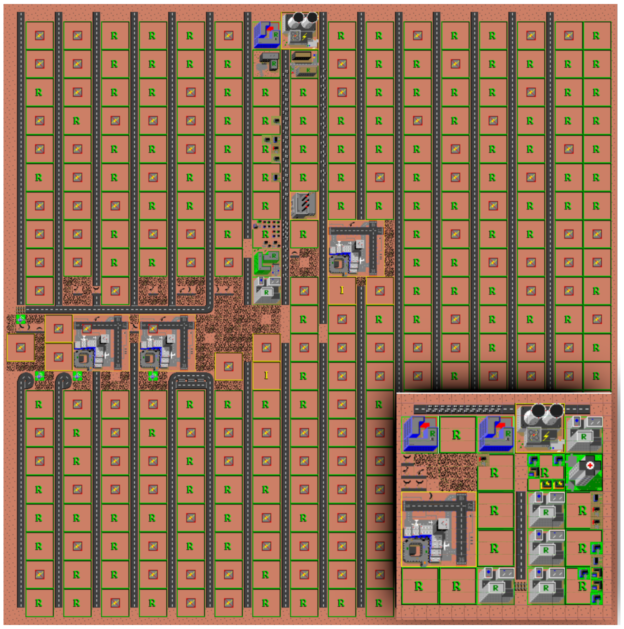

Fig. 4: **Stable state of a representative city built on large and small (lower-right cornder) tile maps by an RL agent, without destruction.** On a large map, the agent plans what would be an effective residential neighbourhood, but fails to connect it to power. This primarily residential pattern is also observed in control experiments on the small and medium maps.

more likely to be grouped together on the map than under different types of destruction, which may thin out populations more uniformly over the board.

With destruction in gym-city, the agent is more likely to deviate from the random start and use regular building patterns to achieve the target population. This is especially common in age-based destruction, where, during the first event, the oldest cells (and those to be deleted) are precisely those belonging to the random starting configuration.The resulting, more regular final configurations, are more compressible.

Interestingly, the differences between catastrophe types with respect to complexity tend to be more statistically significant in gym-city than in Interactive Game of Life. When considering the difference between emergence in the traditional complex systems sense and emergence in games, one might conclude that SimCity is a better testbed for empirical studies of discrete dynamical systems, or at least a complementary one that ALife researchers should consider adding to their digital toolboxes.

An important and interesting question to consider is how the concept of openness in games corresponds to notions of open-endedness in artificial life. At a high level, the goals are the same: to construct a source of unlimited innovation and opportunity for new strategies, or ways of being, that are discovered rather than designed. However, one aspect of open-ended evolution (in the ALife sense) that is *not* captured in at least Jesper Juul's framework is the concept of a shift in individuality or a major/phase transition [41], the lack of which may be a barrier to open-ended evolution [42].

Many open-ended games implicitly or explicitly support and encourage some kind of predetermined progression. This applies to putatively open-ended first-person games such as Minecraft (which in most game modes have some form of quests) and No Man's Sky (which has a very thin but existing questline, and a rather complex progression system). Even third-person open-ended "god games" have some strong steering mechanisms built in. For example, the games in the Civilization series have victory conditions, and most of them finish at a predetermined point. By virtue of being artefacts designed for a particular kind of human-computer interaction (which we for lack of a better expression can call interactive entertainment, or simply "gameplay"), these games offer restrictions, rewards and affordances that push gameplay in a particular direction. It is thus not clear exactly what would constitute an "unbiased" open-ended game. Yet, open-ended evolution (both in nature and in algorithmic form) also necessarily entails some sort of progression, otherwise it would degrade to exploration of a relatively simple possiblity space.

Capturing the spirit of open-endedness in an algorithmic process has long been a focus of the artificial life community. While the exact nature of open-endedness research is difficult to articulate precisely, it essentially aims to replicate the chain of endless innovation achieved by biological evolution in nature. The games community, too, has similarly long aimed to devise such creative algorithmic processes for the purpose of i.e. procedural content generation. The 2016 space exploration game No Man's Sky is an example of such an endeavor; the world consists of over 18 quintillion automatically generated planets each rife with reasonably unique flora and fauna. As described in the New Yorker, "Because the designers are building their universe by establishing its laws of nature, rather than by hand-crafting its details, much about it remains unknown, even to them." [43] Unfortunately, the core PCG algorithm was incapable of producing new *kinds* of things, resulting in what Kate Compton calls the *10,000 Bowls of Oatmeal Problem* [44] wherein generated artifacts may be mathematically distinct (like oats at different positions and orientations in a bowl), but are not distinct in any perceivable or meaningful way.

In artificial life, the salient question is about what constitutes an *individual* in any particular system. In many virtual ALife worlds, an individual is an embodied agent that moves around in a fixed world sometimes with resources that must be collected and predators that must be avoided. This level of description or individuality is commensurate with biology in the natural world and is reminiscent of many video games. Innovation, then, occurs with respect to morphologies and behavioral strategies. It is also possible to define individuality at a much lower level, such as that of an atom or molecule, which is the approach in artificial chemistry. Frankly, here's where things get really interesting for games in particular, as it's less clear what constitutes a fundamental "unit" of gameplay that could be combined in an indefinitely scalable

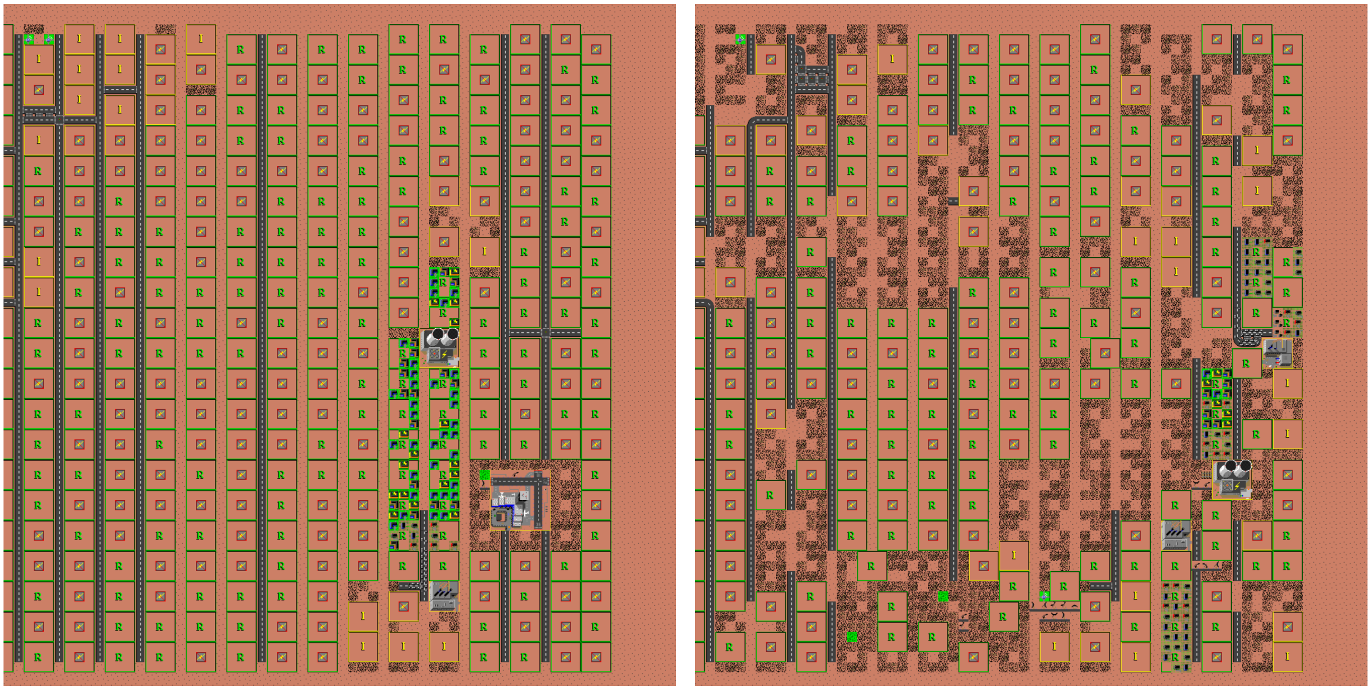

timestep = 1000 timestep = 3000

(a) **Snapshots from a sample gym-city run on a** 64×64 **map with spatial catastrophes occurring every 100 steps.** The usual grid is left incomplete, with lateral connections appearing at destruction sites. Over large time-scales, the residential core and industrial periphery migrate across the map as a result of continued disturbance.

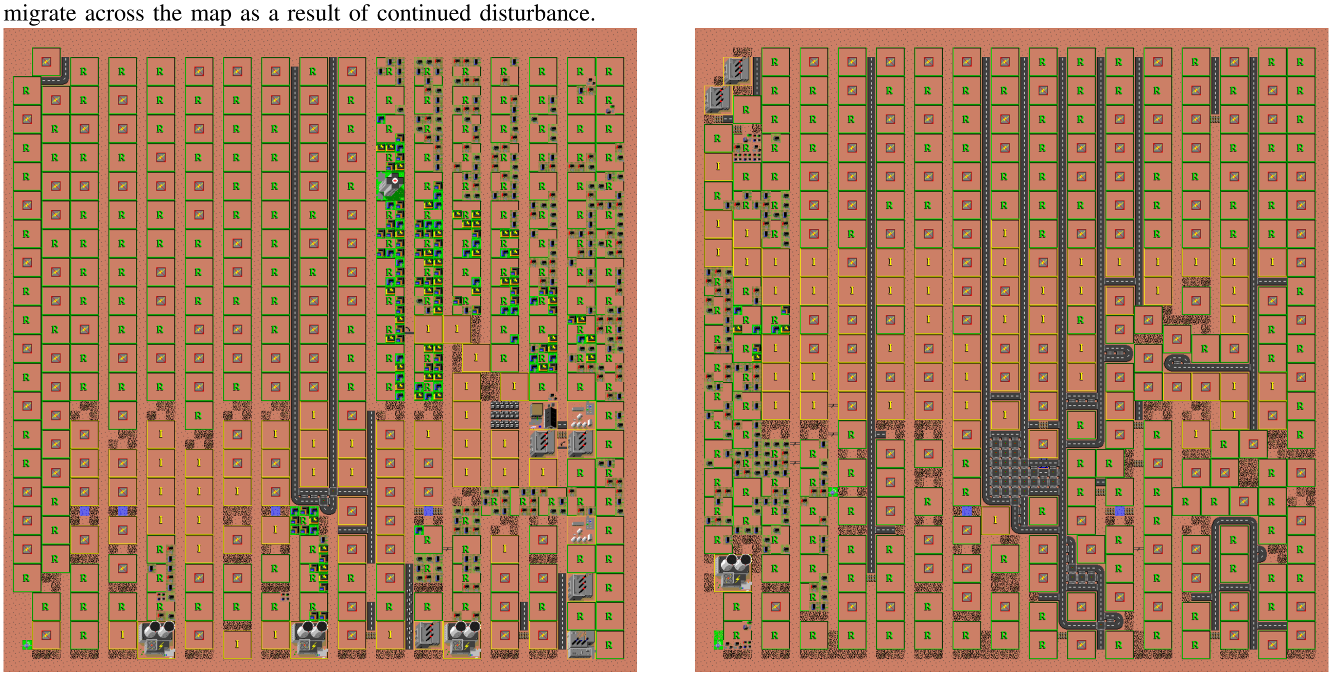

timestep = 1000 timestep = 3000

(b) **Snapshots from a sample gym-city run on a** 64×64 **map with age-based catastrophes occurring every 100 steps.** Large scale and dynamic patterns emerge as the agent's own actions leave a disruptive wake. A band of industrial zones repeatedly "scans" the map upward (mirroring the upward, left-to-right build pattern of the initial residential zones), and a growing cluster of road appears in the center of the map.

Fig. 5: Both spatial and age-based destruction allow an RL agent to explore more complex and diverse states over large timescales and maps. Age-based catastrophes lead to reward comparable with an undisturbed agent.

way, as in the natural world. Still, we haven't yet been able to overcome what might be called the *10,000 bowls of Primordial Soup Problem*.

At the beginning of this paper, we argued that video games can potentially serve as complements to open-endedness research in ALife worlds. Clearly, there is some fuzzy ground at exactly the locus of simulation games such as SimCity, where the goal is to control and environment rather than to participate in it immersively. If nothing else, the fact that compressibility metrics capable of detecting phase transitions in natural systems can be applied to game states should make us consider more seriously where the boundaries between these disciplines truly lay and whether we can all benefit from each other's insights.

VI. Conclusion

This paper investigated the effects of age-based and spatial catastrophic events on cellular-automata-based video games capable of different degrees of emergence. Results showed that spatial catastrophes tend to have a more dramatic effect on system quality and complexity than age-based catastrophes, but that either type of destructive event has a stronger effect than random destructions or none at all. More broadly, these experiments provided an initial demonstration of video games as a complementary domain to artificial life for exploring phenomena related to open-endedness.

References

[1] M. A. Bedau, J. S. McCaskill, N. H. Packard, S. Rasmussen, C. Adami, D. G. Green, T. Ikegami, K. Kaneko, and T. S. Ray, "Open problems in artificial life," *Artificial life*, vol. 6, no. 4, pp. 363–376, 2000.

[2] K. O. Stanley, J. Lehman, and L. Soros, "Open-endedness: The last grand challenge you've never heard of," December 2017.

[3] L. B. Soros, "Necessary conditions for open-ended evolution," Ph.D. dissertation, University of Central Florida, 2018.

[4] P. R. Grant, B. R. Grant, R. B. Huey, M. T. Johnson, A. H. Knoll, and J. Schmitt, "Evolution caused by extreme events," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 372, no. 1723, p. 20160146, 2017.

[5] T. S. Ray, "An approach to the synthesis of life," in *Proc. of Artificial Life II*. Addison-Wesley, 1992, pp. 371–408.

[6] C. Ofria and C. O. Wilke, "Avida: A software platform for research in computational evolutionary biology," *Artificial life*, vol. 10, no. 2, pp. 191–229, 2004.

[7] L. Yaeger, "PolyWorld: Life in a new context," *Proc. Artificial Life*, vol. 3, pp. 263–263, 1994.

[8] A. D. Channon and R. I. Damper, "Towards the evolutionary emergence of increasingly complex advantageous behaviours," *International Journal of Systems Science*, vol. 31, no. 7, pp. 843–860, 2000.

[9] L. Spector, J. Klein, and M. Feinstein, "Division blocks and the open-ended evolution of development, form, and behavior," in *Proc. of the 9th annual conf. on Genetic and evolutionary computation*. ACM, 2007, pp. 316–323.

[10] T. Miconi and A. Channon, "A virtual creatures model for studies in artificial evolution," in *The 2005 IEEE Congress on Evolutionary Computation*, vol. 1. IEEE, 2005, pp. 565–572.

[11] L. B. Soros and K. O. Stanley, "Identifying minimal conditions for open-ended evolution through the artificial life world of chromaria," in *Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*. Cambridge, MA: MIT Press, 2014, pp. 793–800.

[12] K. Ruiz-Mirazo, J. Umerez, and A. Moreno, "Enabling conditions for 'open-ended evolution'," *Biology and Philosophy*, vol. 23, no. 1, pp. 67–85, 2008.

[13] T. Taylor, "From artificial evolution to artificial life," Ph.D. dissertation, University of Edinburgh, 1999.

[14] J. Maynard-Smith and E. Szathmary, *The Major Transitions in Evolution*. New York, NY: Oxford University Press, 1995.

[15] M. Bedau, "Four puzzles about life," *Artificial Life*, no. 4, pp. 125–140, 1998.

[16] J. Holland, *Adaptation in Natural and Artificial Systems*. Cambridge, MA: MIT Press, 1975.

[17] R. K. Standish, "Open-ended artificial evolution," *International Journal of Computational Intelligence and Applications*, vol. 3, no. 02, pp. 167–175, 2003.

[18] J. Lehman and K. O. Stanley, "Abandoning objectives: Evolution through the search for novelty alone," *Evolutionary Computation*, vol. 19, no. 2, pp. 189–223, 2011. [Online]. Available: http://eplex.cs.ucf.edu/papers/lehman_ecj11.pdf

[19] J. Reisinger, K. O. Stanley, and R. Miikkulainen, "Towards an empirical measure of evolvability," in *Proceedings of the 7th annual workshop on Genetic and evolutionary computation*, 2005, pp. 257–264.

[20] M. Pigliucci, "Is evolvability evolvable?" *Nature Reviews Genetics*, vol. 9, no. 1, pp. 75–82, 2008.

[21] J. Clune, J.-B. Mouret, and H. Lipson, "The evolutionary origins of modularity," *Proceedings of the Royal Society b: Biological sciences*, vol. 280, no. 1755, p. 20122863, 2013.

[22] G. P. Wagner and L. Altenberg, "Perspective: complex adaptations and the evolution of evolvability," *Evolution*, vol. 50, no. 3, pp. 967–976, 1996.

[23] M.-L. Dichtel-Danjoy and M.-A. Félix, "Phenotypic neighborhood and micro-evolvability," *Trends in Genetics*, vol. 20, no. 5, pp. 268–276, 2004.

[24] J. Lehman, B. Wilder, and K. O. Stanley, "On the critical role of divergent selection in evolvability," *Frontiers in Robotics and AI*, vol. 3, p. 45, 2016.

[25] J. Lehman and R. Miikkulainen, "Extinction events can accelerate evolution," *PloS one*, vol. 10, no. 8, 2015.

[26] M. E. Palmer and M. W. Feldman, "Spatial environmental variation can select for evolvability," *Evolution: International Journal of Organic Evolution*, vol. 65,

no. 8, pp. 2345–2356, 2011.
[27] F. Veenstra, P. G. de Prado Salas, J. Bongard, K. Stoy, and S. Risi, "Intrinsic mortality governs evolvability," in *Artificial Life Conference Proceedings*. MIT Press, 2018, pp. 242–249.
[28] M. A. Bedau, "Weak emergence," *Noûs*, vol. 31, pp. 375–399, 1997.
[29] J. Juul, "The open and the closed: Games of emergence and games of progression." in *CGDC Conf.*, 2002.
[30] J. Soler-Adillon, "The open, the closed and the emergent: Theorizing emergence for videogame studies," *Game Studies*, vol. 19, no. 2, 2019.
[31] S. Björk and J. Juul, "Zero-player games or: what we talk about when we talk about players," in *Philosophy of Computer Games Conference, Madrid*, 2012.
[32] J. W. Forrester, *Urban Dynamics*. Pegasus Communications, Inc., 1969.
[33] C. O. Gingold, "Play design," Ph.D. dissertation, UC Santa Cruz, 2016.
[34] J. Portugali, *Self-organization and the city*, ser. Springer Series in Synergetics. Springer-Verlag Berlin Heidelberg, 2000.
[35] S. Earle, "Using fractal neural networks to play simcity 1 and conway's game of life at variable scales," *arXiv preprint arXiv:2002.03896*, 2020.
[36] G. Larsson, M. Maire, and G. Shakhnarovich, "Fractalnet: Ultra-deep neural networks without residuals," *arXiv preprint arXiv:1605.07648*, 2016.
[37] S. Earle, M. Edwards, A. Khalifa, P. Bontrager, and J. Togelius, "Learning controllable content generators," 2021.
[38] R. Portelas, C. Colas, K. Hofmann, and P.-Y. Oudeyer, "Teacher algorithms for curriculum learning of deep rl in continuously parameterized environments," *arXiv preprint arXiv:1910.07224*, 2019.
[39] A. Adamatzky and J. Jones, "On using compressibility to detect when slime mould completed computation," *Complexity*, vol. 21, no. 5, pp. 162–175, 2016.
[40] H. Cisneros, J. Sivic, and T. Mikolov, "Evolving structures in complex systems," in *2019 IEEE Symposium Series on Computational Intelligence (SSCI)*, 2019, pp. 230–237.
[41] J. M. Smith and E. Szathmary, *The major transitions in evolution*. Oxford University Press, 1997.
[42] E. Dolson, A. Vostinar, and C. Ofria, "What's holding artificial life back from open-ended evolution," *The Winnower*, vol. 5, p. e152418, 2015.
[43] R. Khatchadourian, "World without end," *The New Yorker*, May 2015.
[44] K. Compton, "So You Want To Build A Generator," https://galaxykate0.tumblr.com/post/139774965871/so-you-want-to-build-a-generator, 2016, [Online; accessed 02-April-2021].