# Understanding Fraternal Transitions in Individuality

Matthew Andres Moreno and Charles Ofria

BEACON Center, Michigan State University, East Lansing, MI 48824
mmore500@msu.edu

## Abstract

The emergence of new replicating entities from the union of existing entities represent some of the most profound events in natural evolutionary history. Facilitating such evolutionary transitions in individuality is essential to the derivation of the most complex forms of life. As such, understanding these transitions is critical for building artificial systems capable of open-ended evolution. Alas, these transitions are challenging to induce or detect, even with computational organisms. Here, we introduce the DISHTINY (DIStributed Hierarchical Transitions in IndividualitY) platform, which provides simple cell-like organisms with the ability and incentive to unite into new individuals in a manner that can continue to scale to subsequent transitions. The system is designed to encourage these transitions so that they can be studied: organisms that coordinate spatiotemporally can maximize the rate of resource harvest, which is closely linked to their reproductive ability. We demonstrate the hierarchical emergence of multiple levels of individuality among simple cell-like organisms that evolve parameters for manually-designed strategies. During evolution, we observe reproductive division of labor and close cooperation between cells, including resource-sharing, aggregation of resource endowments for propagules, and emergence of an apoptosis response to somatic mutation. While a few replicate populations evolved selfish behaviors, many evolved to direct their resources toward low-level groups (behaving like multi-cellular individuals), and many others evolved to direct their resources toward high-level groups (acting as larger-scale multi-cellular individuals). Finally, we demonstrated that genotypes that encode higher-level individuality consistently outcompete those that encode lower-level individuality.

## Introduction

Artificial Life researchers design systems that exhibit properties of biological life in order to better understand their dynamics and, often, to apply these principles toward engineering applications such as artificial intelligence (Bedau, 2003). Studies of evolution have been of particular interest to the community, especially in regard to how organisms are produced with increasing sophistication and complexity (Goldsby et al., 2017). This particular issue is often described as "open-ended evolution." Although precise definitions and measures of open-ended evolution are still being established, this term is generally understood to refer to evolving systems that exhibit the continued production of novelty (Taylor et al., 2016). Evolutionary transitions in individuality, which are key to the complexification and diversification of biological life (Smith and Szathmary, 1997), have been highlighted as key research targets with respect to the question of open-ended evolution (Ray, 1996; Banzhaf et al., 2016). In an evolutionary transition of individuality, a new, more complex replicating entity is derived from the combination of cooperating replicating entities that have irrevocably entwined their long-term fates (West et al., 2015). Eusocial insect colonies and multicellular organisms exemplify this phenomenon (Smith and Szathmary, 1997). Like the definition of open-ended evolution, the notion of what constitutes an evolving individual is not concretely established. Commonly indicated features include: close coordination and cooperation, reproductive division of labor, reproductive bottlenecks, and loss of ability to replicate independently (Ereshefsky and Pedroso, 2015; Bouchard, 2013).

Major challenges in studying evolutionary transitions in individuality include (1) determining the environmental conditions that will promote such a transition and then (2) recognizing that a transition has occurred. In order to begin exploring transitions in individuality, we must devise a system in which we expect such transitions to occur repeatably and in a detectable manner. Once we can consistently induce and observe evolutionary transitions in individuality, we may subsequently proceed to relax aspects of such a system to explore in greater detail what conditions are necessary to induce transitions and how transitions can be detected. For now, we will focus on these initial goals.

To this end, we introduce the DISHTINY (DIStributed Hierarchical Transitions in IndividualitY) platform, which seeks to achieve the evolution if transitions in individuality by explicitly registering organisms in cooperating groups that coordinate spatiotemporally to maximize the harvest of a resource. Detection of such a transition in DISHTINY is accomplished by identifying resource-sharing and reproductive division of labor among organisms registered to the same cooperating group. Our system is designed such that hierarchal transitions across an arbitrary number of levels of

individuality can be selected for and meaningfully detected. We have focused this system on a rigid form of major transition using simple organisms, but the underlying principles can be applied to a wide range of artificial life systems. Furthermore, DISHTINY is decentralized and amenable to massive parallelization via distributed computing. We believe that such scalability — with respect to both concept and implementation — is an essential consideration in the pursuit of artificial systems capable of generating complexity and novelty rivaling that of biological life via open-ended evolution (Ackley and Cannon, 2011; Ackley, 2016).

## Methods

In order to demonstrate that the DISHTINY platform selects for detectable hierarchical transitions in individuality, we performed experiments where cell-like organisms evolved parameters to control manually designed behaviors such as resource-sharing, reproductive decision-making, and apoptosis. We will first cover the design of the DISHTINY platform and then describe the simple cell-like organisms we used to evaluate the platform.

### DISHTINY

DISHTINY allows cell-like organisms to replicate across a toroidal grid. As cells reproduce, they can optionally share signaling channels with their offspring. Over discrete timesteps ("updates"), the cells can collect a continuous-valued resource, either hoarding it or sharing it with others on a signaling channel. Once sufficient resource has been accrued, cells may pay $8.0$ resource to place a daughter cell on an adjoining tile of the toroidal grid (i.e., reproduce), replacing any existing cell already there.

As shown at the top of Figure 1, resources appear at a single point and spread in diamond-shaped waves. Each update the resource wave advances one grid tile outward, disappearing when it reaches a predefined limit. Cells must be in a costly "activated" state to collect resource as it passes. The cell at the starting position of a resource wave is automatically activated, and will send the activate signal to neighboring cells on the same signaling channel. The newly activated cells, in turn, activate their own neighbors registered to the same signaling channel. Neighbors registered to other signaling channels do not activate. Each cell, after sending the activation signal, enters a temporary quiescent state so as not to reactivate from the signal. In this manner, cells sharing a signaling channel activate in concert with the expanding resource wave. As shown Figure $1a, b$, the rate of resource collection for a cell is determined by the size and shape of of its same-channel signaling network; small or fragmented same-channel signaling networks will frequently miss out on resource as it passes by.

Each cell pays a resource cost when it activates. This cost is outweighed by the resource collected such that cells that activate in concert with a resource wave derive a net benefit. Recall, though, that resource waves have a limited extent. Cells that activate outside the extent of a resource wave or activate out of sync with the resource wave (due to an indirect path from the cell that originated the signal) pay the activation cost but collect no resource. Cells that frequently activate erroneously use up their resource and die. In our implementation, organisms that accrue a resource debt of $-11$ or greater are killed. This scenario is depicted in Figure $1c$.

In this manner, "Goldilocks" — not to small and not too big — signaling networks are selected for. Based on a randomly chosen starting location, resource wave starting points (seeds) are tiled over the toroidal grid such that the extents of the resource waves touch, but do not overlap. All waves start and proceed synchronously; when they complete, the next resource waves are seeded. This process ensures that selection for "Goldilocks" same-channel signaling networks is uniformly distributed over the toroidal grid.

Cells control the size and shape of their same-channel signaling group through strategic reproduction. Three choices are afforded: whether to reproduce at all, where among the four adjoining tiles of the toroidal grid to place their offspring, and whether the offspring should be registered to the parent's signaling channel or be given a random channel ID (in the range 1 to $2^{22}$). No guarantees are made about the uniqueness of a newly-generated channel ID, but chance collisions are rare.

Hierarchical levels are introduced into the system through multiple separate, but overlaid, instantiations of this resource wave/channel-signaling scheme. We refer to each independent resource wave/channel-signaling system as a "level." In our experiments, we allowed two resource wave/channel-signaling levels, identified here as level one and level two. On level one, resource waves extended a radius of four toroidal tiles. On level two they extended a radius of twelve toroidal tiles. On both levels, activated cells netted $+1.0$ resource from a resource wave, but suffered an activation penalty of $-5.0$ if no resource was available. Due to the different radii of resource waves on different levels, level one selects for small same-channel signaling networks and level two selects for large same-channel signaling networks.

Cells were marked with two separate channel IDs, one for level one and another for level two. We enforced hierarchical nesting of same-channel signaling networks during reproduction: daughter cells may inherit neither channel ID, just the level-two channel ID, or both channel IDs. Daughter cells may not inherit only the level-one channel ID while having a different level-two channel ID. The distribution of IDs across the level-two and level-one channels can be envisioned by analogy to political countries and territories. Each country (i.e., level-two channel network) may have one or many territories (i.e., level-one channel network). However, no territory spans more than one country. Figure 3 depicts hierarchically nested channel states at the end of three evo-

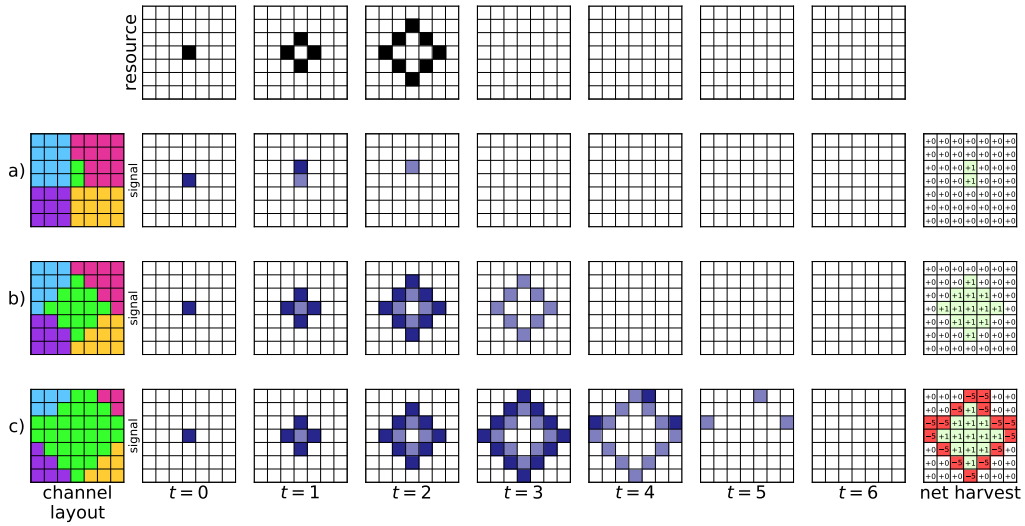


Figure 1: **Activation signaling, and net resource collection for three different channel configurations during a resource wave event.** At the top, a resource wave is depicted propagating over three updates and then ceasing for four updates (left to right). In row $a$, a small channel-signaling group (far left, in green) is activated; tracking the resource wave (middle) yields a small net resource harvest (far right). In row $b$, an intermediate-sized channel-signaling group yields a high net resource harvest. Finally, in row $c$, a large channel-signaling group incurs a net negative resource harvest. In rows $a$, $b$, and $c$, dark purple indicates the active state, light purple indicates the quiescent state, and white indicates the ready state.

lutionary runs.

Channel IDs enable straightforward detection of an evolutionary transition in individuality. Because common channel IDs may only arise systematically through inheritance, common channel IDs indicate a close hereditary relationship in addition to a close cooperative relationship. Because new channel IDs arise first in a single cell, same-channel signaling networks are reproductively bottlenecked, ensuring meaningful reproductive lineages at the level of the same-channel signaling network. To recognize an evolutionary transition in individuality, we therefore evaluate

1. Do cells with the same channel ID choose to share resources (e.g., cooperate)?

2. Is there division of reproductive labor between members of the same channel (e.g., do cells at the interior of a network cede reproduction to those at the periphery?)

If these conditions are met among cells sharing the same level-one channel, we can conclude that a first-level transition in individuality has occurred. Likewise, if these conditions are met among cells sharing the same level-two channel, we can conclude that a second-level transition in individuality has occurred.

## Organisms

We performed our experiments using cell-like organisms composed of 15 floating-point parameters, each controlling a specific strategy component pertinent to transitions in individuality (i.e., reproductive division of labor, resource pooling, apoptosis, propagule generation, and propagule endowment). These particular cell-like organisms are in no way inherent to the DISHTINY platform, but were merely developed to study transitions using as simple a models system as feasible. On reproduction, we applied mutation to each parameter independently with probability $0.00005$.

The **aversion parameters** ($\boldsymbol{A_1}$ and $\boldsymbol{A_2}$) allow cells to avoid reproducing over neighbors sharing the same signaling channel. Specifically, they control the probability that a cell declines to supplant a neighbor sharing the same level-one ($A_1$) or level-two ($A_2$) channel ID. If a cell declines to place its offspring in all four adjoining tiles, it decides not to reproduce. Mutation is performed by a redraw from the uniform distribution $U(-0.5, 1.5)$ clamped to the range $[0, 1]$.

The **resource allocation parameters** control the proportion of resources that go to the cell's stockpile ($\boldsymbol{P_c}$), its level-one channel's resource pool ($\boldsymbol{P_1}$), or its level-two channel's resource pool ($\boldsymbol{P_2}$). These parameters are initialized by a draw from $U(0.0, 1.0)$ and mutated by addition of a normal value drawn from $N(0.0, 0.2)$ with the result clamped to the

range $[0, 1]$. The set $P_c, P_1, P_2$ is always normalized to sum to 1.

Channel resource pools are identical to an organism's individual stockpile, except that any deficit is distributed evenly among the individual organism's stockpile. On every update, cells can spend from their individual stockpile to reproduce or from a channel pool, with priority given to cells nearest to the centroid of that pool's members. As such, pool-funded reproduction fills in a same-channel signaling network from the inside out and help produce diamond-shaped same-channel signaling networks. (Distance is measured using the taxicab metric.)

**Channel cap parameters $C_1$** and **$C_2$** regulate the size of same-channel signaling networks. When an organism reproduces, it checks the size of its level-one signaling network against $C_1$ and the size of its level-two signaling group against $C_2$. If neither cap is met or exceeded, then the organism will produce an offspring sharing both of its channel IDs. If only the $C_1$ cap is exceeded, then the organism will produce an offspring with new level-one channel ID but identical level-two channel ID. Finally, if the $C_2$ cap is exceeded, then the organism will produce an offspring with new IDs for both channels. These parameters are initialized by a draw from $U(0.0, 48.0)$ and mutated by addition of a value drawn from $N(0.0, 24.0)$ with the result clamped to be non-negative.

The **endowment parameters $E_c$**, **$E_1$**, and **$E_2$** determine the amount of resource provided to offspring. This endowment is paid as an additional cost by the cell stockpile (or same-channel resource pool) funding a reproduction. The full amount of the received endowment is divided between the daughter cell's stockpile, level-one same-channel resource pool, and level-two same-channel resource pool according to the offspring's resource allocation parameters. $E_c$ is the endowment amount paid to an offspring that shares both channel IDs of the parent; $E_1$ is the endowment paid to an offspring that shares just the level-two channel ID of the parent; and $E_2$ is the endowment paid to an offspring that shares neither the level-one nor level-two channel ID of the parent. Endowed resources help new-channel propagules to rapidly grow their signaling network in order to begin collecting resource at a rate competitive to other well-established same-channel signaling networks. These parameters are initialized by a draw from $U(0.0, 3.0)$ and are mutated by addition of a value drawn from $N(0.0, 10.0)$ with the result clamped to be non-negative.

Parameters $M_c$, $M_1$, and $M_2$ control the **apoptosis response to mutation**. Each time that a mutation occurs during reproduction, the mutated offspring attempts suicide with probability $M_c$ if it shares both channel IDs of its parent, probability $M_1$ if it shares just the level-two channel ID of its parent, and probability $M_2$ if it shares neither channel ID of the parent. The $M_x$ value applied is from the offspring's genotype after mutation. Attempted suicide succeeds 80% of the time. This capacity enables first- or second-level individuals to combat somatic mutation. Initialization and mutation each of these parameters is performed by a redraw from the distribution $U(-0.5, 1.5)$ clamped to the range $[0, 1]$.

Finally, parameters **$S_1$** and **$S_2$ fine-tune site choice for offspring placement**. If an organism is placing an offspring with identical channel IDs, with probability $S_1$ the four possible sites for offspring placement are considered in order of increasing distance from the centroid of the parent's level-one signaling network. If an organism is placing an offspring with identical level-two channel ID but different level-one channel ID, with probability $S_2$ the four possible sites for offspring placement are considered in order of increasing distance from the centroid of the parent's level-two same-channel signaling network. Otherwise, the four possible sites for offspring placement are considered in a random order. Initialization and mutation are performed by a draw from the distribution $U(-0.5, 1.5)$ clamped to the range $[0, 1]$.

## Experiments

We performed experiments to assess the evolutionary trajectories of populations in the DISHTINY platform environment. We seeded each tile on the $120 \times 120$ toroidal grid with a randomized organism and ran the simulation for 20 million updates. We performed 33 replications of this experiment, each taking approximately 60 hours. Across all successive 10,000 update segments of all replicates, the mean number of cellular generations elapsed per 10,000 updates was 11.3 with a standard deviation of 1.9 cellular generations per 10,000 updates.

We observed evolutionary outcomes that resembled cell-, first-, and second-level individuality. To assess the relative fitness of these evolved organisms, we ran competitions between three genotypes — one selected as the most common genotype from the evolutionary run where the greatest mean $P_c$ was observed (i.e., cell-level individuality was observed), one selected as the most common genotype from the evolutionary run where the greatest mean $P_1$ was observed (i.e., first-level individuality was observed), and the other selected as the most common genotype from the evolutionary run where the greatest mean $P_2$ was observed (i.e., second-level individuality observed). We seeded each competition with three copies of each genotype, uniformly spaced over the $120 \times 120$ toroidal grid with random arrangement. We performed 191 runs in this experiment, each running for 2 million updates with mutation disabled, taking approximately 6 hours.

## Implementation

We implemented our experimental system using the Empirical library for scientific software development in C++, available at `https://github.com/devosoft/`

| | Competitors | | | Mean Dominant ($\pm S.D.$) | | |
|---|---|---|---|---|---|---|
| | $P_c{>}P_{0,1}$ | $P_1{>}P_{c,1}$ | $P_2{>}P_{c,0}$ | $P_c{>}P_{0,1}$ | $P_1{>}P_{c,1}$ | $P_2{>}P_{c,0}$ |
| $n$ | 1 | 1 | 1 | 2 | 16 | 15 |
| $A_1$ | 0.00 | 1.00 | 1.00 | 0.09±0.13 | 0.42±0.47 | 0.27±0.41 |
| $A_2$ | 1.00 | 0.91 | 1.00 | 1.00±0.00 | 0.99±0.02 | 1.00±0.00 |
| $P_c$ | 0.85 | 0.00 | 0.00 | 0.77±0.12 | 0.05±0.04 | 0.00±0.00 |
| $P_1$ | 0.07 | 1.00 | 0.00 | 0.13±0.09 | 0.86±0.15 | 0.00±0.00 |
| $P_2$ | 0.08 | 0.00 | 1.00 | 0.10±0.03 | 0.09±0.15 | 1.00±0.00 |
| $C_1$ | 21.8 | 7.2 | 9.9 | 19.9±2.6 | 10.4±2.5 | 9.9±1.6 |
| $C_2$ | 101.2 | 274.2 | 238.2 | 93.7±10.6 | 221.2±55.9 | 244.0±23.0 |
| $E_c$ | 0.21 | 0.00 | 0.00 | 0.27±0.09 | 0.02±0.05 | 0.00±0.00 |
| $E_1$ | 1.21 | 30.1 | 0.00 | 1.3±0.1 | 3.4±7.4 | 0.046±0.13 |
| $E_2$ | 2.49 | 54.1 | 38.8 | 2.4±0.1 | 29.4±16.9 | 55.4±16.8 |
| $M_c$ | 0.53 | 0.30 | 0.90 | 0.29±0.34 | 0.35±0.40 | 0.95±0.08 |
| $M_1$ | 1.00 | 0.00 | 1.00 | 0.86±0.20 | 0.49±0.40 | 0.67±0.46 |
| $M_2$ | 0.00 | 1.00 | 0.24 | 0.50±0.71 | 0.51±0.47 | 0.48±0.43 |
| $S_1$ | 0.56 | 1.00 | 0.68 | 0.29±0.38 | 0.44±0.46 | 0.68±0.37 |
| $S_2$ | 1.00 | 0.84 | 0.71 | 0.50±0.71 | 0.65±0.40 | 0.45±0.40 |

Figure 2: Enumerations for genotypes used as seeds for competition experiments (left) and enumerations for mean values of the most abundant genotype at the end of evolutionary runs (right), both sorted by resource-caching strategy.

`Empirical`. The code used to perform and analyze our experiments, our figures, data from our experiments, and a live in-browser demo of our system is available via the Open Science Framework at `https://osf.io/ewvg8/`.

## Results and Discussion

Cell-, first-, and second-level individuals were all observed at the conclusion of different runs of our evolutionary simulation (mean generation 22,016; $s = 3,119$). The criteria used to discern these outcomes are described below. Figure 3 shows the level-one and level-two signaling networks at the end of runs where cell-, first-, and second-level individuality evolved, respectively. Figure 4 shows a time series of signaling network snapshots in an evolutionary run where first-level individuality evolved. Cell-level individuals appear to form with comparatively large level-one signaling networks that are arranged into amorphous level-two signaling networks. Zeroth-level individuals appear to form elongated cigar-shaped level-two amalgamations of diverse level-one networks. First-level individuals appear to form highly regular diamond-shaped level-two amalgamations of diverse level-one networks.

Figure 2 describes predominant genotypes observed at the end of our evolutionary simulations. With a single exception, nearly all evolved genotypes had $A_2$ fixed at or very near 1.0 (i.e., population mean $A_2 \geq 0.993$). So, reproduction over cells sharing the same level-two channel was near-universally avoided; genotypes evolved so that cells declined to reproduce when they were located at the interior of level-two same-channel signaling networks.

However, a variety of resource-caching strategies evolved. Most-abundant genotypes at the end of evolutionary runs included strategies where resource was primarily cached in an organism's individual stockpile (i.e., $P_c > P_1, P_2$), strategies where resource was primarily cached in an organism's level-one signaling network's pool (i.e., $P_1 > P_c, P_2$), and strategies where resource was primarily cached in an organism's level-two signaling network's pool (i.e., $P_2 > P_c, P_1$). Among 33 trials, selfish cell-level hoarders dominated at the end of two replicates, level-one resource-sharing dominated in 16 replicates, and level-two resource sharing dominated in 15 replicates.

Given the near-ubiquitous nature of cooperation with regard to reproductive division of labor at the level-two same-channel signaling network, it was on this basis of resource caching strategy that we drew distinctions between cell-, first-, and second-level individuality. (The single predominant genotype with $A_2 = 0.91$ had $P_1 = 1.0$, so was not sharing resource on the level-two same-channel resource pool).

Next, we wanted to compare cell-, first-, and second-level individuals to determine which genotype was the most fit in the DISHTINY platform environment. We ran competition experiments between the the dominant genotypes from the run with greatest mean $P_c$, the run with greatest mean $P_1$, and the run with greatest mean $P_2$. In 22 out of 191 trials performed fixation was reached by update 1.5 million. The cell-level individuality genotype dominated in one trial, the first-level individuality genotype dominated in 12 trials, and the first-level individuality genotype dominated in 178 trials. These results show that in the absence of mutation, first-level individuals tend to exhibit greater fitness than first- and cell-level individuals ($p < 0.0001$; RR 2.8; two-tailed exact test).

In competition experiments, however, higher-level individuals likely benefited from elimination of somatic mutation. To assess the relative fitness of first- and second-level individuals without mutation disabled, we examined the relationship between first- and second-level resource pooling and the rate of cellular reproduction at the end of each of the 33 replicate evolutionary trials performed. We observed a significant negative correlation between mean $P_1$ and cellular reproduction rate ($p < 0.0001$; bootstrap test; Figure 5a) and a significant positive correlation between mean $P_2$ and cellular reproduction rate ($p < 0.0001$; bootstrap test; Figure 5b). This result suggests that second-level individuals tend to collect resource more effectively than first-level individuals. We did not test correlation between $P_c$ and reproduction rate due to the small number of trials where cell-level individuality dominated.

With the viability of cell-, first-, and second-level individuality in the DISHTINY platform environment — and

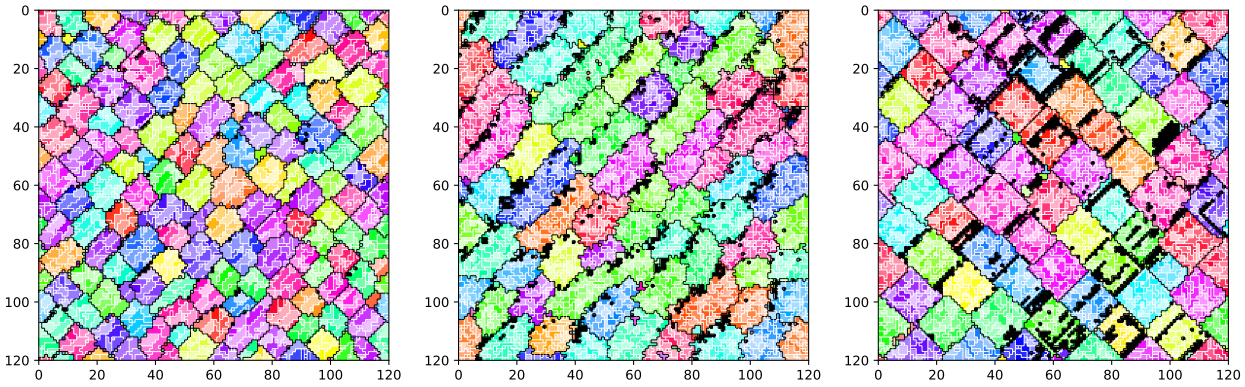

(a) Mean $P_c$=0.77, $P_1$=0.09, $P_2$=0.14; gen. 20,475

(b) Mean $P_1$=1.0; gen. 23,971

(c) Mean $P_2$=1.0; gen. 25,841

Figure 3: End state of same-channel signaling networks in replicates where cell- (3a), first- (3b), and second-level (3c) individuality dominated. (Cell-level individuals are single cells that retain collected resource exclusively for their own use, first-level individuals are level-one same-channel multi-cellular networks that primarily assign collected resource for collective use among level-one channel mates, and second-level individuals are level-two same-channel multi-cellular networks that primarily assign collected resource for collective use among level-two channel mates.) Level-one channels are coded by color saturation and level-two channels are coded by color hue. A single cell-like organism occupies each grid tile except for black tiles, which are empty.

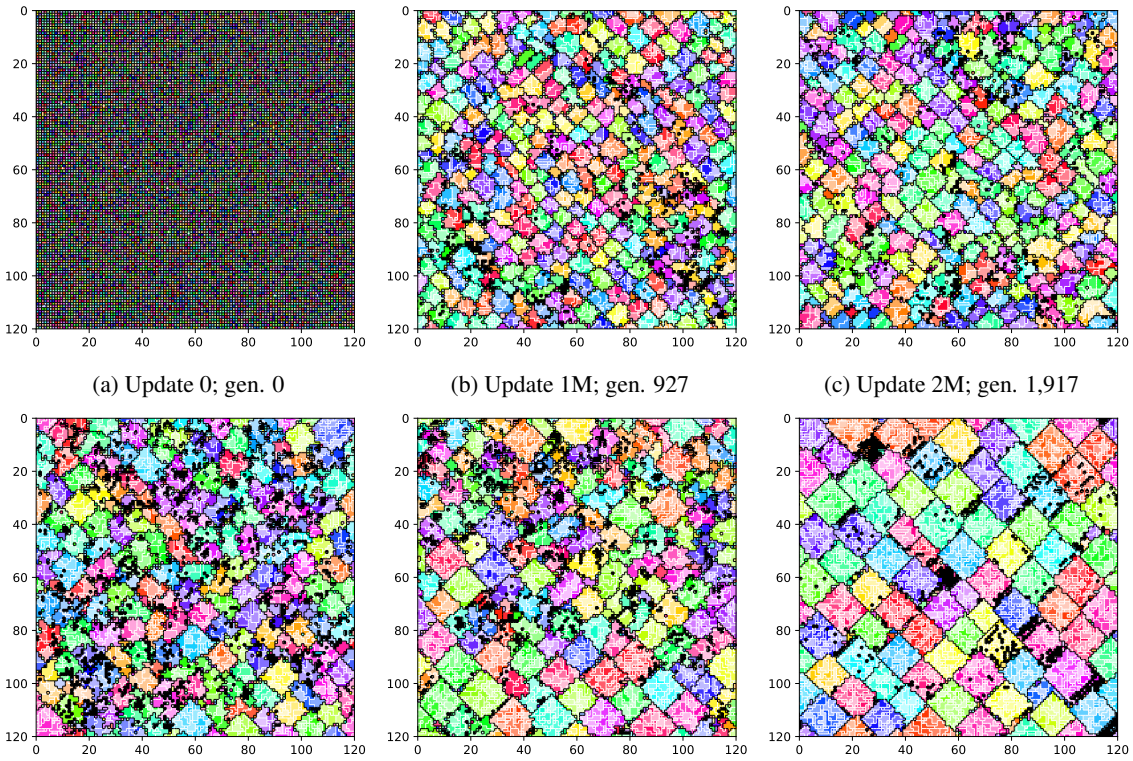

(a) Update 0; gen. 0

(b) Update 1M; gen. 927

(c) Update 2M; gen. 1,917

(d) Update 4M; gen. 4,053

(e) Update 5M; gen. 5,173

(f) Update 7M; gen. 7,312

Figure 4: Progression of of same-channel level-one and level-two signaling networks states in an evolutionary run where second-level individuality evolved. Level-one channels are coded by color saturation and level-two channels are coded by color hue. A single cell-like organism occupies each grid tile except for black tiles, which are empty.

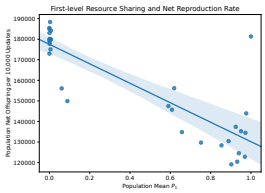


(a) Correlation plot of population mean $P_1$ and population net reproduction rate.

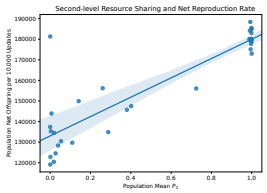


(b) Correlation plot of population mean $P_2$ and population net reproduction rate.

Figure 5: Mean resource caching strategies and net reproduction rate across populations. A bootstrapped 95% confidence interval for the fit is shaded.

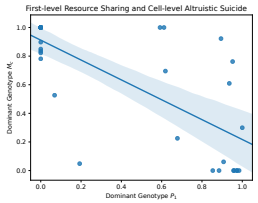


(a) Correlation plot of dominant genotype $P_1$ and dominant genotype $M_c$.

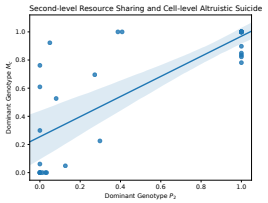


(b) Correlation plot of dominant genotype $P_2$ and dominant genotype $M_c$.

Figure 6: Plots of dominant resource caching strategies and dominant apoptosis strategies. A bootstrapped 95% confidence interval for the fit is shaded.

the greater relative fitness of second-level individuality — established, we were also interested in probing the strategies employed by cell-, first-, and second-level individuals beyond resource caching and reproductive deferment. To assess whether higher-level individuals employed apoptosis to mitigate somatic mutation, we examined the relationship between first- and second-level resource pooling and cellular apoptosis at the conclusion of our 33 replicate evolutionary trials. We observed a significant negative correlation between dominant genotype $P_1$ and $M_c$ ($p < 0.0001$; bootstrap test; Figure 6a) and a significant positive correlation between dominant genotype $P_2$ and $M_c$ ($p < 0.0001$; bootstrap test; Figure 6b). Notably, no genotype encoding second-level individuality was observed with $M_c < 0.5$. This result suggests that second-level individuals, in particular, relied on apoptosis to mitigate somatic mutation, perhaps due to their much larger scale compared to cell- and first-level individuals.

To assess whether higher-level individuals provided larger resource endowments to their propagules (offspring sharing neither the level-one nor the level-two channel ID with the parent), we examined the relationship between first and second-level resource pooling and dominant genotype second-level propagule endowment at the conclusion of our 33 replicate evolutionary trials. We observed a significant negative correlation between dominant genotype $P_1$ and $E_2$ ($p < 0.001$; bootstrap test) and a significant positive correlation between dominant genotype $P_2$ and $E_2$ ($p < 0.0001$; bootstrap test). Second-level individuals might provide larger endowments to propagules simply due to a greater capacity to collect resource or perhaps because of stronger selection for well-endowed offspring when competing against other second-level individuals.

## Conclusion

Using simple organisms that evolve parameters for a set of manually-designed strategies, we have demonstrated that DISHTINY selects for genotypes that exhibit high-level individuality. We observed cell-, first-, and second- level individuality among evolutionary outcomes. Specifically, we observed

1. reproductive division of labor among members of the same channel (i.e., individuals enveloped in a same-channel signaling network ceded reproduction to those at the periphery), and
2. cooperation between members of the same channel (i.e., pooling of resource on same-channel signaling networks).

Competition experiments revealed that second-level individuals usually outcompete first- and cell-level individuals. We observed suppression of somatic mutation through apoptosis correlated with second-level individuality. The magnitude of resource endowment for propagules was also correlated with second-level individuality.

Although shifts in individuality coincident with level-one and level-two signaling networks were both clearly observed, the question of whether these transitions were truly hierarchical in nature is debatable. That is, it is not clear whether level-one individuality was to some extent preserved in or necessary for the emergence of level-two individuality. Given the nature of the manually-designed strategies for resource-pooling and reproductive division of labor, level-two resource pooling and division of labor could readily leapfrog over level-one resource pooling and division of labor and, in many ways, seemed to completely supersede those level-one efforts.

We believe that this is a shortcoming of the design of the simple cell-like organism employed in these experiments, not the DISHTINY platform itself. We have nevertheless demonstrated that DISHTINY ultimately selects for high-level individuality. We are eager to work with more sophisticated cell-like organisms capable of arbitrary computation via genetic programming in order to pursue more open-ended evolutionary experiments. We will also test the implications of relaxing current arbitrary restrictions that artificially promote transitions, such as the hierarchical nesting of same-channel signaling networks and the explicitly-defined signaling networks themselves, leaving these details

to evolution to figure out. Further work will provide valuable insight into scientific questions relating to major evolutionary transitions such as the role of pre-existing phenotypic plasticity (Clune et al., 2007; Lalejini and Ofria, 2016), pre-existing environmental interactions, pre-existing reproductive division of labor, and how transitions relate to increases in organizational (Goldsby et al., 2012), structural, and functional (Goldsby et al., 2014) complexity.

We believe that such an approach also provides a unique opportunity to fundamentally advance Artificial life with respect to open-ended evolution. Fundamental to this goal is scale. The DISHTINY platform trivially scales to select for an arbitrary number of hierarchical levels of individuality (not just the two hierarchical levels explored in these experiments). Importantly, the platform is implemented in a decentralized manner and can comfortably scale as additional computing resources are provided. Parallel computing is widely exploited in evolutionary computing, where subpopulations are farmed out for periods of isolated evolution or single genotypes are farmed out for fitness evaluation (Lin et al., 1994; Real et al., 2017). DISHTINY presents a more fundamental parallelization potential: principled parallelization of the evolving individual phenotype at arbitrary scale (i.e., a high-level individual as a large collection of individual cells on the toroidal grid). Such parallelization will be key to realizing evolving computational systems with scale — and, perhaps, complexity — approaching those of biological systems.

## Acknowledgements

Thanks to members of the DEVOLAB, in particular Michael J. Wiser, for feedback on statistical methods and Heather Goldsby for concept and editing feedback. This research was supported in part by NSF grants DEB-1655715 and DBI-0939454, and by Michigan State University through the computational resources provided by the Institute for Cyber-Enabled Research. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

## References

Ackley, D. H. (2016). Indefinite scalability for living computation. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*.

Ackley, D. H. and Cannon, D. C. (2011). Pursue robust indefinite scalability. In *HotOS*.

Banzhaf, W., Baumgaertner, B., Beslon, G., Doursat, R., Foster, J. A., McMullin, B., De Melo, V. V., Miconi, T., Spector, L., Stepney, S., et al. (2016). Defining and simulating open-ended novelty: requirements, guidelines, and challenges. *Theory in Biosciences*, 135(3):131–161.

Bedau, M. A. (2003). Artificial life: organization, adaptation and complexity from the bottom up. *Trends in cognitive sciences*, 7(11):505–512.

Bouchard, F. (2013). What is a symbiotic superindividual and how do you measure its fitness. *From groups to individuals: evolution and emerging individuality*, 243.

Clune, J., Ofria, C., and Pennock, R. T. (2007). Investigating the emergence of phenotypic plasticity in evolving digital organisms. In *European Conference on Artificial Life*, pages 74–83. Springer.

Ereshefsky, M. and Pedroso, M. (2015). Rethinking evolutionary individuality. *Proceedings of the National Academy of Sciences*, 112(33):10126–10132.

Goldsby, H. J., Dornhaus, A., Kerr, B., and Ofria, C. (2012). Task-switching costs promote the evolution of division of labor and shifts in individuality. *Proceedings of the National Academy of Sciences*, 109(34):13686–13691.

Goldsby, H. J., Knoester, D. B., Ofria, C., and Kerr, B. (2014). The evolutionary origin of somatic cells under the dirty work hypothesis. *PLoS biology*, 12(5):e1001858.

Goldsby, H. J., Young, R. L., Hofmann, H. A., and Hintze, A. (2017). Increasing the complexity of solutions produced by an evolutionary developmental system. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pages 57–58. ACM.

Lalejini, A. and Ofria, C. (2016). The evolutionary origins of phenotypic plasticity. In *Proceedings of the Artificial Life Conference*.

Lin, S.-C., Punch, W. F., and Goodman, E. D. (1994). Coarse-grain parallel genetic algorithms: Categorization and new approach. In *Parallel and Distributed Processing, 1994. Proceedings. Sixth IEEE Symposium on*, pages 28–37. IEEE.

Ray, T. S. (1996). Evolving parallel computation. *Complex Systems*, 10:229–237.

Real, E., Moore, S., Selle, A., Saxena, S., Suematsu, Y. L., Tan, J., Le, Q. V., and Kurakin, A. (2017). Large-scale evolution of image classifiers. In Precup, D. and Teh, Y. W., editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2902–2911, International Convention Centre, Sydney, Australia. PMLR.

Smith, J. and Szathmary, E. (1997). *The Major Transitions in Evolution*. OUP Oxford.

Taylor, T., Bedau, M., Channon, A., Ackley, D., Banzhaf, W., Beslon, G., Dolson, E., Froese, T., Hickinbotham, S., Ikegami, T., et al. (2016). Open-ended evolution: perspectives from the oee workshop in york. *Artificial life*, 22(3):408–423.

West, S. A., Fisher, R. M., Gardner, A., and Kiers, E. T. (2015). Major evolutionary transitions in individuality. *Proceedings of the National Academy of Sciences*, 112(33):10112–10119.