

Toward Open-Ended Fraternal Transitions in Individuality

Matthew Andres Moreno and Charles Ofria

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

Abstract

1 The emergence of new replicating entities from the union of
2 simpler entities represent some of the most profound events
3 in natural evolutionary history. Such transitions in individuality
4 are essential to the evolution of the most complex forms
5 of life. As such, understanding these transitions is critical
6 to building artificial systems capable of open-ended evolution.
7 Alas, these transitions are challenging to induce or detect,
8 even with computational organisms. Here, we introduce the
9 DISHTINY (DIStributed Hierarchical Transitions in Individuality)
10 platform, which provides simple cell-like organisms with the ability
11 and incentive to unite into new individuals in a manner that can
12 continue to scale to subsequent transitions. The system is designed
13 to encourage these transitions so that they can be studied: organisms
14 that coordinate spatiotemporally can maximize the rate of resource
15 harvest, which is closely linked to their reproductive ability.
16 We demonstrate the hierarchical emergence of multiple levels
17 of individuality among simple cell-like organisms that evolve
18 parameters for manually-designed strategies. During evolution,
19 we observe reproductive division of labor and close cooperation
20 among cells, including resource-sharing, aggregation of resource
21 endowments for propagules, and emergence of an apoptosis response
22 to somatic mutation. Many replicate populations evolved to direct
23 their resources toward low-level groups (behaving like multi-cellular
24 individuals) and many others evolved to direct their resources toward
25 high-level groups (acting as larger-scale multi-cellular individuals).

Introduction

29 Artificial life researchers design systems that exhibit properties
30 of biological life in order to better understand their dynamics
31 and, often, to apply these principles toward engineering applications
32 such as artificial intelligence (Bedau, 2003). Studies of evolution
33 have been of particular interest to the community, especially in
34 regard to how organisms are produced with increasing sophistication
35 and complexity (Goldsby et al., 2017). This particular issue is
36 often described as “open-ended evolution.” Although precise
37 definitions and measures of open-ended evolution are still being
38 established, this term is generally understood to refer to evolving
39 systems that exhibit the continued production of novelty (Taylor
40 et al., 2016). Evolutionary transitions in

43 individuality, which are key to the complexification and
44 diversification of biological life (Smith and Szathmary, 1997),
45 have been highlighted as key research targets with respect to
46 the question of open-ended evolution (Ray, 1996; Banzhaf
47 et al., 2016). In an evolutionary transition of individuality,
48 a new, more complex replicating entity is derived from the
49 combination of cooperating replicating entities that have
50 irrevocably entwined their long-term fates (West et al., 2015).
51 In particular, we focus on fraternal transition in individuality,
52 events where closely-related kin come together or stay together
53 to form a higher-level organism (Queller, 1997). Eusocial insect
54 colonies and multicellular organisms exemplify this phenomenon
55 (Smith and Szathmary, 1997). Like the definition of open-ended
56 evolution, the notion of what constitutes an evolving individual
57 is not concretely established. Commonly indicated features include:
58 close coordination and cooperation, reproductive division of labor,
59 reproductive bottlenecks, and loss of ability to replicate
60 independently (Ereshefsky and Pedroso, 2015; Bouchard, 2013).

62 Our appreciation of fraternal transitions in individuality
63 benefits from experimental work probing the origins of multicellularity.
64 In the biological domain, Ratcliff et al. have demonstrated
65 evolution of multicellularity in yeast, deriving fraternal clusters
66 of cells that cling together in order to maximize their settling
67 rate (Ratcliff et al., 2012). The contributions of Goldsby and
68 collaborators are particularly notable among computational Artificial
69 life work on the origins of multicellularity. Their evolutionary
70 experiments track a population composed of demes, distinct
71 spatial domains inhabited by clonal colonies of cells. Two
72 distinct types of reproduction occur: (1) cells reproduce within
73 demes and (2) deme reproduction, where a target deme is
74 sterilized then re-inoculated with genetic material from the
75 parent deme. With such methods, Goldsby et al. have studied
76 division of labor (Goldsby et al., 2010, 2012), the origin of
77 soma (Goldsby et al., 2014), and the evolution of morphological
78 development (Goldsby et al., 2017). We aspire to complement
79 deme-based approaches with a framework where higher level
80 individuality unfolds via cellular reproductions within a single
81 unified space. In particular, we are interested

83 in the potential for such a system to undergo nested hierar- 135
 84 chical transitions. 136

85 Major challenges in studying evolutionary transitions in 137
 86 individuality include (1) determining the environmental con- 138
 87 ditions that will promote such a transition and then (2) rec- 139
 88 ognizing that a transition has occurred. In order to begin 140
 89 exploring transitions in individuality, we must devise a sys- 141
 90 tem in which we expect such transitions to occur repeatably 142
 91 and in a detectable manner. Once we can consistently in- 143
 92 duce and observe evolutionary transitions in individuality, 144
 93 we may subsequently proceed to relax aspects of such a sys- 145
 94 tem to explore in greater detail what conditions are neces- 146
 95 sary to induce transitions and how transitions can be det- 147
 96 ected. For now, we will focus on these initial goals in the 148
 97 context of fraternal transitions in individuality. 149

98 To this end, we introduce the DISHTINY (DIStributed 150
 99 Hierarchical Transitions in Individuality) platform, which 151
 100 seeks to achieve the evolution of transitions in individual- 152
 101 ity by explicitly registering organisms in cooperating groups 153
 102 that coordinate spatiotemporally to maximize the harvest of 154
 103 a resource. Detection of such a transition in DISHTINY 155
 104 is accomplished by identifying resource-sharing and repro- 156
 105 ductive division of labor among organisms registered to the 157
 106 same cooperating group. We designed this system such that 158
 107 hierarchal transitions across an arbitrary number of levels of 159
 108 individuality can be selected for and meaningfully detected. 160
 109 We have focused this system on a rigid form of major tran- 161
 110 sition using simple organisms, but the underlying principles 162
 111 can be applied to a wide range of artificial life systems. Fur- 163
 112 thermore, DISHTINY is decentralized and amenable to mas- 164
 113 sive parallelization via distributed computing. We believe 165
 114 that such scalability — with respect to both concept and im- 166
 115 plementation — is an essential consideration in the pursuit 167
 116 of artificial systems capable of generating complexity and 168
 117 novelty rivaling that of biological life via open-ended evolu- 169
 118 tion (Ackley and Cannon, 2011; Ackley, 2016). 170

119 **Methods** 171

120 In order to demonstrate that the DISHTINY platform selects 173
 121 for detectable hierarchical transitions in individuality, we 174
 122 performed experiments where cell-like organisms evolved 175
 123 parameters to control manually designed behaviors such as 176
 124 resource-sharing, reproductive decision-making, and apop- 177
 125 tosis. We will first cover the design of the DISHTINY plat- 178
 126 form and then describe the simple cell-like organisms we 179
 127 used to evaluate the platform. 180

128 **DISHTINY** 182

129 DISHTINY allows cell-like organisms to replicate across a 183
 130 toroidal grid. Over discrete timesteps (“updates”), the cells 184
 131 can collect a continuous-valued resource. Once sufficient 185
 132 resource has been accrued, cells may pay 8.0 resource to 186
 133 place a daughter cell on an adjoining tile of the toroidal grid 187
 134 (i.e., reproduce), replacing any existing cell already there. 188

As cells reproduce, they can choose to include offspring in 135
 the parent’s cooperating “signaling channel” group or force 136
 offspring to create a new cooperating “signaling channel” 137
 group. 138

As shown at the top of Figure 1, resources appear at 139
 a single point then spread outwards update-by-update in a 140
 diamond-shaped wave, disappearing when the expanding 141
 wave reaches a predefined limit. Cells must be in a costly 142
 “activated” state to collect resource as it passes. The cell at 143
 the starting position of a resource wave is automatically ac- 144
 tivated, and will send the activate signal to neighboring cells 145
 on the same signaling channel. The newly activated cells, 146
 in turn, activate their own neighbors registered to the same 147
 signaling channel. Neighbors registered to other signaling 148
 channels do not activate. Each cell, after sending the acti- 149
 vation signal, enters a temporary quiescent state so as not 150
 to reactivate from the signal. In this manner, cells sharing 151
 a signaling channel activate in concert with the expanding 152
 resource wave. As shown Figure 1a, b, the rate of resource 153
 collection for a cell is determined by the size and shape of 154
 of its same-channel signaling network; small or fragmented 155
 same-channel signaling networks will frequently miss out 156
 on resource as it passes by. 157

Each cell pays a resource cost when it activates. This cost 158
 is outweighed by the resource collected such that cells that 159
 activate in concert with a resource wave derive a net benefit. 160
 Recall, though, that resource waves have a limited extent. 161
 Cells that activate outside the extent of a resource wave or 162
 activate out of sync with the resource wave (due to an in- 163
 direct path from the cell that originated the signal) pay the 164
 activation cost but collect no resource. Cells that frequently 165
 activate erroneously use up their resource and die. In our im- 166
 plementation, organisms that accrue a resource debt of -11 167
 or greater are killed. This erroneous activation scenario is 168
 depicted in Figure 1c. 169

In this manner, “Goldilocks” — not too small and not 170
 too big — signaling networks are selected for. Based on 171
 a randomly chosen starting location, resource wave start- 172
 ing points (seeds) are tiled over the toroidal grid such that 173
 the extents of the resource waves touch, but do not overlap. 174
 All waves start and proceed synchronously; when they com- 175
 plete, the next resource waves are seeded. This process en- 176
 sures that selection for “Goldilocks” same-channel signaling 177
 networks is uniformly distributed over the toroidal grid. 178

Cells control the size and shape of their same-channel sig- 179
 naling group through strategic reproduction. Three choices 180
 are afforded: whether to reproduce at all, where among the 181
 four adjoining tiles of the toroidal grid to place their off- 182
 spring, and whether the offspring should be registered to 183
 the parent’s signaling channel or be given a random chan- 184
 nel ID (in the range 1 to 2^{22}). No guarantees are made about 185
 the uniqueness of a newly-generated channel ID, but chance 186
 collisions are rare. 187

Hierarchical levels are introduced into the system through 188




Figure 1: **Activation signaling, and net resource collection for three different-sized same-channel networks 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 two-cell channel-signaling group (far left, in green) is activated; tracking the resource wave (top) yields a small net resource harvest (far right). In row *b*, an intermediate-sized 13-cell channel-signaling group yields a high net resource harvest. Finally, in row *c*, a large 29-cell 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.

189 multiple separate, but overlaid, instantiations of this re- 211
 190 source wave/channel-signaling scheme. We refer to each 212
 191 independent resource wave/channel-signaling system as a 213
 192 “level.” In our experiments, we allowed two resource 214
 193 wave/channel-signaling levels, identified here as level one 215
 194 and level two. On level one, resource waves extended a ra- 216
 195 dius of three toroidal tiles. On level two they extended a 217
 196 radius of 24 toroidal tiles. On both levels, activated cells 218
 197 netted +1.0 resource from a resource wave, but suffered an 219
 198 activation penalty of -5.0 if no resource was available. Due 220
 199 to the different radii of resource waves on different levels, 221
 200 level one selects for small same-channel signaling networks 222
 201 and level two selects for large same-channel signaling net- 223
 202 works. 224

203 Cells were marked with two separate channel IDs, one 225
 204 for level one and another for level two. We enforced hier- 226
 205 archical nesting of same-channel signaling networks during 227
 206 reproduction: daughter cells may inherit neither channel ID, 228
 207 just the level-two channel ID, or both channel IDs. Daugh- 229
 208 ter cells may not inherit only the level-one channel ID while 230
 209 having a different level-two channel ID. The distribution of 230
 210 IDs across the level-two and level-one channels can be envi- 231

211 sioned by analogy to political countries and territories. Each 212
 213 country (i.e., level-two channel network) may have one or 214
 215 many territories (i.e., level-one channel network). However, 216
 217 no territory spans more than one country. Figure 2 depicts 218
 219 hierarchically nested channel states at the end of three evo- 220
 221 lutionary runs. 222

223 Channel IDs enable straightforward detection of an evolu- 224
 225 tionary transition in individuality. Because common channel 226
 227 IDs may only arise systematically through inheritance, com- 228
 229 mon channel IDs indicate a close hereditary relationship in 230
 231 addition to a close cooperative relationship. Because new 232
 233 channel IDs arise first in a single cell, same-channel sig- 234
 235 naling networks are reproductively bottlenecked, ensuring 236
 237 meaningful reproductive lineages at the level of the same- 238
 239 channel signaling network. To recognize an evolutionary 240
 241 transition in individuality, we therefore evaluate 242

- 243 1. Do cells with the same channel ID choose to share re- 244
 245 sources (e.g., cooperate)? 246
- 247 2. Is there division of reproductive labor between members 248
 249 of the same channel (e.g., do cells at the interior of a net- 250
 251 work cede reproduction to those at the periphery)? 252

232 If these conditions are met among cells sharing the same
 233 level-one channel, a first-level transition in individuality
 234 may have occurred. Likewise, if these conditions are met
 235 among cells sharing the same level-two channel, a second-
 236 level transition in individuality may have occurred. In either
 237 case, observation of altruistic behavior, such as an apoptosis
 238 response to mutation, would further evidence a transition.

239 Organisms

240 We performed our experiments using cell-like organisms
 241 composed of 15 floating-point parameters, each controlling
 242 a specific strategy component pertinent to transitions in indi-
 243 viduality (i.e., reproductive division of labor, resource pool-
 244 ing, apoptosis, propagule generation, and propagule endow-
 245 ment). These particular cell-like organisms are in no way
 246 inherent to the DISHTINY platform, but were merely de-
 247 veloped to study transitions using as simple a model system
 248 as feasible. On reproduction, we applied mutation to each
 249 parameter independently with probability 0.00005.

250 The **aversion parameters** (A_1 and A_2) allow cells to
 251 avoid reproducing over neighbors sharing the same signal-
 252 ing channel. Specifically, they control the probability that a
 253 cell declines to supplant a neighbor sharing the same level-
 254 one (A_1) or level-two (A_2) channel ID. If a cell declines
 255 to place its offspring in all four adjoining tiles, it does not
 256 reproduce. Mutation is performed by a redraw from the uni-
 257 form distribution $U(-0.5, 1.5)$ clamped to the range $[0, 1]$.

258 The **resource allocation parameters** control the propor-
 259 tion of resources that go to the cell's stockpile (P_c), its level-
 260 one channel's resource pool (P_1), or its level-two channel's
 261 resource pool (P_2). These parameters are initialized by a
 262 draw from $U(-1.0, 2.0)$ clamped to the range $[0, 1]$ and mu-
 263 tated by addition of a normal value drawn from $N(0.0, 0.2)$
 264 with the result clamped to the range $[0, 1]$. The set P_c, P_1, P_2
 265 is always normalized to sum to 1.

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

276 **Channel cap parameters** C_1 and C_2 regulate the size
 277 of same-channel signaling networks. When an organism re-
 278 produces, it checks the size of its level-one signaling net-
 279 work against C_1 and the size of its level-two signaling group
 280 against C_2 . If neither cap is met or exceeded, then the or-
 281 ganism will produce an offspring sharing both of its chan-
 282 nel IDs. If only the C_1 cap is exceeded, then the organism
 283 will produce an offspring with new level-one channel ID but
 284 identical level-two channel ID. Finally, if the C_2 cap is ex-

ceeded, then the organism will produce an offspring with
 new IDs for both channels. For level-one caps, these pa-
 rameters are initialized by a draw from $U(0.0, 16.0)$. For
 level-two caps, these parameters are initialized by a draw
 from $U(0.0, 128.0)$. Both are 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 deter-
 mine the amount of resource provided to offspring. This
 endowment is paid as an additional cost by the cell stock-
 pile (or same-channel resource pool) funding a reproduc-
 tion. The full amount of the received endowment is di-
 vided 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 param-
 eters. E_c is the endowment amount paid to an offspring
 that shares both channel IDs of the parent; E_1 is the en-
 dowment 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 the 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.
 In order that adequate resource remain to ensure parental
 stability, endowment was paid out only after twice the en-
 dowment amount had been accrued (leaving an amount of
 resource equal to the endowment remaining with the par-
 ent). Cell level endowments are initialized by a draw from
 $U(0.0, 5.0)$. Level-one endowments are initialized by a
 draw from $U(0.0, 80.0)$. Level-two endowments are initial-
 ized by a draw from $U(0.0, 405.0)$. All endowments are mu-
 tated 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 re-
 sponse 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 chan-
 nel ID of its parent, and probability M_2 if it shares nei-
 ther channel ID of the parent. The M_x value applied is
 from the offspring's genotype after mutation. Attempted
 suicide succeeds 90% of the time. This capacity enables
 first- or second-level individuals to combat somatic muta-
 tion. 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 off-
 spring 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 par-
 ent's level-one signaling network. If an organism is placing
 an offspring with identical level-two channel ID but differ-

ent 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 $[0, 1]$.

347 Treatments

Our standard treatment was designed to assess the evolutionary trajectories of populations in DISHTINY. We seeded each tile on the 120×120 toroidal grid with a randomized organism and ran the simulation for 20 million updates. In order to facilitate turnover, we culled the population intermittently. Starting at update 500,000, and every 50,000 updates thereafter, we randomly selected second-level channel IDs and killed all cells with that channel ID, continuing until at least 5% of grid tiles were empty. We performed 50 replicates within this treatment. On average, each cellular generation took just over 500 updates. Across all successive 10,000 update segments of all replicates, the mean number of cellular generations elapsed per 10,000 updates was 19.2 with a standard deviation of 2.7 cellular generations per 10,000 updates.

In order to detangle the impact of same-channel signaling networks with respect to kin recognition versus cooperation to increase resource collection rate, we performed control evolutionary trials where same-channel signaling networks did not improve cellular resource collection rate. Under control conditions, same-channel signaling networks just helped cells recognize other related cells. In our implementation, this treatment corresponded to resource waves with radius 1 (i.e., the resource wave did not expand beyond its seed) that paid out 1.0 resource units to cells and no cost for erroneous activation. All other aspects of control runs, including the functionality of all lifestyle parameters, were otherwise identical to standard conditions. We performed 50 replications of the control treatment. In control runs, generations progressed much faster, taking only around 50 updates. Across all successive 10,000 update segments of all replicates, the mean number of cellular generations elapsed per 10,000 updates was 211.1 with a standard deviation of 43.8 cellular generations per 10,000 updates. Higher resource inflow rate under the control conditions likely contributed to the faster cellular generation rate compared to control conditions.

In standard evolutionary runs, we observed a spectrum of evolved resource-caching strategies. To assess the relative fitness of these evolved organisms, we ran competitions between the most common genotype from three standard evolutionary runs. The first genotype allocated resource exclusively to its first-level same-channel resource pool (i.e., $P_1 = 1.0$), the second split resource evenly between its first-

level and second-level resource pool (i.e., $P_1 = P_2 = 0.5$), and the third allocated resource primarily to the second-level resource pool (i.e., $P_2 > P_1$). (No most-common genotypes allocated resource exclusively to the second-level resource pool.) We seeded each competition with three copies of each genotype, uniformly spaced over the 120×120 toroidal grid with random arrangement. Each competition lasted 2 million updates. We performed 50 runs in this experiment.

400 Implementation

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

We performed our computational experiments at the Michigan State University High Performance Computing Center. Unfortunately, efforts to prepare the system for an operating system upgrade caused intermittent disk unavailability and checkpoint-restart failure. By cursory inspection of the line count of our most intensively written data files we estimate that 5-10% of content is missing. Lines appear to be missing in contiguous sections (corresponding to contiguous update ranges). Except where explicitly noted, missing data is not relevant to our analyses and visualizations, which primarily describe end-states where data from all runs is available. Due to checkpoint-restart failures, we curtailed our standard-condition evolutionary runs from a planned 25 million to 20 million updates and our control evolutionary runs to 250,000 updates.

Each replicate of standard evolutionary experiments required approximately six days of compute time to reach 20 million updates. Each replicate of control evolutionary experiments expended on the order of one day of compute time to reach 250,000 updates. Control runs were slower than standard runs, likely due to a higher per-update cellular generation rate. Each replicate of competition experiments consumed approximately ten hours of compute time. For standard evolutionary experiments, data processing required approximately four hours of compute time per run. Other data processing was computationally negligible.

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

Standard Evolutionary Experiments


A spectrum of resource allocation strategies ranging from purely allocation to level-one same-channel resource pools to primarily allocation to level-two same-channel resource pools were observed at the conclusion of different runs of our evolutionary simulation (mean generation 37,168 with standard deviation 4,684). We interpret these outcomes as

	Competitors			Mean Dominant ($\pm S.D.$)			Pop Mean ($\pm S.D.$)		Control Pop Mean ($\pm S.D.$)
	$P_1 = 1.0$	$P_2 = P_1$	$P_2 > P_1$	$P_1 = 1.0$	$1.0 > P_1 > P_2$	$P_2 \geq P_1$	<i>all</i>	<i>all</i>	<i>all</i>
<i>Cell Gen.</i>	29920	33852	47507	30841 \pm 3183	35346 \pm 3444	39315 \pm 3346	387 \pm 19	4306 \pm 400	4374 \pm 179
<i>Upd.</i>	20M	20M	20M	20M	20M	20M	200k	2.1M	200k
<i>n</i>	1	1	1	9	7	34	50	48	50
A_1	0.00	0.00	0.89	0.23 \pm 0.35	0.50 \pm 0.47	0.57 \pm 0.46	0.51 \pm 0.14	0.54 \pm 0.33	0.47 \pm 0.31
A_2	1.00	1.00	1.00	1.00 \pm 0.00	1.00 \pm 0.00	1.00 \pm 0.00	1.00 \pm 0.00	1.00 \pm 0.00	1.00 \pm 0.00
P_c	0.00	0.00	0.00	0.00 \pm 0.00	0.00 \pm 0.00	0.03 \pm 0.05	0.07 \pm 0.03	0.01 \pm 0.02	0.00 \pm 0.00
P_1	1.00	0.50	0.00	1.00 \pm 0.00	0.60 \pm 0.07	0.28 \pm 0.16	0.39 \pm 0.11	0.43 \pm 0.22	0.04 \pm 0.08
P_2	0.00	0.50	1.00	0.00 \pm 0.00	0.40 \pm 0.07	0.69 \pm 0.14	0.54 \pm 0.11	0.55 \pm 0.22	0.96 \pm 0.08
C_1	3.13	3.45	2.04	3.90 \pm 0.60	3.38 \pm 0.33	3.03 \pm 0.69	3.38 \pm 0.23	3.24 \pm 0.55	8.08 \pm 5.55
C_2	233.2	238.6	290.2	230.6 \pm 71.1	192.7 \pm 45.3	271.6 \pm 73.6	99.2 \pm 7.4	173.4 \pm 44.0	353.5 \pm 94.7
E_c	0.87	0.14	4.20	0.29 \pm 0.37	0.44 \pm 0.59	0.21 \pm 0.75	1.43 \pm 0.38	1.26 \pm 0.91	2.81 \pm 1.51
E_1	33.4	11.7	4.80	47.2 \pm 21.7	21.3 \pm 12.0	4.62 \pm 7.05	31.5 \pm 6.6	21.6 \pm 15.3	6.13 \pm 8.79
E_2	341.4	397.4	321.1	231.2 \pm 94.3	283.1 \pm 57.0	325.4 \pm 68.9	240.0 \pm 30.0	296.8 \pm 61.8	316.7 \pm 54.6
M_c	0.11	1.00	0.66	0.33 \pm 0.41	0.74 \pm 0.31	0.67 \pm 0.35	0.50 \pm 0.11	0.37 \pm 0.28	0.20 \pm 0.23
M_1	0.00	1.00	0.40	0.52 \pm 0.41	0.65 \pm 0.46	0.68 \pm 0.38	0.50 \pm 0.12	0.51 \pm 0.32	0.49 \pm 0.35
M_2	0.00	0.44	1.00	0.45 \pm 0.39	0.52 \pm 0.37	0.50 \pm 0.42	0.50 \pm 0.13	0.48 \pm 0.31	0.51 \pm 0.33
S_1	0.00	1.00	1.00	0.65 \pm 0.38	0.55 \pm 0.40	0.47 \pm 0.42	0.48 \pm 0.11	0.42 \pm 0.33	0.49 \pm 0.34
S_2	0.00	0.01	0.46	0.51 \pm 0.43	0.35 \pm 0.39	0.45 \pm 0.39	0.49 \pm 0.13	0.46 \pm 0.30	0.43 \pm 0.31


Table 1: The leftmost two table segments enumerate genotypes used as seeds for competition experiments (“Competitors”) and the mean values of the most abundant genotype at the end of evolutionary runs (“Mean Dominant”), both partitioned by resource-caching strategy. The rightmost table segments enumerate the population mean genotype values for standard evolutionary trials (“Pop Mean”) and control treatments (“Control Pop Mean”), matched at both absolute update count and (approximately) elapsed cellular generations. Two observations are missing from the standard evolutionary trial at update 2.1M due to server instability.



(a) Mean $P_c=0.03$, $P_1=0.75$, $P_2=0.23$; cell gen. 29920



(b) Mean $P_c=0.03$, $P_1=0.51$, $P_2=0.49$; cell gen. 33852



(c) Mean $P_c=0.08$, $P_1=0.01$, $P_2=0.90$; cell gen. 47507

Figure 2: End state of same-channel signaling networks in replicates where resource was exclusively allocated to first-level channel pools (2a), was split evenly between first- and second-level channel pools (2b), and was primarily allocated to second-level channel pools (2c). 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.




Figure 3: Progression of of same-channel level-one and level-two signaling networks states in an evolutionary run where level-two resource sharing 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.

443 ranging between individuality at the level of first-level same- 497
444 channel groups to individuality at the level of second-level 498
445 same-channel groups. Figure 2 shows the level-one and 499
446 level-two signaling networks at the end of runs where first-, 500
447 split-, and second-level resource allocation evolved, respec- 501
448 tively. First-level allocators form somewhat irregular level- 502
449 two amalgamations of diverse level-one networks. Second- 503
450 level allocators form highly regular diamond-shaped level- 504
451 two signaling networks. Split-allocation individuals exhibit 505
452 a level-two phenotype of intermediate regularity. Figure 3 506
453 shows a time series of signaling network snapshots in an 507
454 evolutionary run where second-level individuality evolved. 508

455 Table 1 summarizes predominant genotypes observed at 509
456 the end of our evolutionary simulations. All evolved geno- 510
457 types had A_2 fixed at 1.0. So, reproduction over cells shar- 511
458 ing the same level-two channel was universally avoided; 512
459 genotypes evolved so that cells declined to reproduce when 513
460 they were located at the interior of level-two same-channel 514
461 signaling networks. 515

462 However, a variety of resource-caching strategies evolved. 516
463 Most-abundant genotypes at the end of nine evolutionary 517
464 runs exclusively cached resource in organisms' level-one 518
465 signaling network's pool (i.e., $P_1 = 1.0$). We observed 519
466 strategies where resource was primarily, but not entirely, 520
467 cached in an organism's level-one signaling network pool 521
468 (i.e., $1.0 > P_1 > P_2$) as the most-abundant genotype at 522
469 the end of seven evolutionary runs. In one run, the most- 523
470 abundant final genotype split resources evenly between an 524
471 organism's level-one and level-two signaling network pool 525
472 ($P_1 = P_2 = 0.5$). Finally, we observed strategies where 526
473 resource was primarily, but not entirely, cached in an organ- 527
474 ism's level-two signaling network pool (i.e., $1.0 > P_2 >$ 528
475 P_1) as the most-abundant genotype at the end of 33 evolu- 529
476 tionary runs. 530

477 We suspect that a trade-off between growth rate and long- 531
478 term stability prompted the universal allocation of at least 532
479 some resource to level-one pools and/or cell stockpiles. 533
480 Cell- and level-one resource caching might function some- 534
481 thing like saving for a rainy day. Because reproduction 535
482 over level-two channel-mates was universally avoided, cells 536
483 and level-one same-channel networks situated at the inter- 537
484 rior of a larger level-two same-channel network do not ex- 538
485 pend their resource pools unless that larger level-two same- 539
486 channel network is damaged, exposing them to directly- 540
487 adjacent cells of a different level-two channel. Thus, re- 541
488 source accumulates in cell stockpiles and level-one pools un- 542
489 til the level-two same-channel network comes under stress. 543
490 Split allocation might also represent hedging against defec- 544
491 tion of a second-level channel-mate by via somatic mutation. 545

492 Indeed, we did observe selection for apoptosis in the 41 545
493 replicates where the dominant genotype employed second- 546
494 level resource caching. In these replicates, the average pop- 547
495 ulation mean value of M_c was 0.68 with standard devia- 548
496 tion 0.33, significantly greater than the value $M_c = 0.5$ we 549

would expect in the absence of a selective pressure on apop-
tosis response to mutation ($p < 0.001$, bootstrap test).


To assess whether heavy second-level resource alloca-
tors, which we characterize as higher-level individuals, were
more likely to employ apoptosis to mitigate somatic mu-
tation, we examined the relationship between first- and
second-level resource pooling and cellular apoptosis at the
conclusion of our 50 replicate evolutionary trials. We ob-
served a significant negative correlation between dominant
genotype P_1 and M_c ($p < 0.05$; bootstrap test; Figure
4a) and a significant positive correlation between dominant
genotype P_2 and M_c ($p < 0.05$; bootstrap test; Figure 4b).
This result suggests that second-level individuals, in partic-
ular, relied on apoptosis to mitigate somatic mutation.

We also assessed whether higher-level individuals pro-
vided larger resource endowments to their second-level
propagules (offspring sharing neither the level-one nor the
level-two channel ID with the parent). We examined the re-
lationship between first and second-level resource pooling
and dominant genotype second-level propagule endowment
at the conclusion of our 50 replicate evolutionary trials. We
observed a significant negative correlation between domi-
nant genotype P_1 and E_2 ($p < 0.05$; bootstrap test) and a
significant positive correlation between dominant genotype
 P_2 and E_2 ($p < 0.05$; bootstrap test). Second-level individ-
uals might provide larger endowments to propagules simply
due to a greater capacity to collect resource or perhaps be-
cause of stronger selection for well-endowed offspring when
competing against other second-level individuals.

This result prompts the reverse question: do lower-level
individuals provide larger resource endowments to first-level
propagules (offspring that do not share level-one channel ID
with the parent but may or may not share level-two chan-
nel ID with the parent)? Indeed, we observed a significant
positive correlation between first-level resource sharing and
first-level endowment ($p < 0.0001$; bootstrap test) and a sig-
nificant negative correlation between second-level resource
sharing and first-level endowment ($p < 0.0001$; bootstrap
test). Cells that pool resource with their smaller level-one
same-channel group tend to invest more heavily into the di-
rect offshoots of their level-one same-channel group than
cells that pool resource with their larger level-two same-
channel group. This observation suggests that, although
cells do not directly displace their level-one channel-mates,
competitive dynamics between may be at play.

Competition Experiments

Next, we wanted to compare first-, second-, and split-level
allocators to determine which genotype was the most fit.
We ran competition experiments between dominant geno-
types from evolutionary runs representative of each of these
strategies. To prevent further evolution, we disabled mu-
tation for these experiments. To represent first-level allo-
cators, we selected randomly from the nine pure first-level



(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 4: Plots of dominant resource caching strategies and dominant apoptosis strategies. A bootstrapped 95% confidence interval for the fit is shaded. Both correlations are statistically significant ($p < 0.05$; bootstrap test).

550 allocator dominant genotypes we observed. To represent the
 551 split-level allocators, we selected the single dominant geno-
 552 type where resource was partitioned exactly evenly between
 553 first- and second-level channel pools. To represent second-
 554 level allocators, we selected the dominant genotype with the
 555 largest second-level allocation proportion. Table 1 enumerates
 556 the three representative genotypes used. Figure 5 shows
 557 a time series of signaling network snapshots in an competi-
 558 tion experiment run. Colonies of each genotype can be seen
 559 to grow from each seed and then clash, ultimately yielding a
 560 population dominated by second-level allocators.

561 Indeed, the second-level resource caching strategy be-
 562 came most abundant in all 50 trials. Across the 50 replicates,
 563 at update 1.5 million (cellular generation 3489 with stan-
 564 dard deviation 40) the second-level resource caching strategy
 565 constituted 90.2%, with standard deviation 3.8%, of the
 566 competing population of cells. In the absence of mutation,
 567 second-level allocators tend to exhibit greater fitness than
 568 split- and first-level allocators ($p < 0.0001$; two-tailed exact
 569 test).

570 In competition experiments, however, higher-level indi-
 571 viduals likely benefited from elimination of somatic muta-
 572 tion. To assess the relative fitness of first- and second-level
 573 individuals without mutation disabled, we examined the rela-
 574 tionship between first- and second-level resource pooling
 575 and the rate of cellular reproduction at the end of each of
 576 the 50 replicate evolutionary trials performed. We observed
 577 a significant negative correlation between mean P_1 and cellu-
 578 lar reproduction rate ($p < 0.0001$; bootstrap test; Figure
 579 6a) and a significant positive correlation between mean P_2
 580 and cellular reproduction rate ($p < 0.0001$; bootstrap test;
 581 Figure 6b). This result suggests that second-level allocators

582 tend to collect resource more effectively than split- and first-
 583 level allocators.

584 Control Evolutionary Experiments

585 Under control conditions, we observed strong selection for
 586 high-level resource caching. At update 200,000, the aver-
 587 age population mean of P_2 was 0.96 with standard deviation
 588 0.08. For comparison, under the standard treatment the aver-
 589 age population mean of P_2 was 0.54 with standard devi-
 590 ation 0.11 at a time-point matched by absolute elapsed up-
 591 date count and 0.55 with standard deviation 0.22 at a time-
 592 point matched by approximate elapsed cellular generations¹
 593 (Table 1). We also observed strong selection against direct
 594 reproductive competition between channel-mates at update
 595 200,000; all evolved genotypes completely avoided repro-
 596 ducing over level-two channel-mates.

597 The emergence of resource-sharing and competition
 598 avoidance under control conditions suggest kin recognition
 599 alone can prompt some aspects of higher-level individual-
 600 ity. However, we observed selection *against* the apoptosis
 601 response to mutation, M_c , under control conditions. Across
 602 50 replicates of the control treatment, the average population
 603 mean value of M_c was 0.20 with standard deviation 0.23 —
 604 significantly less than the value $M_c = 0.5$ expected with-
 605 out selective pressure against apoptosis response to mutation
 606 ($p < 0.0001$, bootstrap test). Indeed, population mean M_c
 607 for control runs was also significantly reduced compared to
 608 the standard treatment at time-points matched by absolute
 609 elapsed update count ($p < 0.0001$; two-tailed t test) and by

¹This statistic is affected by missing datafile entries for two standard treatment replicates due to server instability.





Figure 5: Progression of same-channel level-one and level-two signaling networks states in a competition run. We seeded the grid with three copies of each of three champion genotypes from evolutionary trials. Then, with mutation disabled to prevent further evolution, the genotypes competed. 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 6: Mean resource caching strategies and net reproduction rate across populations. A bootstrapped 95% confidence interval for the fit is shaded. Both correlations are statistically significant ($p < 0.0001$; bootstrap test).




Figure 7: End state (update 249840, cell gen. 5384) of same-channel signaling networks evolved under the control treatment. 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.

610 approximate elapsed cellular generations ($p < 0.01$; two-
 611 tailed t test).² Perhaps under control conditions, the apop-
 612 tosis response to mutation is disfavored because kin groups
 613 stand to lose less from mutant members (i.e., the resource
 614 penalty for excessive same-channel network expansion is
 615 absent). It appears that, at least in our system, kin recog-
 616 nition alone does not suffice to prompt full-fledged fraternal
 617 transitions in individuality.

618 In the absence of resource penalties for erroneous activa-
 619 tion under control conditions, we also observed the evolu-
 620 tion of larger level-one same-channel groups. Compared to
 621 the standard treatment, control runs exhibited greater mean
 622 level-two same-channel caps C_2 at time-points matched by
 623 absolute elapsed update count ($p < 0.0001$; two-tailed t test)
 624 and approximate elapsed cellular generations ($p < 0.0001$;
 625 two-tailed t test).³ Even at 20 million updates, when evolu-
 626 tion had elapsed around ten times as many cellular gener-
 627 ations in the standard treatment compared to the control
 628 treatment at update 200,000, mean level-two same-channel
 629 caps C_2 reached only 262.9 with standard deviation 72.2
 630 under the standard treatment. This is significantly smaller
 631 than mean C_2 under the control treatment at update 200,000
 632 ($p < 0.0001$; two-tailed t test). Figure 7 depicts the compar-
 633 atively large same-channel level two groups present at the
 634 end of a control run.

²This comparison is affected by missing datafile entries for two standard treatment replicates due to server instability.

³This comparison is affected by missing datafile entries for two standard treatment replicates due to server instability.

635 **Conclusion** 685
 636 Using simple organisms that evolve parameters for a set 686
 637 of manually-designed strategies, we have demonstrated that 687
 638 DISHTINY selects for genotypes that exhibit high-level in- 688
 639 dividuality. We observed a spectrum of first- and second- 689
 640 level individuality among evolutionary outcomes. Specific- 690
 641 ally, we observed 691
 642 1. reproductive division of labor among members of the 692
 643 same channel (i.e., individuals enveloped in a same- 693
 644 channel signaling network ceded reproduction to those at 694
 645 the periphery), 695
 646 2. cooperation between members of the same channel (i.e., 696
 647 pooling of resource on same-channel signaling networks), 697
 648 3. reproductive bottlenecking (i.e., groups of cells sharing a 698
 649 channel ID descend from a single originator of that chan- 699
 650 nel ID), and 700
 651 4. suppression of somatic mutation via apoptosis coincident 701
 652 with second-level individuality. 702
 653 Competition experiments revealed that second-level in- 703
 654 dividuals usually outcompete lower-level individuals. The 704
 655 magnitude of resource endowment for propagules was also 705
 656 correlated with second-level individuality. 706
 657 Although shifts in individuality to level-one and level- 707
 658 two signaling networks were both observed, the question of 708
 659 whether these transitions were truly hierarchical in nature is 709
 660 debatable. That is, it is not clear whether level-one individ- 710
 661 uality was to some extent preserved in or necessary for the 711
 662 emergence of level-two individuality. Given the nature of 712
 663 the manually-designed strategies for resource-pooling and 713
 664 reproductive division of labor, level-two resource pooling 714
 665 and division of labor could readily leapfrog over level-one 715
 666 resource pooling and division of labor and, in many ways, 716
 667 seemed to completely supersede those level-one efforts. 717
 668 We believe that this is a shortcoming of the manual design 718
 669 of behaviors for which simple cell-like organisms evolved 719
 670 parameters, not the DISHTINY platform itself. We have 720
 671 nevertheless demonstrated that DISHTINY ultimately se- 721
 672 lects for high-level individuality. We are eager to work 722
 673 with more sophisticated cell-like organisms capable of arbi- 723
 674 trary computation via genetic programming in order to pur- 724
 675 sue more open-ended evolutionary experiments. We will 725
 676 also test the implications of relaxing current arbitrary re- 726
 677 strictions that artificially promote transitions, such as the hi- 727
 678 erarchical nesting of same-channel signaling networks and 728
 679 the explicitly-defined signaling networks themselves, leav- 729
 680 ing these details to evolution to figure out. Further work 730
 681 will provide valuable insight into scientific questions relat- 731
 682 ing to major evolutionary transitions such as the role of pre- 732
 683 existing phenotypic plasticity (Clune et al., 2007; Lalejini 733
 684 and Ofria, 2016), pre-existing environmental interactions, 734
 735
 736

pre-existing reproductive division of labor, and how transi-
 tions relate to increases in organizational (Goldsby et al.,
 2012), structural, and functional (Goldsby et al., 2014) com-
 plexity.

We believe that such an approach also provides a unique
 opportunity to fundamentally advance Artificial life with re-
 spect 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 exper-
 iments). Importantly, the platform is implemented in a de-
 centralized manner and can comfortably scale as additional
 computing resources are provided. Parallel computing is
 widely exploited in evolutionary computing, where subpop-
 ulations 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 paralleliza-
 tion of the evolving individual phenotype at arbitrary scale
 (i.e., a high-level individual as a large collection of individ-
 ual cells on the toroidal grid). Such parallelization will be
 key to realizing evolving computational systems with scale
 — and, perhaps, complexity — approaching biological sys-
 tems.

Acknowledgements

Thanks to members of the DEVOLAB, in particular Michael
 J. Wisner, 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. This material is based upon work
 supported by the National Science Foundation Graduate Re-
 search Fellowship under Grant No. DGE-1424871. Any
 opinions, findings, and conclusions or recommendations ex-
 pressed in this material are those of the author(s) and do not
 necessarily reflect the views of the National Science Founda-
 tion.

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., Fos-
 ter, J. A., McMullin, B., De Melo, V. V., Miconi, T., Spec-
 tor, L., Stepney, S., et al. (2016). Defining and simulat-
 ing open-ended novelty: requirements, guidelines, and chal-
 lenges. *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.

- 737 Bouchard, F. (2013). What is a symbiotic superindividual and how
738 do you measure its fitness. *From groups to individuals: evolution and emerging individuality*, 243.
739
- 740 Clune, J., Ofria, C., and Pennock, R. T. (2007). Investigating the
741 emergence of phenotypic plasticity in evolving digital organisms. In *European Conference on Artificial Life*, pages 74–83.
742 Springer.
743
- 744 Ereshefsky, M. and Pedroso, M. (2015). Rethinking evolutionary
745 individuality. *Proceedings of the National Academy of Sciences*, 112(33):10126–10132.
746
- 747 Goldsby, H. J., Dornhaus, A., Kerr, B., and Ofria, C. (2012). Task-
748 switching costs promote the evolution of division of labor and
749 shifts in individuality. *Proceedings of the National Academy of Sciences*, 109(34):13686–13691.
750
- 751 Goldsby, H. J., Knoester, D. B., and Ofria, C. (2010). Evolution
752 of division of labor in genetically homogenous groups. In *Proceedings of the 12th annual conference on Genetic and
753 evolutionary computation*, pages 135–142. ACM.
754
- 755 Goldsby, H. J., Knoester, D. B., Ofria, C., and Kerr, B. (2014).
756 The evolutionary origin of somatic cells under the dirty work
757 hypothesis. *PLoS biology*, 12(5):e1001858.
- 758 Goldsby, H. J., Young, R. L., Hofmann, H. A., and Hintze, A.
759 (2017). Increasing the complexity of solutions produced by
760 an evolutionary developmental system. In *Proceedings of the
761 Genetic and Evolutionary Computation Conference Companion*, pages 57–58. ACM.
762
- 763 Lalejini, A. and Ofria, C. (2016). The evolutionary origins of phe-
764 notypic plasticity. In *Proceedings of the Artificial Life Conference*.
765
- 766 Lin, S.-C., Punch, W. F., and Goodman, E. D. (1994). Coarse-
767 grain parallel genetic algorithms: Categorization and new ap-
768 proach. In *Parallel and Distributed Processing, 1994. Proceedings. Sixth IEEE Symposium on*, pages 28–37. IEEE.
769
- 770 Queller, D. C. (1997). Cooperators since life began. *The Quarterly
771 Review of Biology*, 72(2):184–188.
- 772 Ratcliff, W. C., Denison, R. F., Borrello, M., and Travisano, M.
773 (2012). Experimental evolution of multicellularity. *Proceed-
774 ings of the National Academy of Sciences*, 109(5):1595–1600.
- 775 Ray, T. S. (1996). Evolving parallel computation. *Complex Sys-
776 tems*, 10:229–237.
- 777 Real, E., Moore, S., Selle, A., Saxena, S., Suematsu, Y. L., Tan, J.,
778 Le, Q. V., and Kurakin, A. (2017). Large-scale evolution of
779 image classifiers. In Precup, D. and Teh, Y. W., editors, *Pro-
780 ceedings of the 34th International Conference on Machine
781 Learning*, volume 70 of *Proceedings of Machine Learning
782 Research*, pages 2902–2911, International Convention Cen-
783 tre, Sydney, Australia. PMLR.
- 784 Smith, J. and Szathmary, E. (1997). *The Major Transitions in Evo-
785 lution*. OUP Oxford.
- 786 Taylor, T., Bedau, M., Channon, A., Ackley, D., Banzhaf, W.,
787 Beslon, G., Dolson, E., Froese, T., Hickinbotham, S.,
788 Ikegami, T., et al. (2016). Open-ended evolution: per-
789 spectives from the oee workshop in york. *Artificial life*,
790 22(3):408–423.
- 791 West, S. A., Fisher, R. M., Gardner, A., and Kiers, E. T. (2015).
792 Major evolutionary transitions in individuality. *Proceedings
793 of the National Academy of Sciences*, 112(33):10112–10119.