Endosymbiotic mutualism can constrain host diversity and evolved complexity

Delaney Kelley[1,*], Owen Hillary[1,*], Elias Hillman-Emelianoff[2], Ahmad Khalili[2], Ruoyu Zhang[2,4],
Kiara M. Johnson[3], Emily Dolson[3], Anya E. Vostinar[2], and Alexander Lalejini[1]
[1]Grand Valley State University, United States
[2]Carleton College, United States
[3]Michigan State University, United States
[4]University of Michigan, United States
[*]These authors contributed equally to this work.
lalejina@gvsu.edu

Abstract

Coevolutionary arms races between hosts and parasites are known to promote the evolution of complex traits in hosts. However, the evolutionary effects of mutualistic endosymbionts (symbionts that live inside a host) are less well understood. Here, we use populations of self-replicating computer programs (digital organisms) to investigate the effects of trait matching mutualisms on evolution. We replicate previous findings showing that parasitic endosymbionts increase host diversity and promote the evolution of complex traits. Contrastingly, we find that coevolution with mutualistic symbionts decreases host diversity and can constrain the evolution of the host species by imposing stabilizing selection to maintain symbiosis. Under strict trait matching conditions, mutualistic symbionts in our experiments often undergo a form of egalitarian transition in individuality with their hosts, losing the ability to reproduce independent of their host. The constraining effects of mutualism depend on several factors: environmental benefits available to hosts that forgo mutualism, strength of mutualistic interactions, and restrictions on host-mutualist compatibility. Strikingly, when we experimentally remove mutualists from an evolutionarily constrained host population, hosts become unconstrained, which rapidly changes the trajectory of host evolution.

Submission type: **Full Paper**

Data/Code available at: See reference (Kelley et al., 2026a)

Introduction

Symbioses encompass long-term evolutionary relationships between two or more species that benefit at least one of the constituent partners. A symbiont's impact on its partner(s) can range from detrimental (parasitism) to beneficial (mutualism). Symbiotic relationships of all kinds pervade nature, profoundly influencing the evolution of participating species (Moran, 2007; Bennett and Moran, 2015; Cavalier-Smith, 2013). Prior work has shown that parasites can increase diversity and promote the evolution of complex traits in their host species through a combination of evolutionary arms races and population-level "genetic memory" (Zaman et al., 2014, 2011). Those results then raise the question: does mutualism have similar effects on host evolution? In this work, we use coevolving populations of digital hosts and endosymbionts (*i.e.*, symbionts that live inside of their host) to investigate mutualism's effect on host diversity and the evolution of complex traits in hosts.

Strong links exist between mutualism and the origination and maintenance of species diversity (Chomicki et al., 2019). However, mutualism's effect on diversification depends on the mutualist's function, partner dependence, partner specificity, and mechanisms of partner interaction (Chomicki et al., 2019; Yoder and Nuismer, 2010; Raimundo et al., 2014). For example, mutualisms may directly promote diversification when populations experience divergent selection to specialize on different mutualist partners (Whitehead and Peakall, 2014) or when mutualism creates new ecological opportunities for partner species, such as enabling access to new resources (Joy, 2013; Gómez and Verdú, 2012; Salem et al., 2017; Kaltenpoth et al., 2025). Mutualism can also indirectly increase diversity by decreasing extinction risk when symbiosis enables increased range sizes or improves survivability (Lengyel et al., 2009).

Via its effects on diversification, mutualism also indirectly affects evolutionary potentiation and the evolution of complexity, as diversification and the evolution of complexity are linked (Svardal et al., 2014; Dolson et al., 2017; Erwin, 2015). However, mutualism's effects on the evolution of complex traits in symbiotic species can be more direct. According to the symbiotic scaffolding hypothesis, mutualists can provide functional scaffolds to their hosts, enabling the host species to evolve more complex traits before eventually evolving less dependence on their symbiotic partner (Watson et al., 2000). In the extreme, mutualisms can lead to egalitarian major transitions in individuality (symbiogenesis), wherein formerly distinct individuals unite to form a single reproductive entity (Estrela et al., 2016). For example, mitochondria and chloroplasts in eukaryotic cells are recognized to be the result of symbiogenesis (Koonin, 2015). Mutualisms can also cause loss of complexity via evolutionary degradation in obligate endosymbionts (Vostinar et al., 2021) by decreasing selective pressure on vertically transmitted endosymbionts, decreasing opportunities for endosymbionts to purge deleterious mutations, reducing the need for costly genes (Black queen hypothesis, Morris et al. 2012), or hosts protecting endosymbionts from competition. Mutualism might also cause selection against evolutionary change in interacting species if that change would disrupt or destabilize the mutualism. For example, a cooperative mutualism between two genotypes of *Escherichia coli* has been shown to constrain evolutionary adaptation to environmental stress (Pauli et al., 2022). However, such intraspecies cooperation (and more broadly, ectosymbiosis) does not always exert the same evolutionary pressures as endosymbiotic mutualism (Johnson et al., 2022).

In this work, we model trait matching endosymbiotic mutualisms. Trait matching, or phenotype matching, is a common mechanism that mediates symbiotic interactions. In such symbioses, interactions between individuals depend on the compatibility of relevant traits. For example, mutualisms between fruit-eating birds and fleshy-fruited plants depend on the compatibility of bird beak morphology with fruit size and shape (Dehling et al., 2014). Mutualisms that rely on trait matching can impose stabilizing selection on traits involved in the mutualism, as changes to relevant traits may disrupt the symbiotic relationship. This stabilizing selection is hypothesized to constrain the evolution of species engaged in mutualism (Chomicki et al., 2019; Yoder and Nuismer, 2010). Indeed, trait matching has been linked with slower morphological evolution in plant-bird mutualisms, pollination (Davis et al., 2014), and epiphytic ant-plant mutualisms (Chomicki and Renner, 2017).

Here, we test the effects of trait matching endosymbiotic mutualism using the Symbulation digital evolution system (Vostinar et al., 2022). In contrast to the effects of parasites (Zaman et al., 2014, 2011), we find that mutualist endosymbionts constrain the evolution of complex traits in hosts and reduce diversity in host populations. However, this constraining effect depends on the potential for hosts to engage in generalist strategies and the relative benefit of engaging in mutualism versus alternative environmental opportunities. We also observe evolutionary degradation in endosymbionts, wherein symbionts become a fully dependent "organelle" inside of their host.

Methods

Study system

We conducted our experiments using Symbulation, an agent-based system for studying host-symbiont coevolution (Vostinar et al., 2022). Symbulation supports a variety of symbiotic relationship types, modes of host-symbiont interactions, and genetic representations. We used Symbulation to evolve populations of host organisms that compete for space in a well-mixed virtual environment with a maximum carrying capacity of 10,000 hosts. Each host could contain at most one endosymbiont, and endosymbionts could not survive outside of a host. Thus, endosymbionts competed for space in the limited host population. Hereafter, we use the generic term "organism" to discuss shared characteristics of both hosts and endosymbionts, and we use the term "symbiont" to refer to the endosymbionts in our system.

Digital organisms Each organism comprises a linear sequence of program instructions (its genome) and virtual hardware used to express those instructions. We fixed genome length at 100 instructions, and we use the SignalGP genetic representation, which allows for tag-based execution flow control (Lalejini and Ofria, 2018). Organisms execute their genome sequentially, except when the execution of one instruction changes which instruction is executed next (*e.g.*, a `jump` instruction). Genomes are circular, so execution wraps from the end back to the beginning. The instruction set contains operations for input, output, basic computations, flow control, and reproduction. We document the instruction set in our supplemental material (Kelley et al., 2026a).

Symbulation measures time in "updates". On an update, each organism gains 16 "CPU cycles" with which to execute the next 16 instructions in its genome. Each update, we randomize the order that hosts are chosen to execute their genomes. Symbionts always execute their genome immediately prior to their host.

Organisms collect resources by performing any of the following nine Boolean logic functions: NAND, NOT, ORN (or-not), AND, OR, ANDN (and-not), NOR, XOR, and EQU (equals). The resources gained for completing each logic function varied across experiments. We chose these Boolean logic functions to align with the default Logic-9 environment in the Avida Digital Evolution Platform (Ofria and Wilke, 2004), which is well-studied and has been used in related studies on coevolution (Zaman et al., 2014), complexity (Lenski et al., 2003), and diversity (Dolson and Ofria, 2017; Dolson et al., 2020). The Logic-9 functions also allow for a direct measure of evolved trait complexity (Lenski et al., 2003). We include a `nand` instruction in the instruction set, and all other functions can be implemented with different compositions of `nand` instructions. We measure function complexity as the minimum number of `nand` instructions required to encode that function: NAND and NOT require 1 `nand`, ORN and AND require 2, OR and ANDN require 3, NOR and XOR require 4, and EQU requires 5.

In addition to resource collection, these logic functions determine whether a symbiont interacts with its host. To engage in symbiosis, a host and its symbiont must have performed a matching logic function. For most of our experiments, we permit organisms to perform one type of Logic-9 function during their lifetimes. This function can be performed repeatedly for the purpose of resource collection. This restriction creates evolutionary trade-offs for hosts and symbionts, as evolving to perform a different function can disrupt an existing symbiotic relationship.

Reproduction Hosts reproduce asexually by collecting a parameterized number of resources and executing a `reproduce` instruction, consuming resources equal to the reproduction cost. We apply point mutations to offspring: 0.04 per-instruction and per-argument substitution rates and a 0.005 per-bit tag mutation rate. Offspring are placed in a random location, replacing any previous occupant.

Symbionts can transmit *vertically* or *horizontally*. When a host with an symbiont reproduces, that host's symbiont can transmit its offspring into the host's offspring without a resource cost (*vertical transmission*). Symbionts can also attempt to transmit offspring horizontally into a random other host in the population. Symbionts attempt *horizontal transmission* by collecting sufficient resources and executing a `reproduce` instruction. When a symbiont attempts transmission in this way, we select a random host as an infection target. If the infection target has no symbiont, the reproducing symbiont transmits its offspring into the targeted host. If the infection target already has a symbiont, the incoming symbiont can replace the current symbiont if it is a strictly better function match with the infection target. For a single `reproduce` execution by a symbiont, we allow up to five horizontal transmission attempts on randomly selected host targets until an attempt is successful. If the symbiont fails to transmit five times, it fails to produce an offspring and does not incur a resource cost.

Population initialization We initialized populations with NOT-capable hosts, each with a NOT-capable endosymbiont. This set-up ensures that all symbionts initially engage in symbiosis with their hosts. Initial genomes included instructions for performing the NOT function and reproduction. They were otherwise filled with inert no-operation instructions.

Host-endosymbiont interactions Function-matching host-symbiont pairs interact with one another on each update. Symbulation can be configured for mutualism, parasitism, or allow evolution along a spectrum of mutualism and parasitism. In this work, we based interactions between symbionts and their hosts on prior work (Zaman et al., 2014). When interacting, mutualistic symbionts donate a portion of their CPU cycles to their host for the given update, and parasitic symbionts steal a portion of their host's CPU cycles. The number of CPU cycles exchanged (for both mutualists and parasites) is multiplied by a configurable synergy factor, modeling efficiencies that can be gained from division of labor and/or specialization between a host and its mutualist or the asymmetry of resource value between a host and its parasite. For mutualist interactions, the number of CPU cycles received by the host equals the number donated by the symbiont multiplied by the synergy factor. For parasitic interactions, the number of CPU cycles received by the symbiont equals the number stolen from the host multiplied by the synergy factor.

Experimental design

Does mutualism promote the evolution of complex traits in hosts? Using Symbulation, we reproduced the experimental setup from (Zaman et al., 2014) demonstrating that host-parasite coevolution increases host diversity and promotes the evolution of complexity in hosts. We compared evolutionary outcomes across three treatments, running 50 replicates of each: a parasite condition where all symbionts engaged in parasitic interactions, a mutualist condition where all symbionts engaged in mutualistic interactions, and a no-symbiont control. When hosts and parasites interacted each update, parasites stole 75% of their host's CPU cycles. When hosts and mutualists interacted, mutualists donated 75% of their CPU cycles to their host. Across both parasite and mutualist conditions, we used a synergy factor of 4. In treatments with symbionts, any symbiont could reproduce horizontally. In the mutualist condition, hosts vertically transmitted their mutualist partner; however, in the parasite condition, parasites were not vertically transmitted.

We configured each of the Logic-9 functions to provide one resource when performed, imposing no external selection pressure for organisms to perform complex logic functions over simpler functions. In this experiment, more complex functions are disadvantageous as compared to simpler functions: simpler functions require fewer instructions, making them more mutationally robust and more efficient.

We set the cost of host reproduction and symbiont horizontal transmission to one resource, and we imposed minimum lifespans of 100 CPU cycles for hosts and 10 CPU cycles for symbionts, modeling a need to reach maturity before reproducing. In exploratory experiments, we found these minimum execution requirements to be necessary to prevent hosts from rapidly evolving flow control that bypassed executing most of their genome, which made evolution beyond the initial logic function in any condition rare.

Does mutualism constrain the evolution of complex traits in hosts? In our second experiment, we repeat the previous experiment, except we reconfigure the environment and reproduction parameterizations to allow us to test whether mutualism constrains the evolution of complex traits. In our first experiment, failing to observe complex trait evolution does not indicate constraint, as there is no external selection pressure to promote the evolution of complex functions outside of selective forces from coevolving symbionts. For this experiment, we configure increasing resource rewards for the Logic-9 functions such that more complex functions yield more resources than simpler functions: 1 resource gained for NOT and NAND, 2 for ORN and AND, 4 for OR and ANDN, 8 for NOR and XOR, and 16 for EQU. We also configure reproduction to be resource limited: host reproduction costs 256 resources and symbiont horizontal transmission costs 128 resources. For example, a host that performs NOT must perform at least 256 NOT functions to reproduce, whereas a host could perform just 32 XOR functions to accrue sufficient resources to reproduce. Under this configuration, we expect hosts in our no-symbiont control to evolve more complex functions, and symbiont treatments where fewer complex functions evolve can be said to constrain evolution relative to the control.

How does fitness landscape, strength of mutualism, and capacity for generalism influence mutualism's effect on host evolution? We next conducted four sets of experiments to characterize how environmental rewards, strength of mutualism, and capacity for function generalism effect host-mutualist coevolution. For each of these experiments, we use the same experimental setup as in our previous experiment (described above), except for any experiment-specific focal parameters specified below. For each of these experiments, we compared host-mutualist coevolution with no-symbiont controls. We ran 50 replicates of each condition, and all runs across all experiments were independent with unique random number seeds.

Function rewards: We compared evolutionary outcomes under six reward levels for complex functions (given as a multiplier to our previous baseline): x0.5, x1, x2, x4, x8, and x16. For each reward level, the two simplest functions (NOT and NAND) always yield one resource; however, the benefits of all complex functions are increased by the reward-level multiplier. Increasing complex function rewards directly increases selection pressure on organisms to evolve more complex functions, as doing so benefits a host's ability reproduce and a symbiont's ability to horizontally transmit.

Mutualist donate proportion: We tested five mutualist CPU cycle donate proportions: 12.5%, 25%, 50%, 75%, 87.5%, and 100%. Increasing donate proportion increases the impact of mutualist interactions by allowing symbionts to transfer a greater proportion of their CPU cycles to their host, indirectly improving their rate of vertical transmission.

Synergy factor: We tested five synergy factors: 0.5, 1, 2, 4, and 8. A synergy factor of 1 represents no additional synergy between a host and its symbiont: the host receives the number of CPU cycles donated by its mutualist. Synergy factors greater than 1 model efficiencies gained through effective division of labor, and synergy factors less than 1 model net inefficiencies in symbiotic interactions.

Function restrictions: Restricting organisms to one type of Logic-9 function controls their capacity for generalist strategies. We characterize the effects of restricting both hosts and symbionts, restricting just hosts, restricting just symbionts, and not restricting hosts or symbionts.

Metrics and statistical analyses

For all experiments, we tracked function performance, number of interacting host-symbiont pairs, vertical and horizontal transmission rates, generations elapsed, and population-level phenotypic diversity of symbionts and hosts (as counts of function profiles and entropy of function profiles). For brevity, we do not visualize or discuss all metrics in this manuscript; however, all are shown in supplemental material. When measuring evolved function complexity, we identify the highest-complexity function performed in a population over the course of the entire run to account for any cyclic dynamics caused by host-symbiont coevolution. To categorize a function as "performed" by a population, we require at least 2.5% of that population to have performed it. This threshold avoids misattributing rare instances where random outputs from an organism happen to match with the correct output for a complex function given the organism's prior inputs.

When comparing two distributions, we use a Wilcoxon rank-sum test (Wilcoxon, 1992) with a significance threshold of 0.05. When comparing more than two distributions, we screen for statistical differences in distributions using a Kruskal-Wallis test (Kruskal and Wallis, 1952). For comparisons with a significant Kruskal-Wallis test, we perform pairwise Wilcoxon rank-sum tests to identify individual differences. When performing multiple comparisons, we apply a Bonferroni correction as appropriate (Chen et al., 2017).

Code and data availability

Configuration files and code are on GitHub and publicly archived with Zenodo (Kelley et al., 2026a). Our supplemental material documents the version of Symbulation used for experiments in this work. Aggregated and compressed experiment data are archived on OSF (Kelley et al., 2026b).

Results and Discussion

Mutualism fails to promote complex function evolution in hosts

We recreated experiments from (Zaman et al., 2014) showing that host-parasite coevolution can promote complex trait evolution, even when complex traits do not yield greater resource benefits. Here, however, we compare outcomes of hosts coevolved with parasites, mutualists, and a no-symbiont control.

Figure 1 shows the level of host function complexity that evolved across treatments. Overall, hosts coevolved with parasites evolved significantly more complex functions than hosts coevolved with mutualists and hosts evolved without symbionts (Bonferroni-corrected Wilcoxon rank-sum tests, $p < 10^{-4}$). Hosts coevolved with mutualists and hosts evolved in the no-symbiont control failed to evolve any more complex functions than the initial NOT function. In contrast, we observed the evolution of EQU (the most complex function) in 16 of 50 replicates with host-parasite coevolution.

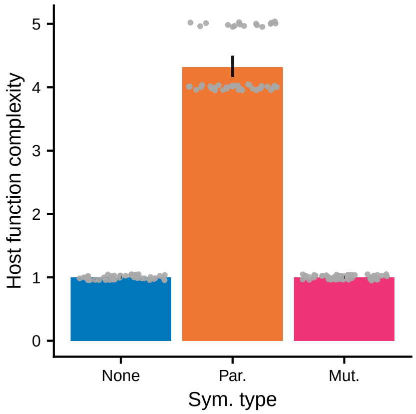

Figure 1: Maximum function complexity evolved by host populations using equal function rewards with no symbionts ("None"), parasites ("Par."), and mutualists ("Mut."). Measurements for individual replicates are given by gray points, bars indicate mean function complexity for each treatment, and lines on the top of each bar indicate a bootstrapped 99% confidence interval around the mean. A Kruskal-Wallis test was statistically significant ($p < 10^{-4}$).

We observed differences in the number of elapsed generations across treatments, indicating variance in the amount of evolution that occurred. Hosts coevolved with mutualists evolved for substantially *more* generations ($\mu = 28125$, $\sigma = 121$) than hosts coevolved with parasites ($\mu = 8621$, $\sigma = 850$) and hosts evolved without symbionts ($\mu = 11600$, $\sigma = 83.4$). Likewise, mutualist symbionts evolved for more generations ($\mu = 30153$, $\sigma = 127$) than parasites ($\mu = 8173$, $\sigma = 1584$). Differences in generational turnover stem from how symbionts interact with their host. Successful parasites steal CPU cycles from hosts, reducing their host's rate of reproduction. Conversely, successful mutualists donate CPU cycles to hosts, which increases their host's rate of reproduction and their own rate of reproduction via vertical transmission. In the absence of other factors, we would expect more complex function evolution in populations evolved for a greater number of generations. However, hosts in the parasite treatment evolved more complex functions with far fewer generations, indicating that the greater functional complexity of these hosts is not due to differences in the rate of evolution observed across treatments.

These results replicate Zaman et al. 2014's core finding in a new study system: coevolution with parasites can promote the evolution of complex traits in host populations. As in (Zaman et al., 2014), we found that host-parasite populations exhibited the greatest phenotypic diversity as measured by both phenotypic richness and Shannon diversity both over time (supplemental material, Kelley et al. 2026a) and at the end of the experiment (corrected Wilcoxon rank-sum tests, $p < 10^{-4}$). Our experiment places mutualism in the context of Zaman et al. 2014's original experiments, showing that mutualism fails to promote the evolution of host complexity as parasitism does.

Mutualism constrains host diversity and complexity

In the previous experiment, complex traits did not benefit organisms more than simple traits, so it is unsurprising that hosts with mutualists did not evolve complex traits. To determine if mutualism would *constrain* the evolution of complex traits in hosts, we next considered an environment where more complex traits were selectively advantageous. To do so, we repeated the previous experimental setup, except we configured more complex Logic-9 functions to confer increasing rewards.

Figure 2A shows evolved function complexity in host populations across treatments. Overall, hosts in the no-symbiont control and parasitism condition evolved more complex functions than hosts coevolved with mutualists (corrected Wilcoxon rank-sum test, $p < 10^{-4}$). Host populations in the mutualism treatment were the least phenotypically diverse as compared to hosts in the no-symbiont control and hosts coevolved with parasites (Figure 2B). All no-symbiont host populations evolved at least one of the second-most complex functions (NOR or XOR). Hosts in seven replicates of the parasitism condition evolved EQU (the most complex function), and all other replicates reached at least NOR or XOR.

In contrast to our previous experiment, we observed symbiont population extinctions in 36 of 50 host-parasite replicates. No mutualist populations were driven to extinction, and in general, all mutualist populations remained close to their maximum carrying capacity for the duration of the experiment. Parasite extinctions occurred when hosts escaped infection vulnerability by evolving a more complex function. Because performing more complex functions conferred more resources than less complex functions, newly immune hosts could rapidly out-compete the still-vulnerable host subpopulation, wiping out infected hosts and their parasites before the parasite population could adapt.

In 48 of 50 replicates of the mutualist condition, mutualists fully constrained the evolution of complex functions in hosts. In the two replicates where mutualists failed to constrain host evolution, the host and mutualist populations were no longer evolutionarily engaged: hosts evolved away from their symbionts by evolving a new function, allowing hosts to gain more resources for performing a more valuable function but losing the benefit of the mutualism. Across all replicates of the mutualism condition, we found that mutualists rapidly lost the ability to reproduce horizontally, relying on vertical transmission for reproduction. In replicates where hosts and mutualists disengaged, the symbionts resembled a "vestigial organelle" in their host; these symbiont lineages rapidly lost the ability to perform their original logic function because it no longer benefited the host's rate of reproduction (and they no longer had the genetic code to attempt horizontal transmission).

As in our previous experiment, we observed differences

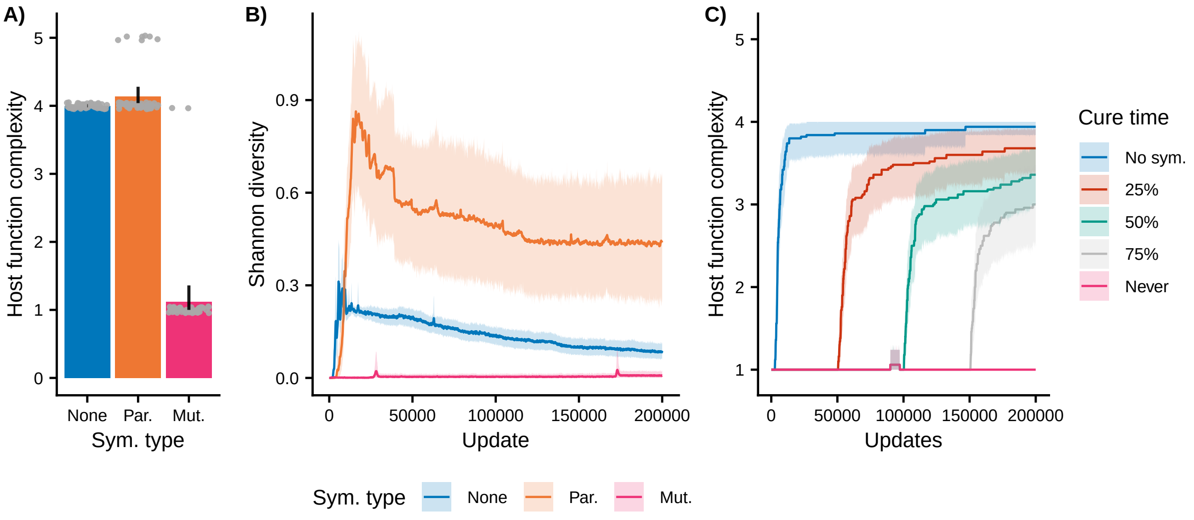

Figure 2: Evolutionary outcomes with differential function rewards. Plot A) shows maximum host function complexity (same format as in Fig. 1; Kruskal-Wallis test, $p < 10^{-4}$), B) shows Shannon diversity of host phenotypes over time, and C) shows host function complexity over time in the cure-mutualist followup experiment. In plots B and C, lines represent means, and shaded regions show corresponding bootstrapped 99% confidence intervals. Note the difference in y-axis ranges for plots A and C.

in the distributions of elapsed generations across treatments. Hosts in the mutualism condition evolved for the greatest number of generations ($\mu = 4333$, $\sigma = 150$; corrected Wilcoxon rank-sum test, $p < 10^{-4}$) relative to hosts evolved with parasites ($\mu = 3069$, $\sigma = 1223$) and hosts evolved without symbionts ($\mu = 3887$, $\sigma = 400$). These results demonstrate that trait-matching mutualisms can constrain the host evolution despite increased generational turnover gained from mutualistic benefits. Further, the symbionts can undergo evolutionary degradation, losing the ability to reproduce independently and locking them to their host.

Eliminating mutualist symbionts unconstrains host evolution Next, we tested whether host evolution would remain constrained if mutualists were removed from the host population. We repeated the mutualist and no-symbiont treatments from the previous experiment, except we eliminated all symbionts in the mutualist treatment 25%, 50%, and 75% of the way through the experiment (Figure 2C). Immediately after eliminating symbionts (regardless of timing), the host population rapidly evolved more complex traits, indicating that external disruptions to an existing mutualism can have dramatic effects on subsequent evolution of any previously partnered species.

The benefits of mutualism and opportunities for host generalism influence evolutionary constraint

Finally, we investigated the impact of factors that could remove the constraining effect of the mutualism: adjusting the benefits of mutualism relative to alternative strategies and varying the capacity for organisms to be generalists.

Hosts escape their mutualists when complex function rewards are sufficiently high Function rewards control the benefit of evolving more complex functions for both hosts and symbionts. We compared evolved function complexity in hosts coevolved with mutualists and evolved without symbionts in six environments with increasing function rewards for complex functions: x0.5, x1, x2, x4, x8, x16. To create these environments, we multiplied the rewards of all functions more complex than NOT and NAND in our standard environment (x1) by the associated multiplier. For example, in the x4 environment, all complex functions provided four times as many resources as in our standard environment, but NOT and NAND still provided one resource each.

Figure 3A shows host maximum function complexity evolved across different environments. Mutualism constrains host evolution when complex functions have smaller rewards (x0.5 and x1). In the x0.5 environment, we observed a bimodal distribution of function evolution in the no-symbiont control; hosts evolved complex functions in 25 of 50 replicates, and hosts failed to evolve beyond the NOT function in the remaining 25 replicates. Increasing function rewards above our standard environment decreased the number of host populations that were constrained by their mutualists. We did not observe significant differences in evolved host complexity between our control and coevolution populations in the x2, x4, x8, and x16 environments. There were two replicates in the x2 environment where mutualists constrained the host population to the NOT function. As complex function reward levels increased beyond x2, we observed no host populations that were constrained to the NOT function. More broadly, these results show that mutualism fails to constrain hosts when the benefits of evolving away

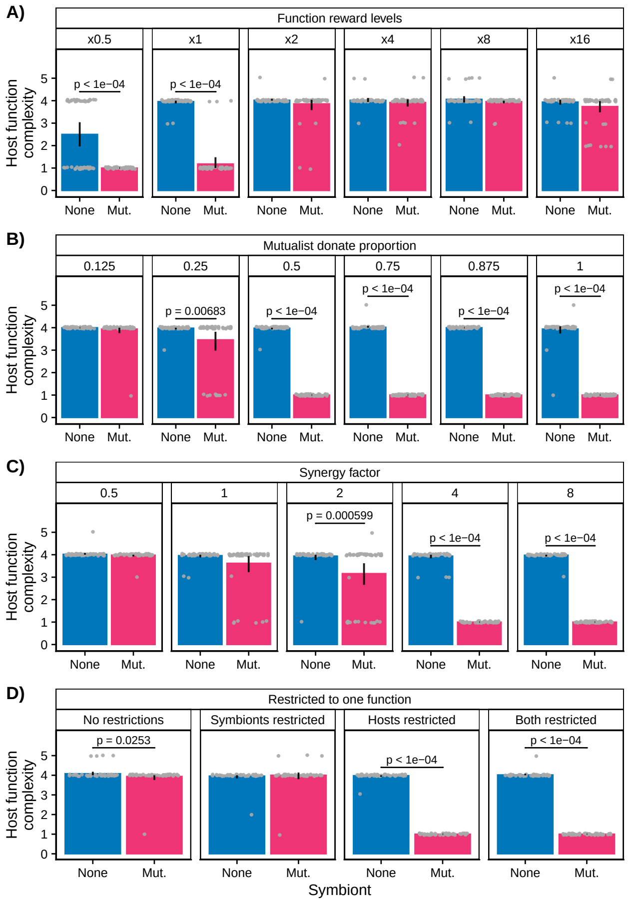

Figure 3: Maximum evolved host function complexity for different A) function reward levels, B) mutualist donate proportions, C) synergy factors, and D) function expression restrictions (*e.g.*, whether hosts and/or symbionts are restricted to specializing on one type of function). Individual bar plots are given in the same format as in Fig. 1, and statistically significant differences (Wilcoxon rank-sum test) between each no-symbiont control and mutualism treatment are indicated above the bars with the given p-value.

from a mutualistic interaction sufficiently outweigh the benefits gained from the interaction.

Across all environments, vertical transmission was the dominant mode of reproduction for symbionts (as in our previous experiment). However, as the reward for more complex functions increased, we observed that mutualist populations were increasingly likely to maintain the ability to reproduce horizontally.

High impact mutualism constrains host evolution We control the strength of mutualism by varying the proportion of CPU cycles that mutualists donate to their host each update. In our prior experiments, mutualists that function match with their host donate 75% of their CPU cycles to their host. To determine the impact of mutualism strength, we compared evolved function complexity in host populations without symbionts and in host populations coevolved with mutualists at each of the following CPU cycle donate proportions: 12.5%, 25%, 50%, 75%, 87.5%, and 100%.

Figure 3B shows host function complexity evolved across different mutualist donate proportions. Overall, we find statistically significant evidence of constraint (relative to no symbiont controls) at all donate proportions 25% and greater. At donate proportions 50% and greater, mutualists constrained host complexity in all replicates. At a 25% donate proportion, just nine replicates with host-mutualist coevolution exhibited constraint. In all other replicates, hosts evolved more complex functions, escaping their mutualist partners. This result indicates that even weak mutualisms can constrain host evolution; however, the likelihood of constraint decreases as the mutualistic benefit decreases.

At donate proportions 75% and greater, all mutualist populations lost the ability to reproduce independently of their hosts. At lower donate proportions, some mutualist populations maintained the ability to reproduce horizontally, but vertical transmission remained the dominant mode of reproduction (see supplemental material, Kelley et al. 2026a).

Increasing host-mutualist synergy increases the likelihood of constraint In our experiments, host-mutualist synergy is modeled as a multiplier on the number of CPU cycles that hosts receive from a function-matching mutualist. We tested host-mutualist coevolution at five synergy factors: 0.5, 1, 2, 4 (used in other experiments), and 8.

Figure 3C shows the maximum function complexity evolved in hosts coevolved with mutualists compared to hosts evolved without mutualists across five synergy factors. Overall, the number of host populations constrained by mutualism increased as we increased the host-mutualist synergy factor. The constraining effects of mutualism were strongest at the largest synergy factors: all host populations coevolved with mutualists failed to evolve functions more complex than NOT at synergy factors 4 and 8. At synergy factor 2, 14 of 50 host populations in the mutualism condition failed to evolve beyond the NOT function. At synergy factor 1, 6 host populations in the mutualism condition failed to evolve beyond the NOT function; however, the absence of additional interaction synergy failed to consistently constrain host populations, as 44 of 50 replicates evolved more complex functions. When the host-mutualist interaction was inefficient (synergy factor 0.5), no host populations were constrained to simple functions by mutualists, and indeed, all of these host-mutualist populations were no longer evolutionarily engaged with one another. This result indicates that the constraining effect of the mutualism depends on its relative benefit, but even a fairly small relative benefit is sufficient to constrain host evolution.

At synergy factors 0.5, 4, and 8, symbiont populations evolved to rely on vertical transmission, rapidly losing the ability to independently reproduce. However, at synergy factors 4 and 8, symbionts were evolutionarily engaged with their hosts for the duration of the experiment, and at synergy 0.5, symbionts rapidly disengaged with their hosts, becoming vestigial free-riders. This result raises questions about the effect of the mutualism on the symbionts' evolutionary trajectory and is a promising direction for future work.

Mutualisms with obligate specialization can constrain host complexity In all prior experiments, hosts and symbionts were required to specialize on one of the Logic-9 functions. That is, once an organism performed a particular logic function, that organism could perform that function repeatedly to earn resources but would not have been credited for performing any other functions. Because function-matching determines whether a symbiont interacts with its host, specialization of function expression can strongly impact coevolution. For example, limiting hosts to a single function can prevent a host lineage from evolving to perform a different function without disrupting an existing mutualism. We expect mutualism to be more easily maintained when hosts are able to be generalists; that is, they are not limited to a single function, as new functions can be gained without disrupting mutualistic interactions. When hosts can evolve to be generalists, we would not expect mutualism to constrain host evolution. To test this hypothesis, we evolved hosts with and without mutualists under four types of function expression restrictions: both hosts and mutualists restricted to expressing a single function (as in prior experiments), hosts restricted and symbionts unrestricted, hosts unrestricted and symbionts restricted, and neither hosts nor symbionts restricted.

Figure 3D shows maximum function complexity evolved in host populations across function restriction configurations. Overall, we find that host function restriction dictates whether mutualism constrains host evolution. When hosts are required to specialize, mutualism constrains host evolution to the simplest function complexity in all replicates, regardless of symbiont specialization requirements. When hosts have no restrictions on the number of different functions that they can perform during their lifetime, hosts consistently evolve to be generalists and perform complex functions when coevolved with mutualists, regardless of whether their symbionts are required to be specialists or not. In these generalist host populations, hosts remain evolutionarily engaged in mutualism by maintaining the NOT function to enable function matching with symbionts.

Across all replicates of each of these experiment configurations, symbionts maintain evolutionary engagement with hosts via the NOT function, not by evolutionarily tracking hosts to more complex functions. It is worth emphasizing that symbionts could have evolved more complex functions, however they rarely did. This indicates that the symbionts' evolutionary potential may be the key to continued coevolutionary adaptation.

Conclusion

Our results demonstrate that trait-matching mutualistic endosymbioses can constrain host evolution. Across our experiments, we found no evidence of instances where mutualism promoted the evolution of complex features or increased host diversification. This result aligns with previous experimental results with intraspecies cooperative mutualisms (Pauli et al., 2022). We additionally found that the mechanistic details of the mutualistic interaction strongly influence whether mutualism can constrain host evolution. Specifically, we found that the benefits of different functions involved in trait matching, the efficiency and strength of mutualistic interactions, and the capacity for hosts to evolve generalist strategies strongly impacted whether mutualism constrained host evolution (Figure 3).

In most experiments, we observed evolutionary degradation in mutualistic endosymbionts: these symbionts rapidly lost the ability to reproduce on their own, relying entirely on vertical transmission into host offspring. This observation demonstrates a simple form of egalitarian major transition in individuality, and few endosymbiont populations (across experiments) regained the ability to reproduce independently. The benefits of mutualism allowed hosts with obligate mutualists to sweep the population, fixing host-symbiont pairings with symbionts that could no longer reproduce on their own. In populations where hosts evolved away from their symbiont partner by evolving to perform a new function, symbionts failed to evolve in lockstep with their hosts to maintain the symbiotic relationship. Instead, these endosymbionts became vestigial organelles inside their host: without the ability to independently reproduce and having no effect on their host's fitness, these endosymbionts evolved via genetic drift. Of course, if we added additional costs to hosting endosymbionts, we might expect the host population to purge evolutionarily disengaged mutualists.

Our results are specific to trait matching symbioses where interactions between individuals depend on the compatibility of relevant traits. Moreover, we focused on studying traits directly involved in interactions. In other circumstances, we might expect mutualism to have more neutral effects on host evolution or promote diversification and the evolution of complexity (as reviewed in Chomicki et al. 2019). For example, it is less clear what mutualism's net effect would be on the evolution of unrelated host traits; data from Figure 3D with unrestricted hosts might suggest that mutualism would be neutral, but future work is needed.

Overall, these results indicate that mutualistic symbiosis could lead to host populations that are unable to adapt to a changing environment if the traits under pressure are involved in a mutualistic interaction. Since most organisms are engaged in mutualistic symbioses of some kind, these results show the importance of identifying the mechanism of the mutualistic interaction to evaluate the risks to host populations as our planet continues to experience rapid change.

Acknowledgements

We thank members of the ECODE lab, CEVO lab, and SymbuLab for thoughtful discussions, ideas, and support. This work was supported by the National Science Foundation under Grants No. 2414125 and No. 2414126. 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. This work was also supported through computational resources and services provided by GVSU's Clipper system and the Institute for Cyber-Enabled Research at MSU.

References

Bennett, G. M. and Moran, N. A. (2015). Heritable symbiosis: The advantages and perils of an evolutionary rabbit hole. *Proceedings of the National Academy of Sciences*, 112(33):10169–10176.

Cavalier-Smith, T. (2013). Symbiogenesis: Mechanisms, Evolutionary Consequences, and Systematic Implications. *Annual Review of Ecology, Evolution, and Systematics*, 44(1):145–172.

Chen, S.-Y., Feng, Z., and Yi, X. (2017). A general introduction to adjustment for multiple comparisons. *Journal of Thoracic Disease*, 9(6):1725–1729.

Chomicki, G. and Renner, S. S. (2017). Partner abundance controls mutualism stability and the pace of morphological change over geologic time. *Proceedings of the National Academy of Sciences*, 114(15):3951–3956.

Chomicki, G., Weber, M., Antonelli, A., Bascompte, J., and Kiers, E. T. (2019). The Impact of Mutualisms on Species Richness. *Trends in Ecology & Evolution*, 34(8):698–711.

Davis, C. C., Schaefer, H., Xi, Z., Baum, D. A., Donoghue, M. J., and Harmon, L. J. (2014). Long-term morphological stasis maintained by a plant–pollinator mutualism. *Proceedings of the National Academy of Sciences*, 111(16):5914–5919.

Dehling, D. M., Töpfer, T., Schaefer, H. M., Jordano, P., Böhning-Gaese, K., and Schleuning, M. (2014). Functional relationships beyond species richness patterns: trait matching in plant–bird mutualisms across scales. *Global Ecology and Biogeography*, 23(10):1085–1093.

Dolson, E., Lalejini, A., Jorgensen, S., and Ofria, C. (2020). Interpreting the Tape of Life: Ancestry-Based Analyses Provide Insights and Intuition about Evolutionary Dynamics. *Artificial Life*, 26(1):58–79.

Dolson, E. and Ofria, C. (2017). Spatial resource heterogeneity creates local hotspots of evolutionary potential. In *Proceedings of the 14th European Conference on Artificial Life ECAL 2017*, pages 122–129, Lyon, France. MIT Press.

Dolson, E. L., Pérez, S. G., Olson, R. S., and Ofria, C. (2017). Spatial resource heterogeneity increases diversity and evolutionary potential. preprint, Ecology.

Erwin, D. H. (2015). Novelty and Innovation in the History of Life. *Current Biology*, 25(19):R930–R940.

Estrela, S., Kerr, B., and Morris, J. J. (2016). Transitions in individuality through symbiosis. *Current Opinion in Microbiology*, 31:191–198.

Gómez, J. M. and Verdú, M. (2012). Mutualism with Plants Drives Primate Diversification. *Systematic Biology*, 61(4):567–577.

Johnson, K., Welch, P., Dolson, E., and Vostinar, A. E. (2022). Endosymbiosis or Bust: Influence of Ectosymbiosis on Evolution of Obligate Endosymbiosis. In *The 2022 Conference on Artificial Life*, Online. MIT Press.

Joy, J. B. (2013). Symbiosis catalyses niche expansion and diversification. *Proceedings of the Royal Society B: Biological Sciences*, 280(1756):20122820.

Kaltenpoth, M., Flórez, L. V., Vigneron, A., Dirksen, P., and Engl, T. (2025). Origin and function of beneficial bacterial symbioses in insects. *Nature Reviews Microbiology*, 23(9):551–567.

Kelley, D., Hillary, O., Hillman-Emelianoff, E., Khalili, A., Zhang, E., Johnson, K. M., Dolson, E., Vostinar, A. E., and Lalejini, A. (2026a). Supplemental material (Archived GitHub Repository). 10.5281/zenodo.16795778. https://github.com/amlalejini/2026-alife-mutualism-constrains-evolution.

Kelley, D., Hillary, O., Johnson, K. M., Dolson, E., Vostinar, A. E., and Lalejini, A. (2026b). Data - mutualism can constrain host diversity and evolution. 10.17605/OSF.IO/J3EZD. https://osf.io/j3ezd.

Koonin, E. V. (2015). Origin of eukaryotes from within archaea, archaeal eukaryome and bursts of gene gain: eukaryogenesis just made easier? *Philosophical Transactions of the Royal Society B: Biological Sciences*, 370(1678):20140333.

Kruskal, W. H. and Wallis, W. A. (1952). Use of Ranks in One-Criterion Variance Analysis. *Journal of the American Statistical Association*, 47(260):583–621.

Lalejini, A. and Ofria, C. (2018). Evolving event-driven programs with SignalGP. In *Proceedings of the Genetic and Evolutionary Computation Conference on - GECCO '18*, pages 1135–1142, Kyoto, Japan. ACM Press.

Lengyel, S., Gove, A. D., Latimer, A. M., Majer, J. D., and Dunn, R. R. (2009). Ants Sow the Seeds of Global Diversification in Flowering Plants. *PLoS ONE*, 4(5):e5480.

Lenski, R. E., Ofria, C., Pennock, R. T., and Adami, C. (2003). The evolutionary origin of complex features. *Nature*, 423(6936):139–144.

Moran, N. A. (2007). Symbiosis as an adaptive process and source of phenotypic complexity. *Proceedings of the National Academy of Sciences*, 104(suppl_1):8627–8633.

Morris, J. J., Lenski, R. E., and Zinser, E. R. (2012). The Black Queen Hypothesis: Evolution of Dependencies through Adaptive Gene Loss. *mBio*, 3(2):e00036–12.

Ofria, C. and Wilke, C. O. (2004). Avida: A Software Platform for Research in Computational Evolutionary Biology. *Artificial Life*, 10(2):191–229.

Pauli, B., Oña, L., Hermann, M., and Kost, C. (2022). Obligate mutualistic cooperation limits evolvability. *Nature Communications*, 13(1):337.

Raimundo, R. L. G., Gibert, J. P., Hembry, D. H., and Guimarães, P. R. (2014). Conflicting Selection in the Course of Adaptive Diversification: The Interplay between Mutualism and Intraspecific Competition. *The American Naturalist*, 183(3):363–375.

Salem, H., Bauer, E., Kirsch, R., Berasategui, A., Cripps, M., Weiss, B., Koga, R., Fukumori, K., Vogel, H., Fukatsu, T., and Kaltenpoth, M. (2017). Drastic Genome Reduction in an Herbivore's Pectinolytic Symbiont. *Cell*, 171(7):1520–1531.e13.

Svardal, H., Rueffler, C., and Doebeli, M. (2014). Organismal complexity and the potential for evolutionary diversification. *Evolution*, 68(11):3248–3259.

Vostinar, A., Welch, P., K-Johnson-Horrigan, Cameron, A., Hafner, L. O., Dolson, E., dorchens, Gormely, C., renztim, Ofria, C., Qin, Y., and mkmmnn (2022). anyaevostinar/symbulationemp: Modes.

Vostinar, A. E., Skocelas, K. G., Lalejini, A., and Zaman, L. (2021). Symbiosis in Digital Evolution: Past, Present, and Future. *Frontiers in Ecology and Evolution*, 9:739047.

Watson, R. A., Reil, T., and Pollack, J. B. (2000). Mutualism, Parasitism, and Evolutionary Adaptation. In *Proceedings of Artificial Life VII*, pages 170–178.

Whitehead, M. R. and Peakall, R. (2014). Pollinator specificity drives strong prepollination reproductive isolation in sympatric sexually deceptive orchids: reproductive isolation in sympatric orchids. *Evolution*, 68(6):1561–1575.

Wilcoxon, F. (1992). Individual Comparisons by Ranking Methods. In Kotz, S. and Johnson, N. L., editors, *Breakthroughs in Statistics*, pages 196–202. Springer New York, New York, NY. Series Title: Springer Series in Statistics.

Yoder, J. B. and Nuismer, S. L. (2010). When Does Coevolution Promote Diversification? *The American Naturalist*, 176(6):802–817.

Zaman, L., Devangam, S., and Ofria, C. (2011). Rapid host-parasite coevolution drives the production and maintenance of diversity in digital organisms. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation - GECCO '11*, page 219, Dublin, Ireland. ACM Press.

Zaman, L., Meyer, J. R., Devangam, S., Bryson, D. M., Lenski, R. E., and Ofria, C. (2014). Coevolution Drives the Emergence of Complex Traits and Promotes Evolvability. *PLoS Biology*, 12(12):e1002023.