



Robots as models of evolving systems

Gao Wang^{a,1}, Trung V. Phan^{b,1}, Shengkai Li^c, Jing Wang^{d,e}, Yan Peng^f, Guo Chen^g, Junle Qu^h, Daniel I. Goldman^c, Simon A. Levin^h, Kenneth Pientaⁱ, Sarah Amendⁱ, Robert H. Austin^{j,2}, and Liyu Liu^{a,2}

Contributed by Robert H. Austin; received November 4, 2021; accepted February 2, 2022; reviewed by Robert Gatenby, Cynthia Reichhardt, and Yuhai Tu

Experimental robobiological physics can bring insights into biological evolution. We present a development of hybrid analog/digital autonomous robots with mutable diploid dominant/recessive 6-byte genomes. The robots are capable of death, rebirth, and breeding. We map the quasi-steady-state surviving local density of the robots onto a multidimensional abstract “survival landscape.” We show that robot death in complex, self-adaptive stress landscapes proceeds by a general lowering of the robotic genetic diversity, and that stochastically changing landscapes are the most difficult to survive.

robotic biology | evolution | adaptable landscapes | stochastic dynamics

Robotics has now reached the stage where artificial life and evolution can be studied in a physical system where robots can be used to ask complex evolutionary questions (1, 2). Further, there have been pioneering efforts to introduce genetics and biology into robotic systems.

The hybrid analog/digital robot community we have created has a collective, evolving behavior and, by design, has a generality which spans analogies from bacterial chemotaxis (3) to many-body soft-matter physics (4), evolutionary biology (5), and multicell collective states (6, 7). We suggest that stochastic time and multiple-chemical chemotherapy rather than a periodic time and monodrug course would target the high mutation rates of cancer cells (8) and be less punishing to cells that have the normal very low mutation rates that natural selection favors in the absence of high stress.

The resource landscapes are three overlapping red–green–blue (RGB) intensity maps generated on a 4.16-m × 4.16-m light-emitting diode light board. The time- and space-varying RGB intensities generated on the board represent three resource landscapes; a given resource landscape is associated with a given color. Although the robots locally deplete resources, the local resource shadow generated by the presence of a robot gradually fades once the robot moves away.

The robots move only in response to local resource gradient created at their position on the light board, which, together with the resource dynamics, creates a complex field drive locomotion mechanism (3). The many-body aspect of the resource landscape gradients is extremely dynamic, since the robots locally deplete the resources; the result is that the gradients are a strong function of both external drivers and the local density of the robots (9). Fig. 1 presents the basic flow of the robot swarm on the landscape.

One of the fundamental innovations which distinguishes our robots from more conventional robots is the closed loop between the mutable genes which control color response phenotype and the four downward looking RGB sensors which determine local color resources and physical response on an interactive dynamical landscape. The sensors are at opposing quadrants on the base of the robot and detect the RGB colors and corresponding intensities from the RGB light board.

An overhead camera determines resource depletion. The presence of a robot on the light board is detected by the camera; this results in a signal sent to the light board controller to progressively decrease the color intensity in an area around the robot. This process has consequences: It provides motility for the robots as they move away from their own developing shadow. The consumption of resources by the robots forms the collective shadow $[S]$, which follows a relaxation dynamics,

$$\partial_t [S] = -\frac{1}{\tau} [S] + \sum_{\text{robots}} C, \quad [1]$$

where C is the consumption rate of each robot centered around its position, and τ is the fixed recovery rate of a resource. In our experiments, τ is a fixed value. The resource value is given by $I = I_{BG} - [S]$ so that, when a robot consumes resources and leaves, the environment will slowly recover to the set background I_{BG} . Details of this can be found in *SI Appendix, sections 1–4*.

Significance

We present a fully realized adaptive resource landscape with diploid three-gene robots presenting interacting roles of population dynamics, mutations, breeding, death, and birth. Although modeling and theory serves as a guide here, the inherent complexity of our robobiology world makes it an experiment in exploring rules of Darwinian natural selection at a level difficult to simulate. We find that the lower the genetic diversity, the lower the survival probability of the robot population. We propose that diploid gene robots can act as avatars of diploid mammalian cells to explore novel programs of administration of drugs.

Author contributions: G.W., T.V.P., J.Q., S.A.L., K.P., S.A., R.H.A., and L.L. designed research; G.W., T.V.P., J.W., Y.P., G.C., and L.L. performed research; G.W., T.V.P., S.L., Y.P., D.I.G., S.A.L., K.P., S.A., and L.L. contributed new reagents/analytic tools; G.W., T.V.P., S.L., J.W., Y.P., G.C., D.I.G., R.H.A., and L.L. analyzed data; and G.W., T.V.P., S.L., D.I.G., K.P., S.A., and R.H.A. wrote the paper.

Reviewers: R.G., Moffitt Cancer Center; C.R., Los Alamos National Laboratory; and Y.T., IBM.

Competing interest statement: K.P. and S.A. were coauthors on a paper with referee R.G.: <https://pubmed.ncbi.nlm.nih.gov/33897809>.

Copyright © 2022 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

See [online](#) for related content such as Commentaries.

¹G.W. and T.V.P. contributed equally to this work.

²To whom correspondence may be addressed. Email: austin@princeton.edu or lyliu@cqu.edu.cn.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2120019119/-DCSupplemental>.

Published March 17 2022.

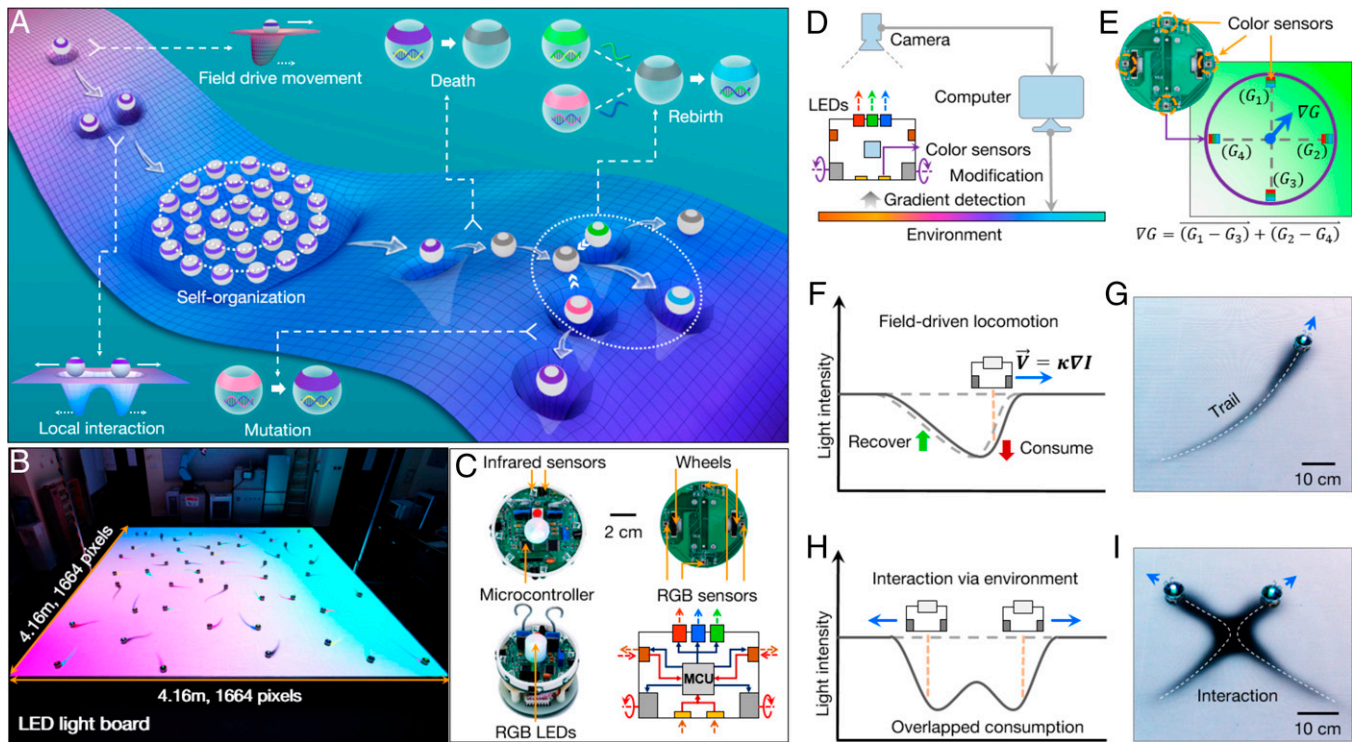


Fig. 1. (A) Depiction of the complex interactions of the robots. Ribbon: The robots move via self-generated field drive on a resource landscape generated by an underlying LED light board, which the robots sense. Field drive is the self-generated movement of the robots in response to local resource depletion. (B) The 4.16-m by 4.16-m LED light board with robots across an RGB nested set of landscapes. (C) The hardware of a robot. (Upper Left) Infrared LEDs and sensors are used for gene exchange. (Upper Right) Two wheels driven by independent motors move the robot over the LED light board, while four RGB downward-looking sensors measure local light board intensities. (Lower Left) RGB LEDs display the state of the 6-byte genome to an overhead camera. (Lower Right) Basic information flow: The downward-looking sensors control wheel movement, the infrared LEDs/sensors control gene exchange, and the upward-transmitting RGB LEDs send the genome state to the overhead camera. (D) The information logic flow for a given robot to determine resource consumption. The upward-transmitting RGB LEDs are read by an overhead camera, which then sends information to a computer which dims the appropriate resource color(s) due to the dominant gene(s), and controls resource recovery. (E) Hardwired computation of resource gradients by a robot. (F) The basic hardware in which a robot moves in response to resource gradients (field drive). (G) The recovering resource “shadow” generated by a moving robot on a white homogeneous landscape due to field drive. (H) The basic repulsive interaction due to field drive between two robots. (I) Soft resource-driven “collision” between two field drive robots on a white landscape.

The robot’s phenotype is determined by a 6-byte diploid genome: (R1, G1, B1) and (R2, G2, B2). Each byte (gene) consists of 8 bits: 1 bit is reserved to flag if it is dominant (1:active) or recessive (0:nonresponsive), as is shown in Fig. 2. Only dominant genes can exhibit sensitivity to the resource gradients.

The remaining 7 bits (0000000 to 1111111) determine the sensitivity to the color associated with that gene and hence the velocity (a two-dimensional vector) with which they climb out of a resource hole following resource gradients. The quantitative relationship between genes and robot velocities is as follows: Three pairs of genomes respectively control robot sensitivity to three different resources (R/G/B), which are represented by R_{gene} , G_{gene} , and B_{gene} ,

$$\begin{aligned}
 R_{gene} &= \frac{R_1^{(flag)} R_1^{(value)} + R_2^{(flag)} R_2^{(value)}}{2 \times 127}, \\
 G_{gene} &= \frac{G_1^{(flag)} G_1^{(value)} + G_2^{(flag)} G_2^{(value)}}{2 \times 127}, \\
 B_{gene} &= \frac{B_1^{(flag)} B_1^{(value)} + B_2^{(flag)} B_2^{(value)}}{2 \times 127}.
 \end{aligned}
 \quad [2]$$

Here the quantities with superscript “(flag)” take value zero for recessive gene and value one for dominant gene, and the quantities with superscript “(value)” take the value of the remaining 7 bits, ranging from 0 to 127 in decimal.

The vector velocity is determined then by the genotype and localized resource gradient,

$$\begin{aligned}
 \vec{V} &= \frac{V_{max}}{3} \\
 &\times (R_{gene} \nabla I_{red} + G_{gene} \nabla I_{green} + B_{gene} \nabla I_{blue}).
 \end{aligned}
 \quad [3]$$

We discovered an unexpected pleiotropy (one gene–multiple phenotypes) in our robot response. We found that pleiotropy in our genomes results from the finite spectral linewidths of the RGB LEDs in the light board which result in the RGB detectors in the robot base responding not only to the primary color (for example, the red detector seeing not only red but also blue and green). *SI Appendix, section 6* gives details of the robot construction and maps out this cross-channel spectral bleeding.

There is no predetermined software algorithm to our robot phenotype (10). Rather, the phenotype is collective, emergent, and hardware driven, yet can be characterized as partly selfish and partly a form of altruism.

The selfish aspect is to exploit (consume) resources where you are, and move to find more resources as you deplete the ones you found, by following the positive gradients of resources created by the robot and other nearby robots (3). Robots which cannot escape from a resource hole die and lose their genes.

The altruism is as follows: If you find a dead robot, copy 1/2 your genome to the dead robot; two successive donations from nonidentical robots gives rise to a rebirthed daughter robot. The altruism here is the survival of the robot species. Since our robots cannot, unfortunately, reconstruct themselves from basic materials

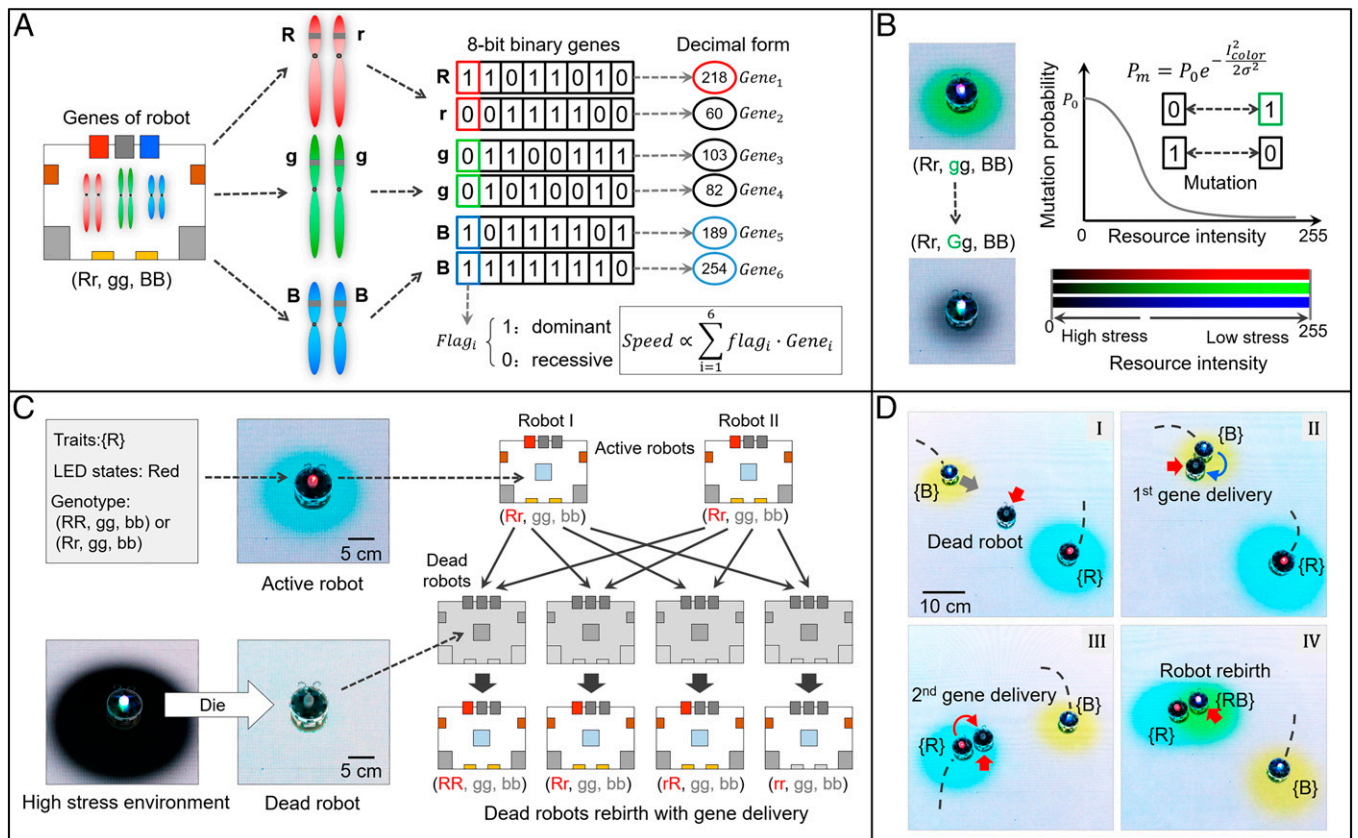


Fig. 2. (A) Genomes of the robots. (Left) Cartoon of the 6-byte genome and analogy to diploid human chromosomes. (Right) The 6-byte genome. The first bit determines dominance or recessive nature, while the remaining 7 bits determine sensitivity to the underlying color on the LED light board. The translational speed of the robot is proportional to the 7-bit number times the state of the dominance/recessive flag. (B) (Left) Green resources untouched by a recessive green gene are consumed once a mutation makes the green gene dominant, giving rise to a black resource hole. (Right) Mutation rates increase with lower levels of resources. (C) Pictorial gene exchange diagram as a dead robot (left) meets an active robot and haploid gene delivery occurs. (D) Images of the rebirth of a dead robot (center) due to two subsequent haploid gene deliveries.

(11), rebirth is our way to prevent inevitable extinction, as Fig. 2 shows. The rebirthed robot is thus genetically related to its parents. It is absolutely possible for us to design different reproduction modes, including asexual rebirth (one cell simply gives genes to a dead one), or just exchange genes between two alive robots, or any combination of these.

Given the plethora of choices, and since we are most interested in fundamental evolution dynamics, we chose the normal mating scheme of (parent-1) + (parent-2) = parents + progeny. The evolution dynamics of asexual reboots of dead robots by a single “parent” is interesting and will be explored in a later paper. Detailed descriptions for robot–resource environment interactions can be found in *SI Appendix, section 4*.

We correlate resource stress with mutation rates (12): More resource stress gives rise to more mutations/time. Thus the stability of a gene in the robot is connected with the resource level: The dimmer the color intensity, the higher the mutation rate of the gene which is associated with that color. This allows a robot to escape, in principle, from a low-resource dim region even if the 7-bit binary number associated with that color is low, and hence there is a low sensitivity to light gradients: The enhanced mutation rate and the already low value of the gene implies a greater probability of drawing a high value and escaping. Likewise, a robot in a bright area will have a low mutation rate; in effect, a robot with low sensitivity (low 7-bit binary number) is rewarded by being more likely to stay in that region until the inevitable depletion of resources happens. Then, there is always the risk of death.

The effective “metabolic” rate of resource consumption by the robots for a given color is fixed to a level proportional to the local

intensity of that color but is independent of the sensitivity of the robot to that color in terms of field drive. In principle, only a robot with all “pure” recessive genes (rr, gg, bb) cannot consume any resources. It will stay still and then either mutate to a dominant gene or die after 5 s. This is the reason why a robot with all recessive genes will die even in a white environment. It has no sensitivity to any resource, so it simply sits in one place, fatally. Thus a robot with more dominant genes will consume more resources, and it will move faster than one with fewer dominant genes. A robot with only recessive genes will consume nothing, remain stationary, and die.

As we noted, there is color leakage (pleiotropy) in this system: The RGB LEDs are not pure monochromatic light sources, so that, for example, a blue sensor on the robot will detect some fraction of the green light board even if no blue light is present; the net result is that robots do not necessarily die in an inappropriate environment. The biological equivalent for this is that recessive genes are not necessarily entirely recessive (13, 14).

Of course, if all the robots die, the population becomes extinct. Hence, breeding to produce an alive robot is essential. The genes of a reborn robot must come from two different alive robots (every robot has a unique ID). Fig. 3 summarizes robot death, rebirth, and random genome dynamics, also showing the power of our ability to obtain detailed genomic dynamics events. However, we cannot easily show, graphically, the underlying stress landscape drivers and the motions of the robot over the landscape for this segment. Details for robot–resource environment interactions and genomic dynamics can be found in *SI Appendix, sections 4–6*.

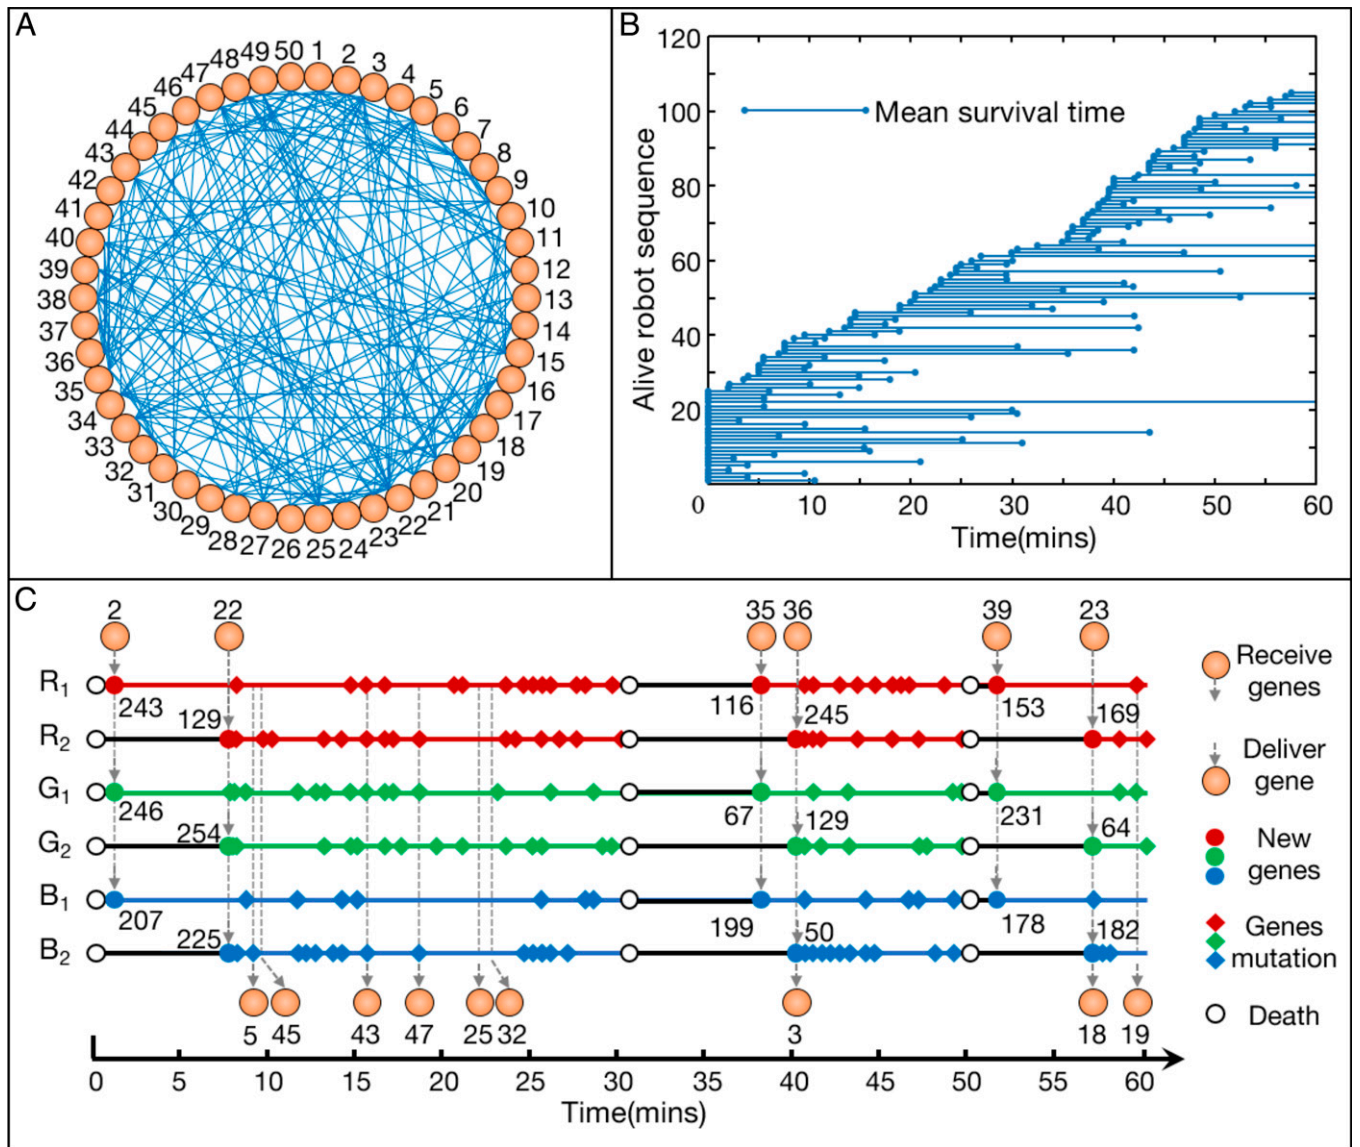


Fig. 3. (A) Interchange of genomes due to death and rebirth between $N_{tot} = 50$ robots in the course of an experiment. A blue line indicates gene exchange has occurred. (B) The complex and unpredictable nature of how long a given robot stays alive even in a uniform, white landscape as a function of time. (C) Detailed local snapshot of a single robot's genome dynamics starting with a death event at $t = 0$ min. The robot becomes alive after two separate haploid donations by robot 2 at $t = 2$ min and robot 22 at $t = 7$ min. The next death events for this robot occur at $t = 31$ min and $t = 50$ min. Note that the robot does nine haploid donations and many single-bit mutations over the course of this time sequence. No mutations can occur while a robot is dead.

A Theoretical Guide to the Experiments

A theoretical model is used to provide a rough guide of where, in the vast parameter space, we should set our robot codes to probe for interesting results.

Five critical parameters that control the survival of our bio-inspired system of robots are 1) the rate of mutation r_m per gene bit, 2) the rate of gene exchange to create offspring at rate r_b which is the only way to increase the number of alive robots in our system, 3) the basal rate at which the robots consume resources, 4) the recovery rate τ of resources in the absence of a robots, and 5) t_d , the time it takes to die when a genetic configuration stays “bad.” We fix t_d , since it is basically intrinsic, defined dimensionless parameters $c_m = r_m t_d$ and $c_b = r_b t_d$, and examined how survival changes with c_m and c_b .

As we discuss in *SI Appendix, section 7*, the metric of survival $\mathcal{S} = N_{st}/N_c(F_{max})$ —where N_{st} is the stationary number of alive robots and $N_c(F_{max})$ is the maximum carrying capacity of robots given the maximum amount of resources F_{max} in the

environment—represents the serviceability of the robot population (15). For an environment with abundant resources, $N_c = N_{tot}$ in which N_{tot} is the total number of robots available. In a fixed white environment, when the space of “good” genetic variations has a much larger volume than that of “bad” genetic variation, the survival is approximately

$$\mathcal{S} \approx 1 - \beta_{g \rightarrow b} e^{-\beta_{b \rightarrow g} c_m \left[\frac{c_m}{c_b} \right]}, \quad [4]$$

where $\beta_{g \rightarrow b}$ and $\beta_{b \rightarrow g}$ are constants. Note that \mathcal{S} cannot be a negative value; thus, when Eq. 4 gives an unphysical result $\mathcal{S} < 0$, then the physical value should be set to $\mathcal{S} = 0$. The survival value \mathcal{S} is minimum when c_{mutate} is equal to the critical value $c_m^{cr} = 1/\beta_{b \rightarrow g}$; therefore, we call the region around the line $c_m = c_m^{cr}$ the mutation meltdown valley. On the two sides, we call the region on the $c_m \ll c_m^{cr}$ side the unchanged hill and call the region on the $c_m \gg c_m^{cr}$ side the undead hill. For the region of parameter space that leads to $\mathcal{S} = 0$ (which means the population goes extinct), we call this the extinction swamp. In

Fig. 4, we qualitatively show the survival S surface (which we call the success landscape) as a function of the dimensionless mutation coefficient c_m and the dimensionless breeding coefficient c_b . The experimental results are given in *SI Appendix, section 7*.

Thus a simplified theoretical analysis shows three possible distinct stationary states for the robot community: surviving with death stabilized (where the resource is abundant, and the stationary density of alive robots are only limited by how fast rebirth can happen to balance out spontaneous deaths due to stochastic resource fluctuations), surviving with resource depletion (where the stationary density of alive robots is directly bounded by the amount of resources), and total extinction (where the population will eventually die, due to either the death rate being too high, the rebirth rate being too low, or the resource recovery being too slow). Note that this is just a guide for a spatially homogeneous external environment with no externally driven time dependence. There are actually many surprises in this robot world.

Experimental Robot Evolutionary Survival Landscapes

We have carried out three different landscape evolution experiments to test to what extent our robot community is able to evolve and adapt to externally driven stresses.

Experiment 1 involves no external time dynamics to resources, and a white landscape with equal resources of red green and blue landscapes. This is the simplest possible landscape, with no explicit time-dependent external drive to the ecology, like Earth's

equatorial climate. However, with time, the landscape will become spatially and temporally complex due to the heterogeneous nature of the robot genotypes, their motile resource exploitation, and their interactions, which result in locally preferential depletion of different resources. *Movie S1* shows robot dynamics on an initial flat, white landscape versus time.

Fig. 5A shows the evolution of a survival landscape as a function of time for a fixed (white) landscape. As the theoretical guide predicted, there exists a mutational rate which leads to minimal survival and near-extinction of the robot community as time progresses. Although some level of mutations is necessary for a population to evolve, presumably, living systems would be driven to minimize mutation rates to the extent it is possible to avoid the mutational risk of extinction in a low-stress environment that occurs with increasing mutations (16).

Experiment 2 involves a spatially uniform but periodically changing landscape. Here we drive the resource landscape with monotonically changing colors, representing, in effect, a well-mixed imposed landscape chemostat evolution experiment (17), but different from the classic chemostat experiment in that the initially homogeneous landscape can be locally modified by the robots. A fixed diameter circle of light homogeneously changes color from red to green to blue. The resource circle changes color with a 4,320-time step period for each color change; see Fig. 5B. *Movie S2* shows mutation and breeding on a periodically changing color landscape.

As we have noted, it is only when robots have both mutations and gene exchange that a robust robot population can exist within




Fig. 4. The theoretical surviving fraction of robots (S , which we call success) landscape as a surface function of the dimensionless mutation coefficient c_m and the dimensionless breeding coefficient c_b for a fixed white resource landscape. The critical line corresponds to $c_m = \alpha_{b \rightarrow g}$. There are four main regions: the mutation meltdown valley, the extinction swamp, the unchanged hill, and the undead hill. Our theory predicts which positions, indicated by four triangle data points (where $r_m \propto P_0$, and r_b depends monotonically on the number of available breeding ports), will have corresponding success fractions of robots in the experiments.

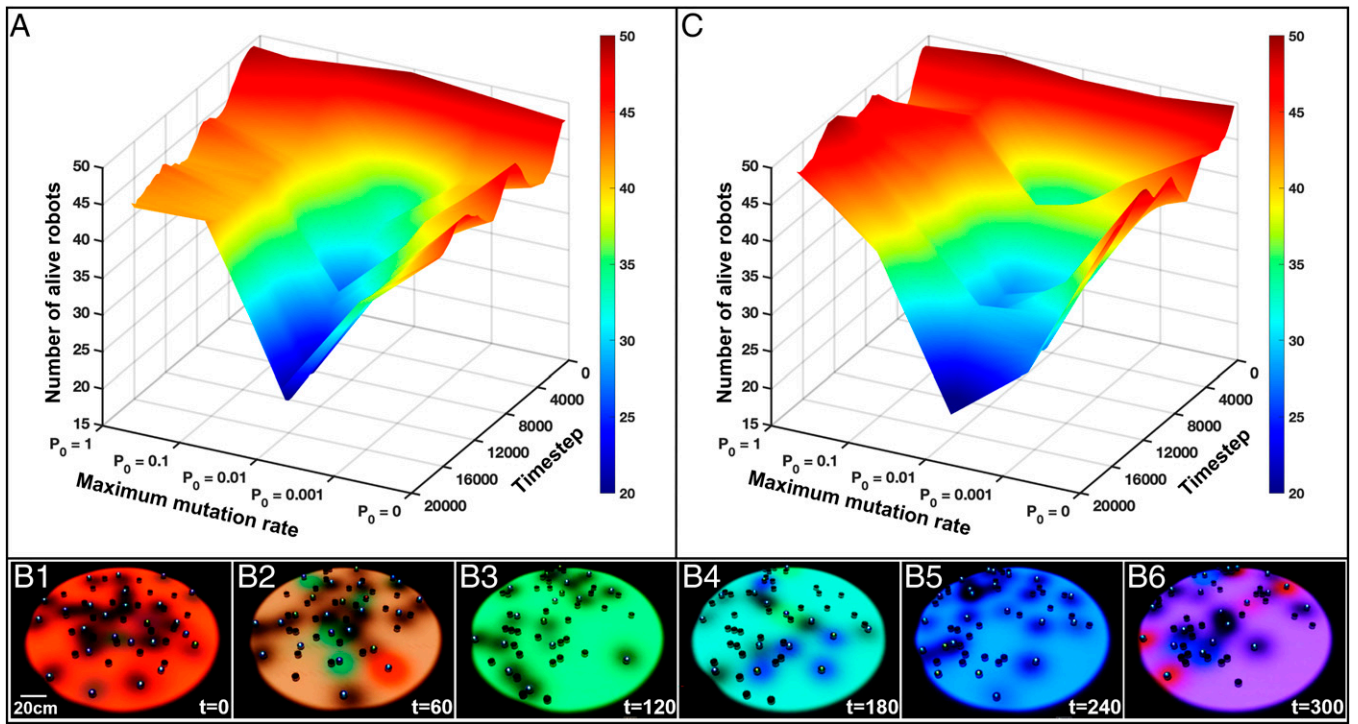


Fig. 5. (A) Surviving robot number versus time for a homogeneous white landscape which has no external time changes. (B) Images of robot positions for a periodically changing resource landscape: RGB. (C) Surviving robot number versus time and mutation rate for a periodically changing environment.

a simple harmonically changing well-mixed landscape. The robot's adaption to dynamic environment is revealed by the phenotypes responding to the detail RGB ingredients of the resource landscape. However, since the environment is changing with time, stress now opens up the mutational meltdown narrow valley and widens it.

Experiment 3 involves stochastic spatial and time-dependent landscapes. To explore truly complex robot resource dynamics,

unlike a typical biological laboratory well-mixed scenario, we created a stochastic, spatially fragmented complex landscape. The details of the algorithm used to create this landscape can be found in *SI Appendix, section 8*. *Movie S3* shows robots mutating and breeding on a stochastic landscape, and *Movie S4* shows the comprehensive introduction of the robot and experiments.

Fig. 6 shows the number of surviving robots versus time on a stochastic, fragmented landscape. Clearly, as the stress imposed

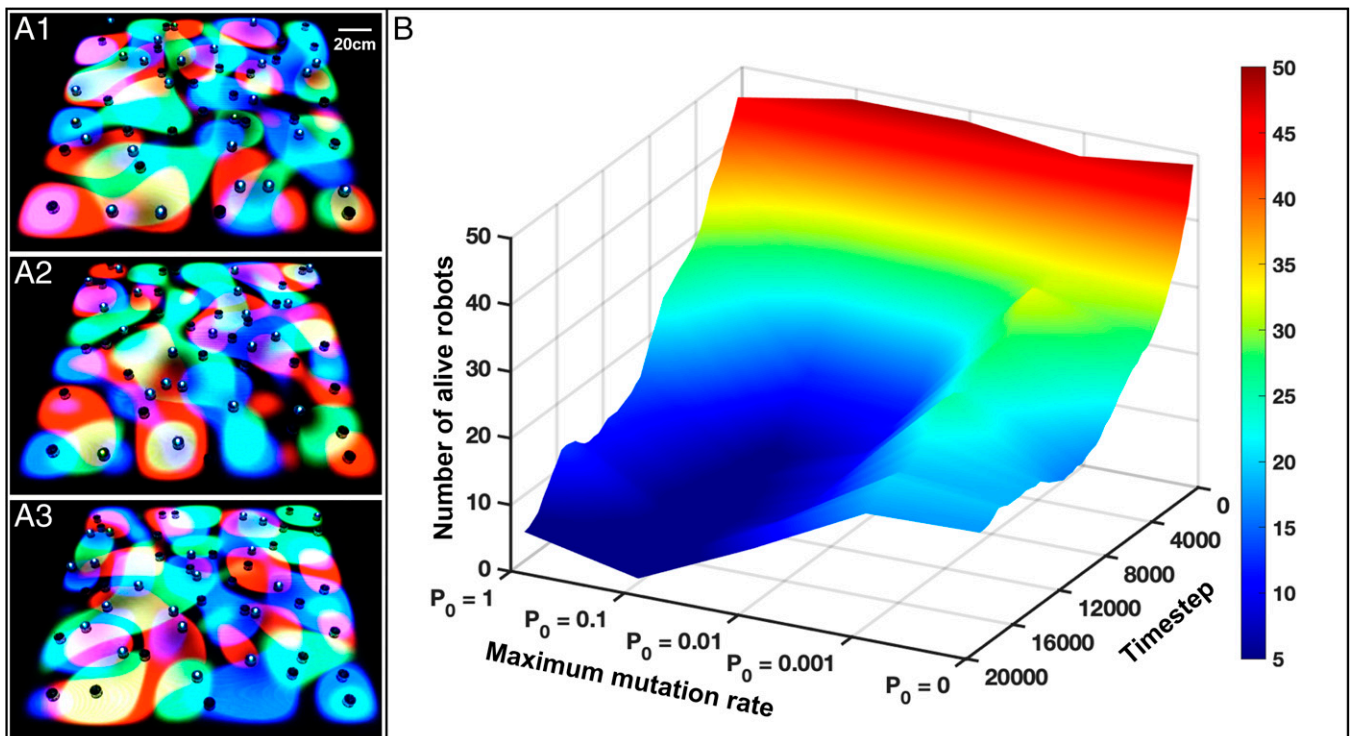


Fig. 6. (A) Images of robot positions for a stochastic landscape which changes pattern randomly with time. (B) Surviving robot number versus time and mutation rate for a stochastic landscape with a stochastic time dependence.

on the robots gets increasingly harder to predict versus time, it becomes increasingly harder for the robots to evolve. The valley of death increases dramatically in width as a function of mutation rate.

As we show in Fig. 6, for the case of pleiotropic genes, $P_0 = 0$ is actually the sole optimum survival point in a stochastic environment with the parameters chosen. Pleiotropy inhibits chasing moving resources which vanish, but it can also result in a fragmented landscape in leaving robots isolated for too long a time and failing to undergo gene exchange, which is critical for survival.

In our toy robot organisms, the three genes are all equivalent in terms of phenotype: They just respond basically to different colors. Yet the system evolves, so what is different in surviving populations that makes them successful, since the phenotypes cannot evolve? We assumed that the evolutionary diversity must be in the relative numbers of bits set high within the genes. A useful metric for this kind of genetic diversity is the Shannon entropy of the dominant genes S (18, 19),

$$S(t) = - \sum_i^{phenotype} p_i(t) \ln[p_i(t)], \quad [5]$$

where $p_i(t)$ is the time-dependent probability of finding one of the $2^3 - 1 = 7$ possible living phenotypes (red, green, blue, red + green = yellow, green + blue = cyan, red + blue = magenta, and red + green + blue = white) as the landscape changes. If all the dominant genes in a collection of N robots are exclusively “green” eventually, for example, then the final Shannon entropy is zero, and there is no dominant gene diversity on the resource landscape. The time dependence of the Shannon entropy is closely related to the Fisher information, since the Fisher information is basically the variance of the Shannon entropy (18). We fit the measured time dependence of $S(t)$ to a logistic equation in Fig. 7A. We choose the logistic equation primarily because of its close connection to ecological population dynamics, namely, the role of the population growth rate R and carrying capacity K of an ecology (20), possibly related to our robot genome rate of change and limiting genomic diversity.

Under, then, the assumption that, with time, the initial Shannon genetic entropy $S(0)$ will evolve to a limiting entropy $S(\infty)$, we determined, by curve fitting the Shannon genetic entropy change $\Delta S = S(\infty) - S(0)$,

$$S(t) = \frac{S(0) \times S(\infty)}{S(0) + \Delta S \times e^{-t/\tau}}, \quad [6]$$

where τ is the relaxation time of $S(t)$. Fig. 7 presents the results primarily for the stochastic system. Note that, as would be expected, in the absence of mutations, the Shannon entropy does not change.

Our evolution analysis indicates two important and surprising points connected to mutation rates: 1) The lower the final Shannon entropy, that is, the lower the genetic diversity, the lower is the survival probability of the robot population. 2) Static or periodically changing environments basically see no net change in genomic diversity with time, even with changing intrinsic mutation rates P_0 , but a stochastically changing environment drives genetic diversity down with time.

Discussion

Our results suggest the basic hypothesis that, while it is key for robots to mutate, exchange genes, and breed to avoid extinction, high mutation rates can be an extinction driver in a sufficiently

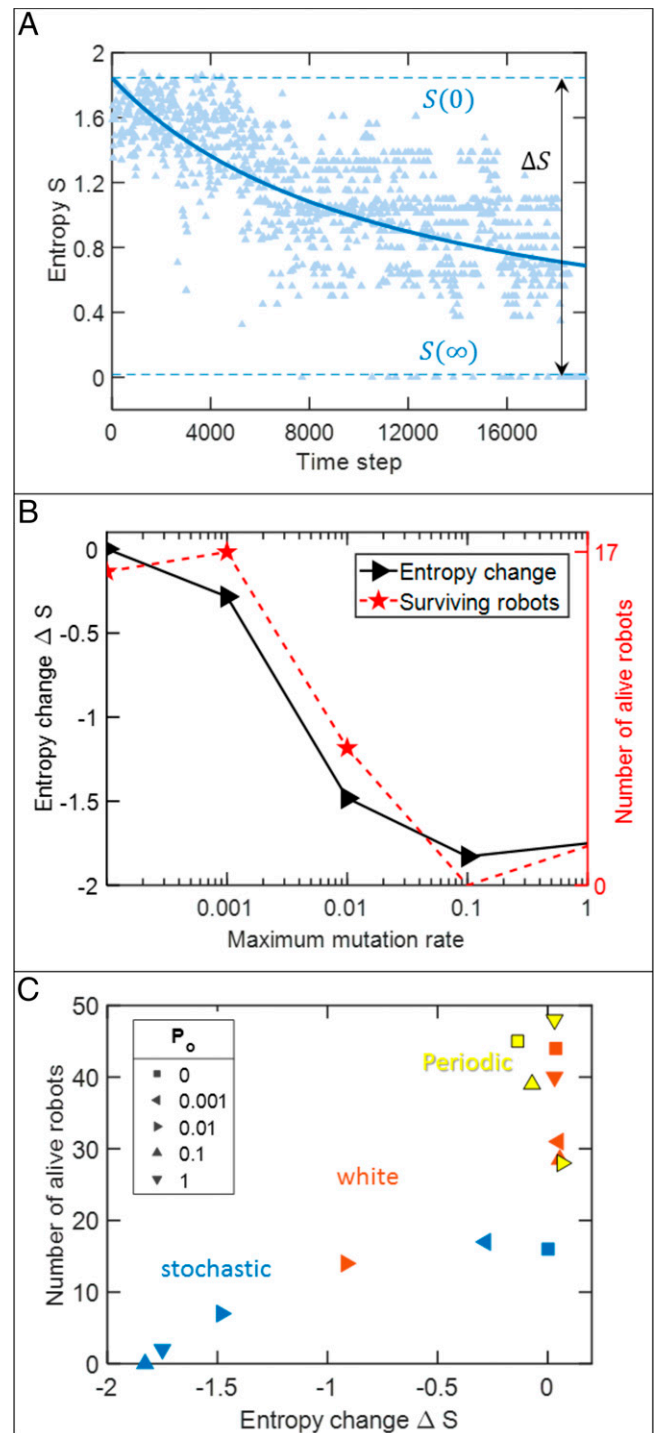


Fig. 7. (A) Example of the logistic curve fit of the Shannon genetic entropy $S(t)$ versus time for $P_0 = 0.1$. (B) Surviving robot population (red) and logistic curve fit of the total Shannon entropy change (black) ΔS as a function of basal mutation rate P_0 for a stochastic landscape. (C) Correlation between Shannon entropy change $S(t)$ and robot survival for three different landscape dynamics, as a function of different intrinsic mutation rate P_0 .

stochastic ecology. The number of robots is large ($N_{tot} = 50$) but not infinite; finite rather than infinite numbers of agents in biology is a meaningful and extremely important constraint in biology (21), especially in fragmented landscapes.

We have attempted to explore complicated bioinspired resource dynamics with real physical robots instead of a more conventional digital approach (agent-based computer simulations), such as was used in the pioneering Sugarscape simulation (22, 23). There are

fundamental differences between these two methodologies, such as the overwhelming combinatoric load (at what point does the number of robots and their interactions become impossible to simulate?) (24), the distinctions between smooth analog time and discrete time in the flow of differential equations (25), and the general breakdown of algorithms in the presence of noise (26). A concrete example is chronological ambiguity in digital operations, because the combination of rules to represent simultaneous processes in discrete time is intrinsically problematic: There usually is a specific order (nonparallel) to how the update of variable values is done. This results in the unsolved replication problem encountered in agent-based simulations, where observations cannot always be faithfully recreated based on just the physical analog description alone (27, 28).

From the success landscape (Fig. 4), we conjecture that absorbing phase transitions (29) play an important role in the emergent behavior of this robotic system. In the extinction swamp, the system has crossed an absorbing phase transition boundary, where the mutation-driven fluctuating state of some “alive” robots is no longer stable, and get absorbed into the state of “dead” robots. Absorbing phase transitions are usually associated with strongly non-Gaussian behavior and diverging time scales (30–32), and there is, of course, room for more exploration here beyond the scope of this paper.

There is a potential clinical aspect to this work based upon our surprising results from stochastic landscapes. From the start of the design of this technology, we viewed our robots as cancer cells, and that the resource landscape over which they move represents changing nutrients, and that chemotherapy is represented by

changing the resource landscape in such a way that the cells, although they are capable of mutation and reproduction, cannot sustain a viable population.

From our results, we predict that a stochastic time and multiple-chemical chemotherapy rather than a periodic time and monodrug course would target the high mutation rates of cancer cells (8) and be less punishing to cells that have the normal very low mutation rates that natural selection favors in the absence of high stress. At present, this is based on a rather abstract stochastic resource landscape, as is clear from Fig. 6. Clearly, we will have to design our resource landscape to more closely resemble a solid tumor, with gradients in stress from an outside perimeter, and time-clocked flow of high stress chemotherapy at the perimeter.

Data Availability. There are no data underlying this work.

ACKNOWLEDGMENTS. This work was supported by The National Natural Science Foundation of China (Grants 11974066 and 12174041), and the US NSF Grant PHY-1659940. We thank Robert Axelrod, Joel Brown, Tuan K. Do, Emma Hammarlund, and Joshua Weitz for helpful comments.

Author affiliations: ^aChongqing Key Laboratory of Soft Condensed Matter Physics and Smart Materials, College of Physics, Chongqing University, Chongqing 401331, China; ^bDepartment of Molecular, Cellular and Developmental Biology, Yale University, New Haven, CT 06520; ^cSchool of Physics, Georgia Institute of Technology, Atlanta, GA 30332; ^dWenzhou Institute, University of Chinese Academy of Sciences, Wenzhou 325011, China; ^eOujiang Laboratory (Zhejiang Lab for Regenerative Medicine, Vision and Brain Health), Wenzhou 325001, China; ^fResearch Institute of USV Engineering, Shanghai University, Shanghai 200444, China; ^gCollege of Physics and Optoelectronic Engineering, Shenzhen University, Shenzhen 518060, China; ^hDepartment of Environmental and Evolutionary Biology, Princeton University, Princeton, NJ 08544; ⁱThe Brady Urological Institute, Johns Hopkins School of Medicine, Baltimore, MD 21287; and ^jDepartment of Physics, Princeton University, Princeton, NJ 08544

1. A. E. Eiben, J. Smith, From evolutionary computation to the evolution of things. *Nature* **521**, 476–482 (2015).
2. W. Aguilar, G. Santamaria-Bonfit, T. Froese, C. Gershenson, The past, present, and future of artificial life. *Front. Robot. AI* **1**, 1–15 (2014).
3. T. V. Phan, G. Wang, L. Liu, R. H. Austin, Bootstrapped motion of an agent on an adaptive resource landscape. *Symmetry (Basel)* **13**, 225 (2021).
4. J. Aguilar *et al.*, A review on locomotion robotics: The study of movement at the intersection of robotics, soft matter and dynamical systems. *Rep. Prog. Phys.* **79**, 110001 (2016).
5. T. Arita, M. Joachimczak, T. Ito, A. Asakura, R. Suzuki, Alife approach to eco-evo-devo using evolution of virtual creatures. *Artif. Life Robot.* **21**, 141–148 (2016).
6. K. N. McGuire, C. De Wagter, K. Tuyls, H. J. Kappen, G. C. H. E. de Croon, Minimal navigation solution for a swarm of tiny flying robots to explore an unknown environment. *Sci. Robot.* **4**, eaaw9710 (2019).
7. S. Nolfi, D. Floreano, *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines* (MIT Press, Cambridge, MA, 2004).
8. D. Hanahan, R. A. Weinberg, The hallmarks of cancer. *Cell* **100**, 57–70 (2000).
9. G. Wang *et al.*, Emergent field-driven robot swarm states. *Phys. Rev. Lett.* **126**, 108002 (2021).
10. L. J. Fogel, A. J. Owens, M. J. Walsh, *Artificial Intelligence through Simulated Evolution* (John Wiley, 1966).
11. A. Dneprov, *Crabs on the Island* (Mir, Moscow, Russia, 1968).
12. D. M. Fitzgerald, P. J. Hastings, S. M. Rosenberg, Stress-induced mutagenesis: Implications in cancer and drug resistance. *Annu. Rev. Cancer Biol.* **1**, 119–140 (2017).
13. D. Monies *et al.*, Autozygosity reveals recessive mutations and novel mechanisms in dominant genes: Implications in variant interpretation. *Genet. Med.* **19**, 1144–1150 (2017).
14. R. Levins, R. Lewontin, *The Dialectical Biologist* (Harvard University Press, Cambridge, MA, 1987).
15. T. V. Phan *et al.*, It doesn't always pay to be fit: Success landscapes. *J. Biol. Phys.* **47**, 387–400 (2021).
16. M. Lynch, The lower bound to the evolution of mutation rates. *Genome Biol. Evol.* **3**, 1107–1118 (2011).
17. L. Rindi, M. D. Bello, L. Dai, J. Gore, L. Benedetti-Cecchi, Direct observation of increasing recovery length before collapse of a marine benthic ecosystem. *Nat. Ecol. Evol.* **1**, 153 (2017).
18. S. A. Frank, Natural selection maximizes Fisher information. *J. Evol. Biol.* **22**, 231–244 (2009).
19. E. V. Koonin, The meaning of biological information. *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.* **374**, 20150065 (2016).
20. R. S. Cantrell, C. Cosner, On the effects of spatial heterogeneity on the persistence of interacting species. *J. Math. Biol.* **37**, 103–145 (1998).
21. S. Wright, *Evolution: Selected Papers* (University of Chicago Press, 1986).
22. J. M. Epstein, R. Axtell, *Growing Artificial Societies: Social Science from the Bottom Up* (Brookings Institution Press, 1996).
23. A. Bigbee, C. Cioffi-Revilla, S. Luke, “Replication of *sugarscape* using Mason” in *Agent-Based Approaches in Economic and Social Complex Systems IV*, T. Terano, H. Kita, H. Deguchi, K. Kijima, Eds. (Springer, 2007), pp. 183–190.
24. S. M. Reia, A. C. Amado, J. F. Fontanari, Agent-based models of collective intelligence. *Phys. Life Rev.* **31**, 320–331 (2019).
25. A. Litman, S. Moran-Schein, Smooth scheduling under variable rates or the analog-digital confinement game. *Theory Comput. Syst.* **45**, 325–354 (2009).
26. F. Antenucci, S. Franz, P. Urbani, L. Zdeborová, Glassy nature of the hard phase in inference problems. *Phys. Rev. X* **9**, 011020 (2019).
27. B. Edmonds, D. Hales, Replication, replication and replication: Some hard lessons from model alignment. *J. Artif. Soc. Soc. Simul.* **6**, 11 (2003).
28. C. Sansores, J. Pavón, “Agent-based simulation replication: A model driven architecture approach” in *MICAI 2005: Advances in Artificial Intelligence*, A. Glebukh, A. de Albornoz, H. Terashima-Marin, Eds. (Springer, 2005), pp. 244–253.
29. C. Reichardt, C. J. Olson Reichardt, Absorbing phase transitions and dynamic freezing in running active matter systems. *Soft Matter* **10**, 7502–7510 (2014).
30. H. Hinrichsen, Non-equilibrium critical phenomena and phase transitions into absorbing states. *Adv. Phys.* **49**, 815–958 (2000).
31. G. Ódor, Universality classes in nonequilibrium lattice systems. *Rev. Mod. Phys.* **76**, 663 (2004).
32. K. A. Takeuchi, M. Kuroda, H. Chaté, M. Sano, Directed percolation criticality in turbulent liquid crystals. *Phys. Rev. Lett.* **99**, 234503 (2007).