JEFF CLUNE

POSTDOCTORAL FELLOW
HOD LIPSON'S CREATIVE MACHINES LABORATORY
DEPARTMENT OF MECHANICAL AND AEROSPACE ENGINEERING
CORNELL UNIVERSITY

JEFFCLUNE@CORNELL.EDU

EDUCATION

Postdoctoral Fellowship, Cornell University, 2010 – Present.
Advisor: Hod Lipson

Postdoctoral Scientist, BEACON Center for the Study of Evolution in Action, Michigan State University, Summer 2010.

Ph.D. in Computer Science, Michigan State University. 2010. 4.0 GPA.
Committee: Charles Ofria, Robert T. Pennock, Richard E. Lenski, Erik D. Goodman

M.A. in Philosophy, Michigan State University. 2005. 4.0 GPA.

Honors B.A. in Philosophy, University of Michigan. 1999. 3.9 GPA.

GRANTS, FELLOWSHIPS, AND AWARDS

- NSF Postdoctoral Research Fellowship in Biology, 2010-2011. $123,000 USD
- Co-author of funded DARPA grant titled "Matter Compiler: Visual Interaction for Rapid Exploratory Design for Manufacturing." $824,063 USD
- Co-author of a successful funding request to the BEACON Center for the Study of Evolution in Action titled "Open-Ended Evolution of Ecologies of Digital Organisms in 3-Dimensions." The grant covered the cost of one graduate student salary for one year.
- Best Paper Award, Genetic and Evolutionary Computation Conference, 2009
- NSF Graduate Research Fellowship, Honorable Mention, 2006
- Co-author of pre-proposal for an Intel Center titled "Augmenting human invention."
- Finalist in the Evolutionary Art Competition (one of four), Genetic and Evolutionary Computation Conference, 2011
- Quantitative Biology and Modeling Initiative Research Fellowship, Michigan State University, 2005
- Dean's Recruitment Fellowship, Michigan State University, 2003-2005
- Branstrom Award, University of Michigan, 1995
- Angell Scholar, University of Michigan, 1999
- Research Travel Fellowship, Michigan State University, 2010
- Genetic and Evolutionary Computation Conference Travel Award, 2010

- European Conference on Artificial Life Student Fellowship, 2009
- Council of Graduate Students Conference Grant, 2009

CHAIRING & TUTORIALS

- Co-chair, Generative and Developmental Systems Track, Genetic and Evolutionary Computation Conference, 2010 & 2011
- The Avida Digital Evolution Platform, Tutorial, Genetic and Evolutionary Computation Conference, 2010
- Leveraging the Avida Digital Evolution Platform for Research in Evolving Cooperation, Tutorial, European Conference on Artificial Life, 2009

JOURNAL PUBLICATIONS

- Clune J, Pennock RT, Ofria C, Lenski RE (2012) Ontogeny tends to recapitulate phylogeny in digital organisms. The American Naturalist, in press.
- Clune J, Baptiste-Mouret J-B, Lipson H (2012) The evolutionary origins of modularity. Submitted.
- Clune J, Stanley KO, Pennock RT, Ofria C (2011) On the performance of indirect encoding across the continuum of regularity. IEEE Transactions on Evolutionary Computation. 15(3): 346-367. 5-year impact factor: 7.6, the top in Evolutionary Computation and 2nd overall in Artificial Intelligence.
- Clune J, Goldsby H, Ofria C, Pennock RT (2011) Selective pressures for accurate altruism targeting: Evidence from digital evolution for difficult-to-test aspects of inclusive fitness theory. Proceedings of the Royal Society. 278: 666-674. 5-year impact factor: 5.4.
- Clune J, Misevic D, Ofria C, Lenski RE, Elena SF, and Sanjuán R (2008) Natural selection fails to optimize mutation rates for long-term adaptation on rugged fitness landscapes. PLoS Computational Biology 4(9): e1000187. 5-year impact factor: 6.4.
- Clune J, Yosinski J, Doan E, Lipson H (2012) Automating the design of physical objects via interactive, crowdsourced evolution based on concepts from developmental biology. In preparation.

PEER REVIEWED CONFERENCE PUBLICATIONS

- Clune J, Lipson H (2011) Evolving three-dimensional objects with a generative encoding inspired by developmental biology. Proceedings of the European Conference on Artificial Life. 144-148.
- Yosinski J, Clune J, Hidalgo D, Nguyen S, Cristobal Zagal J, Lipson H (2011) Evolving robot gaits in hardware: the HyperNEAT generative encoding vs. parameter optimization. Proceedings of the European Conference on Artificial

Life. 890-897.

- Suchorzewski M, Clune J (2011) A novel generative encoding for evolving modular, regular and scalable networks. Proceedings of the Genetic and Evolutionary Computation Conference. 1523-1530.
- Clune J, Beckmann BE, McKinley PK, Ofria C (2010) Investigating whether HyperNEAT produces modular neural networks. Proceedings of the Genetic and Evolutionary Computation Conference. 635-642.
- Clune J, Beckmann BE, Pennock RT, Ofria C (2009) HybrID: A hybridization of indirect and direct encodings for evolutionary computation. Proceedings of the European Conference on Artificial Life. Vol. 2: 134: 141.
- Goldsby HJ, Knoester DB, Clune J, McKinley PK, Ofria C (2009) The evolution of division of labor. Proceedings of the European Conference on Artificial Life. Vol. 2: 10-18.
- Clune J, Pennock RT, and Ofria C (2009) The sensitivity of HyperNEAT to different geometric representations of a problem. Proceedings of the Genetic and Evolutionary Computation Conference. 675-682. **Winner of best paper award**.
- Goldsby HJ, Goings S, Clune J, and Ofria C (2009) Problem decomposition using indirect reciprocity in evolved populations. Proceedings of the Genetic and Evolutionary Computation Conference. 105-112.
- Clune J, Beckmann BE, Ofria C, and Pennock RT (2009) Evolving coordinated quadruped gaits with the HyperNEAT generative encoding. Proceedings of the IEEE Congress on Evolutionary Computing. 2762-2771.
- Clune J, Ofria C, and Pennock RT (2008) How a generative encoding fares as problem-regularity decreases. Proceedings of the 10th International Conference on Parallel Problem Solving From Nature. 358-367.
- Clune J, Ofria C, and Pennock RT (2007) Investigating the emergence of phenotypic plasticity in digital organisms. Proceedings of the European Conference on Artificial Life. 74-83.
- Clune J, Goings S, Goodman ED, and Punch W (2005) Investigations in meta-GAs: panaceas or pipe dreams? Proceedings of the Genetic and Evolutionary Computation Conference. 235-241.
- Goings S, Clune J, Ofria C, and Pennock RT (2004) Kin-Selection: The rise and fall of kin cheaters. Proceedings of the Ninth Conference on Artificial Life. 303-308.

EDITED PUBLICATIONS

- Krasnogor N, Auger A, Meyer-Nieberg S, Bernadó-Mansilla E, Ochoa G, Browne WN, Ong Y-S, Clune J, Pelta D, Coello Coello CA, Poulding S, Collet P, Preuss M, Eiben AE, Raidl GR, Engelbrecht AP, Ritchie M, Freitas AA, Schoenauer M, Gagné C, Sipper M, Gallagher MR, Smith J, Gershenson C, Spector L, Gustafson S, Squillero G, Hansen N, Watson J-P, Hornby GS, Witt C, Landa-Silva D, Wong ML, Lozano JA and Yu T (2011) editors, Proceedings of the Genetic and Evolutionary Computation Conference. ACM Press, New York, New York.

- Branke J, Alba E, Arnold D, Bongard J, Brabazon A, Butz MV, Clune J, Cohen M, Deb K, Engelbrecht A, Krasnogor N, Miller JF, O'Neill M, Sastry K, Thierens D, Vanneschi L, van Hemert J and Witt C (2010) editors, Proceedings of the Genetic and Evolutionary Computation Conference. ACM Press, New York, New York, 2010.

INVITED TALKS

- Paris Descartes Medicine Faculty, Center for Interdisciplinary Research. 2011.
- BEACON Center for the Study of Evolution in Action. Michigan State University. 2011.
- Cornell University, Artificial Intelligence Seminar. 2011
- Institute for Intelligent Robotic Systems, Université Pierre et Marie Curie, Paris 6. 2011
- University of Lausanne, labs of Laurent Keller and Dario Floreano. 2011
- Cornell University, Dynamics Systems and Control Series. 2011
- Cornell University, Computational Systems Laboratory. 2011.
- Stanford University, Mark Feldman's Lab. 2010
- Dalle Molle Institute for Artificial Intelligence, Università della Svizzera Italiana. Jürgen Schmidhuber's Laboratory, Lugano, Switzerland. 2010
- University of California, Irvine. Cognitive Robotics Laboratory. 2009
- Michigan State University. Ecology, Evolution, Biology, and Behavior Colloquium. 2009.
- Cornell. Hod Lipson's Creative Machines Lab. 2008.

TEACHING EXPERIENCE

- Instructor - Integrative Arts and Humanities: Self, Society and Technology. Michigan State University, Fall, 2005. Responsibilities: Managed the entire course, including developing the curriculum, delivering the lectures, and grading the approximately 100 students.
- Teaching Assistant - Introduction to Programming II (C++) Michigan State University, Spring, 2009. Responsibilities: Gave short lectures at the beginning of laboratory sessions, assisted students in completing their weekly laboratory assignments, and graded the students' weekly projects.
- Teaching Assistant - Introduction to Programming I (Python) Michigan State University, Spring, 2008. Responsibilities: Gave short lectures at the beginning of laboratory sessions, assisted students in completing their weekly laboratory assignments, and graded the students' weekly projects and their exams.
- Teaching Assistant - Integrative Arts and Humanities: Freedom in the Modern World. Michigan State University, Fall, 2004. Responsibilities: Led weekly class discussion sections and graded student essays.
- Guest lecture for Artificial Intelligence, Cornell University, 2010 & 2011.

- Guest lecture for Evolutionary Computation, Michigan State University, 2006.

MENTORING

- Jason Yosinski, Ph.D. student evolving robotic gaits and 3D objects
- Anthony Chen, undergraduate researching brain computer interfaces
- Eugene Doan, undergraduate developer of EndlessForms.com
- Nick Donohue, undergraduate researching evolving HyperNEAT topologies
- Chris Heuser, undergraduate researching evolving gaits for legged robots

PRESS COVERAGE

- The New Scientist (featured on cover). 2011. Darwin's robots: A holistic, evolutionary approach means that robots could learn to design themselves.
- Science. 2006 (vol. 311). Darwin's Place on Campus Is Secure—But Not Supreme.
- MSNBC.com. 2011. Intelligent design: Users power evolution in 3-D Web printing.
- Slashdot. 2011. Crowdsourcing speeds evolution of 3D printable objects.
- The New Scientist (featured on cover). 2010.
  - Main article: Artificial life forms evolve basic intelligence
  - Editorial: Digital evolution and the meaning of life
- US News & World Report. 2010. New MSU research sheds light on how we become altruistic.
- The Daily Telegraph. 2010. Computer-simulated life forms evolve intelligence.
- Slashdot. 2010. Artificial life forms evolve basic memory, strategy.
- MIT Technology Review. 2011. 3-D design simplified: a new website could accelerate the adoption of 3-D printing.
- Communications of the ACM. 2010. 'EndlessForms' uses the Web to breed 3D printable objects.
- KurzweilAI.net. 2010. Artificial life forms evolve basic intelligence.
- Discover. 2005 (cover article). Testing Darwin.
- ScienceDaily.com. 2011. No technical know-how needed: Endless Forms Web site helps users 'breed' 3-D printable objects.
- ScienceDaily.com. 2010. Research sheds light on altruism.
- Hacker News (front page). 2011. Breed 3D printable objects, no technical know-how needed.
- Research also covered in the following: Lansing State Journal, Jerusalem Post, Innovation News Daily, LiveScience.com, PhysOrg.com, eCampusNews, NewsWise.com, BigThink.com, Business News Daily, Cornell Chronicle, Cornell Daily Sun, State News (twice), Shapeways blog, Thingiverse blog, Carl Zimmer's blog, 3Dprinter.net, PlasticsToday.com, Biota Live Podcast, Impact Radio, City Pulse, OneIndia.com, Computerra.com, TodayOnline.com,

Creativity Online, Heise.de, ZeitNews.org, MyScience.cc, Ponoko.com, TheHighLow.com, and 40+ other media outlets.

OUTREACH

- Led team that built EndlessForms.com, a website where non-technical users can design 3-D, printable objects with evolutionary algorithms based on concepts from developmental biology. The site also enables the public to learn about evolution and see its ability to create complexity. To date, over 2 million objects have been evaluated by nearly 40 thousand visitors from over 140 countries and all 50 US states. Video tours of the site have been viewed nearly 10,000 times.
- One of three designers and developers of Avida-ED, a software package used in university biology classes to teach evolution. Avida-ED enables students to conduct research in experimental evolution by testing evolutionary hypotheses and getting immediate feedback. Avida-ED was discussed in Science magazine (2006: 311) and has been used in many universities worldwide. The NSF grant overseer for Avida-ED described it as "one of the most successful science education materials projects with which I am acquainted. The product is excellent, dissemination is already successful, and the assessment plan is outstanding."

SERVICE

- Reviewer:
  - IEEE Transactions on Neural Networks
  - Journal of Machine Learning Research
  - IEEE Transactions on Evolutionary Computation
  - Evolutionary Computation
  - Neural Computation
  - SIGGRAPH Conference
  - Adaptive Behavior
  - Astrobiology
  - Soft Computing
  - Genetic and Evolutionary Computation Conference
  - Artificial Life Conference
  - AAAI AI Video Competition
- Robotics/AI and Diplomacy Board Member, The Lifeboat Foundation
- Computer Science and Engineering Graduate Student Association, Liaison to Computer Science Department Faculty Meetings, MSU, 2009-2010
- Computer Science and Engineering Advisory Committee, MSU, 2009-2010
- College Hearing Board, MSU, 2009-2010

- Panelist & Speaker, Graduate Student Orientation and Recruitment
- Founder, Meteorite, University of Michigan Undergraduate Journal of Philosophy
- Associate Editor, Michigan Journal of Political Science, 1996-1998

POSTERS

- Strelioff C, Clune J, Ofria C, Lenski R, Epstein C (2009) Evolution of mutation rates. Workshop for Young Researchers in Mathematical Biology. Columbus, OH.
- Clune J, Ofria C, and Pennock RT (2008) How generative encodings fare on less regular problems. Poster at the Genetic and Evolutionary Computation Conference. Atlanta, GA.
- Smith JJ, Pennock RT, Clune J, Armstrong E, Braverman M, Brady C (2006) Adapting avida as an evolution education tool: development of model lesson plans. Sigma Xi, The Scientific Research Society Annual Meeting and Student Research Conference. Detroit, MI.
- Smith JJ, Pennock RT, Clune J, Armstrong E, Braverman M, Brady C (2006) Adapting avida as an evolution education tool: development of model lesson plans. Society for the Study of Evolution. Stony Brook, NY.

WORK EXPERIENCE

- Co-founder and owner, RoomSimple, a house rental company in Lansing, MI. 2003-Present.
- Analyst Relations Manager, SoftAd, an Internet Software Company. San Francisco, California. 2000-2001.
- Business Development Strategist, SoftAd, an Internet Software Company. Dearborn, MI and San Francisco, CA. 1999-2000.
- Marketing Intern, Ford Motor Company. Dearborn, MI. Summer, 1998 & 1999.
- Co-founder and owner, Full Immersion, a website development company. Ann Arbor, MI. 1997-1998.

PROFESSIONAL SOCIETIES

- Sigma Xi
- Phi Beta Kappa
- Phi Kappa Phi
- Golden Key National Honor Society
- Association for Computing Machinery

- IEEE

INTERESTS

- Travel (over 50 countries on 6 continents)
- Sports (surfing, kitesurfing, rock climbing, hockey, whitewater kayaking, ultimate frisbee, hiking, mountain climbing, running)
- Literature (Borges, Kundera, Calvino, Penn Warren, Dostoyevsky, DeLillo, Marquez, Card, Tolkien, Tolstoy, Carroll, Pirsing, Stephenson)
- Spanish language (fluent)
- Other (music, writing, philanthropy)

REFERENCES

- Hod Lipson, Associate Professor, Departments of Mechanical and Aerospace Engineering and Computing and Information Science, Cornell University. Postdoctoral advisor. hod.lipson@cornell.edu. (607) 254-8940.
- Charles Ofria, Associate Professor, Department of Computer Science and Engineering, Michigan State University. Ph.D. and postdoctoral advisor. ofria@cse.msu.edu. (517) 355-8389.
- Robert T. Pennock, Professor, Lyman Briggs College and Department of Computer Science and Engineering, Michigan State University. Ph.D. advisor. pennock5@msu.edu. (517) 432-7701.
- Richard E. Lenski, National Academy of Sciences member, Hannah Distinguished Professor, Department of Microbiology and Molecular Genetics, Michigan State University. Ph.D. committee member. lenski@msu.edu. (517) 355-6463 ext.1603.
- Eric Wieschaus, Nobel Prize Winner and Professor, Department of Molecular Biology, Princeton University. efw@princeton.edu. (609) 258-5383.
- Erik Goodman, Director of BEACON: An NSF Center for the Study of Evolution in Action, Professor of Electrical and Computer Engineering, Mechanical Engineering, and Computer Science and Engineering, Michigan State University. Ph.D. committee member. goodman@egr.msu.edu. (517) 355-6453.
- Kenneth O. Stanley, Associate Professor, School of Electrical Engineering and Computer Science, University of Central Florida. kstanley@cs.ucf.edu. (407) 473-0072.

Understanding the Evolution of Intelligence by Re-Evolving It

Research Statement for Jeff Clune

Summary

I experimentally investigate the evolution of complex biological traits and have shed light on the evolutionary origins of intelligence[1–8], phenotypic plasticity[9], modularity[10], evolvability[11,12], and altruism[13–16]. I conduct experiments in computational simulations of evolution because they enable unprecedented levels of control, speed, and data. For example, in a paper currently under review in *Science*, my colleagues and I performed experiments with computational evolution that would be impossible in a natural system[10]. These experiments enabled us to identify a key reason why biological networks, such as the neural networks that make up animal brains, evolve the important property of modularity. I will focus my career on understanding the evolution of intelligence by isolating the factors that promote its evolution. Such research not only answers fundamental biological questions regarding the evolutionary origins of intelligence, but it also enables engineers to create synthetic evolutionary processes to produce more sophisticated designs, such as artificially intelligent robots. Porting biological insights to engineering also further sheds light on cognition because we learn a surprising amount about something by trying to build it.

To study the evolution of intelligence I evolve digital brains called artificial neural networks[17]. Although neural networks evolved with current algorithms show impressive levels of sophistication, and often outperform human engineers[8,18], they pale in comparison to the complexity of natural brains. I investigate the forces that give natural bodies and brains important properties such as modularity, regularity, and hierarchy, and have shown that when artificial intelligence algorithms are modified to generate these properties they produce significantly more complex designs and more intelligent behaviors[1–8,10,19]. I have also shown that more sophisticated intelligence evolves when evolution is combined with concepts from developmental biology. Adding a developmental phase that captures key aspects of how nature grows complex organisms from genomic information enables the evolution of complex, regular neural networks that increase the performance and coordination of robotic behaviors[1–8] (Figure 1). This research also informs developmental biology, by enabling us to experimentally investigate which features of development are important in the evolution of complex traits.

I will continue to use computational evolution to investigate how natural evolution produced complex, structurally organized brains. I will test the validity of the generated models in simulated and physical robots, with the goal of eventually making robots as smart and capable as natural animals. Specifically, I will 1) scale evolved neural networks up to the scale of animal brains (millions of neurons), which is only now possible due to the breakthroughs described below, 2) combine these advances with neural learning algorithms to investigate the interplay between innate and learned behaviors, and to improve robotic artificial intelligence 3) identify additional factors that promote structural organization in neural networks 4) deploy evolved neural networks on simulated and physical robots to study the evolution of intelligence and behavior.

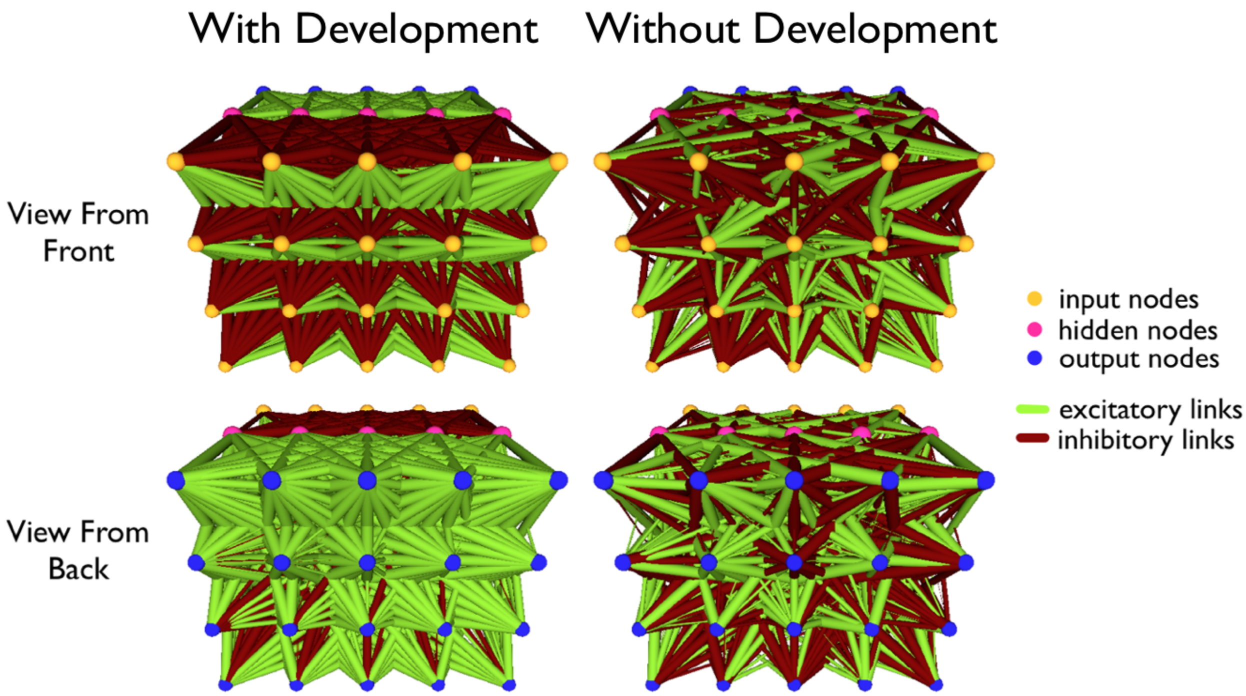

Figure 1 | Robot neural networks evolved with concepts from developmental biology (left) versus a traditional, non-developmental encoding (right)[1]. The developmental neural network is more structurally organized, with regular, repeating motifs, symmetries, and graded variation. In contrast, the non-developmental neural network is disordered and shows none of the structural organization seen in natural brains.

Combining Evolution with Concepts from Developmental Biology

A powerful way to understand the evolution of intelligence is to attempt to reproduce it in computational simulations of evolution. To evolve artificial intelligence that rivals the intelligence of animals we need to design powerful genetic *encodings*–the way information is stored in a genome and how that information is mapped to a phenotype. Traditional *direct encodings*–where each element in the genotype specifies an independent component of the phenotype–are infeasible at natural scales, where brains can have trillions of connections. Evolution cannot effectively search such high-dimensional spaces. The more natural *generative encodings*–where genomic information can influence multiple aspects of a phenotype–enable the evolution of large-scale neural networks and facilitate the evolution of structural organization, including modularity, regularity, and hierarchy[1–3,5,19–21].

Simply using generative encodings is not enough, however, as many generative encodings are instantiated at a level of abstraction that is either too low, with prohibitive computational costs, or too high, producing overly simple regularities. In 2007, a new generative encoding called a Compositional Pattern Producing Network (CPPN) was introduced that captures some of the power of natural developmental processes without expensive computation[22,23]. The key insight is that the fate of organismal components, such as cells, depends on their location in coordinate frames built in geometric space. Developing organisms produce increasingly complex coordinate frames, starting with simple gradients (e.g. anterior-posterior) and then build up complexities, such as symmetries (e.g. bilateral) and repeating themes (e.g. segments). While embryos build these gradients and coordinate frames with gene regulation and diffusing proteins, CPPNs achieve the same effect with networks of computationally fast math functions.

Visually comparing the phenotypes produced by CPPNs and previous generative encodings illustrates the advance CPPNs represent in terms of producing entities that look complex and natural, instead of simple and formulaic (Figure 2). I have shown that CPPNs generate complex three-dimensional designs, and led a team that built the website EndlessForms.com to enable people who lack technical skills to design compelling objects via CPPN-driven interactive evolution[7,24]. MIT's *Technology Review* quoted Neri Oxman of the MIT Media Lab as saying that the user-friendliness of the EndlessForms approach "could help drive the broader adoption of 3-D printing technologies, similar to how easy-to-use image editors fueled the growth of digital photography and graphic manipulation...this could ultimately have an impact on design similar to the impact that blogs and social media have had on journalism, opening the field to the general public."[25] In its first few months, over 2.5 million objects have been evaluated on the site by visitors from 145 countries and all 50 states. The project was also covered by MSNBC.com, the *New Scientist*, *Slashdot*, *Communications of the ACM*, and over 50 other media outlets. Additionally, EndlessForms.com helps educate the public about how evolution can generate complex designs through a series of small steps.

CPPNs can create similarly complex geometric patterns in the neuronal connections of neural networks, which has proved beneficial in problem domains such as checkers[28], coordinating teams of predators[29], and a pattern-recognition task where CPPNs produced a functional neural network with over eight million neural connections[23]. I have demonstrated that neural networks evolved with CPPNs increasingly outperform directly encoded neural networks as the regularity of problems increase, including on the difficult problem of evolving a neural network capable of producing running gaits for a simulated legged robot[1–3]. The neural networks evolved to control this robot exhibited desirable features of organization, such as symmetries and repeated neural motifs (Figure 1). I also showed that engineers can use CPPNs to inject domain knowledge or preferences into the design algorithm by altering the geometric representation of phenotypic components, such as requesting gaits with certain symmetries (e.g. front-back)[5]. Neural networks evolved with CPPNs produced similarly impressive results on a physical robot, where they produced gaits faster than those designed by traditional machine learning algorithms and one hand-designed by the robot's designer[8]. This work highlights the role development plays in creating structural organization in brains and was featured on the cover of the *New Scientist* twice, in the *Daily Telegraph*, on *KurzweilAI.net*, and on *Slashdot*.

I also found that CPPNs are overly biased toward regular phenotypes, and have problems creating certain exceptions to the patterns they generate. To remedy this bias toward regularity, I introduced a new algorithm called HybrID that combines the best attributes of generative and direct encodings[1,4]. HybrID yielded performance improvements of up to 40% over CPPNs, and its advantage was most pronounced on problems with intermediate regularity, which is representative of real-world problems. HybrID's success led me to hypothesize that a way toward evolving higher levels of artificial intelligence is to follow biology in having generative encodings create innate neural wiring patterns and then allow neural learning algorithms to adjust those patterns to the

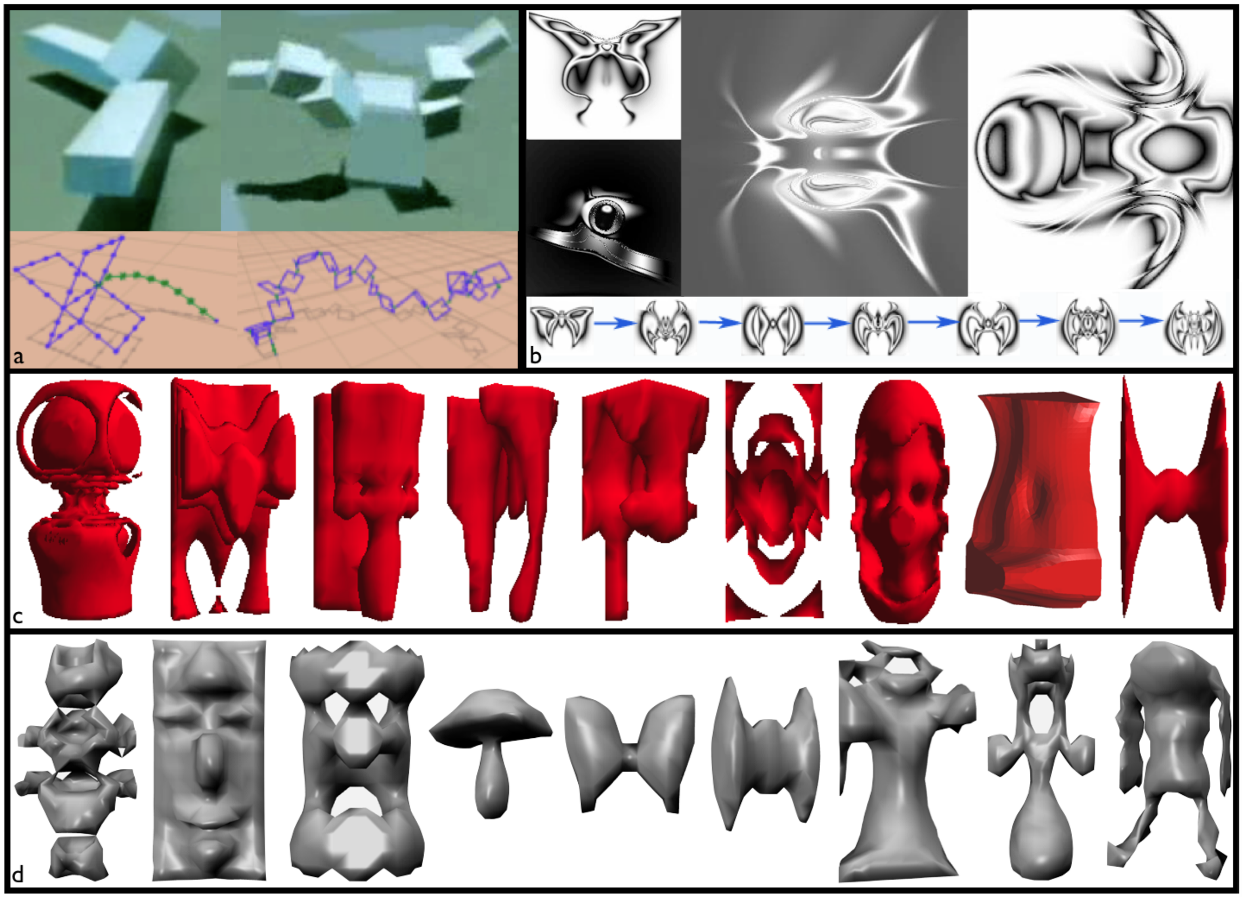

Figure 2 | Phenotypes evolved with previous generative encodings (a)[21,26] versus CPPNs (b-d). (b) Images evolved on PicBreeder.org[27]. (c) Objects I evolved to investigate the ability of CPPNs to design interesting three-dimensional forms[7]. (d) Objects evolved on EndlessForms.com, a crowdsourced exploration of the space of CPPN designs that non-technical visitors can use to create objects for 3D printing[24].

needs of specific environmental challenges[1]. A preliminary independent study recently supported this hypothesis[30], revealing a fruitful research area that I will explore early in my career.

A study of mine uncovered that CPPNs tend to evolve less modular neural networks than other classes of generative encodings[6,19]. That led me to investigate and discover a force that causes modularity in natural networks, such as metabolic, gene regulatory, and neural networks[10]. In future research I will harness this force to study the effects of modular organization on the level of intelligence that evolution can produce. I have also begun work discovering how to evolve other network properties, such as being scale-free[19], and will continue such research in the future, especially focusing on evolving hierarchically organized networks.

Future Work

I will use simulated evolution and physical robots to pursue the following research goals:

1. Scaling neural networks to natural scales (millions of neurons) - Traditional algorithms evolve

small networks with hundreds of neurons. In contrast, CPPNs have produced neural networks with millions of neurons[23], but their connectivity is relatively simple. By identifying how to evolve networks with structural organization, my research will enable the evolution of complex, functional neural networks at natural scales, which can reach billions of neurons. Such work will open new fields of research using controlled experiments to isolate the factors that promote the evolution of different levels and types of intelligence.

2. Studying geometrically patterned neural learning - Different parts of natural brains learn in different ways. Generative encodings enable the evolution of organized patterns in neural plasticity (learning rules) and can provide innate connectivity regularities that provide good starting points for neural network learning algorithms. I will begin by combining deep learning algorithms[31] with CPPN-evolved neural networks.

3. Identifying additional drivers of structural organization in biological networks - I have shed light on what causes the evolution of networks that are regular, modular, and scale-free. I will continue to investigate how best to design neural networks with these and other important attributes, such as hierarchy and having small-world connectivity.

4. Evolving robot morphologies - I will co-evolve robotic morphologies along with their neural networks to allow the objects in Figure 2c-d to come to life and run, fly, and swim. Evolution will control developmental processes that will deposit different materials (e.g. muscle, tendons, bone), increasing the freedom to design soft-bodied robots that are more similar to animals. This work will enable studies into the evolution of different morphologies and the co-evolution of brains and bodies. It is often argued that having a body is a key to evolving intelligence, and this system will provide a powerful way to test this hypothesis. The resultant designs will also be a convincing demonstration of the power of evolution and could be combined with 3D printing technology to allow evolved creatures to leave the virtual world and enter our own.

Conclusions

The human brain has billions of neurons and trillions of connections, yet it is encoded with less than 25,000 genes. Moreover, it exhibits modularity, regularity, and hierarchy. To better understand the evolutionary origins of intelligence, we need to be able to produce organized neural network models at this scale and with equivalent levels of structural organization. Being able to evolve large-scale, functional, structurally organized neural network models would enable new types of science to better understand those networks and the conditions under which they are produced. My research has demonstrated that it is possible to make progress toward these goals by identifying how natural evolution produced key organizational properties in animal brains and by implementing abstractions of them in evolutionary models. I will continue to use this strategy to improve our understanding of the evolution of intelligence.

References

1. Clune, J., Stanley, K., Pennock, R., and Ofria, C. (2011) On the performance of indirect encoding across the continuum of regularity. *IEEE Transactions on Evolutionary Computation*, **15**, 346–367.
2. Clune, J., Ofria, C., and Pennock, R. (2008) How a generative encoding fares as problem-regularity decreases. *Parallel Problem Solving from Nature*, pp. 358–367, Springer.
3. Clune, J., Beckmann, B., Ofria, C., and Pennock, R. (2009) Evolving coordinated quadruped gaits with the HyperNEAT generative encoding. *Proceedings of the IEEE Congress on Evolutionary Computation*, pp. 2764–2771.
4. Clune, J., Beckmann, B., Pennock, R., and Ofria, C. (2009) HybrID: A Hybridization of Indirect and Direct Encodings for Evolutionary Computation. *Proceedings of the European Conference on Artificial Life*.
5. Clune, J., Ofria, C., and Pennock, R. (2009) The sensitivity of HyperNEAT to different geometric representations of a problem. *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 675–682.
6. Clune, J., Beckmann, B., McKinley, P., and Ofria, C. (2010) Investigating whether HyperNEAT produces modular neural networks. *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 635–642, ACM.
7. Clune, J. and Lipson, H. (2011) Evolving three-dimensional objects with a generative encoding inspired by developmental biology. *Proceedings of the European Conference on Artificial Life*, pp. 144–148.
8. Yosinski, J., Clune, J., Hidalgo, D., Nguyen, S., Zagal, J., and Lipson, H. (2011) Evolving robot gaits in hardware: the hyperneat generative encoding vs. parameter optimization. *Proceedings of the European Conference on Artificial Life*, pp. 890–897.
9. Clune, J., Ofria, C., and Pennock, R. (2007) Investigating the emergence of phenotypic plasticity in evolving digital organisms. *Proceedings of the European Conference on Artificial Life*, pp. 74–83, Springer.
10. Clune, J., Mouret, J.-B., and Lipson, H. The evolutionary origins of modularity. *Science, in review.*.
11. Clune, J., Misevic, D., Ofria, C., Lenski, R., Elena, S., and Sanjuán, R. (2008) Natural selection fails to optimize mutation rates for long-term adaptation on rugged fitness landscapes. *PLoS Computational Biology*, **4**, e1000187.
12. Clune, J., Goings, S., Punch, B., and Goodman, E. (2005) Investigations in meta-gas: panaceas or pipe dreams? *Proceedings of the Genetic and Evolutionary Computation Conference Workshops*, pp. 235–241, ACM.
13. Clune, J., Goldsby, H., Ofria, C., and Pennock, R. (2011) Selective pressures for accurate altruism targeting: evidence from digital evolution for difficult-to-test aspects of inclusive fitness theory. *Proceedings of the Royal Society B: Biological Sciences*, **278**, 666–674.
14. Goings, S., Clune, J., Ofria, C., and Pennock, R. (2004) Kin selection: The rise and fall of kin-cheaters. *Proceedings of the Ninth International Conference on Artificial Life*, pp. 303–308.
15. Goldsby, H., Goings, S., Clune, J., and Ofria, C. (2009) Problem decomposition using indirect reciprocity in evolved populations. *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 105–112, ACM.
16. Goldsby, H., Knoester, D., Clune, J., McKinley, P., and Ofria, C. (2011) The evolution of division of labor. *Proceedings of the European Conference on Artificial Life*, vol. 2, pp. 10–18, Springer.
17. Stanley, K. and Miikkulainen, R. (2002) Evolving neural networks through augmenting topologies. *Evolutionary Computation*, **10**, 99–127.
18. Hornby, G., Takamura, S., Yamamoto, T., and Fujita, M. (2005) Autonomous evolution of dynamic gaits with two quadruped robots. *IEEE Transactions on Robotics*, **21**, 402–410.
19. Suchorzewski, M. and Clune, J. (2011) A novel generative encoding for evolving modular, regular and scalable networks. *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 1523–1530.
20. Stanley, K. and Miikkulainen, R. (2003) A taxonomy for artificial embryogeny. *Artificial Life*, **9**, 93–130.

21. Hornby, G., Lipson, H., and Pollack, J. (2003) Generative representations for the automated design of modular physical robots. *IEEE Transactions on Robotics and Automation*, **19**, 703–719.
22. Stanley, K. (2007) Compositional pattern producing networks: A novel abstraction of development. *Genetic Programming and Evolvable Machines*, **8**, 131–162.
23. Stanley, K., D'Ambrosio, D., and Gauci, J. (2009) A hypercube-based encoding for evolving large-scale neural networks. *Artificial Life*, **15**, 185–212.
24. Clune, J., Yosinski, J., Doan, E., and Lipson, H. Automating the design of physical objects via interactive, crowdsourced evolution based on concepts from developmental biology. *In prep.*.
25. Cass, S. (2011) 3-d design simplified: A new website could accelerate the adoption of 3-d printing. *MIT Technology Review*.
26. Sims, K. (1994) Evolving 3D morphology and behavior by competition. *Artificial Life*, **1**, 353–372.
27. Secretan, J., Beato, N., D'Ambrosio, D., Rodriguez, A., Campbell, A., Folsom-Kovarik, J., and Stanley, K. (2011) Picbreeder: A Case Study in Collaborative Evolutionary Exploration of Design Space. *Evolutionary Computation*, **19**, 373–403.
28. Gauci, J. and Stanley, K. (2010) Autonomous evolution of topographic regularities in artificial neural networks. *Neural Computation*, **22**, 1860–1898.
29. D'Ambrosio, D. and Stanley, K. (2008) Generative encoding for multiagent learning. *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 819–826, ACM.
30. Tonelli, P. and Mouret, J.-B. (2011) On the relationships between synaptic plasticity and generative systems. *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 1531–1538, ACM.
31. Hinton, G. and Salakhutdinov, R. (2006) Reducing the dimensionality of data with neural networks. *Science*, **313**, 504.

Teaching for me involves transforming students from passive receivers of information to active learners that take charge of their own learning processes. I focus on two key components of this transformation: I teach students general approaches to problem solving, especially how to obtain missing information, and I train them to critically evaluate new information before integrating it into their existing knowledge framework. I also have students constantly practice communication skills, both verbal and written, and provide real-world examples of computer science in action to help students see the relevance of what they are learning. While most students are hesitant at first to take charge of their own education, they quickly flourish in their new role and change their perceptions of learning: instead of a boring process that must be endured, they soon see learning as an exciting opportunity to explore the world, expand their skill set, and create novel solutions to challenging problems.

Early in a semester I have students reflect on why they receive the grades they do. Many students feel that their poor grades are due to capricious and unfair teachers. That attitude evaporates once they are put in the position of grading their peers. My students grade the work of an anonymous classmate and write a short justification of that grade. I do not use the grades the students give each other, but I do grade their justifications. After this process, the students' quality of work rises substantially. Many students report that this exercise causes them to think about how they would grade their own work. They suddenly find new flaws and fix them, as they want to turn in something that will be thought well of, now that they realize that their work represents who they are.

I teach my students to take the initiative to get answers to their questions. To facilitate this behavior, I create a class environment where students constantly ask and answer questions. I have the students meet a new classmate each week for the first month, so they feel comfortable asking each other questions, and I encourage them to ask questions when I am speaking. If I encounter silence when asking for a class response, I use simple but effective techniques to switch the class back into a state wherein questions and comments flow freely. For example, when I hit a wall of silence, I have every student in the room raise their hand for a few seconds. Once all the hands go down, I repeat my question and invariably a number of hands go up. I have found that creating an environment where students are active participants produces an effective learning environment, whether the class size is less than 10 or over 100.

It is rewarding to see students transition from expecting to be fed answers to answering their own questions. I facilitate this transformation by not answering a student's question, but showing them how to answer it for themselves. For example, when students get stuck while writing code, I show them techniques for debugging, looking up documentation, or searching relevant Internet forums. Soon, when students ask me a question, they start by explaining the different ways they tried to solve the problem. They do not want just the solution, but look forward to learning a new process for knowledge discovery. I also let advanced students take pride in their skills by pairing them with students that are struggling. This practice helps the struggling students, who learn from their peers, and the advanced students, who learn by teaching.

An important step toward becoming an active learner is identifying how one best takes in knowledge. To catalyze this process, I assess student attitudes towards different class activities and ask how they would improve them. These assessments help me rebalance the class to meet student needs and get the students reflecting on how they and others learn. I then share the results with the class, and many students are surprised to see that others benefit from parts of the class they dislike, which encourages them to better use those portions of the class.

A critical step in intellectual maturity is realizing that knowledge alone is insufficient if you cannot communicate your ideas to others. To teach this point, I divide classes into groups and have them propose solutions to a problem. Students first have to communicate their individual ideas to team members and then work together to orally present their group's proposal. The class then votes on which proposed solution to adopt. After the vote, students often comment that the best proposal lost only because it had not been explained well enough. This realization emphasizes the lesson that communication can be as important as innovation.

Students need to know the relevance of what they learn, especially in computer science, where the nuts and bolts of the discipline are often abstracted in a way that obscures their connection to the real world. At the start of each class I present a real-world problem related to the class material and ask the students how they would solve it. We then go through how that problem has been solved, and debate the merits of different approaches. This phase of the class period is a great way to introduce students to cutting-edge computer science research and motivates them to think about computer science more broadly. Instead of seeing computer science as simply software programming, they learn that it is a broad field involved in solving complicated and interesting problems. Many of these undergraduates volunteer to perform research and express an interest in graduate school. Assessments I have conducted reveal that the vast majority of my students find this section of my classes interesting and worthwhile.

My teaching reviews are consistently high because students realize that I care deeply about them maturing into capable active learners. Here are anonymous reviews that are typical of the comments I receive from students.

> *[Jeff] is without a doubt one the best teachers I have had during my four years at MSU. He gives fascinating, thought-provoking reading material and challenges the way you think through class discussion and short (but intellectually challenging) reaction paper topics. Highly, highly recommended!*
>
> *Jeff always made the lab interesting and was great at making our problem solving skills better.*
>
> *Jeff Clune was probably one of the best teachers I have ever had at Michigan State. You WILL Work ALOT in his class. … I have learned so much about me and the subject. If Jeff is teaching a class, I would highly recommend him to any of you. Overall, Great Teacher!*

What I am most proud of in these reviews is that the students enjoyed a challenging class because they became excited about learning. My favorite moment in teaching is at the end of a semester when I stand back and watch groups of students solving problems on their own and clearly explaining their solutions. That is the reward for the hard work of not answering questions, but showing students how they can find those answers themselves. It is the reward for setting high expectations and demanding clear explanations. Above all, it is the reward for not only teaching the material in a single course, but helping students become independent learners who actively develop the skill sets they need to accomplish the goals they set for themselves.