

The Case for a Mixed-Initiative Collaborative Neuroevolution Approach

Sebastian Risi, Jinhong Zhang, Rasmus Taarnby, Peter Greve, Jan Piskur, Antonios Liapis, Julian Togelius

IT University Copenhagen, 2300 Copenhagen, Denmark
{sebr, jinh, reta, pgre, japi, anli, juto}@itu.dk

Abstract

It is clear that the current attempts at using algorithms to create artificial neural networks have had mixed success at best when it comes to creating large networks and/or complex behavior. This should not be unexpected, as creating an artificial brain is essentially a design problem. Human design ingenuity still surpasses computational design for most tasks in most domains, including architecture, game design, and authoring literary fiction. This leads us to ask which the best way is to combine human and machine design capacities when it comes to designing artificial brains. Both of them have their strengths and weaknesses; for example, humans are much too slow to manually specify thousands of neurons, let alone the billions of neurons that go into a human brain, but on the other hand they can rely on a vast repository of common-sense understanding and design heuristics that can help them perform a much better guided search in design space than an algorithm. Therefore, in this paper we argue for a mixed-initiative approach for collaborative online brain building and present first results towards this goal.

With around 200 billion neurons and 125 trillion synapses the human brain is the most complex system known to exist. The brain's structural complexity, with intricate synaptic motifs that repeat throughout it, gives rise to our unique mental abilities. Additionally, the brain's plasticity, allows us to learn new abilities throughout our life (e.g. learning a new language) and to change our behavior based on past experience.

Therefore, creating similar artificial structures by recapitulate the process that created intelligence on earth, is an intriguing possibility. In this context, neuroevolution (i.e. evolving artificial neural networks (ANNs) through evolutionary algorithms) has shown promising results in a variety of different domains (Floreano et al., 2008; Stanley and Miikkulainen, 2004; Yao, 1999). However, these results still pale in comparison to the capabilities of natural brains. The reasons for this are manifold. Especially the problem of deceptive fitness landscapes (i.e. mutations increase fitness but actually lead further away from the final objective) has limited the scope of problems amenable to evolutionary algorithms.

Research in circumventing this problem has mainly fo-

cused on two different ideas. First, novelty search (Lehman and Stanley, 2011), a method that rewards novel behaviors instead of rewarding objective performance, has shown promise and significantly outperforms objective-based approaches in a variety of different domains. Other approaches are based on interactive evolutionary computation (IEC) methods, wherein the human user guides evolution by repeatedly choosing from a set of candidates. Woolley and Stanley (2011) showed that interactive evolution can help discovering artifacts that are very hard to evolve with traditional objective-based methods. Recently Woolley and Stanley (2014) combined IEC with novelty search, demonstrating that the approaches complement each other and together address some of the challenges that each method struggles with by itself (e.g. novelty search can get lost in large search spaces, interactive evolution is limited by user fatigue). In a related approach, Bongard and Hornby (2013) recently demonstrated that human input and objective based search can also be combined synergistically to solve challenging robotic domains.

While these *human in the loop* approaches have shown promising results, we argue in this paper that they do not exploit the whole range of human intuition, ability to identify promising stepping stones, or collaborative problem solving skills that could ultimately allow us to create more brain-like artificial structures. For example, collaborative games like Foldit hint at the power of crowdsourcing the human brain's natural abilities for specific tasks (e.g. pattern matching, spatial reasoning), which are hard to solve by computational approaches. However, in traditional IEC applications the role of the user is often reduced to solely judging the created artifacts and only influencing evolution by deciding between a discrete choice of candidates. In other words, only the computer creates content (e.g. images, ANNs, etc.) and the role of the human is to guide evolution to content they prefer.

An approach that does require significantly more input from the user is a *mixed-initiative process* (Liapis et al., 2014), wherein both the computer and the human take turns in creating content. Yet while mixed-initiative based approaches have shown promise in the context of procedurally

generated content for games (Liapis et al., 2014, 2013), they have not yet been applied to the evolution of large-scale ANNs for complex tasks. However, the promise of such an approach is a system that can benefit from the different skill sets of a human and a computational creator. For example, while a human user might develop an intuition about promising domain-dependent network topographies, a computational method is likely more effective at fine-tuning specific synaptic weights. Therefore, we argue for a collaborative mixed-initiative approach in which both, the computer and collaborating human users can take initiative and propose changes to an evolving neural networks. That is, at any point the human can revert to just doing selection and let the evolutionary algorithm serve up new content to judge, or decide to jump in and have a more active role in the design of the ANN.

Our recent work has taken steps towards this goal by focusing on two parallel lines of research. First, before introducing a mixed-initiative approach it is useful to determine how good we as humans are at building complex ANN for certain tasks (without evolution) and if collaborating with other users proves useful. Insights from this experiment should in turn provide useful clues about the strengths of human networks designers and most importantly, non-intuitive aspects of the design process we tend to struggle with (i.e. aspects which would benefit most from an assistant computational creator). In this context we recently introduced *BrainCrafter* (BrainCrafter.dk), which is an online application that allows the user to build ANNs for specific control programs (e.g. a robot that must traverse a maze) by adding neurons and connections in a drag&drop like fashion (Figure 1). While building ANNs the users can observe the resulting simulated robot behaviors in real-time, providing insights into the effects of different network modifications. Importantly, BrainCrafter also allows collaborations between users by building on high scoring solutions created by other people.

While the ongoing BrainCrafter experiment should provide insights into our abilities to collaboratively construct complex networks, more challenging domains will likely require the orchestrated effort of both human *and* computational creators. Therefore extending BrainCrafter to support the users collaborative engineering efforts through a mixed-initiative approach is an important next goal. To allow seamlessly switching between interactive evolution and manual ANN design will require an ANN representations that (1) can be edited easily by the user on a local (e.g. individual neurons) and global level (e.g. overall network topography and topology), (2) is evolvable and compact, and (4) allows the system to produce meaningful suggestions based on the users input.

In this context we are building on the a generative encoding called *compositional pattern producing networks* (CPPNs; Stanley (2007)) and the recently introduced con-

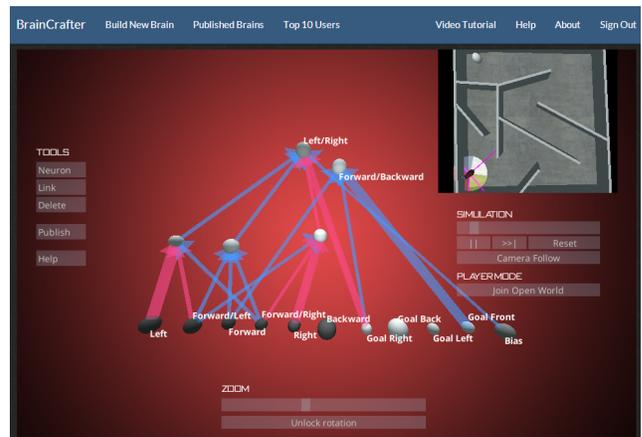


Figure 1: BrainCrafter Web Interface. Beta version online at: braincrafter.dk

cept of a CPPN-Compiler (Risi, 2013). CPPNs, which are based on principles of how natural organisms develop, allow the compact encoding of complex patterns, from two-dimensional images (Secretan et al., 2011) and three-dimensional forms (Clune and Lipson, 2011) to large-scale ANNs (Stanley et al., 2009; Clune et al., 2011; Risi and Stanley, 2012). The idea behind the CPPN-Compiler is to allow the user to directly compile a high-level description of the desired starting structure into the CPPN itself. We are currently exploring the benefits of this approach through the collaborative interactive evolution of images (Figure 2), in which the user can draw a vector image and annotate it with important regularities like symmetry. The CPPN-Compiler then compiles this high-level description into the CPPN itself. Because the compiled CPPN now directly embodies the annotated domain regularities (e.g. bilateral symmetric arms), the produced offspring show meaningful variations that nevertheless share common features. Based on the insight that the CPPNs can also produce large-scale ANNs (Stanley et al., 2009) and can be modified to create certain neural topologies (ES-HyperNEAT; Risi and Stanley (2012)), opens up the intriguing possibility of a neural CPPN-Compiler. Our current efforts focus on such a compiler that will form the backbone for our mixed-initiative BrainCrafter application.

The envisioned collaborative mixed-initiative system is depicted in Figure 3. The users can collaboratively construct neural networks and annotate them with regularities (e.g. symmetry, repetition, etc.), which are used by the computational creator to construct the internal CPPN model and in turn propose meaningful variations to the user. At any point in the process the human can revert to just doing selection or decide to directly edit the ANN, wherein phenotypic edits are directly compiled back into the genotype. The promise of the proposed system is that it could allow a variety of tasks to be solved by many people online within a

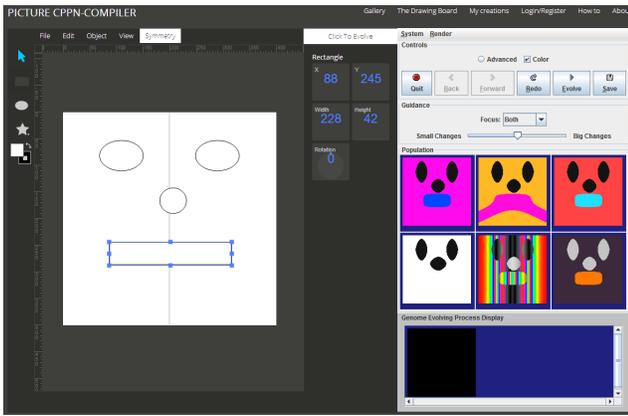


Figure 2: Picture CPPN-Compiler User Interface. Beta version online at: rasmustaarby.dk/thesis

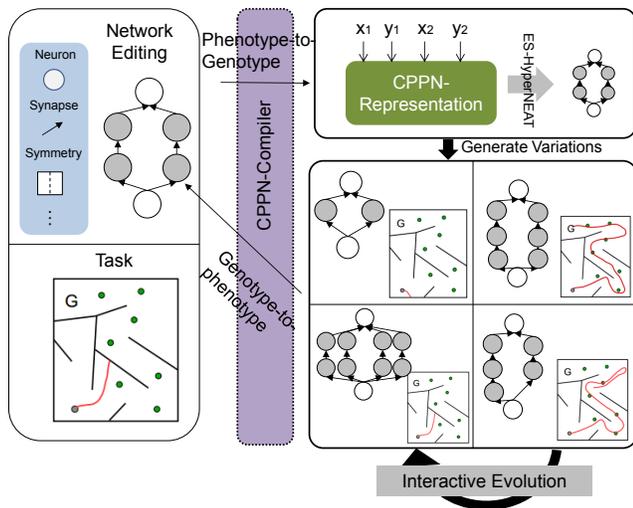


Figure 3: Mixed-initiative Neuroevolution Framework

mixed-initiative environment, which have heretofore proven to difficult. We expect this project will profoundly impact the fields of ANN research and to potentially also deepen our understanding of the way biological neural networks solve certain problems.

References

Bongard, J. C. and Hornby, G. S. (2013). Combining fitness-based search and user modeling in evolutionary robotics. In *Proceeding of the fifteenth annual conference on Genetic and evolutionary computation conference*, pages 159–166. ACM.

Clune, J. and Lipson, H. (2011). Evolving three-dimensional objects with a generative encoding inspired by developmental biology. In *Proceedings of the European Conference on Artificial Life*, See <http://EndlessForms.com>.

Clune, J., Stanley, K. O., Pennock, R. T., and Ofria, C. (2011). On the performance of indirect encoding across the continuum

of regularity. *Evolutionary Computation, IEEE Transactions on*, 15(3):346–367.

Floreano, D., Dürr, P., and Mattiussi, C. (2008). Neuroevolution: from architectures to learning. *Evol. Intelligence*, 1(1):47–62.

Lehman, J. and Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2):189–223.

Liapis, A., Shaker, N., and Smith, G. (2014). Mixed-initiative approaches. In Shaker, N., Togelius, J., and Nelson, M. J., editors, *Procedural Content Generation in Games: A Textbook and an Overview of Current Research*. Springer.

Liapis, A., Yannakakis, G., and Togelius, J. (2013). Sentient sketchbook: Computer-aided game level authoring. In *Proceedings of ACM Conference on Foundations of Digital Games (FDG)*.

Risi, S. (2013). A compiler for cppns: Transforming phenotypic descriptions into genotypic representations. In *2013 AAAI Fall Symposium Series*.

Risi, S. and Stanley, K. O. (2012). An enhanced hypercube-based encoding for evolving the placement, density, and connectivity of neurons. *Artificial Life*, 18(4):331–363.

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(3):373–403.

Stanley, K. O. (2007). Compositional pattern producing networks: A novel abstraction of development. *Genetic Programming and Evolvable Machines Special Issue on Developmental Systems*, 8(2):131–162.

Stanley, K. O., D’Ambrosio, D. B., and Gauci, J. (2009). A hypercube-based indirect encoding for evolving large-scale neural networks. *Artificial Life*, 15(2):185–212.

Stanley, K. O. and Miikkulainen, R. (2004). Competitive coevolution through evolutionary complexification. *JAIR*, 21:63–100.

Woolley, B. G. and Stanley, K. O. (2011). On the deleterious effects of a priori objectives on evolution and representation. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pages 957–964. ACM.

Woolley, B. G. and Stanley, K. O. (2014). A novel human-computer collaboration: Combining novelty search with interactive evolution. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2014)*, New York, NY, USA. ACM.

Yao, X. (1999). Evolving artificial neural networks. *Proceedings of the IEEE*, 87(9).