

Creature Academy: A System for Virtual Creature Evolution

Marcin L. Pilat and Christian Jacob

Abstract—In this paper, we present Creature Academy, a virtual laboratory that allows for the evolution of form and function within simulated physical 3D environments. Creature Academy can be used to explore evolutionary mechanisms, design, learning and other processes studied in artificial life simulations. Our system allows to perform hierarchical evolutionary experiments and ecosystem-inspired setups to investigate bodied creatures that interact, compete, adapt, and evolve. As a first proof of concept, we use Creature Academy to evolve morphologies and motion strategies of virtual creatures that walk and jump. We then present results that compare hierarchical evolution scenarios to generate creatures that excel in both walking and jumping, demonstrating how to evolve from creature specialists to generalists.

I. INTRODUCTION

Form plays an important role in biological life. Form enables function which, in turn, is an important factor for survival. Most life forms evolved some method of locomotion, dependent on the environment the organisms live in and their ability to adapt and control their morphologies through evolution. Wake argues that the study of morphologies (body forms) and development (the process by which morphologies are created) will provide clues to answer many open questions in biology [15].

In this paper, we describe Creature Academy, a virtual laboratory to study various aspects of the evolution of form and function. The system allows for exploration of co-evolution and large scale ecosystem evolution using a wide range of complex, 3D physics-based simulation environments. The system can be used to study evolution, design, learning, and other processes that are of interest in artificial life research.

The Creature Academy is composed of a set of independent simulation environments—*training grounds*—which can be used for isolated studies of specific evolutionary algorithms, environments, and/or fitness criteria (Fig. 1). The training grounds can be distributed to different computers and can support exchange of creature information. *Morphocosm*, a special type of shared training ground simulates a large-scale, distributed virtual ecosystem where creatures evolve and learn through interaction with each other and the environment.

We use the evolution of morphologies and controllers of virtual creatures (revisiting the work of Sims [10]) to demonstrate the features of the Creature Academy system. We evolve walking and jumping virtual creatures with diverse morphologies and behaviours. Furthermore, using Creature Academy we explore different control parameters in order to further evolve creatures using a combined walk-and-jump

fitness function. The experimental setup enables us to quickly generate and run new experiments to gain insights into the specific mechanisms behind virtual creature evolution.

In Section II, we outline previous and recent research in the area of virtual creature evolution. We present our Creature Academy system in Section III. In Section IV, we describe our implementation of Sims' creature model, the experimental setup, and the results of our exploration of virtual creature evolution. Finally, we conclude and offer ideas of future investigations in the area.

II. EVOLVING VIRTUAL CREATURES

Karl Sims in the mid 1990s [10], [11] pioneered the evolution of morphologies and controllers of virtual creatures. Sims used a genetic algorithm (running on a parallel machine) to co-evolve the morphologies and controllers of articulated virtual creatures made from 3D blocks (cuboids). His graph-based generative encoding naturally produced modularity and symmetry in the morphologies. Neural network controllers were embedded into the creature modules. The obtained results showed various behaviours for independently evaluated learning tasks, such as swimming, walking, jumping, and following a light source. This work provided key insights into the interplay between shape adaptability and physical constraints in biological life, which is still one of the main premises within the artificial life community.

Taylor and Massey [14] evolved swimming and crawling creatures following Sims' original work and demonstrated that it is now possible to reproduce Sims' work to a certain extent on current personal computers. Recently, with the availability of powerful open source physics and graphics engines, research into the area of morphological evolution has gained new momentum. The Framsticks system by Komosinski [3] evolved stick creatures composed of cylindrical body parts and joints. Control of each creature was accomplished using an evolved neural controller. The ability to simulate multiple creatures at the same time was available, however, the artificial world was only a reduced model of an ecosystem with emphasis on simulation of only a few individuals at a time. Komosinski has used the Framsticks system to evolve creatures showing simple swimming and walking behaviours.

Shim et al. [9], [8] have used a modified Sims-like creature representation to evolve flying and swimming morphologies using a GA-optimized controller of piecewise sinusoidal functions. Miconi and Channon have successfully evolved virtual creatures for the tasks of locomotion [4] and the co-evolutionary task of box-grabbing [5]. In contrast to Sims' work, their model used a tree-like genotype (with self-loops allowed) and standard McCulloch-Pitts neurons. Chaumont

et al. [1] have used Sims’ creature model to evolve walking and block-throwing (catapults) creatures. Their model used a subset of Sims’ neurons and body modules specifically modified to fit the block-throwing task. Their work also discusses prevention of exploitative behaviour and the importance of carefully crafting the fitness function.

III. CREATURE ACADEMY

Creature Academy revisits the evolution of 3D morphologies and controllers, but goes beyond the previous work cited above. The system allows for quick creation of virtual worlds—represented as training grounds—that are used for experimenting with various evolutionary approaches (genetic algorithms, neural network learning algorithms), scenarios (genotype representations, fitness criteria, phenotypical models), and environments (landscapes, aquatics, gravity). The idea is analogous to athletes training for competitions by using different training tasks in specific environments in order to improve their skills.

Creature Academy then allows for the integration of and interaction among training grounds through a shared “ecology” (Fig. 1). The level of integration depends on the purpose of a specific experimental setup. No training ground interaction equates to an evolutionary system where various experiments are run independently from each other (see Parallel in Fig. 1). These training grounds can be thought of as creature generators that continuously produce fit creatures to be used in other, usually more complex, training grounds. Constant exchange of creatures between training grounds increases the level of interaction in the system (see Interactive in Fig. 1). This realizes an island model of an evolutionary algorithm where sub populations evolve separately, but with intermittent exchanges of individuals.

Training grounds of the Creature Academy can also be structured into a tiered (hierarchical) system akin to rounds in a sports tournament (see Tiered in Fig. 1). Upon finishing an evolutionary run, training grounds can send their best performing creatures to a higher-tier training ground where the creatures are further improved via evolution, learning, etc. Once all training grounds have completed, the remaining best individuals can be thought of as having successfully passed all the previous training grounds. The challenge here is to ensure that previously learned behaviours are not forgotten. Due to the higher complexity of this setup, it is up to the experimenter to choose the training ground structure that performs meaningful experiments during the whole life cycle of the Creature Academy.

A. Training Grounds

Each Training Ground can be thought of as a separate experiment in virtual creature evolution. In its simplest form, it enables experimenters to create pools of genetic material yielding creatures specialized for a particular task. In a more complex scenario, it is part of an interconnected evolutionary training system that exchanges virtual creatures between the various training grounds. A training ground can be

divided into three components: environment, population, and algorithm (Fig. 1).

1) *Training Environment*: The training environment is composed of a simulated physical world (including 3D landscapes and inanimate objects) and physics engine parameters (such as gravity, viscosity, friction and dynamical forces due to collisions). The environment also specifies interaction rules used by the virtual creatures to interact with the environment and among themselves. Such interaction rules add more intricate fitness scenarios and allow for the evolution of more complex tasks and behaviours. For visualization, the environment also specifies the meshes, textures, and other graphics settings used to render objects.

2) *Population Container*: Virtual creatures are stored inside a population (or a container in a more generic terminology). Each virtual creature has a genotype representation and a corresponding simulated physical phenotype. Selection of a genotype encoding is up to the user, but generative encodings, which incorporate the idea of development into the genotype, are preferred since they can provide greater scalability through the re-use of self-similar and hierarchical structures [12]. The user must also decide on the phenotypic building blocks used to make up the final virtual creature.

3) *Adaptation Algorithm*: A typical adaptation algorithm is an evolutionary algorithm (e.g., a genetic algorithm) that evolves the creatures using specific fitness criteria. The fitness function is considered to be part of the algorithm and can vary between different training grounds. When studying a learning task, such as neural network learning for virtual creatures, a learning algorithm can be used. The type of algorithm used in a training ground is up to the experimenter, which provides the freedom to study various aspects of artificial life simulations.

B. Morphocosm

Taylor [13] argues that “a fruitful, and as yet largely unexplored, avenue for artificial life research lies in modelling organisms (specifically, phenotypes) and environment as a single dynamical system.” Morphocosm is a special type of training ground that allows for the simulation of a large-scale physically simulated virtual ecology of creatures. It can be used as an efficient experimental environment to study complex topics in evolution that require ecosystem-scale simulations.

As an artificial life system, Morphocosm contains a population of virtual creatures interacting with each other and their environment. Interaction rules (including sensing abilities) must be specified by the user such that the system provides enough evolutionary selection pressures. Each creature has a finite life cycle. In contrast to a standard GA, the mating of creatures, and hence their recombination of genotypes, is triggered by spatial proximity.

IV. SIMS REVISITED

The virtual creature model presented by Sims [10] is an excellent benchmark for the testing of various settings and ideas related to virtual creatures and virtual ecologies. With

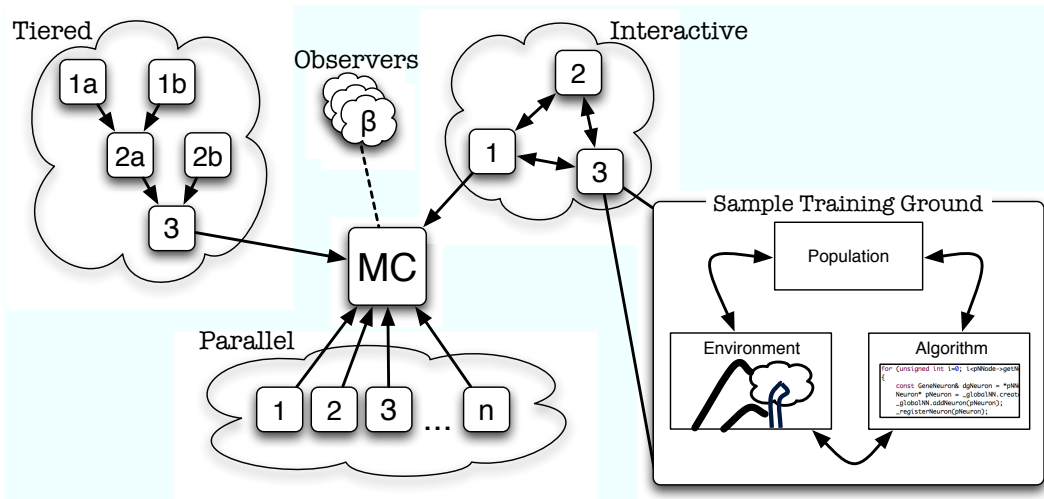


Fig. 1. Schematic view of the Creature Academy system. Training grounds are shown as boxes with arrows showing the flow of creatures between training grounds. The central Morphocosm (MC) creature ecology receives creatures from training grounds with different exploration schemes: tiered, interactive, and parallel. A set of observer modules visualizes the MC via different methods, such as cameras, parameter probes, etc. The components of a sample training ground are shown on the right.




Fig. 2. Morphologies evolving from complex to simple. Examples showing the changes to morphologies of a creature due to evolution in our presented 2-tier Creature Academy. The creature (a) is contained in the initial population for tier-2 (arrived from tier-1 jumper experiments). Creatures (b) and (c) evolve from two separate runs of the tier-2 training ground using the walk-and-jump fitness criteria. The genotype graph, phenotype graph, and a wireframe of each creature are shown. Genotype and phenotype graphs generated with Graphviz[2].

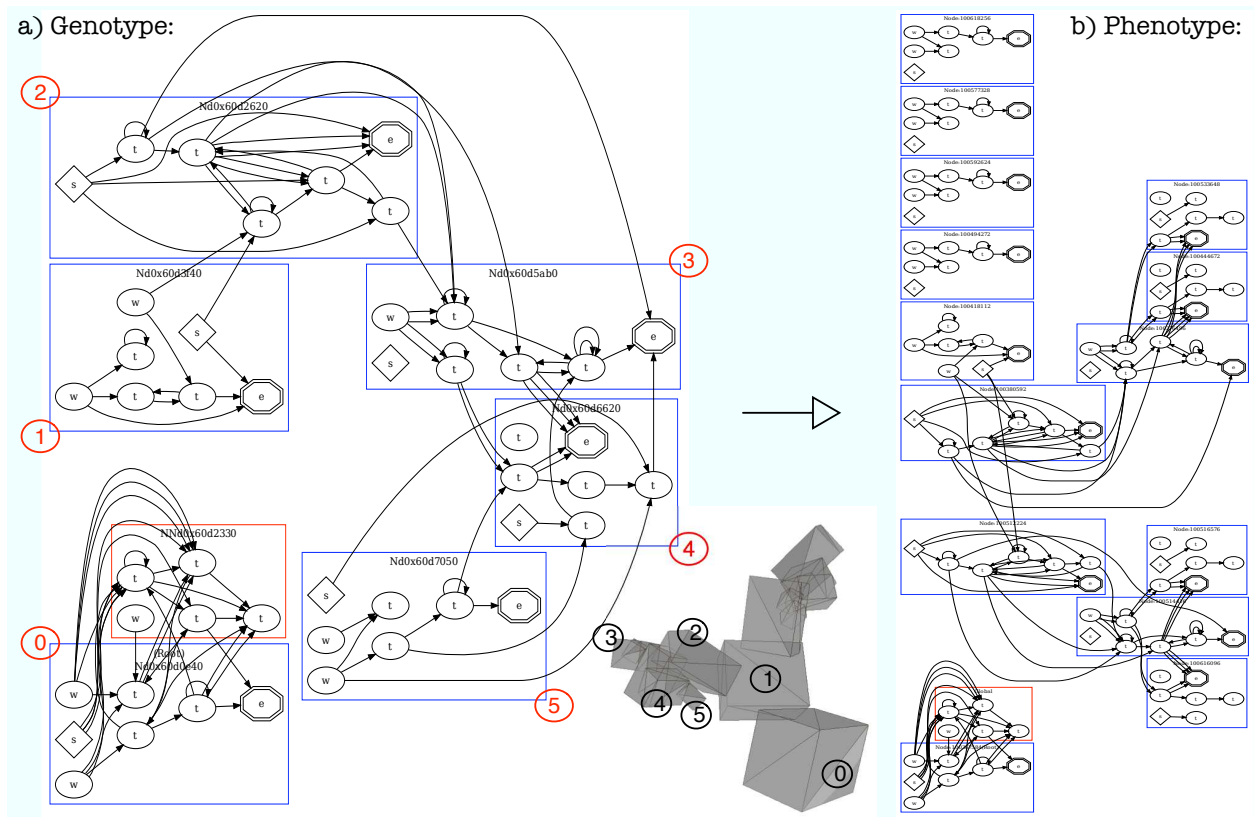


Fig. 3. Example of an evolved control network: The neural network architecture of the creature from Fig. 2(a) is illustrated. The NN genotype (a) is shown together with the corresponding phenotype (b). The numbered boxes correspond to the nodes in the shown morphology. Sensor neurons are shown with diamond shape, effector neurons with double octagon shape, and hidden neurons are represented as circles. Parameters of the neurons and links are omitted to simplify the diagram. Diagrams generated with Graphviz[2].

the recently regained interest in Sims' work, researchers that use the model are able to compare their findings.

A. Model

We model the phenotype of a virtual creature as a rooted tree. Nodes of the tree describe body parts of the creature, whereas the links represent connections between body parts. The morphological genotype, however, has the form of a directed graph with possible cycles and self loops (Fig. 2). The graph representation is a generative encoding which specifies instructions on how the phenotype is built [12]. This encoding allows for the reuse of components and development of recursive structures. Figure 2 illustrates some examples of the genotype-to-phenotype mapping for morphologies.

A genotype body node stores information used to create the corresponding body part. The body type (a cuboid/block shape) has specific dimensions (width, height, depth), a recursion limit (how many times the node will be used to create phenotypic body parts while in a recursive cycle), and a local neural network, which is part of the neural controller. For this work, we use hinge joint types which provide each body part with one degree of freedom with respect to its parent.

Each genotype graph link specifies the position of the contact point on the parent body, orientation of the child node

relative to the parent, and a parameter which specifies how the size of the child is scaled. Scale is relative to the parent scale and can be used to generate increasingly smaller body parts. Negative scale denotes a reflection about the main axis of the parent node and allows for the generation of symmetric arrangements. A phenotypical joint is stored in the child node and denotes the connection of the child to its parent.

The control of a virtual creature is accomplished via a recurrent neural network (NN) made up of neurons and neural links specifying the directed graph NN topology. Each body part contains a list of sensor, effector, and hidden neurons. The network is encoded into the genotype of the creature and is copied into the phenotypical body parts during development. Connections between neurons in different local networks are only allowed between neighbouring body parts. A global neural network node, which is not associated with a physical body, stores neurons that can be connected to any valid neuron in the NN. The global node is used to facilitate communication between the local neural networks and provides a form of centralized control. Figure 3 shows an example of an evolved NN controller.

Sims used a diverse list of neuron types in his work. However, it is not clear from the published results if such a wide variety of neurons is necessary for the evolution of interesting creature behaviours. Miconi et al. [4] have

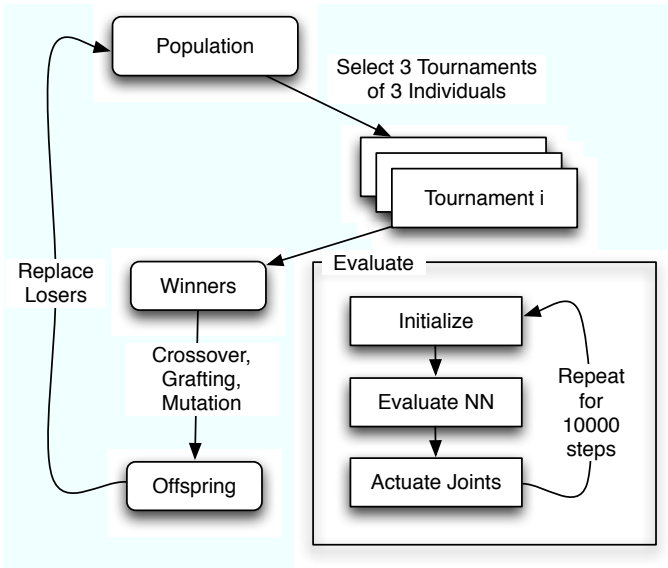


Fig. 4. Schematic view of our parallel tournament-selection genetic algorithm. The creature evaluation loop is shown in the lower right.

demonstrated that simple McCulloch-Pitts neurons, usually when coupled with a periodic function, are enough for the evolution of a wide range of virtual creature behaviours. Chaumont et al. [1] have used a subset of the neurons applied by Sims in order to demonstrate similar results. For this work, we have used two types of hidden neurons: simple McCulloch-Pitts neurons, with a hyperbolic tangent transfer function, and sinusoidal periodic neurons. The return value of each neuron is in the range $[-1, 1]$ to standardize the output/input values of all neurons in the system.

Each body part contains a sensor neuron and an effector neuron. A sensor neuron stores the values of its body’s joint angle, with respect to its axis, scaled to the $[-1, 1]$ range. An effector neuron, which uses the hyperbolic tangent transfer function on its input, represents the output of the local neural network and is used to calculate the desired angular velocity along the hinge joint axis. The input degrees of hidden neurons are controlled by a configuration parameter.

B. Implementation

We use a standard genetic algorithm with tournament selection (of size 3) on a population of creature genotypes. Physical evaluation of the resulting creature phenotypes is performed simultaneously for three tournaments. This form of a parallel GA is able to process multiple small tournaments (e.g. among 9 creatures) at roughly the same time as processing one creature (due to the nature of the physics engine[7]). It also allows for creature interaction as part of their evaluation. For the presented work, we have disabled creature interactions and only allowed collisions with the ground.

Our evolutionary algorithm is implemented as a state machine (Fig. 4). Evaluation of each creature is performed in steps, where each step consists of: two passes through the neural network, actuation of the joints using the output values

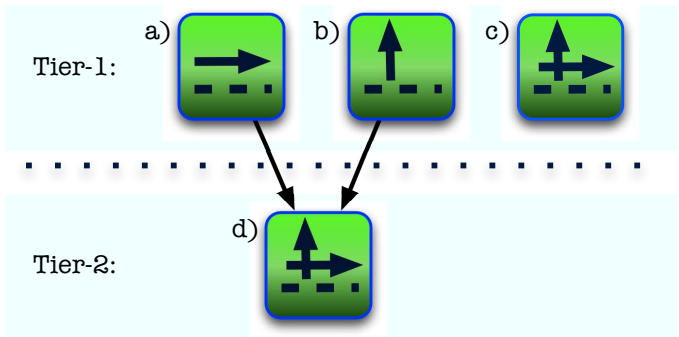


Fig. 5. Two-tier Creature Academy experimental setup used for the presented experiment. Three Tier-1 training grounds evolved virtual creatures from random populations for the (a) walking, (b) jumping, and (c) walk-and-jump fitness criteria. The tier-2 training ground (d) selects initial individuals from results of the tier-1 walking and jumping runs to further evolve creatures using the walk-and-jump fitness.

computed by the neural network, and interaction stepping of the physics engine. The integrity of each simulated creature was constantly checked, using joint separation errors, to prevent joint explosions by disabling invalid creatures.

For our experimental setup, we created a two-tier Creature Academy with the intent to evolve virtual creature locomotion (Fig 5). Tier-1 training grounds used random initial populations of creature genotypes, whereas the tier-2 training grounds used initial populations consisting of individuals evolved through the tier-1 experiments. This allowed us to try various evolutionary parameter values on pre-evolved populations.

We looked at three fitness criteria (Fig 6): walking, jumping, and a combination of both (denoted as walk-and-jump). The walking fitness, Ω_{walk} , is specified as the $x - y$ distance travelled by a creature over a specified number of time steps. The jumping fitness, Ω_{jump} , is defined as the maximum height difference of the main body block during creature evaluation, with the initial position as the baseline height. The walk-and-jump fitness criterion awards both walking distance and jumping height: $\Omega_{wj} = \Omega_{walk} * \Omega_{jump}$.

Initially, we randomly generate a population of creature genotypes which are developed into physical phenotypes. Invalid phenotypes (those that are created outside of the specified parameters) are flagged as invalid creatures and are not evaluated. We then randomly choose creatures, with no replacement, for each tournament set consisting of 3 tournaments of 3 creatures and evaluate them at the same time. For each tournament, we choose the two fittest individuals, mate them to create a child offspring and replace the worst performing tournament member with the offspring.

As with Sims’ work, mating consists of three genetic operators: mutation, crossover, and grafting. Mutation is applied to the winner of a tournament to form an offspring. Each parameter of the creature genotype is mutated with a specified low probability. New nodes, links, neural nodes, and neural links are also created, or old ones are removed with the same probability. Both the crossover and grafting operators are applied on the best and second best individual of a




Fig. 6. Schematic diagrams of the fitness calculations for (a) walking (Ω_{walk}), (b) jumping (Ω_{jump}), and (c) walking-and-jumping (Ω_{wj}).

tournament to produce one child. Crossover is implemented as a single point GA-style crossover treating each genotype as an indexed list of body nodes. Similarly to GP crossover, the grafting operator points a random link of a parent to a node of the other parent.

Every n th tournament, as defined in a configuration file, the whole population of genotypes is stored, together with statistical data, in XML format. Thus, we are able to restore any generation, view it, or use it as an initial population of a new simulation run. During the run, various statistics are stored (tournament fitness, population fitness, population diversity measures, etc.) and can be easily visualized. We also provide the option of converting a creature’s representation (genotype or phenotype with various levels of detail) into a format which can be used to generate network diagrams via Graphviz [2] (cf. Fig. 2 and 3). These diagrams have been instrumental in the testing of our implementations.

C. Results

The initial sets of experiments to evolve walking, jumping, and walk-and-jump behaviours were run on tier-1 training grounds as described in the previous section (Fig 5(a)-(c)). Table I presents experimental parameter values used in our experiments. Typical experimental runs lasted from 3 to 4 hours of full-time CPU use with some runs of complex creatures lasting as much as 7 hours¹.

TABLE I
EXPERIMENTAL PARAMETERS USED FOR EVOLUTION OF VIRTUAL CREATURES IN TIER-1 AND TIER-2 TRAINING GROUNDS.

Parameter	Value (tier-1)	Value (tier-2)
Algorithm	parallel tournament (3x3) GA	parallel tournament (3x3) GA
Population Size	300 (random)	200 (from tier-1)
Mutation p.	40%	40%, 0%, 100%
Crossover p.	30%	30%, 100%, 0%
Grafting p.	30%	30%, 0%, 0%
Tournaments	5000	3000
Fitness Criteria	a) Walking (Ω_{walk}) b) Jumping (Ω_{jump}) c) Walk-and-Jump (Ω_{wj})	Walk-and-Jump (Ω_{wj})

Experiments for all three fitness criteria produced a large variety of creature morphologies (Fig. 7). Creatures evolved

¹Time based on a Mac Mini with 1.66Ghz Intel Core Duo processor and 1GB of RAM.

under different fitness criteria were sometimes similar in structure and in behaviour. However, many of the observed morphologies and behaviours evolved unique solutions to a fitness criterion. Some runs developed behaviours that can be viewed as exploiting the simulated physical environment, however, most exploits have been eliminated due to the evaluation procedure and creature validity checks.

The control neural networks of evolved creatures showed great complexity (Fig. 3). The function of each network is unclear by visual inspection of its topology. Detailed examination of the resulting NNs is necessary in order to understand the range of motions and behaviours seen in the evolved creatures.

1) *Tier-1 Evolved Walkers*: Several good behaviours evolved for the walking task using an initial population of random creatures (Fig. 7). Simple 2- or 3-block phenotypes usually employed a spinning or swinging strategy. *Spinners* quickly rotated their body parts, which were usually similar in size, to gain momentum that would make the creatures move along the surface with a jumping or rolling motion. *Swingers* used a smaller body part to swing it back and forth without touching the ground, thus allowing them to move slightly off the ground on one side creating a sliding motion. More complex morphologies, with five or more blocks, employed a pushing or lifting strategy. *Pushers* used one or more body parts to push themselves along the surface while *lifters* moved with a more complex motion requiring the (sometimes synchronized) movement of multiple connected blocks we call limbs.

2) *Tier-1 Evolved Jumpers*: Morphologies of the creatures evolved using the jumping task were noticeably more complex in the number of parts, length of limbs, and use of symmetric body parts. Two main behaviours observed were spinning and lifting. *Spinners* were usually simpler in their morphology and similar in behaviour compared to the ones evolved for the walking task but with a more noticeable vertical jumping component. *Lifters* were complex with multiple limbs and gained high fitness by pushing their main body part high off the ground by extending their limb segments. Some creatures resembled naturally occurring morphologies and movements, such as long worms, crabs with symmetric limbs, or turtles pushing their large bodies with small legs over the ground.

3) *Tier-1 Evolved Walkers and Jumpers*: Runs of tier-1 training grounds with the walk-and-jump fitness function




Fig. 7. A selection of evolved virtual creatures from our experimental runs. Rows 1 and 2 show examples from the tier-1 runs to reward walking or jumping. Row 3 displays creatures from the tier-2 runs using the walk-and-jump fitness. Interesting miscellaneous creature morphologies are shown in the last row. The highlighted creatures show variations of the “Dancer” creature from Figures 2 and 3.

showed similar morphologies and behaviours to the ones for walkers and jumpers. Many of the evolved creatures had worm-like morphologies consisting of segments of body parts. The most often observed behaviour was spinning. Best-of-run fitness varied on each separate run, averaging 64,901 at tournament 5000, as shown in Fig. 8(b).

Apart from the standard (0.4*m*/0.3*c*/0.3*g*) genetic operator frequencies, we have also run tier-1 experiments using high mutation and high crossover rates. Very high crossover rate (0.2*m*/0.8*c*/0.0*g*) produced more complex individuals than the standard, with best-of-run fitness values averaging 70,382 at tournament 5000. Using only crossover (0.0*m*/1.0*c*/0.0*g*) resulted in more complex individuals with many body parts but with a lower best-of-run fitness averaging 35,072 at tournament 5000. Various different behaviours were seen including morphologies consisting of worm-like segments of body parts showing a lifting and/or spinning strategy.

In runs with only mutation (1.0*m*/0.0*c*/0.0*g*), the resulting morphologies were simpler than those from standard runs.

All of the evolved creatures had simple two or three body part morphologies employing some form of spinning behaviour. Many of the small spinners moved by launching themselves into the air thus generating very high fitness values. The best-of-run fitness averaged 164,150 at tournament 5000. Although much higher fitness was achieved by using a very high mutation rate, the morphological and behavioural diversity was visibly lower compared to using standard genetic operator probabilities.

4) *Tier-2 Evolved Walkers and Jumpers:* Tier-2 experiments (Fig 5(d)) were performed from an initial population of 200 individuals selected from the fittest creatures evolved through tier-1 walking or jumping training grounds (Fig 5(a) and (b)). The most fit individual from the initial population had a walk-and-jump fitness of 51,756. We have set up various tier-2 experiments by varying the probabilities of the genetic operators, probabilities and effect of mutations, and by addition of random individuals to the initial crafted population. The results were compared using fitness, morphology, and behaviour to the results of the initial population and of

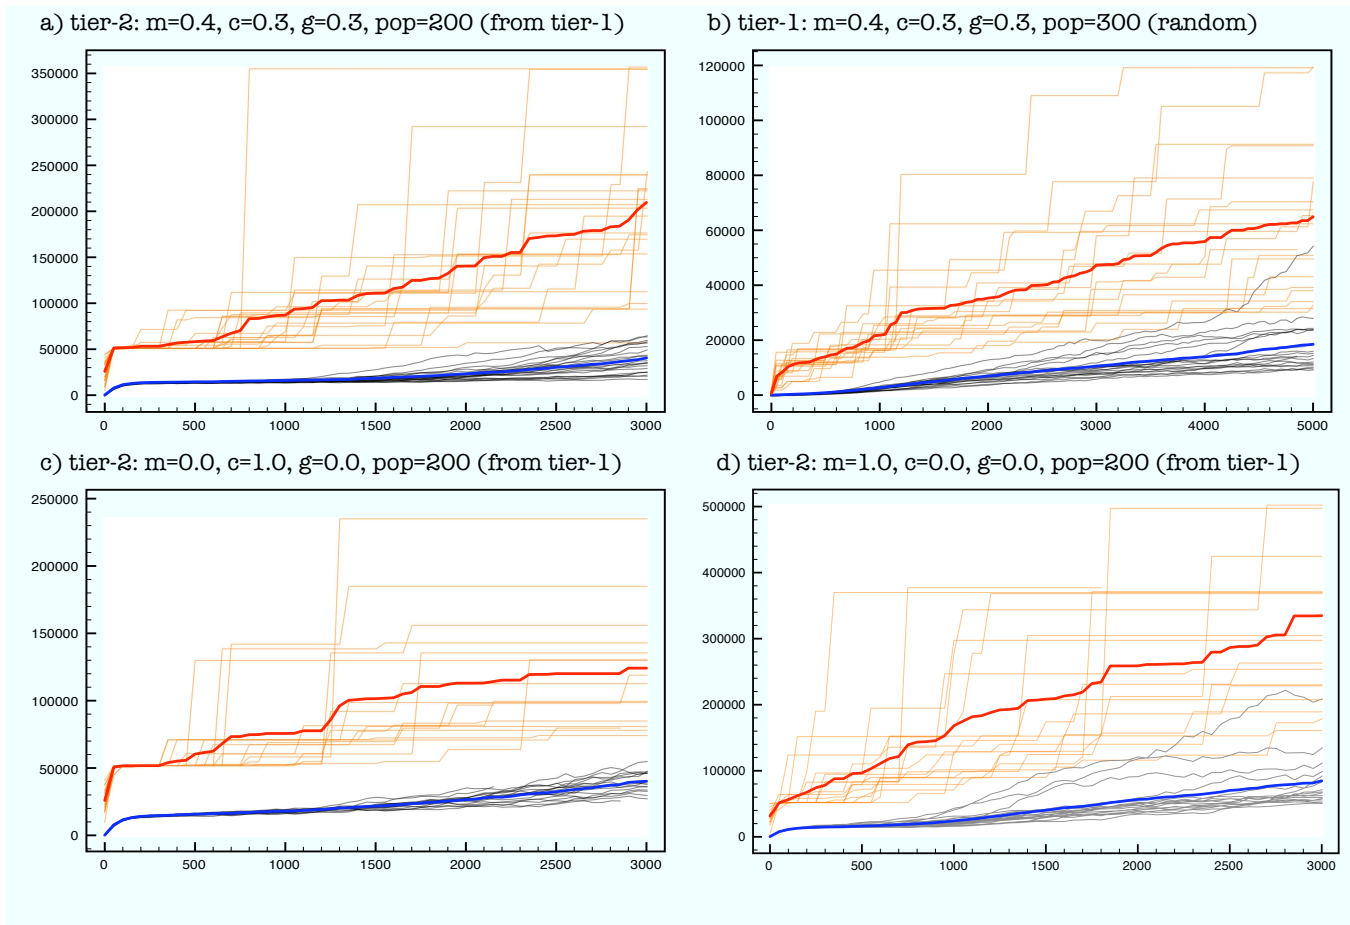


Fig. 8. Fitness plots showing maximum population fitness (top curves) and population average fitness (bottom curves) over several runs using the walk-and-jump fitness function. Plots (a), (c) and (d) are from tier-2 of the presented creature evolution Creature Academy with the initial population composed of the best individuals from 100 tier-1 runs. Plot (b) is from tier-1 runs starting from initial random populations. Averages are shown in thicker lines. Note that the creatures evolved through the incremental tier-2 experiments display two to three times the fitness, compared to the tier-1 results.

tier-1 runs with the walk-and-jump fitness function.

Experiments with the standard genetic operator probabilities ($0.4m/0.3c/0.3g$) produced very fit individuals compared to the initial population (Fig. 8(a)), with best-of-run fitness average of 209,682 at tournament 3000. The fitness values were also much higher than those of tier-1 runs using the walk-and-jump fitness function (compare Fig. 8(a) and (b)). The evolved morphologies were mostly simple with 2 or 3 bodies showing a spinning behaviour. Similar morphologies and behaviours were observed with a higher crossover rate and no grafting ($0.2m/0.8c/0.0g$), with best-of-run fitness average of 181,167 at tournament 3000. Only 20% of the initial population individuals had simple 2- or 3-body morphologies.

With only crossover ($0.0m/1.0c/0.0g$), best-of-run fitness average was 124,165 at tournament 3000 (Fig 8(c)). Hence, crossover produced lower fit individuals compared to using mutation. However, the resulting morphologies were quite complex and more diverse than those from mutation runs. They included simpler morphologies such as 2-block spinners and variations on the highest fit initial individual (as shown in Fig. 2). These variations are of importance since

they provide us with the ability to compare the original creature of the initial population to its variants in evolved results of separate tier-2 runs. Figure 7 shows an example of such variation on the highest fit initial individual named “Dancer”.

In experiments with only the mutation operator ($1.0m/0.0c/0.0g$), the best-of-run fitness averaged 334,767 at tournament 3000 (Fig 8(d)). These fitness values were significantly higher than runs with other genetic operator probabilities. However, all the evolved creatures had simple 2-body morphologies and spinning movement. Experiments with high grafting probability generated little improvement to the fitness of the initial population and to the morphologies of the creatures. The poor performance of grafting might be due to the creatures reaching maximum phenotypical size while using the operator. This seems evident from the lower overall population fitness compared to runs with standard operator values (since failed grafting produces random child individuals).

These results indicate that a balance of genetic operator probabilities needs to be achieved in order to evolve highly fit but still distinct morphologies and corresponding behaviours.

High mutation rates lower the morphological and behavioural diversity of the evolved creatures, while crossover generates good diversity at a cost of lower fitness.

Alternatively, we have run experiments with varying effects of mutation, which produced slightly fitter individuals in some cases, but more experimentation is required to find the optimal parameter values. Addition of random individuals to the initial population provided some improvement to the fitness of resulting creatures by introducing diversity of genes into the mating pool. However, more testing is required to verify the usefulness of this approach.

V. CONCLUSION

In this paper, we have introduced the Creature Academy experimentation system for evolution of form and function. The system can be used to run large sets of independent experiments as well as complex hierarchical tiered and interacting experimental platforms. The ease of use and extendability makes this system an attractive experimentation solution for projects ranging from evolutionary design to virtual ecosystems.

Creature Academy is an on-going development project with an extensive list of future additions and modifications. The addition of Morphocosm-style training grounds and process communication will further extend the range of possibilities of this evolutionary system. Since the project is open-source, we welcome collaboration and interest in development, testing, and experimentation using Creature Academy.

The open-source Creature Academy system is written in platform-independent C++. It makes use of other open source projects such as OGRE [6] for graphical output, and ODE [7] for physical simulation and realism. All the experiments in this study have been performed on a cluster of 50 Apple Mac minis, each with a 1.66Ghz Intel Core Duo processor and 1GB of RAM. More information about the Creature Academy, including images, videos, and experimental data, can be found at the project website: <http://www.morphids.com>.

In this work, we were able to reproduce Sims' results for walking and jumping creatures using similar experimental settings. The evolved virtual creatures showed a large range of morphologies and behaviours. Furthermore, we created a 2-tier Creature Academy for evolution of virtual creatures using a combined walk-and-jump fitness function. This setup allowed us to pre-evolve creatures using walking and jumping fitness criteria and later further evolve them with different experimental parameter values.

By using a 2-tier system, we were able to produce fitter individuals than in a comparable 1-tier system akin to a regular GA run. This performance increase was due to varying the genetic operator probabilities and is similar to performance gains in GA systems with dynamic operator frequencies. However, our approach allows the experimenter not only to modify frequencies but also to change other aspects of the experimental runs such as initial population

makeup, different experimental environments, and changing fitness criteria.

The presented results indicate a number of questions that we would like to investigate further. Our experiments have generated hundreds of gigabytes of creature data. We would like to visualize this data through population- and run-based visualization tools in order to gain an insight into the larger scale evolutionary system. It is also important to study the details of the complex neural networks generated for each creature and how these networks change over evolutionary time.

We will also experiment with different fitness criteria in a more complex Creature Academy setup that includes the Morphocosm virtual ecology. Ecological scale models and simulations would help to answer difficult questions and open problems in evolutionary biology, developmental biology, and artificial life. Such simulations will allow us to formulate and test the models of the mechanisms responsible for natural evolution of form and function and how they can be applied to areas of computer science, robotics, and design.

REFERENCES

- [1] N. Chaumont, R. Egli and C. Adami. "Evolving Virtual Creatures and Catapults". *Artificial Life*, vol. 13, no. 2, pp. 139–157, 2007.
- [2] Graph Visualization Software (Graphviz). <http://www.graphviz.org>.
- [3] M. Komosinski, "The World of Framsticks: Simulation, Evolution, Interaction". *2nd International Conference on Virtual Worlds (VW2000)*, Paris, France, pp. 214–224, July 2000.
- [4] T. Miconi and A. D. Channon, "A Virtual Creatures Model for Studies in Artificial Evolution". *2005 IEEE Congress on Evolutionary Computation (CEC 2005)*, Edinburgh, UK, pp. 565–572, September 2005.
- [5] T. Miconi and A. D. Channon, "An Improved System for Artificial Creatures Evolution". *Tenth International Conference on the Simulation and Synthesis of Living Systems (ALife X)*, Bloomington, Indiana, USA, pp. 255–261, 2006.
- [6] Object-Oriented Graphics Rendering Engine (OGRE). <http://www.ogre3d.org>.
- [7] Open Dynamics Engine (ODE). <http://www.ode.org>.
- [8] Y.-S. Shim and C.-H. Kim, "Evolving Physically Simulated Flying Creatures for Efficient Cruising", *Artificial Life*, vol. 12, no. 4, pp. 561–591, 2006.
- [9] Y.-S. Shim, S.-Y. Shin and C.-H. Kim, "Two-Step Evolution Process for Path-Following Virtual Creatures". *Computer Animations and Social Agents 2004*, Geneva, Switzerland, pp. 85–93, July 2004.
- [10] K. Sims. "Evolving Virtual Creatures". *21st Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH 94)*, New York, NY, USA, pp.15–22, July 1994.
- [11] K. Sims. "Evolving 3d Morphology and Behavior by Competition". *Fourth International Workshop on the Synthesis and Simulation of Living Systems*, Cambridge, Massachusetts, USA, pp. 28–39, 1994.
- [12] K. O. Stanley and R. Miikkulainen. "A Taxonomy for Artificial Embryogeny". *Artificial Life*, vol. 9, no. 2, pp. 93–130, 2003.
- [13] T. Taylor. "Redrawing the Boundary between Organism and Environment". *Ninth International Conference on the Simulation and Synthesis of Living Systems (ALife IX)*, Boston, Massachusetts, USA, pp. 268–273, September 2004.
- [14] T. Taylor and C. Massey. "Recent Developments in the Evolution of Morphologies and Controllers for Physically Simulated Creatures". *Artificial Life*, vol. 7, no. 1, pp. 77–87, 2000.
- [15] M. H. Wake. "Bodies and Body Plans, and How They Came to Be". In *A New Century of Biology*, edited by W. J. Kress and G. W. Barrett, pp. 29–52, Smithsonian Books, 2001.