

Evolvability signatures of generative encodings: beyond standard performance benchmarks

Danesh Tarapore and Jean-Baptiste Mouret*

Evolutionary robotics is a promising approach to autonomously synthesize machines with abilities that resemble those of animals, but the field suffers from a lack of strong foundations. In particular, evolutionary systems are currently assessed solely by the fitness score their evolved artifacts can achieve for a specific task, whereas such fitness-based comparisons provide limited insights about how the same system would evaluate on different tasks, and its adaptive capabilities to respond to changes in fitness (e.g., from damages to the machine, or in new situations). To counter these limitations, we introduce the concept of “evolvability signatures”, which picture the post-mutation statistical distribution of both behavior diversity (how different are the robot behaviors after a mutation?) and fitness values (how different is the fitness after a mutation?). We tested the relevance of this concept by evolving controllers for hexapod robot locomotion using five different genotype-to-phenotype mappings (direct encoding, generative encoding of open-loop and closed-loop central pattern generators, generative encoding of neural networks, and single-unit pattern generators (SUPG)). We observed a *predictive* relationship between the evolvability signature of each encoding and the number of generations required by hexapods to adapt from incurred damages. Our study also reveals that, across the five investigated encodings, the SUPG scheme achieved the best evolvability signature, and was always foremost in recovering an effective gait following robot damages. Overall, our evolvability signatures neatly complement existing task-performance benchmarks, and pave the way for stronger foundations for research in evolutionary robotics.

1 Introduction

Evolutionary robotics (ER) is a promising approach to achieve one of the prominent long-term goals of artificial intelligence research: creating machines with the adaptive and cognitive abilities of animals. Since the eighties, the ER field has made amazing progress to both design sophisticated artifacts and to endow machines with impressive adaptive abilities. For instance, it allows for the automated construction of modular, three-dimensional, physically locomoting robots, (Hornby et al., 2003), to synthesize neural networks to control robot behaviors (e.g., Lehman and Stanley (2011b); Mouret and Doncieux (2012); Kubota (2005); Santos et al. (2001)), and discover a multitude of walking gaits for mul-

tilegged robots following unforeseen mechanical damages (Bongard et al., 2006; Koos et al., 2013; Cully and Mouret, 2013). However, even the most advanced evolved artifacts are still far behind the state of the art in mainstream robotics (Kajita and Espiau, 2008; Bongard, 2013): conventionally engineered robots are capable of operating successfully in a wide variety of indoor and outdoor environments (e.g., locomotion with the BigDog quadruped robot, Raibert et al. (2008)), whereas the best evolved robots are still only capable of simplistic behaviors (e.g., walking in a straight line on a flat terrain, or avoiding obstacles in an enclosed indoor arena). To progress further, ER needs to go beyond the mere “stamp collecting” of proofs of concept, evident in the infancy of many scientific fields (Hayes, 2004), and build strong theoretical and methodological foundations for future research. The objective of the present study is to move in this direction.

In most ER studies, fitness comparison is the main instrument used to compare different evolutionary systems and assess their progress. Such a benchmark-based comparative approach has led to incremental improvements in the robot’s performance in specific tasks (e.g., for multilegged robot locomotion, the inclusion of evolved gaits on the commercial release of Sony’s AIBO Hornby et al. (2005); Valsalam and Miikkulainen (2008), and the progressive improvements in walking speed of the QuadraBot Yosinski et al. (2011); Lee et al. (2013)), and is sufficient if excelling at the given function is the ultimate goal for the robot. Nonetheless, if the evaluated task is treated as a tool to compare different evolutionary systems, and as a stepping stone to harder problems, then a mere comparison of performance does not suffice. This is because such a methodology of comparison only provides a very limited amount of information about the behavior of the system. In particular, it does not provide any insights on, (i) how efficiently does the evolutionary process explore the search space (e.g., can it also lead to solutions for other similar tasks, or is it biased to the type of solutions useful only for a very specific task?), and (ii) what capabilities are provided to the evolved population to respond to novel situations (e.g., an unexpected breakage of the multilegged robot’s limbs, or changes to its weight distribution). Furthermore, while adaptive evolutionary systems utilize a variety of population-diversity maintenance methods to operate in changing environments Jin and Branke (2005), they are mostly concerned with numerical optimization problems (e.g., Morrison and De Jong (1999)), and constrained to fitness-based indices to evaluate available approaches Weicker (2002). In summary, there is a need for additional metrics when comparing evolutionary systems, especially if one is interested in the adaptive abilities provided by evolution.

*Danesh Tarapore and Jean-Baptiste Mouret are with the ISIR, Université Pierre et Marie Curie-Paris 6, CNRS UMR 7222, F-75252, Paris Cedex 05, France. Contact: danesh.tarapore@gmail.com

In benchmark-based comparative approaches, the fitness value in an evolutionary system is often used as a proxy for the *evolvability* provided by the system Gruau (1994); Komosiński and Rotaru-Varga (2001); Hornby et al. (2003); Clune et al. (2009a) — the capacity of the evolved population to rapidly adapt to novel environments Hu and Banzhaf (2010). However, such a fitness-based proxy provides little information on the potential of the evolutionary system to generate novel phenotypes, and consequently rapidly adapt to new, untested environments. While fitness landscape models can provide interesting insights on search difficulty in the Genotype-to-Fitness map (Provine, 1989; Wright, 1932), the models 3D landscape can be deceptive when analyzing highly multidimensional genotypes (Kaplan, 2008; Gavrillets, 1997; McCandlish, 2011). Additionally, in NK fitness landscape models (Kauffman, 1993; Tomassini et al., 2008), the value of K that controls the degree of epistasis is not easily transferable to more complex and open-ended Genotype-to-Fitness mappings. Also, the individual solutions in all these models are positioned in the landscape solely based on their measured fitness. In the present paper, to counter the limitations of the fitness measure, we introduce a new evolvability metric that features both the quality and quantity of phenotypic variation following genetic change. With this new metric, we can *visualize* evolvability in the behavior-diversity/performance space and *predict* the performance of the population in previously untested environments¹. Such predictive insights on the adaptive characteristics of evolved individuals is particularly important, since it is difficult if not impossible to consider and evaluate a priori every possible scenario the robot may encounter during its operation. We employ our new approach to “signature” evolvability to compare many different encodings of controllers extracted from the literature. Numerous encodings have been proposed in ER, taking inspiration from natural developmental processes, in particular, to evolve control systems for robots (e.g., Gruau (1994); Kodjabachian and Meyer (1998); Clune et al. (2009a); Cheney et al. (2013); Lee et al. (2013); Lewis et al. (1992); Morse et al. (2013)). Given the multitude of available encodings, it is crucial to compare them and understand their differences, so that the ER community can focus on the most promising ones. In the selection of encodings investigated in our study, both direct and generative schemes are considered. Direct encodings encompass a one-to-one mapping between genes and phenotypic traits, and are the simplest form of encoding thus serving as a reference for comparison (e.g., Koos et al. (2013)). We also evaluate the more complex generative encodings characterized by a one-to-many mapping between genes and phenotypic traits, i.e., a single gene describes several phenotypic traits (Stanley and Miikkulainen, 2002; Stanley, 2007). These state of the art encodings are expected to exploit geometric information of the robot morphology to generate regular and modular phenotypic patterns (e.g., (Stanley et al., 2009; Clune et al., 2011; Morse et al., 2013)).

Overall, we investigate five encodings for the classical ER problem of legged robot locomotion Lewis et al. (1992); Gruau (1994); Hornby et al. (2005); Clune et al. (2009a); Bongard et al. (2006); Clune et al. (2009b, 2011); Yosinski et al. (2011); Koos et al. (2013); Lee et al. (2013): (1) open-loop central pattern generator (CPG) evolved with a direct encoding,

(2) open-loop CPG based on non-linear oscillators Crespi et al. (2013), evolved with a Compositional Pattern Generator (CPPN) Stanley (2007), (3) closed-loop CPG evolved with a CPPN, (4) artificial neural network (ANN) evolved with CPPN, inspired by HyperNEAT Stanley et al. (2009); Clune et al. (2011), and (5) the recently introduced single-unit pattern generator (SUPG) Morse et al. (2013). For all these encodings, the pertinent questions are the same: are these encodings facilitating evolvability, and are the encoded individuals capable of adapting rapidly to novel situations? Furthermore, does the inclusion of a sensory feedback mechanism improve the evolvability provided, and the adaptive capabilities of the individual? To both answer these questions and evaluate the relevance of our measure of evolvability, our experiments are divided into two phases: first, we compare the evolvability signature obtained with each encoding, and consequently predict their adaptability to novel scenarios, then we evaluate the accuracy of our predictions by analyzing the ability of each encoding to effectively deal with the new scenarios (here, when some of the robot’s legs are damaged).

2 Related work

The section reviews the different Genotype-to-Phenotype mapping schemes implemented in evolutionary computation studies. We also review the functional characteristics of encodings that facilitate evolvability both in natural and artificial systems, and the empirical methods available to estimate evolvability. The encodings compared in this study are described in detail in Sec. 5.

2.1 Encoding schemes

A mapping from the Genotype-to-Phenotype is a model of the process that “develops” an individual’s phenotype from the information available in the genotype (Johannsen, 1911). In natural organisms, the elaborate and intricate developmental systems comprise intra-cellular mechanisms of transcription and translation of proteins (Freeland and Hurst, 1998; Novozhilov et al., 2007) and regulation of gene expression (Alberts et al., 1994), and the inter-cellular mechanisms of cell differentiation and specialization (Michod and Roze, 2001; Grosberg and Strathmann, 2007). Furthermore, these complex developmental processes are a result of billions of years of natural selection (Javaux et al., 2010).

Many different genetic encodings have been implemented in the field of evolutionary computation, from the early deterministic and fixed Genotype-to-Phenotype mappings (Friedberg, 1959; Koza, 1992), to the recent, plastic and more naturalistic schemes (Stanley, 2007; Tonelli and Mouret, 2013). In this review, we broadly classify the available encodings based on their degree of genotypic reuse, i.e., the extent to which genes are allowed to be reused in developing the phenotype, into direct encodings, and generative or developmental encodings (for detailed review see Stanley and Miikkulainen (2003)).

2.1.1 Direct encodings

Direct encodings are characterized by the complete absence of any genotypic reuse, wherein each gene of the genotype is utilized at most once to determine the phenotype (Hornby

¹A preliminary study on our approach to visualize evolvability is published in a conference paper Tarapore and Mouret (2014).

and Pollack, 2002). Such a mapping has been used since the advent of evolutionary computation and genetic programming research (Friedberg, 1959), and is the most widely used encoding scheme in the field. Moreover, direct encodings are easy to implement, and have been applied successfully for the evolution of various robot behaviors such as locomotion gaits for multilegged and tensegrity platforms (e.g., Lewis et al. (1992); Téllez et al. (2006); Koos et al. (2013); Iscen et al. (2013)), navigation and obstacle avoidance for wheeled robots (e.g., Nolfi and Floreano (2001)), body-brain evolution in artificial life systems (Hornby and Pollack, 2002), and cooperative foraging in robot swarms (e.g., Waibel et al. (2009); Dorigo et al. (2005); Kernbach (2013)).

Researchers in evolutionary robotics have recently sought to evolve increasingly complex artifacts (Eiben, 2014). However, in using the direct encoding scheme, as the size of the genotype grows linearly, the possible combination of allelic values for the genotype (the solution search space) grows exponentially (Yao, 1999). The consequent scalability problem prevents the usage of a direct Genotype-to-Phenotype mapping to evolve solutions for complex problems. For example, the human brain consists of approximately 86 billion neurons and 100 trillion neural connections (Azevedo et al., 2009), and therefore successfully evolving a directly encoded *in silico* brain-like cognitive device is almost surely impossible.

2.1.2 Generative encodings

Natural systems are much more sophisticated than the state of the art solutions from artificial evolution. Inspired by the complexity of biological systems (Lipson, 2007), researchers have abstracted the underlying developmental processes, to formulate generative or developmental Genotype-to-Phenotype maps for artificial systems (e.g., Bongard and Pfeifer (2001); Komosiński and Rotaru-Varga (2001); Gruau (1994); Hornby and Pollack (2002); Stanley (2007); Mouret et al. (2010); Devert et al. (2011)). The generative encodings can be broadly defined as, *encodings that map the genotype to the phenotype through a process of growth from a simple genotype in a low-dimensional search space to a complex phenotype in a high-dimensional space*. These encodings have been applied successfully to various application problems, from allowing a computer to design antennas for satellites (Hornby et al., 2006), designing tables (Hornby, 2005, 2004) and creating tessellating tile shapes (Bentley and Kumar, 1999), to evolving locomotion gaits for both soft robots (e.g., Cheney et al. (2013)) and conventional multilegged robots (e.g., Clune et al. (2009a)).

The many implemented generative encodings are all commonly characterized by a one-to-many mapping between an element of the genotype and many elements of the phenotype. In the resulting Genotype-to-Phenotype mappings, each gene of the genotype encodes for many phenotypic traits, thus allowing for *genotype reuse*. Consequent to the capability of the encoding to reuse parts of the genotype to affect different phenotypes, generative encodings in comparison to direct encodings, (i) exhibit higher efficiency in representing complex phenotypes, (ii) operate in a more tractable solution search space, (iii) scale well to large phenotypic spaces, and (iv) are capable of generating regular and modular control architectures (Stanley and Miikkulainen, 2003; Clune et al., 2009b; Stanley et al., 2009). However, while direct encodings are easy to implement, generative encodings

follow more complex implementations, determined by their level of abstraction of development.

Existing generative encodings model the underlying natural developmental processes at several different levels of abstraction, ranging from the low-level cell chemistry simulations to the high-level grammatical approaches (Stanley, 2007). The cell chemistry methods simulate the local intracellular and inter-cellular interactions between genes and protein products, modulated by signals from a gene regulatory network (e.g., Furusawa and Kaneko (1998)). These microscopic approaches are based on the philosophy that the vital functions that allow development to assemble complex phenotypes are located in the low-level cellular interactions occurring in a developing embryo. By contrast, the grammatical approaches simulate development with a set of high-level symbol replacement rules (e.g., Mjolsness et al. (1991)). These approaches grow a final structure from a single seed symbol, by the repeated application of grammar rules on specified target symbols. Consequently, local interactions and temporal unfolding control the phenotype development through the grammars (Stanley and Miikkulainen, 2003).

In most studies comparing encoding schemes, the generative encoded individuals frequently outperform their directly encoded counterparts for a range of diverse tasks such as, designing 3D objects (e.g., Hornby (2005); Clune and Lipson (2011)), game playing (e.g., Reisinger and Miikkulainen (2007); Gauci and Stanley (2010)), pattern matching (e.g., Clune et al. (2011)), and multilegged robot locomotion (e.g., Hornby and Pollack (2002); Seys and Beer (2007)). The generatively encoded individuals achieve a better task performance, and a faster rate of evolution (e.g., see Gruau (1994); Komosiński and Rotaru-Varga (2001)). Furthermore, generative encodings also attain a higher proportion of beneficial genetic mutations (e.g., Clune et al. (2009a); Reisinger and Miikkulainen (2007)), and are capable of exploring a larger range of phenotypes from genetic change (e.g., Reisinger et al. (2005); Lehman and Stanley (2013)).

2.2 Evolvability

The process of evolution in natural systems comes from the cooperation of, (i) exploratory genotypic variation, (ii) the corresponding phenotypic variation, and (iii) selection operators that preserves the improvements in heritable phenotypic traits over previous generations. The crucial coordination between these three forces yields the evolvability of an evolutionary system (Alberch, 1991; Hu and Banzhaf, 2010; Pavlicev and Wagner, 2012).

The first formal definition of evolvability stems from research in computer science. In experiments with optimization algorithms using genetic programming, Lee Altenberg defined evolvability as “the ability of a population to produce variants fitter than any yet existing” (Altenberg, 1994). In natural evolutionary systems, Kirschner and Gerhart (Kirschner and Gerhart, 1998) describe evolvability, also called evolutionary adaptability, as “the capacity to generate heritable, selectable phenotypic variation”. Marrow (Marrow et al., 1999) considers evolvability as a characteristic relevant to both artificial and natural evolutionary systems, and viewed as the capability of a population to evolve. In a summary of results from both evolutionary biology and evolutionary computer science, Wagner and Altenberg (Wagner and Altenberg, 1996) view evolvability as “the ability

of random variations to sometimes produce improvements". These incremental improvements are critically dependent on the Genotype-to-Phenotype encoding. Mappings facilitating evolvability, confer on the individual a robustness to lethal mutations, and exhibit a modular architecture wherein genes preferably only affect traits with the same function (Pavlicev and Wagner, 2012).

Although the concept of evolvability is still very much under discussion, for our study we adopt the definition pertaining to adaptability, and the generation of major phenotypic breakthroughs Pigliucci (2008); Clune et al. (2013): *Evolvability is the capability of a population to rapidly adapt to novel and challenging environments.*

The measurement of evolvability conferred by an encoding is a complex and difficult problem. Phenotypic fitness or task performance is a directly observable measure, and a criteria for selection. However, the *potential* to generate a better fitness, evolvability, is a less tangible type of observable and is more difficult to measure (Hu and Banzhaf, 2010). While a formal method to quantify evolvability has not yet been agreed upon in the literature, some empirical methods have been proposed notwithstanding.

In Gerhart and Kirschner’s theory of facilitated variation Gerhart and Kirschner (2007), which unifies most earlier findings of cellular and developmental processes with characteristics of evolvability, the capacity of an individual to evolve is considered to have two functional components: (i) to curtail the proportion of lethal mutations; and (ii) to decrease the number of mutations necessary to evolve diverse or novel phenotypes. Nonetheless, most studies measuring evolvability focus mainly on one only of these two aspects. Most comparisons estimate evolvability solely as the proportion of mutations that are beneficial to an individual (e.g., Hornby et al. (2003); Clune et al. (2011); Hornby and Pollack (2002); Clune et al. (2009a); Reisinger and Mikkulainen (2007)), and irrespective of the phenotypic novelty of the resultant offspring. In other work, Ibáñez-Marcelo and Alarcón Ibáñez-Marcelo and Alarcón (2014) characterize evolvability as the number of viable mutations required to reach different phenotypes, but do not take into account the novelty or the amount of diversity between the different phenotypes. By contrast, Lehman and Stanley (2011b); Reisinger et al. (2005); Lehman and Stanley (2013) quantify evolvability only on the basis of the amount of phenotypic diversity resulting from genetic change, usually without considering the quality of the change. However, both factors are essential to quantify evolvability, to discount for, (i) mutations that generate very diverse phenotypes, but prove lethal to the organism, and (ii) mutations resulting in small increments in performance that improve on a trait, but may not be able to generate novel phenotypes. Therefore, for our comparison between encodings, evolvability is visualized by characterizing both the nature of the genetic mutation, and the quantity of generated phenotypic variation.

3 Evolvability signature

In this study, the evolvability provided by a Genotype-to-Phenotype mapping is described by a distinct signature featuring information on the effect of genetic mutations on both, the *quality* of mutated individuals (their viability), and the *quantity* of generated phenotypic variation. Our signature pictures these two features as a statistical distribution

of fitness (quality-feature) and behavior diversity (quantity-feature), following multiple independent, randomly sampled mutations. The two features are treated separately instead of being combined into a single quantitative measure of evolvability, to consider the trade-offs between them in their individual influence on evolvability Deb (2001).

Feature 1: Deleteriousness of mutations. The first feature in our signature of evolvability is computed as the proportion decrease in the fitness of a mutated individual.

For an individual i and the mutant i' , we have,

$$f_1 = \frac{F'_i - F_i}{F_i} \quad (1)$$

where F_i and F'_i , are the fitness values before and after the application of a random genetic mutation, respectively.

The feature f_1 reflects the behavior quality following beneficial ($f_1 > 0$), neutral ($f_1 \approx 0$), and deleterious ($f_1 < 0$) genetic change. Additionally, mutations that prove lethal are associated with f_1 values less than -1 , reflecting a 100% or larger decrease in individual fitness.

Feature 2: Diversity of behaviors. Following the theory of facilitated variation Gerhart and Kirschner (2007), the second feature in our signature of evolvability evaluates the diversity of phenotypes than can be reached from a given individual. The phenotype can here be understood in two ways: in an evolutionary biology perspective, the phenotype can describe both morphological traits and behaviors Arnold (1992); Dawkins (1999), whereas in a evolutionary robotics perspective, only morphological traits are considered to be parts of the phenotype (e.g., the parameters and the topology of an evolved networks form the phenotype) Stanley and Mikkulainen (2003). The distinction between phenotype and behaviors avoids potential confusions when working on developmental encodings (genotype-phenotype maps), which focus on morphological traits, or when working on selective pressures, which often focus more on the behavior than on the representation Doncieux and Mouret (2014).

In the present study, we focus on the diversity of *behaviors*, as done in evolutionary biology, because it best distinguishes promising individuals from the poor performers when evolving robot controllers Lehman and Stanley (2011a); Mouret and Doncieux (2012). For instance, all the neural-networks that are not connected to the robot’s actuators lead to the same stopped-robot behavior, whereas the morphological traits (synaptic parameters and topology of the neural-network) can be widely different. A second advantage of looking at behaviors instead of morphological traits is that the behavior representation can be independent of the implementation of the controller, thus allowing us to compare the evolvability of very different controllers like CPGs, neural networks, and SUPG controllers.

Measuring behavioral differences recently received a lot of attention in evolutionary robotics because several experiments showed that explicitly encouraging the diversity of evolved behaviors helps to mitigate the issue of premature convergence Mouret (2011); Mouret and Doncieux (2012); Doncieux and Mouret (2013, 2014). It is also the main driving force in the Novelty Search algorithm, which leads to high-performing individuals in deceptive domains by only searching for novel behaviors and disregarding task-fitness values Lehman and Stanley (2011a). Following this interest in measuring behavioral differences, many behavioral diversity metrics have been proposed, ranging from task-specific metrics (e.g., difference between end points of a robot’s tra-

jectory), to more task-agnostic measures (e.g., differences in the robot sensory-motor flow), and various information theoretic measures (detailed review in Doncieux and Mouret (2014) and in Mouret and Doncieux (2012)).

Among the investigated measures, the mutual information diversity metric provides a general approach to compute a non-linear, non-monotonic relationship between behaviors, that is applicable to numerical and symbolic behavioral representations, both in the continuous and discrete domains Cover and Thomas (1991); Kraskov et al. (2004). For our signature of evolvability, we compute the behavioral diversity as the normalized mutual information between behaviors of an individual, before and after its genome is mutated.

Assuming that the behavior of an individual i can be represented as a discrete vector B_i (details in Mouret and Doncieux (2012)), for the behaviors B_i and B'_i , of individual i and mutant i' , we have:

$$H(B_i) = - \sum_{b_i \in B_i} p(b_i) \log p(b_i) \quad (2a)$$

$$H(B_i, B'_i) = - \sum_{b_i \in B_i} \sum_{b'_i \in B'_i} p(b_i, b'_i) \log p(b_i, b'_i) \quad (2b)$$

$$f_2 = 1 - \frac{H(B_i) + H(B'_i) - H(B_i, B'_i)}{\max(H(B_i), H(B'_i))} \quad (2c)$$

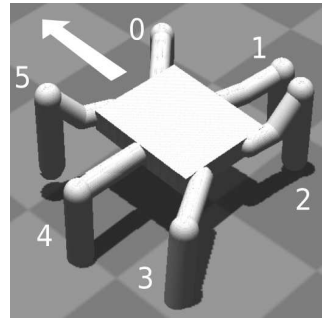
where $H(B_i)$ is the entropy of the behavior B_i comprising the individual states b_i with probability $p(b_i)$, $H(B_i, B'_i)$ is the joint entropy between behaviors B_i and B'_i with joint probability density function $p(b_i, b'_i)$, and f_2 denotes the inverse of the normalized mutual information between the two behaviors.

The entropy and joint entropy are computed by first approximating $p(b_i)$ and $p(b_i, b'_i)$, by counting the number of instances of each behavior state. Systematic errors in the probability estimates, consequent to the limited number of available data samples, is compensated for by adding a corrective term E to the computed entropy: $E = (S_i - 1) / 2T$ (where T is the size of the temporal window over which the entropy is computed, and S_i is the number of states for which $p(b_i) \neq 0$), and $E = (S_i + S_{i'} - S_{i,i'}) / 2T$ to the joint entropy (where $S_i, S_{i'}, S_{i,i'}$, and T have an analogous meaning to the previous case). These corrective terms compensate for the systematic and random errors in the observed entropy of a series, that systematically bias downwards the expected value of the observed entropy from the true entropy (derived in Roulston (1999)). Integrating the corrective term to the equations for entropy and joint entropy, we have:

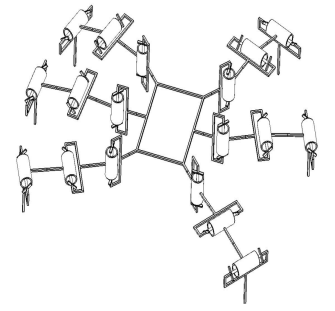
$$H(B_i) = - \sum_{b_i \in B_i} p(b_i) \log p(b_i) + \frac{S_i - 1}{2T} \quad (3a)$$

$$H(B_i, B'_i) = - \sum_{b_i \in B_i} \sum_{b'_i \in B'_i} p(b_i, b'_i) \log p(b_i, b'_i) + \frac{S_i + S_{i'} - S_{i,i'} - 1}{2T} \quad (3b)$$

Estimates of the corrected entropy (eq. 3a) and joint entropy (eq. 3b) are then used to update the mutual information distance between behaviors. The resulting feature f_2 represents the quantity of behavioral variation following genetic change, and is indicative of the ability of the evolutionary system to produce novel behaviors.



(a) Hexapod robot



(b) Kinematic scheme

Figure 1: (a) Snapshot of an 18-DOF simulated hexapod robot walking on a horizontal surface, with contacts simulated. (b) Kinematic scheme of the robot, with cylinders representing actuated pivot joints. The three servos on each leg, s_1 , s_2 and s_3 , are labeled in increasing order of distance to robot torso.

4 Hexapod robot locomotion problem

The evolution of locomotion gaits for multilegged robots is a classical problem in evolutionary robotics, addressed in many studies utilizing both direct and generative encodings, on bipedal (e.g., Liu and Iba (2004)), quadrupedal (e.g., Clune et al. (2011); Hornby et al. (2005); Risi and Stanley (2013); Téllez et al. (2006); Valsalam and Miikkulainen (2008)), and hexapedal robots (e.g., Valsalam and Miikkulainen (2008); Zykov et al. (2004); Barfoot et al. (2006)) – employed here for the comparison of different encodings. In most existing studies on evolved locomotion gaits, the performance of an individual is analyzed solely by its walking speed and the required number of generations of evolution. The rate of evolution and evolved performance has also been linked to evolvability provided by the encoding scheme, wherein controllers achieving a higher task fitness and requiring fewer generations to evolve are considered more evolvable (e.g., see Hornby et al. (2003); Clune et al. (2009a); Gruau (1994); Komosiński and Rotaru-Varga (2001)). While these approaches provide interesting insights on the performance of the Genotype-to-Phenotype mapping, they largely ignore its capabilities to generate viable phenotypic variations (diverse gaits in case of legged robots). However, the diversity of evolved walking gaits is important for the legged robot to recover rapidly from faults such as, the loss of one or more limbs, or motor malfunctions (Koos et al., 2013), and for the robot to adapt to previously unencountered environmental changes. Furthermore, an efficient recovery is particular relevant for hexapedal legged robots, wherein the probability of component failure is high, consequent to the large number of moving parts.

Hexapod platform details: The hexapod robot is simulated on a flat, horizontal surface (Fig. 1a), with the Open Dynamics Engine² (ODE) physics simulator. The robot has 18 Degrees of Freedom (DOF), 3 for each leg (Fig. 1b), and each DOF is actuated by a single servo. The first servo on each leg (s_1) actuates the horizontal orientation of the leg within range $[-\pi/8, \pi/8]$ radians. The second (s_2) and third (s_3) servos control the leg elevation and extension, respectively, each within the range of $[-\pi/4, \pi/4]$ radians.

Angular positions are sent to the 18 servos once every

²<http://www.ode.org>

15 ms. Furthermore, in order to maintain the last subsegment of each leg vertical (for enhanced stability), the control signal for the third servo (s_3) is always in antiphase to that of the second servo (s_2). Consequently, the robot is reduced to a 12 DOF system, despite being actuated by 18 motors.

Hexapod gait representation: The behavioral diversity in our signature of evolvability corresponds to the inter-gait diversity in the hexapod robot locomotion problem. For this diversity, a hexapod gait is represented using a gait diagram (Kajita and Espiau, 2008, p. 379), comprising a binary matrix C of leg-surface contacts:

$$C_{il} = \begin{cases} 1 & \text{if leg } i \text{ makes surface contact at time-step } t, \\ 0 & \text{otherwise.} \end{cases}$$

where $t \in \{0 \dots T\}$, the gait is evaluated for T time-steps, and the hexapod legs $l \in \{0 \dots 5\}$.

The hexapod gait for an individual i is represented by binary vector B_i , comprising the contacts in C concatenated in row-major order, $B_i = [C_{00}, C_{10} \dots C_{T5}]$. Diversity between two gaits is measured as the normalized mutual information between the corresponding gait vectors (eq. 2c).

5 Encoding schemes analyzed

Generative encodings for evolving our hexapod locomotion controllers are based on CPPNs (Stanley, 2007). The CPPN abstracts the processes of embryonic development by determining the attributes of phenotypic components as a function of their geometric location in the individual, instead of simulating complex inter-cellular interactions and chemical morphogen gradients to determine component location (Carroll, 2005). In nature, cells differentiate themselves into different lineages, influenced by their immediate environment (the epigenetic landscape, Goldberg et al. (2007)). Analogously, the CPPN genome outputs the fate of an organismal component as a function of its geometric coordinates in the individual.

The CPPN genome is represented as a directed graph, comprising a set of *Sine*, *Gaussian*, *Sigmoid*, and *Linear* type of nodes, connected by weighted links. The node type indicates the activation function applied to the sum of its weighted inputs, to compute the node output. Selected activation functions can succinctly encode a wide variety of phenotypic patterns, such as symmetry (e.g., a Gaussian function) and repetition (e.g., a Sine function), that evolution can exploit. Mutations to the CPPN genome can change the connection weights and node type, and add or remove nodes from the graph. Consequently, the topology of the CPPN is unconstrained, open-ended, and can represent any possible relationship between the input coordinates of the phenotypic component and its output attributes (see details in Stanley (2007)).

In this study, the CPPN genotype is mapped to four very different phenotypes to control hexapod locomotion, open-loop CPGs, closed-loop CPGs, ANNs (minimal HyperNEAT), and SUPGs (summarized in Table 1). The SUPG is a new type of macro-neuron introduced by Morse et al. Morse et al. (2013) to genetically encode spatio-temporal oscillatory patterns.

5.1 Encoding CPGs with CPPNs

A CPG is a biological neural network, comprising a distributed array of neurons located along the spinal cord of vertebrates (Frigon and Rossignol, 2006). The spinal neural centers exhibit rhythmic activity, generating complex high-dimensional signals for the control of coordinated periodic movements required for animal locomotion. In recent years, bio-inspired engineering approaches have led researchers to model these CPGs as coupled dynamical oscillatory systems (Ijspeert, 2008), to generate locomotion control policies for biped (Aoi and Tsuchiya, 2005; Taga, 1994), quadruped (Righetti and Ijspeert, 2008, 2006; Fukuoka et al., 2003) and hexapod robots (Ren et al., 2014), as well as modular (Sproewitz et al., 2008) and swimming/walking salamander-like robots (Ijspeert et al., 2007; Crespi et al., 2013).

The design of CPG-based control policies for robot locomotion offers many advantages (Ijspeert, 2008). The main advantage of using such controllers is the stable limit cycle behaviors of the coupled dynamical system. Additionally, the CPGs maintain a smooth transient response to changes to external inputs (e.g., from sensory feedback). Another advantage of CPGs is the intrinsic properties of inter-oscillator synchronization, allowing for coordinated interactions between the robot and its environment. Finally, using CPGs for generating coordinated policies reduces the search space of the control problem, as opposed to using other approaches such as continuous-time recurrent neural networks (Beer and Gallagher, 1992).

In most studies investigating CPGs for robot locomotion, the parameters of the model are hand-tuned by the designer (e.g., Aoi and Tsuchiya (2005); Righetti and Ijspeert (2008); Crespi et al. (2013)). While such an approach does provide interesting coordinated robot movements, the system is unable to autonomously explore a rich diversity of different locomotion gaits. Furthermore, the parameters of the model are required to be re-tuned in response to changes in the environment, or to damages incurred by the robot. To overcome these issues, we seek to evolve parameters of generatively encoded CPGs, thus combining the advantages of CPG models with that of generative encodings. The following two CPPN-encoded CPG models are investigated: (i) an open-loop model wherein the control policies are not influenced by any sensory feedback; (ii) a closed-loop model, with sensory feedback from touch sensors on each of six hexapod legs influencing the resultant gait.

5.1.1 Open-loop CPG

The open-loop generatively encoded CPG system is comprised of 12 coupled amplitude-controlled phase oscillators (Ijspeert et al., 2007), governing the actuation of the 12 servos (s_1 and s_2 , on each of 6 robot legs). Each oscillator is modeled by the following set of ordinary differential equations:

$$\dot{\theta}_i = 2\pi F_i + \sum_j \alpha_j w_{i,j} \sin(\theta_j - \theta_i - \phi_{i,j}) \quad (4a)$$

$$\ddot{\alpha}_i = \beta \left(\frac{\beta}{4} (A_i - \alpha_i) - \dot{\alpha}_i \right) \quad (4b)$$

$$\gamma_i = \alpha_i \cos(\theta_i) \quad (4c)$$

where α_i and θ_i denote the amplitude and phase of the oscillator i , A_i represents its intrinsic amplitude (in radians),

Table 1: Summary of encoding schemes to evolve controllers for hexapod robot locomotion.

Encoding	Signal generator	Feedback	Summary	References
Open-loop CPG	Phase oscillator	✗	An amplitude-controlled phase oscillator is used to generate locomotion gaits. The 12 amplitude parameters, and 11 inter-oscillator phase bias parameters of the oscillator, are generatively encoded with a CPPN.	Ijspeert et al. (2007); Righetti and Ijspeert (2006); Crespi et al. (2013); Ijspeert (2008)
Closed-loop CPG	Phase oscillator	✓	A generatively encoded amplitude-controlled phase oscillator is extended with the inclusion of a phase resetting mechanism. The extension introduces a sensory feedback mechanism that modulates the oscillations produced by the CPG, adapting the oscillation period depending on the locomotion gait and terrain.	Aoi and Tsuchiya (2005); Fukuoka et al. (2003); Righetti and Ijspeert (2008)
HyperNEAT	ANN	✗	A widely used generative encoding scheme, used to evolve large-scale ANNs. The weights of the neural network are encoded with a CPPN. Evolved ANNs have been successfully deployed to generate symmetric and coordinated gaits for both simulated and physical quadruped robots.	Stanley et al. (2009); Clune et al. (2011); Yosinski et al. (2011); Lee et al. (2013)
SUPG	CPPN	✓	A recently developed encoding scheme, wherein a CPPN encodes the attributes of a SUPG. The SUPG is a macro-neuron, that upon receiving an external trigger, produces a single cycle of an oscillatory signal. Repeated triggering of the SUPG generates a periodic gait for a multilegged robot. In previous work, the SUPG outperformed HyperNEAT encodings in evolving locomotion gaits for a simulated quadruped robot, thus encouraging further study.	Morse et al. (2013)
Direct	Phase oscillator	✗	A simple locomotion controller, used as reference for comparison between encoding schemes. The robot is controlled with an amplitude controlled phase oscillator, the 12 amplitude parameters, and 11 inter-oscillator phase bias parameters, are directly encoded.	Koos et al. (2013); Cully and Mouret (2013)

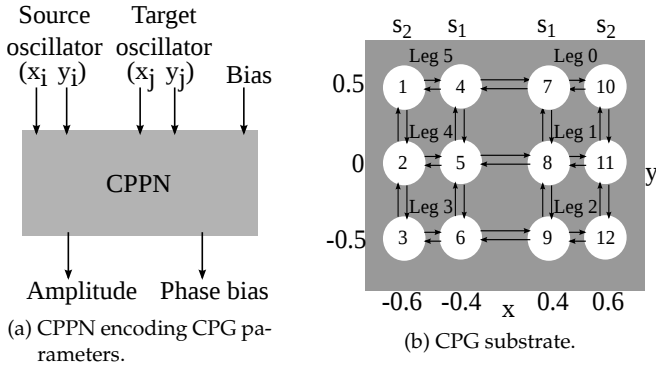


Figure 2: Encoding CPGs with CPPNs. The intrinsic amplitude A_i of each oscillator is encoded by the CPPN as a function of its position in the substrate. Phase biases $\phi_{i,j}$ of inter-oscillator couplings are determined by querying the CPPN with the coordinates of every pair of adjacent source (x_i, y_i) and target (x_j, y_j) oscillators. For every pair of adjacent oscillators, the query is made only once as $\phi_{i,j} = -\phi_{j,i}$.

and F_i the intrinsic frequency (in Hz). The oscillator amplitude α_i converges to the intrinsic amplitude A_i at a rate determined by positive constant β (set to 10 rad/s for rapid convergence). Couplings between oscillators are defined by weights $w_{i,j}$, and phase biases $\phi_{i,j}$. The coupling weights define the strength of the coupling, and influence the time to synchronize between oscillators (all weights prefixed at 20 to enable rapid inter-oscillator synchronization). The phase bias $\phi_{i,j}$ determines the phase difference between the oscillators i and j . The output of the oscillator is denoted by γ_i , and is computed using the Euler integration method with a step-size of 20 ms.

Oscillator equations (eq. 4a,b and c) were designed such that the output of the oscillator γ_i exhibits a limit cycle behavior, producing a stable periodic output. The first equation (eq. 4a) determines the time evolution of the phase θ_i of

the oscillators. In this study, all 12 oscillators of the CPG have the same frequency ($F_i = 1$ Hz), and the bidirectional coupling between oscillators is such that $\phi_{i,j} = -\phi_{j,i}$. Furthermore, we ensure that the sum of all the phase biases in every closed loop of inter-oscillation couplings is a multiple of 2π , so that all the phase biases in a loop are consistent. In such a parameter regime, the oscillator phases grow linearly at a common rate F_i and with an inter-oscillator phase difference of $\phi_{i,j}$. The second equation (eq. 4b) is a critically damped second order linear differential equation, with A_i as the stable fixed point. Therefore, from any initial condition the oscillator amplitude α_i will asymptotically and monotonically converge to the intrinsic amplitude A_i , allowing the amplitude to be smoothly modulated. Finally, utilizing the amplitude and the phase of the oscillator (eq. 4c), the output of the oscillator γ_i governs servo actuation.

In this first generative encoding evaluated, a CPPN encodes the intrinsic amplitudes A_i and inter-oscillator phase biases $\phi_{i,j}$ of 12 oscillators of the CPG (Fig. 2a). The oscillators are placed in a 2-D Cartesian grid termed the *substrate*, so that each oscillator has a distinct (x, y) coordinate, and so as to reflect the hexapod robot morphology (Fig. 2b). The intrinsic amplitude of each oscillator i is obtained by inputting to the CPPN the coordinates (x_i, y_i) , and setting the inputs (x_j, y_j) to 0. Amplitudes output are scaled to the allowable angular range of the corresponding motors. In the CPG, adjacent oscillators are coupled together (Fig. 2b). The phase bias for every pair of adjacent oscillators i and j is obtained by querying the CPPN with inputs (x_i, y_i) and (x_j, y_j) , and scaling the output to range $[0, 2\pi]$. Furthermore, the following two constraints are introduced: (i) couplings are bilaterally symmetrical, i.e., $\phi_{i,j} = -\phi_{j,i}$. For every pair of adjacent oscillators, the phase bias is queried only once; (ii) phase biases $\phi_{2,1}$, $\phi_{2,3}$, $\phi_{7,4}$, $\phi_{9,6}$, $\phi_{10,11}$ and $\phi_{12,11}$ are not queried, but computed such that the sum of phase biases in every closed loop of the CPG is a multiple of 2π (oscillators numbered in Fig. 2b). Therefore, the total number of CPG parameters generatively encoded by the CPPN is 23 (12 intrinsic amplitude and 11 phase bias parameters).

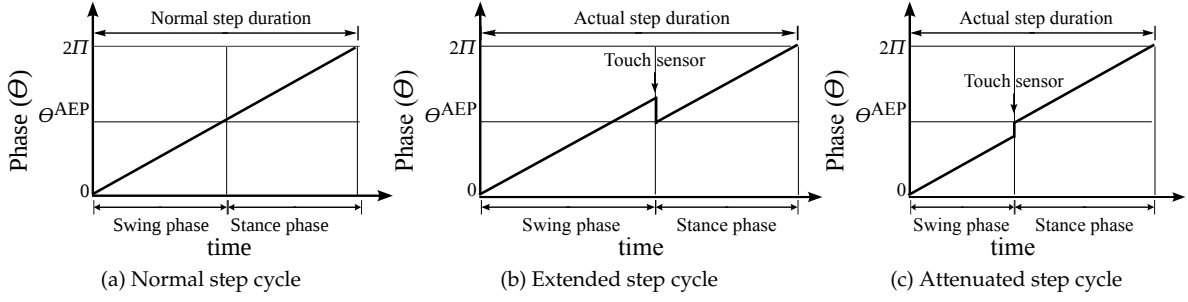


Figure 3: An illustration of the normal and actual phase θ_i of oscillator i over time. The phase of an oscillator governing leg horizontal orientation (servo s_1) is reset to θ^{AEP} following a signal from the touch sensor, resulting in changes to the step cycle.

5.1.2 Closed-loop CPG

The second generative encoding scheme evaluated is an extension of the open-loop scheme (see Sec. 5.1.1). While the generatively encoded CPG parameters and the CPPN encoding remains the same as in the open-loop model (Fig. 2), the modification introduced is a sensory feedback mechanism that modulates the oscillations produced by the CPG. In this closed-loop encoding scheme, feedback signals from the touch sensors attached to each of the six legs of the hexapod trigger a phase-resetting mechanism Aoi and Tsuchiya (2005), that adapts the oscillation period depending on the gait and the terrain. For the phase-resetting mechanism, two extreme positions of the horizontal orientation of the robot leg are introduced with respect to the robot trunk, (i) the anterior extreme position (AEP), where the swing phase transitions to the stance phase, and (ii) the posterior extreme position (PEP), where the stance phase transitions to the swing phase.

The phase of the oscillator at the AEP, θ^{AEP} is determined by the following equations.

$$\beta = \frac{T_{ST}}{T_{ST} + T_{SW}}$$

$$\theta^{AEP} = 2\pi(1 - \beta)$$

where T_{SW} is the duration of the swing phase, T_{ST} is the duration of the stance phase, and β ($0 < \beta < 1$) is the duty ratio, i.e., the ratio between the stance phase duration and the total step cycle (Kajita and Espiau, 2008). In this study, β is prefixed at 0.5, resulting in θ^{AEP} to be at π .

Utilizing θ^{AEP} for the oscillators governing horizontal orientation on each of six legs of the robot (oscillators $i \in \{4, 5 \dots 9\}$, see Fig. 2), the phase equation (eq. 4a) is modified while keeping the equations for the amplitude (eq. 4b) and the oscillator output (eq. 4c) unchanged.

$$\dot{\theta}_i = 2\pi F_i + \sum_j \alpha_j w_{i,j} \sin(\theta_j - \theta_i - \phi_{i,j}) + g_i \quad (5a)$$

$$g_i = (\theta^{AEP} - \theta_i \pmod{2\pi}) \delta(t_i^{land} - t) \quad (5b)$$

where t_i^{land} is the time when the foot of the leg being horizontally oriented by oscillator i lands on the ground, and $\delta(\cdot)$ is the Dirac delta function ($[\cdot] : \mathbb{R} \rightarrow [0, \infty)$) defined as 0 for all arguments except at the origin, where it is 1. The function g_i resets the phase θ_i of oscillator i to the normal value θ^{AEP} instantaneously when the foot of the leg lands on the ground.

When the foot touches the ground, the phase of the corresponding oscillator is reset to θ^{AEP} , the normal phase value where the transition from the swing phase to the stance phase is supposed to occur. Consequently, in the oscillator's step cycle the actual duration of the swing phase is not fixed, but is dependent on the timing of the foot touching the ground (see Fig. 3), and the resulting oscillator's step cycle is thus modulated by sensory feedback.

5.2 Encoding ANNs with CPPNs (minimal HyperNEAT)

The third generative encoding scheme evaluated is a simplified version of HyperNEAT indirect encoding³ In previous work, the CPPN has been used successfully to evolve modular and regular patterns in the connection space of the ANN, resulting in symmetric and coordinated gaits for both simulated and physical quadruped robots (Stanley et al., 2009; Clune et al., 2011; Yosinski et al., 2011; Lee et al., 2013). The results encourage us to include the HyperNEAT encoding in our comparative study.

The CPPNs encode the weights of a fixed topology, single-layer feedforward ANN, featuring 2-D Cartesian grids of inputs, hidden and output neurons (Fig. 4). Neurons of the ANN are positioned in the substrate, in accordance with the hexapod robot morphology. Using the encoding, the CPPN is iteratively queried the positions of all source (x_1, y_1) and target (x_2, y_2) neurons in proximal layers, along with a constant bias, and it outputs the corresponding weights of the input-hidden and hidden-output neuron connections.

The ANN receives as input the previously requested angles (actual angles unknown) for each of the 12 pivot joints of the hexapod robot (s_1 and s_2 , for 6 legs). In addition, sine and cosine waves of frequency 1 Hz are also input to the ANN, to facilitate periodic oscillations at the output neurons. The output from the ANN at each time-step are 12 numbers (one for each of s_1 and s_2 , on each of 6 legs) in interval $[-1, 1]$, that are scaled to the allowable angular range of the corresponding motors, and indicate the next position of each motor.

In preliminary experiments, the HyperNEAT encoding evolved ANNs that exhibited high frequency output oscillations (in excess of 20 Hz). In the resultant gaits, the robot could traverse large distances by vibrating its legs rapidly, and in unison. Such high frequency pronking gaits would

³The CPPN is evolved with a simple multiobjective evolutionary algorithm, instead of the NEAT method (details in Tonelli and Mouret (2013)).

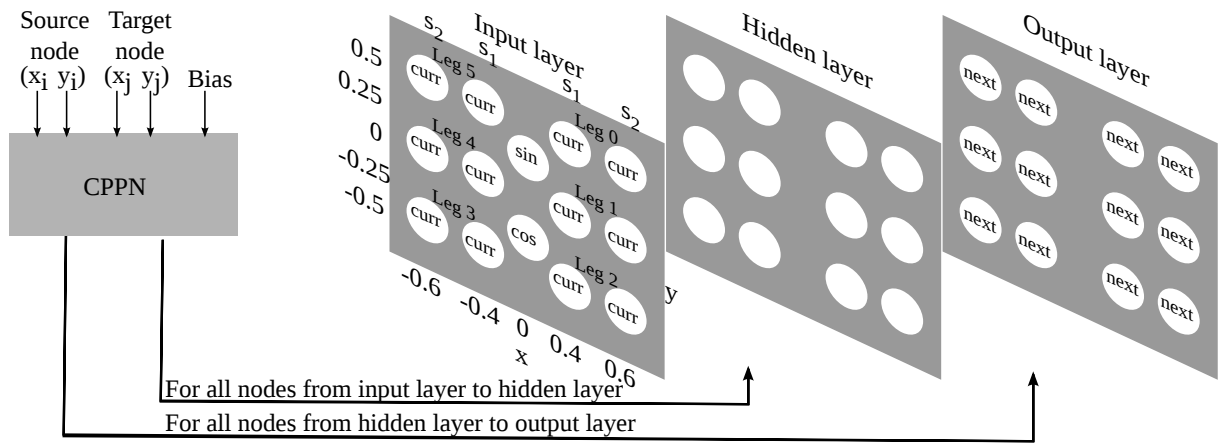


Figure 4: Encoding ANNs with CPPNs (inspired by Clune et al. (2011); Yosinski et al. (2011); Lee et al. (2013)). Inter-neuron connection weights are encoded as function of the position of source and target neurons of each neural connection. The CPPN outputs the weights of input-hidden and hidden-output neuron connections, for each source (x_i, y_i) and target (x_j, y_j) neuron in proximal layers. The ANN is input the requested angles of the previous time-step for the first two servos (s_1 and s_2) on each leg, and a sine and cosine wave. The output neurons specify the new joint angles for the current time-step.

clearly not be replicable on the physical robots, and result in overly taxed servos Yosinski et al. (2011). To resolve this problem, and as suggested by Lee et al. (2013), we generated joint angular positions with a time-step of 15 ms, by averaging over four consecutive pseudo-positions generated at 3.75 ms intervals. The number of evolved ANN controllers outputting high frequency oscillation was thus reduced, with oscillations at 2.44 ± 1.95 Hz (Median \pm IQR in 20 replicates, with the frequency at each replicate averaged across 18 servos).

5.3 Encoding SUPGs with CPPNs

In the fourth generative encoding scheme evaluated, the CPPN encodes the attributes of a SUPG. The SUPG is a macro-neuron that upon receiving an external trigger, produces a single cycle of a CPPN encoded oscillatory signal. Consequently, the repeated triggering of the SUPG generates a periodic pattern, used to govern the actuation of a multi-legged robot. In a previous work, the SUPG outperformed HyperNEAT encodings in evolving locomotion gaits for a simulated quadruped robot (Morse et al., 2013). The resultant SUPG gaits appeared faster and steadier in extended evaluations, encouraging us to study the encoding both in terms of performance, and the evolvability provided.

In this encoding, the CPPN is input the position (x, y) of the SUPG in the substrate, and the elapsed time since the SUPG was last triggered (Fig. 5a). The elapsed time is recorded by an internal timer, individual to each SUPG, and ramps upwards from an initial value of 0 to a maximum value of 1, in one period of the SUPG's output signal (Fig. 5b). Therefore, the SUPG's output is a function of both, its position in the substrate, and the time since its last cycle was triggered. Applying the SUPGs for hexapod locomotion, the substrate comprises 12 SUPGs positioned to reflect the robot morphology (Fig. 5c). The outputs of the SUPGs at each time-step specify the desired angles for the first and second servos (s_1 and s_2), on each leg of the robot.

The internal timer of the SUPG can be restarted from 0, following the occurrence of an external trigger event (Fig. 5b). Consequently, the SUPG cycle does not need to match the

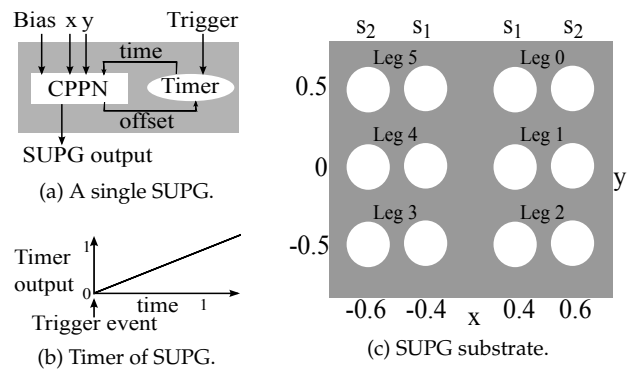


Figure 5: Encoding SUPGs with CPPNs (inspired by Morse et al. (2013)). (a) The SUPG output is a function of its coordinates (x, y) in the substrate, and the elapsed time since last trigger (output of Timer). The time of first trigger is determined by an offset. (b) Once triggered, the SUPG timer ramps upward linearly from 0 to 1 and stays there, until it is re-triggered. (c) Positions of the 12 SUPGs in the substrate, outputting the desired angles for the first two servos (s_1 and s_2), on each leg of the hexapod.

length of an optimal robot-walking step. Rather, the oscillation period can be adjusted depending on the gait and the terrain by restarting the SUPG whenever its associated foot touches the ground, thus producing a closed-loop control. In our hexapod robot, the two SUPGs actuating each leg (Fig. 5c), are simultaneously triggered by the corresponding foot touching the ground.

At the start of the simulation, all six legs of the robot are in contact with the ground, resulting in all the SUPGs being triggered simultaneously. To avoid the resulting hopping gaits, the first trigger to each SUPG is delayed by an offset. The offset output of the CPPN is determined for the s_1 SUPG on each leg by supplying its coordinates as input, and setting the time input to 0. The same offset value is also applied to the s_2 SUPG on the leg, allowing both the oscillators on each leg to start simultaneously.

5.4 Direct encoding

Locomotion controllers evolved with direct encoding are designed to be simple, wherein the actuation along each DOF of the robot is governed by the periodic signal of an open-loop amplitude controlled phase oscillator. With this encoding, hexapod leg actuation is governed by the differential equation model (eq. 4a, b and c). There are 12 evolved parameters for the intrinsic amplitudes of oscillators, governing the actuation of the two servos s_1 and s_2 , on each of six legs of the hexapod. In addition, 11 inter-oscillator phase bias parameters are also evolved (see Fig. 2b, and Sec. 5.1.1 for details on constraints on phase biases). Consequently, a directly encoded controller for the hexapod robot is fully represented by 23 parameters.

6 Experiments

We conducted 8,000 generations of artificial selection in populations consisting of 100 individuals ($8,000 \times 100$ evaluations). Our aim was to evolve controllers for the hexapod robot to walk forward, evaluated for a period of 5 s (334 time-steps). The Nondominated sorting genetic algorithm II (Deb et al., 2002) was used to simultaneously optimize the following three objectives:

$$\text{Maximize} \begin{cases} -F_i \\ -|\Theta_i| \\ \frac{1}{N} \sum_{j=0}^N D(B_i, B_j) \end{cases} \quad (6)$$

where for individual i in the population, F_i is the fitness computed as the distance between the final position of i and a goal located 25 m directly in front of the robot's initial position, Θ_i denotes the angle of the robot's trajectory with respect to the normal forward walking direction, $D(B_i, B_j)$ is the hamming distance between the binary gait vectors of individual i and j , and N is the size of the population.

In eq. 6, the first and second objectives reward individuals to walk forward large distances towards a goal, unattainable by the robot within the experiment evaluation time. The third objective is introduced to facilitate the exploration of diverse solutions and avoid premature convergence to suboptimal solutions at local minima Mouret and Doncieux (2012).

Artificial selection was conducted in 20 independent replicates, for the **Direct** encoding, and the four generative encodings, (i) **CPG** (open-loop controller), (ii) **CPG-f/b** (closed-loop controller), (iii) **ANN** (minimal HyperNEAT), and (iv) **SUPG**⁴. Performance and evolvability analysis are reported for the best individual of each replicate, selected to have the highest fitness in the population, and with an angle of trajectory in the range of $\pm 1^\circ$ (simulation source code can be downloaded from http://pages.isir.upmc.fr/evorob_db).

6.1 Performance

In all five encodings, the performance of the best individuals rapidly increased with a quasi-stable equilibrium being reached with less than 5,000 generations of selection (Fig. 6a). Additionally, individuals with evolved CPGs

(Direct, CPG and CPG-f/b) converged more rapidly as compared to those encoded with the ANN and SUPG schemes (Fig. S4, generations 0 to 8,000). After 8,000 generations, the performance in forward displacement of the Direct, CPG, CPG-f/b, ANN and SUPG encodings was 1.92 ± 0.19 , 1.79 ± 0.08 , 1.68 ± 0.13 , 2.93 ± 1.60 and 2.78 ± 1.43 m, respectively (Median \pm IQR, see Fig. 6b, Kruskal-Wallis test: $p < 0.001$). The ANN and SUPG schemes achieved the highest performance values across all five encodings ($d.f. = 38$, all $p < 0.001$, using Matlab's Mann-Whitney test, which is the default statistical test unless otherwise specified), but with no significant difference in performance between them. Furthermore, amongst the Direct, CPG, and CPG-f/b encodings, a decrease in performance was detected with a generative encoding, and with the inclusion of a feedback mechanism ($d.f. = 38$, $p < 0.001$). However, this decrease in performance was small, and did not exceed 12% (detailed statistical comparison provided in Table S1).

Importantly, intrinsic inter-encoding differences existed in the frequencies of oscillation governing leg actuation. The frequency of the CPG oscillations was prefixed at 1 Hz, irrespective of the sensory feedback provided, and the direct or generative nature of the encoding. By contrast, the individuals evolved with ANN and SUPG schemes were capable of expressing higher frequency oscillations (2.44 ± 1.95 Hz for ANN, and 3.81 ± 0.73 Hz for SUPG), and the frequency of the gait may itself be under selection. *Consequently, an assessment of the encodings solely on the basis of the performance is biased, and other measures are needed to compare encodings.*

6.2 Evolvability analysis

The evolvability provided by the encoding schemes is analyzed by mutating the best individual of each replicate at generation 8,000, and reporting the following: (i) The proportion decrease in performance consequent to the mutation (eq. 1); and (ii) The gait diversity, computed as the mutual information between gait vectors of the original and mutated individual (eq. 2c). The individual is mutated at a predetermined mutation rate as used during selection, in 1,000 separate and independent instances. Finally, a kernel density estimation⁵ (Scott, 2009) is used to visualize the resultant landscape of 20,000 data points (1,000 mutations \times 20 replicates), pooled together from all replicates.⁶

Distinct evolvability signatures were exhibited by the Direct, CPG, CPG-f/b, ANN, and SUPG encoding schemes, after 8,000 generations of selection (see Fig. S2). In evolvability analysis with a direct encoding, a conservative exploration of the phenotype, limited to solutions close to the unmutated individuals was found (11.9% and 0.33, median decrement in performance and gait diversity, respectively). A generative encoding of the CPG model had only a minor effect on the evolvability provided (performance decrement of 12.7% and gait diversity of 0.36). The inclusion of a feedback mechanism in the generatively encoded CPG model resulted in more diverse gaits (0.74), but with not much change in the performance loss (16.6%) following mutations. By contrast, the generative encoded ANN and SUPG schemes were much more aggressive in the exploration of the phenotypic land-

⁵A commonly used non-parametric technique for the estimation of the probability density function of a random variable, from a finite data sample.

⁶Bivariate density estimation, with Gaussian type kernels over a grid of 100×100 equidistant points.

⁴A single evolution replicate required about 24h of computational time on a 8-cores Intel Xeon E5520 at 2.27 Ghz.

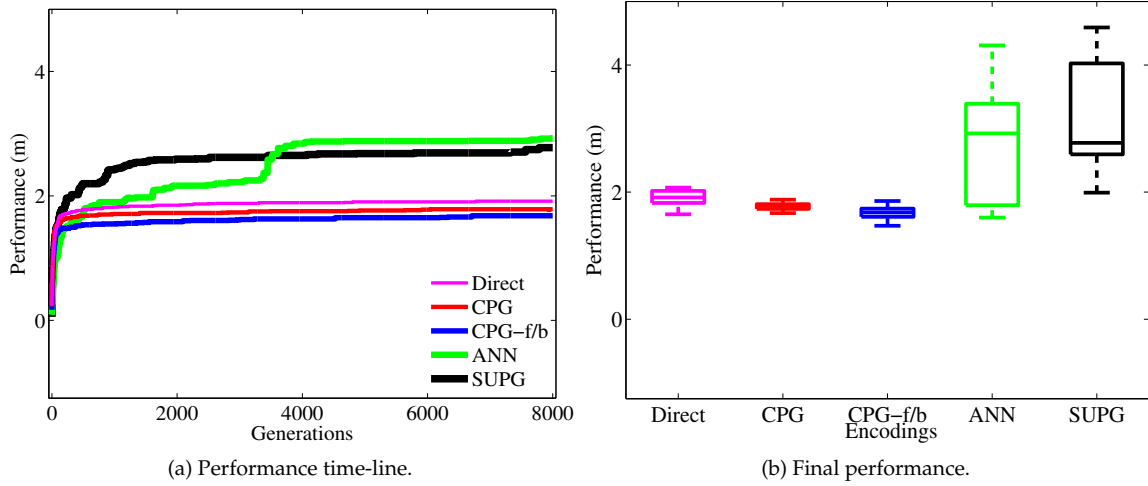


Figure 6: Performance in forward displacement for the Direct, CPG, CPG-f/b, ANN and SUPG encoding schemes: (a) Median performance for 8,000 generations of selection; and (b) Performance of encodings at generation 8,000.⁷

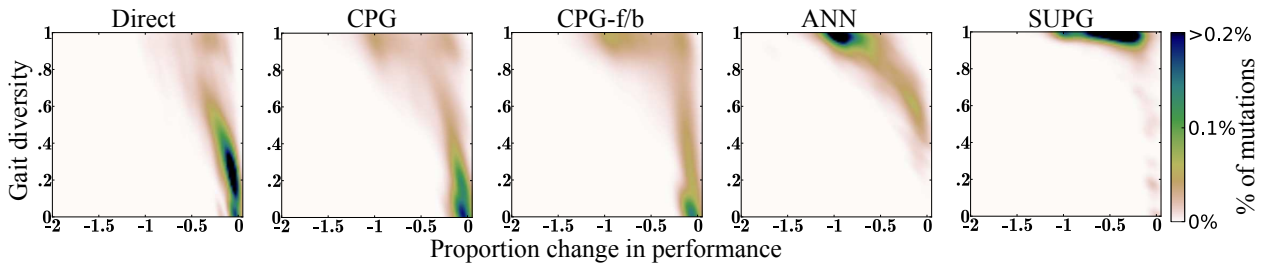


Figure 7: Evolvability signatures for different encodings: Gait diversity and the proportion decrease in performance, following 20,000 independent mutations of the best individuals, for the Direct, CPG, CPG-f/b, ANN, and SUPG encoding schemes after 8,000 generations of selection, pooled from all 20 replicates. In good evolvability signatures, mutations are located in the upper-right corner of the signature space, indicating high gait diversity, and a robustness to deleterious and lethal mutations.

scope, with the gait diversity of mutated individuals at 0.95 for ANN, and 0.99 for the SUPG encodings. However, differences existed in the severity of negative effects of mutations amongst the two encoding schemes. The ANN encoded individuals were sensitive to the effects of deleterious mutations, resulting in a 78.9% drop in performance. In comparison, individuals evolved with the SUPG encoding were much more resilient to the negative effects of mutations, with a smaller decrement of 43.1% in performance following mutation.

The evolvability provided by the encodings is further analyzed by computing the number of mutations in the evolvability signature (Fig. S2), that are both non-lethal and result in diverse locomotion gaits. Mutations are classified as lethal if they result in performance decrement in excess of 100% ($f_1 < -1$, see eq. 1), corresponding to the failure of any forward movement by the robot. Similarly, a mutation is considered to generate a diverse gait, if the inter-gait diversity exceeds 0.5 ($f_2 > 0.5$, see eq. 2c).

The proportion of viable and diverse-gait generating mutations was affected by the encoding scheme (see Fig. S1 and Table S2, Kruskal-Wallis test: $p < 0.001$). Across the five encodings, the SUPG scheme was most efficient at generating such mutations ($d.f. = 38$, all $p < 0.001$). Both the ANN and the CPG-f/b encodings led to an intermediate number of viable and diversity generating mutations (ANN significantly higher than CPG-f/b, and both different from all other encodings, all $p < 0.001$). The lowest mutation count was

achieved by the Direct and CPG encoding schemes (not significantly different from each other $p = 0.03$, but different from the three other encodings, all three $p < 0.001$). Thus, across the five encoding schemes, the SUPG approach provided the highest evolvability, with the capability to explore very different but viable gaits. Additionally, with a stricter definition of viable mutants resulting in no more than 50% drop in performance, the SUPG still achieved the highest beneficial mutation count (all $p < 0.001$), with no difference between the other four encodings (see Fig. S1).

6.2.1 Evolvability under varying mutation intensities

In this section, we study the sensitivity of our signature of the evolvability provided by the Direct, CPG, CPG-f/b, ANN and SUPG encodings with respect to the parameters of the variation operator used to generate mutants. The main questions are, if and how differences in the mutation operator affect our conceived signature of evolvability? We ran a series of experiments to access evolvability, with genetic mutants generated at different intensities. Mutations were considered at the standard mutation rate and mutation step-size

⁷On each box, the mid-line marks the median, and the box extends from the lower to upper quartile below and above the median. Whisker outside the box generally indicate the maximum and minimum values, except in case of outliers, which are shown as crosses. Outliers are data points outside of 1.5 times the interquartile range from the border of the box.

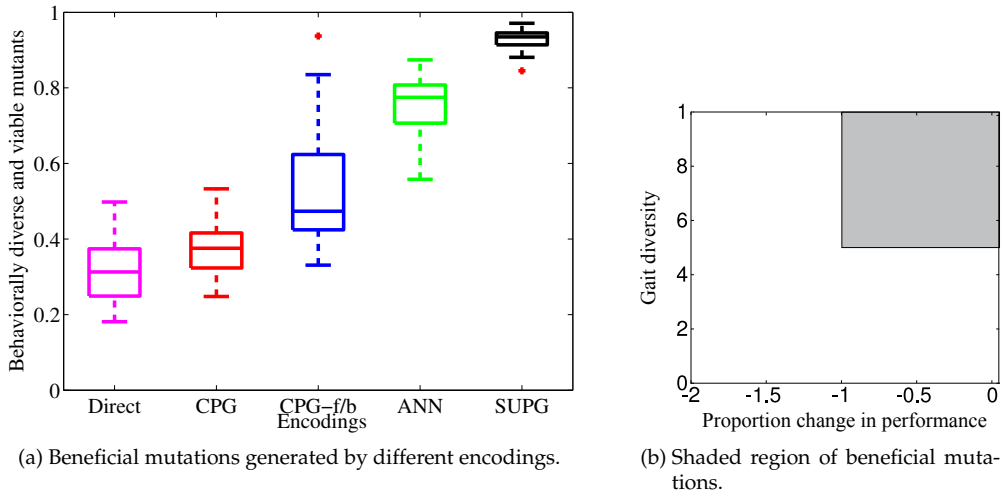


Figure 8: Proportion of viable mutants with gait diversity in excess of 0.5, from 1,000 independent mutations of the best individuals at generation 8,000 in each of 20 replicates, for the Direct, CPG, CPG-f/b, ANN and SUPG encoding schemes. These beneficial mutations are counted from the shaded region of each encoding’s evolvability signature.

as used during selection (medium intensity), and at a four-fold decrease (low intensity) and a four-fold increase (high intensity) of the standard mutation operator parameters of both rate and step-size.

In order to analyze the effect of the variation operator on our signature of evolvability, in Fig. 9 we outline the perimeter of the evolvability signatures generated from low, medium and high intensity mutations (see Fig. S2 for interior of signature). In all five encodings, the distribution of mutants shift towards more diverse gaits and is accompanied by larger loss in performance, following increments in the mutation intensity. A 16-fold increment in the mutation rate and step-size (low to high mutation intensity) resulted in a 0.94, 0.96 and 0.98 difference in mutant gait diversity for the Direct, CPG and CPG-f/b encodings respectively (see Fig. S3a, b and c). By contrast, the ANN and SUPG schemes achieved highly diverse gaits at the lowest mutation intensity (0.94 for ANN and 0.99 for SUPG), and an increment from low to high mutation intensity only resulted in a 6.0% for ANN, and 0.4% for SUPG further increase in the mutated gaits diversity (Fig. S3d and e). The increment from low to high mutation intensity also resulted in a drop in mutant performance of 63.7%, 90.5%, 93.5%, 98.6% and 54.9%, for the Direct, CPG, CPG-f/b, ANN, and SUPG encodings respectively (Fig. S3f-j). In summary, the SUPG encoding facilitates the exploration of diverse gaits even when mutants are generated at a low mutation intensity. Furthermore, across all five encodings, the SUPG scheme provides the most resilience against the deleterious nature of high intensity mutations.

6.3 Damage recovery

The significance of our evolvability signatures of the Direct, CPG, CPG-f/b, ANN and SUPG encodings was investigated by analyzing the adaptation of the evolved robot’s gait, following the removal of one or more of its legs. We expect that for the encodings registering a better evolvability signature, the corresponding evolved individuals would require fewer generations to recover an effective walking gait.

In these experiments, the new (damage recovery) populations were comprised of 100 mutated individuals of the best individual of each replicate at generation 8,000 of selection.

In separate preliminary experiments, the use of the entire population at generation 8,000 (instead of the best individuals) did not change our results of the adaptability provided by different encodings. Individuals in the damage recovery population were mutated at the standard mutation rate and step size used during selection. A further 10,000 generations of artificial selection was conducted on the populations of amputee hexapods for each of the following three damage scenarios: (i) an asymmetrical damage, following the removal of one leg of the robot (leg 1, Fig. 10a); (ii) a symmetrical damage occurs, wherein the two middle legs on either side of the robot are removed (legs 1 and 4, Fig. 10b); and (iii) a highly asymmetrical damage occurs consequent to the removal of the middle leg on one side and the rear leg on the opposing side of the hexapod (legs 1 and 3, Fig. 10c). The number of generations required to regain an effective gait and the proportion of the original performance (undamaged robot’s performance at generation 8,000) recovered for each of the three damage scenarios is analyzed.

In the 10,000 generations of selection, all the five encodings were capable of recovering a majority of their original performance in forward displacement, irrespective of the damage to the hexapod robot (see Fig. S4). After 10,000 generations post robot damage, the Direct, CPG, CPG-f/b, ANN and SUPG schemes all recovered the highest proportion of their original performance in the first damage scenario (0.89 ± 0.07 , Median \pm IQR across all encodings), followed by an intermediate recovery in the second (0.78 ± 0.07), and third (0.72 ± 0.05) scenarios (Fig. 11a, b and c, Kruskal-Wallis test: $d.f. = 4$, $p < 0.001$ for first two scenarios, and $p = 0.018$ for third scenario). Across all five encodings, the SUPG was most efficient in recovering its original performance in the first (1.02 ± 0.14) and second (0.99 ± 0.23) scenarios (for both, all $p < 0.001$), while in both scenarios no significant difference in recovery was registered between the remaining four encodings (Fig. 11a and b). In the third scenario which was the hardest, the SUPG again achieved the highest performance recovery (0.88 ± 0.3), although performance was no longer significantly different between encodings (Fig. 11c).

In order to analyze the time required by damaged hexapods to recover an effective gait, in Fig. 12 we have plot-




Figure 9: Contours of the evolvability signature for the Direct, CPG, CPG-f/b, ANN, and SUPG encoding schemes with mutations generated at low (L), medium (M), and high (H) intensity, of the best individuals at generation 8,000 of selection. The low, medium, and high mutation intensities were quantitatively 0.25, 1, and 4 times respectively, the mutation rate and step-size used during selection. Individual contours encompass a 0.025% or higher density of mutants.




Figure 10: The three damage scenarios imposed on the hexapod robot (undamaged robot in Fig. 1).

ted the number of generations required to restore 85% of the original performance in forward displacement. Across the three robot-damage scenarios, amputee hexapods in the first scenario achieved the highest recovery rate (13 ± 3.25 of 20 replicates restored performance in 10,000 generations, across all encodings), followed by an intermediate recovery rates in the second (4 ± 7.75 replicates), and third scenarios (1 ± 8.25 replicates). In the first two scenarios (Kruskal-Wallis test: $d.f. = 4$, both $p < 0.001$), the SUPG encoded individuals recovered at least an order of magnitude faster (373 and 957 generations in scenarios 1 and 2, respectively) than individuals with the Direct, CPG, CPG-f/b and ANN encodings (see Fig. 12a and b, all $p < 0.001$). In both scenarios, no significant difference in recovery was registered between these four encoding schemes. In the third scenario (Kruskal-Wallis test: $d.f. = 4$, $p < 0.001$), the SUPG encoded amputee hexapods continued to exhibit the fastest recovery (8466.5 generations), although it was no longer significantly different from ANN encoded hexapods (see Fig. 12c, $p = 0.25$). Furthermore, in this scenario, all three CPG based encodings (Direct, CPG, and CPG-f/b) performed poorly, with only one replicate making the 85% mark in the 10,000 generations of selection. In summary, across all five encoding schemes, the SUPG encoded individuals had the fastest recovery, despite being in increasingly difficult robot-damage scenarios wherein most of the individuals encoded by the other encodings failed to recover an effective walking gait (detailed comparison of damage recovery in Tables S3, S4 and S5).

7 Discussion

In this study, we introduce a methodology to compare genetic encoding schemes, utilizing our novel approach to signaturize evolvability. The evolvability provided by the generatively encoded CPG, CPG-f/b, ANN (Clune et al., 2011; Yosinski et al., 2011) and SUPG (Morse et al., 2013) schemes and a direct encoding, is characterized by quantifying both, (i) the quality of the genetic mutations, and (ii) the quantity of behavioral variation generated from genetic change. Our evolvability based inter-encoding comparison methodology provides insights on the adaptability of evolved individuals to novel scenarios. To this effect, the significance of our conceived signature of evolvability is evaluated by the individual adaptation response to morphological changes previously unencountered by the hexapod robot.

Our results revealed a direct relationship between the estimated evolvability provided by the encodings, and the capability of the evolved individuals to adapt to severe changes in morphology, simulated by the amputation of one or more of the hexapod legs. Amongst the five encodings evaluated, the SUPGs (Morse et al., 2013) had the best evolvability signature, and their encoded individuals were also foremost to recover following sustained damages. In both the easy and the intermediate robot-damage scenarios (scenarios 1 and 2), the SUPG encoded individuals were capable of recovering 85% of their performance on the undamaged robot in all but two replicates, and did so more than an order of magnitude faster than the other four encodings. Furthermore, even in the most difficult robot-damage scenario (scenario 3), the SUPG scheme achieved the fastest recovery in the majority of the evaluated replicates.

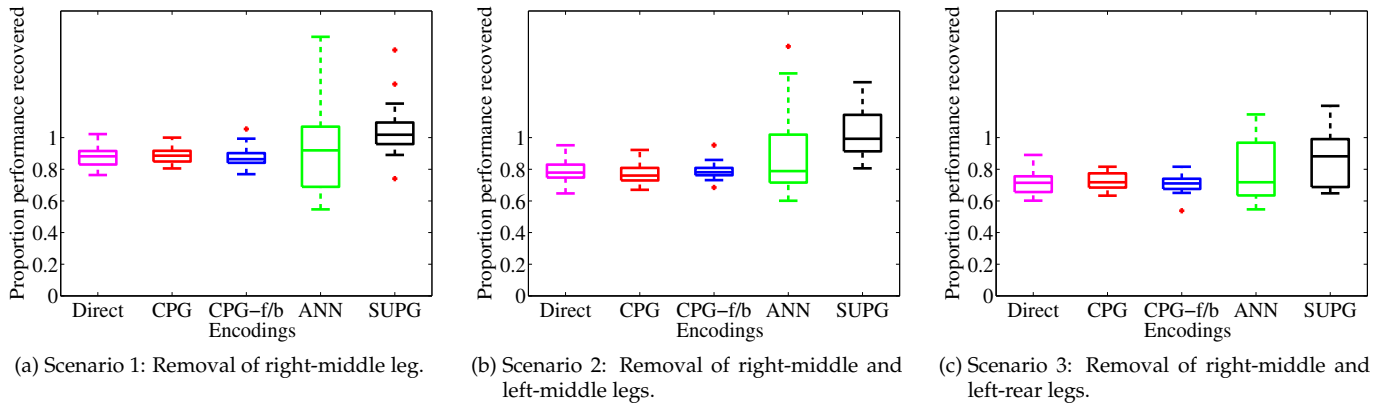


Figure 11: Proportion of the original performance in forward displacement restored 10,000 generations after the three damage scenarios, across 20 replicates, for the Direct, CPG, CPG-f/b, ANN and SUPG encoding schemes.

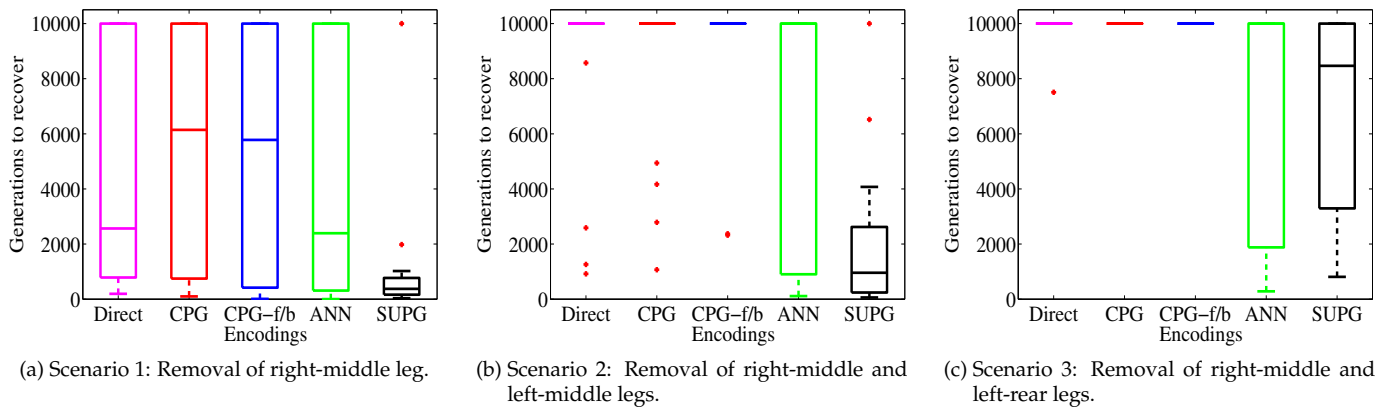


Figure 12: The number of generations of selection required to restore 85% of the original performance of the undamaged hexapod in forward displacement across 20 replicates, for the Direct, CPG, CPG-f/b, ANN, and SUPG encoding schemes. In replicates unable to attain the 85% mark, the recovery time was set to the upper limit of 10,000 generations.

The ANN encoding scheme (minimal HyperNEAT) was capable of producing highly diverse hexapod gaits, following genetic mutations. The high behavioral diversity generated is consistent with the earlier use of this encoding to evolve gaits for the QuadraBot robot (Clune et al., 2009a, 2011). However, in our ANN implementation, most of genetic mutations producing diverse gaits were highly deleterious, and resulted in little forward hexapod movement. Consequently, an estimate of evolvability solely on the basis of the generated behavioral diversity (Lehman and Stanley, 2011b; Reisinger et al., 2005; Lehman and Stanley, 2013) is not reliable, and both the quality (individual viability) and quantity of phenotypic variation consequent to genetic change is required to characterize evolvability. Furthermore, the poor evolvability signature for the ANN encoding scheme is reflected in its poor recovery from sustained robot damages.

In studies on evolution of multilegged robot locomotion, the generative encodings exploit the symmetry of the robot morphology to generate regular and coordinated gait patterns that often outperform gaits evolved with direct encodings (e.g., Clune et al. (2011)). Generative encodings also facilitate scalability, wherein evolution in the low-dimensional genetic search space is capable of evolving complex phenotypes comprising of many more dimensions (Stanley and Mikkulainen, 2003). However, no difference in performance was registered between the directly and generatively en-

coded CPGs for our hexapod locomotion problem, perhaps consequent to the already low-dimensional search space for the directly encoded locomotion controllers. For example, our directly encoded CPGs for hexapod locomotion comprise 23 amplitude and phase bias parameters, in contrast to the 800 *Fixed-Topology NEAT* (FT-NEAT) encoded neural weight parameters for quadruped locomotion controllers (Clune et al., 2009a). Thus, the potential benefits of phenotypic scalability in utilizing generative encoding schemes are reduced in our study.

For our signature of evolvability, we mutated the individuals with a predetermined mutation rate, tuned to allow a speedy convergence of the evolved solutions. This is a critical consideration as variations to the mutation rate can affect the viability and gait diversity of generated mutants. A comparison of the evolvability provided by encodings at low and high mutation intensities suggests that with an increase in mutation rate, the peak of the distribution of mutants shifts towards more diverse gaits with a larger decrease in task performance. However, the overall shape of the distribution, highlighting desirable regions of the evolvability landscape, remained the same for all five encodings. Importantly, across the five encodings, the SUPG scheme continued to provide the highest resilience to deleterious genetic change, despite a 16-fold increase in mutation intensity.

The high evolvability and rapid recovery provided by the

SUPGs may be consequent to the closed-loop control ingrained in the encoding scheme. Such a feedback mechanism provides an adaptive period of the SUPG oscillations, that can be adjusted to the new step size of a gait more appropriate for a four or five legged robot, after the hexapod has suffered damages. Alternatively, the better performance of SUPGs may be consequent to the open-ended encoding of control signals by these oscillators. In contrast to the simple sinusoidal waves of the CPG-based schemes wherein only the signal amplitude and phase difference is encoded, no constraints are imposed on the CPPNs encoding the SUPG output signals. The resulting unconstrained encoding may help for example in adjusting the duty ratio for each oscillator to match the new swing and stance phase durations of the remaining undamaged hexapod legs. Furthermore, since the inclusion of feedback in the CPG encoding (CPG-f/b) registered no improvement in task-performance, evolvability or damage recovery, a combination of the SUPG closed-loop system and the open-ended encoding of its oscillatory signal may be responsible for its high performance and adaptive capabilities.

In our evolvability signatures, the phenotypic variation from genetic change was associated with the mutual information between hexapod gaits. The diversity may also be computed for the gaits of bipedal and quadrupedal type of robots. Similarly, behavioral diversity may be computed for other benchmark problems such as, the final position of a robot in a maze navigation task (Lehman and Stanley, 2011b; Mouret and Doncieux, 2012), the final positions of balls in an arena for the robot ball-collecting task (Doncieux and Mouret, 2009, 2010), and a vector of board piece moves in game playing tasks (Reisinger and Miikkulainen, 2007; Gauci and Stanley, 2010). Consequently, our approach to estimate evolvability is easily applicable to a wide range of tasks, commonly used in evolutionary robotics experiments.

The systematic building and organizing of knowledge, a requirement in any scientific discipline, can not be achieved without a wide assortment of quantifiable measures to compare and contrast concepts, hypotheses, testable explanations and predictions. In the field of evolutionary robotics, task-performance has been prominently and often solely used as such a quantifiable measure to form links between the different available evolutionary systems. However, the evaluated fitness by its very nature, is limited to the specific problem for which the individual solutions are tested. By contrast, an estimate of evolvability facilitated by an evolutionary process, may be applicable to a much broader scope of scenarios, allowing the formation of a more generic relationship between existing evolutionary system implementations. An evolvability-based approach of comparison between evolutionary processes may help to extrapolate to un-evaluated problem regions in-between existing evaluations, and consequently better lend itself to building a strong theoretical foundation for future research in the field.

8 Conclusion

A novel approach to compare encoding schemes is proposed in this paper, using evolvability instead of task-fitness, and demonstrated for the problem of evolving locomotion gaits for a hexapod robot. From the different direct and generative encodings evaluated, the single-unit pattern generator had the best evolvability signature. Furthermore, its en-

coded individuals were foremost in adapting to various new and challenging scenarios. Our biologically-relevant “fingerprint” of the evolvability provided by an encoding accounts for both the quantity and quality of phenotypic variation resulting from genetic change, and results in distinct evolvability signatures for the different encodings, despite similar task-fitness values. These evolvability signatures allow for a discrimination between encoding schemes, based on their capability of producing highly adaptable individuals for novel, and importantly, a priori unknown scenarios. In summary, our evolvability signatures for encoding comparison serves as a comprehensive analysis tool beyond the standard task-fitness performance benchmarks commonly employed in the evolutionary robotics field.

Acknowledgment

This study was supported by the ANR Creadapt project (ANR-12-JS03-0009).

Supplemental data

Movies of the hexapod robot walking behaviors are available online at <http://goo.gl/uyY1RX>.

References

- P. Alberch. From genes to phenotype: dynamical systems and evolvability. *Genetica*, 84(1):5–11, 1991.
- B. Alberts, D. Bray, J. Lewis, M. Raff, K. Roberts, and J. D. Watson. *Molecular biology of the cell*. Garland, New York, 1994.
- L. Altenberg. The evolution of evolvability in genetic programming. *Advances in Genetic Programming*, 3:47–74, 1994.
- S. Aoi and K. Tsuchiya. Locomotion control of a biped robot using nonlinear oscillators. *Autonomous Robots*, 19(3):219–232, 2005.
- S. J. Arnold. Constraints on phenotypic evolution. *American Naturalist*, pages S85–S107, 1992.
- F. A. Azevedo, L. R. Carvalho, L. T. Grinberg, J. M. Farfel, R. E. Ferretti, R. E. Leite, R. Lent, S. Herculano-Houzel, et al. Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain. *Journal of Comparative Neurology*, 513(5):532–541, 2009.
- T. D. Barfoot, E. J. Earon, and G. M. D’Eleuterio. Experiments in learning distributed control for a hexapod robot. *Robotics and Autonomous Systems*, 54(10):864–872, 2006.
- R. D. Beer and J. C. Gallagher. Evolving dynamical neural networks for adaptive behavior. *Adaptive Behavior*, 1(1): 91–122, 1992.
- P. Bentley and S. Kumar. Three ways to grow designs: A comparison of evolved embryogenies for a design problem. In *Proc. GECCO*, pages 35–43. ACM, New York, NY, 1999.

- J. Bongard, V. Zykov, and H. Lipson. Resilient machines through continuous self-modeling. *Science*, 314(5802):1118–1121, 2006.
- J. C. Bongard. Evolutionary robotics. *Communications of the ACM*, 56(8):74–83, 2013.
- J. C. Bongard and R. Pfeifer. Repeated structure and dissociation of genotypic and phenotypic complexity in artificial ontogeny. In *Proc. GECCO*, pages 829–836. ACM, New York, NY, 2001.
- S. B. Carroll. *Endless forms most beautiful: The new science of evo devo and the making of the animal kingdom*. Number 54. WW Norton & Company, 2005.
- N. Cheney, R. MacCurdy, J. Clune, and H. Lipson. Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. In *Proc. GECCO*, pages 167–174. ACM, New York, NY, 2013.
- J. Clune and H. Lipson. Evolving 3D objects with a generative encoding inspired by developmental biology. *SIGEVolution*, 5(4):2–12, 2011. ISSN 1931-8499.
- J. Clune, B. E. Beckmann, C. Ofria, and R. T. Pennock. Evolving coordinated quadruped gaits with the hyperneat generative encoding. In *Proc. CEC*, pages 2764–2771. IEEE Press, Piscataway, NJ, 2009a.
- J. Clune, C. Ofria, and R. T. Pennock. The sensitivity of hyperneat to different geometric representations of a problem. In *Proc. GECCO*, pages 675–682, New York, NY, USA, 2009b. ACM. ISBN 978-1-60558-325-9.
- J. Clune, K. O. Stanley, R. T. Pennock, and C. Ofria. On the performance of indirect encoding across the continuum of regularity. *IEEE Transactions on Evolutionary Computation*, 15(3):346–367, 2011.
- J. Clune, J.-B. Mouret, and H. Lipson. The evolutionary origins of modularity. *Proceedings of the Royal Society B: Biological Sciences*, 280(1755):20122863, 2013.
- T. M. Cover and J. A. Thomas. *Elements of information theory*. John Wiley & Sons, Inc., 1991.
- A. Crespi, K. Karakasiliotis, A. Guignard, and A. Ijspeert. Salamandra robotica II: An amphibious robot to study salamander-like swimming and walking gaits. *IEEE Transactions on Robotics*, 29(2):308–320, April 2013. ISSN 1552-3098.
- A. Cully and J.-B. Mouret. Behavioral repertoire learning in robotics. In *Proc. GECCO*, pages 175–182. ACM, New York, NY, 2013.
- R. Dawkins. *The extended phenotype: The long reach of the gene*. Oxford University Press, 1999.
- K. Deb. *Multi-objective optimization using evolutionary algorithms*. Wiley, 2001. ISBN 978-0-471-87339-6.
- K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002.
- A. Devert, N. Bredeche, and M. Schoenauer. Robustness and the halting problem for multicellular artificial ontogeny. *IEEE Transactions on Evolutionary Computation*, 15(3):387–404, June 2011.
- S. Doncieux and J.-B. Mouret. Behavioral diversity with multiple behavioral distances. In *Proc. CEC*, pages 1427–1434. IEEE Press, 2013.
- S. Doncieux and J.-B. Mouret. Single step evolution of robot controllers for sequential tasks. In *Proc. GECCO*, pages 1771–1772. ACM, New York, NY, 2009. ISBN 978-1-60558-325-9.
- S. Doncieux and J.-B. Mouret. Behavioral diversity measures for evolutionary robotics. In *Proc. CEC*, pages 1–8. IEEE Press, Piscataway, NJ, July 2010.
- S. Doncieux and J.-B. Mouret. Beyond black-box optimization: a review of selective pressures for evolutionary robotics. *Evolutionary Intelligence*, 7(2):71–93, 2014. ISSN 1864-5909.
- M. Dorigo, E. Tuci, R. Groß, V. Trianni, T. H. Labella, S. Nouyan, C. Ampatzis, J.-L. Deneubourg, G. Baldassarre, S. Nolfi, F. Mondada, D. Floreano, and L. M. Gambardella. The swarm-bots project. In *Proc. SAB (SAB-04 Workshop on Swarm Robotics)*, pages 31–44, Berlin, Heidelberg, 2005. Springer-Verlag.
- A. Eiben. Grand challenges for evolutionary robotics. *Frontiers in Robotics and AI*, 1(4), 2014.
- S. J. Freeland and L. D. Hurst. The genetic code is one in a million. *Journal of Molecular Evolution*, 47(3):238–248, 1998.
- R. M. Friedberg. A learning machine: Part II. *IBM Journal of Research and Development*, 3:183–191, 1959.
- A. Frigon and S. Rossignol. Experiments and models of sensorimotor interactions during locomotion. *Biological Cybernetics*, 95(6):607–627, 2006.
- Y. Fukuoka, H. Kimura, and A. H. Cohen. Adaptive dynamic walking of a quadruped robot on irregular terrain based on biological concepts. *International Journal of Robotics Research*, 22(3-4):187–202, 2003.
- C. Furusawa and K. Kaneko. Emergence of multicellular organisms with dynamic differentiation and spatial pattern. *Artif. Life*, 4(1):79–93, 1998.
- J. Gauci and K. O. Stanley. Autonomous evolution of topographic regularities in artificial neural networks. *Neural Comput.*, 22(7):1860–1898, 2010.
- S. Gavrillets. Evolution and speciation on holey adaptive landscapes. *Trends in Ecol. & Evol.*, 12(8):307–312, 1997.
- J. Gerhart and M. Kirschner. The theory of facilitated variation. *Proc. Natl. Acad. Sci. USA.*, 104(Suppl 1):8582–8589, 2007.
- A. D. Goldberg, C. D. Allis, and E. Bernstein. Epigenetics: a landscape takes shape. *Cell*, 128(4):635–638, 2007.
- R. K. Grosberg and R. R. Strathmann. The evolution of multicellularity: a minor major transition? *Annual Review of Ecology, Evolution, and Systematics*, pages 621–654, 2007.

- F. Gruau. Automatic definition of modular neural networks. *Adapt. Behav.*, 3(2):151–183, 1994.
- B. Hayes. Undisciplined science. *American Scientist*, 92(4):306–310, 2004.
- G. S. Hornby. Functional scalability through generative representations: the evolution of table designs. *Environment and Planning B*, 31(4):569–588, 2004.
- G. S. Hornby. Measuring, enabling and comparing modularity, regularity and hierarchy in evolutionary design. In *Proc. GECCO*, pages 1729–1736. ACM, New York, NY, 2005.
- G. S. Hornby and J. B. Pollack. Creating high-level components with a generative representation for body-brain evolution. *Artif. Life*, 8(3):223–246, 2002.
- G. S. Hornby, H. Lipson, and J. B. Pollack. Generative representations for the automated design of modular physical robots. *IEEE Trans. Robot. Automat.*, 19(4):703–719, 2003.
- G. S. Hornby, S. Takamura, T. Yamamoto, and M. Fujita. Autonomous evolution of dynamic gaits with two quadruped robots. *IEEE Trans. Robot.*, 21(3):402–410, 2005.
- G. S. Hornby, A. Globus, D. S. Linden, and J. D. Lohn. Automated antenna design with evolutionary algorithms. In *Proc. AIAA Space Conference*, pages 1–8. American Institute of Aeronautics and Astronautics, Reston, VA, 2006.
- T. Hu and W. Banzhaf. Evolvability and speed of evolutionary algorithms in light of recent developments in biology. *J. Artif. Evol. App.*, 2010:1:1–1:28, January 2010. ISSN 1687-6229.
- E. Ibáñez-Marcelo and T. Alarcón. The topology of robustness and evolvability in evolutionary systems with genotype–phenotype map. *Journal of Theoretical Biology*, 356:144–162, 2014.
- A. J. Ijspeert. Central pattern generators for locomotion control in animals and robots: a review. *Neural Networks*, 21(4):642–653, 2008.
- A. J. Ijspeert, A. Crespi, D. Ryczko, and J.-M. Cabelguen. From swimming to walking with a salamander robot driven by a spinal cord model. *Science*, 315(5817):1416–1420, 2007.
- A. Iscen, A. Agogino, V. SunSpiral, and K. Tumer. Controlling tensegrity robots through evolution. In *Proc. GECCO*, pages 1293–1300. ACM, New York, NY, 2013.
- E. J. Javaux, C. P. Marshall, and A. Bekker. Organic-walled microfossils in 3.2-billion-year-old shallow-marine siliclastic deposits. *Nature*, 463(7283):934–938, 2010.
- Y. Jin and J. Branke. Evolutionary optimization in uncertain environments—a survey. *IEEE Transactions on Evolutionary Computation*, 9(3):303–317, June 2005.
- W. Johannsen. The genotype conception of heredity. *American Naturalist*, 45(531):129–159, 1911.
- S. Kajita and B. Espiau. Springer handbook of robotics. In B. Siciliano and O. Khatib, editors, *Springer handbook of robotics*, chapter 16. Springer Berlin Heidelberg, 2008.
- J. Kaplan. The end of the adaptive landscape metaphor? *Biology & Philosophy*, 23(5):625–638, 2008.
- S. A. Kauffman. *The origins of order: Self-organization and selection in evolution*. Oxford university press, 1993.
- S. Kernbach. *Handbook of Collective Robotics: Fundamentals and Challenges*. Pan Stanford, 2013. ISBN 9789814316422.
- M. Kirschner and J. Gerhart. Evolvability. *Proc. Natl. Acad. Sci. USA.*, 95(15):8420–8427, 1998.
- J. Kodjabachian and J.-A. Meyer. Evolution and development of neural controllers for locomotion, gradient-following, and obstacle-avoidance in artificial insects. *IEEE Transactions on Neural Networks*, 9(5):796–812, 1998.
- M. Komosiński and A. Rotaru-Varga. Comparison of different genotype encodings for simulated three-dimensional agents. *Artificial Life*, 7(4):395–418, 2001.
- S. Koos, A. Cully, and J.-B. Mouret. Fast damage recovery in robotics with the T-resilience algorithm. *Int. J. Robot. Res.*, 32(14):1700–1723, 2013.
- J. R. Koza. *Genetic programming: on the programming of computers by means of natural selection*, volume 1. MIT press, 1992.
- A. Kraskov, H. Stögbauer, and P. Grassberger. Estimating mutual information. *Physical Review E*, 69(6):066138, 2004.
- N. Kubota. Computational intelligence for structured learning of a partner robot based on imitation. *Information-Sciences*, 171(4):403 – 429, 2005.
- S. Lee, J. Yosinski, K. Glette, H. Lipson, and J. Clune. Evolving gaits for physical robots with the hyperneat generative encoding: The benefits of simulation. In A. Esparcia-Alcázar, editor, *Applications of Evolutionary Computation*, volume 7835 of LNCS, pages 540–549. Springer Berlin Heidelberg, 2013. ISBN 978-3-642-37191-2.
- J. Lehman and K. O. Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2):189–223, 2011a.
- J. Lehman and K. O. Stanley. Improving evolvability through novelty search and self-adaptation. In *Proc. CEC*, pages 2693–2700. IEEE Press, Piscataway, NJ, 2011b.
- J. Lehman and K. O. Stanley. Evolvability is inevitable: Increasing evolvability without the pressure to adapt. *PloS One*, 8(4):e62186, 2013.
- M. A. Lewis, A. H. Fagg, and A. Solidum. Genetic programming approach to the construction of a neural network for control of a walking robot. In *Proc. IEEE ICRA*, pages 2618–2623. IEEE Press, Piscataway, NJ, 1992.
- H. Lipson. Principles of modularity, regularity, and hierarchy for scalable systems. *Journal of Biological Physics and Chemistry*, 7(4):125–128, 2007.
- H. Liu and H. Iba. A hierarchical approach for adaptive humanoid robot control. In *Proc. CEC*, volume 2, pages 1546–1553 Vol.2. IEEE Press, Piscataway, NJ, June 2004.

- P. Marrow, M. Heath, and I. I. Re. Evolvability: Evolution, computation, biology. In *Proc. GECCO (GECCO-99 Workshop on Evolvability)*, pages 30–33. ACM, New York, NY, 1999.
- D. M. McCandlish. Visualizing fitness landscapes. *Evolution*, 65(6):1544–1558, 2011.
- R. E. Michod and D. Roze. Cooperation and conflict in the evolution of multicellularity. *Heredity*, 86(1):1–7, 2001.
- E. Mjolsness, D. H. Sharp, and J. Reintz. A connectionist model of development. *Journal of Theoretical Biology*, 152(4):429–453, 1991.
- R. Morrison and K. De Jong. A test problem generator for non-stationary environments. In *Proc. CEC*, volume 3, pages 2047–2053, 1999.
- G. Morse, S. Risi, C. R. Snyder, and K. O. Stanley. Single-unit pattern generators for quadruped locomotion. In *Proc. GECCO*, pages 719–726. ACM, New York, NY, 2013.
- J.-B. Mouret and S. Doncieux. Encouraging behavioral diversity in evolutionary robotics: An empirical study. *Evol. Comput.*, 20(1):91–133, 2012.
- J.-B. Mouret. Novelty-based multiobjectivization. In S. Doncieux, N. Bredche, and J.-B. Mouret, editors, *New Horizons in Evolutionary Robotics*, volume 341 of *Studies in Computational Intelligence*, pages 139–154. Springer Berlin Heidelberg, 2011.
- J.-B. Mouret, S. Doncieux, and B. Girard. Importing the computational neuroscience toolbox into neuro-evolution-application to basal ganglia. In *Proc. GECCO*, pages 587–594. ACM, New York, NY, 2010.
- S. Nolfi and D. Floreano. *Evolutionary Robotics. The Biology, Intelligence, and Technology of Self-organizing Machines*. MIT Press, 2001.
- A. S. Novozhilov, Y. I. Wolf, and E. V. Koonin. Evolution of the genetic code: partial optimization of a random code for robustness to translation error in a rugged fitness landscape. *Biology Direct*, 2(24):1–24, 2007.
- M. Pavlicev and G. P. Wagner. Coming to grips with evolvability. *Evol. Educ. Outreach*, 5(2):231–244, 2012.
- M. Pigliucci. Is evolvability evolvable? *Nature Reviews Genetics*, 9(1):75–82, 2008.
- W. B. Provine. *Sewall Wright and evolutionary biology*. University of Chicago Press, 1989.
- M. Raibert, K. Blankespoor, G. Nelson, R. Playter, et al. Bigdog, the rough-terrain quadruped robot. In *Proc. IFAC*, pages 10823–10825. IFAC Publications Office, Elsevier Ltd., Oxford, UK, 2008.
- J. Reisinger and R. Miikkulainen. Acquiring evolvability through adaptive representations. In *Proc. GECCO*, pages 1045–1052. ACM, New York, NY, 2007.
- J. Reisinger, K. O. Stanley, and R. Miikkulainen. Towards an empirical measure of evolvability. In *Proc. GECCO*, pages 257–264. ACM, New York, NY, 2005.
- G. Ren, W. Chen, S. Dasgupta, C. Kolodziejski, F. Wörgötter, and P. Manoonpong. Multiple chaotic central pattern generators with learning for legged locomotion and malfunction compensation. *InformationSciences*, 2014.
- L. Righetti and A. J. Ijspeert. Design methodologies for central pattern generators: An application to crawling humanoids. In *Proc. RSS*. MIT Press, Cambridge, 2006.
- L. Righetti and A. J. Ijspeert. Pattern generators with sensory feedback for the control of quadruped locomotion. In *Proc. IEEE ICRA*, pages 819–824. IEEE Press, Piscataway, NJ, 2008.
- S. Risi and K. O. Stanley. Confronting the challenge of learning a flexible neural controller for a diversity of morphologies. In *Proc. GECCO*, pages 255–262. ACM, New York, NY, 2013.
- M. S. Roulston. Estimating the errors on measured entropy and mutual information. *Physica D: Nonlinear Phenomena*, 125(3):285–294, 1999.
- J. Santos, R. Duro, J. Becerra, J. Crespo, and F. Bellas. Considerations in the application of evolution to the generation of robot controllers. *InformationSciences*, 133(34):127 – 148, 2001.
- D. W. Scott. *Multivariate density estimation: theory, practice, and visualization*, volume 383. John Wiley & Sons, Inc., 2009.
- C. W. Seys and R. D. Beer. Genotype reuse more important than genotype size in evolvability of embodied neural networks. In *Proc. ECAL*, pages 915–924. Springer Berlin Heidelberg, 2007.
- A. Sproewitz, R. Moeckel, J. Maye, and A. J. Ijspeert. Learning to move in modular robots using central pattern generators and online optimization. *International Journal of Robotics Research*, 27(3-4):423–443, 2008.
- K. Stanley and R. Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10(2):99–127, 2002.
- K. O. Stanley. Compositional pattern producing networks: A novel abstraction of development. *Genetic Programming and Evolvable Machines*, 8(2):131–162, 2007.
- K. O. Stanley and R. Miikkulainen. A taxonomy for artificial embryogeny. *Artif. Life*, 9(2):93–130, 2003.
- K. O. Stanley, D. B. D’Ambrosio, and J. Gauci. A hypercube-based encoding for evolving large-scale neural networks. *Artif. Life*, 15(2):185–212, 2009.
- G. Taga. Emergence of bipedal locomotion through entrainment among the neuro-musculo-skeletal system and the environment. *Physica D: Nonlinear Phenomena*, 75(1):190–208, 1994.
- D. Tarapore and J.-B. Mouret. Comparing the evolvability of generative encoding schemes. In *Proc. ALife*, volume 14, pages 1–8, 2014.
- R. A. Téllez, C. Angulo, and D. E. Pardo. Evolving the walking behaviour of a 12 DOF quadruped using a distributed neural architecture. In A. Ijspeert, T. Masuzawa, and S. Kusumoto, editors, *Biologically Inspired Approaches*

- to *Advanced Information Technology*, volume 3853 of LNCS, pages 5–19. Springer Berlin Heidelberg, 2006. ISBN 978-3-540-31253-6.
- M. Tomassini, S. Verel, and G. Ochoa. Complex-network analysis of combinatorial spaces: The NK landscape case. *Physical Review E*, 78(6):066114, 2008.
- P. Tonelli and J.-B. Mouret. On the relationships between generative encodings, regularity, and learning abilities when evolving plastic artificial neural networks. *PloS One*, 8(11):e79138, 2013.
- V. K. Valsalam and R. Miikkulainen. Modular neuroevolution for multilegged locomotion. In *Proc. GECCO*, pages 265–272. ACM, New York, NY, 2008.
- G. P. Wagner and L. Altenberg. Perspective: Complex adaptations and the evolution of evolvability. *Evolution*, 50(3):967–976, 1996.
- M. Waibel, L. Keller, and D. Floreano. Genetic team composition and level of selection in the evolution of cooperation. *IEEE Transactions on Evolutionary Computation*, 13(3):648–660, June 2009. ISSN 1089-778X.
- K. Weicker. Performance measures for dynamic environments. In *Parallel Problem Solving from Nature – PPSN VII*, pages 64–73. Springer Berlin Heidelberg, 2002.
- S. Wright. The roles of mutation, inbreeding, crossbreeding, and selection in evolution. In *Proc. of the Sixth International Congress on Genetics*, volume 1, pages 355–366, 1932.
- X. Yao. Evolving artificial neural networks. *Proceedings of the IEEE*, 87(9):1423–1447, 1999.
- J. Yosinski, J. Clune, D. Hidalgo, S. Nguyen, J. Zagal, and H. Lipson. Evolving robot gaits in hardware: the hyperneat generative encoding vs. parameter optimization. In *Proc. ECAL*, pages 890–897. MIT Press, Cambridge, 2011.
- V. Zykov, J. Bongard, and H. Lipson. Evolving dynamic gaits on a physical robot. In *Proc. GECCO*, volume 4. ACM, New York, NY, 2004.

Supplemental data

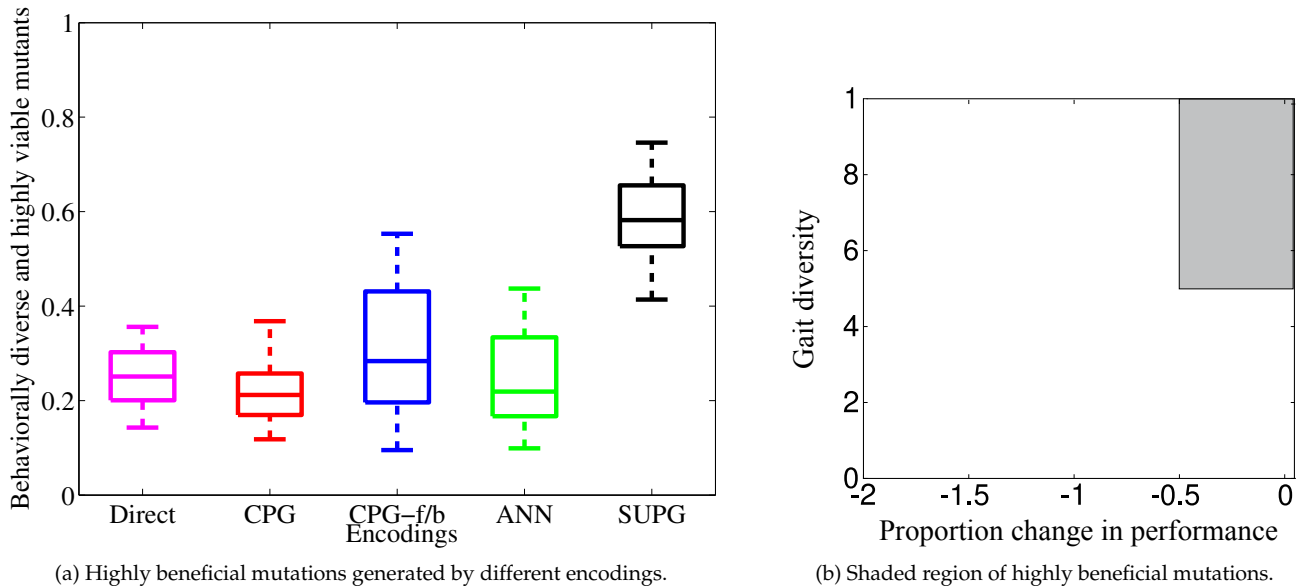


Figure S1: Proportion of mutants generating gait diversity in excess of 0.5, and performance loss not exceeding 50%, from 1,000 independent mutations of the best individuals at generation 8,000 in each of 20 replicates, for the Direct, CPG, CPG-f/b, ANN and SUPG encoding schemes. These highly beneficial mutations are counted from the shaded region of each encoding's evolvability signature.

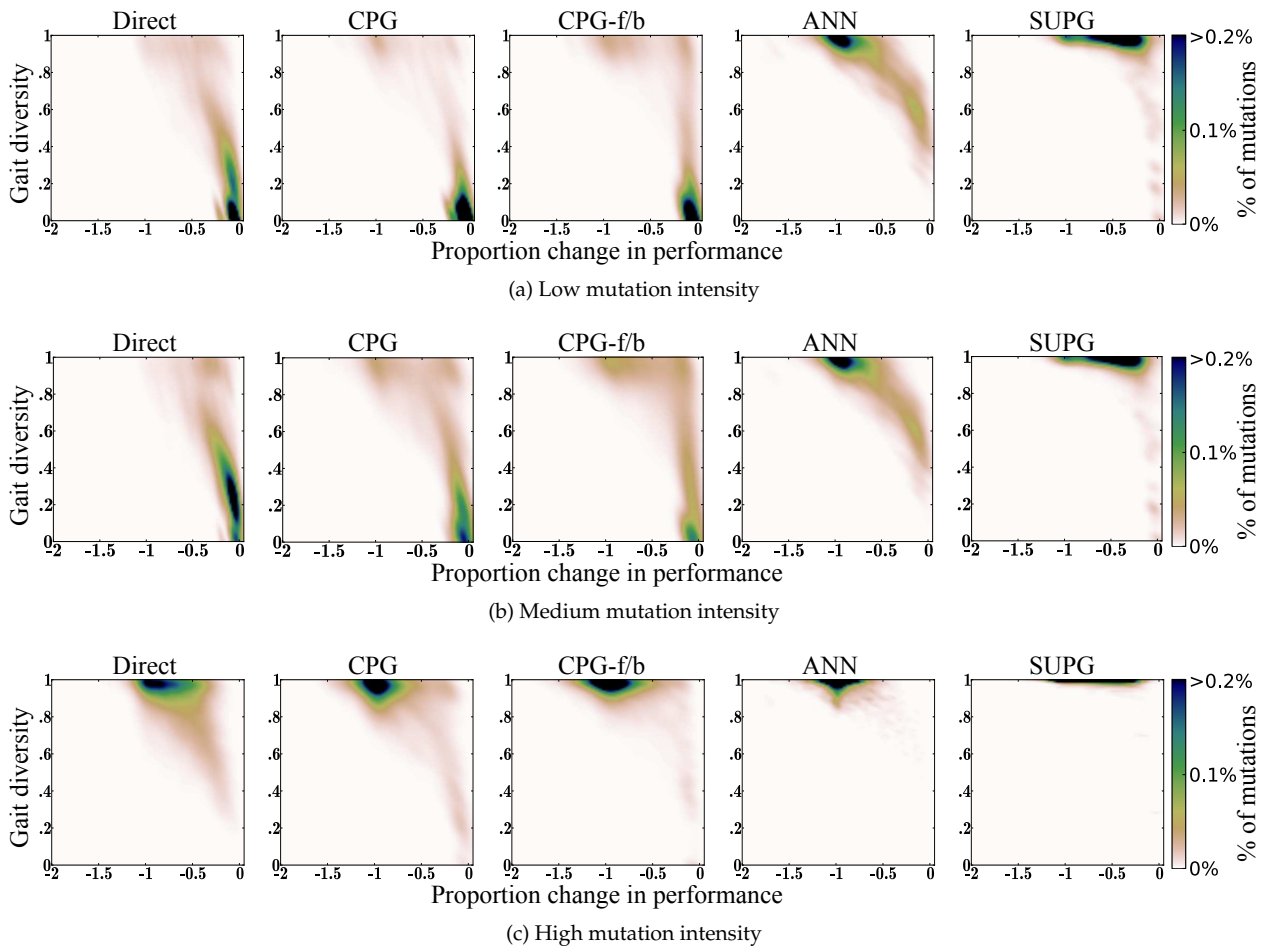


Figure S2: Gait diversity and the proportion decrease in performance, following 20,000 independent mutations of different intensities, for the Direct, CPG, CPG-f/b, ANN and SUPG encoding schemes after 8,000 generations of selection, pooled from all 20 replicates. Mutations at low, medium and high intensity, corresponded to 0.25, 1 and 4 times, respectively, the mutation rate and step-size used during selection.

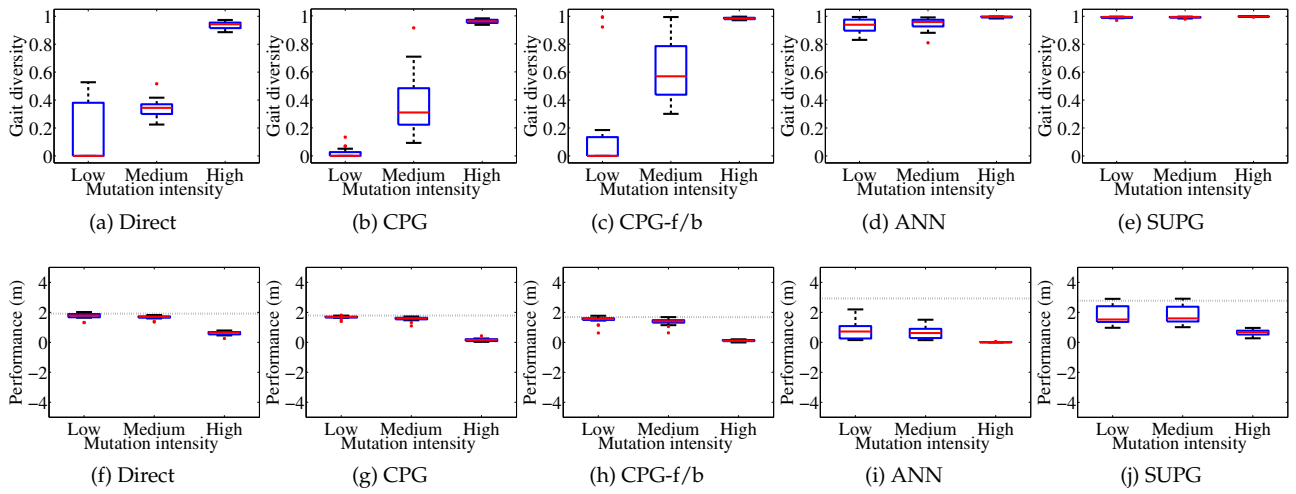


Figure S3: Mutant performance in forward displacement and the gait diversity resulting from mutation, for the Direct, CPG, CPG-f/b, ANN and SUPG encoding schemes. Each box corresponds to the median of 1,000 mutants, in each of 20 replicates. Mutations were generated at low ($\times 0.25$), medium ($\times 1$) and high ($\times 4$) times mutation rate and step-size used during selection, and applied to the best individuals at generation 8,000 of selection. Dotted lines (f-j) indicate the median performance of best individuals across 20 replicates, after 8,000 generation of selection.

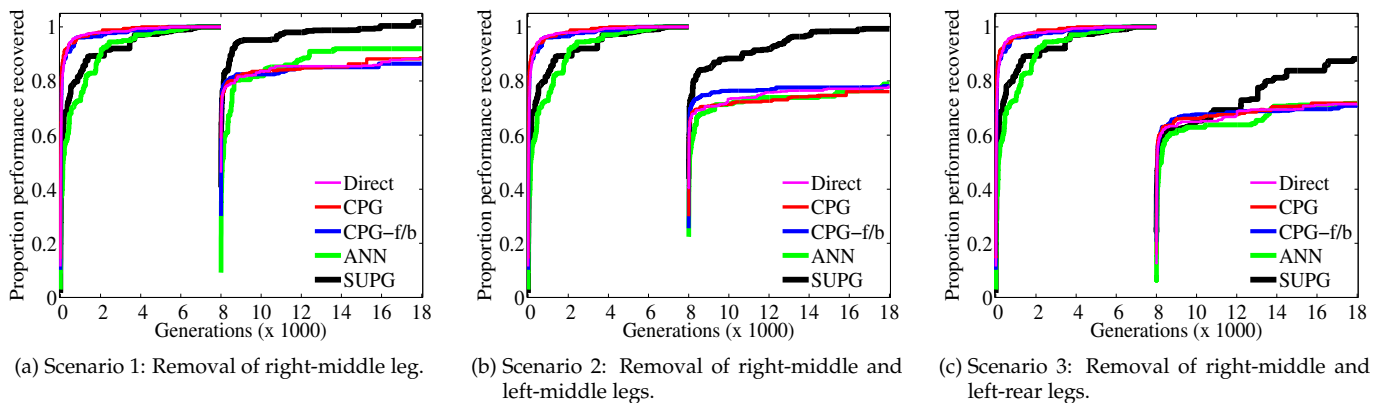


Figure S4: The proportion of performance in forward displacement before (generations 0 to 8,000) and after (generations 8001 to 18,000) the three robot-damage scenarios, for the Direct, CPG, CPG-f/b, ANN and SUPG encoding schemes.

Table S1: Performance comparison of the Direct, Open-loop CPG, Closed-loop CPG, HyperNEAT, and SUPG encoding schemes.

Encodings	<i>p</i> -Value*	Effect size**
Direct	8.0×10^{-5}	0.85
Open-loop CPG		
Direct	7.0×10^{-7}	0.90
Closed-loop CPG		
Direct	0.0142	0.30
HyperNEAT		
Direct	1.2×10^{-7}	0.02
SUPG		
Open-loop CPG	3.0×10^{-4}	0.81
Closed-loop CPG		
Open-loop CPG	4.2×10^{-4}	0.21
HyperNEAT		
Open-loop CPG	1.7×10^{-8}	0.0
SUPG		
Closed-loop CPG	4.8×10^{-6}	0.09
HyperNEAT		
Closed-loop CPG	3.4×10^{-8}	0.0
SUPG		
HyperNEAT	0.1860	0.42
SUPG		

*Significance computed with the Mann-Whitney-Wilcoxon non-parametric test, at 0.05 significance level.

**Effect size estimated as ratio of Mann-Whitney *U* statistic and the product of two sample sizes.

Table S2: Comparison of the number of beneficial mutations generated by the Direct, Open-loop CPG, Closed-loop CPG, HyperNEAT, and SUPG encoding schemes.

Encodings	<i>p</i> -Value*	Effect size**
Direct	0.0304	0.30
Open-loop CPG		
Direct	1.5×10^{-5}	0.10
Closed-loop CPG		
Direct	6.8×10^{-8}	0.0
HyperNEAT		
Direct	6.7×10^{-8}	0.0
SUPG		
Open-loop CPG	5.1×10^{-4}	0.18
Closed-loop CPG		
Open-loop CPG	6.8×10^{-8}	0.0
HyperNEAT		
Open-loop CPG	6.8×10^{-8}	0.0
SUPG		
Closed-loop CPG	2.5×10^{-4}	0.16
HyperNEAT		
Closed-loop CPG	3.4×10^{-7}	0.03
SUPG		
HyperNEAT	9.1×10^{-8}	0.01
SUPG		

*Significance computed with the Mann-Whitney-Wilcoxon non-parametric test, at 0.05 significance level.

**Effect size estimated as ratio of Mann-Whitney *U* statistic and the product of two sample sizes.

Table S3: Comparison of number of generations of selection required to restore 85% of the original performance of the undamaged hexapod, for the Direct, CPG, CPG-f/b, ANN, and SUPG encoding schemes, in the first damage scenario (removal of right-middle leg).

Encodings	<i>p</i> -Value*	Effect size**
Direct	0.4182	0.48
Open-loop CPG		
Direct	0.3538	0.47
Closed-loop CPG		
Direct	0.4393	0.52
HyperNEAT		
Direct	0.0003	0.81
SUPG		
Open-loop CPG	0.4505	0.49
Closed-loop CPG		
Open-loop CPG	0.3045	0.55
HyperNEAT		
Open-loop CPG	0.0010	0.81
SUPG		
Closed-loop CPG	0.0045	0.55
HyperNEAT		
Closed-loop CPG	3.4×10^{-7}	0.79
SUPG		
HyperNEAT	9.1×10^{-8}	0.74
SUPG		

*Significance computed with the Mann-Whitney-Wilcoxon non-parametric test, at 0.05 significance level.

**Effect size estimated as ratio of Mann-Whitney *U* statistic and the product of two sample sizes.

Table S4: Comparison of number of generations of selection required to restore 85% of the original performance of the undamaged hexapod, for the Direct, CPG, CPG-f/b, ANN, and SUPG encoding schemes, in the second damage scenario (removal of right-middle and left-middle legs).

Encodings	<i>p</i> -Value*	Effect size**
Direct	0.4768	0.50
Open-loop CPG		
Direct	0.1979	0.45
Closed-loop CPG		
Direct	0.0286	0.65
HyperNEAT		
Direct	3.9×10^{-6}	0.90
SUPG		
Open-loop CPG	0.2231	0.46
Closed-loop CPG		
Open-loop CPG	0.0286	0.65
HyperNEAT		
Open-loop CPG	1.8×10^{-6}	0.91
SUPG		
Closed-loop CPG	0.0052	0.69
HyperNEAT		
Closed-loop CPG	4.8×10^{-7}	0.93
SUPG		
HyperNEAT	0.0026	0.76
SUPG		

*Significance computed with the Mann-Whitney-Wilcoxon non-parametric test, at 0.05 significance level.

**Effect size estimated as ratio of Mann-Whitney *U* statistic and the product of two sample sizes.

Table S5: Comparison of number of generations of selection required to restore 85% of the original performance of the undamaged hexapod, for the Direct, CPG, CPG-f/b, ANN, and SUPG encoding schemes, in the third damage scenario (removal of right-middle and left-rear legs).

Encodings	<i>p</i> -Value*	Effect size**
Direct	0.1711	0.48
Open-loop CPG		
Direct	0.1711	0.47
Closed-loop CPG		
Direct	0.0082	0.67
HyperNEAT		
Direct	0.0001	0.78
SUPG		
Open-loop CPG	0.5	0.50
Closed-loop CPG		
Open-loop CPG	0.0023	0.67
HyperNEAT		
Open-loop CPG	3.3×10^{-5}	0.80
SUPG		
Closed-loop CPG	0.0023	0.68
HyperNEAT		
Closed-loop CPG	3.3×10^{-5}	0.80
SUPG		
HyperNEAT	0.2460	0.56
SUPG		

*Significance computed with the Mann-Whitney-Wilcoxon non-parametric test, at 0.05 significance level.

**Effect size estimated as ratio of Mann-Whitney *U* statistic and the product of two sample sizes.