Multi-Robot System Evolution with a Cost of Complexity

Alexander Furman
University of Cape Town
Cape Town, South Africa
alex.cleo@gmail.com

ABSTRACT

This paper investigates the impact of imposing a cost on morphological complexity in the co-evolution of control policy (behaviour) and sensory configuration (morphology) for robot teams in different environments. Namely, we investigate (1) evolving lower morphological complexity without sacrificing behavioural competence in a given environment and (2) the relationship between selection pressure for morphological complexity and environmental difficulty. Two experiment sets use direct-encoded Neuro-evolution to evolve team controller-morphology couplings in three environments of increasing difficulty. A single Task Performance objective is maximised in the first set of experiments, while both Task Performance and a Morphological Simplicity objective (that is, a cost of complexity) are maximised in the second set. Overall a cost of complexity was found to be beneficial for Multi-Robot System design. Results indicate that, in a given environment, evolution with a cost on complexity produces teams which are morphologically simpler and as behaviourally competent as teams evolved without a complexity constraint. Additionally, with a cost of complexity, evolution maintained a constant selection pressure for morphological complexity across all environments.

1 INTRODUCTION

A general topic of interest in the field of Evolutionary Robotics (ER) [17] is the co-evolution of controller and morphology. Traditionally, a control policy is evolved for a fixed sensory configuration (morphology) that the designer pre-specifies [58, 59]. While this unburdens the designer from having to linearly model complex behaviours [24], the optimal morphology is often also impractical to ascertain *a priori*. This can result in unnecessary expenditure on both design and physical parts (such as unneeded sensors) [28]. Research on the coupled dynamics of controller, morphology and the environment indicates that controller-morphology co-evolution with a *cost* imposed on morphological complexity can result in evolved robots which are both morphologically cheaper, and as behaviourally competent, as robots evolved without a complexity constraint [4, 5].

For problems which are better solved by co-operative robot teams (Multi-Robot Systems / MRS) over individual robots, such as toxic waste cleanup [49], minefield clearing [27] and search-and-rescue missions [43], these improvements to the design process are particularly desirable. As the control policy and morphology of each robot in a MRS contributes to the emergence of the system's co-operative behaviour, it is even more impractical to make *a priori* design decisions.

While a significant amount of work has been done on co-evolving behaviour and morphology for individual robots [5, 8, 11, 20, 23, 29, 30, 36, 37, 50, 53] much less work has been done on co-evolving behaviour and morphology for co-operative robot teams [2, 10, 28, 47], and - to the best of our knowledge - no work has been done on the co-evolution of behaviour and morphology for robot teams with a constraint on morphological complexity. A barrier to this area of research is the fact that standard ER methods - most notably Neuro-Evolution of Augmenting Topologies (NEAT) [55] - tend to be designed for single-objective optimisation, whereas optimising behaviour and morphological complexity is a multi-objective problem requiring convergence to a set of trade-off solutions.

This study contributes to research on automated controller-morphology design methods for MRS solutions to collective behaviour tasks. Specifically, we investigate research on constraining morphological complexity without sacrificing behavioural competence, and moreover how selection pressure for morphological complexity changes across environments. We equate competent behaviour with optimal performance at *collective gathering* [7], a benchmark task in co-operative robotics [17]. Collective Gathering requires a team of robots to locate and push a set of *resources* (blocks) to a *gathering zone*. We define *co-operation* as the number of robots required to push a block (based on block-size), and *task difficulty* as a function of both *number of blocks* in the environment and *degree of co-operation*.

As in related work, each robot in the team is controlled by an Artificial Neural Network (ANN) and morphologically comprised of a sensory configuration (modeled after the *Khepera* [39] robot) [28, 58, 59], both of which are evolved by *NEAT-M*, an extended NEAT implementation which genotypically encodes the morphology in addition to the controller [28]. We define *task performance* as number of gathered blocks and *morphological complexity* as a function of number of sensors as well as the FOV and Range of each sensor. To impose a *cost* (constraint) on morphological complexity during evolution, we integrate NEAT-M with the *NEAT-MODS* [1] technique for multi-objective optimisation with NEAT. Namely, the integrated method is used to both maximise task performance and minimise morphological complexity, converging to a pareto-optimal [14] set of solutions providing different compromises between morphological complexity and task performance. Due to the high computational complexity of evolving *heterogeneous* teams (different controller-morphology couplings) as reported in related work [28], this study evolves behaviourally and morphologically *homogenous* teams (same controller-morphology coupling across team) for all experiments.

We hypothesise the following:

- ***H0:*** *With a cost on morphological complexity, evolution produces teams which are less morphologically complex, but at least as competent for task performance, as teams evolved with no cost on morphological complexity.*

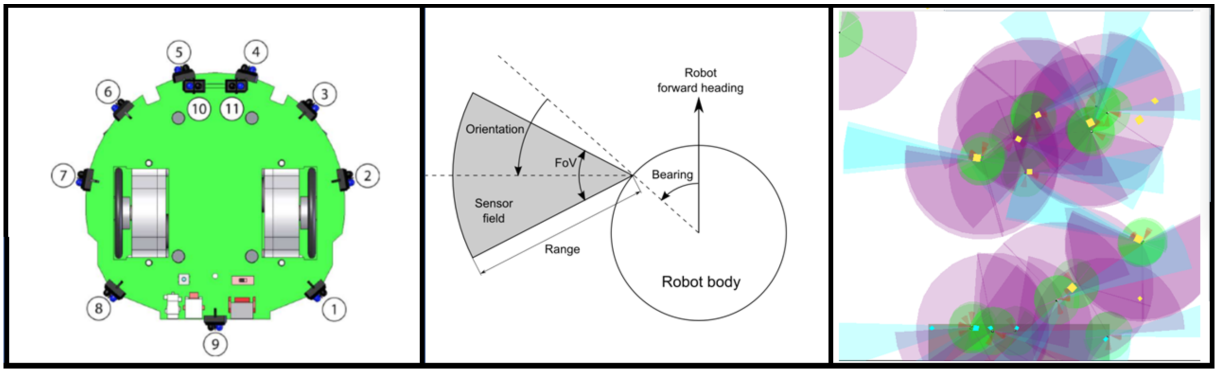

Figure 1: **Left:** Khepera-III morphology which the simulated robot model we employ is based on [39]. **Middle:** The simulated robot model. A conical field-of-view protrudes from each sensor [28]. Sensory input is received by an Artificial Neural Network controller which actuates left and right turning of the wheel. **Right:** Example simulation of collective gathering. Robots attempt to locate and cooperatively push the yellow blocks to the rectangular *gathering zone* at the bottom of the environment. The colour of a conical field-of-view corresponds to a different type of sensor in use.

- ***H1:*** *Evolution will not respond to increased environmental difficulty by increasing selection pressure for morphological complexity.*

Both hypotheses are motivated by the results of related work. Namely, the first hypothesis is supported by the results of **(author?)** [4, 5] which demonstrate that evolution with a cost of complexity can produce individual robots which are morphologically simpler and as behaviourally competent as robots evolved with no complexity constraint. A crucial difference, however, is that we are evolving teams rather than individual robots. Support for the second hypothesis is provided by similar work on co-adaptation of controller and morphology for robot teams. Namely, it has been found that increased environmental difficulty does not necessarily imply a need for greater morphological complexity [58, 59]. These studies have not, however, explored this relationship with a cost imposed on morphological complexity.

2 BACKGROUND

2.1 Neuro-evolution of Augmenting Topologies (NEAT)

Neuro-evolution [62] combines the relative strengths of *Artificial Neural Networks* [24] and *Evolutionary Algorithms* [19] to provide a robust parallel search of the space of candidate network solutions for a given problem. As such, it is an effective controller design technique for non-linear problem spaces where other common approaches fail [24].

NEAT [55] is a neuro-evolution method which evolves both connection weights and topology, based on three principal ideas: *historical marking*, *speciation* and *complexification. Historical Marking* is an ordered numbering system for topological innovations during evolution. Whenever a parent genotype is mutated, a record (*innovation number*) is kept of that mutation. This facilitates topological comparison of networks without the need for computationally-expensive graph traversal. *Speciation* protects new topological innovations by grouping homologous ANNs into species using a *compatibility function.* Solutions compete and reproduce within their species rather than with the population at large, affording young solutions the chance to refine over time. *Complexification* is the idea that ANN solutions should start off with minimal topologies and grow incrementally, *complexifying* over the number of epochs towards a desirable solution.

In addition to applications of NEAT in research areas such as autonomous vehicle control [26, 41, 44, 48, 54, 57] and games development [32, 51, 52], it is particularly well-suited for the non-linearity of problems in Evolutionary Robotics [17] and the related fields of *collective robotics* [31] and *swarm robotics* [6]. Within these fields, NEAT has been used for research on the evolution of co-operation [40, 56], communication [22, 61], predator-prey behaviour [45, 46], morphological complexity [3, 5, 9] as well as embodied cognition [13].

2.2 Evolutionary Multi-objective Optimisation

A multi-objective problem (MOP) consists of multiple (often conflicting) *objectives* that must be optimised simultaneously, yielding a set of trade-off solutions among the objectives [14]. The set is generated according to the notion of *Pareto Optimality*, which states that any solution to a MOP is pareto optimal if none of the objective functions can be better optimised without degrading another of the objective functions in value. Since different solutions might perform better for different subsets of the objectives, the principle of *dominance* is used for comparison of solutions. A solution x^* *dominates* another solution x if the following conditions hold [18]:

1. the solution x^* is not worse than x with respect to all objectives;
2. the solution x^* is strictly better than x with respect to at least one objective.

Thus, no two solutions in the final set of trade-off solutions (the *pareto front*) dominate one another, but every solution on the pareto front dominates every solution outside of the pareto front. While mathematical programming has historically been used to solve MOPs [25, 38], EAs are often a superior alternative. EAs relax the contextual requirement of differentiable objective functions, provide better handling of concavity on the pareto front, and can produce the pareto-optimal set in a single evolutionary run [14].

2.3 Multi-Objective NEAT

A number of problems which initially appear well-suited to neuroevolution, notably within the field of Evolutionary Robotics [17], later prove to require the optimisation of multiple conflicting objectives. Several approaches to a multi-objective implementation of NEAT have been reported in the literature, but most come at the cost of discarding core innovations of the original algorithm to make room for conflicting multi-objective procedures [1].

NEAT's *speciation* mechanism is particularly difficult to preserve in a multi-objective implementation; NSGA-II [16] and other state-of-the-art EMOA techniques sub-divide the population based on metrics such as *non-domination*, while NEAT sub-divides the population into species based on topological homology. Here we describe several approaches to multi-objective NEAT which attempt to preserve the core features of NEAT.

2.3.1 *Weighted Fitness Function.*

A naive approach to optimising multiple objectives with NEAT is to design the fitness function as a *weighted* function of the multiple objectives [14], without having to modify any core features of NEAT. This approach, however, does not take into account trade-offs among the objectives, making it impossible to produce a pareto-optimal set of solutions.

2.3.2 *NEAT-PS.*

The *NEAT-PS* [60] approach optimises separate objective functions but transforms each individual's score into a single scalar value (based on *Pareto Strength* as is done in the *SPEA-II* [63] EMOA). This scalar value is treated as a fitness score and passed directly to NEAT in the usual manner, with no other modifications to the algorithm. However, since NEAT is *non-elitist* by design [55], **(author?)** [1] point out that this approach does not guarantee monotonic evolution of all objectives. While *elitism* allows both child and parent solutions to propagate to the next generation, *non-elitism* only selects child solutions [16]. As such, parent solutions in NEAT-PS which should have remained in the pareto-optimal set stand a chance of being lost during evolution.

2.3.3 *Auxiliary Diversity Objective.*

A number of studies have modified NEAT for multi-objective optimisation by replacing speciation with an auxiliary *genotypic diversity* objective which incentivises topological diversity in the network solution space (in similiar spirit to speciation) [33, 35, 42]. The population can thus be sub-grouped into a set of pareto fronts (rather than species) via non-dominated sort, where each individual in the front with the lowest index dominates all other fronts (as in the NSGA-II [16] algorithm). However, the reliability of this approach as a replacement for speciation has not been verified quantitatively, nor have any efforts been made to generalise the approach into a generic multi-objective NEAT method. Another possible drawback of an auxiliary objective is increased dimensionality of the objective space [34].

2.3.4 *NEAT-MODS.*

(author?) [1] recently developed a novel approach to multi-objective NEAT which preserves speciation. Based on NSGA-II, this method incorporates non-domination rank and crowding distance into a modified version of NEAT's selection process which ensures both genotypic diversity and elitism in the context of pareto-optimality. NEAT-MODS has been quantified against NEAT-PS and is intended to be used as a generic multi-objective NEAT method [1].

3 METHODS

This section outlines the methods used for evolving solutions to the collective gathering task. The first method, NEAT-M [28], provides the core model for co-evolving behaviour and morphology. The second method, NEAT-M-MODS, extends this core model to facilitate multi-objective optimisation.

3.1 NEAT-M

NEAT-M [28] follows the same core evolutionary steps as NEAT, but rather than exclusively evolving a single *controller genotype* for each individual in the population, it also evolves a direct genotypic encoding of the morphology, which is in essence a sensory configuration. This is achieved by creating a relationship between the controller and morphology, such that each (sensory) input node of the controller contains a direct encoding of the parameter-set for a corresponding sensor on the morphology. Namely, the parameter-set for each sensory input node includes a *Sensor Type, Field of View, Range, Bearing,* and *Orientation.* To facilitate evolution of the sensory configuration, there is a corresponding genetic operator for each of these values; namely: *Mutate Field of View, Mutate Range, Mutate Bearing,* and *Mutate Orientation.* Additionally, *speciation* is adapted to account for sensory evolution [28]. If, for instance, evolution discards an input node from the controller, the corresponding sensor is removed from the morphology. A more thorough treatment of NEAT-M is provided by **(author?)** [28].

3.2 NEAT-M-MODS

The *NEAT-M-MODS* method we introduce combines the controller-morphology co-evolution facilitated by NEAT-M [28] with the NEAT-MODS [1] approach to optimising multiple objectives with NEAT. The algorithm initialises (*generation 0*) by generating a minimal parent population, computing each parent's score vector, speciating the population, and then computing a *rank* for each individual based on non-dominated sort and crowding distance comparison. For additional generations, it proceeds as follows:

(1) Apply mutation and crossover to the parent population, producing the child population.
(2) Compute each child's score vector (note parents have already been scored).
(3) Combine the parent and child population
(4) Speciate the combined population
(5) Compute a rank for each parent based on non-dominated sort and crowding distance comparison
(6) Select individuals from the combined population (*size 2N*) to create the population for the next generation (*size N*).
    - *Selection Phase 1 (Select Species):* The combined population is traversed, in order of individual rank, to select a list of species. A limiting function is used during the traversal to ensure that the list of chosen species is both representative of the *genotypic diversity space* and the *elitist pareo-optimal space.*
    - *Selection Phase 2 (Select Individuals):* The list of species is traversed via *serial progression* in order to select a set of *N* individuals to constitute the new parent population. The serial progression ensures that the list of selected individuals

	Number of Blocks of Size:		
	Small	Medium	Large
Environment 1	10	5	0
Environment 2	5	5	5
Environment 3	0	5	10

Table 1: Simulated Task Environments. Environments are numbered from simplest to most difficult. Small, Medium and Large blocks (*resources*) require one, two and three robots, respectively, to be pushed.

Experiment Set ID	**Objectives**	**Environments**		
		1	2	3
MO	M, T	MO1	MO2	MO3
SO	T	SO1	SO2	SO3

Table 2: Experiment Sets. Experiment sets evolve controller and morphology (sensory configuration) for homogenous robot teams in three different environments (Table 1). In each MO experiment, both Morphological Simplicity *M* (equation 2) and Task Performance *T* (equation 1) are maximised. In each SO experiment, only Task Performance is maximised. Each experiment per set is repeated five times.

is both genotypically diverse and elitist in the context of pareto-optimality.

(7) Repeat steps 1 to 6.

4 EXPERIMENTAL SETUP

We employ a multi-robot collective gathering simulator [28] for the experiments. This section provides the simulator configuration as well as parameter choices for the evolutionary methods which were used to conduct the experiments[1]. Two sets of three experiments are designed. In both sets of experiments, controller-morphology solutions for collective gathering are evolved for three environments of increasing difficulty (provided in Table 1). In the first set of experiments, evolution only maximises a task performance objective. In the second set of experiments, both task performance and a morphological simplicity objective are maximised. Each experiment set is repeated five times.

4.1 Simulated Robot Configuration

Simulated robots constitute the co-operative team solutions which we evolve for the experiments. Namely, we evolve *homogenous* team solutions such that each robot uses the same control policy and sensory configuration. As in related work [28, 58, 59], the general morphology for each robot is based on the *Khepera III* [39] robot and illustrated in Figure 1. An Artificial Neural Network (ANN) control policy governs robot behaviour by providing output values for left and right movement actuators. These values are computed by input nodes which receive environmental information from corresponding sensors on the morphology. The number of sensors can be configured manually or, as in our experiments, via automatic methods. Notably, we extend the sensory model designed by **(author?)** [28] to include a simulated *Ultrasonic Sensor* and *Low Resolution Camera* for a more accurate model of the actual *Khepera III* robot [39]. For a full list of robot specifications used throughout the experiments, including the comparative advantages of different sensors, see Table 3. We note that the simulated robots use several heuristics to reduce the computational complexity of interacting with the environment [28].

4.2 Environment Configuration

The simulated environment is a (1.0 x 1.0) two-dimensional continuous plane which contains a set of robots, resources and a gathering zone. Resources are block shapes which vary in size and can be pushed or pulled by robots. Namely, larger blocks are abstracted to be heavier and thus require more robots to be moved (values provided in Table 3). The gathering zone (0.5 x 0.2) is a rectangular portion of the environment, inside of which re-located blocks are considered to be 'collected'. For a given run of the simulator, robots and resources are initially placed at random positions, and with random orientations, outside of the target area.

We configure the three environments which the experiments are repeated for such that a simpler environment implies mostly smaller blocks to be moved and a more difficult environment contains larger blocks. Table 1 provides the three environment configurations.

4.3 Fitness Functions

In Experiment 1, a single *Task Performance* objective is maximised, thereby placing a selection pressure on behavioural competence. In Experiment 2, both Task Performance and a *Minimal Morphological Complexity* objective are maximised. The Minimal Morphological Complexity objective places an additional selection pressure towards lower morphological complexity, thereby imposing a *cost* on morphological complexity. This section provides the fitness functions for the objectives.

4.3.1 Task Performance. Task Performance (or behavioural competence) is the team fitness of a candidate solution taken as an average over five simulated task trials of that solution in a given generation. **(author?)** [28] define v_c as total value of resources in the gathering zone, v_t as total value of all resources in the environment, s_e as the number of simulation timesteps elapsed, and s_t as number of trial evaluations per individual ($n = 5$, Table 3). As such, Task Performance T is maximised according to equation 1:

$$T = 100 \times \frac{v_c}{v_t} + 20 \times (1.0 - \frac{s_e}{s_t}) \tag{1}$$

4.3.2 Minimal Morphological Complexity. Minimal Morphological Complexity is defined as a function of the number of sensors n ($n \in [0, 10]$) on a candidate solution as well as the Field of View (FOV) value f_i and Range value r_i of each sensor S_i in the set of n selected sensors. The values f_i and r_i are constrained by the *sensor type* of S_i. Namely, $\vee F_i$ and $\wedge F_i$, and $\vee R_i$ and $\wedge R_i$, are the maximum and minimum possible values of f_i and r_i, respectively, for S_i's *sensor type* (see Table 3). Thus, Morphological Simplicity M is maximised according to equation 2:

$$M = 100 - 5 \times \sum_{i=1}^{n} \left(\frac{f_i - \wedge F_i}{\vee F_i - \wedge F_i} + \frac{r_i - \wedge R_i}{\vee R_i - \wedge R_i} \right) \tag{2}$$

---

[1] The multi-robot simulator, Neuro-evolution methods and source-code used for the experiments can be found at: https://github.com/rudolfbono/honours-project

Table 3: Neuro-Evolution and Simulator Paramaters

Neuro-Evolution Paramaters	
Generations per experiment	250
Trial evaluations per phenotype	5
Population size	150
ANN connection weight range	$[-1.0, 1.0]$
Sensor Mutation	0.08
Add Sensor Mutation	0.07
Sensor FOV / Range / Bearing / Orientation Perturb	Cauchy mutation $(0, 5)$
Connection weight mutation probability	0.335
Initial Connection Density	0.5
Initial Sensory Input Nodes / Output Nodes	5 / 2
Output Nodes	2
Crossover / Mutation	0.32 / 0.34
Simulation Parameters	
Timesteps per simulated trial evaluation	10000
Robot team size	20
Robot size (diameter) / Gripping distance	0.004 / 0.002 (Portion of environment size)
Maximum robot movement per timestep	0.013 (Portion of environment size moved per iteration)
Initial robot / block positions	Random (Outside gathering zone)
Environment width x height / Gathering zone size	1.0 x 1.0 / 0.5 x 0.2
Small / Medium / Large block size (Width / Height)	0.01 x 0.01 / 0.015 x 0.015 / 0.02 x 0.02
Ultrasonic sensor Range / FOV	$(0.0, 1.0]$ / $(0.0, \pi)$
Infrared Proximity Range / FOV	$(0.0, 0.4]$ / $(\pi/6, 5\pi/6)$
Colour Sensor Range / FOV	$(0.0, 0.4]$ / $(\pi/6, 5\pi/6)$
Low Res Camera Range / FOV	$(0.0, 0.8]$ / $(\pi/9, 8\pi/9)$
Bottom Proximity	Downward-facing
Sensor Bearing Range	$[-\pi, \pi]$ Radians
Sensor Orientation Range	$[-\pi/2, \pi/2]$ Radians

- $\dfrac{f_i - \wedge F_i}{\vee F_i - \wedge F_i}$ is the fraction of total possible FOV used by S_i
- $\dfrac{r_i - \wedge R_i}{\vee R_i - \wedge R_i}$ is the fraction of total possible Range used by S_i.

5 RESULTS AND DISCUSSION

This section provides an analysis and discussion of results from the experiment sets (Table 2) in the context of the hypothesised outcomes (see Introduction). In the experiments, evolution produced homogenous team controller-morphology couplings for a collective gathering task. Specifically, NEAT-M-MODS (multi-objective) evolved *MO solutions* with a cost of morphological complexity, and NEAT-M (single-objective) evolved *SO solutions* without a cost on morphological complexity.

5.1 Validation of Hypothesis 0

H0: *With a cost on morphological complexity, evolution produces teams which are less morphologically complex, but at least as competent for task performance, as teams evolved with no cost on morphological complexity.*

Figure 2 shows the best evolved MO and SO solutions for each of the three environments with respect to task performance and morphological simplicity. Namely, the best MO solution for a given environment was a pareto front of non-dominated trade-off solutions (controller-morphology couplings) and the best SO solution was a single controller-morphology coupling. As each experiment was repeated five times (see Table 2), the final pareto front for each environment was generated via non-dominated sort of the five evolved pareto fronts for each MO experiment [1]. The best evolved SO solution was generated by averaging the best SO task performance and morphological simplicity scores over the five SO evolutionary runs in each environment.

Validation of H0 required a quantitative metric for comparing the best SO and MO solutions for each of the three environments. Namely, it needed to be shown that (1) SO and MO produce equivalent task performance and (2) that MO produces greater morphological simplicity than SO. However, there is no generic quantitative method reported in the literature for selecting the ideal solution from a pareto front for comparison with a single solution [5].

In a related study, **(author?)** [5] formulate a metric for comparing pareto fronts to points by combining several proposed methods from the literature for computing the *ideal* point on the front, and showing that each method produces the same conclusion within a specified margin of error. We adopt a similar approach and compute the following *ideal* points for each run of MO ($n = 5$) in each of the three environments:

(1) $\bar{x}$**:** Average value of front objectives [5]

	Environment 1 *(MO vs SO)*		Environment 2 *(MO vs SO)*		Environment 3 *(MO vs SO)*	
SO compared with	Task Performance	Morphological Simplicity	Task Performance	Morphological Simplicity	Task Performance	Morphological Simplicity
$\bar{x}$	SO: [5, 10]	MO: [30, 50]	SO: [5, 10]	MO: [30, 50]	SO: [5, 10]	MO: [30, 50]
x_{max}	MO: [1, 5]	MO: [5, 10]	MO: [1, 5]	MO: [0, 20]	=	MO: [20, 30]
x_{ϵ}	SO: [1, 5]	MO: [30, 50]	=	MO: [30, 50]	=	MO: [30, 50]
X_{knee}	SO: [0.5, 1]	MO: [30, 50]	=	MO: [30, 50]	=	MO: [30, 50]

Table 4: Range of differences between best evolved SO solution (average) and set of *ideal* MO pareto front points for both task performance and morphological simplicity in three different environments. Best SO Task Performance and Morphological Simplicity are taken as average over five runs of a given SO experiment. Best set of ideal MO points are computed from a final pareto front resulting from non-dominated sort of the five fronts of each MO experiment. "=" indicates there was no difference between average SO point and given ideal pareto front point. All tests gave *p<0.05* (Shapiro-Wilk test).

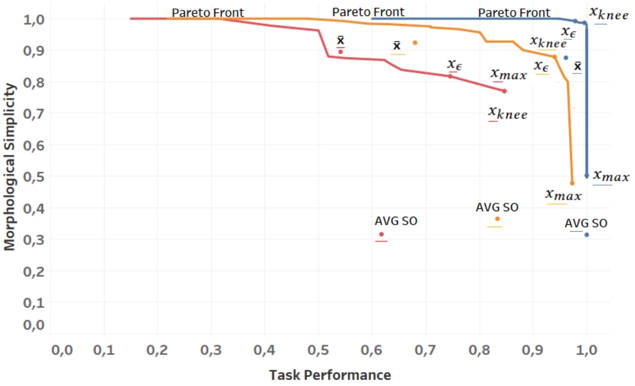

Figure 2: Best-evolved MO and SO solutions for each of the three environments with respect to task performance and morphological simplicity. Environments 1, 2 and 3 are blue, orange and red respectively (1 is simplest, 3 is most difficult). The best MO solution for a given environment was a pareto front of non-dominated trade-off solutions (controller-morphology couplings) and the best SO solution was a single controller-morphology coupling. Lowest attainable morphological simplicity ($M = 0$, equation 2) corresponds to a robot body with 10 sensors and the most powerful parameters for each sensor. Highest attainable morphological simplicity ($M = 1$, equation 2) corresponds to a robot body with no sensors.

(2) x_{max}: Solution on front with highest task performance [5]
(3) x_{ϵ}: Point which maximises morphological simplicity and is up to 10% lower than x_{max} for task performance [21]
(4) x_{knee}: The knee point of the front; solution with the highest value of both objectives [15]

For each environment we conduct an independent samples t-test to test for a significant difference in task performance between the five best SO solutions and each point in the set $\{\bar{x}, x_{\epsilon}, x_{max}, x_{knee}\}$ for each of the five final pareto fronts. Normality of the data was confirmed using the Shapiro-Wilk test. Results of the t-tests (shown in Table 4) indicate that for each environment, SO and MO mean task performance do not differ statistically by more than 10% ($p < 0.05$). As the average standard deviation of best task performance across the five runs of MO and SO was σ*=0.1*, we accept a difference of 10% to meet the criteria for approximate equivalence of MO and SO task performance, thereby satisfying the first condition of H0. Indeed, exact equivalence is highly unlikely in practice given the stochastic nature of neuro-evolution [5].

A second set of t-tests (also shown in Table 4) returned a significant difference in morphological simplicity between SO and MO in each environment, showing that MO produces simpler morphology in all cases ($p < 0.05$). Namely, the mean difference in morphological simplicity between MO and SO is above 30% for the majority of cases. Thus, MO produces greater morphological simplicity than SO in all tested cases, which satisfies the final condition of H0.

As such, we conclude that the results provide corroborative evidence for H0; namely that team controller-morphology couplings evolved with a cost of complexity are morphologically simpler and as behaviourally competent (that is, produce similar task performance) as solutions evolved without a cost on complexity.

These results are consistent with related work [58, 59] which demonstrates that higher sensory complexity does not necessarily result in higher task performance for robot teams. It also consistent with **(author?)** [4] who show that decreasing *mechanical* complexity results in higher task performance (note this was shown for a single robot rather than a team).

Notably, the difference in mean SO and MO task performance tends towards 0 as the environment increases in difficulty. Specifically, Table 4 indicates that SO produces marginally higher task performance in simpler environments, but equal task performance (that is, both methods gathered same amount of resources) in more difficult environments. The underlying reason for this will be investigated in future research.

5.2 Validation of Hypothesis 1

H1: *Evolution will not respond to increased environmental difficulty by increasing selection pressure for morphological complexity.*

From H0, it follows that with a cost of morphological complexity, evolution produces degrees of morphological simplicity which are closer to the absolute minimum morphological requirements for competent behaviour in the given environment than what would be produced without a cost of complexity. Namely, this indicates that there is a relationship between morphological complexity and the environment.[2] However, further analysis is necessary to understand

---

[2] While we do not provide proof that team morphology is directly influenced by the environment rather than by the underlying mechanisms of NEAT (such as Complexification [55]), our claim is supported by similar work on single-robot systems, in which a neutral shadow model was employed to prove that morphology is influenced by the environment even with random selection enabled [5]. In future work we will employ a neutral shadow model for this measure.

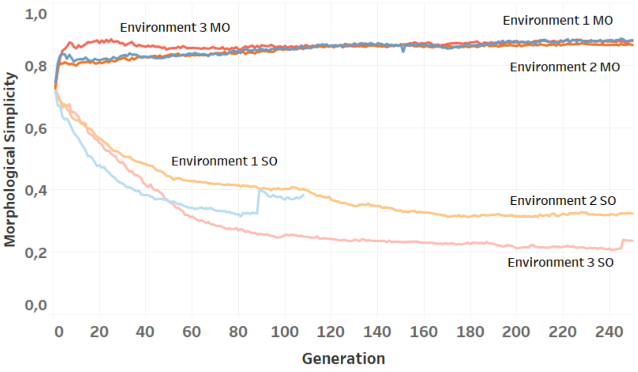

Figure 3: Progression of Morphological Simplicity produced by MO and SO in the three environments over evolutionary time (generations). Lowest attainable morphological simplicity ($M = 0$, equation 2) corresponds to a robot body with 10 sensors and the most powerful parameters for each sensor. Highest attainable morphological simplicity ($M = 1$, equation 2) corresponds to a robot body with no sensors. Note that in *Environment 1 SO* (Experiment SO1, Table 2) evolution converged prior to the 250th generation as maximum task performance was attained.

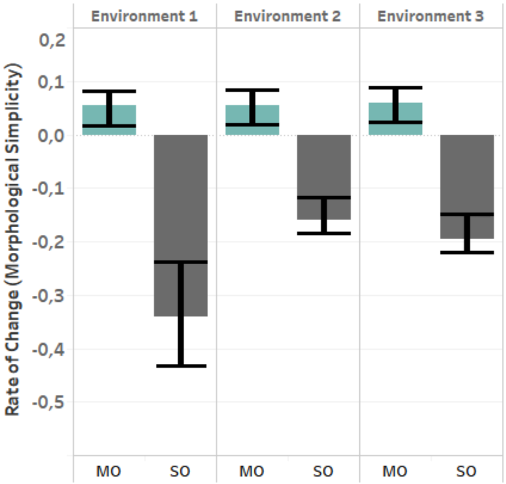

Figure 4: Average rates at which Mean Morphological Simplicity changes over evolutionary time for MO and SO in increasingly difficult environments. Environments are numbered in order of difficulty (see Table 1). Rates of Change correspond to each of the variables in Figure 3. A drop in Rate of Change across environments corresponds to lowering of selection pressure for morphological simplicity in response to environmental difficulty. Comparatively high standard error in SO1 (Table 2) is expected to be a result of premature convergence (maximum task performance attained).

this relationship *across different environments* (namely from simple to more difficult environments). Indeed, related research indicates that (with no cost of complexity) increased environmental difficulty does not necessarily equate to a need for higher morphological complexity [28, 58, 59].

Figure 1 shows the evolutionary progression of morphological simplicity in the three environments produced by MO (evolution with a cost of complexity) and SO (no cost on complexity) over 250 generations. Note that experiment SO1 (Table 2) evolved maximum task performance, and thus converged, prior to the 250th generation (also that convergence was disabled for MO). We firstly observe that, for all environments, MO produces simpler morphology over time while SO produces more complex morphology over time. This finding is in line with H0; namely that evolution with a cost of complexity places higher selection pressure on morphological simplicity than evolution without a cost of complexity.

Moreover, Figure 1 appears to show that MO's selection pressure for morphological simplicity is approximately the same *across all three environments*; namely that an increase in environmental difficulty does not induce a change in selection pressure for morphological simplicity. On the other hand, SO's selection pressure for morphological simplicity differs across the environments over time, but there appears to be no relationship between morphological simplicity and environmental difficulty over time.

We investigate these observations further using Figure 4, which illustrates the average rate of change of morphological simplicity over evolutionary time for MO and SO in each of the three environments. Figure 4 shows that MO's rate of change of morphological simplicity over time is the same in all three environments, which builds on our observation from Figure 3 that MO selection pressure for morphological simplicity is the same across environments. We additionally observe that for SO, the rate of change of morphological simplicity increases from Environment 1 (simple) to Environment 2 (medium), but then decreases from Environment 2 (medium) to Environment 3 (hard). This suggests further that there is no relationship between SO's selection pressure for team morphological simplicity and environmental difficulty [28, 58, 59].

Overall, results indicate that for SO there appears to be no relationship between selection pressure for team morphological simplicity and environmental difficulty[3] (as also found in related work [28, 58, 59]) and that MO (cost of complexity) maintains a constant selection pressure for team morphological simplicity across environments. These findings provide suitable corroborative evidence for H1; namely that evolution will not respond to increased environmental difficulty by increasing selection pressure for morphological complexity.

Our finding that evolution with a cost of complexity selects for simpler team sensory configurations in difficult environments might appear to be counter-intuitive. Indeed, **(author?)** [5] find that individual robots exhibit higher morphological complexity when evolved in more difficult environments. However, while we have loosely made use of the term *morphological complexity*, complexity comes in a number of forms, many of which are inversely correlated [5]. Indeed, **(author?)** [5] find that evolution selects for simpler *mechanical* (range of mechanical joint motion) complexity in the same difficult environments that induce high morphological (triangle mesh body parts) complexity. We propose that lower *sensory* complexity (the subject of our research) is more desirable for co-operative robot teams in difficult environments on account of the high *controller* (neural) complexity required for the emergence of co-operative behaviour. Namely, we suggest that competent co-operative behaviour in difficult environments requires greater

---

[3] We note, however, that premature convergence of SO1 may have interfered with this result.

evolutionary emphasis on optimising the control policy rather than optimising for the least expensive morphology. Validation of this proposition would require controller complexity to be studied in addition to morphological complexity, and will be the subject of future work.

6 CONCLUSIONS

This study investigated the impact of imposing a cost of morphological (sensory) complexity during evolution of sensory-controller couplings for homogenous robot teams in different environments. Morphological complexity was a function of the number of sensors on a given robot, as well as the parameters evolved for each sensor.

Experiments used direct-encoding Neuro-evolution to evolve team solutions for a benchmark collective behaviour task, requiring robots to locate and gather resources co-operatively in a range of environments. Difficult environments contained large resources which required more robots (greater *co-operation)* to be pushed than the smaller resources in simpler environments. A single *Task Performance* objective was maximised in the first set of experiments, while both Task Performance and a *Morphological Simplicity* objective (that is, a cost of complexity) were maximised simultaneously in the second set.

Overall, the results indicated that imposing a cost on morphological complexity is beneficial for Multi-Robot System design. It was firstly found that, in a given environment, adding a cost of morphological complexity leads to simpler team morphology without sacrificing behavioural competence. This suggests that, for real-world co-operative robot tasks, competent solutions can be evolved while also automatically reducing design costs which would have been spent on unnecessary sensors. It was additionally found that, with a cost of complexity, evolution maintained a constant selection pressure for morphological complexity across all environments. This suggests that, contrary to intuition, robot teams do not require greater sensory complexity for competent behaviour in increasingly difficult environments. For real world multi-robot system design, this could facilitate more economical spending on sensors for task environments of varying difficulty in which additional sensors are wrongly expected to be necessary. While additional research is necessary to understand the underlying reason for this finding, we suggest that evolution places higher selection pressure on controller than on morphology in order to facilitate the emergence of co-operative behaviour.

In addition to benefits for Multi-Robot System design, this work is a preliminary step towards clearer understandings of the coupled dynamics between control, morphology and the environment in evolving co-operative systems.

7 FUTURE WORK

While there are many avenues for future research given the preliminary scale of this work and the various types of task environments, complexity and evolutionary models that can be studied [5], we propose the following prioritised subjects for future research:

Firstly, we plan to use a *neutral shadow model* as validation for the direct influence of the environment on team morphological complexity which we report in the current work. A neutral shadow model is a repeated run of a given experiment, except that individuals are selected randomly for each new generation. By comparing the result of a given experiment to a shadow model of that experiment, it is possible to ascertain whether changes to morphological complexity were truly a product of environmental influence rather than due to the underlying mechanisms of the evolutionary method itself (such as Complexification in NEAT [55]).

Additionally, we plan to investigate complexity of the controller in addition to complexity of the morphology in homogenous robot teams. This might help to validate our hypothesis that, with a cost of complexity, evolution prioritises selection pressure for controller complexity rather than for morphological complexity to facilitate the emergence of co-operative behaviour.

Future work should also attempt to reproduce our experiments using different multi-objective NEAT approaches to NEAT-MODS [1], such as to replace speciation with a Genotypic Diversity Function and place NEAT's selection mechanism under the control of NSGA-II [16] (as is done in related work [4, 5, 12, 34]).

Lastly, should the researchers have access to the necessary computational resources, future work should attempt to evolve behaviourally and morphologically *heterogenous* (different controllers and morphologies) rather than homogenous robot teams (same controller and morphology across team).

8 ACKNOWLEDGEMENTS

The author would like to thank Dr Geoff Nitschke for his supervision and support throughout the project. Additional thanks go to Danielle Nagar for collaborating with the author on this research. Computations were performed using facilities provided by the Centre for High Performance Computing: https://chpc.ac.za/ as well as the ICTS High Performance Computing centre: http://hpc.uct.ac.za/.

REFERENCES

[1] O. Abramovich and A. Moshaiov. Multi-objective topology and weight evolution of neuro-controllers. In *Evolutionary Computation (CEC), 2016 IEEE Congress on*, pages 670–677. IEEE, 2016.

[2] Y. Asai and T. Arita. Coevolution of morphology and behavior of robots in a multiagent environment. In *Proceedings of the SICE 30th Intelligent System Symposium*, pages 61–66, 2003.

[3] J. Auerbach and J. Bongard. Evolving complete robots with cppn-neat: the utility of recurrent connections. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pages 1475–1482. ACM, 2011.

[4] J. Auerbach and J. Bongard. On the relationship between environmental and mechanical complexity in evolved robots. In *Artificial Life 13*, number EPFL-CONF-191280, pages 309–316. MIT Press, 2012.

[5] J. Auerbach and J. Bongard. Environmental influence on the evolution of morphological complexity in machines. *PLoS computational biology*, 10(1):e1003399, 2014.

[6] G. Beni. From swarm intelligence to swarm robotics. In *International Workshop on Swarm Robotics*, pages 1–9. Springer, 2004.

[7] E. Bonabeau, M. Dorigo, G. Théraulaz, et al. *Swarm intelligence: from natural to artificial systems*. Number 1. Oxford university press, 1999.

[8] J. Bongard. The utility of evolving simulated robot morphology increases with task complexity for object manipulation. *Artificial life*, 16(3):201–223, 2010.

[9] J. Bongard. Morphological change in machines accelerates the evolution of robust behavior. *Proceedings of the National Academy of Sciences*, 108(4):1234–1239, 2011.

[10] G. Buason, N. Bergfeldt, and T. Ziemke. Brains, bodies, and beyond: Competitive co-evolution of robot controllers, morphologies and environments. *Genetic Programming and Evolvable Machines*, 6(1):25–51, 2005.

[11] M. Bugajska and A. Schultz. Co-evolution of form and function in the design of autonomous agents: Micro air vehicle project. Technical report, NAVAL RESEARCH LAB WASHINGTON DC CENTER FOR APPLIED RESEARCH IN ARTIFICIAL INTELLIGENCE, 2000.

[12] N. Cheney, J. Bongard, and H. Lipson. Evolving soft robots in tight spaces. In *Proceedings of the 2015 annual conference on Genetic and Evolutionary Computation*, pages 935–942. ACM, 2015.

[13] N. Cheney, J. Clune, and H. Lipson. Evolved electrophysiological soft robots. In *ALIFE*, volume 14, pages 222–229, 2014.

[14] C. Coello. Evolutionary multi-objective optimization: a historical view of the field. *IEEE computational intelligence magazine*, 1(1):28–36, 2006.

[15] K. Deb and S. Gupta. Understanding knee points in bicriteria problems and their implications as preferred solution principles. *Engineering optimization*, 43(11):1175–1204, 2011.

[16] 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.

[17] S. Doncieux, N. Bredeche, J. Mouret, and E. Eiben. Evolutionary robotics: what, why, and where to. *Frontiers in Robotics and AI*, 2:4, 2015.

[18] S. Doncieux and J. Mouret. Beyond black-box optimization: a review of selective pressures for evolutionary robotics. *Evolutionary Intelligence*, 7(2):71–93, 2014.

[19] A. Eiben and J. Smith. *Introduction to Evolutionary Computing*. Springer, Berlin, Germany, 2003.

[20] K. Endo, T. Maeno, and H. Kitano. Co-evolution of morphology and walking pattern of biped humanoid robot using evolutionary computation. consideration of characteristic of the servomotors. In *Intelligent Robots and Systems, 2002. IEEE/RSJ International Conference on*, volume 3, pages 2678–2683. IEEE, 2002.

[21] J. Ferreira, C. Fonseca, and A. Gaspar-Cunha. Methodology to select solutions from the pareto-optimal set: a comparative study. In *Proceedings of the 9th annual conference on Genetic and evolutionary computation*, pages 789–796. ACM, 2007.

[22] D. Floreano, S. Mitri, S. Magnenat, and L. Keller. Evolutionary conditions for the emergence of communication in robots. *Current biology*, 17(6):514–519, 2007.

[23] P. Funes and J. Pollack. Evolutionary body building: Adaptive physical designs for robots. *Artificial Life*, 4(4):337–357, 1998.

[24] F. J. Gomez. *Robust non-linear control through neuroevolution*. PhD thesis, 2003.

[25] S. Greco, J. Figueira, and M. Ehrgott. Multiple criteria decision analysis. *Springer's International series*, 2005.

[26] S. Hallé, J. Laumonier, and B. Chaib-Draa. A decentralized approach to collaborative driving coordination. In *Intelligent Transportation Systems, 2004. Proceedings. The 7th International IEEE Conference on*, pages 453–458. IEEE, 2004.

[27] N. Hazon and G. Kaminka. Redundancy, efficiency and robustness in multi-robot coverage. In *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*, pages 735–741. IEEE, 2005.

[28] J. Hewland and G. Nitschke. The benefits of adaptive behavior and morphology for cooperation. In *Computational Intelligence, 2015 IEEE Symposium Series on*, pages 1047–1054. IEEE, 2015.

[29] G. Hornby and J. Pollack. Body-brain co-evolution using l-systems as a generative encoding. In *Proceedings of the 3rd Annual Conference on Genetic and Evolutionary Computation*, pages 868–875. Morgan Kaufmann Publishers Inc., 2001.

[30] G. Hornby and J. B. Pollack. Creating high-level components with a generative representation for body-brain evolution. *Artificial life*, 8(3):223–246, 2002.

[31] S. Kernbach. Introduction to collective robotics: Reliability, flexibility, and scalability, 2013.

[32] Z. H. Laster et al. Evolving agents using neat to achieve human-like play in fps games. 2014.

[33] J. Lehman, S. Risi, D. D'ambrosio, and K. Stanley. Rewarding reactivity to evolve robust controllers without multiple trials or noise. *Artificial Life*, 13:379–386, 2012.

[34] J. Lehman, S. Risi, D. D'Ambrosio, and K. Stanley. Encouraging reactivity to create robust machines. *Adaptive Behavior*, 21(6):484–500, 2013.

[35] J. Lehman, K. Stanley, and R. Miikkulainen. Effective diversity maintenance in deceptive domains. In *Proceedings of the 15th annual conference on Genetic and evolutionary computation*, pages 215–222. ACM, 2013.

[36] H. Lipson and J. Pollack. Automatic design and manufacture of robotic lifeforms. *Nature*, 406(6799):974, 2000.

[37] H. Lund. Co-evolving control and morphology with lego robots. In *Morpho-functional machines: the new species*, pages 59–79. Springer, 2003.

[38] K. Miettinen. Nonlinear multiobjective optimization, volume 12 of international series in operations research and management science, 1999.

[39] F. Mondada, E. Franzi, and P. Ienne. Mobile robot miniaturisation: A tool for investigation in control algorithms. In *Experimental robotics III*, pages 501–513. Springer, 1994.

[40] J. Montanier and N. Bredeche. Surviving the tragedy of commons: Emergence of altruism in a population of evolving autonomous agents. In *European Conference on Artificial Life*, 2011.

[41] D. Moriarty and P. Langley. Distributed learning of lane-selection strategies for traffic management. *Daimler-Benz Res. Technol. Center, Palo Alto, CA, Tech. Rep*, pages 98–2, 1998.

[42] J. Mouret and S. Doncieux. Encouraging behavioral diversity in evolutionary robotics: An empirical study. *Evolutionary computation*, 20(1):91–133, 2012.

[43] R. Murphy, S. Tadokoro, D. Nardi, A. Jacoff, P. Fiorini, H. Choset, and A. Erkmen. Search and rescue robotics. In *Springer Handbook of Robotics*, pages 1151–1173. Springer, 2008.

[44] A. Narayan, E. Tuci, and F. Labrosse. Simulated road following using neuroevolution. In *Artificial Life and Intelligent Agents Symposium*, pages 17–30. Springer, 2014.

[45] G. Nitschke. Emergence of cooperation in a multiple predator, single prey game. In *FLAIRS Conference*, pages 234–238, 2003.

[46] G. Nitschke and L. Langenhoven. Neuro-evolution for competitive co-evolution of biologically canonical predator and prey behaviors. In *Nature and Biologically Inspired Computing (NaBIC), 2010 Second World Congress on*, pages 546–553. IEEE, 2010.

[47] R. O'Grady, A. Christensen, and M. Dorigo. Swarmorph: Morphogenesis with self-assembling robots. In *Morphogenetic Engineering*, pages 27–60. Springer, 2012.

[48] A. Parker and G. Nitschke. How to best automate intersection management. In *Evolutionary Computation (CEC), 2017 IEEE Congress on*, pages 1247–1254. IEEE, 2017.

[49] L. Parker. Alliance: An architecture for fault tolerant multirobot cooperation. *IEEE transactions on robotics and automation*, 14(2):220–240, 1998.

[50] J. Pollack, H. Lipson, G. Hornby, and P. Funes. Three generations of automatically designed robots. *Artificial Life*, 7(3):215–223, 2001.

[51] S. Risi and J. Togelius. Neuroevolution in games: State of the art and open challenges. *IEEE Transactions on Computational Intelligence and AI in Games*, 9(1):25–41, 2017.

[52] J. Schrum and R. Miikkulainen. Evolving multimodal behavior with modular neural networks in ms. pac-man. In *Proceedings of the 2014 annual conference on genetic and evolutionary computation*, pages 325–332. ACM, 2014.

[53] K. Sims. Evolving virtual creatures. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 15–22. ACM, 1994.

[54] K. Stanley, N. Kohl, R. Sherony, and R. Miikkulainen. Neuroevolution of an automobile crash warning system. In *Proceedings of the 7th annual conference on Genetic and evolutionary computation*, pages 1977–1984. ACM, 2005.

[55] K. Stanley and R. Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2):99–127, 2002.

[56] M. Waibel, D. Floreano, and L. Keller. A quantitative test of hamilton's rule for the evolution of altruism. *PLoS biology*, 9(5):e1000615, 2011.

[57] Y. Wang and B. Schreiber. Creating a traffic merging behavior using neuroevolution of augmenting topologies. 2015.

[58] J. Watson and G. Nitschke. Deriving minimal sensory configurations for evolved cooperative robot teams. In *Evolutionary Computation (CEC), 2015 IEEE Congress on*, pages 3065–3071. IEEE, 2015.

[59] J. Watson and G. Nitschke. Evolving robust robot team morphologies for collective construction. In *Computational Intelligence, 2015 IEEE Symposium Series on*, pages 1039–1046. IEEE, 2015.

[60] W. Willigen, E. Haasdijk, and L. Kester. A multi-objective approach to evolving platooning strategies in intelligent transportation systems. In *Proceedings of the*

15th annual conference on Genetic and evolutionary computation, pages 1397–1404. ACM, 2013.

[61] S. Wischmann, D. Floreano, and L. Keller. Historical contingency affects signaling strategies and competitive abilities in evolving populations of simulated robots. *Proceedings of the National Academy of Sciences*, 109(3):864–868, 2012.

[62] X. Yao. Evolving artificial neural networks. *Proceedings of the IEEE*, 87(9):1423–1447, 1999.

[63] E. Zitzler, M. Laumanns, and L. Thiele. Spea2: Improving the strength pareto evolutionary algorithm. *TIK-report*, 103, 2001.