

Evolution of a Functionally Diverse Swarm via a Novel Decentralised Quality-Diversity Algorithm

Emma Hart
Edinburgh Napier University
Edinburgh, Scotland, UK
e.hart@napier.ac.uk

Andreas S.W. Steyven
Edinburgh Napier University
Edinburgh, Scotland, UK
a.steyven@napier.ac.uk

Ben Paechter
Edinburgh Napier University
Edinburgh, Scotland, UK
b.paechter@napier.ac.uk

ABSTRACT

The presence of functional diversity within a group has been demonstrated to lead to greater robustness, higher performance and increased problem-solving ability in a broad range of studies that includes insect groups, human groups and swarm robotics. Evolving group diversity however has proved challenging within Evolutionary Robotics, requiring reproductive isolation and careful attention to population size and selection mechanisms. To tackle this issue, we introduce a novel, decentralised, variant of the MAP-Elites illumination algorithm which is hybridised with a well-known distributed evolutionary algorithm (mEDEA). The algorithm simultaneously evolves multiple diverse behaviours for multiple robots, with respect to a simple token-gathering task. Each robot in the swarm maintains a local archive defined by two pre-specified functional traits which is shared with robots it come into contact with. We investigate four different strategies for sharing, exploiting and combining local archives and compare results to mEDEA. Experimental results show that in contrast to previous claims, it is possible to evolve a functionally diverse swarm without geographical isolation, and that the new method outperforms mEDEA in terms of the diversity, coverage and precision of the evolved swarm.

CCS CONCEPTS

• **Computing methodologies** → **Evolutionary robotics**;

KEYWORDS

Functional Diversity, Evolutionary Robotics, Map-Elites

ACM Reference Format:

Emma Hart, Andreas S.W. Steyven, and Ben Paechter. 2018. Evolution of a Functionally Diverse Swarm via a Novel Decentralised Quality-Diversity Algorithm. In *GECCO '18: Genetic and Evolutionary Computation Conference, July 15–19, 2018, Kyoto, Japan*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3205455.3205481>

1 INTRODUCTION

In both natural and artificial systems, the benefits of *functional diversity* within a species or group is well understood. In ecology for example, functional diversity within bee pollinators (such as flower height preference, daily time of flower visitation and within-flower behaviour) has been shown to lead to increased pollination rates [10] which in turn sustain the bee population. Within groups

of humans, functional diversity reflects differences in how people represent problems and how they go about solving them. A number of studies indicate the benefit of functional diversity within a group with studies using agent-based modelling suggesting that collective diversity trumps individual ability under certain conditions in groups of problem solvers [11]. In Computer Science, a body of theory supports the requirement for behavioural diversity within ensemble classifiers [3] while more recent work suggests that diversity may also benefit ensembles of optimisation algorithms [9]. A current review of diversity preservation in EAs is given in [20].

We hypothesise that the presence of functional diversity within a robotic swarm would result in similar benefit, specifically providing *robustness* with respect to operation across a range of environmental conditions, potentially resulting in increased longevity of the swarm. A natural question then arises: *how can we evolve functional diversity within a large swarm?* Existing research with Evolutionary Robotics tends towards that view that this is a challenging proposition [14], noting that very specific conditions concerning reproductive isolation must be met or at least careful choice of operators is required. However, an alternative paradigm known as *Quality Diversity* (QD) for evolving diversity has recently been introduced: these algorithms are designed to locate a maximally diverse collection of individuals, according to some low-dimensional set of behavioural characteristics, in which each individual is as high-performing as possible. Examples of QD algorithms include Novelty Search [12] and MAP-Elites [15], with seminal applications in robotics domains, for instance, discovering a set of morphologies for successful walking virtual creatures [13], or diverse behaviours for maze-navigation [18].

Typically, QD algorithms maintain a discrete archive (map) formed by binning a low-dimensional representation of feature vectors, with a single *elite* individual stored within each bin (i.e. the best found so far with those features). The archive can be completely external to the breeding population, having no influence over evolution [19]. Alternatively, as in MAP-Elites [15], the archive is used to select genomes for variation and is an intrinsic part of the evolutionary process.

Central to all the existing QD algorithms just described is the maintenance of a *single* archive, that is accessed and updated by a population at each generation. However, in a real-world robotic swarm, environmental conditions mean that communication is often limited to local interactions. Hence in their current form, the practical applicability of QD algorithms as a mechanism for evolving *swarm* diversity appears limited, given that individual members of the swarm may be unable to access or update a centralised archive.

GECCO '18, July 15–19, 2018, Kyoto, Japan

© 2018 Association for Computing Machinery.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *GECCO '18: Genetic and Evolutionary Computation Conference, July 15–19, 2018, Kyoto, Japan*, <https://doi.org/10.1145/3205455.3205481>.

We therefore propose a novel decentralised QD algorithm for maximising behavioural diversity with an evolving swarm where each robot swarm is tasked with the same single objective. The eventual goal of this research is to study whether this will lead to increased robustness in dynamic environments. The first key step, described here, is to propose and demonstrate a method that evolves the required diversity across the swarm. Our approach hybridises an open-ended distributed evolutionary algorithm mEDEA [2] with MAP-Elites [15] and is referred to as *EDQD - Embodied Distributed Quality Diversity*. The algorithm simultaneously evolves multiple diverse behaviours across a swarm, through local communication only. Its goal is to evolve a swarm that maximises functional diversity, i.e. maximises the number of distinct approaches to achieving the same task. We speculate that evolving a behaviourally diverse swarm will “future proof” the swarm, providing robustness to a broad spectrum of potential environmental conditions – even if some of the evolved behaviours are sub-optimal with respect to the *current* environment. The specific goals of the paper are as follows:

- Evaluate the extent to which a decentralised QD algorithm can be used to obtain a swarm exhibiting functional diversity w.r.t a simple foraging task
- Evaluate the quality of each evolved behaviour in the swarm w.r.t the estimated optimal performance for that behaviour and the reliability with which high-performing behaviours are discovered
- Compare the diversity and relative quality of behaviours evolved using the QD method to a standard distributed evolutionary algorithm (mEDEA).

The paper makes a number of contributions. Firstly, it describes the first completely *decentralised* version of a QD algorithm for *simultaneous evolution of multiple diverse behaviours within a swarm*. This differentiates it from existing work with robotics in which a single archive is used to evolve multiple behaviours for a *single* robot. Secondly, it proposes and evaluates four different methods of sharing and updating map-archives amongst the swarm. These different methods alter the amount of information shared in an attempt to maximise diversity. Experimental results show that the method significantly outperforms a distributed EA in terms of simultaneously evolving diverse behaviours. Thirdly, it extends the usage of the QD paradigm, which is typically directed towards either *illumination* of a feature-space [15] or to providing a range of pre-evolved options for selecting a single behaviour to match current environmental conditions [5]. It represents the first step in a longer-term vision to show that a functionally-diverse swarm will be robust to environmental change, and that it can continue to learn over time – key qualities of practical robotic swarms.

2 BACKGROUND

As the terms *behavioural/functional diversity* are used in a broad range of contexts with evolutionary robotics, it is helpful to position existing research with the field with the help of a simple taxonomy, described in figure 1.

A clear distinction arises between work that evolves multiple distinct behaviours for a *single* robot and that which attempts to simultaneously evolve multiple behaviours within a *swarm*. The former has been tackled extensively in recent years through the

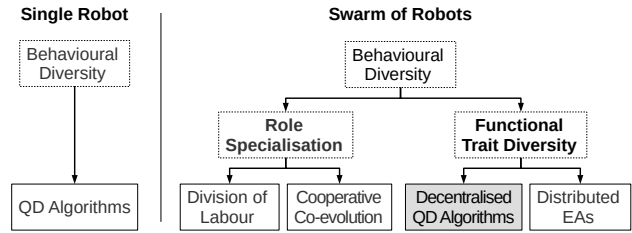


Figure 1: A taxonomy of research focusing on behavioural diversity in Evolutionary Robotics. The novel contribution of this paper is shaded in grey

use of *quality-diversity* algorithms: a new class of algorithms that return an archive of diverse, high-quality behaviours in a single run. An overview of recent work is provided in by Pugh *et al* [18]. On the one hand, QD algorithms encourage exploration of the search space in order to find better solutions, overcoming deception [19]. On the other, they can be used to develop an archive of behaviours that can be for example used in future to guide a trial-and-error learning algorithm to select new appropriate behaviour in the face of damage [5].

This article concerns the latter category, evolution of a *swarm*. Here, a further distinction can be made between evolving *role specialisation* and *functional trait diversity* with the swarm. The first category is commonly referred to as *division of labour* and covers scenarios in which swarms divide into sub-groups, each accomplishing a sub-task, in order to solve a complex problem. This kind of role-specialisation, observed in social insects [4] and more rarely in mammals [6] has been a focus of much effort with robotics and more generally in problem solving [17].

However, our work is positioned within the sub-category classified as *functional trait diversity*: the term refers to the presence of multiple behavioural traits within a group that result in multiple strategies for achieving the same goal. Montanier *et al* [14] find that behavioural diversity in this case is very hard to achieve, noting that reproductive isolation is necessary, whether such isolation is due to geographic constraints or particular mating strategies, and that large population sizes help. Haasdijk *et al* [7] considered evolution of co-existing foraging behaviours within a population and found they needed to introduce a market mechanism favouring sub-groups in order to evolve diversity. Trueba *et al*. [22] conducted an in-depth empirical study of behavioural specialization within the *same* geographic location and showed that judicious choice of evolutionary operators could be used to enforce behavioural specialisation but the study was greatly simplified in that robots simply selected from three predefined behaviours.

We address the challenges just described in obtaining functional trait diversity by hybridising MAP-Elites, a QD algorithm shown to work well in applications involving a *single* robot, with a distributed evolutionary algorithm for evolving behaviours within a *swarm*. The new algorithm – *Embodied Distributed Quality Diversity* algorithm, EDQD is now described.

Table 1: Variants of the EDQD Algorithm

Acronym	Creation of SelectMap	Update of MemoryMap	Update of LocalMap
R	random selection from ReceivedMapList	-	Executed Genome
M1	merge(ReceivedMapList)	-	Executed Genome
M2	merge(ReceivedMapList, MemoryMap)	ReceivedMapList	Executed Genome
M3	merge(ReceivedMapList, LocalMap),	-	Executed Genome, ReceivedMapList

3 THE EMBODIED DISTRIBUTED QUALITY DIVERSITY ALGORITHM (EDQD)

EDQD hybridises a distributed environment-driven evolutionary adaptation algorithm called mEDEA [2] with a novel decentralised version of a MAP-Elites algorithm [15]. The mEDEA algorithm has been extensively studied in previous works and provides a good baseline for embodied evolution. In the original version of this algorithm, there is no explicit fitness function. Therefore, no selection pressure w.r.t. fitness value is applied, but selection pressure w.r.t. ability to spread one’s own genome is still at work. Later versions introduced an additional fitness mechanism [16] to regulate the trade-off between the exploitation of a fitness function and the exploration of solutions allowing the survival of robots. We adopt this version here due to the fact we also introduce an explicit task.

In brief, in mEDEA, for a fixed period (lifetime), robots move according to their control algorithm (a neural network, specified by weights in the genome). As they move, they broadcast their genome, which is received and stored by any robot within range. At the end of this period, a robot selects a random genome from its list of collected genomes and applies a variation operator. This takes the form of a Gaussian random mutation operator that can be easily tuned through a σ parameter. Robots that have not collected any genomes become inactive, thus temporarily reducing the population size.

In EDQD, each robot stores an n -dimensional discrete map defined by functional traits related to the specified task termed the *LocalMap* (LM). In contrast to mEDEA in which a robot broadcasts its current *genome*, in EDQD, each robot broadcasts its *LocalMap*, an elite archive of the genomes it has previously evaluated. All robots in range receive a copy of the broadcast map and store it in a temporary list. At the start of each new generation, every robot selects a new genome from the maps received during the previous generation. A variation operator is applied as in mEDEA and the new controller evaluated over the defined robot-lifetime. The *LocalMap* is then updated based on the fitness of the evaluated phenome (where update is contingent on the discovery of a new elite). These fitness evaluations are inherently noisy, as in all online embodied algorithms, which arises from the fact that agents are evaluated in a shared environment and influence each other. The algorithm is defined in Algorithm 1.

We explore four variants of the algorithm which are differentiated by the method in which the robots select a new genome from maps they have received, and the method in which the *LocalMap* is updated. Four variants of the EDQD algorithm are proposed. These are summarised in table 1 and in figure 2, and described below. The following definitions apply throughout:

- **Local Map** Each robot maintains a local map (*LM*) consisting of a 2-dimensional map containing the best solutions found so far at each point in a space defined by the two dimensions of variation. Each individual *LM* contains elites discovered by the robot by executing genomes in its own body, with the exception of a single variation (EDQD-M3, see below) in which a *LM* is updated with elites obtained from other robots. This map is constructed as described by Mouret *et al* [15], i.e. following execution, a genome with fitness x is mapped to a feature vector b : the genome replaces the current occupant of the cell corresponding to b if it has higher fitness¹ (see Algorithm 2). Note however, that unlike MAP-Elites, there is no initialisation phase in which a map is first created by sampling \mathcal{G} random genomes.
- **Received Map List** A set of maps collected during one lifetime through encounters
- **Select Map** At the end of each lifetime, each robot has a *ReceivedMapList*. We consider four methods by which the robot can condense this information into a single map: the *SelectMap* (*SM*). A genome is selected at random from the *SelectMap* to be executed in the next generation.
- **Memory Map** In variant EDQD-M2, a robot keeps an additional map in memory which combines elites from all LMs it has ever received across all generations of the algorithm.

EDQD-R. At the end of each lifetime, each robot attempts to update its own local map with the fitness of its current phenome, according to Algorithm 2. Following this, in similar vein to mEDEA, the robot simply selects a random map from the *ReceivedMapList* which becomes the *SelectMap*. A random genome is then chosen from the *SelectMap*. The robot then empties its *ReceivedMapList*.

EDQD-M1. As above, each *LocalMap* is first updated with the result of current phenome. The robot forms the *SelectMap* by merging the maps contained in the *ReceivedMapList* (Algorithm 3). As above, it then selects a random genome from the *SelectMap*, before deleting the merged map and emptying its *ReceivedMapList*. This provides more selection-pressure than variant **EDQD-R**, as the merged map will potentially have more cells covered than any single map and contain the genome with the highest phenotypic fitness for overlapping cells.

EDQD-M2. Similar to above, however in this case, at iteration 0 the robot creates an empty *MemoryMap*. At the end of each generation, a merged map is created (Algorithm 3), and this is then merged with the *Memory Map*. This forms the *SelectMap* from which the robot selects a random genome. The *MemoryMap* is updated to be

¹If the new phenome has *equal* fitness to the current occupant, then the phenome with behavioural vector b closest to the cell centre is retained; this method is used in *code* provided by [15]

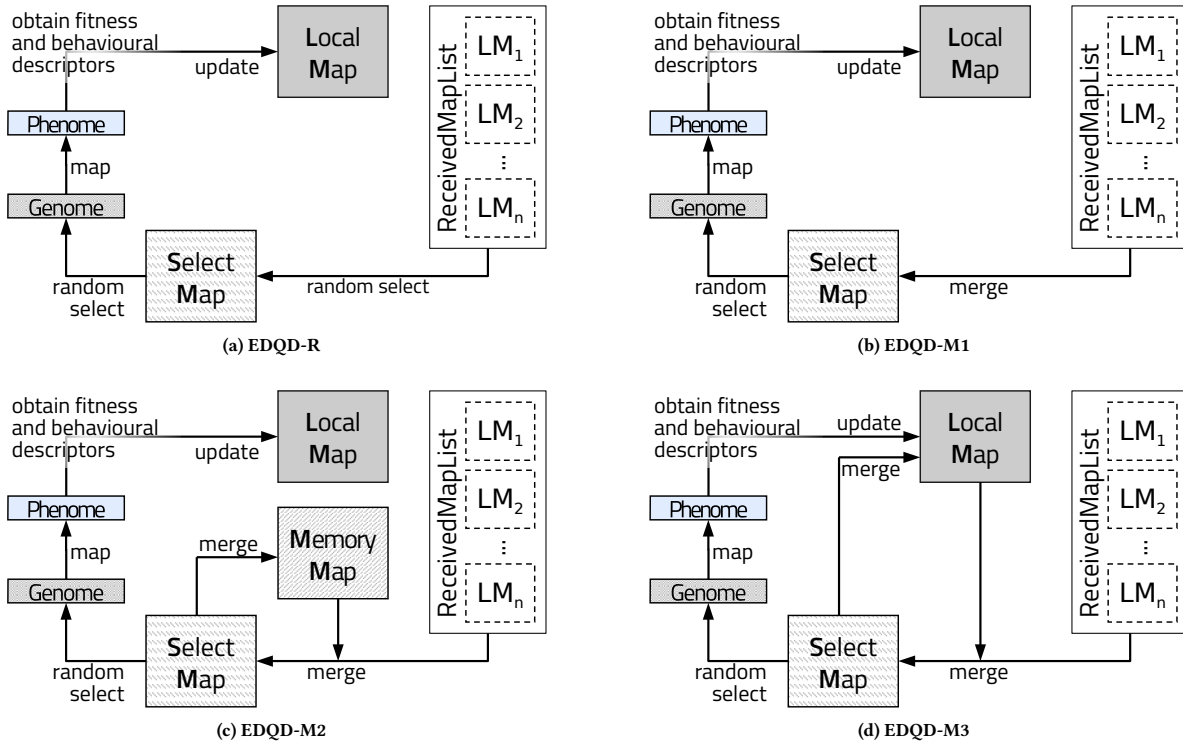


Figure 2: Variants of the EDQD Algorithm

equivalent to the newly merged map. Essentially, this results in the robot maintaining a list of all elites collected from every robot it meets during its lifetime that it can use to *select* a new genome from. This further increases selection pressure, as the *MemoryMap* stores all elites known to all robots ever encountered.

EDQD-M3. In this case, the *ReceivedMapList* is merged with the robot's own *LocalMap* and a genome selected at random from this map. The robot's *LocalMap* is also updated with this information, unlike in the methods above. This alters the information *broadcast* by each robot (the LM), resulting in a more global sharing of information across the population, as each robot now receives a map containing information from genomes executed by many robots. This is expected to provide most selection pressure as it combines information from elites found by an individual robots with all elites found by all encountered robots.

4 EXPERIMENTS

We define a simple foraging task in which a population of robots is placed in a circular arena containing equal numbers of blue and red tokens. Tokens are coloured cylinders of the same size as robots and are collected by a robot through contact. The fitness function is defined as the *total* number of tokens collected during a fixed interval called the *lifetime* (irrespective of token colour). An archive is defined by two functional traits: (1) the maximum Euclidean distance from starting point within a lifetime and (2) the ratio of red:blue tokens collected by the robot. The first encourages

```

genome.randomInitialise();
localMap.create();
while (generations < maxGen) do
    for iteration = 0 to lifetime do
        if agent.hasGenome() then
            agent.move();
            broadcast(localMap);
        end
        mapList ← receivedMaps
    end
    localMap ← updateLocalMap(genome, fitness);
    genome.empty();
    if mapList.size() > 0 then
        createSelectMap();
        genome = applyVariation(selectrandom(selectMap));
        if collectedMapMemory.isForget() then
            mapList.empty();
        end
    end
end
end

```

Algorithm 1: EDQD Algorithm: hybridised from mEEDA [2] and MAP-Elites [15]

diversity w.r.t the extent of the arena explored by the robot, while the second encourages diversity in the *type* of token foraged. Traits

```

 $x \leftarrow \text{random\_selection}(X);$ 
 $x' \leftarrow \text{random\_variation}(x);$ 
 $b' \leftarrow \text{featureDescriptor}(x');$ 
 $p' \leftarrow \text{performance}(x');$ 
if  $\mathcal{P}(b') = \emptyset$  or  $\mathcal{P}(b') < p'$  then
   $\mathcal{P}(b') \leftarrow p';$ 
   $X(b') \leftarrow x';$ 
end
return  $\text{map}(\mathcal{P}, X)$ 

```

Algorithm 2: Update of cells within a *LocalMap* (as defined in [15])

```

 $M1 \leftarrow (X_1, \mathcal{P}_1);$ 
 $M2 \leftarrow (X_2, \mathcal{P}_2);$ 
forall cells  $c$  in  $\text{map do}$ 
  if  $\mathcal{P}_{M2}(c) > \mathcal{P}_{M1}(c)$  then
     $\mathcal{P}_{M1}(c) \leftarrow \mathcal{P}_{M2}(c);$ 
     $X_{M1}(c) \leftarrow X_{M2}(c);$ 
  end
end
return  $M1$ 

```

Algorithm 3: Merge two maps: if there is a list of maps, then process is repeated until all maps are merged

can only be calculated retrospectively at the end of each lifetime. The two-dimensional map is discretised into 15 bins per dimension, therefore contains 225 discrete locations (approximately equal to the number of robots in the population).

All experiments are conducted in simulation using the Roborobo simulator (version 3) [1]. The arena has a diameter of 956 pixels. It contains 150 red tokens and 150 blue tokens uniformly distributed throughout the environment. When a robot consumes a token, a token of the same colour is regenerated at a random location. Robots are cylindrical, with 12 sensors (7 toward the front uniformly covering 90° with the remaining 5 evenly spaced around the remainder of the body) and two motors. A genome defines 126 weights of a feed-forward neural network with 63 inputs corresponding to 3 inputs for the RGB ground colour, the 5 values for each of the 12 sensors (the distance to the nearest object and whether that object is a robot, a wall or either of the 2 type of token) and 2 outputs corresponding to translational and rotational speeds.

A fixed population size of 200 robots is used in all experiments. In all algorithms tested, a Gaussian mutation is used as the variation operator, in which σ is initialised to 0.1 and subsequently evolves. Each robot evaluates its genome over a lifetime defined as 800 iterations. *EDQD* algorithms are compared to *mEDEA_{fps}* which uses fitness proportionate selection to select a new genome from the list of genomes collected during a lifetime; otherwise the parameters of this algorithm are identical to *EDQD*. Each treatment is allocated 1000 generations; for each treatment, 30 independent runs are performed.

Statistical analysis was conducted based on the method in [21] using a significance level of 5%. The distributions of two results were checked using a Shapiro-Wilk test. A Kruskal-Wallis rank sum test was performed to determine the p-value if one of the results

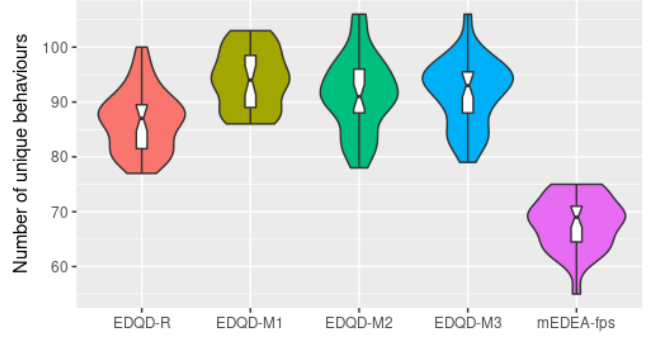


Figure 3: Swarm Diversity: number of unique behaviours exhibited by the swarm at the end of each treatment (over 30 runs)

followed a non-Gaussian distribution. Otherwise Levene’s test for homogeneity of variances was performed. For unequal variances the p-value was determined using a Welch test, otherwise using an ANOVA test.

5 RESULTS

The following section describes results obtained with respect to evaluating the four variants of *EDQD* in terms of diversity of behaviours evolved, and the relative quality of those behaviours.

5.1 Diversity of expressed behaviours within a swarm

At the end of each run (1000 generations), the behaviour of the last-executed phenotype x of each active² robot in the swarm is mapped to the corresponding feature descriptor b . A new map is constructed using algorithm 2, and the number of occupied cells n_{occ} counted. This is repeated for each of the 30 runs: figure 3 shows a violin plot of the distribution of n_{occ} for each treatment. Statistical tests (table 2) show that the *EDQD* variants all outperform *mEDEA_{fps}*; *EDQD-R* (that selects a random genome from a random map) is outperformed by the other three *EDQD* variants; however there are no significant differences between variants M1,M2,M3.

	EDQD-M1	EDQD-M2	EDQD-M3	mEDEA-fps
EDQD-R	< 9.21e-06	< 2.67e-03	< 1.16e-03	> 1.28e-11
EDQD-M1		= 1.05e-01	= 1.4e-01	> 1.28e-11
EDQD-M2			= 7.03e-01	> 1.28e-11
EDQD-M3				> 1.29e-11

Table 2: p-value and directionality obtained from pairwise comparison of number of unique behaviours exhibited by the active swarm at the final generation. Bold indicates significance at 5% level

An alternative perspective is given in figure 4 which indicates which cells in the map-archive are filled by the genomes expressed

²Any robot that has not collected any local-maps in the previous generation is considered as inactive

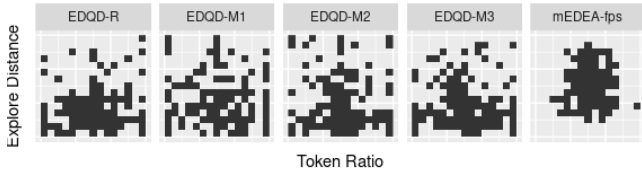


Figure 4: Map-archive of behaviours expressed by robots in a single swarm at the final generation (from a single run)

in the swarm in the final generation. The maps shown are generated from a single run of each treatment. Table 3 shows the number of cells filled from the single run shown as well as the median cells filled per map across 30 runs. Note firstly that the EDQD variants exhibit more unique behaviours and secondly, that these behaviours are more widely *spread* across the behaviour space. In contrast, mEDEA_{fps} tends to find a cluster of similar behaviours in the centre of the map.

Exp. name	Count	Median Count (over 30 runs)
1 EDQD-R	86	87
2 EDQD-M1	94	94
3 EDQD-M2	91	91
4 EDQD-M3	94	93
5 mEDEA-fps	71	69

Table 3: Number of unique behaviours expressed by robots in a swarm at the final generation

Figure 5 shows the variation in diversity within the population (i.e. n_{occ} as defined above) over the course of a single run of each treatment. The rapid loss of diversity using mEDEA is clear. In contrast, the EDQD treatments increase diversity over time. EDQD-R (select from random map) produces least diversity: the random nature of the map-selection method here limits spread of information through the population. On the other hand, the three other methods that encourage faster map propagation increase diversity rapidly.

5.2 Quality of Evolved Behaviours

As stated in section 1, the motivation of this paper is to determine the extent to which it is possible to evolve a diverse set of behaviours within a swarm for a single task. The previous sections clearly demonstrate that EDQD evolves swarms that exhibit functional diversity. However, it is important to understand the impact of encouraging *diversity* on the *quality* of behaviours discovered.

Figure 6 shows the same map-archive of behaviours expressed by the swarm at the end of the final generation of a single run as shown above, but now with cells coloured according to fitness. Note that mEDEA results in a cluster of high-fitness cells in the centre of the map. The behaviours discovered by the EDQD variants, particularly around the edges of the map, are of lower fitness, although EDQD-M2 and EDQD-M3 also locate several high-performing behaviours. For the specific environment in which experiments are conducted, it is inevitable that high-fitness behaviours are likely to occur when a controller results in a robot consuming all tokens, i.e. the resulting token ration trait is 1:1. A robot whose current controller results in it

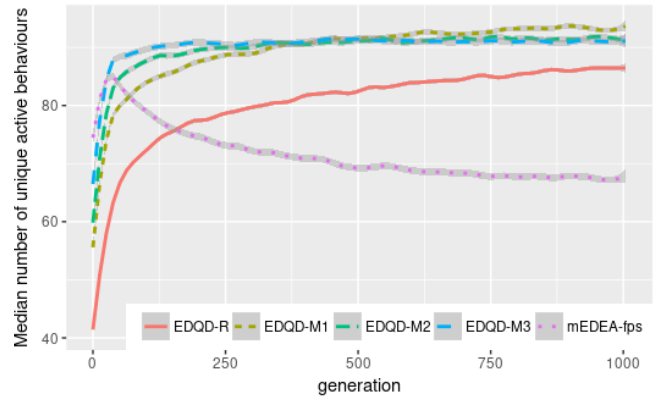


Figure 5: Evolution of the median number of unique behaviours exhibited by the swarm over the course of a single run (lines have been smoothed for better visualisation)




Figure 6: Map-archive of behaviours expressed by robots in a single swarm at the final generation (from a single run), coloured according to fitness

only reacting to red tokens and moving in (for example) tight circles will not collect as many tokens as a robot whose controller results in a behaviour that reacts to both types of token and moves with the same circular pattern, or indeed one that happens to explore more of its environment. However, although such behaviours are *sub-optimal* in the current environment, they may be *optimal* in a different environment, e.g. one in which blue tokens suddenly disappear. Hence there is value in capturing them within the swarm to proof against future change.

For a given environment, it is useful to understand *opt-in reliability* or *precision* of each treatment, using the terminology introduced in [15]. For each run, if (and only if) a run creates a solution in a cell, then the precision is calculated as the average across all such cells of the highest performing solution produced for that cell divided by the optimal solution for that cell. Essentially, this reflects the trust we can have that, if an algorithm returns a solution in a cell, that solution will be high-performing relative to *what is possible* for that cell. As the metric is only measured over filled-cells, it is expected that traditional objective-based evolutionary algorithms should fare well on this criterion, as they explore only a few cells but should produce high-performing solutions in those cells.

The precision of each local map is calculated as just described. As the optimal value is usually unknown, we follow [15] in estimating this value as the best fitness found for a vector b in any run of any treatment. The median precision per swarm of local maps is then calculated for each of the 30 runs. A violin plot of this information

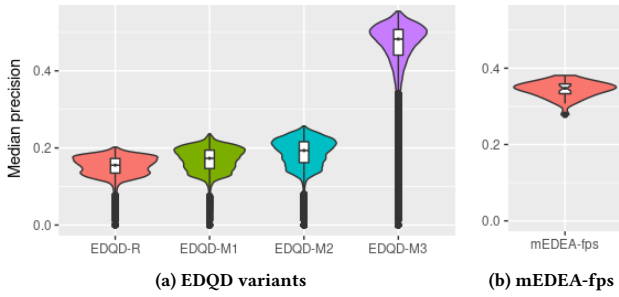


Figure 7: Median precision of a) each swarm at end generation (calculated as median precision of the 200 local-maps in each of 30 runs of EDQD treatments) b) expressed genomes in last generation from mEDEA (over 30 runs)

is given in Figure 7. Although mEDEA does not produce a local map, we project the 200 fitness of the behaviours expressed by the swarm at the end generation onto a map-archive and calculate the same precision metric; this is displayed in fig. 7 for comparison. The results of calculating statistical significance for the EDQD treatments is also given in table 4.

EDQD-M3 significantly outperforms the other EDQD variants. The median swarm precision is also higher than the precision of mEDEA across a swarm. This indicates confidence in the opt-in reliability of EDQD in that it is capable of finding high-quality behaviours, even if some of those behaviours are sub-optimal for the current conditions.

	EDQD-M1	EDQD-M2	EDQD-M3
EDQD-R	< 2.48e-10	< 1.34e-11	< 1.34e-11
EDQD-M1		< 1.35e-09	< 1.34e-11
EDQD-M2			< 1.34e-11

Table 4: p-values obtained from pairwise comparison of median of (median) precision of 200 local-maps (over 30 runs). Bold used to indicate significance at 5% level

5.3 Archive of functional diversity: the swarm-map

The analysis conducted so far has focused on information contained within a distributed swarm as it evolves online. Given that experiments are conducted in simulation, we can extract additional information by collating the *LocalMaps* held by the 200 individual robots at any point t , and merging these into a single *swarm-map*. It should be clear that such a map must be created externally and is *never* accessible to the swarm during the course of a run³. However, in a practical scenario, it is conceivable that robots might regroup following an extended period of online evolution and that the swarm-map could be created and distributed across the swarm before evolution restarts. The external archive would then provide a comprehensive library of diverse behaviours for robots to draw

³Strictly speaking, it could in fact be created by EDQD-M3 if every robot met every other robot during the course of a single lifetime

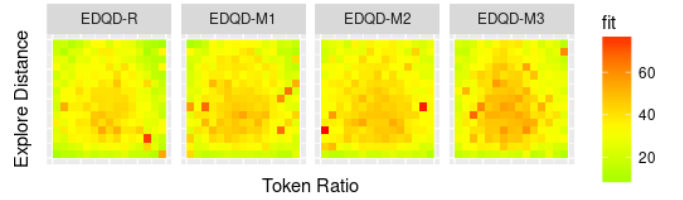


Figure 8: Swarm-map: archive of best individuals from collation of all *LocalMaps* across the swarm of each of the EDQD algorithms at the final generation (from a single run), coloured according to fitness

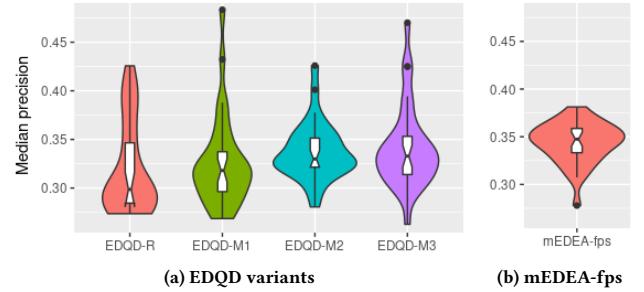


Figure 9: Precision of a) swarm-maps at end generation b) active genomes in last generation (over 30 runs)

on in future, in the face of varying environmental conditions (e.g. as described by [5] for a single robot). Figure 8 shows the *swarm-map* created from a single run of each of the EDQD algorithms at generation 1000. We note that *all* variants fill the entire archive, i.e. the swarm-map is maximally diverse. The *quality* of the behaviours found however appear better using variants M2 and M3.

As described in section 5.2, the precision of a map indicates the reliability with which cells are filled with high-quality solutions. The precision of each swarm-map created from each of 30 runs is calculated and plotted in figure 9(a). As in the previous section, this is also compared to the map obtained by plotting genomes expressed in robots evolved using mEDEA at the final generation (figure 9(b)). Table 5 displays the corresponding statistical analysis obtained by comparing EDQD treatments. Combining the 200 local-maps into a single swarm-map blurs the distinction between the EDQD variants. EDQD-R is outperformed by variants M2 and M3, and EDQD-M2 outperforms M1. Otherwise, there is no statistical difference observed.

	EDQD-M1	EDQD-M2	EDQD-M3
EDQD-R	= 2.63e-01	< 8.29e-03	< 9.78e-03
EDQD-M1		< 2.63e-02	= 5.12e-02
EDQD-M2			= 9.31e-01

Table 5: p-values obtained from pairwise comparison of median swarm-map precision (30 runs). Results that are significant at 5% level shown in bold

6 DISCUSSION AND CONCLUSIONS

A body of evidence from both natural and computational intelligence underlines the benefits that can be derived from groups which exhibit functional diversity. We have hypothesised that a robotic swarm might enjoy the same benefits, in terms of being *robust* to future changes in the environment in which it operates. If a subset of the swarm can continue to function in the face of change (e.g. a sudden disappearance of a particular token type) then the presence of diversity should ensure that at least a subset of the swarm can continue to function. Furthermore, appropriate behaviours can be passed to other robots, assuming the use of an open-ended distributed evolutionary algorithm.

As a step towards achieving this goal, we have proposed a novel decentralised quality diversity algorithm, hybridising a distributed evolutionary algorithm mEDEA with MAP-Elites – EDQD. To the best of our knowledge, this is the first decentralised version of a QD algorithm that simultaneously evolves *multiple* behaviours across *multiple* robots. Experimental results have shown that this approach is capable of evolving a diverse swarm, with multiple behaviours being exhibited by active robots. Further to this, we have provided new insights into how to maximise diversity, showing that amalgamating maps and exploiting memory facilitates spread of information across the swarm. We also noted that if the local archives distributed across a swarm can be aggregated into a *swarm-map* then a comprehensive library of potential behaviours could be obtained for future use. This approach could also be applied to evolve a library of behaviours for a single robot in *parallel* with the caveat that the environment each individual robot experiences is shared and manipulated by all the other robots.

The paper provides new evidence that behavioural diversity *can* be generated across a swarm without requiring reproduction isolation or the addition of a market-mechanism [7, 14]. We recognise that this comes with a caveat that functional traits appropriate to the required task must be pre-defined, and hence the designer is required to have *a priori* knowledge of the traits that might be useful. However, we believe this represents a useful first step in the quest to find a mechanism to evolve swarms that are robust to changing environments. Recent work by Hamann [8] discusses how selective pressure towards diversity might be generated based on an understanding of speciation dynamics in biology without having to pre-define appropriate traits. Although this is applied to generating diverse behaviours for a single robot navigating a maze, it may provide pointers for future work.

In contrast to some previous work in the QD domain in which archives are generated from millions of evaluations [5], we conduct experiments using relatively few evaluations: 1000 generations and 200 robots. Mouret *et al* [15] highlight that for a given budget of function evaluations, an EA allocates all of its evaluations to very few cells, and thus find good solutions for those cells. On the other hand, MAP-Elites has to distribute the same budget across a much larger number cells which can result in lower fitness values being obtained. Running the EDQD treatments for more evaluations would likely improve EDQD precision, importantly without danger of convergence, due to its propensity to enforce diversity.

Finally, having shown that we can evolve a diverse swarm using EDQD, work is already under way to evaluate the central hypothesis

driving this work, i.e. that the evolved swarm will be robust to dynamic changes in its environment. A natural extension of this would be to then show that in gradually changing environments, the repertoire of behaviours known to the swarm could gradually and continually increase, i.e. that the swarm could demonstrate lifelong learning.

CODE

Code repository: <https://github.com/asteyven/EDQD-GECCO2018>

REFERENCES

- [1] Nicolas Bredèche, Jean-Marc Montanier, Berend Weel, and Evert Haasdijk. 2013. Roborobo! a Fast Robot Simulator for Swarm and Collective Robotics. *CoRR* abs/1304.2888 (2013).
- [2] Nicolas Bredèche and Jean-Marc Montanier. 2010. Environment-driven embodied evolution in a population of autonomous agents. In *International Conference on Parallel Problem Solving from Nature*. Springer, 290–299.
- [3] Leo Breiman. 2001. Random forests. *Machine learning* 45, 1 (2001), 5–32.
- [4] Lars Chittka and Helene Muller. 2009. Learning, specialization, efficiency and task allocation in social insects. *Communicative & Integrative Biology* 2, 2 (2009), 151–154.
- [5] Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. 2015. Robots that can adapt like animals. *Nature* 521, 7553 (2015), 503.
- [6] Stefanie K Gazda, Richard C Connor, Robert K Edgar, and Frank Cox. 2005. A division of labour with role specialization in group-hunting bottlenose dolphins (*Tursiops truncatus*) off Cedar Key, Florida. *Proceedings of the Royal Society of London B: Biological Sciences* 272, 1559 (2005), 135–140.
- [7] Evert Haasdijk, Nicolas Bredèche, and AE Eiben. 2014. Combining environment-driven adaptation and task-driven optimisation in evolutionary robotics. *PLoS one* 9, 6 (2014), e98466.
- [8] Heiko Hamann. 2015. Lessons from Speciation Dynamics: How to Generate Selective Pressure Towards Diversity. *Artificial Life* 21, 4 (2015), 464–480.
- [9] Emma Hart and K. Sim. 2017. On Constructing Ensembles for Combinatorial Optimisation. *Evolutionary Computation* (2017).
- [10] Patrick Hoehn, Teja Tschardt, Jason M Tylanakis, and Ingolf Steffan-Dewenter. 2008. Functional group diversity of bee pollinators increases crop yield. *Proceedings of the Royal Society of London B: Biological Sciences* 275, 1648 (2008), 2283–2291.
- [11] Lu Hong and Scott E. Page. 2004. Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences* 101, 46 (2004), 16385–16389. <https://doi.org/10.1073/pnas.0403723101>
- [12] Joel Lehman and Kenneth O Stanley. 2008. Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE*. 329–336.
- [13] Joel Lehman and Kenneth O Stanley. 2011. Evolving a diversity of virtual creatures through novelty search and local competition. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*. ACM, 211–218.
- [14] Jean-Marc Montanier, Simon Carrignon, and Nicolas Bredèche. 2016. Behavioral specialization in embodied evolutionary robotics: Why so Difficult? *Frontiers in Robotics and AI* 3 (2016), 38.
- [15] Jean-Baptiste Mouret and Jeff Clune. 2015. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909* (2015).
- [16] Iñaki Fernández Pérez, Amine Boumaza, and François Charpillet. 2014. Comparison of Selection Methods in On-line Distributed Evolutionary Robotics. In *ALife2014*. MIT Press, 282–289.
- [17] Mitchell A Potter and Kenneth A De Jong. 2000. Cooperative coevolution: An architecture for evolving coadapted subcomponents. *Evolutionary computation* 8, 1 (2000), 1–29.
- [18] Justin K Pugh, Lisa B. Soros, and Kenneth O. Stanley. 2016. Quality Diversity: A New Frontier for Evolutionary Computation. *Frontiers in Robotics and AI* 3 (2016), 40.
- [19] Justin K Pugh, Lisa B Soros, and Kenneth O Stanley. 2016. Searching for quality diversity when diversity is unaligned with quality. In *International Conference on Parallel Problem Solving from Nature*. Springer, 880–889.
- [20] Giovanni Squillero and Alberto Tonda. 2016. Divergence of character and premature convergence: A survey of methodologies for promoting diversity in evolutionary optimization. *Information Sciences* 329 (feb 2016), 782–799. <https://doi.org/10.1016/j.ins.2015.09.056>
- [21] Andreas Steyven, Emma Hart, and Ben Paechter. 2017. An Investigation of Environmental Influence on the Benefits of Adaptation Mechanisms in Evolutionary Swarm Robotics. In *Proceedings of GECCO '17*. ACM, 155–162.
- [22] Pedro Trueba, Abraham Prieto, Francisco Bellas, Pilar Caamaño, and Richard J Duro. 2013. Specialization analysis of embodied evolution for robotic collective tasks. *Robotics and Autonomous Systems* 61, 7 (2013), 682–693.