Distributed Embodied Evolution over Networks

Anil Yaman

Department of Bio and Brain Engineering
Korea Advanced Institute of Science and Technology
Daejeon, 34141, Republic of Korea

Giovanni Iacca

Department of Information Engineering and Computer Science
University of Trento
Via Sommarive 9, 38123 Povo (Trento), Italy

Abstract

In several network problems the optimum behavior of the agents (i.e., the nodes of the network) is not known before deployment. Furthermore, the agents might be required to adapt, i.e. change their behavior based on the environment conditions. In these scenarios, offline optimization is usually costly and inefficient, while online methods might be more suitable. In this work, we use a distributed Embodied Evolution approach to optimize spatially distributed, locally interacting agents by allowing them to exchange their behavior parameters and learn from each other to adapt to a certain task within a given environment. Our results on several test scenarios show that the local exchange of information, performed by means of crossover of behavior parameters with neighbors, allows the network to conduct the optimization process more efficiently than the cases where local interactions are not allowed, even when there are large differences on the optimal behavior parameters within each agent's neighborhood.

Keywords: Embodied Evolution, Distributed Evolution, Genetic Algorithms, Networks, IoT.

1. Introduction

Networks of agents, such as sensor networks, wireless networks, swarms of drones or terrestrial robots, etc. are used nowadays in many tasks in the context of environment exploration, monitoring, and Internet of Things (IoT) applications. Typically, the behavior of the network as a whole

Email addresses: `anilyaman@kaist.ac.kr` (Anil Yaman), `giovanni.iacca@unitn.it` (Giovanni Iacca)

derives from the *local behavior* of each single agent. In this work, we consider a *local behavior* any form of agent-local decision based on the environment, be it an action depending on external stimuli (in the case of agents with actuation capabilities), or the update of a data-driven model used e.g. for classification or forecast tasks. Modelling the possible mutual interactions of such agent-local behaviors (which can be very different across different parts of the network), as well as the interactions of the agents with the environment, can be difficult. In fact, the environment conditions might not be known *a priori*, i.e. before deployment, and may change dynamically and unexpectedly. As such, finding the optimal behavior of network agents offline, analytically, can be challenging: on the one hand, an analytical formulation is rarely available; on the other, collecting all the agents to tune their parameters when needed can be inefficient -and expensive- especially when the number of agents is large. Moreover, in some cases the agents might not even be accessible for collection or data extraction, e.g. if their position is not known or they are placed in hard-to-access environments such as underground tunnels or pipes [1].

For all these reasons, a large body of literature has investigated various concepts such as *autonomic*, or *self-adaptive* networking [2, 3], in the quest for solutions to make networks capable of adapting automatically to the environment and thus optimizing autonomously, at runtime, their behavior. In this area, bio-inspired techniques -especially distributed Evolutionary Algorithms (EAs)- have attracted a great attention [4, 5] due to their intrinsic adaptivity. It is worth mentioning though that the idea of distributed EAs was initially introduced mainly in the context of numerical optimization, in the form of structured populations of solutions evolved in parallel over multiple cores or over computing networks [6, 7, 8]. More recently, distributed versions of Differential Evolution have been presented for instance in [9, 10], and [11], among which the latter work introduced a distributed optimization framework over P2P networks. A special flavour of distributed EAs is Cellular Evolution (CE), originally devised for Alife studies with explicit Cellular Automata (CA) models [12], and later applied e.g. to evolvable hardware [13] and other numerical optimization problems [14], where a population of individuals is restricted to interact (i.e., mate) locally. This can help preserve the population diversity longer. Of note, in CE all individuals optimize the same fitness function.

As an alternative paradigm to using a network as a collection of computing nodes for solving an optimization problem, more recently distributed EAs have been used also in physical or simulated networks to evolve agent-local parameters. One field of application of this concept is Wireless Sensor

2

Networks (WSNs), for which previous works have proposed e.g. CA [15, 16], distributed EAs [17, 18] and distributed Genetic Programming (GP) [19, 20] to evolve the sensor nodes' parameters and functioning logic. However, robotics is the area where distributed EAs have shown their greatest potential so far. In this context, Embodied Evolution (EE), i.e. the idea of running an EA distributedly on a group of agents, has been successfully applied to optimize the behavior of robotic swarms in various works, see e.g. [21, 22, 23, 24]. In particular, *environment-driven* EE, a form of artificial ecology where the environment conditions guide the evolutionary process of a collection of agents, has emerged as a promising way to obtain truly self-adaptive swarms [25]. As such, this research area has been very active in recent years, advancing not only on the algorithmic aspects [25, 26, 27, 28, 29, 30], but also on the application side: for instance, an interesting application on indoor surveillance and location has been proposed in [31, 32].
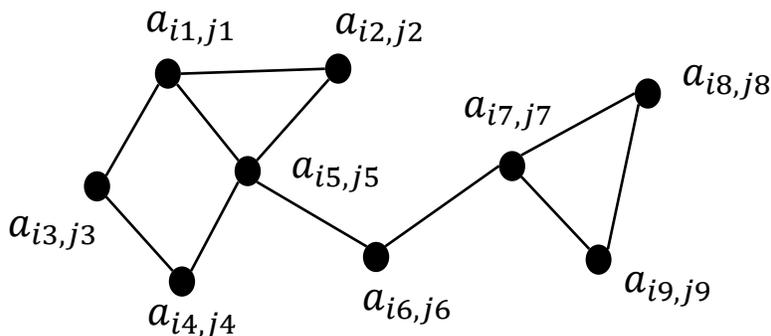

Figure 1: An illustration of a network of spatially distributed, locally connected agents.

In this work, we continue this research stream by investigating the effect of *local information exchange* on distributed Embodied Evolution. However, differently from most of the existing literature on EE focusing on robotics, here we focus specifically on networks. We consider networks of agents distributed at fixed locations within a certain environment, see Figure 1 for an example. We assume that each agent is required to perform a certain local behavior, that is described by some agent-local parameters, and that in general the optimal local behaviors are different across the agents as they depend on the local environmental conditions. Nevertheless, in some cases the behaviors of the agents in close proximity can be similar, since they may share similar environmental conditions. We also assume that the parameters for the optimal local behaviors are not known at the time of deployment. Therefore, in our setup the agents are required to optimize their local

parameters autonomously while they operate within the environment. To achieve that, we use a distributed embodied EE, allowing each agent to run a population-less EA *in situ*, to optimize its own behavior. In contrast to CE, we assume that the fitness function of each agent can be different.

We introduce local information exchange in terms of exchange of behavior parameters with local neighbors: we consider various kinds of parameter exchange, such as 1) copying the parameters from the best or a random agent in the neighborhood, with some mutation, or 2) exchange the parameters partially, i.e. by means of crossover of parameters with neighbors. We hypothesize that using parameter exchange as crossover and local perturbations for mutation leads to a social learning process through which it is possible to adapt a network of agents to a task at runtime, and in a distributed fashion. While it may be intuitive to exchange parameters with neighbors when their local conditions are similar, we are also interested in investigating the effect of parameter exchange in the cases where the optimal behaviors of the neighboring agents are different. Therefore, our main research questions are:

1. What are the conditions under which it is more appropriate to exchange local parameters across nodes in an EE setting?

2. What are the most appropriate parameter exchange mechanisms, and parameters thereof?

To answer these questions, we devise three test problems –of which one is a real-world IoT application taken from the literature– with different network sizes and different levels of similarities in terms of optimal behavior within a neighborhood. Additionally, these three problems show how the generic definition of behavior introduced above encompasses different kinds of applications, ranging from distributed actuation (in a synthetic illumination task) to a distributed data-driven model of human presence and activity (in a realistic IoT scenario).

Overall, our results show that local information exchange is beneficial in any case -even with large differences of the optimal parameters within a neighborhood- while the optimization performance can be largely affected by the frequency of the parameter exchange (i.e., the crossover probability) and the mutation rate. To summarize, our main contributions are the following:

- We apply EE to networks, with a particular emphasis on IoT systems such as sensor networks and alike: we believe that this represents a novel (and promising) application area of EE, which so far has been applied almost exclusively to swarm robotics.

4

- We specifically include in our experiments cases where the local fitness functions are different. This is especially important because sharing local parameters may not be necessarily advantageous in those cases. To the best of our knowledge, this aspect is not well-studied in previous works on EE in robotics, where usually the fitness function is assumed to be the same for all robots (e.g. their energy level, such as in [21]), but also in most of the literature on CE, where little or no attention has been given to possible differences on the local fitness functions.

- We perform extensive tests using several different variants of local information exchange strategies, also studying their parametrization, and assess their effect on the evolutionary process with a thorough statistical analysis (see the Appendix). We find that this in-depth analysis of the parameters' effect on EE schemes is also unprecedented in the previous related literature.

The remaining of this paper is structured as follows. In the next section, we introduce the general scheme for distributed Embodied Evolution over networks. Then, Section 3 describes the experimental setup, while Section 4 presents the numerical results. Finally, in Section 5 we give the conclusions and highlight possible future research directions.

2. Methods

Without loss of generality, we consider a system composed of a collection of k agents, spatially distributed on a 2-dimensional plane at fixed locations indexed by i, j, with i and $j \in \mathbb{N}_0$. Each agent $a_{i,j}$ must show a behavior such that it optimizes a certain local function. The performance of each agent is measured by its local fitness value $f_{i,j}$, that measures how good is the agent's behavior w.r.t. the desired local function. Furthermore, each agent can communicate (bi-directionally) with its local neighbors defined by a neighborhood function N. An example of such a system is shown in Figure 1: for instance, the neighborhood function of $a_{i5,j5}$ is defined by $N = \{(i1, j1), (i2, j2), (i3, j3), (i4, j4), (i6, j6)\}$, i.e. the set of indices of the agents that are connected to $a_{i5,j5}$.

We employ an Embodied Evolution approach [24] where each agent runs a (population-less) evolutionary algorithm to optimize its own behavior parameters, which represent the agent's genotype. This can be achieved with or without information sharing (i.e. sharing all or some of the elements of the genotypes) with the neighboring agents. In the case where there is no information sharing, each agent can iteratively test a randomly perturbed version of its genotype, and store the

genotype that performs the best (in the following, we will refer to this approach as HillClimbing). In the case where information sharing is allowed, the agents can copy and/or perform crossover with the genotypes of their neighbors. As discussed earlier, the optimal behavior of each agent can be potentially different from that of its neighbors, and studying the possible advantages/disadvantages of sharing information in this condition is one of the main focuses of this work.

We define two helper functions, $best(\boldsymbol{X}^{i,j})$ and $rand(\boldsymbol{X}^{i,j})$, that return respectively the genotype of the best and the genotype of a (uniformly selected) random agent in the neighborhood of a given agent $\boldsymbol{X}^{i,j}$. In real-world scenarios, this may be implemented by allowing each agent to communicate its fitness and genotype with its neighbors. In this case, and assuming that the local fitness values are comparable, it would be possible for each agent to pick the best genotype, or a random genotype, to perform evolutionary operators.

Finally, we use two evolutionary operators, namely *uniform crossover* and *Gaussian mutation*, applied in this order. The crossover operator generates a new genotype $\boldsymbol{X}'$ from two given parent genotypes, $\boldsymbol{X}^{i,j}$ and $\boldsymbol{X}^{k,l}$, as follows: first, $\boldsymbol{X}^{k,l}$ is copied into $\boldsymbol{X}'$; then, each element of $\boldsymbol{X}^{i,j}$ is copied into the corresponding element of $\boldsymbol{X}'$ with a crossover rate $cr \in [0, 1]$. The crossover probability $cp \in [0, 1]$ specifies the probability of performing the crossover operator. If the genotype of the agents consists of a single parameter, then we use *arithmetic crossover*, that computes the mean of the parents' parameters, i.e. $x' = (x^{i,j} + x^{k,l})/2$ (note that in this case cr is not needed). After crossover, the mutation operator perturbs each element of $\boldsymbol{X}'$ using a Gaussian mutation (sampled independently for each element) $\mathcal{N}(0, \sigma)$ with zero mean and standard deviation σ (in the following, we will refer to σ as the mutation rate, mr).

For illustration purposes, we provide in Algorithm 1 a pseudo-code of a version of the algorithm in which we apply crossover with a random neighbor. We dub this algorithm as XoverRand. The procedure randomInitialize() indicates a random initialization of the initial behavior parameters $\boldsymbol{X}^{i,j}$, where each element is uniformly sampled within its range $[lb, ub]$, being lb and ub the lower and upper bound respectively, while g and g_{max} indicate, respectively, the current and the maximum number of generations.

3. Experimental setup

In this section, we provide the details of our experimental setup. We are mainly interested in investigating the performance of Embodied Evolution with and without exchanging information, i.e.

6

Algorithm 1 Embodied evolution of an agent $a_{i,j}$ (and its corresponding behavior $\boldsymbol{X}^{i,j}$) with crossover applied with a random neighbor (XoverRand).

```
 1: procedure Evolve
 2:      X^{i,j} ← randomInitialize()
 3:      f_{i,j} ← eval(X^{i,j})
 4:      g ← 0
 5:      while g < g_max do
 6:          X' ← xover(X^{i,j}, rand(X^{i,j}), cp, cr)        ▷ Crossover with random neighbor
 7:          X' ← X' + 𝒩(0, σ)                                ▷ Gaussian mutation
 8:          f' ← eval(X')
 9:          if f' < f_{i,j} then                              ▷ Update (assuming minimization)
10:              f_{i,j} ← f'
11:              X^{i,j} ← X'
12:          end if
13:          g ← g + 1
14:      end while
15: end procedure
```

exchanging each agent's genotype (behavior parameters) with its neighbors, as well as the effect of the frequency of information sharing. Therefore, we test different versions of the Embodied Evolution algorithm introduced in the previous section, including variants in which it is allowed to exchange information across agents, and variants in which this exchange does not occur.

In the following, we first define the three test problems we used in our experimentation (and, for each problem, the relative scenarios). Then, we describe the five versions of the algorithm we tested on these problems.

In the first two problems (dubbed as "imitation" and "illumination", respectively), we use a 2-dimensional environment represented as an $m \times n$ grid. Each cell i,j on this grid is occupied by an agent $a_{i,j}$. Of note, in this setup the grid topology makes the EE model similar to Cellular Evolution, although as discussed earlier differently from CE here we also consider cases where the local fitness functions are different. We use the following neighborhood function: $N_{Moore} = \{(i-1,j-1),(i,j-1),(i+1,j-1),(i-1,j),(i+1,j),(i-1,j+1),(i,j+1),(i+1,j+1)\}$. This function, known as the Moore Neighborhood [33], allows each agent to communicate with its nearest horizontal, vertical and diagonal neighbors. The agents that are located at the border of the environment can communicate only with their existing neighbors.

In the third problem, concerning a distributed model of indoor human presence and activity, we instead use three different irregular (i.e., not grid-shaped) network topologies, shown in Figure 5, where each node in the network corresponds to an agent and the neighborhood function is the 1-hop neighborhood according to the given topology.

(a) $t = 1$ (b) $t = 50$ (c) $t = 100$

Figure 2: Imitation problem: Visualization of three images (ground truth) at $t = 1$, $t = 50$ and $t = 100$.

3.1. Imitation problem

We define as *"imitation problem"* a generic problem in which each agent in a network is required to "imitate" a desired pattern (the ground truth) over time. We demonstrate this by optimizing a network of agents to imitate a sequence of 100 images taken from the well-known MNIST hand-written digits dataset [34]. For this problem, we devise two scenarios, differing for the number of agents in the network and the number of parameters describing each agent's behavior.

3.1.1. 28×28 scenario

In this scenario, the network consists of 28×28 agents, each one corresponding to one pixel of the MNIST images. The grayscale intensity in each cell is controlled by agent $a_{i,j}$ at time $t \in \{1, 2, \ldots, 100\}$, where t corresponds to the t-th image. Thus, the genotype of $a_{i,j}$ is $\boldsymbol{X}^{i,j} = (x_1^{i,j}, x_2^{i,j}, \ldots, x_{100}^{i,j}) \in [0,1]^{100}$, where each t-th element determines the grayscale intensity (the phenotype) corresponding to the agent's cell at time t. A visualization of some selected images (ground truth) at indices $t = 1$, $t = 50$ and $t = 100$ is shown in Figure 2.

In the initialization phase, the genotype of each agent is generated by randomly sampling its genes in $[0, 1]$, with uniform probability. We use then Embodied Evolution to find the optimal parameters of each agent that can collectively "imitate" the sequence of 100 images as close as possible. Thus, in total, there are 7.84×10^4 parameters to be optimized in the network (28×28 agents $\times 100$ parameters). The fitness value of each agent is calculated as follows:

$$f_{i,j} = \frac{1}{100} \sum_{t=1}^{100} \mid I_{i,j}(t) - x_t^{i,j} \mid \tag{1}$$

8

where $I_{i,j}(t)$ denotes the desired i,j pixel value of the t-th image selected from the MNIST dataset (pixel values are scaled in the range $[0,1]$), and $x_t^{i,j}$ represents the corresponding t-th element of the agent's genotype $\boldsymbol{X}^{i,j}$. The *collective fitness* of the network at generation g, F_g, is then computed as the average fitness across the agents ($F_g = \sum_i \sum_j f_{i,j}$). This value is used for comparing various versions of the Embodied Evolution algorithm with different parameter settings. We run each algorithm for 10 independent runs, each one consisting of 20000 generations.

3.1.2. 7×7 scenario

One of the key aspects of Embodied Evolution approaches is to make use of the neighboring agents during the evolutionary process. As seen earlier, an agent can copy for instance the behavior parameters of the best performing agent in its neighborhood, with some mutation. However, not always these parameters can benefit the agent, since the optimal behavior of each agent can be very different from that of its neighbors. The imitation problem demonstrates this point in that there can be large differences between adjacent cells, e.g. at the edge between the background and the digit. Moreover, these differences change from one digit to another. To further stress these differences, we define an additional scenario where we reduce the number of agents to 7×7, by assigning to each agent the control of a tile made of of 4×4 pixels. As such, in this case each agent controls the intensity level of 16 pixels (rather than just one as in the 28×28 scenario), again for a sequence of 100 images. Thus, the number of parameters per agent increases from 100 to $4 \times 4 \times 100 = 1600$. Also in this case the behavior parameters are encoded in the genotype of each agent in a real-valued vector. However, differently from the previous scenario the spatial differences of adjacent tiles are larger, such that the difference between the optimal parameters of two neighboring agents is in turn larger.

To make this point clearer, in Figure 3 we show the average (across 100 MNIST images) differences of the optimal parameters of the agents with their neighbors for the two scenarios (28×28 vs 7×7 agents). For each cell, we first find the average (across genes) differences with each of its neighbors, and then we take the average (across neighbors) of these differences. The two resulting matrices are finally normalized by scaling each cell value w.r.t. the maximum value among all cells in the two scenarios: in fact, the maximum value in the 7×7 scenario (Figure 3b) is about twice as big as the maximum value in the 28×28 scenario (Figure 3a), thus both matrices are scaled w.r.t. the maximum value in the 7×7 scenario. A higher grayscale intensity in a cell indicates a higher difference with its neighbors. It can be observed that the differences are much higher in the

9

middle of the images (this is a consequence of the different shapes of the digits, while the cells at the border are more similar since they encode the background), and in general the differences in the 7×7 scenario are much higher (this is a consequence of the lower agent density, which leads to each agent controlling more pixels, such that the desired behavior can be quite different between adjacent tiles).

In this scenario, we run various versions of the Embodied Evolution algorithm for 10 independent runs, each one consisting of 100000 generations: compared to the 28×28 scenario, we increase the number of generations since the number of parameters per agent is 16 times higher. We use again the collective fitness F_g for comparisons among the tested algorithms.

(a) 28×28 scenario (b) 7×7 scenario

Figure 3: Imitation problem: Average differences (represented in grayscale) between the optimal parameters of each agent and those of the neighboring agents in the two scenarios.

3.2. Illumination problem

This problem illustrates a possible application of the imitation problem. In this case, a network of agents is required to learn to optimally illuminate an environment. We consider a 2-dimensional grid of $25 \times 50 = 1250$ agents, where each agent controls the illumination level in its cell. We assume that different locations in the environment receive different levels of natural light during the day. Each agent $a_{i,j}$ is then required to learn the optimal local illumination, based on the hour of the day, $t \in \{0, 1, \ldots, 23\}$, and its location, i, j. We define the optimal illumination (ground truth) for each cell location as follows:

$$l_{i,j}(t) = sin\left(\frac{2\pi j}{n} + \frac{2\pi t}{24}\right) \tag{2}$$

10

$l_{i,j}(t) = sin\left(\frac{2\pi j}{n} + \frac{2\pi t}{24}\right)$, where n is the maximum value of j. To convert $l_{i,j}(t)$ to pixel values, we scale its values to $[0, 1]$ (i.e. $l_{i,j}(t) = (l_{i,j}(t) + 1)/2$). A visualization of this environment at $t = 0$, $t = 7$ and $t = 15$ is shown in Figure 4.

In this case we further specify the problem in two scenarios, one in which each agent's behavior depends on a single real-valued parameter, and one in which the behavior is represented with a vector of real values.

(a) $t = 0$ (b) $t = 7$ (c) $t = 15$

Figure 4: Illumination problem: Visualization of the optimal illumination (ground truth) at $t = 0$, $t = 7$ and $t = 15$.

3.2.1. Single parameter scenario

In this scenario, each agent has only one real-valued parameter, $x^{i,j} \in [0, 50]$. The agent's behavior is then obtained by replacing in Equation (2) the desired location parameter, j, with the agent parameter $x^{i,j}$, and setting $n = 50$, as follows:

$$a_{i,j}(t) = sin\left(\frac{2\pi x^{i,j}}{50} + \frac{2\pi t}{24}\right) \tag{3}$$

At the beginning of the algorithm, $x^{i,j}$ is randomly initialized in $[0, 50]$. The fitness value of each agent $f_{i,j}$ is calculated as the average the difference from the optimal illumination during the day, namely:

$$f_{i,j} = \frac{1}{24} \sum_{t=0}^{23} | l_{i,j}(t) - a_{i,j}(t) | \tag{4}$$

We use the Embodied Evolution approach to find the optimal parameters of each agent that can collectively produce the desired illumination pattern. The whole network has in this case 25×50 parameters in total. We use again the collective fitness F_g (defined as for the imitation problem) to perform comparisons among various versions of the Embodied Evolution algorithm. In this case, we perform 10 independent runs for each version of the Embodied Evolution algorithm for 5000 generations.

3.2.2. Vector representation scenario

In this scenario, we encode the agents' behavior using a 24-dimensional real-valued vector (similarly to the imitation problem), $\boldsymbol{X}^{i,j} = (x_1^{i,j}, x_2^{i,j}, \ldots, x_{24}^{i,j}) \in [0,1]^{24}$, to control the agent-local level of illumination at each time step $t = \{0, 1, \ldots, 23\}$. Here, the values in $[0,1]$ encode the amount of light, from no light (0) to the maximum amount of light (1). The main difference w.r.t. the single parameter scenario is then that in this case we do not enforce a sinusoidal pattern in each agent, but this should be discovered by the Embodied Evolution process. At the beginning of the algorithm, the vector $\boldsymbol{X}^{i,j}$ of each agent is randomly initialized in $[0,1]^{24}$. Similarly to the single parameter scenario, the fitness value of each agent $f_{i,j}$ is the average difference from the optimal illumination during the day, namely:

$$f_{i,j} = \frac{1}{24} \sum_{t=0}^{23} \mid l_{i,j}(t) - x_t^{i,j} \mid \tag{5}$$

In this scenario, there are in total 3×10^4 parameters (25×50 agents $\times 24$ parameters). We perform 10 independent runs for each version of the algorithm for 20000 generations (the number of generations is increased w.r.t. the single parameter scenario due to the larger number of parameters). We use again the collective fitness F_g to perform comparisons among the tested algorithms.

3.3. Distributed model of indoor human presence and activity

IoT is a growing application area of Machine Learning and distributed algorithms. Typical applications in this field involve, for instance, distributed anomaly detection [35, 36, 37, 38, 39], transmission power control [40] or other forms of distributed intelligence [41, 42]. Here, we consider an IoT-based distributed model in which a number of sensors nodes are deployed in an indoor environment in order to detect human presence and activity, an urgent matter in terms of security and privacy. One important difference w.r.t. the two previous problems is that in this case the local behavior is not intended as a node-local actuation affecting the environment (like setting the local intensity or illumination), but rather the output of a node-local data-driven model that does not affect the environment (in this case, the detected presence/activity). This somehow reflects the flexibility and general-purposeness of the Embodied Evolution scheme.

In our experiments, we use a real-world dataset taken from [43], which consists of temperature and humidity sensor readings of three to five sensor nodes distributed in three different rooms. The sensor readings are associated with ground truth labels on the human presence (if the room

is occupied or not by a human) and activity (one of four human activities, i.e.: reading, standing, walking, and working on a PC).

For training and testing the model, we use a similar experimental setting as the original paper [43], i.e. we split the sensor readings of each node in windows of 150 samples each, where each window is obtained by shifting the previous one by 30 samples. Overall, we obtain about 800, 1400 and 450 windows for the human presence task, and 300, 450 and 150 windows for the human activity task in room A, B and C respectively. We use 80% of these windows for computing the fitness of the nodes' model (training accuracy) and the rest for testing. We perform testing at each generation during the evolutionary process, in order to collect the test accuracy values used for comparing the different versions of the EE algorithm.

As for the model, we employ on each node a feed-forward neural network (FFNN) architecture. We concatenate the raw measurements of the temperature and humidity sensors within a window and feed these input data to the FFNN. Therefore, the input layer of the networks has 300 neurons (150 each for temperature and humidity). The temperature and humidity measurements are scaled (separately) in the range $[0, 1]$, by subtracting the minimum value and dividing by the maximum value. We use one hidden layer with 100 hidden neurons.

We consider the human presence detection and activity detection as two separate tasks, i.e. we train and test a separate FFNN on each node for each of the two tasks. In the case of presence detection, the output layer consists of two neurons for deciding if there is a human in the room or not. In the case of activity detection, we use four output neurons, one for each activity. In both cases the output neuron with the highest activation level is selected as the final output of the network.

It is important to remark that in [43] no communication occurs between nodes. Rather, the authors extract the relevant features from the sensor readings and use these features to train a model –separately for each node– in order to detect the human presence and activity. Differently from this approach, here we assume instead that there is communication between nodes. However, we use communication only to allow parameter exchange for the copy and crossover operators, i.e., we do not use communication to transfer sensor data across the nodes. The network topologies (one per room) we assumed in our experimentation are shown in Figure 5.

In this problem, the behavior parameters of each node are the weights (in $[-1, 1]$) of its FFNN. Based on the size of the networks used, there are in total of $(300 + 1) \times 100 + (101 + 1) \times 2 = 30302$

13

Figure 5: Network topology for the three different rooms described in [43]: room A, B and C (from left to right).

and $(300 + 1) \times 100 + (101 + 1) \times 2 = 30504$ parameters in the human presence and human activity detection tasks respectively (+1 indicates bias neurons to hidden and output layers). These parameters are optimized using the distributed Embodied Evolution approach, by performing local mutations and/or crossover by means of each node's parameters exchange with its neighbors. As a note, since the nodes are positioned at different locations in each room, their local readings are usually different, although there is a certain spatial continuity. Therefore, the optimum parameter settings of their FFNNs may be different and as such our goal is to test if parameter exchange can be beneficial also in this application scenario.

Similarly to the imitation and illumination problems, also in this case we perform 10 independent runs for each version of the EE algorithm for 10000 generations, for each room and task. As said earlier, we use the test accuracy to perform comparisons among the different tested algorithms.

3.4. Versions of the Embodied Evolution algorithm

We tested in total five versions of the Embodied Evolution algorithm presented in the previous section, configured as follows:

- **XoverRand:** this version of the algorithm is the one given in Algorithm 1.

- **XoverBest:** this version is similar to the previous one, the only difference being that the function $rand(\boldsymbol{X}^{i,j})$ in Line 6 of Algorithm 1 is replaced with $best(\boldsymbol{X}^{i,j})$. In other words, this algorithm performs crossover with the best agent in each agent's neighborhood.

- **HillClimbing:** in this version, we do not use any neighborhood function, therefore the agents

do not share any information with their neighbors. We modify Algorithm 1 by deleting Line 6 (where the crossover operator is used), and changing Line 7 with: $\boldsymbol{X}' \leftarrow \boldsymbol{X}^{i,j} + \mathcal{N}(0, \sigma)$.

- **CopyBest:** in this version, each agent copies the genotype of the best agent in its neighborhood and then applies the mutation operator. Similar to HillClimbing, crossover is not used. Thus, we modify Algorithm 1 by deleting Line 6 and changing Line 7 with: $\boldsymbol{X}' \leftarrow best(\boldsymbol{X}^{i,j}) + \mathcal{N}(0, \sigma)$.

- **CopyRand:** in this version, each agent copies the genotype of a randomly selected agent in its neighborhood and then applies the mutation operator. Similarly to CopyBest, also in this case crossover is not used. Thus, we modify Algorithm 1 by deleting Line 6 and changing Line 7 with: $\boldsymbol{X}' \leftarrow rand(\boldsymbol{X}^{i,j}) + \mathcal{N}(0, \sigma)$.

4. Experimental Results

On the three problems and related scenarios introduced above, we tested the five versions of the EE algorithm with different combinations of mutation rates (mr) and crossover probability (cp)[1]. We use uniform crossover ($cr = 0.5$) for all experiments. In the case of the imitation and illumination problems with vector representation (both with domain $[0, 1]$), we use $mr \in \{0.0001, 0.001, 0.01\}$, while for the illumination problem with single parameter (with domain $[0, 50]$) we use $mr \in \{0.005, 0.05, 0.5\}$. In the case of the distributed model of indoor human presence and activity (with domain $[-1, 1]$), we use $mr \in \{0.0001, 0.001, 0.01\}$.

It should be noted that the imitation and illumination problems are formulated as *minimization* problems, while the distributed model problem is a *maximization* problem.

4.1. Imitation problem

In the following, we analyze separately the numerical results on the imitation problem in the two scenarios described above: 28×28 agents and 7×7 agents.

4.1.1. 28×28 scenario

We tested the five Embodied Evolution algorithms, with $cp \in \{0.2, 0.5, 1.0\}$, $cr = 0.5$, and $mr \in \{0.0001, 0.001, 0.01\}$. Figure 6a shows the average (across 10 runs per algorithm) collective fitness (F_g) trends obtained with the tested algorithms. The shaded areas represent the range $\pm$

---

[1] We report a detailed comparative analysis of the different algorithms and parameter settings in the Appendix.

std. dev. In the case of CopyBest, CopyRand, XoverBest and XoverRand we show the fitness trend obtained with the best parameter setting according to the Nemenyi test [44], see Appendix A.

We observe that performing crossover with a random neighbor performs better than performing crossover with the best neighbor, which in turn performs better than not doing any crossover. This indicates that exchanging partial components of the genotype helps the optimization process, even though the optimal parameters of the neighbors are different. The fact that crossing over with a random neighbor is more efficient than doing that with the best neighbor is likely due to a higher chance of selecting a neighbor whose optimal behavior is closer to that of the focal agent, which might not be the case of the best neighbor (e.g. if the focal agent is on the digit while its best neighbor is on the background).

We further observe that the versions that copy the neighbors perform better than HillClimbing. Also, when crossover is not used, lower mutation rates appear to slow the convergence: being the algorithm population-less, the only way to converge faster is to perform larger mutations. However, we observe the opposite effect when crossover is used: since mutation is performed *after* crossover, in this case too high mutation rates might obliterate any advantage obtained from the neighbor's exchanged genes. Overall, XoverRandCP05 (this notation, used also in the following, indicates $cp = 0.5$) appears to perform the best. However, as we show in the Appendix, different mutation rates perform differently ($mr = 0.01$ performs the worst).

Figure 7 shows a visualization of the results obtained by HillClimbing, CopyBest and Xover-RandCP05 (with $mr = 0.001$) during the evolutionary process[2]. We observe that different versions of the algorithm show different behaviors. For instance, with HillClimbing the reconstructed image appears very noisy: whereas some agents appear to imitate well the ground truth, there are many isolated agents that are not optimized. This is expected since each agent is trying to optimize its own local fitness function without any parameter exchange. On the other hand, in the case of CopyBest and XoverRand we observe that the images are much smoother and less noisy. Also, CopyBest seems unable to optimize the agents in the middle of the image: this is due to the fact that those agents have much larger differences (in terms of their optimal parameters) w.r.t. their neighbors. Therefore, copying the neighbors provides little or no benefit.

---

[2]Video of the evolutionary process available at: `https://youtu.be/DfjSvKA6KNI`.

(a) Imitation problem (28×28 scenario)

(b) Imitation problem (7×7 scenario)

(c) Illumination problem (Single parameter scenario)

(d) Illumination problem (Vector representation scenario)

Figure 6: Imitation and illumination problems: Collective fitness (mean $\pm$ std. dev. across 10 runs) obtained with different EE algorithms during the evolutionary process. For CopyBest, CopyRand, XoverBest and XoverRand we show the fitness trend obtained with the best parameter setting.

4.1.2. 7×7 scenario

We tested the five Embodied Evolution algorithms with the same parameter settings tested for the 28×28 scenario. Figure 6b shows the average (across 10 runs per algorithm) collective fitness (F_g) trends obtained with the tested algorithms and the corresponding std. dev. In the case of CopyBest, CopyRand, XoverBest and XoverRand, again we show the fitness trend obtained with the best parameter setting according to the Nemenyi test, see Appendix A.

Similarly to the 28×28 scenario, the versions of the algorithm that do not employ crossover perform worse than those that use it. However in this case the difference between XoverRand and

17

Figure 7: Imitation problem (28×28 scenario): Visualization of the evolutionary process at generations $g \in \{100, 1000, 10000, 20000\}$ for each row from left to right (ground truth shown in Figure 2a). The rows present the results of HillClimbing, CopyBest and XoverRand ($mr = 0.001$) respectively.

XoverBest is much larger, with XoverBest obtaining similar results to those of the versions without crossover. This confirms our previous hypothesis, since in this case it is even more unlikely that the best agent in the neighborhood of one agent (which has more parameters) might provide a benefit during crossover. In any case, these results show that even if the optimal parameters of the neighbor agents are quite different (as we have shown in Figure 3), sharing behavior parameters with randomly selected neighbors helps the optimization process.

In Figure 8, we show the visualization of the results obtained by HillClimbing, CopyBest and XoverRandCP02 (with $mr = 0.001$) at generations $g = \{100, 1000, 10000, 100000\}$: also here we observe that HillClimbing produces a much noisier image, while XoverRandCP02 performs the best.

18

Figure 8: Imitation problem (7×7 scenario): Visualization of the evolutionary process at generations $g \in \{100, 1000, 10000, 100000\}$ for each row from left to right (ground truth shown in Figure 2a). The rows present the results of HillClimbing, CopyBest and XoverRand ($mr = 0.001$) respectively.

4.2. Illumination problem

In the following, we analyze separately the numerical results on the illumination problem in the two scenarios described above: agents controlled by a single parameter, and agents controlled by a vector of parameters.

4.2.1. Single parameter scenario

As seen earlier, in this scenario each agent aims to optimize a parameter corresponding to the phase shift of the sine function used for the ground truth, see Equations (2) and (3).

We tested again the five Embodied Evolution algorithms, with $cp \in \{0.2, 0.5, 1.0\}$ and $mr \in \{0.005, 0.05, 0.5\}$ (in this case we use arithmetic crossover, so cr is not needed). Figure 6c shows

19

the average (across 10 runs per algorithm) collective fitness (F_g) trends obtained with the tested algorithms, and the corresponding std. dev. Also in this case, for CopyBest, CopyRand, XoverBest and XoverRand we show the fitness trend obtained with the best parameter setting according to the Nemenyi test, see Appendix A.

The main observation is that in this case CopyRand obtains the best results, better than the algorithms with crossover. This might be due to the fact that, with one parameter, mutations cancel out the effect of the arithmetic crossover (which is essentially another kind of mutation), thus producing a slower convergence (or even stagnation) if crossover is applied. So, it is just more efficient to copy a random neighbor (rather than doing crossover) and apply a small mutation.

Figure 9 shows a visualization of the results obtained by HillClimbing, CopyBest and Xover-RandCP02 ($mr = 0.05$) during the evolutionary process. We observe that with no parameter sharing (HillClimbing), the results are very noisy, with several isolated agents that are non-optimal (Figures 9a-9d). With crossover, results are clearly better, see Figures 9i-9l. However, we observe in the results of XoverRand that at $g = 100$ there are many non-optimal agents at the right and left edges of the image, see Figure 9i. This is due to the problem formulation and the arithmetic crossover operator: in fact, there are two optimal values for the left and right edge of the image (0 and 50), due to the periodicity of the sine function. Thus, agents are optimized towards one of these values, which may be different from that of their neighbors. Moreover, arithmetic crossover in this case does not help because the fitness of the average of these two values is worse.

4.2.2. Vector representation scenario

We tested the five Embodied Evolution algorithms with the same parameter settings tested for the imitation problem. Figure 6d shows the average (across 10 runs per algorithm) collective fitness (F_g) trends obtained with the tested algorithms, and the corresponding std. dev. Also in this case, for CopyBest, CopyRand, XoverBest and XoverRand we show the fitness trend obtained with the best parameter setting according to the Nemenyi test, see Appendix A.

Similarly to imitation problem, we observe that the versions that use the crossover operator tend to perform better. Moreover, XoverRand performs the best for different mutation rates, as it does in the imitation problem. Also, we observe that smaller mutation rates lead to slower convergence in all cases.

Figure 10 shows a visualization of the results obtained by HillClimbing, CopyBest and Xover-

Figure 9: Illumination problem (single parameter scenario): Visualization of the evolutionary process at generations $g \in \{100, 1000, 3000, 5000\}$ in each row from left to right (ground truth shown in Figure 4a). The rows present the results of HillClimbing, CopyBest and XoverRandCP02 ($mr = 0.05$) respectively.

RandCP02 ($mr = 0.001$) during the evolutionary process[2]. The results of HillClimbing appear more noisy, with many isolated, non-optimal agents. Also, we observe vertical "stripes" in the results of CopyBest: this is due to the fact that the optimal parameters of the agents in each column are the same, therefore they tend to share the same parameters. Once again, we observe that crossover helps the optimization process even though the optimal parameters of neighboring agents are different.

4.3. Distributed model of indoor human presence and activity

Differently from the two previous problems, in this case many versions of the algorithm produced statistically equivalent results[3]. We show in Figure 11 the fitness trends (test accuracy) obtained with the Embodied Evolution algorithms for each room and task, averaged across 10 runs per algorithm, with the corresponding std. dev. As we did for the previous two problems, in the case

_____

[2]Video of the evolutionary process available at: https://youtu.be/5HZ-SMLyv_E.
[3]See Tables A.8-A.13 in the Appendix for pairwise comparisons of p-values.
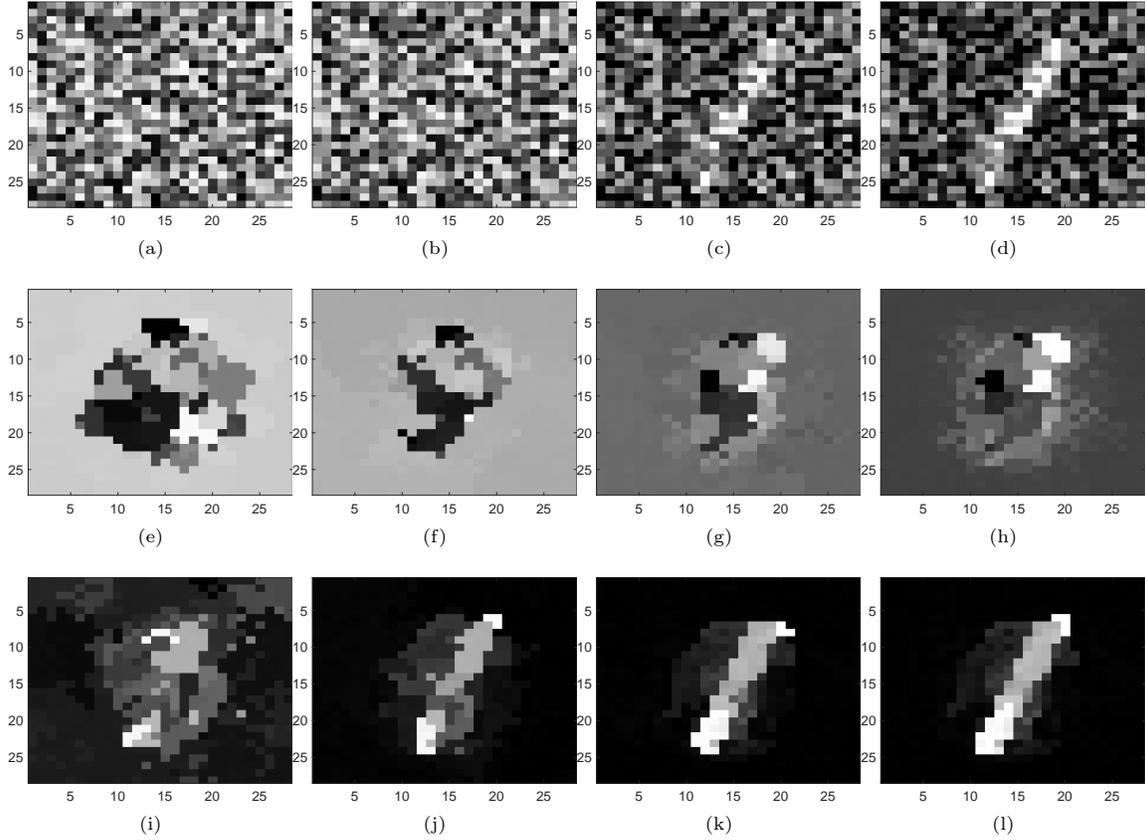
Figure 10: Illumination problem (vector representation): Visualization of the evolutionary process at generations $g \in \{100, 1000, 10000, 20000\}$ for each row from left to right (ground truth shown in Figure 4a). The rows present the results of HillClimbing, CopyBest and XoverRand ($mr = 0.001$) respectively.

of CopyBest, CopyRand, XoverBest and XoverRand we show the fitness trend obtained with the best parameter setting according to the Nemenyi test, see Appendix A for details. Overall, also on this problem the algorithms that use crossover perform better. In particular, XoverRand performs best in all rooms on the human presence task, and in room A and C on the human activity task. In the remaining case (human activity task in room B), XoverBest performs slightly better than XoverRand. However, as we discuss below this task appears intrinsically harder than the others, and the difference in performance across different algorithms is less clear. In all cases, HillClimbing shows the worst results, followed by CopyBest and CopyRand which are statistically equivalent in most cases except for the human presence task in room A, where CopyRand performs much better than CopyBest. These results confirm then the general trends that we observed on the imitation and illumination problems.

For completeness, in Tables 1-3, we report the model results (in terms of training and test accuracy) obtained with the best Embodied Evolution algorithm (according to the Nemenyi test)

on each node and each room for both tasks (human presence and human activity), in comparison with those reported in [43]. Please note that concerning the original results on the human activity task, except for the nodes A1 and C1 (i.e., the nodes showing the best and worst accuracy value, respectively) all the other results are reported only graphically, therefore the exact accuracy values are not available. Also, in [43] only the test accuracy results are reported.

From the results, it can be seen that the Embodied Evolution algorithm (in particular the XoverRand/XoverBest versions, depending on the room and task) is quite effective especially on the human activity task, although with different performances depending on the room and node. Limiting our analysis to the test accuracy, it can be noted that:

- On the human presence task, the results are –in general– slightly worse than those reported in [43]. This is likely due to the different Machine Learning model used (in [43] the model was a Naïve Bayes classifier, while in our case it is a FFNN), and to the fact that in [43] the sensor data are pre-processed (with statistical analysis and feature selection), while in our case we simply use the raw data, properly scaled to feed the Neural Network. Nevertheless, the accuracy values are high in all cases (better than the probability of random guessing, 50%). Furthermore, in the case of room C our method is able to find an accuracy value on nodes C3 and C4 that is even higher than that reported in [43] on node C1. Finally, as in [43] we observe that the task on room B seems to be more difficult (showing lower accuracy values) w.r.t. the other two rooms. This is probably due to the comparatively larger size of this room and the layout of the sensor network.

- On the human activity task, the results are much better than those reported in [43] in the case of room A and C, while in the case of room B they are quite worse. This finding is especially interesting because this multi-class classification problem is supposed to be harder than the human presence problem, that is instead binary (either a human is present in the room or not). In particular, we can observe that our method produces consistently higher accuracy values on all nodes in room A and C. In the case of room B, it produces better accuracy only on node B3, while in the other two nodes the accuracy is much lower. Again, this is likely due to the layout of room B and its corresponding sensor network, and to the fact that this network is composed of only three nodes, such that the parameter exchange is probably too limited to be effective.

To conclude, this IoT problem shows how Embodied Evolution can be easily applied to real-world network cases, where leveraging the parameter exchange can lead to improved performances. On

the other hand, the case of room B shows one possible limitation of our scheme, whose performance might be affected by the number as well as the position of the nodes, although these aspects might also depend on cost and logistic constraints of the specific application. As a side note, this observation is consistent with the literature on evolutionary algorithms, where various works, such as [45], have shown that "micro-population" algorithms (i.e., algorithms with a very small population size) tend to get stuck at local optima more easily.

Table 1: Distributed model (room A): Training/test accuracy at each node found by the best performing Embodied Evolution algorithm. The boldface indicates the best accuracy value (across the nodes) obtained on each task.

Node	Human presence			Node	Human activity		
	Training [ours]	Test [ours]	Test [43]		Training [ours]	Test [ours]	Test [43]
A1	85.46	**85.56**	**93.5**	A1	55.93	**57.47**	**56.8**
A2	**87.73**	85.51	89.7	A2	**59.09**	51.32	~ 35
A3	83.87	80.47	84.2	A3	57.94	52.04	~ 30
A4	83.44	81.45	78.4	A4	54.76	51.32	~ 24

Table 2: Distributed model (room B): Training/test accuracy at each node found by the best performing Embodied Evolution algorithm. The boldface indicates the best accuracy value (across the nodes) obtained on each task.

Node	Human presence			Node	Human activity		
	Training [ours]	Test [ours]	Test [43]		Training [ours]	Test [ours]	Test [43]
B1	**74.55**	**62.78**	**88.5**	B1	**35.59**	30.65	~ 54
B2	72.00	60.61	81.3	B2	33.48	28.85	~ 35
B3	73.11	62.56	66.8	B3	33.43	**39.35**	~ 30

Table 3: Distributed model (room C): Training/test accuracy at each node found by the best performing Embodied Evolution algorithm. The boldface indicates the best accuracy value (across the nodes) obtained on each task.

Node	Human presence			Node	Human activity		
	Training [ours]	Test [ours]	Test [43]		Training [ours]	Test [ours]	Test [43]
C1	92.13	86.67	**91.0**	C1	52.97	45.00	23.9
C2	92.73	91.11	88.6	C2	48.31	**56.67**	~ 24
C3	**93.59**	**92.22**	86.0	C3	55.60	48.33	~ 34
C4	90.76	**92.22**	88.7	C4	50.43	45.00	~ 24
C5	86.59	86.67	90.7	C5	**58.97**	48.33	~ 31

(a) Human presence (room A)

(b) Human activity (room A)

(c) Human presence (room B)

(d) Human activity (room B)

(e) Human presence (room C)

(f) Human activity (room C)

Figure 11: Distributed model: Test accuracy (mean ± std. dev. across 10 runs) obtained with different EE algorithms during the evolutionary process. For CopyBest, CopyRand, XoverBest and XoverRand we show the fitness trend obtained with the best parameter setting.

25

5. Conclusions

In many network applications, the optimal behavior of the agents is not known before deployment, and the agents often need to adapt to environment changes. Performing the optimization process offline, by collecting each agent to change its behavior parameters whenever it is needed, would be costly and ineffective. In this work, we used a distributed Embodied Evolution approach to optimize the behavior of networks of agents at runtime.

We assumed a collection of spatially distributed agents that can communicate locally. We used the local communication capability to exchange the agents' behavior parameters, so to copy the neighbors' parameters and/or use them for local perturbations and crossover. Intuitively, the agents with close proximity may be required to perform similarly (thus may have similar optimal parameters) since they are likely to share similar environmental conditions. In this case, it would make sense to share their parameters to learn from one another. Nevertheless, even when the optimal parameters are drastically different within a neighborhood, parameter exchange may still be helpful in the optimization process. To test these two cases, we devised a number of experimental scenarios with various levels of differences between the optimal parameters of any given agent and those of its neighbors. We further compared these results with the case where each agent optimizes its behavior without any parameter exchange.

We found that parameter exchange (through crossover) works the best in all cases -except with arithmetic crossover and small mutation rates- including when the differences of the optimal parameters of the agents with their neighbors are large. We also observed a certain variation of the performance depending on the parameter settings of the algorithm (i.e., mr and cp), therefore in the future it might be useful to evaluate (self-)adaptive evolutionary approaches. Among the various versions of the algorithm we tested, XoverRand demonstrated the best performance in general. Since this version does not require the fitness of the neighbors (due to random selection), in practical scenarios it might possible to broadcast only some randomly selected components of the agents' genotypes. In terms of limitations, the performance of the proposed EE scheme seems to be affected by the number as well as the position of the nodes, as we have observed in the IoT case.

This work was mainly aimed at demonstrating the effectiveness of the Embodied Evolution approach over networks. Therefore, we made some simplifying assumptions, in particular we considered fixed topologies with perfect communication. In future works, we will extend this approach

to more realistic cases, i.e. with dynamic connections and possibly noisy communication. Another assumption of this work was that the local fitness functions are in general different (in terms of optimum) across agents, but still comparable. However, this might not be the case of all applications: therefore, it may be interesting to see if parameter exchange is beneficial also in those cases where the local fitness functions are not directly comparable, and how embodied evolution could be applied (for instance, function similarity could be used to define the neighborhood function, to create niches/communities, or techniques from multi-factorial optimization [46, 47] could be applied). Another possibility would be to hybridize Embodied Evolution schemes with distributed gradient-based optimization [48] and other forms of diffusion strategies [49]. Finally, we plan to validate our approach with physical networks.

References

[1] A. Hallawa, G. Iacca, C. Sariman, T. Rahman, M. Cochez, G. Ascheid, Morphological evolution for pipe inspection using Robot Operating System (ROS), Mater. Manuf. Process. 35 (6) (2020) 714–724.

[2] G. Bouabene, C. Jelger, C. Tschudin, S. Schmid, A. Keller, M. May, The autonomic network architecture (ANA), IEEE J. Sel. Area Comm. 28 (1) (2009) 4–14.

[3] Y. Xiao, Bio-inspired computing and networking, CRC Press, 2016.

[4] T. Nakano, Biologically inspired network systems: A review and future prospects, IEEE T. Syst. Man Cyb. 41 (5) (2010) 630–643.

[5] F. Dressler, O. B. Akan, A survey on bio-inspired networking, Comput. Netw. 54 (6) (2010) 881–900.

[6] E. Alba, M. Tomassini, Parallelism and evolutionary algorithms, IEEE T. Evolut. Comput. 6 (5) (2002) 443–462.

[7] M. G. Arenas, P. Collet, A. E. Eiben, M. Jelasity, J. J. Merelo, B. Paechter, M. Preuß, M. Schoenauer, A framework for distributed evolutionary algorithms, in: International Conference on Parallel Problem Solving from Nature (PPSN), Springer, 2002, pp. 665–675.

[8] F. Maqbool, S. Razzaq, J. Lehmann, H. Jabeen, Scalable Distributed Genetic Algorithm Using Apache Spark (S-GA), in: International Conference on Intelligent Computing (ICIC), Springer, 2019, pp. 424–435.

[9] I. De Falco, D. Maisto, U. Scafuri, E. Tarantino, A. Della Cioppa, Distributed differential evolution for the registration of remotely sensed images, in: International Conference on Parallel, Distributed and Network-Based Processing (PDP), EUROMICRO, 2007, pp. 358–362.

[10] I. De Falco, U. Scafuri, E. Tarantino, A. Della Cioppa, A distributed differential evolution approach for mapping in a grid environment, in: International Conference on Parallel, Distributed and Network-Based Processing (PDP), EUROMICRO, 2007, pp. 442–449.

[11] M. Biazzini, A. Montresor, P2POEM: Function optimization in P2P networks, Peer Peer Netw. Appl. 6 (2) (2013) 213–232.

[12] M. Sipper, Studying artificial life using a simple, general cellular model, Artif. Life 2 (1) (1994) 1–35.

[13] M. Sipper, Evolution of parallel cellular machines, Vol. 4, Springer, 1997.

[14] E. Alba, J. M. Troya, Cellular evolutionary algorithms: Evaluating the influence of ratio, in: International Conference on Parallel Problem Solving from Nature (PPSN), Springer, 2000, pp. 29–38.

[15] Y. Wang, Z. Qian, D. Sun, C. Zhou, A cellular automata model for wireless sensor networks, in: International Conference on Systems (ICONS), 2012, pp. 1–6.

[16] S. Choudhury, K. Salomaa, S. G. Akl, A cellular automaton model for wireless sensor networks, J. Cell. Autom. 7 (3).

[17] G. Iacca, Introducing DOWSN: distributed optimization in wireless sensor networks, in: UK Workshop on Computational Intelligence (UKCI), IEEE, 2012, pp. 1–8.

[18] G. Iacca, Distributed optimization in wireless sensor networks: an island-model framework, Soft Comput. 17 (12) (2013) 2257–2277.

[19] D. M. Johnson, A. M. Teredesai, R. T. Saltarelli, Genetic programming in wireless sensor networks, in: European Conference on Genetic Programming (EuroGP), Springer, 2005, pp. 96–107.

[20] P. Valencia, P. Lindsay, R. Jurdak, Distributed genetic evolution in WSN, in: International Conference on Information Processing in Sensor Networks (IPSN), ACM/IEEE, 2010, pp. 13–23.

[21] R. A. Watson, S. G. Ficici, J. B. Pollack, Embodied evolution: Distributing an evolutionary algorithm in a population of robots, Robot Auton. Syst. 39 (1) (2002) 1–18.

[22] A. L. F. Perez, G. Bittencourt, M. Roisenberg, Embodied evolution with a new genetic programming variation algorithm, in: International Conference on Autonomic and Autonomous Systems (ICAS), 2008, pp. 118–123.

[23] A. Eiben, E. Haasdijk, N. Bredeche, Embodied, on-line, on-board evolution for autonomous robotics, in: Symbiotic Multi-Robot Organisms: Reliability, Adaptability, Evolution, Springer, 2010, pp. 361–382.

[24] N. Bredeche, E. Haasdijk, A. Prieto, Embodied evolution in collective robotics: A review, Front. Robot. AI 5 (2018) 12.

[25] N. Bredeche, J.-M. Montanier, W. Liu, A. F. Winfield, Environment-driven distributed evolutionary adaptation in a population of autonomous robotic agents, Math. Comp. Model. Dyn. 18 (1) (2012) 101–129.

[26] E. Haasdijk, N. Bredeche, A. Eiben, Combining environment-driven adaptation and task-driven optimisation in evolutionary robotics, PloS ONE 9 (6).

[27] J.-M. Montanier, S. Carrignon, N. Bredeche, Behavioral specialization in embodied evolutionary robotics: why so difficult?, Front. Robot. AI 3 (2016) 38.

[28] E. Hart, A. Steyven, B. Paechter, Improving survivability in environment-driven distributed evolutionary algorithms through explicit relative fitness and fitness proportionate communication, in: Genetic and Evolutionary Computation Conference (GECCO), 2015, pp. 169–176.

[29] P. Zahadat, H. Hamann, T. Schmickl, Evolving diverse collective behaviors independent of swarm density, in: Genetic and Evolutionary Computation Conference (GECCO) Companion, 2015, pp. 1245–1246.

[30] I. F. Pérez, S. Sanchez, Influence of local selection and robot swarm density on the distributed evolution of GRNs, in: International Conference on the Applications of Evolutionary Computation (Part of EvoStar), Springer, 2019, pp. 567–582.

[31] P. Trueba, A. Prieto, F. Bellas, R. J. Duro, Embodied evolution for collective indoor surveillance and location, in: International Work-Conference on the Interplay Between Natural and Artificial Computation (IWINAC), Springer, 2015, pp. 138–147.

[32] A. Prieto, F. Bellas, P. Trueba, R. J. Duro, Real-time optimization of dynamic problems through distributed embodied evolution, Integr. Comput-Aid. Eng. 23 (3) (2016) 237–253.

[33] L. Gray, A mathematician looks at Wolfram's new kind of science, Not. Am. Math. Soc. 50 (2) (2003) 200–211.

[34] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition, Proc. IEEE 86 (11) (1998) 2278–2324.

[35] H. H. Bosman, A. Liotta, G. Iacca, H. J. Wörtche, Anomaly detection in sensor systems using lightweight machine learning, in: IEEE International Conference on Systems, Man, and Cybernetics (SMC), IEEE, 2013, pp. 7–13.

[36] H. H. Bosman, A. Liotta, G. Iacca, H. Wortche, Online extreme learning on fixed-point sensor networks, in: IEEE International Conference on Data Mining (ICDM), IEEE, 2013, pp. 319–326.

[37] H. H. Bosman, G. Iacca, H. J. Wörtche, A. Liotta, Online fusion of incremental learning for wireless sensor networks, in: IEEE International Conference on Data Mining (ICDM), IEEE, 2014, pp. 525–532.

[38] H. H. Bosman, G. Iacca, A. Tejada, H. J. Wörtche, A. Liotta, Ensembles of incremental learners to detect anomalies in ad hoc sensor networks, Ad Hoc Netw. 35 (2015) 14–36.

[39] H. H. Bosman, G. Iacca, A. Tejada, H. J. Wörtche, A. Liotta, Spatial anomaly detection in sensor networks using neighborhood information, Inform. Fusion 33 (2017) 41–56.

[40] P. Pace, G. Fortino, Y. Zhang, A. Liotta, Intelligence at the edge of complex networks: The case of cognitive transmission power control, IEEE Wireless Communications 26 (3) (2019) 97–103.

[41] G. Dartmann, H. Song, A. Schmeink, Big data analytics for cyber-physical systems: machine learning for the internet of things, Elsevier, 2019.

[42] D. P. Kumar, T. Amgoth, C. S. R. Annavarapu, Machine learning algorithms for wireless sensor networks: A survey, Inform. Fusion 49 (2019) 1–25.

[43] P. Morgner, C. Müller, M. Ring, B. Eskofier, C. Riess, F. Armknecht, Z. Benenson, Privacy implications of room climate data, in: European Symposium on Research in Computer Security, Springer, 2017, pp. 324–343.

[44] P. Nemenyi, Distribution-free multiple comparisons, in: Biometrics, Vol. 18, International Biometric Society, 1962, p. 263.

[45] A. Yaman, G. Iacca, F. Caraffini, A comparison of three differential evolution strategies in terms of early convergence with different population sizes, in: AIP Conference Proceedings, Vol. 2070, 2019, p. 020002.

[46] A. Gupta, Y. Ong, L. Feng, K. C. Tan, Multiobjective multifactorial optimization in evolutionary multitasking, IEEE T. Cyb. 47 (7) (2017) 1652–1665.

[47] A. Gupta, Y. Ong, L. Feng, Multifactorial evolution: Toward evolutionary multitasking, IEEE T. Evolut. Comput. 20 (3) (2016) 343–357.

[48] M. Rabbat, R. Nowak, Distributed optimization in sensor networks, in: IEEE International Symposium on Information Processing in Sensor Networks (IPSN), 2004, pp. 20–27.

[49] J. Chen, A. H. Sayed, Diffusion adaptation strategies for distributed optimization and learning over networks, IEEE T. Signal Proces. 60 (8) (2012) 4289–4305.

[50] F. Wilcoxon, Individual comparisons by ranking methods, in: Breakthroughs in statistics, Springer, 1992, pp. 196–202.

Appendix A. Statistical analysis

In the following, we report the results of the pairwise comparisons (based on Wilcoxon Rank-Sum test [50], $\alpha = 0.05$), visually represented in a symmetric matrix format, and a post-hoc multiple-comparisons analysis (based on Nemenyi test [44], $\alpha = 0.05$), visually represented with Critical Difference plots. We performed both tests on the five versions of the Embodied Evolution algorithm described in the main text, with different parameter settings, as we tested them on the imitation problem (both 28×28 scenario and 7×7 scenarios), the illumination problem (both single parameter and vector representation scenarios), and the distributed model of indoor human presence and activity (for the three rooms and both tasks).

Appendix A.1. Wilcoxon Rank-Sum test

As for the imitation and illumination problems, in Tables A.4-A.7 it can be noted that, apart from a few occasional cases (slightly more frequent in the case of the imitation problem, 7×7 scenario), the null-hypothesis H_0 (statistical equivalence) is rejected in most cases, highlighting the fact that in general the different parameter settings and algorithm versions are indeed different (pairwise) in statistically significant terms.

Concerning the distributed model problem, in Tables A.8-A.13 it can be seen that there are "clusters" of algorithm versions that are pairwise equivalent (in particular the various parameter settings of XoverRand and XoverBest). It is also confirmed that the two tasks (human presence and human activity) appear more difficult in the case of room B, where most algorithms result equivalent.

Appendix A.2. Nemenyi test

In Figures A.12-A.21 algorithms that are considered statistically equivalent w.r.t. the Critical Difference (CD) calculated by the Nemenyi test (depicted as a segment on top of the plot) are graphically connected by a thick black line. Better algorithms get a lower rank. With this notation in mind, the main results of the Nemenyi test can be summarized as follows.

As for the imitation and illumination problems, it results that:

- In Figure A.12 (imitation problem, 28×28 scenario), XoverRand versions consistently rank before XoverBest, CopyRand/CopyBest and HillClimbing versions, in this order. Furthermore, all XoverRand versions and two XoverBest versions result statistically equivalent w.r.t. the CD.

The versions without crossover (CopyRand/CopyBest and HillClimbing) perform better (i.e. they get lower ranks) with higher mr values, while the contrary happens on the versions with crossover.

- In Figure A.13 (imitation problem, $7{\times}7$ scenario), the trend is similar to the $28{\times}28$ scenario, with XoverRand versions resulting statistically equivalent w.r.t. their parametrization, and ranking before the other versions (being CopyBest with $mr = 0.001$ the worst-performing algorithm). Similar observations on the effect of mr on the versions with/without crossover hold true also in this case.

- In Figure A.14 (illumination problem, single parameter scenario), CopyRand with $mr = 0.005$ ranks first, but it results statistically equivalent (w.r.t. the CD) to CopyRand with $mr = 0.05$ and $mr = 0.5$, as well as XoverBest and XoverRand versions with mr larger than 0.005. Also in this case, HillClimbing (with the lowest mr value) obtains the worst rank.

- In Figure A.15 (illumination problem, vector representation scenario), most of the XoverRand versions rank before XoverBest and the versions without crossover (although the general trend is somehow less clear than the previous cases).

As for the distributed model problem, from Figures A.16-A.21 it results that, in all cases (except the activity detection task in room B), XoverRand is assigned better ranks w.r.t. the other algorithms. In the case of the human activity detection task in room B, most of the algorithms appear to produce similar results (which, as we have seen in the main text, correspond to low accuracy values).

Overall, the above observations obtained from the multiple-comparisons analysis are in line what we discussed in Section 4 of the main text.

Table A.4: Wilcoxon Rank-sum test ($\alpha = 0.05$) on the pairwise comparisons between the algorithm versions tested on the imitation problem (28×28 scenario). "=" indicates that the null-hypothesis H_0 (statistical equivalence) is accepted (omitted on the diagonal cells). Empty cells indicate that the null-hypothesis is rejected.

		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27
HillClimbingMR01	1																											
CopyBestMR01	2											=																
CopyRandMR01	3																											
XoverBestCP1MR01	4																											
XoverBestCP05MR01	5																								=			
XoverBestCP02MR01	6																								=			
XoverRandCP1MR01	7																											
XoverRandCP05MR01	8																											
XoverRandCP02MR01	9																								=			
HillClimbingMR001	10																											
CopyBestMR001	11		=																									
CopyRandMR001	12																											
XoverBestCP1MR001	13																											
XoverBestCP05MR001	14															=												
XoverBestCP02MR001	15														=													
XoverRandCP1MR001	16																									=		
XoverRandCP05MR001	17																											
XoverRandCP02MR001	18																											
HillClimbingMR0001	19																											
CopyBestMR0001	20																											
CopyRandMR0001	21																											
XoverBestCP1MR0001	22																								=			
XoverBestCP05MR0001	23					=	=			=																		
XoverBestCP02MR0001	24																						=					
XoverRandCP1MR0001	25																=											
XoverRandCP05MR0001	26																											
XoverRandCP02MR0001	27																											

Figure A.12: Critical Difference plot based on the Nemenyi post-hoc test ($\alpha = 0.05$) on the algorithm versions tested on the imitation problem (28×28 scenario).

Table A.5: Wilcoxon Rank-sum test ($\alpha = 0.05$) on the pairwise comparisons between the algorithm versions tested on the imitation problem (7×7 scenario). "=" indicates that the null-hypothesis H_0 (statistical equivalence) is accepted (omitted on the diagonal cells). Empty cells indicate that the null-hypothesis is rejected.

		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27
HillClimbingMR01	1											=										=						
CopyBestMR01	2							=															=	=	=			
CopyRandMR01	3											=		=										=	=			
XoverBestCP1MR01	4																											
XoverBestCP05MR01	5						=																					
XoverBestCP02MR01	6				=																							
XoverRandCP1MR01	7																=	=	=									
XoverRandCP05MR01	8																=											
XoverRandCP02MR01	9																											=
HillClimbingMR001	10		=																				=	=	=			
CopyBestMR001	11	=																				=						
CopyRandMR001	12			=										=														
XoverBestCP1MR001	13			=								=												=	=			
XoverBestCP05MR001	14															=												
XoverBestCP02MR001	15														=													
XoverRandCP1MR001	16							=	=									=	=									
XoverRandCP05MR001	17							=									=		=									
XoverRandCP02MR001	18							=									=	=										
HillClimbingMR0001	19																											
CopyBestMR0001	20																											
CopyRandMR0001	21	=										=																
XoverBestCP1MR0001	22		=								=													=	=			
XoverBestCP05MR0001	23		=	=							=			=									=		=			
XoverBestCP02MR0001	24		=	=							=			=									=	=				
XoverRandCP1MR0001	25																										=	=
XoverRandCP05MR0001	26																									=		=
XoverRandCP02MR0001	27									=																=	=	

Figure A.13: Critical Difference plot based on the Nemenyi post-hoc test ($\alpha = 0.05$) on the algorithm versions tested on the imitation problem (7×7 scenario).

Table A.6: Wilcoxon Rank-sum test ($\alpha = 0.05$) on the pairwise comparisons between the algorithm versions tested on the illumination problem (single parameter scenario). "=" indicates that the null-hypothesis H_0 (statistical equivalence) is accepted (omitted on the diagonal cells). Empty cells indicate that the null-hypothesis is rejected.

		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27
HillClimbingMR05	1																											
CopyBestMR05	2																											
CopyRandMR05	3																											
XoverBestCP1MR05	4																											
XoverBestCP05MR05	5																											
XoverBestCP02MR05	6									=																		
XoverRandCP1MR05	7																											
XoverRandCP05MR05	8																											
XoverRandCP02MR05	9						=																					
HillClimbingMR005	10																											
CopyBestMR005	11																					=		=				
CopyRandMR005	12																											
XoverBestCP1MR005	13																						=					
XoverBestCP05MR005	14																											
XoverBestCP02MR005	15																		=									
XoverRandCP1MR005	16																											
XoverRandCP05MR005	17																											
XoverRandCP02MR005	18													=														
HillClimbingMR0005	19																											
CopyBestMR0005	20										=													=				
CopyRandMR0005	21																											
XoverBestCP1MR0005	22												=															
XoverBestCP05MR0005	23										=								=									
XoverBestCP02MR0005	24																											
XoverRandCP1MR0005	25																											
XoverRandCP05MR0005	26																											
XoverRandCP02MR0005	27																											

Figure A.14: Critical Difference plot based on the Nemenyi post-hoc test ($\alpha = 0.05$) on the algorithm versions tested on the illumination problem (single parameter scenario).

Table A.7: Wilcoxon Rank-sum test ($\alpha = 0.05$) on the pairwise comparisons between the algorithm versions tested on the illumination problem (vector representation scenario). "=" indicates that the null-hypothesis H_0 (statistical equivalence) is accepted (omitted on the diagonal cells). Empty cells indicate that the null-hypothesis is rejected.

		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27
HillClimbingMR01	1																											
CopyBestMR01	2																											
CopyRandMR01	3																											
XoverBestCP1MR01	4																											
XoverBestCP05MR01	5																											
XoverBestCP02MR01	6																											
XoverRandCP1MR01	7																											
XoverRandCP05MR01	8																											
XoverRandCP02MR01	9																											
HillClimbingMR001	10																											
CopyBestMR001	11																											
CopyRandMR001	12																									=		
XoverBestCP1MR001	13																						=					
XoverBestCP05MR001	14															=												
XoverBestCP02MR001	15													=														
XoverRandCP1MR001	16																											
XoverRandCP05MR001	17																											
XoverRandCP02MR001	18																											
HillClimbingMR0001	19																											
CopyBestMR0001	20																											
CopyRandMR0001	21																											
XoverBestCP1MR0001	22												=															
XoverBestCP05MR0001	23																											
XoverBestCP02MR0001	24											=																
XoverRandCP1MR0001	25																										=	=
XoverRandCP05MR0001	26																									=		=
XoverRandCP02MR0001	27																									=	=	

Figure A.15: Critical Difference plot based on the Nemenyi post-hoc test ($\alpha = 0.05$) on the algorithm versions tested on the illumination problem (vector representation scenario).

Table A.8: Wilcoxon Rank-sum test ($\alpha = 0.05$) on the pairwise comparisons between the algorithm versions tested on the distributed model (room A, human presence task). "=" indicates that the null-hypothesis H_0 (statistical equivalence) is accepted (omitted on the diagonal cells). Empty cells indicate that the null-hypothesis is rejected.

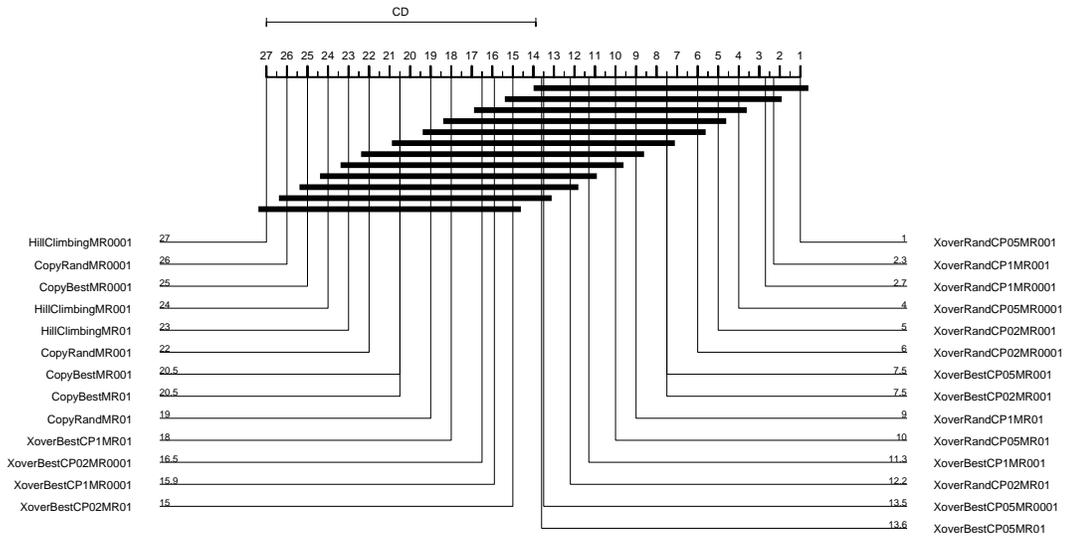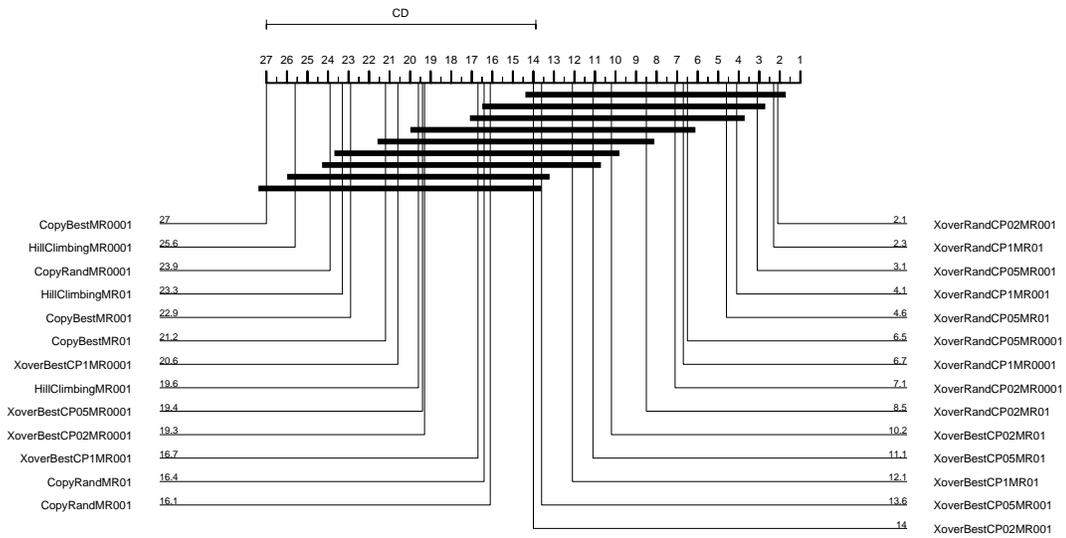|		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27
HillClimbingMR01	1		=								=		=							=		=						
CopyBestMR01	2	=									=	=	=	=	=					=		=	=	=	=			
CopyRandMR01	3													=	=	=	=	=	=				=			=	=	=
XoverBestCP1MR01	4					=	=	=	=	=									=									=
XoverBestCP05MR01	5				=		=	=	=	=									=									
XoverBestCP02MR01	6				=	=		=	=	=							=		=							=		=
XoverRandCP1MR01	7				=	=	=		=	=																		
XoverRandCP05MR01	8				=	=	=	=		=									=									
XoverRandCP02MR01	9				=	=	=		=								=	=	=							=		=
HillClimbingMR001	10	=	=										=							=		=						
CopyBestMR001	11		=										=							=	=	=						
CopyRandMR001	12	=	=								=	=								=		=						
XoverBestCP1MR001	13		=	=											=													
XoverBestCP05MR001	14		=	=										=			=		=				=	=	=		=	
XoverBestCP02MR001	15		=												=		=	=	=							=	=	=
XoverRandCP1MR001	16		=				=			=					=	=		=	=							=	=	=
XoverRandCP05MR001	17		=							=						=	=		=							=	=	=
XoverRandCP02MR001	18			=	=	=	=		=	=					=	=	=	=								=	=	=
HillClimbingMR0001	19	=	=								=	=	=									=						
CopyBestMR0001	20											=																
CopyRandMR0001	21	=	=								=	=	=							=								
XoverBestCP1MR0001	22		=											=	=									=	=			
XoverBestCP05MR0001	23		=	=										=	=								=		=			
XoverBestCP02MR0001	24		=											=	=								=	=				
XoverRandCP1MR0001	25						=			=						=	=	=	=								=	=
XoverRandCP05MR0001	26			=											=	=	=	=	=							=		=
XoverRandCP02MR0001	27			=	=		=			=						=	=	=	=							=	=	

Figure A.16: Critical Difference plot based on the Nemenyi post-hoc test ($\alpha = 0.05$) on the algorithm versions tested on the distributed model (room A, human presence task).

39

Table A.9: Wilcoxon Rank-sum test ($\alpha = 0.05$) on the pairwise comparisons between the algorithm versions tested on the distributed model (room A, human activity task). "=" indicates that the null-hypothesis H_0 (statistical equivalence) is accepted (omitted on the diagonal cells). Empty cells indicate that the null-hypothesis is rejected.

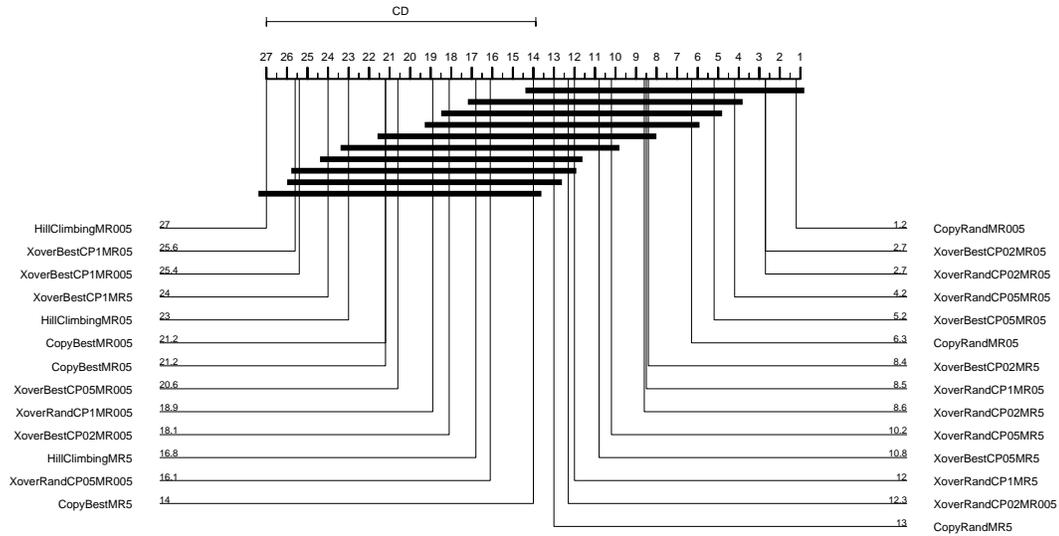|		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27
HillClimbingMR01	1			=								=	=									=						
CopyBestMR01	2			=								=	=								=	=						
CopyRandMR01	3	=	=									=	=								=	=						
XoverBestCP1MR01	4													=	=	=			=				=	=	=			
XoverBestCP05MR01	5						=	=	=	=					=		=	=	=							=	=	=
XoverBestCP02MR01	6				=			=	=	=					=		=	=	=							=	=	=
XoverRandCP1MR01	7				=	=			=	=					=		=	=	=							=	=	=
XoverRandCP05MR01	8				=	=	=			=					=		=	=	=							=	=	=
XoverRandCP02MR01	9				=	=	=	=																		=		
HillClimbingMR001	10																			=								
CopyBestMR001	11	=	=	=									=								=	=						
CopyRandMR001	12	=	=	=								=									=	=						
XoverBestCP1MR001	13				=										=	=	=		=				=	=	=		=	
XoverBestCP05MR001	14				=							=		=		=	=	=	=				=	=	=	=	=	=
XoverBestCP02MR001	15				=									=	=		=	=	=				=	=	=		=	
XoverRandCP1MR001	16					=	=	=	=					=	=	=			=				=	=	=	=	=	=
XoverRandCP05MR001	17					=	=	=	=						=	=			=							=	=	=
XoverRandCP02MR001	18				=	=	=	=						=	=	=	=	=					=	=	=	=	=	=
HillClimbingMR0001	19										=																	
CopyBestMR0001	20		=	=								=	=											=				
CopyRandMR0001	21	=	=	=								=	=							=								
XoverBestCP1MR0001	22				=									=	=	=	=		=					=	=		=	
XoverBestCP05MR0001	23				=									=	=	=	=		=				=		=		=	
XoverBestCP02MR0001	24				=									=	=	=	=		=				=	=			=	
XoverRandCP1MR0001	25					=	=	=	=	=					=		=	=	=								=	=
XoverRandCP05MR0001	26					=	=	=	=					=	=	=	=	=	=				=	=	=	=		=
XoverRandCP02MR0001	27					=	=	=	=						=		=	=	=							=	=	

Figure A.17: Critical Difference plot based on the Nemenyi post-hoc test ($\alpha = 0.05$) on the algorithm versions tested on the distributed model (room A, human activity task).

Table A.10: Wilcoxon Rank-sum test ($\alpha = 0.05$) on the pairwise comparisons between the algorithm versions tested on the distributed model (room B, human presence task). "=" indicates that the null-hypothesis H_0 (statistical equivalence) is accepted (omitted on the diagonal cells). Empty cells indicate that the null-hypothesis is rejected.

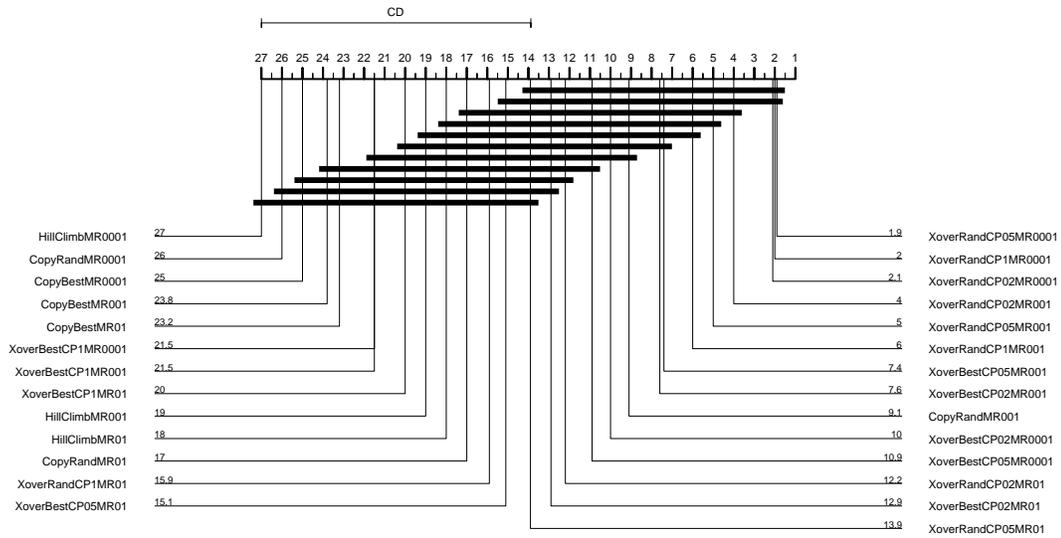|		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27
HillClimbingMR01	1			=	=						=	=		=	=	=	=		=	=			=	=	=		=	=
CopyBestMR01	2			=								=	=							=	=	=						
CopyRandMR01	3	=	=		=						=	=		=	=	=	=		=				=	=	=	=	=	=
XoverBestCP1MR01	4	=		=		=	=	=	=		=	=		=	=	=	=		=				=	=	=	=	=	=
XoverBestCP05MR01	5				=		=	=	=	=				=	=		=	=	=						=	=	=	=
XoverBestCP02MR01	6			=	=	=		=	=	=				=			=	=							=	=	=	=
XoverRandCP1MR01	7			=	=		=		=	=				=	=		=	=							=	=	=	=
XoverRandCP05MR01	8			=	=	=	=	=		=				=			=	=					=	=		=	=	=
XoverRandCP02MR01	9				=	=	=	=	=								=	=										
HillClimbingMR001	10	=		=								=		=	=	=	=		=	=			=	=	=		=	=
CopyBestMR001	11	=	=	=	=						=			=	=		=		=	=	=	=	=	=	=			=
CopyRandMR001	12		=																		=	=						
XoverBestCP1MR001	13	=		=	=						=	=				=	=		=	=			=	=	=		=	=
XoverBestCP05MR001	14	=		=							=	=		=		=	=		=	=			=	=	=		=	=
XoverBestCP02MR001	15	=		=	=	=		=	=		=	=		=	=		=		=	=			=	=	=		=	=
XoverRandCP1MR001	16	=		=	=		=	=			=	=		=	=	=			=	=			=	=	=		=	=
XoverRandCP05MR001	17			=	=	=	=	=	=					=					=	=							=	=
XoverRandCP02MR001	18	=		=	=	=	=	=	=		=	=		=	=	=	=			=			=	=	=		=	=
HillClimbingMR0001	19	=	=	=	=						=	=		=	=	=	=		=		=	=	=	=	=		=	=
CopyBestMR0001	20		=	=							=	=	=							=		=						
CopyRandMR0001	21		=	=							=	=	=							=	=							
XoverBestCP1MR0001	22	=		=	=						=	=		=	=	=	=		=	=				=	=		=	=
XoverBestCP05MR0001	23	=		=	=						=	=		=	=	=	=		=	=			=		=		=	=
XoverBestCP02MR0001	24	=		=	=	=			=		=	=		=	=	=	=		=	=			=	=			=	=
XoverRandCP1MR0001	25			=	=	=	=	=						=		=	=		=				=	=			=	=
XoverRandCP05MR0001	26	=		=	=	=	=	=	=		=			=	=	=	=		=	=			=	=	=	=		=
XoverRandCP02MR0001	27	=		=	=	=	=	=	=		=	=		=	=	=	=		=	=			=	=	=	=	=	

Figure A.18: Critical Difference plot based on the Nemenyi post-hoc test ($\alpha = 0.05$) on the algorithm versions tested on the distributed model (room B, human presence task).

CD

Left-side (ranks):
- CopyRandMR001 — 24.45
- CopyRandMR0001 — 22.45
- CopyBestMR01 — 22.4
- CopyBestMR0001 — 21.95
- CopyBestMR001 — 18.45
- HillClimbingMR0001 — 17.2
- HillClimbingMR01 — 17.05
- XoverBestCP1MR001 — 16.75
- CopyRandMR01 — 16.7
- HillClimbingMR001 — 16.05
- XoverBestCP1MR0001 — 15.8
- XoverBestCP02MR0001 — 15.5
- XoverBestCP05MR001 — 14.35

Right-side (ranks):
- XoverRandCP02MR01 — 4.55
- XoverBestCP02MR01 — 6.7
- XoverBestCP05MR01 — 7.25
- XoverRandCP05MR001 — 7.3
- XoverRandCP05MR01 — 7.4
- XoverRandCP1MR01 — 7.6
- XoverRandCP1MR0001 — 10
- XoverRandCP02MR001 — 10.7
- XoverRandCP05MR0001 — 11.25
- XoverBestCP02MR001 — 11.55
- XoverRandCP1MR001 — 13.25
- XoverRandCP02MR0001 — 13.35
- XoverBestCP1MR01 — 14
- XoverBestCP05MR0001 — 14

41

Table A.11: Wilcoxon Rank-sum test ($\alpha = 0.05$) on the pairwise comparisons between the algorithm versions tested on the distributed model (room B, human activity task). "=" indicates that the null-hypothesis H_0 (statistical equivalence) is accepted (omitted on the diagonal cells). Empty cells indicate that the null-hypothesis is rejected.

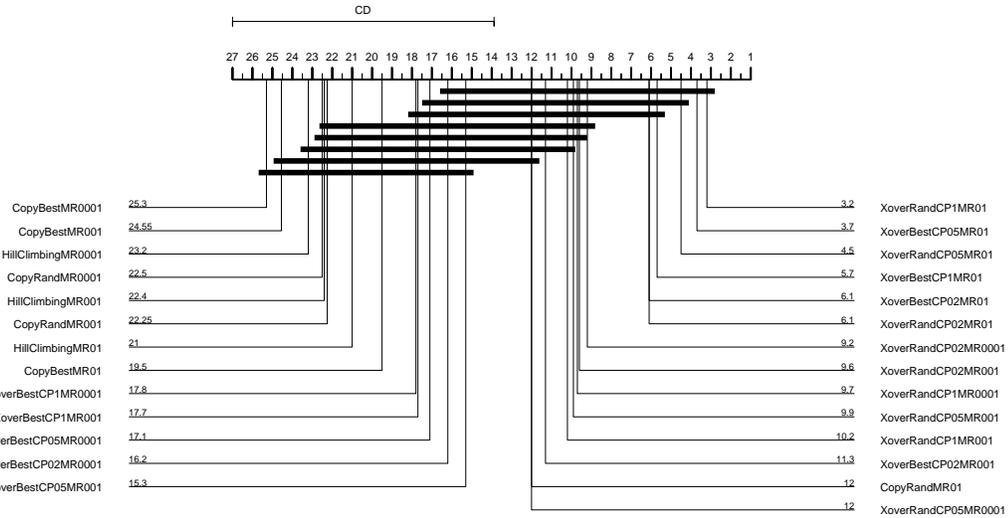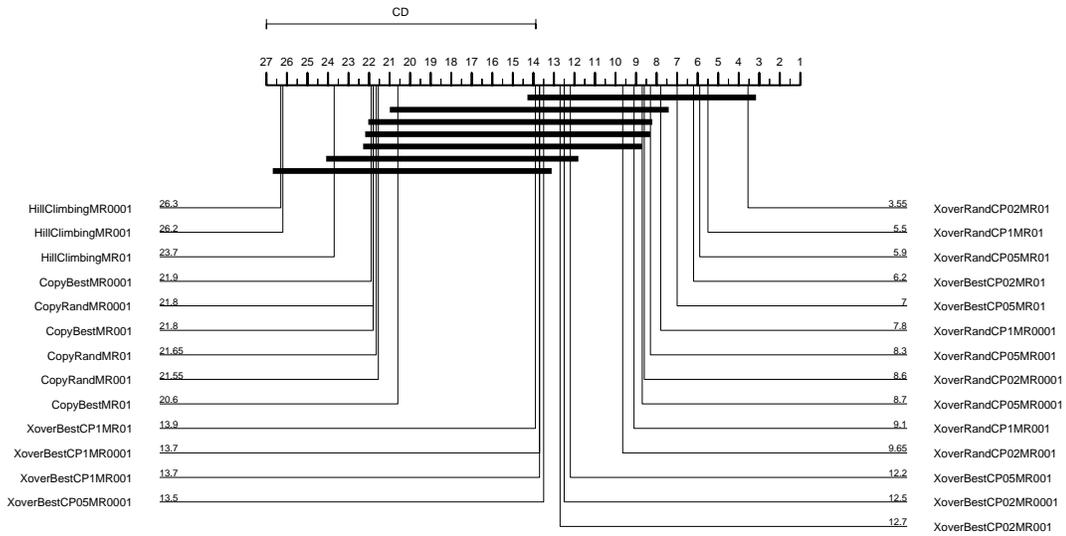|		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27
HillClimbingMR01	1										=									=	=							
CopyBestMR01	2			=	=	=	=	=	=	=			=		=	=	=	=	=		=	=	=	=		=	=	
CopyRandMR01	3		=		=	=	=	=	=	=		=	=		=		=		=		=	=	=	=	=		=	=
XoverBestCP1MR01	4		=	=		=	=	=	=			=	=	=	=	=	=	=	=			=	=	=	=	=	=	=
XoverBestCP05MR01	5		=	=	=		=	=	=			=	=	=	=	=	=	=	=			=	=	=	=	=	=	=
XoverBestCP02MR01	6		=	=	=	=		=	=			=	=	=	=	=	=	=	=		=		=	=	=	=	=	=
XoverRandCP1MR01	7		=	=	=	=	=		=	=		=	=	=	=	=	=	=	=		=		=	=	=	=	=	=
XoverRandCP05MR01	8		=	=	=	=	=	=		=			=	=	=	=	=	=	=		=		=	=	=	=	=	=
XoverRandCP02MR01	9		=	=	=	=		=	=			=	=		=	=	=	=	=		=		=	=	=	=	=	=
HillClimbingMR001	10	=																		=	=							
CopyBestMR001	11		=					=		=			=						=		=	=	=				=	
CopyRandMR001	12		=	=	=		=	=	=	=		=					=		=		=	=		=		=	=	
XoverBestCP1MR001	13			=	=	=									=		=		=									
XoverBestCP05MR001	14			=	=	=	=	=	=					=		=	=	=					=	=	=	=		=
XoverBestCP02MR001	15		=	=	=	=	=	=	=						=		=	=	=				=	=	=	=	=	=
XoverRandCP1MR001	16		=		=	=	=	=	=	=				=	=	=		=	=				=	=	=	=	=	=
XoverRandCP05MR001	17		=	=	=	=	=	=	=					=	=	=	=		=				=	=	=	=	=	=
XoverRandCP02MR001	18		=	=	=	=	=	=	=			=		=		=	=	=					=	=	=	=	=	=
HillClimbingMR0001	19	=									=										=							
CopyBestMR0001	20	=	=	=			=	=	=	=	=	=	=					=	=	=		=		=			=	
CopyRandMR0001	21		=		=		=	=	=			=	=					=	=		=					=	=	
XoverBestCP1MR0001	22		=	=	=	=	=	=	=				=		=	=	=	=	=					=	=	=	=	=
XoverBestCP05MR0001	23		=	=	=	=	=	=	=				=		=	=	=	=	=		=		=		=	=	=	=
XoverBestCP02MR0001	24			=	=	=	=	=	=						=	=	=	=	=				=	=		=		=
XoverRandCP1MR0001	25		=	=	=	=	=	=	=				=	=	=	=	=	=	=				=	=	=		=	=
XoverRandCP05MR0001	26		=	=	=	=	=	=	=			=	=			=	=	=	=		=		=	=		=		=
XoverRandCP02MR0001	27			=	=	=	=	=	=	=				=	=	=	=	=	=				=	=	=	=	=	

Figure A.19: Critical Difference plot based on the Nemenyi post-hoc test ($\alpha = 0.05$) on the algorithm versions tested on the distributed model (room B, human activity task).

CD

HillClimbingMR0001	24.95	
HillClimbingMR01	24.8	
HillClimbingMR001	24.75	
CopyBestMR001	20.05	
CopyBestMR0001	19.7	
CopyRandMR0001	17.15	
CopyRandMR001	16.8	
XoverRandCP05MR0001	16.3	
XoverRandCP02MR001	15.6	
CopyRandMR01	15.55	
XoverRandCP02MR01	14.95	
CopyBestMR01	14.65	
XoverRandCP05MR01	14.05	

5.85	XoverBestCP1MR001
7.65	XoverBestCP02MR0001
8.35	XoverRandCP1MR001
9	XoverRandCP02MR0001
9.2	XoverBestCP05MR001
9.45	XoverBestCP02MR01
9.95	XoverBestCP02MR001
10.4	XoverBestCP1MR0001
10.45	XoverBestCP05MR01
10.65	XoverBestCP05MR0001
10.75	XoverBestCP1MR01
11	XoverRandCP05MR001
12.1	XoverRandCP1MR0001
13.9	XoverRandCP1MR01

Table A.12: Wilcoxon Rank-sum test ($\alpha = 0.05$) on the pairwise comparisons between the algorithm versions tested on the distributed model (room C, human presence task). "=" indicates that the null-hypothesis H_0 (statistical equivalence) is accepted (omitted on the diagonal cells). Empty cells indicate that the null-hypothesis is rejected.

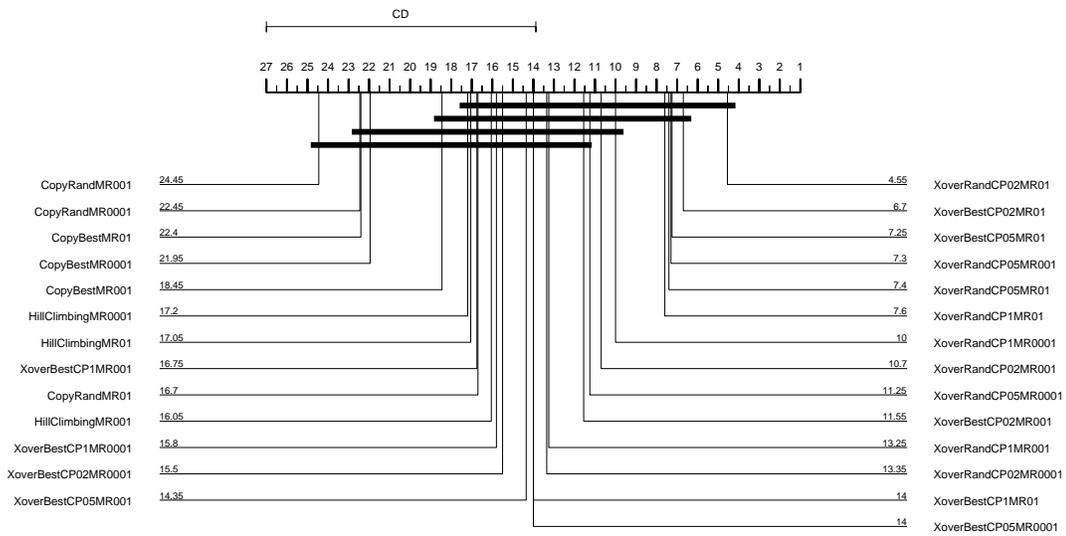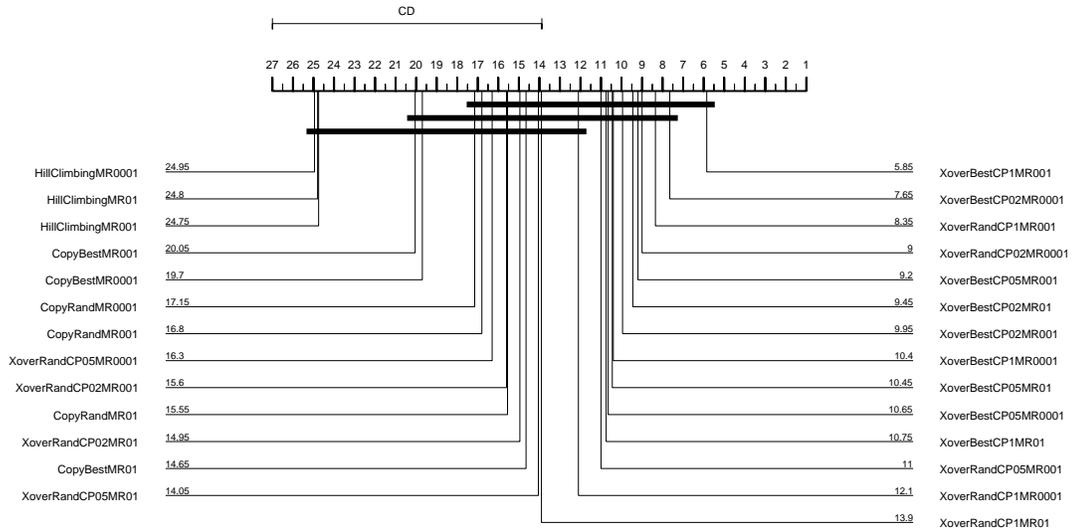|		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27
HillClimbingMR01	1		=	=								=	=									=						
CopyBestMR01	2	=		=								=	=								=	=						
CopyRandMR01	3	=	=									=	=	=	=						=	=	=	=	=			
XoverBestCP1MR01	4					=	=		=	=					=	=							=			=		=
XoverBestCP05MR01	5				=		=		=	=						=	=	=	=							=		=
XoverBestCP02MR01	6				=	=			=	=						=	=	=	=							=		=
XoverRandCP1MR01	7									=						=	=	=	=							=	=	=
XoverRandCP05MR01	8				=	=	=			=						=	=	=	=							=		=
XoverRandCP02MR01	9				=	=	=	=	=							=	=	=	=							=	=	=
HillClimbingMR001	10																			=								
CopyBestMR001	11	=	=	=										=									=	=				
CopyRandMR001	12	=	=	=																			=	=				
XoverBestCP1MR001	13			=								=				=	=						=	=	=			
XoverBestCP05MR001	14			=	=											=							=		=			
XoverBestCP02MR001	15				=	=	=	=	=	=				=	=		=		=							=		=
XoverRandCP1MR001	16					=	=	=	=	=				=		=		=	=							=	=	=
XoverRandCP05MR001	17					=	=	=	=	=						=	=		=							=	=	=
XoverRandCP02MR001	18					=	=	=	=	=						=	=	=								=	=	=
HillClimbingMR0001	19										=										=							
CopyBestMR0001	20		=	=																=								
CopyRandMR0001	21	=	=	=								=	=															
XoverBestCP1MR0001	22			=	=							=	=	=	=										=	=		
XoverBestCP05MR0001	23			=								=	=	=											=			
XoverBestCP02MR0001	24			=										=	=								=	=				
XoverRandCP1MR0001	25				=	=	=	=	=	=						=	=	=	=								=	=
XoverRandCP05MR0001	26							=		=							=	=	=							=		=
XoverRandCP02MR0001	27				=	=	=	=	=	=						=	=	=	=							=	=	

Figure A.20: Critical Difference plot based on the Nemenyi post-hoc test ($\alpha = 0.05$) on the algorithm versions tested on the distributed model (room C, human presence task).

CD

Left-side (rank, algorithm):
- 26 — HillClimbingMR001
- 25.6667 — HillClimbingMR0001
- 23.7778 — CopyBestMR0001
- 22.6111 — CopyRandMR001
- 22.3889 — CopyBestMR001
- 21.9444 — HillClimbingMR01
- 21.6667 — CopyRandMR0001
- 20.6111 — CopyBestMR01
- 19.2778 — CopyRandMR01
- 17.1667 — XoverBestCP05MR0001
- 16.4444 — XoverBestCP02MR0001
- 15.4444 — XoverBestCP1MR0001
- 14.6111 — XoverBestCP1MR001

Right-side (rank, algorithm):
- 3.8889 — XoverRandCP05MR0001
- 3.9444 — XoverRandCP1MR01
- 5.2222 — XoverRandCP05MR001
- 6.6667 — XoverRandCP02MR001
- 6.8889 — XoverRandCP1MR001
- 6.9444 — XoverRandCP02MR01
- 7.6111 — XoverRandCP1MR0001
- 7.6667 — XoverRandCP02MR0001
- 8.2778 — XoverRandCP05MR01
- 8.5 — XoverBestCP05MR01
- 9.8333 — XoverBestCP02MR01
- 10.5556 — XoverBestCP1MR01
- 10.9444 — XoverBestCP02MR001
- 13.4444 — XoverBestCP05MR001

Table A.13: Wilcoxon Rank-sum test ($\alpha = 0.05$) on the pairwise comparisons between the algorithm versions tested on the distributed model (room C, human activity task). "=" indicates that the null-hypothesis H_0 (statistical equivalence) is accepted (omitted on the diagonal cells). Empty cells indicate that the null-hypothesis is rejected.

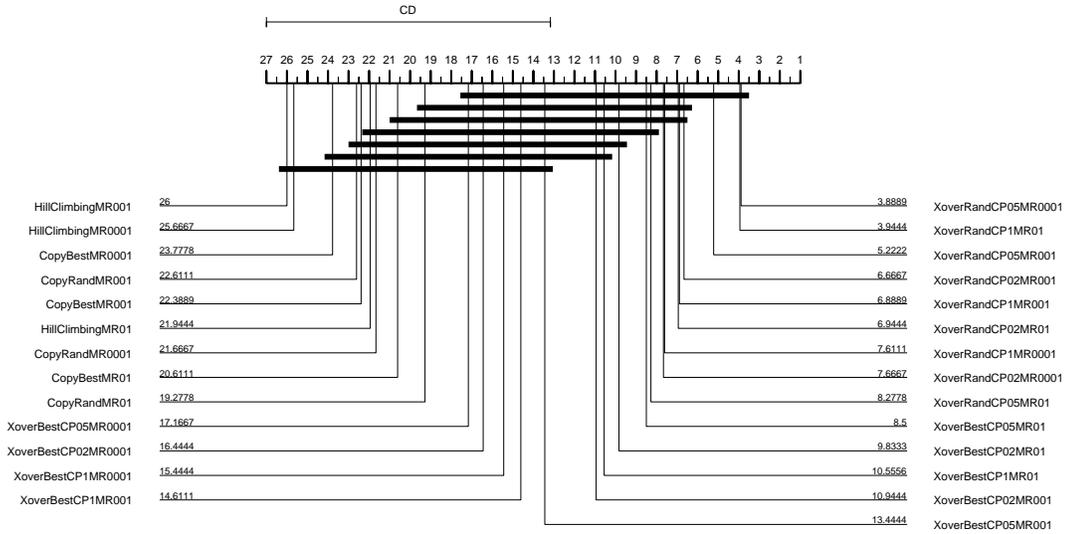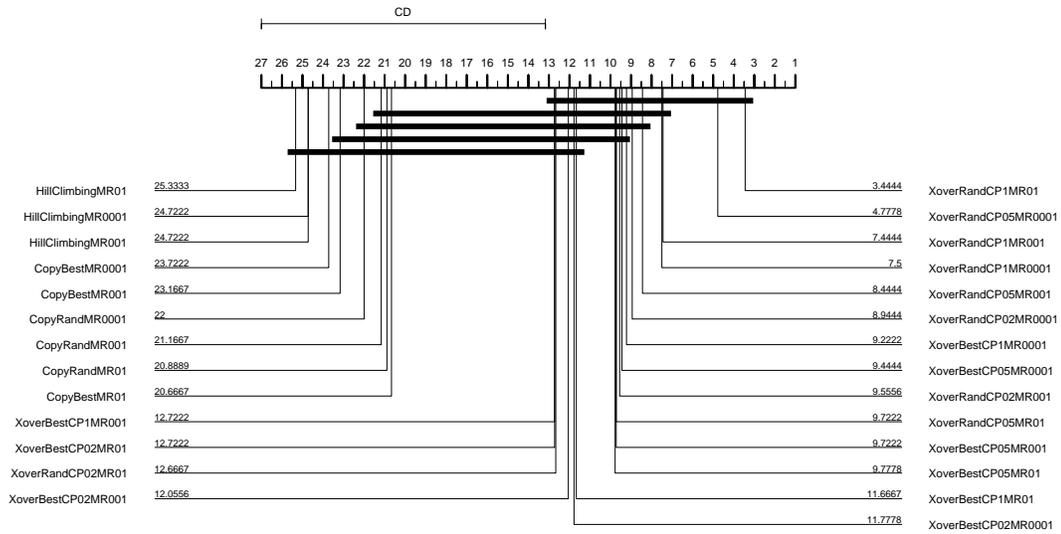|		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27
HillClimbingMR01	1										=		=							=	=	=						
CopyBestMR01	2			=								=	=									=						
CopyRandMR01	3		=									=	=									=						
XoverBestCP1MR01	4					=	=		=	=				=	=	=	=	=	=				=	=	=	=		=
XoverBestCP05MR01	5				=		=	=	=	=				=	=	=	=	=	=				=	=	=	=	=	=
XoverBestCP02MR01	6				=	=			=	=				=	=	=	=	=	=				=	=	=	=		=
XoverRandCP1MR01	7					=													=								=	
XoverRandCP05MR01	8				=	=	=			=				=	=	=	=	=	=				=	=	=	=		=
XoverRandCP02MR01	9				=	=	=		=					=	=	=		=	=				=	=	=			=
HillClimbingMR001	10	=											=							=	=	=						
CopyBestMR001	11	=	=	=										=						=	=	=						
CopyRandMR001	12		=	=										=						=	=							
XoverBestCP1MR001	13				=	=	=		=	=					=		=	=	=				=	=	=	=		=
XoverBestCP05MR001	14				=	=	=		=	=				=		=	=	=	=				=	=	=	=		=
XoverBestCP02MR001	15				=	=	=		=	=				=	=		=	=	=				=	=	=	=		=
XoverRandCP1MR001	16				=	=			=					=	=			=	=				=	=	=	=	=	=
XoverRandCP05MR001	17				=	=	=	=	=	=				=	=	=	=		=				=	=	=	=	=	=
XoverRandCP02MR001	18				=	=	=		=	=				=	=	=	=	=					=	=	=	=	=	=
HillClimbingMR0001	19	=									=	=									=	=						
CopyBestMR0001	20	=									=	=	=							=		=						
CopyRandMR0001	21	=	=	=							=	=	=							=	=							
XoverBestCP1MR0001	22				=	=	=		=	=				=	=	=	=	=	=					=	=	=	=	=
XoverBestCP05MR0001	23				=	=	=		=	=				=	=	=	=	=	=						=	=		=
XoverBestCP02MR0001	24				=	=	=		=	=				=	=	=	=	=	=					=		=		=
XoverRandCP1MR0001	25				=	=	=		=					=	=		=	=	=				=	=	=			=
XoverRandCP05MR0001	26				=	=		=									=	=	=				=					
XoverRandCP02MR0001	27				=	=	=		=	=				=	=	=	=						=	=	=	=		

Figure A.21: Critical Difference plot based on the Nemenyi post-hoc test ($\alpha = 0.05$) on the algorithm versions tested on the distributed model (room C, human activity task).