**ORIGINAL RESEARCH**

# AutoMoDe-`Cedrata`: Automatic Design of Behavior Trees for Controlling a Swarm of Robots with Communication Capabilities

Jonas Kuckling[1] · Vincent van Pelt[1] · Mauro Birattari[1]

## Abstract

Behavior trees are a control architecture that has gained recent attention in AI and robotics. Previous research on the use of behavior trees in swarm robotics has shown the necessity for the behaviors to have proper return values, instead of running indefinitely. This work extends our previous work in which we defined AutoMoDe-`Cedrata`, an automatic modular design that makes use of modules that have been explicitly defined for behavior trees. While the search space is sufficiently large to include well-performing solutions, `Cedrata` had problems discovering communication-based strategies. In this work, we extend `Cedrata` by introducing `Cedrata-GP` and `Cedrata-GE` which are based on genetic programming and grammatical evolution, respectively. We test these design methods on two missions and compare the performance of the automatic design methods against the performance of solutions created by human designers. The results show that the structure of `Cedrata` allows for well-performing solutions that are reliably found by human designers. However, the automatic design methods fail to discover the same communication strategies as the human designers.

**Keywords** Swarm robotics · Design by optimization · AutoMoDe · Genetic programming · Grammatical evolution

## Introduction

Swarm robotics is a research area that combines robotics and swarm intelligence, and that is recognized as a promising approach for controlling large groups of robots [1–11]. Robot swarms are self-organizing decentralized systems, consisting of relatively simple robots that cooperate to achieve a goal that would not be achievable for each individual robot alone. The collective behavior of the robot swarm emerges from the interactions between the robots themselves and between the robots and the environment [12]. One challenge of swarm robotics is the difficulty of designing control software for the individual robots, so that the desired collective behavior emerges [13].

✉ Jonas Kuckling
jonas.kuckling@ulb.be

Mauro Birattari
mbiro@ulb.ac.be

1  IRIDIA, Université libre de Bruxelles, Brussels, Belgium

One approach to the design of control software for robot swarms is manual design, in which a human designer creates the control software. However, only few and limited principled approaches to manual design exist [14–22] and no general methodology has yet been proposed. As a result, most manual design approaches rely on trial and error, a time-consuming, costly, and often error-prone strategy [23, 24].

Other approaches rely on the use of an optimization algorithm and can be broadly categorized into two categories: semi-automatic design and fully automatic design (although hybrid approaches exist) [10]. In semi-automatic design, a human designer uses an optimization algorithm as a tool to design the control software. The designer specifies the problem and defines the parameters of the optimization algorithm. They observe the optimization process and adjust the problem specification or the parameters of the optimization algorithm until the result is satisfactory. While the semi-automatic approach alleviates some drawbacks of manual design, the involvement of a human designer still entails similar challenges: as long as no general principled approach exists, much of the performance depends on the experience and domain knowledge of the human designer.

In contrast, in fully automatic design, the role of the human designer is reduced to the problem specification. After receiving the problem specification, the fully automatic design process searches for a satisfactory solution without any further human intervention [10]. This lack of human intervention also implies that no mission-specific domain knowledge can be incorporated into the design process. Indeed, any fully automatic design method needs to be able to address not only a single mission, but a class of missions [13].

Fully automatic design often produces the control software off-line, i.e., the software is designed using simulations and only the final resulting control software is uploaded onto the real robots for evaluation. While this approach offers many advantages, like speeding up the design process through faster-than-real-time simulations and parallelization of simulation processes and no need for hardware availability for the design process, it suffers from one major drawback, the reality gap. The reality gap is the inherent difference between the simulation and the real environment and often manifests itself in the form of a performance drop [25]. Not all methods are affected equally by the reality gap [25].

Francesca et al. proposed to look at the reality gap problem akin to the bias-variance trade-off [26]. They hypothesized that design methods with a very large and fine-grained action space ("low bias") are more prone to overfit the simulation context ("high variance"). By restricting the space of possible behaviors ("introducing bias"), it should be possible to produce software that is more robust to the reality gap. Based on this hypothesis, they proposed Auto-MoDe, a class of automatic modular design methods. In automatic modular design, a set of pre-defined modules is assembled and fine-tuned into more complex control software by an optimization algorithm. The first method of this class is `Vanilla`, an automatic modular design approach that crosses the reality gap satisfactorily [26]. `Chocolate` extends `Vanilla` by using Iterated F-race [27] as the optimization algorithm to assemble a finite-state machine with up to four states and sixteen transitions from a set of six behavioral modules (mapped to the states of the finite-state machine) and six conditions (mapped to the transitions of the finite-state machine) [28]. Other AutoMoDe methods vary or extend the capabilities of `Chocolate`. `Gianduja` [29], `TuttiFrutti` [30], or `Arlequin` [31] introduce new software modules, that extend the capabilities of the robotic platform, e.g., by enabling direct communication, color detection or the use of artificial pheromones. `Waffle` [32] allows the design process not only to control aspects of the control software but also of the hardware capabilities of the robot. `IcePop` [33] investigates the use of local search-based optimization algorithms. `Maple` [34] and `Cedrata` [35] are design methods that use behavior trees [36] as the target architecture.

The work presented in this paper extends [35], which introduced `Cedrata`. Our previous work showed that the modules and architecture of `Cedrata` allowed for well-performing solutions, but the optimization algorithm (Iterated F-race) had problems finding these solutions. In this work, we investigate additionally the use of two other optimization algorithms, namely genetic programming [37] and grammatical evolution [38]. In Table 1 we list the abbreviations used in this paper.

## Related Work

Behavior trees are a control architecture that originates from video games [39], but which since has found applications in fields such as artificial intelligence or robotics [40]. In this work, we follow the behavior tree definition of Marzinotto et al. [36].

In this framework, behavior trees are a control architecture whose structure can be described as a directed acyclic graph and that operate on a tick that is created with a fixed frequency $f_{tick}$ by an implicitly defined root node. Every time a tick is generated, it traverses the tree, activating the nodes that it visits. The inner nodes of the tree are called control-flow nodes and control the way that the tick takes through the tree. The leaf nodes can be either an action node that executes a single time step of a behavior or a condition node that checks a condition of the environment.

**Table 1** Abbreviations and symbols used in this paper

| Abbreviation | Meaning |
| --- | --- |
| ? | Selector node |
| → | Sequence node |
| ?* | Selector* node |
| →* | Sequence* node |
| GP | Genetic programming |
| GE | Grammatical evolution |
| Exp | Exploration behavior |
| Stop | Stop behavior |
| Group | Grouping behavior |
| Isol | Isolation behavior |
| Meet | Meeting behavior |
| Ack | Acknowledgement behavior |
| ESig | Emit Signal behavior |
| Bflr | Black Floor condition |
| Gflr | Grey Floor condition |
| Wflr | White Floor condition |
| Ngb | Neighborhood Count condition |
| INgb | Inverted Neighborhood Count condition |
| FP | Fixed Probability condition |
| RSig | Receiving Signal condition |

After activating, each node in a behavior tree returns the tick to its parent along with one of three possible return values (*success*, *failure*, *running*) that determine the further way that the tick takes through the tree. Condition nodes can return *success* or *failure*, depending on whether their associated condition is met. Action nodes usually return *running*, *success* or *failure*, if the action takes longer than one time step or if the robot is in a state that can be classified as success or failure with regard to the associated action or behavior. The control-flow nodes receive the return value of one of their children nodes and determine if they either pass the tick to another child or to its parent, together with a return value. In this work, we consider the following control-flow nodes: selector (?), sequence (→), selector* (?*), sequence* (→*). For a detailed and formal definition of all nodes, see [36].

Ligot et al. [34] have proposed `Maple`, an automatic modular design method that assembles modules into a restricted behavior tree architecture. `Maple` utilizes the same modules as `Chocolate` [28]. As these modules have originally been conceived for finite-state machine, they do not offer any states that could be characterized as *success* or *failure* and can therefore only return *running*. The authors propose a very restricted architecture that can successfully incorporate these modules. Results showed that `Maple` could produce solutions that performed adequately, but that the architecture limited the space of possible solutions and that some well-performing solutions found by `Chocolate` could not be represented within the restricted architecture of `Maple`. This work highlighted the need for modules that have proper return values. Communication-based behaviors have been used in many works in swarm robotics, for example [41–47]. With `Gianduja`, Hasselmann et al. [29] were able to show that an automatic design process can automatically assign a semantic to messages that do not have meaning a priori.

Other works that have applied behavior trees to swarm robots include the works by Jones et al. [48, 49]. In Ref. [48], Jones et al. evolved behavior trees for a swarm of kilobots.

In that work, the authors defined only atomic actions that are executed for a single tick and then always return *success*. This necessitates the inclusion of a *repeat* node that can repeatedly tick its children, allowing these actions to be executed more than once consecutively. In another work, Jones et al. evolved behavior trees onboard of a swarm of Xpucks [49]. The actions in these behavior trees are also atomic, as they perform a singular write to the blackboard. For the onboard evolution, however, each robot performs a design process using genetic programming to create behavior trees. At regular intervals, the best performing behavior tree is selected as the current control software for the robot. Another approach to the design of control software in the context of swarm robotics was proposed by Neupane and Goodrich [50]. They used distributed grammatical evolution to design behavior trees for a swarm of simulated robots.

GESwarm is another design approach that uses grammatical evolution [51]. In that work, the authors proposed an automatic design method that can design a foraging behavior for a swarm of footbot robots. The control software is represented as a set of policies that map conditions and the current behavior to actions of the robot. The behaviors that the robot can select exhibit similar properties as the behaviors of `Chocolate`, namely that they are simple and can run potentially endlessly, without any implicit success or failure states.

## AutoMoDe-`Cedrata`

Originally, we presented `Cedrata` in Ref. [35]. For the convenience of the reader, we describe again the method here.

### Reference Model

The reference model RM2.2, on which `Cedrata` is based, is shown in Table 2 [52]. The robot is equipped with eight proximity sensors, three ground sensors and one range-and-bearing board for sensing and two sets of actuators: the range-and-bearing board to send messages and two wheels with differential drive. The reference model formalizes the way that the control software has access to these sensors and actuators. The proximity sensors can detect obstacles up to 30cm away, the ground sensors can sense the floor color on a grey scale and the range-and-bearing board can transmit messages up to 50cm. The control software can set the speed of the two wheels of the robot independently. It also always sends a signal value *s*, that can be equal to 0, which is a special value that means *no signal* and that is sent by default, or an integer in $\{1, ..., 6\}$. Similar to Ref. [29], signal values do not have a particular semantic, instead, it is the role of

**Table 2** The E-puck reference model RM2.2 used in `Cedrata` [52]

| Sensors | Variables |
| --- | --- |
| Proximity | $prox \in [1, 8]$, $\angle q \in [0, 2\pi]$ |
| Ground | $gnd \in \{0, 0.5, 1\}$ |
| Range-and-bearing | $n \in \mathbb{N}$, $r \in [0.5, 20]$, $\angle b \in [0, 2\pi]$ |
| | $n_s, r_s, \angle b_s$, for $s \in \{1, ..., 6\}$ |
| Actuators | Variables |
| Signal broadcast | $s \in \{0, 1, ..., 6\}$ |
| Wheels | $v_l, v_r \in [-v, v]$, with $v = 0.16$m/s |

Control cycle period: 100ms

the design process to assign semantics to the signals. For the sensors, the reference model provides an aggregated vector (in the form of magnitude and direction) over all proximity readings and a single aggregated ground reading. The reference model also provides access to the number of neighboring robots $n$ and a vector to their center of mass. Similarly, it provides the number of messaging robots and a vector to the center of mass of the messaging robots, for each signal $s \in S$. The control cycle period is 100ms, that is, every 100ms the sensors are updated and the control software is invoked, generating a new tick in the behavior tree.

## Modules

Based on the reference model RM2.2, we defined fourteen modules—seven behavior modules and seven condition modules (see Table 3). In the following descriptions of the signal-based conditions and behaviors, the set of signals $\{1, ..., 6\}$ will be denoted $S$. Some modules can use a special value *any* that is activated if any of the signals in $S$ is received. The set $S^* = S \cup \{any\}$ will denote the sets used by these modules. The design process is free to choose several instances of the same module in an instance of control software and can tune the parameters independently for each instance of a module.

Behaviors are associated to action nodes and allow the robot to interact with the environment. The action nodes can return *success* or *failure* if the behavior ends in a state that it considers being a success or a failure. Otherwise, they return *running*. The behavior modules are defined as follows:

**Table 3** Behavior and condition modules and their parameters used in Cedrata

| Behavior | Short | Parameters |
|---|---|---|
| Exploration | Exp | $\tau$ |
| Stop | Stop | |
| Grouping | Group | $N_{max}, N_{min}, \alpha$ |
| Isolation | Isol | $N_{max}, N_{min}, \alpha$ |
| Meeting | Meet | $s, d_{min}$ |
| Acknowledgement | Ack | $s, t_{max}$ |
| Emit Signal | ESig | $s$ |
| Condition | Short | Parameters |
| Black Floor | Bflr | $\beta$ |
| Grey Floor | Gflr | $\beta$ |
| White Floor | Wflr | $\beta$ |
| Neighborhood Count | Ngb | $\eta, \xi$ |
| Inverted Neighborhood Count | INgb | $\eta, \xi$ |
| Fixed probability | FP | $\beta$ |
| Receiving signal | RSig | $s$ |

Exploration    The robot performs a random walk. It moves straight until it perceives an obstacle in front of itself. Then the robot turns on the spot for a random number of time steps in $\{0, ..., \tau\}$, where $\tau \in \{1, ..., 100\}$ is a tunable parameter. This behavior always returns *running*.

Stop    The robot stays still. This behavior always returns *running*.

Grouping    The robot tries to get closer to its neighbors by moving towards the geometric center of its neighbors. If the number of neighbors becomes greater than $N_{max}$, the behavior returns *success*, where $N_{max}$ is a tunable parameter. If the number of neighbors becomes smaller than $N_{min}$, the behavior returns *failure*, where $N_{min}$ is a tunable parameter. Otherwise, it returns *running*. The speed of convergence is controlled by the tunable parameter $\alpha \in [1, 5]$. The robot moves in the direction $w = w' - kw_0$, where $w'$ is the target component and $kw_0$ is the obstacle avoidance component. If robots are perceived, then $w' = w_{r\&b} = (\alpha \cdot r, \angle b)$, otherwise $w' = (1, \angle 0)$. $kw_0$ is the obstacle avoidance component, with $k$ being a constant fixed to 5 and $w_0$ defined as $w_0 = (prox, \angle q)$.

Isolation The robot tries to move away from its neighbors by moving in the opposite direction of the geometric center of its neighbors. If the number of neighbors becomes smaller than $N_{min}$, the behavior returns *success*, where $N_{min}$ is a tunable parameter. If the number of neighbors becomes greater than $N_{max}$, the behavior returns *failure*, where $N_{max}$ is a tunable parameter. Otherwise, it returns *running*. The speed of divergence is controlled by the tunable parameter $\alpha \in [1, 5]$. The Isolation behavior uses the same embedded collision avoidance as in Grouping, but with $w'$ defined as: $w' = -w_{r\&b}$ if robots are perceived, where $w_{r\&b}$ is defined as in the Grouping behavior. Otherwise $w' = (1, \angle 0)$.

Meeting The robot listens for a signal $s \in S^*$ emitted by other robots and moves towards the geometrical centre of the emitters. The behavior returns *success* if the distance between the robot and the geometrical centre is smaller than a distance $d_{min}$, where $d_{min}$ is a tunable parameter. The behavior returns *failure* if the robot does not perceive any robot sending the expected signal. Otherwise, the behavior returns *running*. The Meeting behavior uses the same embedded collision avoidance as in Grouping, but with $w'$ defined as: $w' = w_{r\&b} = (\alpha \cdot r_s, \angle b_s)$ if robots emitting $s$ are perceived. Otherwise $w' = (1, \angle 0)$.

Acknowledgement    The robot sends a signal $s \in S$ and waits for an answer in the form of the same signal, where $s$ is a tunable parameter. The behavior returns *success* if the signal is received or *running* if not. After $t_{max}$ ticks, the behavior returns *failure* if the signal is

still not received, where $t_{max}$ is a tunable parameter. This behavior also sets the velocity of both wheels to zero.

Emit Signal       The robot sets its emitted signal to $s \in S$ for the current tick, where $s$ is a tunable parameter. This behavior always returns *success*. This behavior also sets the wheel velocity to zero.

Conditions are associated to condition nodes and check an aspect of the environment. The condition nodes return *success*, when their condition is met, or *failure*, otherwise. The condition modules are defined as follows:

Black Floor       When all ground sensors detect a black floor, the condition returns *success* with probability $\beta$, where $\beta$ is a tunable parameter.

Grey Floor       When all ground sensors detect a grey floor, the condition returns *success* with probability $\beta$, where $\beta$ is a tunable parameter.

White Floor       When all ground sensors detect a white floor, the transition is enabled with probability $\beta$, where $\beta$ is a tunable parameter.

Neighborhood Count       Returns *success* with probability $z(n) = \frac{1}{1+e^{\eta(\xi-n)}}$ where $n$ is the number of robots in the neighborhood, $\eta \in [0, 20]$ and $\xi \in \{0, 1, ..., 10\}$ are tunable parameters.

Inverted Neighborhood Count       Same as Neighborhood Count but with probability $1 - z(n)$.

Fixed Probability Returns *success* with probability $\beta$, where $\beta$ is a tunable parameter.

Receiving Signal Returns *success* if the robot has perceived a neighbor sending $s \in S^*$ in the last 10 ticks, where $s$ is a tunable parameter.

## Architecture

In Cedrata, the optimization process can create a tree that has a maximum of three levels and a maximum of three children per node. The top-level node must be a control-flow node. Nodes of the second level can be control-flow nodes, action nodes or condition nodes. If it is an action node or a condition node, then the node can have no children itself. Not all branches are forced to have the same depth: the top-level node could have some children that are control-flow nodes and some that are action or condition nodes. Nodes on the third level can only be action nodes or condition nodes. The structure of such trees is depicted in Fig. 1. The optimization process can choose any control-flow node type to be either a sequence, sequence*, selector or selector* node. For a formal definition of these nodes, see Marzinotto et al. [36]. Prior research has shown that high complexity in automatic design methods can increase

the difficulties in crossing the reality gap [11, 26, 28]. To match the complexity of other AutoMoDe methods [53], the tree may have at most four action nodes and four condition nodes. The constraints on the depth and on the number of children implicitly impose that the tree contains no more than four control nodes.

## Optimization Algorithm

The optimization algorithm of Cedrata is Iterated F-race [27]. Iterated F-race works over several iterations, each reminiscent of a race [54]. In each iteration, Iterated F-race samples a set of candidate solutions. The first iteration samples randomly from all possible candidate solutions, subsequent iterations sample around the survivors of the preceding one. In the context of Cedrata, these candidate solutions are representations of behavior trees according to the constraints described in the previous sections. These candidate solutions are evaluated incrementally over an increasing number of instances. In the case of Cedrata, each instance is equivalent to different (random) starting positions and orientations of the robots in the mission. However, all candidate solutions that are evaluated on the same instance will be provided with the same starting positions and orientations. If at one point a candidate solution is statistically worse than another one (determined by a Friedman test), it is discarded. By discarding inferior solutions, Iterated F-race frees up the design budget for more promising solutions. The iteration ends if either the allocated budget for this iteration is exhausted or all but a fixed number of candidate solutions are discarded. The following iteration then samples its set of candidate solutions around the elites of the previous iteration and continues the race until the remaining budget is too small to conduct another iteration.

Cedrata-GP and Cedrata-GE use the same reference model, modules and architecture as Cedrata. They differ only in the optimization algorithm employed. Cedrata-GP uses genetic programming [37] as the



**Fig. 1** The possible behavior tree structure for Cedrata. In Cedrata, the top-level node can be any control-flow node. Underneath it the tree can have between one and three nodes, chosen among control-flow nodes, action nodes and condition nodes. If a control-flow node is chosen, then it can have between one and three children, which are either action nodes or condition nodes

optimization algorithm. The parameters of this design method are those used in the work of Jones et al. [48] and summarized in Table 4. We use the genetic programming implementation of the DEAP library [55]. `Cedrata-GE` uses grammatical evolution [38] as the optimization algorithm. The parameters of this design method are those used in the work of Neupane and Goodrich [50] and summarized in Table 4. We use the grammatical evolution implementation of PonyGE2 [56].

## Design Process

In the context of fully automatic design [10, 13], the design process generates control software without any human intervention (besides the mission specification). Once the mission and the experimental protocol are specified, the design process generates the control software for the robots using simulations to determine the performance of candidate solutions. The design process is free to choose the structure, the modules and the parameters of the modules within the constraints described in sections "Modules" and "Architecture".

One parameter of the experimental protocol is the design budget. The design budget poses an upper limit on the number of simulations that the design process can run before the final instance of control software is returned. It serves a similar role as limiting the computation time available to the design process, while being independent of the computational hardware.

## Experimental Setup

In this work, we test `Cedrata` and related design methods on a set of two missions. The experimental setup is equivalent to the one described in Ref. [35], but we describe it again for the convenience of the reader. All code and data is available from the supplementary material [57].

### Missions

We consider two missions: MARKER AGGREGATION and STOP. These missions must be performed in a dodecagonal arena (see Fig. 2) and last 250s.

In the mission MARKER AGGREGATION (see Fig. 2a), the robots must aggregate within the dotted area. The area itself is not perceivable by the robots. Instead, a black spot is placed in the middle of the aggregation area that can serve as a marker. The objective function for this mission is the cumulative time that the robots spend within the aggregation area: $F_{MA} = \sum_{i=0}^{2500} N_A^i$, where $N_A^i$ is the number of robots in the aggregation area at time step $i$. The higher the score of the objective function, the better the robots perform the mission.

In the mission STOP (see Fig. 2b), the robots must find a white spot and then stop as soon as possible. A robot is considered moving, if it has travelled more than 5mm in the last time step. The objective function for this mission is reduced for each robot that is not moving at any given time step before the white spot has been found and for each robot that is moving after the white spot has been found and additionally for the time that the swarm needed to discover the white spot: $F_{S}top = 100000 - \left( \bar{t}N + \sum_{t=1}^{\bar{t}} \sum_{i=1}^{N} \bar{I}_i(t) + \sum_{\bar{t}}^{2500} \sum_{i=1}^{N} I_i(t) \right)$,

**Table 4** Parameters for genetic programming and grammatical evolution

| Parameter | Genetic programming | Grammatical evolution |
|---|---|---|
| Initialization | Half-and-half | Uniform_tree |
| Selection strategy | Tournament selection | Truncation |
| Tournament size | 3 | – |
| Selection proportion | – | 50% |
| Crossover | One-point crossover | One-point crossover |
| Population size | 25 | 100 |
| Number of elites | 3 | 1 |
| Crossover probability | 0.8 | 0.9 |
| Uniform mutation probability | 0.05 | – |
| Shrink mutation probability | 0.1 | – |
| Node replacement mutation probability | 0.5 | – |
| Ephemeral mutation probability | 0.5 | – |
| Flip per codon mutation probability | – | 0.01 |
| Codon size | – | 1000 |

Parameters for genetic programming are those used in the work of Jones et al. [48] and parameters for grammatical evolution are those used in the work of Neupane and Goodrich [50]
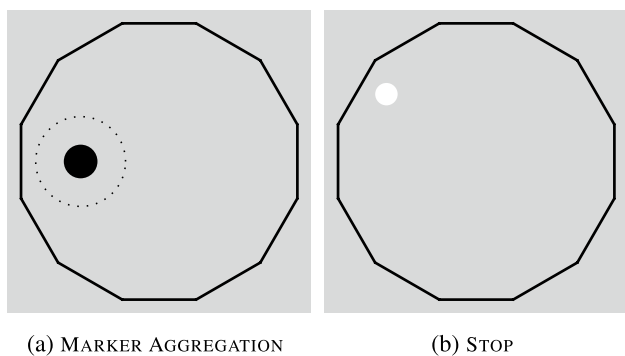
(a) MARKER AGGREGATION              (b) STOP

**Fig. 2** Layouts of the arena for the missions considered

where $\bar{t}$ is the time step during which the white spot was discovered, $I_i(t)$ is an indicator that a robot $i$ has moved in time step $t$ and $\bar{I}_i(t)$ is an indicator that a robot $i$ has not moved in time step $t$. The higher the score of the objective function, the better the robots perform the mission.

## Design Methods

We consider `Cedrata`, as described in section "Auto-MoDe-`Cedrata`". As `Cedrata` had problems in producing communication-based strategies for the mission considered, we performed experiments with additional design methods: `Cedrata-GP` and `Cedrata-GE`. We also performed several manual designs. For the manual designs, we asked human designers—with prior experience in swarm robotics, but not with behavior trees—to design control software within the same constraints as `Cedrata`, that is, with the same modules and architecture. The human designers had access to the AutoMoDe Editor [58], a tool that allows the designers to visualize and manipulate the behavior trees and to launch simulations of the designed behavior tree. The human designers received feedback about their designed behavior tree through the objective function and a visual representation of the arena and the behavior of the swarm.

Lastly, we include a reference design as an additional point of reference for the reader. These reference designs are not part of the experimental protocol and have been designed by us. They are not optimized and do not aim to be the best performing solutions for each mission, but simply to provide a sensible solution. These designs serve to highlight particular strategies that we expected to be discovered in each mission. They were not known to the human designers prior to their manual designs.

## Reference Designs

The reference design for the mission MARKER AGGREGATION is shown in Fig. 3a. In this design, robots explore the arena

until they find the marker. Then, using the signal framework, they will attract their neighbors to the aggregation area. At any given time step, the tick traverses the three subtrees from left to right. The left subtree handles the case where the robot is on the marker. If the condition Black Floor evaluates true, then the tick is passed on to the action node, which invokes the Emit Signal behavior. Since Emit Signal always returns *success* and the action node is the last child of the sequence node, this subtree then returns *success* as well. This will cause the selector node to also return *success*. If the condition Black Floor is not met, then the tick is passed into the middle subtree, which handles the case where the robot is on the grey floor and perceives at least one signaling neighbor. Here, if the condition Grey Floor is met, the robot executes on time step of the Meeting behavior. If Meeting returns *success* or *running*, then the tick will leave the tree. If either Meeting or Grey Floor return *failure*, then the tick is passed to the last subtree. This subtree only consists of an action node with the Exploration behavior.

The reference design for the mission STOP is shown in Fig. 3b. In this design, robots will send and forward signals to their neighbors to transmit the information that the white spot has been discovered. If a robot receives a signal, it stops; if it does not receive any signal, it explores the arena to find the white spot. At any given time step, the tick traverses the three subtrees from left to right. The left subtree handles the case in which the robot is on the white spot. While the condition White Floor evaluates true, the robot executes the behavior Emit Signal to signal the other robots the discovery of the spot. If the condition White Floor was not met, then the tick is passed to the middle subtree that forwards received signals. If the condition Receiving Signal is met, then the tick will be passed to the Emit Signal behavior that emits the same signal as is checked for in the Receiving Signal condition. If the Receiving Signal condition is not met, then the tick is passed to the right subtree, which consists only of an action node with the Exploration behavior.



(a) MARKER AGGREGATION              (b) STOP

**Fig. 3** The reference designs for the two missions. The conditions and actions names have been abbreviated in the following way: Exp: Exploration; Meet: Meeting; ESig: Emit Signal; Bflr: Black Floor; Gflr: Grey Floor; Wflr: White Floor; RSig: Receiving Signal

## Protocol

For each mission, `Cedrata` is executed with different budgets: 20,000, 50,000, 100,000 and 200,000 simulation runs. The budget specifies the number of simulations that the design process is allowed to perform before it returns the best control software produced. Additionally, `Cedrata-GP` and `Cedrata-GE` are tested on a budget of 200,000 simulation runs. For each combination of method, mission and budget, 10 independent runs of the methods are performed, leading to 10 instances of control software. The manual designs are done by four human designers per mission, with a maximum design duration of 4 h.

Simulations are performed in a realistic and physics-based simulation environment, based on the ARGoS simulator [59]. The simulated robots have a real world counter-part, and the simulator has been used in the past with this robotic platform and comparisons between simulated performance and real-world performance have been made. In accordance with the consensus in the literature, a realistic noise model is applied to the simulation (see Table 5). The generated instances of control software of all designs methods are assessed in pseudo-reality to investigate the impact of the reality gap. Ligot and Birattari [25] have shown that the effect of the reality gap can be mimicked in simulation-only environments by testing the control software with a different noise model than it was originally designed for.

## Results

In this section, we describe the results obtained by the experimental setup described in section "Experimental Setup". In the supplementary material [57], we also include an extended study, where we include another method that is based on another reference model which provides us with some additional insights tangential to the work presented here.

Figure 4 shows the results for the missions STOP and MARKER AGGREGATION. Results are shown for both the performance in simulation and pseudo-reality. Each box in the box plot represents the performances of the final instances of control software generated by the independent runs of the design methods.

Figure 4a shows the development of the performance of `Cedrata` in the mission MARKER AGGREGATION. There is a clear trend of increasing performance with increasing budget. A detailed investigation of the generated control software reveals that `Cedrata` develops two general solution strategies: one strategy is based on the communication framework, while the other is not. In the communication-less strategy (for an example, see Fig. 5a), the robots explore the arena until they discover the black spot, at which point

**Table 5** Design and pseudo-reality noise models

| Sensor/actuator | Design model | Pseudo-reality model |
|---|---|---|
| Proximity | 0.05 | 0.05 |
| Light | 0.05 | 0.90 |
| Ground | 0.05 | 0.05 |
| Range-and-bearing | 0.85 | 0.90 |
| Wheels | 0.05 | 0.15 |

they usually stop. In the communication-based strategy (see Fig. 5b), however, the robots make use of the communication behaviors to quickly aggregate within the target area. The communication-based designs are similar in that regard to the reference design. The performance of `Cedrata` for each budget then seems to primarily depend on the ratio of the two strategies. Indeed, for design budgets of 20,000 and 50,000 simulation runs, `Cedrata` only produces control software that uses the communication-less strategy. For a budget of 100,000 simulation runs, `Cedrata` produces a single solution that follows the communication-based strategy and for a budget of 200,000 simulation runs, four designs make use of that strategy. It appears that the ratio of communication-less to communication-based strategies depends on the available budget. Indeed, as the communication-based strategy requires at least two modules to interact correctly, Iterated F-race is more likely to discover such a combination the more often it samples new solutions, which depends on the number of iterations and therefore the budget.

In Fig. 4b, we can see the comparison of performances across all considered design methods. All manual designers found solutions that make use of communication. Their control software performs similar well as the communication-based behavior trees generated by `Cedrata` and better than the reference design, which was not meant to be the best performing solution, but just to highlight the general strategy. The human designers were therefore not only able to discover the strategy but also to find a reasonable tuning for the parameter.

`Cedrata-GP` and `Cedrata-GE` both fail to generate any solution making use of the communication modules, even for a budget of 200,000 simulation runs. Interestingly, both design methods generate solutions that, under the right circumstances, perform nearly as good as the best instances of control software generated by `Cedrata`. However, this appears to be mostly due to the initial starting position favoring quick aggregation within the target zone and in total both `Cedrata-GP` and `Cedrata-GE` perform worse than `Cedrata`.

**Fig. 4** Results for the mission MARKER AGGREGATION (top) and STOP (bottom). The left plots show the development of the performance over increasing budget for Cedrata. The right plots show the comparison of all design methods under consideration for a budget of 200,000 simulation runs. The thin plots present the results in simulation, the thick plots the results in pseudo-reality
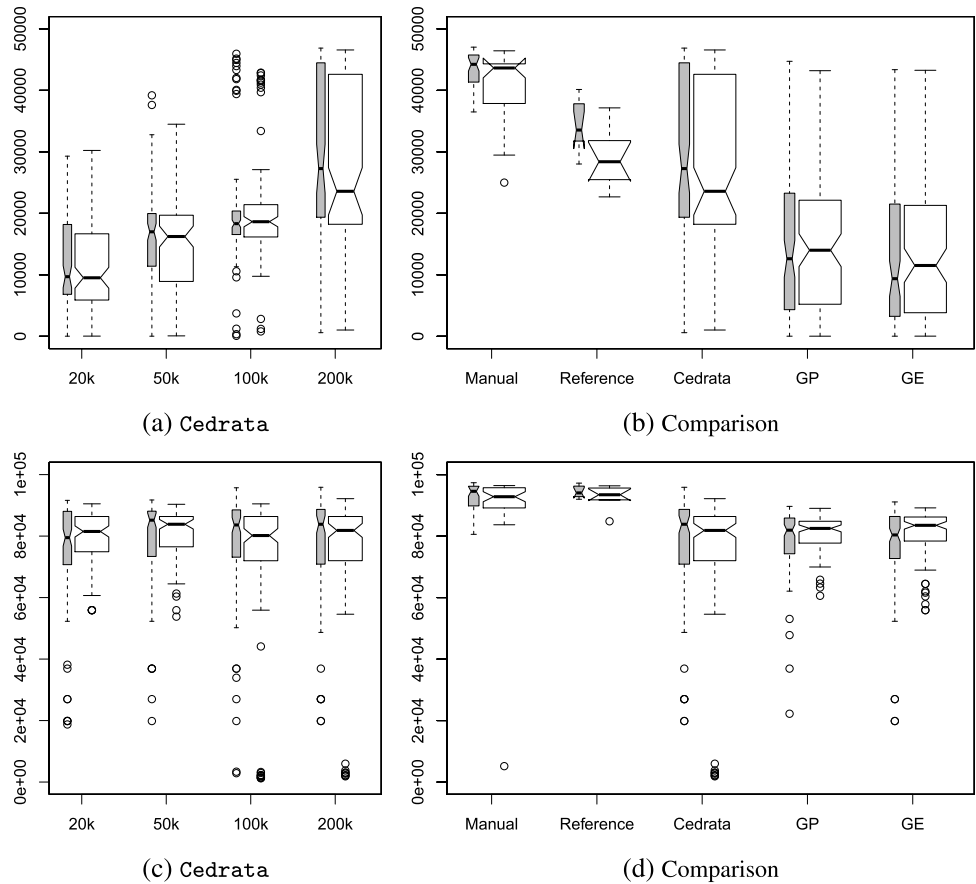
(a) Cedrata

(b) Comparison

(c) Cedrata

(d) Comparison

**Fig. 5** Typical behavior trees generated by Cedrata

(a) MARKER AGGREGATION, without communication

(b) MARKER AGGREGATION, communication-based

(c) STOP
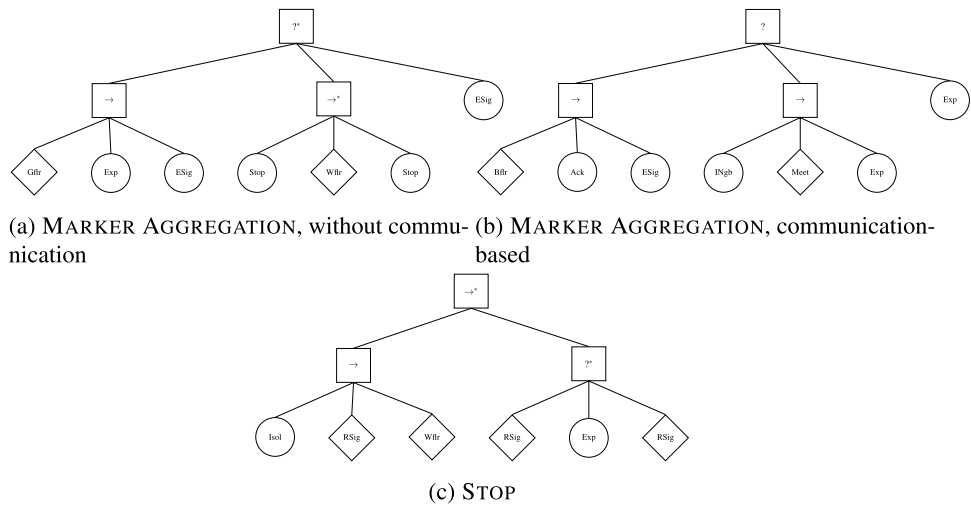
Figure 4c shows the development of performance over budget for Cedrata in the mission STOP. Unlike in the mission MARKER AGGREGATION, there is no improvement for increasing budgets. Instead, the performance remains relatively stable. Investigation of the generated behavior trees reveals that Cedrata fails to make use of the communication modules for this mission. All generated behavior trees employ a strategy, where the robots are using the Isolation behavior (for an example, see Fig. 5c). As a result, the swarm expands and, with high probability, a robot passes over the white spot. At the end of the expansion phase, the robots slow down and move relatively little, often falling below the threshold of 5mm per time step. Some behavior trees also include an Exploration module for cases when no neighbors are detected.

Figure [4]d displays a comparison of the performances of all design methods in the mission STOP. The manual designs, just like the reference design, make use of the communication framework and show the best performance. Both `Cedrata-GP` and `Cedrata-GE` find solutions that follow the same Isolation-based strategy as `Cedrata` and achieve similar performances. For all design methods, there are some runs where the performance is relatively close to 0. Often, in these runs, the control software fails to find the white spot.

We made some observations that hold for all considered missions: The first observation is that all design methods show a relatively small pseudo-reality gap. That is, they experience only a small drop in performance when assessing the control software in pseudo-reality. We believe that this is a first indicator that `Cedrata` and the design methods based on it might transfer well into reality as well. A second observation is that all behavior trees generated by `Cedrata`, `Cedrata-GP` and `Cedrata-GE` contain many modules that will never be ticked by the behavior tree. We believe this to be because of the reduced restrictions in the architecture, which allow modules to be easily placed in the tree in a way that ensures they will never receive a tick. The design process has no explicit way of distinguishing necessary and superfluous modules and all techniques that aim at generating new behavior trees (random sampling around elites, cross-overs, mutations) are therefore highly likely to transfer some of the superfluous modules into the newly generated behavior tree. This poses a challenge to the automatic design process. Namely, that the design process will spend some resources on tuning these superfluous modules, which have no influence on the behavior of the swarm, thus effectively wasting a part of the allocated budget. Lastly, we observed that the automatic design process had difficulties generating communication-based behaviors. In both missions, MARKER AGGREGATION and STOP, the human designers found well performing solutions that made use of the communication framework. Only in the mission MARKER AGGREGATION was `Cedrata` able to generate at least a few solutions following a similar strategy. Our initial hypothesis was that this might have been caused by some properties of the underlying optimization algorithm, Iterated F-race. We have therefore replaced Iterated F-race with two different optimization algorithms, whose parametrization we have taken from other works in the swarm robotics literature. Unfortunately, both `Cedrata-GP` and `Cedrata-GE` appeared to have even greater difficulties generating communication-based behaviors than `Cedrata`. We believe that this could be due to the fact that communication requires two corresponding modules, a sender and a receiver, while all other strategies can rely on a single module.

## Conclusion

In this work, we have extended AutoMoDe-`Cedrata`, by implementing two variants `Cedrata-GP` and `Cedrata-GE`, based on genetic programming and grammatical evolution. We have investigated the performance of these automatic design methods over a set of two missions and compared them to solutions found by human designers, following the same constraints. The results generated by the human designers show that the modules and constraints of `Cedrata` are sensible, as the human designers were able to design control software that performed satisfactorily. Furthermore, as the human designers had no prior experience with behavior trees, this seems to be an indicator that behavior trees are an intuitive control architecture to design for. The automatic design method `Cedrata`, on the other hand, was not able to generate communication-based behaviors. We hypothesized that this might have been due to some property of the optimization algorithm Iterated F-race, and therefore we created `Cedrata-GP` and `Cedrata-GE`, two variants of `Cedrata` that are based on genetic programming and grammatical evolution, respectively. Neither of these two variants was able to generate communication-based strategies either.

For future work, we would like to investigate in more detail how an automatic design process can discover meaningful communication-based strategies and why the approach taken in this work failed. The results of this work indicate that simply tuning the parameters of an optimization algorithm would probably not be enough. Nevertheless it would be interesting to investigate the effects of different parameters on the performance of generated solutions, especially with respect to the exploration-exploitation trade-off. Another issue for investigation could be the mapping of behavior trees into representations that can be manipulated by the genetic programming and grammatical evolution implementations. One possible approach to create communication-based behaviors could be to create an interleaved optimization process. Starting from a minimal communicating solution, we alternate between fixing the sending or the receiving part of the behavior tree and optimizing the remaining part of the tree. Another approach to solve this problem could be cooperative co-evolution. We could possibly create two distinct populations that are given a sending or receiving module, respectively. This ensures the existence of communication from the starting population. Subsequent generations could then refine the communication protocol and integrate it with the other modules. Additionally, the results presented here showed that `Cedrata` and its variants were able to perform satisfactorily also in pseudo-reality. While this is an indicator that the design approaches

might cross the reality gap well, we would like to confirm this hypothesis by performing real robot experiments.

**Author contributions** The experiments were designed and performed by JK and VP. The paper was drafted by JK and edited by MB; all authors read and commented the final version. The research was directed by MB.

## Declarations

**Conflict of interest** The authors declare that they have no conflict of interest.

## References

1. Rubenstein M, Cornejo A, Nagpal R. Programmable self-assembly in a thousand-robot swarm. Science. 2014;345(6198):795–9. https://doi.org/10.1126/science.1254295.
2. Werfel J, Petersen K, Nagpal R. Designing collective behavior in a termite-inspired robot construction team. Science. 2014;343(6172):754–8. https://doi.org/10.1126/science.1245842.
3. Yang GZ, Bellingham J, Dupont PE, Fischer P, Floridi L, Full R, Jacobstein N, Kumar V, McNutt M, Merrifield R, Nelson BJ, Scassellati B, Taddeo M, Taylor R, Veloso M, Wang ZL, Wood R. The grand challenges of science robotics. Sci Robot. 2018;3(14):eaar7650. https://doi.org/10.1126/scirobotics.aar7650.
4. Garattoni L, Birattari M. Autonomous task sequencing in a robot swarm. Sci Robot. 2018;3(20):eaat0430. https://doi.org/10.1126/scirobotics.aat0430.
5. Slavkov I, Carrillo-Zapata D, Carranza N, Diego X, Jansson F, Kaandorp J, Hauert S, Sharpe J. Morphogenesis in robot swarms. Sci Robot. 2018;3(25):eaau9178. https://doi.org/10.1126/scirobotics.aau9178.
6. Yu J, Wang B, Du X, Wang Q, Zhang L. Ultra-extensible ribbon-like magnetic microswarm. Nat Commun. 2018;9(1):3260. https://doi.org/10.1038/s41467-018-05749-6.
7. Li S, Batra R, Brown D, Chang HD, Ranganathan N, Hoberman C, Rus D, Lipson H. Particle robotics based on statistical mechanics of loosely coupled components. Nature. 2019;567(7748):361–5. https://doi.org/10.1038/s41586-019-1022-9.
8. Xie H, Sun M, Fan X, Lin Z, Chen W, Wang L, Dong L, He Q. Reconfigurable magnetic microrobot swarm: multimode transformation, locomotion, and manipulation. Sci Robot. 2019;4(28):eaav8006. https://doi.org/10.1126/scirobotics.aav8006.
9. Dorigo M, Theraulaz G, Trianni V. Reflections on the future of swarm robotics. Sci Robot. 2020;5:eabe4385. https://doi.org/10.1126/scirobotics.abe4385.
10. Birattari M, Ligot A, Hasselmann K. Disentangling automatic and semi-automatic approaches to the optimization-based design of control software for robot swarms. Nat Mach Intell. 2020;2(9):494–9. https://doi.org/10.1038/s42256-020-0215-0.
11. Hasselmann K, Ligot A, Ruddick J, Birattari M. Empirical assessment and comparison of neuro-evolutionary methods for the automatic off-line design of robot swarms. Nat Commun. 2021;12:4345. https://doi.org/10.1038/s41467-021-24642-3.
12. Dorigo M, Birattari M, Brambilla M. Swarm robotics. Scholarpedia. 2014;9(1):1463. https://doi.org/10.4249/scholarpedia.1463.
13. Birattari M, Ligot A, Bozhinoski D, Brambilla M, Francesca G, Garattoni L, Garzón Ramos D, Hasselmann K, Kegeleirs M, Kuckling J, Pagnozzi F, Roli A, Salman M, Stützle T. Automatic off-line design of robot swarms: a manifesto. Front Robot AI. 2019;6:59. https://doi.org/10.3389/frobt.2019.00059.
14. Hamann H, Wörn H. A framework of space-time continuous models for algorithm design in swarm robotics. Swarm Intell. 2008;2(2–4):209–39. https://doi.org/10.1007/s11721-008-0015-3.
15. Kazadi S. Model independence in swarm robotics. Int J Intell Comput Cybern. 2009;2(4):672–94. https://doi.org/10.1108/17563780911005836.
16. Berman S, Kumar V, Nagpal R. Design of control policies for spatially inhomogeneous robot swarms with application to commercial pollination. In: 2011 IEEE international conference on robotics and automation (ICRA). Piscataway: IEEE; 2011. pp. 378–385. https://doi.org/10.1109/ICRA.2011.5980440
17. Beal J, Dulman S, Usbeck K, Viroli M, Correll N. Organizing the aggregate: languages for spatial computing. In: Marjan M, editor. Formal and practical aspects of domain-specific languages: recent developments. Hershey: IGI Global; 2012. pp. 436–501. https://doi.org/10.4018/978-1-4666-2092-6.ch016
18. Brambilla M, Brutschy A, Dorigo M, Birattari M. Property-driven design for swarm robotics: a design method based on prescriptive modeling and model checking. ACM Trans Auton Adapt Syst. 2014;9(4):17:1-17:28. https://doi.org/10.1145/2700318.
19. Reina A, Valentini G, Fernández-Oto C, Dorigo M, Trianni V. A design pattern for decentralised decision making. PLOS ONE. 2015;10(10): e0140950. https://doi.org/10.1371/journal.pone.0140950.
20. Lopes YK, Trenkwalder SM, Leal AB, Dodd TJ, Groß R. Supervisory control theory applied to swarm robotics. Swarm Intell. 2016;10(1):65–97. https://doi.org/10.1007/s11721-016-0119-0.
21. Pinciroli C, Beltrame G. Buzz: a programming language for robot swarms. IEEE Softw. 2016;33(4):97–100. https://doi.org/10.1109/MS.2016.95.
22. Hamann H. Swarm robotics: a formal approach. Cham, Switzerland: Springer; 2018. https://doi.org/10.1007/978-3-319-74528-2.
23. Brambilla M, Ferrante E, Birattari M, Dorigo M. Swarm robotics: a review from the swarm engineering perspective. Swarm Intell. 2013;7(1):1–41. https://doi.org/10.1007/s11721-012-0075-2.
24. Francesca G, Birattari M. Automatic design of robot swarms: achievements and challenges. Front Robot AI. 2016;3(29):1–9. https://doi.org/10.3389/frobt.2016.00029.
25. Ligot A, Birattari M. Simulation-only experiments to mimic the effects of the reality gap in the automatic design of robot swarms. Swarm Intell. 2019. https://doi.org/10.1007/s11721-019-00175-w.
26. Francesca G, Brambilla M, Brutschy A, Trianni V, Birattari M. AutoMoDe: a novel approach to the automatic design of control software for robot swarms. Swarm Intell. 2014;8(2):89–112. https://doi.org/10.1007/s11721-014-0092-4.
27. López-Ibáñez M, Dubois-Lacoste J, Pérez Cáceres L, Birattari M, Stützle T. The irace package: iterated racing for automatic algorithm configuration. Oper Res Perspect. 2016;3:43–58. https://doi.org/10.1016/j.orp.2016.09.002.
28. Francesca G, Brambilla M, Brutschy A, Garattoni L, Miletitch R, Podevijn G, Reina A, Soleymani T, Salvaro M, Pinciroli C, Mascia F, Trianni V, Birattari M. AutoMoDe-Chocolate: automatic design

of control software for robot swarms. Swarm Intell. 2015;9(2–3):125–52. https://doi.org/10.1007/s11721-015-0107-9.

29. Hasselmann K, Birattari M. Modular automatic design of collective behaviors for robots endowed with local communication capabilities. PeerJ Comput Sci. 2020. https://doi.org/10.7717/peerj-cs.291.

30. Garzón Ramos D, Birattari M. Automatic design of collective behaviors for robots that can display and perceive colors. Appl Sci. 2020;10(13):4654. https://doi.org/10.3390/app10134654.

31. Ligot A, Hasselmann K, Birattari M. AutoMoDe-Arlequin: neural networks as behavioral modules for the automatic design of probabilistic finite state machines. In: Dorigo M, Stützle T, Blesa MJ, Blum C, Hamann H, Heinrich MK, Strobel V, editors. Swarm intelligence: 12th international conference, ANTS 2020, Lecture Notes in Computer Science, vol. 12421. Cham: Springer; 2020. pp. 109–122. https://doi.org/10.1007/978-3-030-60376-2_21

32. Salman M, Ligot A, Birattari M. Concurrent design of control software and configuration of hardware for robot swarms under economic constraints. PeerJ Comput Sci. 2019;5:e221. https://doi.org/10.7717/peerj-cs.221.

33. Kuckling J, Ubeda Arriaza K, Birattari M. AutoMoDe-IcePop: automatic modular design of control software for robot swarms using simulated annealing. In: Bogaerts B, Bontempi G, Geurts P, Harley N, Lebichot B, Lenaerts T, Louppe G, editors. Artificial Intelligence and Machine Learning: BNAIC 2019, BENELEARN 2019, Communications in Computer and Information Science, vol. 1196. Cham, Switzerland: Springer; 2020. p. 3–17.

34. Ligot A, Kuckling J, Bozhinoski D, Birattari M. Automatic modular design of robot swarms using behavior trees as a control architecture. PeerJ Comput Sci. 2020;6:e314. https://doi.org/10.7717/peerj-cs.314.

35. Kuckling J, van Pelt V, Birattari M. Automatic modular design of behavior trees for robot swarms with communication capabilities. In: Castillo PA, Jiménez Laredo JL, editors. Applications of evolutionary computation: 24th international conference, EvoApplications 2021, Lecture Notes in Computer Science, vol. 12694. Cham: Springer; 2021. pp. 130–145.

36. Marzinotto A, Colledanchise M, Smith C, Ögren P. Towards a unified behavior trees framework for robot control. In: 2014 IEEE international conference on robotics and automation (ICRA), Piscataway: IEEE; 2014. pp. 5420–5427. https://doi.org/10.1109/ICRA.2014.6907656

37. Koza JR. Genetic programming: on the programming of computers by means of natural selection, first edn. MIT Press, Cambridge, MA, USA. 1992. A Bradford Book

38. O'Neill M, Ryan C. Grammatical evolution: evolutionary automatic programming in an arbitrary language, 1st ed. Genetic programming series. Boston: Springer; 2003. https://doi.org/10.1007/978-1-4615-0447-4

39. Isla D. Handling complexity in the Halo 2 AI. In: Game developers conference, GDC 2005, vol. 12. London: Game Developers Conference (GDC). 2005.

40. Colledanchise M, Ögren P. Behavior trees in robotics and AI: an introduction, 1st ed. In: Chapman & Hall/CRC artificial intelligence and robotics series. Boca Raton: CRC Press; 2018. https://doi.org/10.1201/9780429489105

41. Nolfi S, Floreano D. Evolutionary robotics: the biology, intelligence, and technology of self-organizing machines, 1st ed. A Bradford Book. Cambridge: MIT Press. 2000.

42. Trianni V, Labella Thomas H, Dorigo M. Evolution of direct communication for a swarm-bot performing hole avoidance. In: Dorigo M, Birattari M, Blum C, Gambardella LM, Mondada F, Stützle T, editors. Ant colony optimization and swarm intelligence: 4th international workshop, ANTS 2004, Lecture Notes in Computer Science, vol. 3172. Berlin: Springer; 2004. pp. 130–141. https://doi.org/10.1007/978-3-540-28646-2_12

43. Jones C, Matarić MJ. Automatic synthesis of communication-based coordinated multi-robot systems. In: 2004 IEEE/RSJ international conference on intelligent robots and systems (IROS), vol. 1. Piscataway: IEEE; 2004. pp. 381–387. https://doi.org/10.1109/IROS.2004.1389382

44. Wischmann S, Pasemann F. The emergence of communication by evolving dynamical systems. In: Nolfi S, Baldassarre G, Calabretta R, Hallam J, Marocco D, Meyer JA, Miglino O, Parisi D, editors. From animals to animats 9: 9th international conference on simulation of adaptive behavior, SAB 2006, Lecture Notes in Computer Science, vol. 4095. Berlin: Springer; 2006. pp. 777–788. https://doi.org/10.1007/11840541_64

45. Marocco D, Nolfi S. Self-organization of communication in evolving robots. In: Rocha LM, Yaeger LS, Bedau MA, Floreano D, Goldstone RL, Vespignani A, editors. Artificial life X: proceedings of the tenth international conference on the simulation and synthesis of living systems, complex adaptive systems. Cambridge: MIT Press; 2006.

46. Wischmann S, Floreano D, Keller L. Historical contingency affects signaling strategies and competitive abilities in evolving populations of simulated robots. Proc Natl Acad Sci USA. 2012;109(3):864–8. https://doi.org/10.1073/pnas.1104267109.

47. Uno R, Marocco D, Nolfi S, Ikegami T. Emergence of protosentences in artificial communicating systems. IEEE Trans Auton Mental Dev. 2011;3(2):146–53. https://doi.org/10.1109/TAMD.2011.2120608.

48. Jones S, Studley M, Hauert S, Winfield A. Evolving behaviour trees for swarm robotics. In: Groß R, Kolling A, Berman S, Frazzoli E, Martinoli A, Matsuno F, Gauci M, editors. Distributed autonomous robotic systems: the 13th international symposium, Springer Proceedings in Advanced Robotics, vol. 6. Cham; Springer; 2018. pp. 487–501. https://doi.org/10.1007/978-3-319-73008-0_34

49. Jones S, Winfield A, Hauert S, Studley M. Onboard evolution of understandable swarm behaviors. Adv Intell Syst. 2019;1(6):1900031. https://doi.org/10.1002/aisy.201900031.

50. Neupane A, Goodrich M. Learning swarm behaviors using grammatical evolution and behavior trees. In: Kraus S, editor. Proceedings of the twenty-eighth international joint conference on artificial intelligence, IJCAI-19. CA, USA; IJCAI Organization; 2019; pp. 513–520. https://doi.org/10.24963/ijcai.2019/73

51. Ferrante E, Duéñez-Guzmán EA, Turgut AE, Wenseleers T. GESwarm: grammatical evolution for the automatic synthesis of collective behaviors in swarm robotics. In: Blum C, editor. GECCO'13: proceedings of the 15th annual conference on genetic and evolutionary computation. New York: ACM; 2013. pp. 17–24. https://doi.org/10.1145/2463372.2463385

52. Hasselmann K, Ligot A, Francesca G, Garzón Ramos D, Salman M, Kuckling J, Mendiburu FJ, Birattari M. Reference models for AutoMoDe. Tech. Rep. TR/IRIDIA/2018-002, IRIDIA, Brussels. 2018.

53. Kuckling J, Ligot A, Bozhinoski D, Birattari M. Search space for AutoMoDe-Chocolate and AutoMoDe-Maple. Tech. Rep. TR/IRIDIA/2018-012, IRIDIA, Brussels. 2018.

54. Maron O, Moore AW. The Racing Algorithm: model selection for lazy learners. Artif Intell Rev. 1997;11(1–5):193–225. https://doi.org/10.1023/A:1006556606079.

55. Fortin FA, De Rainville FM, Gardner MA, Parizeau M, Gagné C. DEAP: evolutionary algorithms made easy. J Mach Learn Res. 2012;13:2171–5.

56. Fenton M, McDermott J, Fagan D, Forstenlechner S, Hemberg E, O'Neill M. PonyGE2: grammatical evolution in Python. arxiv:1703.08535. 2017.

57. Kuckling J, van Pelt V, Birattari M. AutoMoDe-Cedrata: automatic design of behavior trees for controlling a swarm of robots

with communication capabilities: supplementary material. http://iridia.ulb.ac.be/supp/IridiaSupp2021-004/. 2021.

58. Kuckling J, Hasselmann K, van Pelt V, Kiere C, Birattari M. Auto-MoDe Editor: a visualization tool for AutoMoDe. Tech. Rep. TR/IRIDIA/2021-009, IRIDIA, Brussels. 2021.

59. Pinciroli C, Trianni V, O'Grady R, Pini G, Brutschy A, Brambilla M, Mathews N, Ferrante E, Di Caro GA, Ducatelle F, Birattari M, Gambardella LM, Dorigo M. ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems. Swarm Intell. 2012;6(4):271–95. https://doi.org/10.1007/s11721-012-0072-5.