

1 **Evolution of Self-Organized Task Specialization in Robot Swarms**

2 Eliseo Ferrante¹, Ali Emre Turgut³, Edgar Duéñez-Guzmán¹, Marco Dorigo² & Tom
3 Wenseleers¹

4

5 ¹ Laboratory of Socio-Ecology and Social Evolution, Zoological Institute, KULeuven,
6 Naamsestraat 59, B-3000, Leuven, Belgium

7 ² IRIDIA – CoDE, Université Libre de Bruxelles, 50, Av. F. Roosevelt, CP 194/6
8 B-1050 Brussels, Belgium

9 ³ Mechanical Engineering Department, Middle East Technical University, 06800, Ankara,
10 Turkey

11

12 * To whom correspondence should be addressed. Current address: Laboratory of Socio-
13 ecology and Social Evolution, Zoological Institute, KULeuven, Naamsestraat 59, B-3000,
14 Leuven, Belgium, Phone: (+32) 16 3 23964, email: eliseo.ferrante@bio.kuleuven.be

15

16 Running head: Evolution of Task Specialization in Robot Swarms

17

18 **Abstract**

19 Division of labor is ubiquitous in biological systems, as evidenced by various forms of
20 complex task specialization observed in both animal societies and multicellular organisms.
21 Although clearly adaptive, the way in which division of labor first evolved remains enigmatic,
22 as it requires the simultaneous co-occurrence of several complex traits to achieve the
23 required degree of coordination. Recently, evolutionary swarm robotics has emerged as an
24 excellent test bed to study the evolution of coordinated group-level behavior. Here we use this
25 framework for the first time to study the evolutionary origin of behavioral task specialization
26 among groups of identical robots. The scenario we study involves an advanced form of
27 division of labor, common in insect societies and known as “task partitioning”, whereby two
28 sets of tasks have to be carried out in sequence by different individuals. Our results show that
29 task partitioning is favored whenever the environment has features that, when exploited,
30 reduce switching costs and increase the net efficiency of the group, and that an optimal mix of
31 task specialists is achieved most readily when the behavioral repertoires aimed at carrying
32 out the different subtasks are available as pre-adapted building blocks. Nevertheless, we also
33 show for the first time that self-organized task specialization could be evolved entirely from
34 scratch, starting only from basic, low-level behavioral primitives, using a nature-inspired
35 evolutionary method known as Grammatical Evolution. Remarkably, division of labor was
36 achieved merely by selecting on overall group performance, and without providing any prior
37 information on how the global object retrieval task was best divided into smaller subtasks. We
38 discuss the potential of our method for engineering adaptively behaving robot swarms and
39 interpret our results in relation to the likely path that nature took to evolve complex sociality
40 and task specialization.

41

42 **Author summary**

43 Many biological systems execute tasks by dividing them into finer sub-tasks first. This is seen
44 for example in the advanced division of labor of social insects like ants, bees or termites. One
45 of the unsolved mysteries in biology is how a blind process of Darwinian selection could have
46 led to such hugely complex forms of sociality. To answer this question, we used simulated
47 teams of robots and artificially evolved them to achieve maximum performance in a foraging
48 task. We find that, as in social insects, this favored controllers that caused the robots to
49 display a self-organized division of labor in which the different robots automatically
50 specialized into carrying out different subtasks in the group. Remarkably, such a division of
51 labor could be achieved even if the robots were not told beforehand how the global task of
52 retrieving items back to their base could best be divided into smaller subtasks. This is the first
53 time that a self-organized division of labor mechanism could be evolved entirely de-novo. In
54 addition, these findings shed significant new light on the question of how natural systems
55 managed to evolve complex sociality and division of labor.

56

57 **Introduction**

58 The “major transitions in evolution”, whereby cells teamed up to form multicellular organisms
59 or some animals went on to live in societies, are among the keys to the ecological success of
60 much life on earth [1]. The efficiency of both organisms and animal societies frequently
61 depends on the presence of an advanced division of labor among their constituent units [2-4].
62 The most celebrated examples can be found in social insects, which exhibit astonishing levels
63 of social organization and are ecologically dominant in many natural ecosystems [5,6].
64 Through division of labor, social insects can perform complex tasks by dividing them up into
65 smaller sub-tasks carried out by different sets of individuals [7-10]. Although the adaptive
66 benefits of division of labor are evident, the way in which it can evolve is more enigmatic,
67 since an effective division of labor requires the simultaneous co-occurrence of several
68 complex traits, including self-organized mechanisms to decompose complex tasks into
69 simpler subtasks, mechanisms to coordinate the execution of these tasks, mechanisms to
70 allocate an appropriate number of individuals to each task, and the ability of individuals to
71 effectively carry out each of the subtasks [4]. The complexity of this co-evolutionary problem
72 is further exacerbated by the fact that division of labor should also be flexible to be able to
73 cope with changing environmental conditions [4,10,11].

74 To date, most analytical and individual-based simulation models of division of labor
75 [4,9,10,12-16] have focused merely on determining the optimal proportion of individuals
76 engaging in different tasks [12] or on determining optimal task allocation mechanisms
77 [4,9,10,13,16], sometimes in relation to particular levels of intragroup genetic variation [14,15].
78 These studies implicitly assume that pre-optimized behaviors to carry out each of the different
79 subtasks, which we refer to as “pre-adapted behavioral building blocks”, are already present
80 in nonsocial ancestors [17], and that division of labor merely involves the rewiring of these
81 behaviors. Empirical support for this hypothesis can be found for example in the somatic cell
82 differentiation in multicellular organisms, which is derived from a genetic switch involved in the
83 induction of diapause during stress periods in unicellular ancestors [2,18]. Similarly, in insect
84 societies, worker brood care is thought to be derived from ancestral parental care [19], and
85 reproductive division of labor as well as worker task specialization may be derived from
86 mechanisms that initially regulated reproduction and foraging in solitary ancestors [17,20-22].

87 A limitation of traditional analytical modeling approaches to division of labor [4,10],
88 however, is that they can only consider a finite and pre-specified number of behavioral
89 strategies. In recent years, artificial evolution of teams of embodied agents has been used to
90 enable the study of social traits in more detail, taking into account more realistic physical
91 constraints and a much larger set of allowable behaviors and strategies [23-25]. In
92 evolutionary swarm robotics, for example, this framework has been used to study the
93 evolution of the origin of communication [26,27], collective transport [28], collective motion
94 [29], aggregation [30-32] and chain formation [33] (reviewed in [23,24,34-37]). Nevertheless,
95 to date, no study in evolutionary swarm robotics has succeeded in evolving complex, self-
96 organized division of labor entirely de novo [38,39]. This may be due to the fact that most
97 evolutionary robotics studies have made use of neural network-based approaches [23-25,36],
98 which have been shown to scale badly to more complex problems [38,40].

99 The main aim of our study was to test if other nature-inspired evolutionary methods
100 than traditionally used in evolutionary swarm robotics would be able to achieve complex task
101 specialization in social groups. Analogously to the situation in nature where subtask behaviors
102 may or may not be recycled from pre-adapted behavioral building blocks, we do this using
103 one of two approaches, in which we either do or do not pre-specify the behaviors required for
104 carrying out the different subtasks. Evidently, we expected that task specialization could
105 evolve much more easily when pre-adapted behavioral building blocks were present, but we
106 were also interested to see if a self-organized mechanism of task specialization could be
107 evolved entirely de-novo using our recently developed method of Grammatical Evolution [41].
108 This nature-inspired evolutionary method allows a set of low-level behavioral primitives to be
109 recombined and evolved into complex, adaptive behavioral strategies through the use of a
110 generative encoding scheme that is coupled with an evolutionary process of mutation,
111 crossover and selection [41].

112 The type of division of labor we consider in our set-up is known as “task partitioning”,
113 and requires different tasks to be carried out in sequence by different sets of individuals [7]. In
114 particular, our experimental scenario was inspired by a spectacular form of task partitioning
115 found in some leafcutter ants, whereby some ants (“droppers”) cut and drop leaf fragments
116 into a temporary leaf storage cache and others (“collectors”) specialize in collecting and

117 retrieving the fragments back to the nest [42,43] (Fig. 1). In our analogous robotics setup, we
118 used a team of robots [44] simulated in-silico using an embodied swarm robotics simulator
119 [45] (Fig. 2) and required the robots to collect items and bring them back to the nest in either
120 a flat or sloped environment (see Figs. 1 and 2b and Material and Methods). In this setup,
121 task specialization should be favored whenever some features of the environment (in our
122 case, the presence of a slope) can be exploited by the robots to achieve faster foraging
123 (“economic transport”, [46]) and reduce switching costs between different locations [9,47].
124 The results of these experiments show for the first time that complex, self-organized task
125 specialization and task allocation could be evolved in teams of robots. Nevertheless, a fitness
126 landscape analysis also demonstrates that task specialization was much easier to evolve
127 when pre-evolved behavioral building blocks were present. We use these findings as a
128 starting point to speculate about the likely path that nature took to evolve complex sociality
129 and division of labor. Furthermore, we discuss the potential of our nature-inspired
130 evolutionary method for the automated design of swarms of robots displaying complex forms
131 of coordinated, social behavior.

132

133 **Materials and Methods**

134 *The task and the environment*

135 Our experimental setup is inspired by the type of task partitioning observed in *Atta* leafcutter
136 ants [42,43], that collect leaves and other plant material as a substrate for a fungus that is
137 farmed as food (Fig. 1a). In these insects, particularly in species that harvest leaves from
138 trees, leaf fragments are retrieved in a task partitioned way, whereby some ants (“droppers”)
139 specialize in cutting and dropping leaf fragments to the ground, thereby forming a leaf cache,
140 and others specialize in collecting leaves from the cache to bring them back to the nest
141 (“collectors”) [42,43]. In addition, another strategy is known whereby the whole leaf cutting
142 and retrieval task is carried out by single individuals (“generalists”), without any task
143 partitioning [42,43]. Task partitioning in this scenario is thought to be favored particularly in
144 situations where the ants forage on leaves from trees, due to the fact that the leaf fragments
145 can then be transported purely by gravity, which saves the ants the time to climb up and down
146 the tree, and the fact that there are few or no costs associated with material loss thanks to the

147 large supply of leaves [7,43,48] (Fig. 1a). This theory is supported by the fact that species
148 living in more homogeneous grassland usually retrieve leaf fragments in an unpartitioned
149 way, without first dropping the leaves (Fig. 1c), particularly at close range to the nest [43,49].

150 In the corresponding robotic setup, we substituted the tree with a *slope area* and leaves
151 with cylindrical items. A team of robots then had to collect these items from what we call the
152 *source area* and bring them back to what we refer to as the *nest area* (Fig. 1b). Simulations
153 were carried out using the realistic, physics-based simulator ARGoS [45]. As demonstrated in
154 the past, controllers developed within ARGoS can be directly transferred to real robots with
155 minimal or no intervention [50,51]. The robots involved in the experiments were a simulated
156 version of the foot-bot robot [44], which is a differential-drive, non-holonomic, mobile robot
157 (Fig. 2a). A screen-shot of a simulation instant is shown in Figure 2b. We used a setup
158 whereby 5 items were always present in the source area. The 5 items were replaced and put
159 in a random position within the source area each time a robot picked up one of them. This is
160 justified by the fact that leaf availability in the natural environment is often virtually unlimited. A
161 light source was placed at a height of 500 m, 500 m away from the nest, in the direction of the
162 source area. The light allowed the robots to navigate in the environment, since phototaxis
163 allowed them to go towards the item source, whereas anti-phototaxis allowed them to return
164 to the nest. The slope area had an inclination of about 8 degrees. The linear velocity of the
165 robots on the flat part of the arena was 0.15 m/s, but this reduced to a maximum speed of
166 0.015 m/s when they had to climb up the slope, and increased to 0.23 m/s when they came
167 down from the slope. If an item was dropped in the slope area, it slid down the slope at a
168 speed of 1 m/s until it reached the *cache area*, where it was stopped due to friction and their
169 impact with other items in the cache. This was done to simulate leaves being dropped from
170 the tree, as in Figure 1a. In addition, in some of the experiments, we considered a flat
171 environment of the same length and width as the one described above (Fig. 1d), to mirror the
172 case in nature where ants forage in a flat, homogeneous environment (Fig. 1c).

173

174 *Evolution of task-partitioning from pre-adapted building blocks*

175 In a first set of experiments, we assumed that the behavioral strategies required to carry out
176 each of the subtasks (dropper or collector behavior, as well as generalist, solitary foraging)

177 were available to the robots as pre-adapted behavioral building blocks and then determined
178 the optimal mix of each of the strategies [12]. This setup, therefore, matched some
179 evolutionary scenarios proposed for the origin of division of labor in biological systems based
180 on co-opting pre-adapted behavioral patterns [2,17-22]. In addition, this scenario allowed us
181 to determine under which environmental conditions task partitioning is favored, and provided
182 a fitness benchmark for the second scenario below, where task partitioning was evolved
183 entirely de-novo.

184 In this first set of experiments, dropper, collector and generalist foraging strategies
185 were implemented as follows:

186 (1) **Dropper strategy:** A dropper robot is a robot that climbs the slope area and never
187 descends it again, continuously collecting items from the source area and dropping
188 them to the slope area.

189 (2) **Collector strategy:** A collector robot is a robot that never climbs the slope area.
190 Instead, it continuously collects items from the cache (when present) and brings them
191 back to the nest. If it cannot find any items, the collector robot keeps exploring the
192 cache area by performing random walk, until an item is found.

193 (3) **Generalist strategy:** A generalist robot is a robot that performs a standard foraging
194 task. It climbs the slope and explores the source area, collects items, and brings them
195 all the way back to the nest. The generalist robot does not explore the cache area,
196 but in case it finds an item at the cache while going towards the source, it collects it
197 and brings it back to the nest.

198 The rules that we employed to implement these strategies are shown in Supplemental Table
199 S1. We also assumed that the robots would specialize for life in each of these available
200 strategies according to a particular evolved allocation ratio. This was equivalent to assuming
201 that in nature, these behavioral strategies would already have evolved due to selection in their
202 ancestral environment, and that natural selection would favor a particular hard-wired
203 individual allocation between the different sets of tasks, e.g. through fine-tuning of the
204 probability of expression of the gene-regulatory networks coding for the different behavioral
205 patterns. For these experiments, we used teams of 4 robots, to match the evolutionary
206 experiments with fine-grained building blocks (cf. next section). Subsequently, a fitness

207 landscape analysis was used to determine the optimal mix between the three strategies in
208 one of two possible environments, a flat or a sloped one (Fig. 1b, d). This was done via
209 exhaustive search, that is, by testing all possible ratio combinations and determining the
210 corresponding fitness values in the two environments, rather than using an evolutionary
211 algorithm. This was possible due to the relatively small search space, which gave access to
212 the full fitness landscape. Group performance, measured by the total number of items
213 retrieved to the nest over a period of 5,000 simulated seconds, for each possible mix of the
214 three strategies, was measured in 10 simulated runs and then averaged.

215

216 *Evolution of task-partitioning from first principles*

217 In a second set of experiments, we considered an alternative scenario where both task
218 specialization and task allocation could evolve entirely de-novo, starting only from basic, low-
219 level behavioral primitives. These primitives were simply navigational behaviors allowing
220 robots to either go towards the source or towards the nest, as well as a random walk
221 behavior:

- 222 (1) **PHOTOTAXIS**: uses the light sensor to make the robot go towards the direction with the
223 highest perceived light intensity.
- 224 (2) **ANTI-PHOTOTAXIS**: uses the light sensor to make the robot go towards the lowest
225 perceived light intensity.
- 226 (3) **RANDOM WALK**: makes the robot move forward for a random amount of time and then
227 turn to a random angle, repeating this process while the block is activated, without using
228 any sensors.

229 In addition, a mechanism of obstacle avoidance, based on the robot's range and bearing and
230 proximity sensors, was switched on at all times to avoid that the robots would drive into each
231 other or into the walls of the foraging arena. Finally, two instantaneous actions were allowed,
232 namely picking up and dropping an item. To be able to evolve adequate behavioral switching
233 mechanisms, we allowed the robots to perceive their position in space, that is, whether they
234 were in the source, slope, cache or nest, based on sensorial input from the ground and light
235 sensors, as well as perceive whether or not they were currently holding an item.

236 The fine-grained behavioral building blocks were combined together using a method
 237 known as grammatical evolution [52] as implemented in GESwarm [41], in order to evolve
 238 rule-based behaviors representing more complex strategies. GESwarm was developed for the
 239 automatic synthesis of individual behaviors consisting of rules leading to the desired collective
 240 behavior in swarm robotics. These rules were represented by strings, which in turn were
 241 generated by a formal grammar. The space of strings of such a formal grammar was used as
 242 a behavioral search space, and mutation, crossover and selection were then used to favor
 243 controllers that displayed high group performance.

244 The individual behavior of a given robot was expressed by a set \mathbf{R} of an arbitrary
 245 number n_R of rules R_i :

$$246 \quad \mathbf{R} = \{R_i\}, i \in \{1, \dots, n_R\}.$$

247 Each rule was composed of three components:

$$248 \quad R_i = P_i \times B_i \times A_i,$$

249 where B_i denotes a subset of all possible fine-grained behavioral building blocks (phototaxis,
 250 anti-phototaxis and random walk), A_i denotes a subset of all possible instantaneous *actions*
 251 (pickup, drop, change behavior or change an internal state variable) and P_i denotes a subset
 252 of all possible *preconditions*. The preconditions were specified as logical conditions with
 253 respect to the current value of a number of state variables, which included both sensorial
 254 input (the environment they were in and whether or not they were carrying an item) and
 255 internal state variables (a state variable that specified whether they wanted to pick up an item
 256 or not and two memory state variables, with evolvable meaning).

257 If all the preconditions in P_i were met, and if a given robot was executing any of the
 258 low-level behaviors present in B_i , all actions contained in A_i were executed with evolvable
 259 probability p_i . In this way, we could allow the evolution of probabilistic behaviors, which have
 260 been extensively used both in the swarm robotics literature [53,54] and as microscopic

261 models of the behavior of some social animals [55,56]. Finally, each robot executed all rules
262 and actions in their order of occurrence.

263 To be able to generate the rules above, we devised a grammar using the Extended
264 Backus-Naur Form notation [57]. Within the framework of grammatical evolution [41,52], a
265 genotype represented a sequence of production rules to be followed to produce a valid string
266 (in our case a set of rules) starting from that grammar. Mutation and crossover acted at the
267 level of this genotype, modifying the sequence of production rules. The full grammar of
268 GESwarm is described in [41].

269 Biologically speaking, our GESwarm rule-based controllers can be considered
270 analogous to gene-regulatory networks or to a brain, whereby internal memory state variables
271 in our model can be seen as an analogy to epigenetic states or memory states in the brain.
272 Furthermore, as in biological systems, we use a generative encoding (a string coding for a
273 series of conditional rules, similar to a DNA sequence coding for conditionally expressed gene
274 regulatory networks) and evolve our system using mutation and crossover. One departure in
275 our setup from biological reality, however, was that we used genetically homogeneous teams,
276 as is common in evolutionary swarm robotics [58], but different from the situation in most
277 social insects, where sexual reproduction tends to be the norm. This choice was made
278 because homogeneous groups combined with team-level selection has been shown to be the
279 most efficient approach to evolve tasks that require coordination [28]. Nevertheless, this setup
280 can still be considered analogous to the genetically identical cells of multicellular organisms
281 [59] or the clonal societies of some asexually reproducing ants [60] that both display complex
282 forms of division of labor.

283 We executed a total of 22 evolutionary runs on a computer cluster, of which we used
284 100 to 200 nodes in parallel. The number 22 was chosen to meet the total amount of
285 computational resources we had at our disposal (3 months of cluster time) and was
286 statistically speaking more than adequate. All evolutionary runs were carried out for 2,000
287 generations using 100 groups of 4 robots and were each evaluated 3 times. This relatively
288 small number of robots was chosen to limit the computational burden of the evolutionary runs.
289 Nevertheless, we also verified if the evolved controllers could be scaled to larger teams of 20
290 robots each. In this case, the foraging arena was scaled in direct proportion with the number

291 of robots. We used single-point crossover with crossover probability 0.3 and mutation
292 probability 0.05. We chose a generational replacement with 5% elitism, in order to exploit
293 parallel evaluation of multiple individuals on a computer cluster. We used roulette-wheel
294 selection, that is, the probability to select a given genotype was proportional to its fitness
295 relative to the average fitness of all genotypes in the population. As fitness criterion we used
296 group performance, measured as the total number of items retrieved to the nest over a period
297 of 5,000 seconds. During post-evaluation, this same fitness criterion was used to evaluate the
298 evolved controllers. We also assessed the average absolute linear speed of the robots along
299 the long axis of the arena, measured as a percentage of the theoretical maximum speed, and
300 the degree of task specialization, measured by counting how many items have been retrieved
301 through the action of more than one robots (i.e. via specialists) and by dividing this number by
302 the total number of objects retrieved (i.e. via specialists + via only one generalist).

303

304 **Results**

305 *Evolution of task-partitioning from pre-adapted building blocks*

306 In the first set of simulations, we assumed that each robot could specialize for life to one
307 among the three possible preexisting behavioral strategies required for task partitioning,
308 dropper, collector and generalist, and determined the optimal mix between the three
309 strategies based on an exhaustive search on the full fitness landscape (Fig. 1b, d). These
310 simulations were performed both in a flat and a sloped environment. As proposed for natural
311 systems [7,43,48], our a priori hypothesis was that task partitioning would be favored
312 particularly in the sloped environment, and that maximal group performance would be
313 achieved when some robots would drop items in a cache and others specialized in collecting
314 items from the cache. This is because, in such an environment, some of the robots would be
315 able to avoid the costly traversal of the slope area (i.e. avoid switching costs) and the fact that
316 gravity could also help to move items across the slope, thereby resulting in more economical
317 transport (Fig. 1).

318 Examination of the obtained fitness landscapes reveals that there was one globally
319 attracting optimum in each of the two environments considered (Fig. 3a, b). As expected, this
320 optimum involved task partitioning in the sloped environment (Fig. 3b), with a mix of 50%

321 droppers and 50% collectors being most efficient, but only generalist foraging in the flat
322 environment (Fig. 3a, Videos S1 and S2). In addition, our fitness landscape analysis showed
323 that when pre-adapted behavioral building blocks can be used in the evolutionary process,
324 the fitness landscape tends to be very smooth, thereby making task specialization easily
325 evolvable, without the risk of the system getting trapped in suboptimal local optima. It should
326 also be noted that in our setup, the absolute group performance was significantly higher (t -
327 test, $t=-16.6$, $d.f.=18$, $p<10^{-11}$) in the sloped environment (144.1 ± 4.3 S.D. items collected in
328 5,000 s, $n=10$) than in the flat one (120.2 ± 1.4 S.D. items collected in 5,000 s), due to the fact
329 that in the first case, gravity helped to move the items towards the source.

330

331 *Evolution of task-partitioning from first principles*

332 In a second set of experiments, we used GESwarm [41] to evolve task specialization and task
333 allocation entirely de-novo, starting only from basic, low-level behavioral primitives (see
334 Materials and Methods). Surprisingly enough, these evolutionary experiments demonstrated
335 that task partitioning and fully self-organized task specialization and task allocation could also
336 emerge entirely from scratch by selecting purely on overall group performance (number of
337 items retrieved to the nest). In particular, our experiments show that in 59% (13 out of 22) of
338 the runs, the majority of the items were retrieved by the robots in a task-partitioned way in the
339 final evolved controller obtained after 2,000 generations (Fig. 4, Videos S3 and S4.1-S4.22).
340 In these cases, most of the items were first dropped by one robot and later picked up by
341 another one. In contrast to the case with predefined behavioral strategies, however, the task
342 specialization that was seen in these controllers did not entail fixed roles, but instead was
343 characterized by a dynamic allocation in response to the size of the cache. An example of a
344 controller (nr. 2) displaying such behavior is shown in Supplemental Video S3, where the
345 majority of the robots first exploit the source to act as droppers, but then move down the slope
346 as the cache fills up to act as collectors (the evolved rules of this controller are shown in
347 Table S2). The robots shown in these simulations used simple probabilistic rules to switch
348 from the source to the cache area, while the cache itself was exploited to switch from the
349 cache area back to the source area. We observed that the latter mechanism was also very
350 simple and based on stigmergy, i.e. robots would collect from the cache whenever objects

351 were found on the way, but would continue all the way to the source when cache items were
352 not encountered. Thanks to these simple mechanisms, the robots could dynamically switch
353 roles in response to the size of the cache. The same adaptive specialization dynamics are
354 apparent in Figure 5a, where the density of the robot positions across the arena is shown
355 across the 30 runs used for post-evaluation of the same controller, and in Figure 5b, which
356 displays the individual trajectories of the four robots in a typical evaluation run (the spatial
357 segregation and robot trajectories for all other evolved controllers are shown in Fig. S1).

358 That such self-organized task specialization and task allocation could evolve from first
359 principles by selecting purely on group performance is significant, given that we started from a
360 random controller that barely achieved any foraging during the first few generations (Fig. 4,
361 Video S1). As in the case without pre-adapted building blocks that we considered in the
362 previous section, also here, the presence of a slope was sufficient for the evolution of task
363 partitioning. Indeed, when we conducted the very same experiments in a flat environment,
364 none of the controllers evolved task partitioning and generalist foraging was the favored
365 strategy [41].

366 Significantly, the evolved rules for both generalist foraging [41] and task partitioned
367 object retrieval scaled very well also to larger teams of robots. An example is shown in Video
368 S1, where one of the evolved controllers from a 4 robot team is implemented in a team of 20
369 robots. In this case, the achieved group performance scaled almost perfectly with the
370 increase in group size (457 ± 72 S.D. in the 20 robot team vs. 103 ± 24 S.D. in the 4 robot
371 one). Excellent scalability properties were also shown by the fact that for the 8 best evolved
372 controllers, the performance ratio of the rules when implemented in the 20 robot teams
373 relative to that in the 4 robot ones in which the rules were first evolved was very close to the
374 expected linear scaling factor of 5 (4.4 , S.D. 0.14 , see Table S3).

375 Although the lack of fixed roles precluded an analysis in terms of behavioral roles
376 similar to that presented in the section above, it turned out that both increased amounts of
377 task partitioning and higher average linear speeds significantly increased group fitness
378 (multiple regression analysis, $p < 0.01$ and $p < 10^{-5}$, respectively, $n = 22$, Fig. 6). In fact, all 8
379 evolved controllers displaying a high group performance (top 35%, >ca. 100 items collected)
380 had very high levels of task partitioning ($92\% \pm 0.08$ S.D. of all items retrieved in a task

381 partitioned way) and achieved a high average linear speed ($31\% \pm 0.6\%$ S.D. of the
382 theoretical maximum). Significantly, out of these 8, the performance of the best evolved
383 controller (135 ± 14 S.D., $n=30$ items retrieved) was not significantly different from the optimal
384 2 dropper-2 collector mix obtained in the experiment using hand-coded behavioral strategies
385 above (144.1 ± 4.3 S.D., t -test, $t=2.01$, d.f.=38, $p > 0.05$). Among these 8 best controllers,
386 between 4 and 11 rules were used to switch between the different allowed behaviors and
387 instantaneous actions (cf. evolved rules shown in Supplemental Table S2). Interestingly, in 3
388 of these best controllers, the rules employed as a precondition a memory state variable that
389 was increased or decreased as a result of some of the actions performed in other rules. In
390 principle, the use of these state variables could have allowed for the evolution of mechanisms
391 akin to the response threshold model, which has been extensively used in studies on division
392 of labor [4,9,10,16]. Nevertheless, none of our controllers succeeded in evolving this
393 particular mechanism, and task allocation instead appeared to be based purely on
394 probabilistic and stigmergic switching, as explained above.

395 A detailed analysis of the fitness and behavior of the final evolved controllers
396 demonstrated that there was one global optimum characterized by a high level of task
397 partitioning and high linear speed (Fig. 6). Nevertheless, some runs were trapped in
398 suboptimal regions of the search space. For example, some controllers merely displayed
399 generalist foraging, which was suboptimal in our setup (Fig. 6, bottom right points). Similarly,
400 other controllers were characterized by defective locomotory skills, even if some actually
401 achieved task partitioning (Fig. 6, left blue points). Finally, two evolved controllers were
402 characterized by a high degree of task partitioning and a decent speed, but nevertheless had
403 low overall performance due to the use of a suboptimal dropping strategy, characterized by a
404 continuous dropping and picking up in all the regions of the environment, which affected
405 performance but not speed and degree of task partitioning (Fig. 6, two blue points in the
406 upper-right corner). These outliers, however, did not change the fact that fitness was strongly
407 correlated with both the degree of task specialization and the linear speed of the robots.

408 Despite the variation in performance of the final evolved controllers, an analysis of
409 fitness and degree of task partitioning over the course of the evolutionary runs (Fig. 4) clearly

410 demonstrates that high task partitioning was evolutionarily stable, since any transition to high
411 task partitioning never reverted back to generalist foraging in later generations.

412

413 **Discussion**

414 One of the unsolved mysteries in biology is how a blind process of Darwinian selection could
415 have led to the hugely complex forms of sociality and division of labor as observed in insect
416 societies [4]. In the present paper, we used simulated teams of robots and artificially evolved
417 them to achieve maximum team performance in a foraging task. Remarkably, we found that,
418 as in social insects, this could favor the evolution of a self-organized division of labor, in which
419 the different robots automatically specialized into carrying out different subtasks in the group.
420 Furthermore, such a division of labor could be achieved merely by selecting on overall group
421 performance and without pre-specifying how the global task of retrieving items would best be
422 divided into smaller subtasks. This is the first time that a fully self-organized division of labor
423 mechanism could be evolved entirely de-novo. Overall, these findings have several important
424 implications. First, from a biological perspective, they yield novel evidence for the adaptive
425 benefits of division of labor and the environmental conditions that select for it [4], provide a
426 possible mechanistic underpinning to achieve effective task specialization and task allocation
427 [4] and provide possible evolutionary pathways to complex sociality. Second, from an
428 engineering perspective, our nature-inspired evolutionary method of Grammatical Evolution
429 clearly has significant potential as a method for the automated design of adaptively behaving
430 teams of robots.

431 In terms of the adaptive benefits of division of labor and the environmental conditions
432 that select for it, our results demonstrated that task partitioning was favored only when
433 features in the environment (in our case a slope) could be exploited to achieve more
434 economic transport and reduce switching costs, thereby causing specialization to increase the
435 net efficiency of the group. Previous theoretical work has attributed the evolution of task
436 specialization to several ultimate factors, some of which are hard to test empirically [61].
437 Duarte et al. [4], for example, reviewed modeling studies that showed that the adaptive
438 benefits of a behaviorally-defined division of labor could be linked to reduced switching costs
439 between different tasks or locations in the environment, increased individual efficiency due to

440 specialization, increased behavioral flexibility or reduced mortality in case only older
441 individuals engage in more risky tasks (“age polyethism”). Out of these, there is widespread
442 agreement on the role of switching costs and positional effects as key factors in promoting
443 task specialization [4,10,47,62], and our work confirms this hypothesis. Indeed, in our set-up,
444 task partitioning greatly reduced the amount of costly switching required between
445 environmental locations. Furthermore, our work also confirms the economic transport
446 hypothesis, i.e. that task partitioning results in more economical transport, which in our case
447 was due to the fact that gravity acted as a helping hand to transport the items. Previously, this
448 hypothesis had also found significant empirical support [7,43,46,48], e.g. by the fact that in
449 leafcutter ants, species that collect leaves from trees tend to engage in task partitioned leaf
450 retrieval, whereas species living in more homogeneous grassland usually retrieve leaf
451 fragments in an unpartitioned way, without first dropping the leaves, particularly at close
452 range to the nest [43,49].

453 A surprising result in our evolutionary experiments was that adaptive task
454 specialization was achieved despite the fact that the robots in each team all had identical
455 controllers encoded by the same genotype. This implies that a combination of individual
456 experience, stigmergy and stochastic switching alone were able to generate adaptive task
457 specialization, akin to some of the documented mechanisms involved in behavioral task
458 specialization in some asexually reproducing ants [63] and in cell differentiation in
459 multicellular organisms and clonal bacterial lineages [59,64,65]. The choice of using
460 homogeneous, clonal groups of robots with an identical morphology precluded other
461 mechanisms of division of labor observed in nature from evolving, based, for instance, on
462 morphological [4,12] or genetic [4] role specialization. Such mechanisms, however, could be
463 considered in the future if one allowed for genetically heterogeneous robot teams [28] or
464 evolvable robot morphologies. Lastly, the grammar we used did not specifically allow for
465 recruitment signals to evolve, such as those observed in leafcutting ants, where both trail
466 pheromones and stridulation are used as mechanisms to recruit leaf cutters [66,67], or the
467 ones in honeybees, where the tremble dance is used to regulate the balance between
468 number of foragers and nectar receivers inside the colony [68,69]. Nevertheless, including
469 low-level primitives for communication behavior into the grammar, which we plan to do in

470 future work, would readily allow for the evolution of such mechanisms, and would likely boost
471 the performance of the evolved controllers even further (cf. [26,27]).

472 In terms of the mechanisms of task specialization and task allocation evolved, our
473 work is important in that alleviates one of the limitations of existing models on the evolution of
474 task specialization, namely, that they normally take pre-specified subtasks and an existing
475 task allocation model (e.g. the response threshold model) as point of departure [4], thereby
476 greatly constraining the path of evolution. Our work is an important cornerstone in
477 establishing, at the best of our knowledge, the first model that bridges the gap between self-
478 organization and evolution without significantly constraining the behavioral strategies and
479 coordination mechanisms that can be obtained to achieve optimal task specialization and task
480 allocation. In fact, compared to other previous studies on evolution of task specialization
481 [47,62,70-72], our work is the first time to consider non-predefined sub-tasks that could evolve
482 de-novo and combine into complex individual behavioral patterns.

483 Although our experiments demonstrate that division of labor and behavioral
484 specialization in teams of identical robots could evolve in both the scenarios we considered,
485 fitness landscape analyses showed that optimal task allocation could be achieved most easily
486 if optimized behaviors capable of carrying out the different subtasks were available as pre-
487 adapted behavioral building blocks. This leads us to suggest that when building blocks are
488 solidified in earlier stage of the evolution, complex coordination strategies such as task
489 specialization are more likely to evolve as the fitness landscape becomes smoother and also
490 easier to explore due to its greatly reduced size. In addition, it brings further support for the
491 hypothesis that, in nature, the evolution of division of labor in social groups and other
492 transitions in the evolution of sociality also tends to be based on the co-option of pre-existing
493 behavioral patterns, as opposed to requiring the de-novo evolution of many entirely new
494 social traits [17]. Our results, therefore, match and can be integrated with available evidence
495 with respect to the importance of preadaptations in the origin of advanced forms of sociality
496 [2,17-22,73], where, for example, reproductive division of labor and worker task specialization
497 are thought to be derived from mechanisms that initially regulated reproduction and foraging
498 in solitary ancestors [17,20-22], sibling care is thought to be derived from ancestral parental
499 care [19], and reproductive altruism (i.e., a sterile soma) in some multicellular organisms

500 evolved via the co-option of a reproduction-inhibiting gene expressed under adverse
501 environmental conditions [73]. Furthermore, it confirms other studies that have examined the
502 building block hypothesis with various digital systems, for example in the context of genetic
503 algorithms [74], evolution of single robot morphologies [75] and the open-ended evolution of
504 simple computer programs [76].

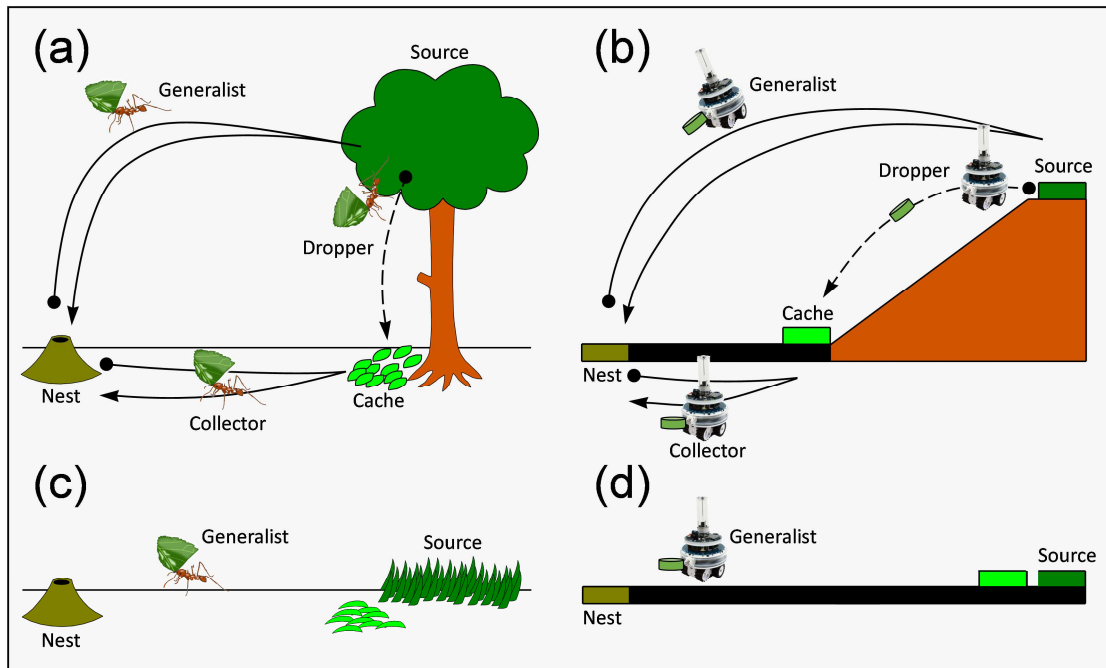
505 From an engineering perspective our study is the first to achieve a complex form of
506 division of labor using an evolutionary swarm robotics approach, and the first to use the
507 method of Grammatical Evolution to evolve complex, non-trivial behavioral patterns. This
508 result is novel in the field of evolutionary swarm robotics, where, few exceptions aside, most
509 studies have used non-incremental and non-modular approaches, e.g. based on monolithic
510 neural networks [38,77]. In fact, previously, the only other studies which evolved a
511 rudimentary task allocation in swarms of robots were those of Tuci *et al.* [78], which used a
512 neural network controller combined with a fitness function favoring a required preset task
513 allocation [78], of Duarte *et al.* [40], which used evolved neural network controllers capable of
514 carrying out particular subtasks, which were then combined with a manually engineered
515 decision tree, and the work of refs. [79-81], which used open-ended evolution and a simplified
516 robotic scenario to evolve heterogeneous behaviors for collective construction [79,80] and
517 pursuit [81] tasks in presence of a pre-specified set of three sub-tasks. Typically, the
518 behavioral complexity that could be reached in these artificial neural network-based studies
519 was quite limited, making the evolution of self-organized task specialization in homogeneous
520 groups out of reach for these methods. In fact, the evolution of self-organized task
521 specialization would clearly require a non-standard neural network approach, involving
522 recurrent neural connections to keep track of the internal state (e.g. the current direction of
523 motion to be able to perform phototaxis), a mechanism to achieve modularity and a
524 mechanism to switch stochastically between these modules. Extending the neural network
525 approach used in evolutionary swarm robotics to this level of complexity would be an
526 interesting task for the future. Other studies on task allocation and task partitioning in swarm
527 robotics typically used traditional, manually engineered approaches [82-88] (reviewed in [89]).
528 All these methods are significantly less general than ours, given that we used a nature-
529 inspired automatic design method with a single fitness criterion, group performance, without

530 any pre-engineered decision-making mechanisms, and simultaneously evolved a self-
531 organized task decomposition and task allocation mechanism as well as optimized behaviors
532 to carry out each of the evolved subtasks. We therefore believe that GESwarm and
533 grammatical evolution will play a key role in the future of evolutionary swarm robotics.

534 Overall, our work and the results we obtained are therefore important both to explain
535 the origin of division of labor and complex social traits in nature, as well as to advance the
536 field of evolutionary swarm robotics, as we showed that the novel methodological and
537 experimental tools we developed were able to synthesize controllers that were beyond the
538 level of complexity achieved to date in the field.

539


540



542

543 **Figure 1.** Task partitioning in insects and robots. **(a)** Task partitioned retrieval of leaf
 544 fragments, as found in most *Atta* leafcutter ants that harvest leaves from trees [7,43]. Dropper
 545 ants cut leaves which then accumulate in a cache, after which the leaves are retrieved by
 546 collectors and brought back to the nest, where they serve as a substrate for a fungus which is
 547 farmed as food. Ants also occasionally use a generalist strategy whereby both tasks are
 548 performed by the same individuals. **(b)** Analogous robotics setup, whereby items have to be
 549 transported across a slope using the coordinated action of droppers, collectors and possibly
 550 generalists. **(c)** Grasscutting leafcutter ants cutting leaf fragments in a flat environment
 551 without task partitioning, using a generalist foraging strategy [49]. **(d)** Analogous robotics
 552 setup, with robots being required to collect items in a flat arena.

553

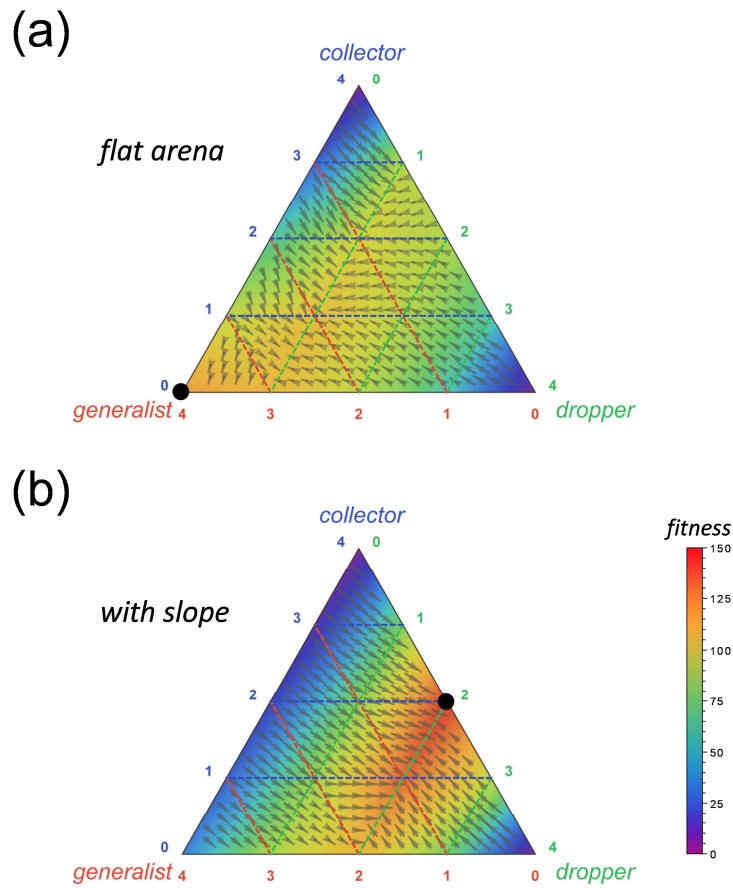


554

555 **Figure 2.** Foot-bot robots and ARGoS simulation platform. **(a)** The foot-bot robot [44] and its
 556 sensors and actuators. **(b)** A snapshot of the ARGoS [45], the physics-based simulator used
 557 in our experiments. The snapshot shows the different elements composing our experimental
 558 setup.

559

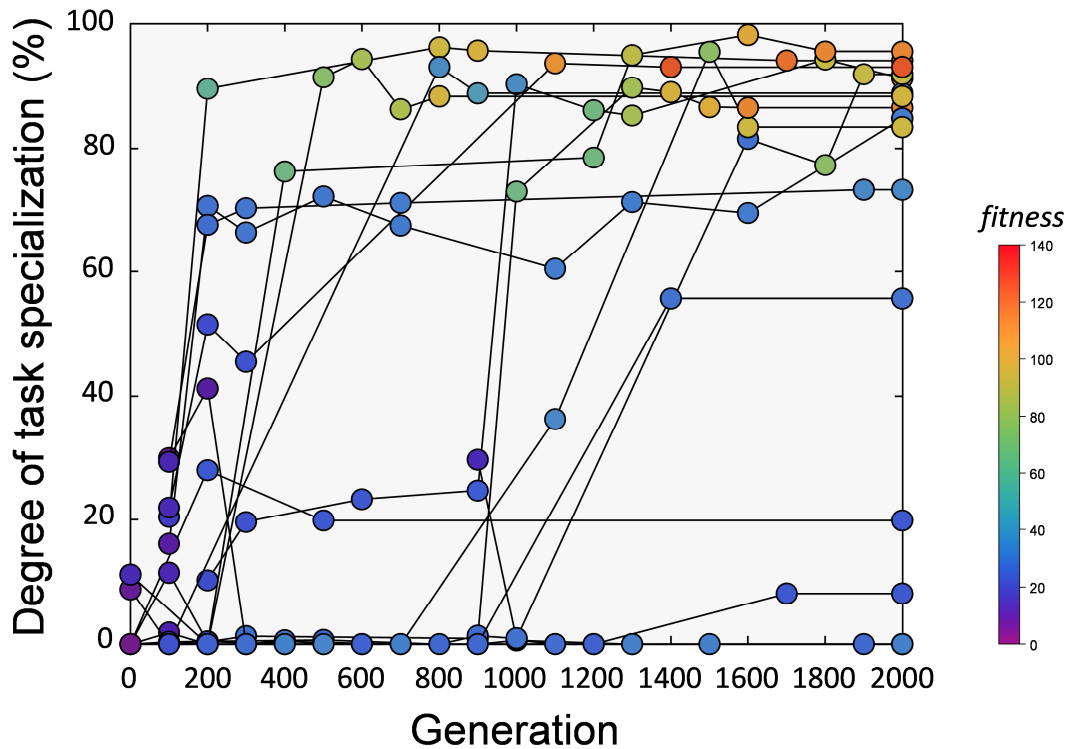
560



561

562 **Figure 3.** Optimal group composition in 4 robot teams using pre-adapted dropper, collector or
 563 generalist foraging strategies (cf. handcoded rules shown in Table S1). Ternary plots show
 564 group performance (total number of items retrieved to the nest over a period of 5,000
 565 simulated seconds averaged over 10 simulation runs, color coded) as a function of the
 566 number of collectors (blue), droppers (green) and generalist foragers (red) in the 4 robot
 567 teams (black dot=optimum). In a flat environment **(a)**, teams of generalist foragers achieve
 568 optimal performance (cf. Supplemental Video S2), whereas in a sloped arena **(b)**, a mix of 2
 569 droppers and 2 collectors is most optimal (cf. Video S1). Both of these optima are global
 570 attractors in their respective fitness landscapes (cf. vectors which represent the phase
 571 portrait).

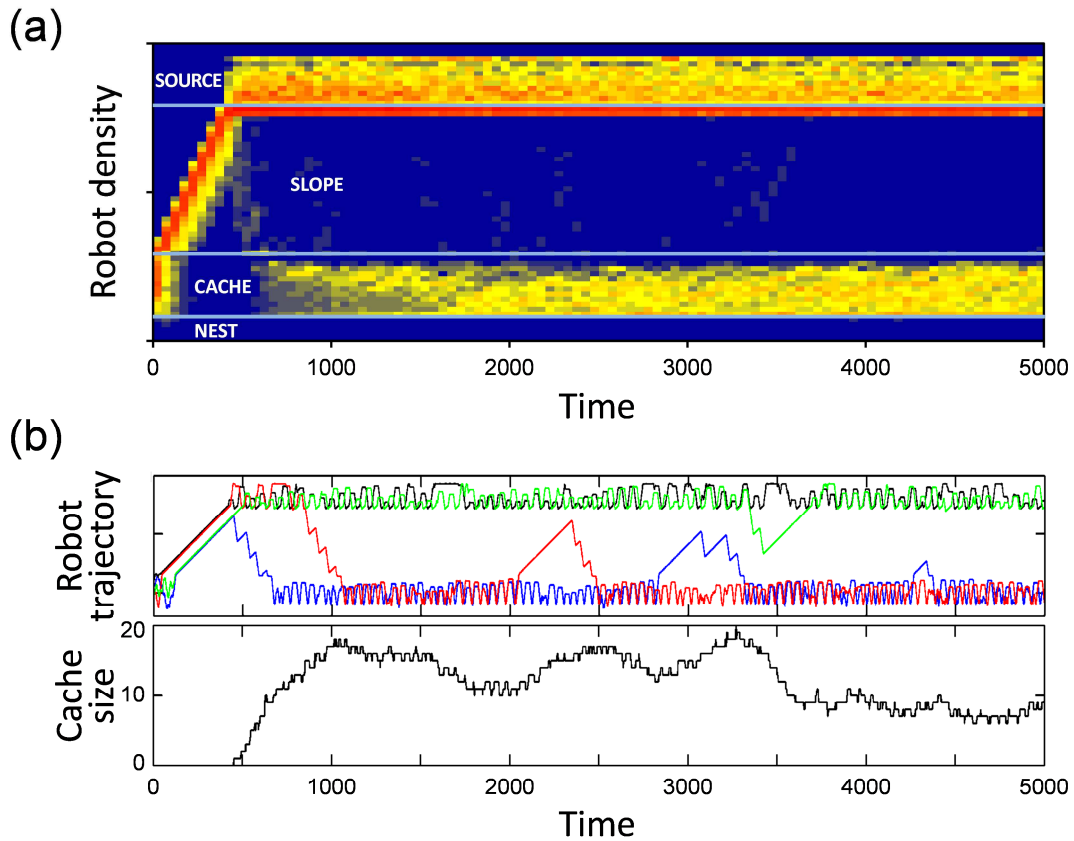
572



573

574 **Figure 4.** Group performance and degree of task specialization displayed by 4 robot teams
 575 over subsequent generations for each of the 22 evolutionary runs. The degree of task
 576 specialization (Y axis) is measured as the proportion of items retrieved by more than one
 577 robot (task-partitioned) over the total number of items retrieved. The group fitness (color-
 578 coded) is the total number of items retrieved to the nest over a period of 5,000 simulated
 579 seconds averaged over 2 simulation runs. The degree of task specialization and the group
 580 fitness of the best evolved controller in each generation is shown over subsequent
 581 generations for each of the 22 evolutionary runs. High task partitioning was evolutionarily
 582 stable, since any transition to high task partitioning never reverted back to generalist foraging
 583 in later generations. Some controllers, however, did not evolve task partitioning as a result of
 584 being trapped in local optima.


585



586

587 **Figure 5.** Self-organized task specialization and task allocation displayed by a controller
 588 evolved from first principles using Grammatical Evolution (cf. Video S3 and evolved rules
 589 shown in Table S2). **(a)** Robot densities in the experimental arena of as a function of time
 590 (average of 30 runs). Despite having identical controllers, robots segregate quickly between
 591 the source and cache areas, thereby avoiding the costly traverse of the slope. **(b)** Robot
 592 trajectory on the arena and cache size in a typical evaluation run. All robots first move to the
 593 source to collect items, but after 500-1000 s into the simulation, the robot teams self-organize
 594 to have two droppers pushing items off the slope and two robots collecting items from the
 595 cache, without these tasks having been explicitly rewarded during the evolutionary runs.

596



597

598 **Figure 6.** The effect of the degree of task specialization (Y axis, proportion of items retrieved
 599 through the action of multiple robots) and average linear speed (absolute average linear
 600 speed of the robots along the long axis of the arena as a percentage of the theoretical
 601 maximum speed) on the fitness performance of the 22 controllers evolved from first principles.
 602 A high degree of task partitioning and high speed significantly increased group fitness (color
 603 code, multiple regression analysis: $p < 0.01$ and $p < 10^{-5}$; color gradient represents the best-fit
 604 plane, average of 30 runs).

605

606

607 **Legends Supplemental Material**

608 **Figure S1.** The different types of dynamics displayed by all 22 controllers evolved from first
609 principles using Grammatical Evolution (cf. Videos S4 and evolved rules shown in Table S2).
610 The figures are ordered based on performance, from the best to the worst. (a) Robot densities
611 in the experimental arena as a function of time (average of 30 runs). (b) Robot trajectory on
612 the arena and cache size in a typical evaluation run.

613

614 **Table S1.** Rules used to encode the dropper, collector and generalist foraging strategies in
615 the experiments with pre-adapted building blocks. Most of the rules are used by more than
616 one behavioral building block (rules R1 and R4-R6 are used by droppers, rules R2-R3, R5
617 and R7-R8 are used by collectors and rules R1, R4-R5 and R7-R8 are used by generalists).
618 For each rule: the first row contains the list of preconditions, each denoted by the syntax
619 $P_{NAME} = True|False$ where $NAME$ is the intuitive name of the precondition; the second row
620 contains the list of fine-grained behavioral building blocks (B_{RANDOM_WALK} , $B_{PHOTOTAXIS}$, B_{ANTI-}
621 $PHOTOTAXIS}$, c.f. Materials and Methods); the remaining rows contain the list of actions (one per
622 row), where the first column indicates the type of the action (A_B are actions that change the
623 currently-executed behavior, while A_{IS} are all other actions), the second column indicates the
624 execution probability, and the third column indicates the effect of the action (either the new
625 behavior to switch to in case of A_B or the new value of the internal state IS_{NAME} in case of A_{IS}).
626 Memory states were set as follows: $P_{STAY_DOWN} = True$ and $P_{STAY_UP} = False$ for collectors,
627 $P_{STAY_DOWN} = False$ and $P_{STAY_UP} = True$ for droppers and $P_{STAY_DOWN} = False$ and $P_{STAY_UP} = False$ for
628 generalists.

629

630 **Table S2.** Rules evolved via Grammatical Evolution in the 22 evolutionary runs. Controllers
631 are sorted from high to low group performance.

632

633 **Table S3.** Performance of the 22 evolved controllers and degree of task partitioning observed
634 in the 4 robot teams and in the 20 robot ones used during post-validation. Controllers are
635 sorted from high to low group performance.

636

637 **Supplemental Video S1.** Video of the optimal behavior displayed by the controller with pre-
638 adapted building blocks in the sloped environment. In this case, an allocation of 50%
639 droppers and 50% collectors resulted in maximal group performance.

640

641 **Supplemental Video S2.** Video of the optimal behavior displayed by the controller with pre-
642 adapted building blocks in the flat environment. In this case, an allocation of 100% generalist
643 foragers resulted in maximal group performance.

644

645 **Supplemental Video S3.** Example of task partitioning behavior evolved during evolutionary
646 run number 2. From initially random behavior, the robots first evolve generalist foraging after
647 500 generations. Subsequently, after 500 more generations, the robots evolve task
648 partitioning, which gets further perfected over the following 1000 generations. We conclude by
649 showing how the controller evolved in a 4 robot team scaled up when tested in a swarm of 20
650 robots. The full HD video is available at <https://www.youtube.com/watch?v=8mIHxcCNzjg>.

651

652 **Supplemental Videos S4.1-S4.22.** Behavior displayed by the 22 evolved controllers. Videos
653 are sorted from high to low group performance.

654

655 **References**

- 656 1. Maynard Smith J, Szathmary E (1997) The major transitions in evolution: Oxford University
657 Press, USA.
- 658 2. Kirk DL (2005) A twelve-step program for evolving multicellularity and a division of labor.
659 Bioessays 27: 299-310.
- 660 3. Bourke AFG (2011) Principles of social evolution. Oxford ; New York: Oxford University
661 Press. 267 p.
- 662 4. Duarte A, Weissing FJ, Pen I, Keller L (2011) An Evolutionary Perspective on Self-
663 Organized Division of Labor in Social Insects. Annu Rev Ecol Evol Syst 42: 91-110.
- 664 5. Holldobler B, Wilson EO (1990) The Ants: Springer.
- 665 6. Holldobler B, Wilson E (2009) The superorganism: the beauty, elegance, and strangeness
666 of insect societies: WW Norton & Company.

- 667 7. Ratnieks FLW, Anderson C (1999) Task partitioning in insect societies. *Insect Soc* 46: 95-
668 108.
- 669 8. Anderson C, Franks NR (2003) Teamwork in animals, robots, and humans. *Adv Stud*
670 *Behav* 33: 1-48.
- 671 9. Duarte A, Pen I, Keller L, Weissing FJ (2012) Evolution of self-organized division of labor in
672 a response threshold model. *Behav Ecol Sociobiol* 66: 947-957.
- 673 10. Beshers SN, Fewell JH (2001) Models of division of labor in social insects. *Annu Rev*
674 *Entomol* 46: 413-440.
- 675 11. Passera L, Roncin E, Kaufmann B, Keller L (1996) Increased soldier production in ant
676 colonies exposed to intraspecific competition. *Nature* 379: 630-631.
- 677 12. Oster GF, Wilson EO (1979) *Caste and ecology in the social insects*: Princeton Univ Pr.
- 678 13. Duarte A, Scholtens E, Weissing FJ (2012) Implications of Behavioral Architecture for the
679 Evolution of Self-Organized Division of Labor. *PLoS Comput Biol* 8: e1002430.
- 680 14. Tarapore D, Floreano D, Keller L (2010) Task-dependent influence of genetic architecture
681 and mating frequency on division of labour in social insect societies. *Behav Ecol*
682 *Sociobiol* 64: 675-684.
- 683 15. Waibel M, Floreano D, Magnenat S, Keller L (2006) Division of labour and colony
684 efficiency in social insects: effects of interactions between genetic architecture,
685 colony kin structure and rate of perturbations. *Proc R Soc B-Biol Sci* 273: 1815-1823.
- 686 16. Lichocki P, Tarapore D, Keller L, Floreano D (2012) Neural Networks as Mechanisms to
687 Regulate Division of Labor. *Am Nat* 179: 391-400.
- 688 17. Page Jr RE, Linksvayer TA, Amdam G (2009) Social life from solitary regulatory networks:
689 a paradigm for insect sociality. In: Gadau J, Fewell J, editors. *Organization of insect*
690 *societies: from genome to sociocomplexity*. Harvard: Harvard University Press. pp.
691 357-367.
- 692 18. Nedelcu AM (2009) Environmentally induced responses co-opted for reproductive
693 altruism. *Biol Lett* 5: 805-808.
- 694 19. Ross L, Gardner A, Hardy N, West SA (2013) Ecology, Not the Genetics of Sex
695 Determination, Determines Who Helps in Eusocial Populations. *Curr Biol* 23: 2383-
696 2387.

- 697 20. Hunt JH, Kensinger BJ, Kossuth JA, Henshaw MT, Norberg K, et al. (2007) A diapause
698 pathway underlies the gyne phenotype in *Polistes* wasps, revealing an evolutionary
699 route to caste-containing insect societies. *Proc Natl Acad Sci U S A* 104: 14020-
700 14025.
- 701 21. Hunt JH, Amdam GV (2005) Bivoltinism as an antecedent to eusociality in the paper wasp
702 genus *Polistes*. *Science* 308: 264-267.
- 703 22. Oldroyd BP, Beekman M (2008) Effects of selection for honey bee worker reproduction on
704 foraging traits. *PLoS Biol* 6: 463-470.
- 705 23. Nolfi S, Floreano D (2000) *Evolutionary robotics: The biology, intelligence, and technology*
706 *of self-organizing machines*: MIT press Cambridge, MA.
- 707 24. Floreano D, Mattiussi C (2008) *Bio-inspired artificial intelligence: theories, methods, and*
708 *technologies*: The MIT Press.
- 709 25. Ray TS (1992) *An Approach to the Synthesis of Life*. In: Langton C, Taylor C, Farmer JD,
710 Rasmussen S, editors. *Artificial Life II*. pp. 371-408.
- 711 26. Floreano D, Mitri S, Magnenat S, Keller L (2007) Evolutionary conditions for the
712 emergence of communication in robots. *Curr Biol* 17: 514-519.
- 713 27. Mitri S, Floreano D, Keller L (2009) The evolution of information suppression in
714 communicating robots with conflicting interests. *Proc Natl Acad Sci U S A* 106:
715 15786.
- 716 28. Waibel M, Keller L, Floreano D (2009) Genetic Team Composition and Level of Selection
717 in the Evolution of Cooperation. *IEEE Trans Evol Comput* 13: 648-660.
- 718 29. Baldassarre G, Nolfi S, Parisi D (2003) Evolving mobile robots able to display collective
719 behaviors. *Artif Life* 9: 255-267.
- 720 30. Trianni V, Groß R, Labella T, Şahin E, Dorigo M (2003) Evolving Aggregation Behaviors in
721 a Swarm of Robots. In: Banzhaf W, Ziegler J, Christaller T, Dittrich P, Kim J, editors.
722 *Advances in Artificial Life*: Springer Berlin Heidelberg. pp. 865-874.
- 723 31. Bahçeci E, Sahin E. Evolving aggregation behaviors for swarm robotic systems: a
724 systematic case study; 2005. pp. 333-340.
- 725 32. Dorigo M, Trianni V, Sahin E, Gross R, Labella TH, et al. (2004) Evolving self-organizing
726 behaviors for a swarm-bot. *Auton Robot* 17: 223-245.

- 727 33. Sperati V, Trianni V, Nolfi S (2011) Self-organised path formation in a swarm of robots.
728 Swarm Intell 5: 97-119.
- 729 34. Harvey I, Di Paolo E, Wood R, Quinn M, Tuci E (2005) Evolutionary robotics: A new
730 scientific tool for studying cognition. Artif Life 11: 79-98.
- 731 35. Floreano D, Lausanne S, Husbands P, Nolfi S, Roma I (2008) Chapter 61: Evolutionary
732 Robotics. Handbook of Robotics. Berlin: Springer. pp. 1423-1451.
- 733 36. Floreano D, Keller L (2010) Evolution of Adaptive Behaviour in Robots by Means of
734 Darwinian Selection. PLoS Biol 8: e1000292.
- 735 37. Hauert S, Mitri S, Keller L, Floreano D (2010) Evolving Cooperation: From Biology to
736 Engineering. The Horizons of Evolutionary Robotics. Cambridge, USA: MIT Press.
- 737 38. Brambilla M, Ferrante E, Birattari M, Dorigo M (2013) Swarm robotics: a review from the
738 swarm engineering perspective. Swarm Intell 7: 1-41.
- 739 39. Francesca G, Brambilla M, Brutschy A, Trianni V, Birattari M (2014) AutoMoDe: A novel
740 approach to the automatic design of control software for robot swarms. Swarm Intell
741 8: 89-112.
- 742 40. Duarte M, Oliveira SM, Christensen AL. Evolution of Hierarchical Controllers for Multirobot
743 Systems; 2014.
- 744 41. Ferrante E, Duéñez-Guzmán E., Turgut AE, Wenseleers T (2013) GESwarm:
745 Grammatical Evolution for the Automatic Synthesis of Collective Behaviors in Swarm
746 Robotics. Proceedings of the 15th International Conference on Genetic and
747 Evolutionary Computation Conference Companion (GECCO 2013). New York: ACM.
748 pp. 17-24.
- 749 42. Fowler HG, Robinson SW (1979) Foraging by *Atta sexdens* (Formicidae: Attini): seasonal
750 patterns, caste and efficiency. Ecol Entomol 4: 239-247.
- 751 43. Hart AG, Anderson C, Ratnieks FL (2002) Task partitioning in leafcutting ants. Acta Ethol
752 5: 1-11.
- 753 44. Bonani M, Longchamp V, Magnenat S, Rétornaz P, Burnier D, et al. (2010) The MarXbot,
754 a miniature mobile robot opening new perspectives for the collective-robotic research.
755 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2010):
756 IEEE Press. pp. 4187-4193

- 757 45. Pinciroli C, Trianni V, O'Grady R, Pini G, Brutschy A, et al. (2012) ARGoS: a Modular,
758 Parallel, Multi-Engine Simulator for Multi-Robot Systems. *Swarm Intell* 6: 271-295.
- 759 46. Röschard J, Roces F (2011) Sequential load transport in grass-cutting ants (*Atta*
760 *vollenweideri*): maximization of plant delivery rate or improved information transfer?
761 *Psyche: A Journal of Entomology* 2011.
- 762 47. Goldsby HJ, Dornhaus A, Kerr B, Ofria C (2012) Task-switching costs promote the
763 evolution of division of labor and shifts in individuality. *Proc Natl Acad Sci U S A* 109:
764 13686-13691.
- 765 48. Hubbell SP, Johnson LK, Stanislav E, Wilson B, Fowler H (1980) Foraging by bucket-
766 brigade in leaf-cutter ants. *Biotropica* 12: 210-213.
- 767 49. Röschard J, Roces F (2003) Cutters, carriers and transport chains: Distance-dependent
768 foraging strategies in the grass-cutting ant *Atta vollenweideri*. *Insect Soc* 50: 237-244.
- 769 50. Ferrante E, Turgut AE, Huepe C, Stranieri A, Pinciroli C, et al. (2012) Self-organized
770 flocking with a mobile robot swarm: a novel motion control method. *Adapt Behav* 20:
771 460-477.
- 772 51. Ferrante E, Turgut AE, Stranieri A, Pinciroli C, Birattari M, et al. (2014) A self-adaptive
773 communication strategy for flocking in stationary and non-stationary environments.
774 *Nat Comput* 13: 225-245.
- 775 52. O'Neill M, Ryan C (2003) Grammatical evolution: evolutionary automatic programming in
776 an arbitrary language. London, UK: Springer.
- 777 53. Soysal O, Bahçeci E, Şahin E (2007) Aggregation in swarm robotic systems: evolution
778 and probabilistic control. *Turk J Electr Eng Co* 15: 199-225.
- 779 54. Labella TH, Dorigo M, Deneubourg JL (2006) Division of Labor in a Group of Robots
780 Inspired by Ants' Foraging Behavior. *ACM Trans Auton Adapt Syst* 1: 4-25.
- 781 55. Jeanson R, Rivault C, Deneubourg JL, Blanco S, Fournier R, et al. (2005) Self-organized
782 aggregation in cockroaches. *Anim Behav* 69: 169-180.
- 783 56. Theraulaz G, Bonabeau E, Deneubourg JL (1998) Response threshold reinforcement and
784 division of labour in insect societies. *Proc R Soc B-Biol Sci* 265: 327-332.
- 785 57. Wirth N (1977) What can we do about the unnecessary diversity of notation for syntactic
786 definitions? *Commun ACM* 20: 822-823.

- 787 58. Trianni V, Nolfi S (2011) Engineering the Evolution of Self-Organizing Behaviors in Swarm
788 Robotics: A Case Study. *Artif Life* 17: 183-202.
- 789 59. Buss LW (1987) The evolution of individuality. Princeton, N.J.: Princeton University Press.
790 201 p.
- 791 60. Ravary F, Lecoutey E, Kaminski G, Chaline N, Jaisson P (2007) Individual experience
792 alone can generate lasting division of labor in ants. *Curr Biol* 17: 1308-1312.
- 793 61. Jeanson R, Weidenmüller A (2014) Interindividual variability in social insects – proximate
794 causes and ultimate consequences. *Biol Rev* 89: 671-687.
- 795 62. Rueffler C, Hermisson J, Wagner GP (2012) Evolution of functional specialization and
796 division of labor. *Proc Natl Acad Sci U S A* 109: E326-E335.
- 797 63. Ravary F, Lecoutey E, Kaminski G, Châline N, Jaisson P (2007) Individual experience
798 alone can generate lasting division of labor in ants. *Curr Biol* 17: 1308-1312.
- 799 64. Kurakin A (2005) Self-organization vs Watchmaker: stochastic gene expression and cell
800 differentiation. *Dev Genes Evol* 215: 46-52.
- 801 65. Veening J-W, Smits WK, Kuipers OP (2008) Bistability, epigenetics, and bet-hedging in
802 bacteria. *Annu Rev Microbiol* 62: 193-210.
- 803 66. Roces F, Holldobler B (1995) Vibrational communication between hitchhikers and
804 foragers in leaf-cutting ants (*Atta Cephalotes*). *Behav Ecol Sociobiol* 37: 297-302.
- 805 67. Roces F, Holldobler B (1996) Use of stridulation in foraging leaf cutting ants: Mechanical
806 support during cutting or short range recruitment signal? *Behav Ecol Sociobiol* 39:
807 293-299.
- 808 68. Seeley TD (1995) The wisdom of the hive: The social physiology of honey bee colonies:
809 Harvard University Press, Kittridge Hall, 79 Garden St., Cambridge, Massachusetts
810 02138-9904, USA; Harvard University Press, London, England, UK. 295 p.
- 811 69. Seeley TD, Kuhnholz S, Weidenmuller A (1996) The honey bee's tremble dance
812 stimulates additional bees to function as nectar receivers. *Behav Ecol Sociobiol* 39:
813 419-427.
- 814 70. Goldsby H, Knoester D, Ofria C (2010) Evolution of Division of Labor in Genetically
815 Homogenous Groups.

- 816 71. Goldsby HJ, Knoester DB, Kerr B, Ofria C (2014) The Effect of Conflicting Pressures on
817 the Evolution of Division of Labor. *PLoS One* 9: e102713.
- 818 72. Gavrillets S (2010) Rapid transition towards the division of labor via evolution of
819 developmental plasticity. *PLoS Comput Biol* 6.
- 820 73. Nedelcu AM, Michod RE (2006) The Evolutionary Origin of an Altruistic Gene. *Mol Biol*
821 *Evol* 23: 1460-1464.
- 822 74. Forrest SaM, M. (1996) Relative Building-Block Fitness and the Building-Block
823 Hypothesis. In: Whitley D, editor. *Foundations of Genetic Algorithms 2: Morgan*
824 *Kaufmann*.
- 825 75. Bongard JC (2011) Morphological and environmental scaffolding synergize when evolving
826 robot controllers: artificial life/robotics/evolvable hardware. *Proceedings of the 13th*
827 *annual conference on Genetic and evolutionary computation*. Dublin, Ireland: ACM.
828 pp. 179-186.
- 829 76. Lenski RE, Ofria C, Pennock RT, Adami C (2003) The evolutionary origin of complex
830 features. *Nature* 423: 139-144.
- 831 77. Matarić M, Cliff D (1996) Challenges in evolving controllers for physical robots. *Robot*
832 *Auton Syst* 19: 67-83.
- 833 78. Tuci E (2014) *Evolutionary Swarm Robotics: Genetic Diversity, Task-Allocation and Task-*
834 *Switching*. *Swarm Intell: Springer*. pp. 98-109.
- 835 79. Trueba P, Prieto A, Caamaño P, Bellas F, Duro RJ (2011) Task-Driven Species in
836 Evolutionary Robotic Teams. In: Ferrández J, Álvarez Sánchez J, de la Paz F, Toledo
837 FJ, editors. *Foundations on Natural and Artificial Computation: Springer Berlin*
838 *Heidelberg*. pp. 138-147.
- 839 80. Nitschke GS, Schut MC, Eiben AE (2012) Evolving behavioral specialization in robot
840 teams to solve a collective construction task. *Swarm and Evolutionary Computation 2:*
841 25-38.
- 842 81. Nitschke GS, Eiben AE, Schut MC (2012) Evolving team behaviors with specialization.
843 *Genet Program Evol Mach* 13: 493-536.
- 844 82. Krieger MJB, Billeter JB, Keller L (2000) Ant-like task allocation and recruitment in
845 cooperative robots. *Nature* 406: 992-995.

- 846 83. Lerman K, Jones C, Galstyan A, Matarić MJ (2006) Analysis of dynamic task allocation in
847 multi-robot systems. *Int J Rob Res* 25: 225-241.
- 848 84. Khaluf Y, Birattari M, Hamann H (2014) A Swarm Robotics Approach to Task Allocation
849 under Soft Deadlines and Negligible Switching Costs. *From Animals to Animats 13*:
850 Springer. pp. 270-279.
- 851 85. Brutschy A, Pini G, Pinciroli C, Birattari M, Dorigo M (2014) Self-organized task allocation
852 to sequentially interdependent tasks in swarm robotics. *Auton Agents Multi-Agent*
853 *Syst* 28: 101-125.
- 854 86. Pini G, Brutschy A, Frison M, Roli A, Dorigo M, et al. (2011) Task partitioning in swarms of
855 robots: an adaptive method for strategy selection. *Swarm Intell* 5: 283-304.
- 856 87. Pini G, Brutschy A, Pinciroli C, Dorigo M, Birattari M (2013) Autonomous task partitioning
857 in robot foraging: an approach based on cost estimation. *Adapt Behav* 21: 118-136.
- 858 88. Pini G, Brutschy A, Scheidler A, Dorigo M, Birattari M (2014) Task Partitioning in a Robot
859 Swarm: Object Retrieval as a Sequence of Subtasks with Direct Object Transfer. *Artif*
860 *Life* 20: 291-317.
- 861 89. Gerkey BP, Matarić MJ (2004) A formal analysis and taxonomy of task allocation in multi-
862 robot systems. *Int J Rob Res* 23: 939-954.
- 863