

# Beyond pheromones: evolving error-tolerant, flexible, and scalable ant-inspired robot swarms

Written by Joshua P. Hecker and Melanie E. Moses

Presented by Nitin Bhandari, Antonio Griego, Jacob  
McCullough, and Noah Lewis

# Topics to be covered

- Introduction and Related Work
- Methods
- Results
- Discussion

# Introduction

# What is swarm robotics?



- Swarm robotics is an approach to the coordination of multiple robots as a system which consist of large numbers of mostly simple physical robots.
- A collective behaviour emerges among the robots via interactions among themselves and with the environment
- We don't have to make individual robots more intelligent but make them capable to forming a collective intelligent behaviour


# Desert Harvester Ants



- Emulate their foraging strategies.
- They have evolved to collect many seeds as quickly as possible without exhaustively collecting all.
- They use site fidelity and pheromones

# iAnt Robot

(a)



- Made up of inexpensive components.
- Can be multiplied to produce swarms.
- Are robust to communication errors.
- Inexpensive components leads to increased sensor errors and a higher likelihood of hardware failure.

# Foraging Strategies via robots

1. Testing was done with iAnt robots.
2. Robot behaviours were specified by central-placed foraging algorithm (CPFA), that mimics the behaviours of seed-harvester ants.
3. The performance of CPFA was optimised using GA by evolving the movement, sensing and communication with the help of environment evaluation.


With this we are not just evaluating ant behaviour of foraging, but also the evolutionary process that combines these behaviours into integrated strategies.

- Ants collect seeds more quickly when the seeds were clustered.
- Foraging in heterogeneous clustering requires more complex communication, memory and environmental sensing strategies which are the common problems faced by animals in natural environment.
- Evolutionary Robotics
- Using neural networks
- Genetic algorithms and reinforcement learning for switching behaviours in robots



# Insights

1. Success of a foraging strategy depends strongly on spatial distribution of resources that are being collected.
2. Site fidelity and pheromones are critical components for foraging strategies when resources are clustered.




# Methods

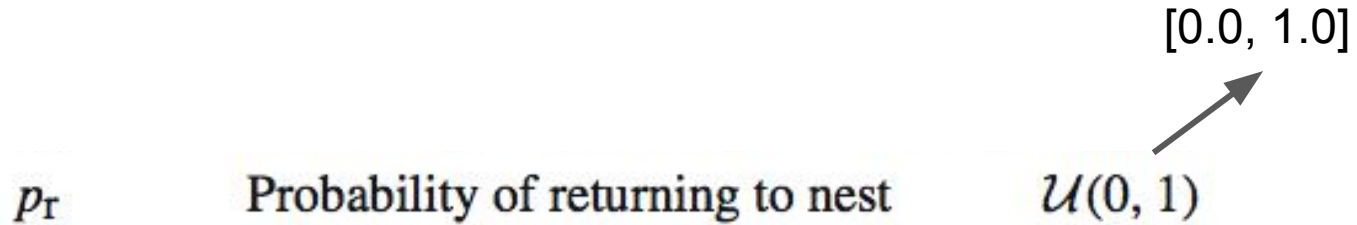
1. CPFA Parameters
2. CPFA Algorithm
3. Genetic Algorithm
4. Experimental Setup
5. Measuring Performance

# Central Place Foraging Algorithm Parameters

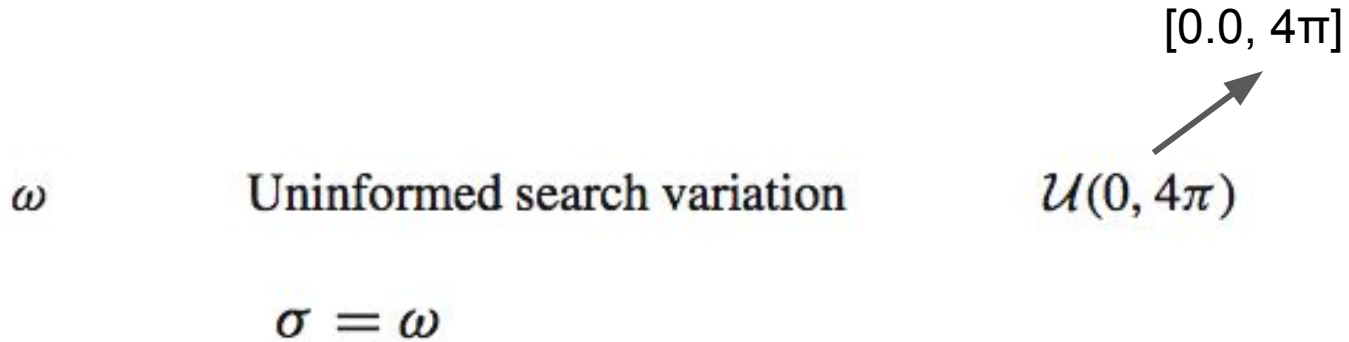
Parameter	Description	Initialization function
$p_s$	Probability of switching to searching	$\mathcal{U}(0, 1)$
$p_r$	Probability of returning to nest	$\mathcal{U}(0, 1)$
$\omega$	Uninformed search variation	$\mathcal{U}(0, 4\pi)$
$\lambda_{id}$	Rate of informed search decay	$\exp(5)$
$\lambda_{sf}$	Rate of site fidelity	$\mathcal{U}(0, 20)$
$\lambda_{lp}$	Rate of laying pheromone	$\mathcal{U}(0, 20)$
$\lambda_{pd}$	Rate of pheromone decay	$\exp(10)$

$p_s$  Probability of switching to searching  $\mathcal{U}(0, 1)$   [0.0, 1.0]

As a robot moves to a search location, it may give up traveling and instead begin searching from its current location. This parameter short circuits absurdly long trips to found resources in the hopes of discovering something closer.



Robots that are currently searching for resources may give up their search and return to the nest. This gives them the chance to follow pheromones or return to a previous site fidelity location.



When uninformed, robots travel by (1) randomly selecting a turning angle in the range  $[0, \omega]$ , (2) turning, and (3) moving a fixed step size. Low values of  $\omega$  produce straighter paths that cover long distance versus high values of  $\omega$  that produce sharp turns that exhaustively search a local region.

$\lambda_{id}$ 

Rate of informed search decay

 $\exp(5)$  $[0.0, e^{5.0}]$   


$$\sigma = \omega + (4\pi - \omega)e^{-\lambda_{id}t}$$

When informed, robots search a local area thoroughly by making sharper turns in between travel steps. That is,  $\omega$  is temporarily increased in value and decays to its original value over time. This parameter throttles the speed of the decay.



$\lambda_{sf}$ **Rate of site fidelity**


$$\text{POIS}(k, \lambda) = e^{-\lambda} \sum_{i=0}^{\lfloor k \rfloor} \frac{\lambda^i}{i!}$$

 $\mathcal{U}(0, 20)$  $[0.0, 20.0]$   


After finding a resource and returning to the nest, a robot may return to the last location it found a resource with the probability defined by this parameter in the Poisson CDF where  $\lambda$  = the rate of site fidelity and  $k$  = the resource density: the number of resources detected by the robot when it discovered a resource and scanned the immediate area.

$\lambda_{lp}$ 

Rate of laying pheromone

 $\mathcal{U}(0, 20)$   
  
[0.0, 20.0]

$$\text{POIS}(k, \lambda) = e^{-\lambda} \sum_{i=0}^{\lfloor k \rfloor} \frac{\lambda^i}{i!}$$

After finding a resource and returning to the nest, a robot may lay a pheromone with the probability defined by this parameter used in the Poisson CDF where  $\lambda$  = the rate of laying a pheromone and  $k$  = the resource density. In other words, we calculate the likelihood of finding at least  $k$  additional resources.

$\lambda_{pd}$ 

Rate of pheromone decay

$$\gamma = e^{-\lambda_{pd}t}$$

exp(10)

[0.0, e^10.0]



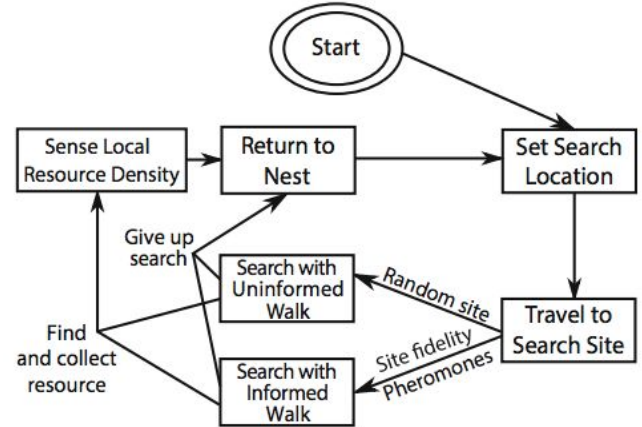
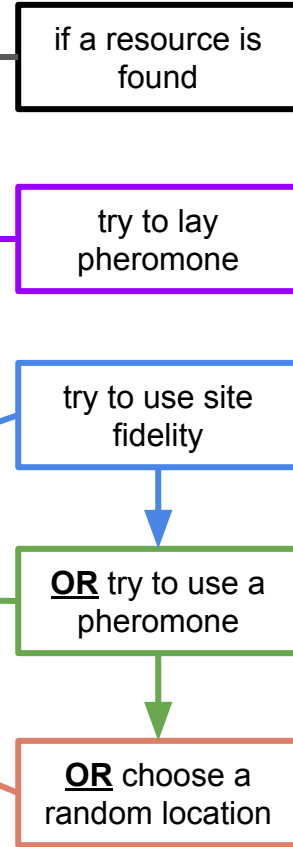
When a robot “lays” a pheromone, it produces an (x, y) coordinate point with an associated weight value and stores it in a list. This weight (defined as  $\gamma$ ) decays over time, and the rate of this decay is throttled by this parameter. Pheromones with a weight below the threshold (defined as 0.001) are deleted.

# What the heck is a digital pheromone?

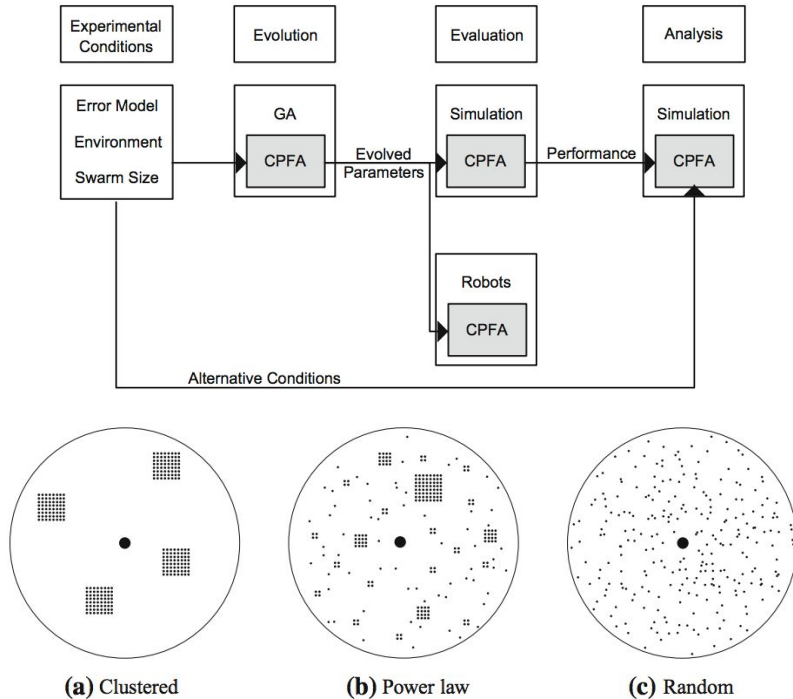
1. a robot finds a resource
2. the robot records its position:  $P = (x, y)$  and counts the number of other local resources by spinning in a circle and observing the immediate area
3.  $P$  = the site fidelity waypoint a robot will return to if it uses site fidelity
4.  $P$  also = the pheromone waypoint location shared with the swarm
  - a. if a robot follows a pheromone, it navigates to this  $(x, y)$  position
  - b. pheromones decay over time and are eventually deleted
  - c. position recording and physical navigation is a NOISY process
5. how a robot chooses a specific pheromone is not explicitly defined

## Algorithm 1 Central-Place Foraging Algorithm

- 1: Disperse from nest to random location
- 2: **while** experiment running **do**
- 3: Conduct uninformed correlated random walk
- 4: **if** resource found **then**
- 5: Collect resource
- 6: Count number of resources  $c$  near current location  $l_f$
- 7: Return to nest with resource
- 8: **if**  $\text{POIS}(c, \lambda_{lp}) > U(0, 1)$  **then**
- 9: Lay pheromone to  $l_f$
- 10: **end if**
- 11: **if**  $\text{POIS}(c, \lambda_{sf}) > U(0, 1)$  **then**
- 12: Return to  $l_f$
- 13: Conduct informed correlated random walk
- 14: **else if** pheromone found **then**
- 15: Travel to pheromone location  $l_p$
- 16: Conduct informed correlated random walk
- 17: **else**
- 18: Choose new random location
- 19: **end if**
- 20: **end if**
- 21: **end while**



# Genetic Algorithm



- GA fitness = total number of resources collected in a finite time period
- the GA evolves the 7 CPFA parameter values for three types of distributions:
  - clustered
  - power law
  - random

# Genetic Algorithm

- 1 evolutionary process = 100 simulated robot swarms run for 100 generations
  - 1 generation = 8 simulation runs
  - 10 evolutionary processes are run in total
- gene values: CPFA parameters
- recombination (AKA crossover): combine the gene values of two parents to produce new offspring
- mutation: altering one more more gene values from its initial state
- tournament selection: choosing the best gene value set by running the same experiments (I.E., the same resource distributions) for each gene set and choosing the best
- elitism: copy the best gene value set, unaltered, into the next generation

# Experimental Setup

## physical experiments

- run for 1 hour
- 100 square meter arena
- 256 resources
- lamp beacon for finding the nest
- robots transmit position data one-way over WiFi
- a central server saves, updates, and shares pheromones

## simulated experiments

- run for 1 simulated hour
- 125 x 125 cellular grid, each cell representing an 8 x 8 cm square
- 256 resources
- no collisions
- simulated sensor and localization error



# Measuring Performance

Efficiency is the total number of resources collected in a fixed (1 hour) time period.

## Error Tolerance

$$\frac{E_2 - E_1}{E_1} \times 100 \%$$

E1 is the efficiency of a strategy evolved assuming no error.

E2 is the efficiency of a strategy evolved in the presence of error.

Measures how well robots mitigate the effects of error inherent in hardware (or simulated error).

## Flexibility

$$\frac{E_2}{E_1} \times 100 \%$$

E1 is the efficiency of the BEST strategy evolved for a given resource distribution.

E2 is the efficiency of an ALTERNATIVE strategy evolved for a different resource distribution tested on E1's resource distribution.

Scalability uses this formula and measures the number of robots instead of strategies.

# Results

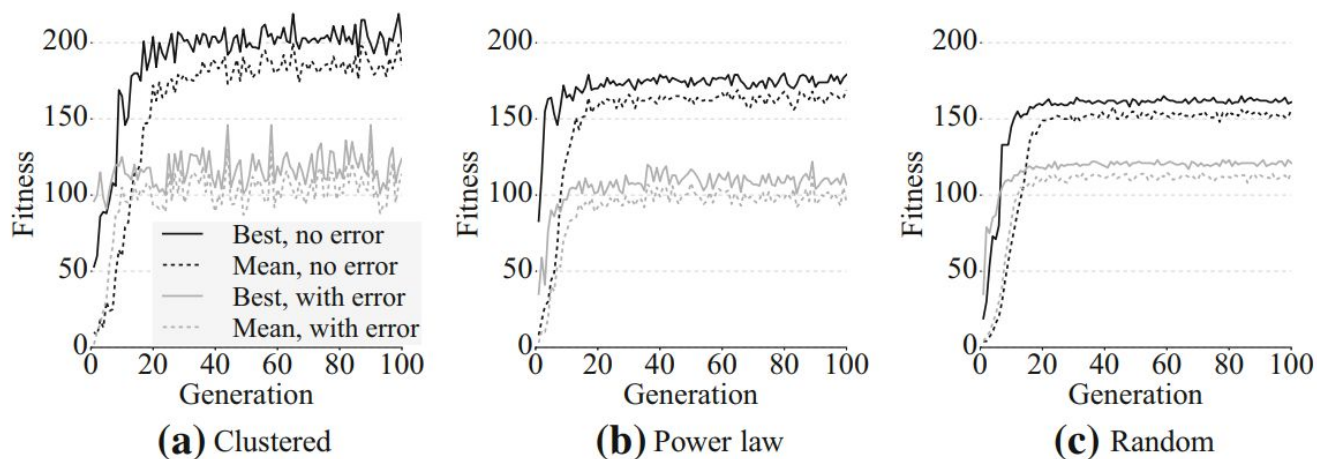
1. Error Tolerance
2. Flexibility
3. Scalability

# Error Tolerance

Does introducing error to the world affect the efficiency of an evolved foraging strategy?

It is interesting to note that after approximately 20 generations, the fitness stabilizes for all three distributions.

This shows that robots with error are always less efficient than robots without error.



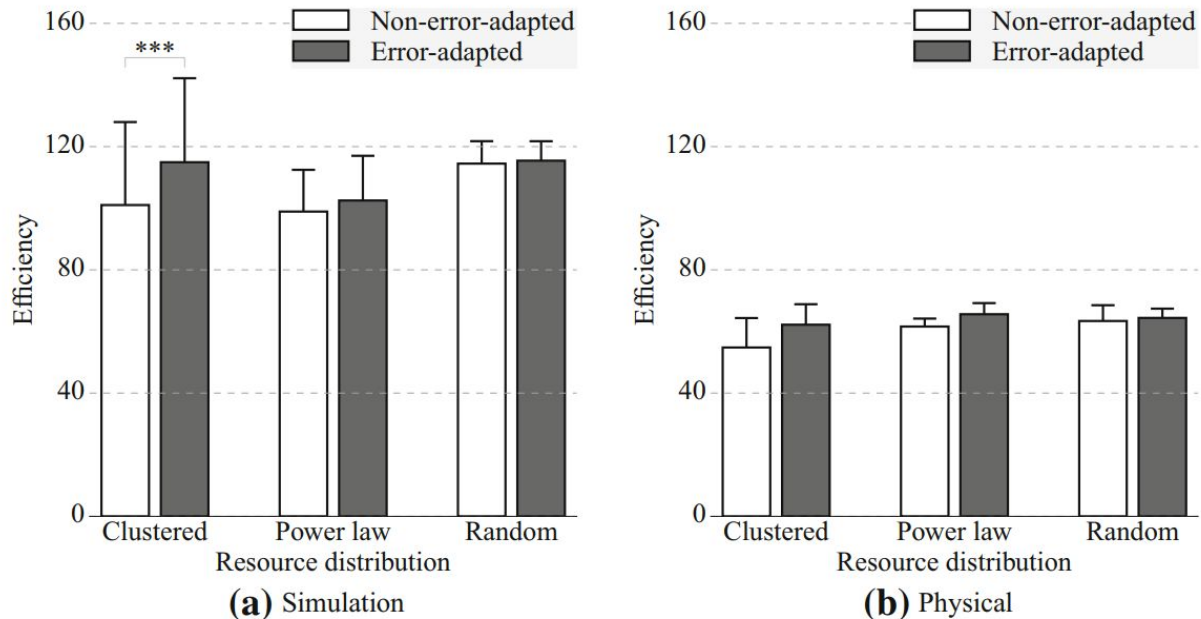
**Fig. 5** Best and mean fitness, measured as foraging efficiency (resources collected per hour, per swarm) for simulated swarms foraging on **a** clustered, **b** power law, and **c** random resource distributions with and without real-world sensor error. Results are for 100 replicates

# Error Tolerance

Adapting to error allows for an increase in efficiency.

Error-adapted swarms actually outperform non-error-adapted swarms on the clustered and power law distributions.

Random distributions did not see a significant statistical change.



**Fig. 6** Foraging efficiency (resources collected per hour, per swarm) using error-adapted and non-error-adapted parameters for **a** 6 robots foraging in a simulation that includes sensor error and **b** 6 physical robots. Asterisks indicate a statistically significant difference ( $p < 0.001$ )

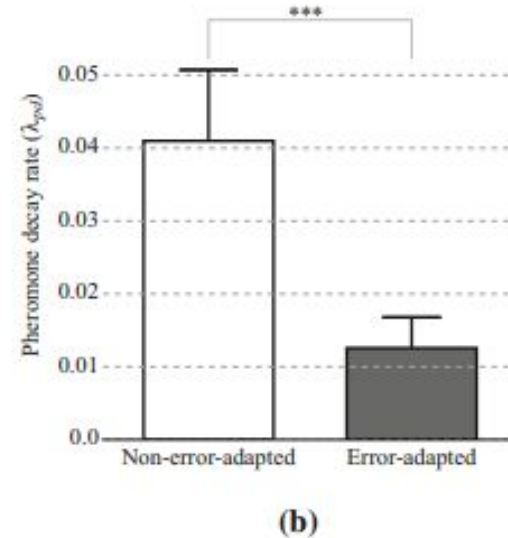
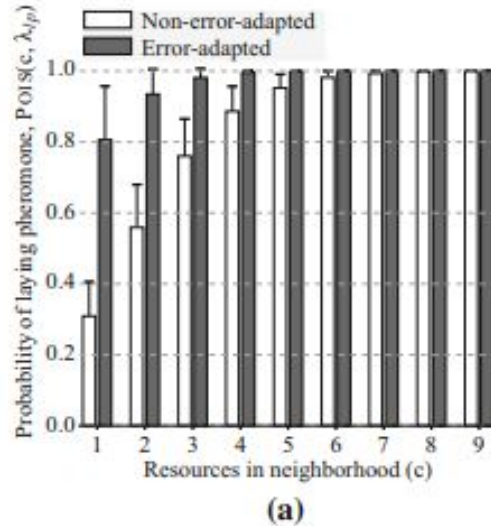
# Error Tolerance

The individual robot's sensor errors are compensated for by the evolved strategy.

This results in a significantly higher probability that pheromones are used at lower values of  $c$ .

A small number of detected tags indicates the presence of nearby undetected tags.

Another form of compensation is to lower the rate of pheromone decay.



**Fig. 7** For error-adapted and non-error-adapted swarms foraging on clustered resources, **a** the probability of laying pheromone as a function of the count  $c$  of resources in the neighborhood of the most recently found resource (Eq. 4:  $k \leftarrow c, \lambda \leftarrow \lambda_{ip}$ ), and **b** the pheromone waypoint decay rate ( $\lambda_{pd}$ ). Asterisks indicate a statistically significant difference ( $p < 0.001$ )

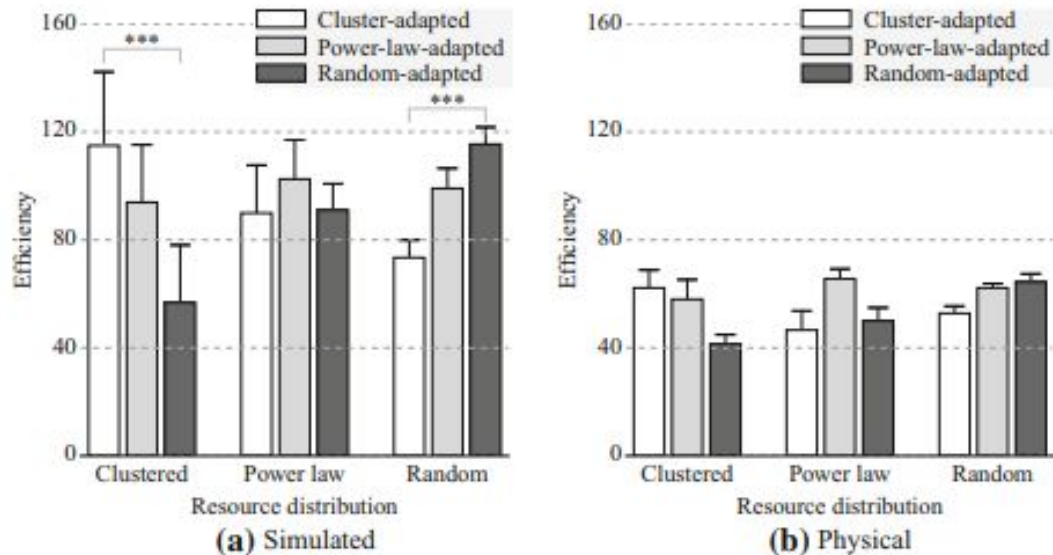
# Flexibility

As expected, each evolved strategy is best at its own type of distribution.

Both specialist and generalist strategies are evolved.

The power-law-adapted strategy is sufficiently flexible on both of the other distributions.

If the distribution of resources is known a priori, a swarm would use a specialist strategy. Otherwise, it should use the most general strategy.



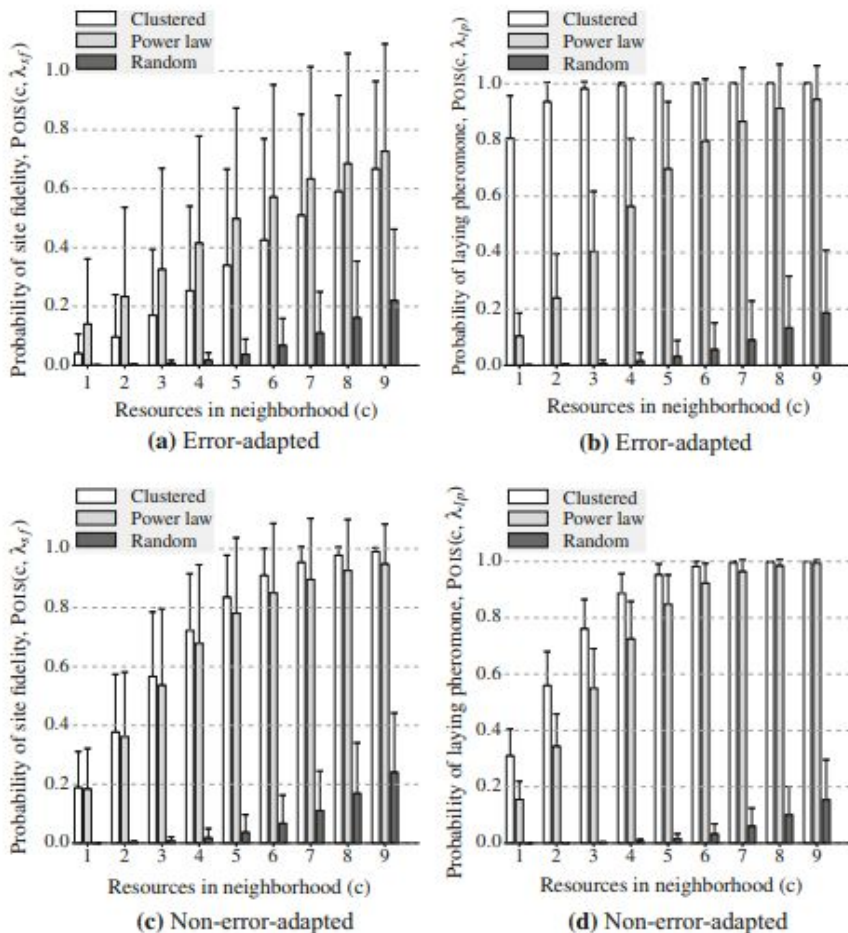
**Fig. 8** Foraging efficiency (resources collected per hour, per swarm) using parameters adapted to different resource distributions for **a** 6 robots foraging in a simulation that includes sensor error and **b** 6 physical robots. Asterisks indicate a statistically significant difference ( $p < 0.001$ )

# Flexibility

Each evolved strategy has tuned its parameters in ways that are to be expected.

In clustered distributions, it makes sense that pheromones are more likely to be laid.

The power-law-adapted strategy shows the most variation, which mimics the variation in resource pile size.



**Fig. 9** For error-adapted swarms (top) and non-error-adapted swarms (bottom), **a, c** the probability of returning to a site (Eq. 4:  $k \leftarrow c, \lambda \leftarrow \lambda_{sf}$ ) and **b, d** the probability of laying pheromone (Eq. 4:  $k \leftarrow c, \lambda \leftarrow \lambda_{lp}$ ) given the number of resources  $c$  in the neighborhood of a found resource


# Scalability

Swarm efficiency increases as swarm size increases.

However, individual robot efficiency decreases as swarm size increases.

The simulated swarms increasingly overestimate swarm efficiency.

This could be from inter-robot interference that is introduced in the physical tests. However, researchers found that collisions are not a cause for the overestimation.



**Fig. 10** Foraging efficiency (resources collected per hour, per robot) of 1, 3, and 6 robots foraging on a power law distribution for **a** swarms in a simulation that includes sensor error and **b** physical swarms. All results are statistically different ( $p < 0.001$ )



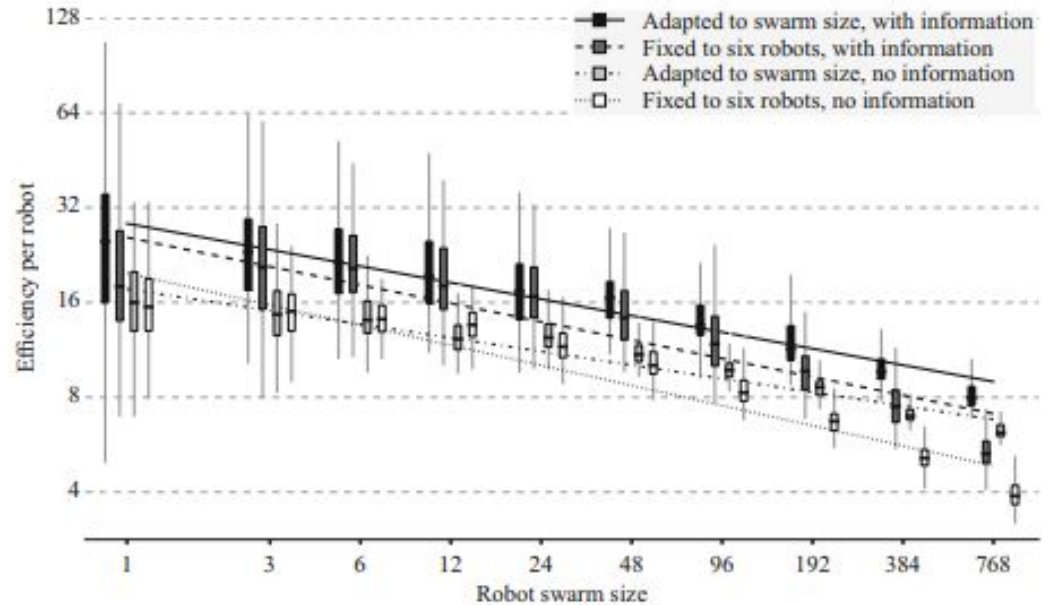
# Scalability

Robots that are part of a larger swarm have to travel greater distances to collect more resources, reducing individual efficiency.

The GA compensates for the reduced individual efficiency.

This compensation is what allows the swarm to gain back some of the lost individual efficiency.

The use of information increases efficiency.



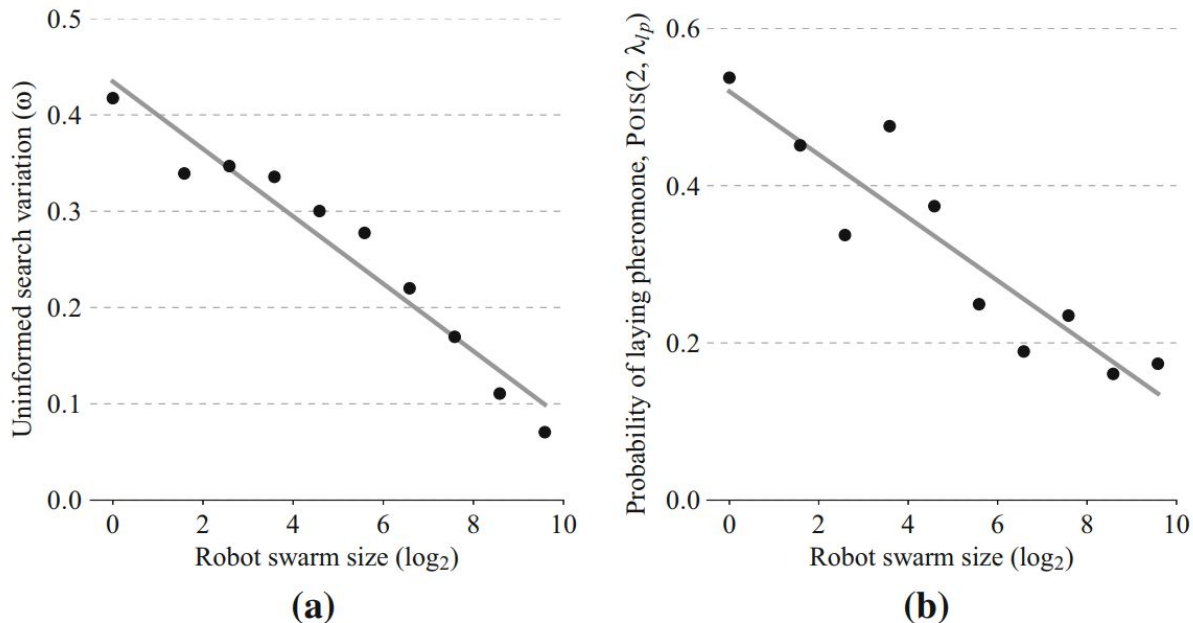
**Fig. 11** Foraging efficiency (resources collected per hour, per robot) in simulated swarms of 1–768 robots foraging without sensor error. Data are shown on a log scale, and linear regression lines are shown for log-transformed data. Per-robot efficiency is shown for four cases: using the full CPFA parameter set adapted to swarm size (slope =  $-0.17$ ,  $R^2 = 0.96$ ), using the full CPFA with parameters fixed to values evolved for a swarm size of 6 (slope =  $-0.19$ ,  $R^2 = 0.83$ ), using parameters adapted to swarm size without information (i.e., the CPFA without memory and communication; slope =  $-0.14$ ,  $R^2 = 0.95$ ), and using parameters fixed to values evolved for a swarm size of 6 without information (slope =  $-0.21$ ,  $R^2 = 0.91$ ). All linear fits are statistically significant ( $p < 0.001$ )

# Scalability

As swarm size increases, the variation in uninformed searches decreases. This makes up for the increases possibility of more robots being nearby.

Pheromone usage is also related to swarm size. Bigger swarms rely less on pheromones.

This prevents over-exploitation of resources by not recruiting too many robots to harvest a single resource pile.



**Fig. 12** **a** Swarm size versus best evolved uninformed search variation ( $\omega$ ) (slope =  $-0.035$ ,  $R^2 = 0.94$ ,  $p < 0.001$ ) (see Fig. 14 in supplementary material for statistical distribution). **b** Swarm size versus best evolved probability of laying pheromone when two resources are found in the resource neighborhood (Eq. 4:  $k \leftarrow 2$ ,  $\lambda \leftarrow \lambda_{lp}$ ) (slope =  $-0.040$ ,  $R^2 = 0.84$ ,  $p < 0.001$ ) (see Fig. 14, supplementary material)

# Discussion

1. Conclusion
2. Interpretation
3. Going Forward


# Summary

The system evolved successfully to collectively adapt to a variety of conditions using only 7 parameters and, importantly, by selecting individual behaviors.


Specifically, the importance of pheromone communication was sensitive to navigation and sensing error, resource distribution, and swarm size.

Condition-specific evolution produced the highest rate of foraging efficiency, whereas condition-specific evolutions applied to different distributions showed reduced efficiency in all cases.

The focus on optimizing combinations of parameters for the GAs was not only effective, but also mirrored natural evolution.



# Interpretation



This system allows for new insight into how memory, communication, and movement change based on foraging conditions, which is not available to experiments under natural environments.

In these experiments, individuals within smaller swarms were more likely to lay pheromones than those in larger swarms. This conflicts with current hypotheses that communication is positively correlated with colony size. This is possibly due to what kinds of environments are preferred by collectives of different sizes.

# Going Forward

This system could be used to test current biological hypotheses, or generate new ones.

Possibly, this system could be used to test the balance between communication and memory for different resource distributions.

The relationship between communication and colony size is also an avenue for future study, due to the conflicts between this research and current hypotheses.

