Agent-Based Modeling

A Third Way of Doing Science

Two traditional ways of doing science

- Induction: inferring from particular data a general theory
- Deduction: reasoning from first principles to a general theory

Third Way

- Generative: using first principles to generate a particular set of data that can create a general theory

Ant Corpse Piles

- How does it work?
- Is it swarm intelligence?
- What is the simplest model able to explain the process?

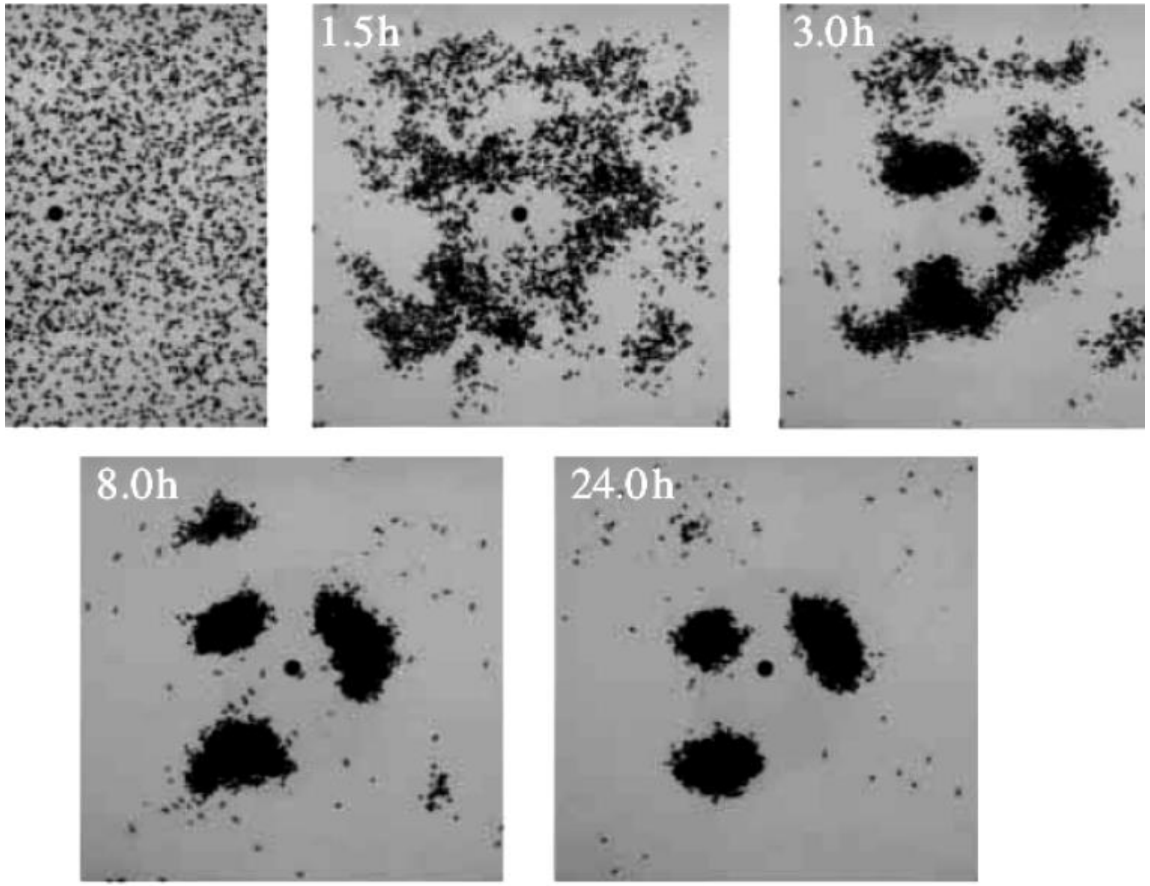

Jost *et al.*, J. R. Soc. Interface, 2007

Agent-Based Models (ABMs)

Modeling the basic entities as individuals and observe the global *emergent* behavior.

The Traffic Basic Model

Basic Components of ABM

- Agents
- Environment
- Interactions

Agents

Agents are the fundamental entities of ABM

Concept introduced in the Artificial Intelligence field

Autonomous and decentralized

Interact with an environment

Agent Schools

- Artificial intelligence
  - Agents as autonomous entities solving problems
- Multi-agent systems
  - Distributed control of systems
- Agent-based modeling (and simulation)
  - Simulating (real world) phenomena

Agents

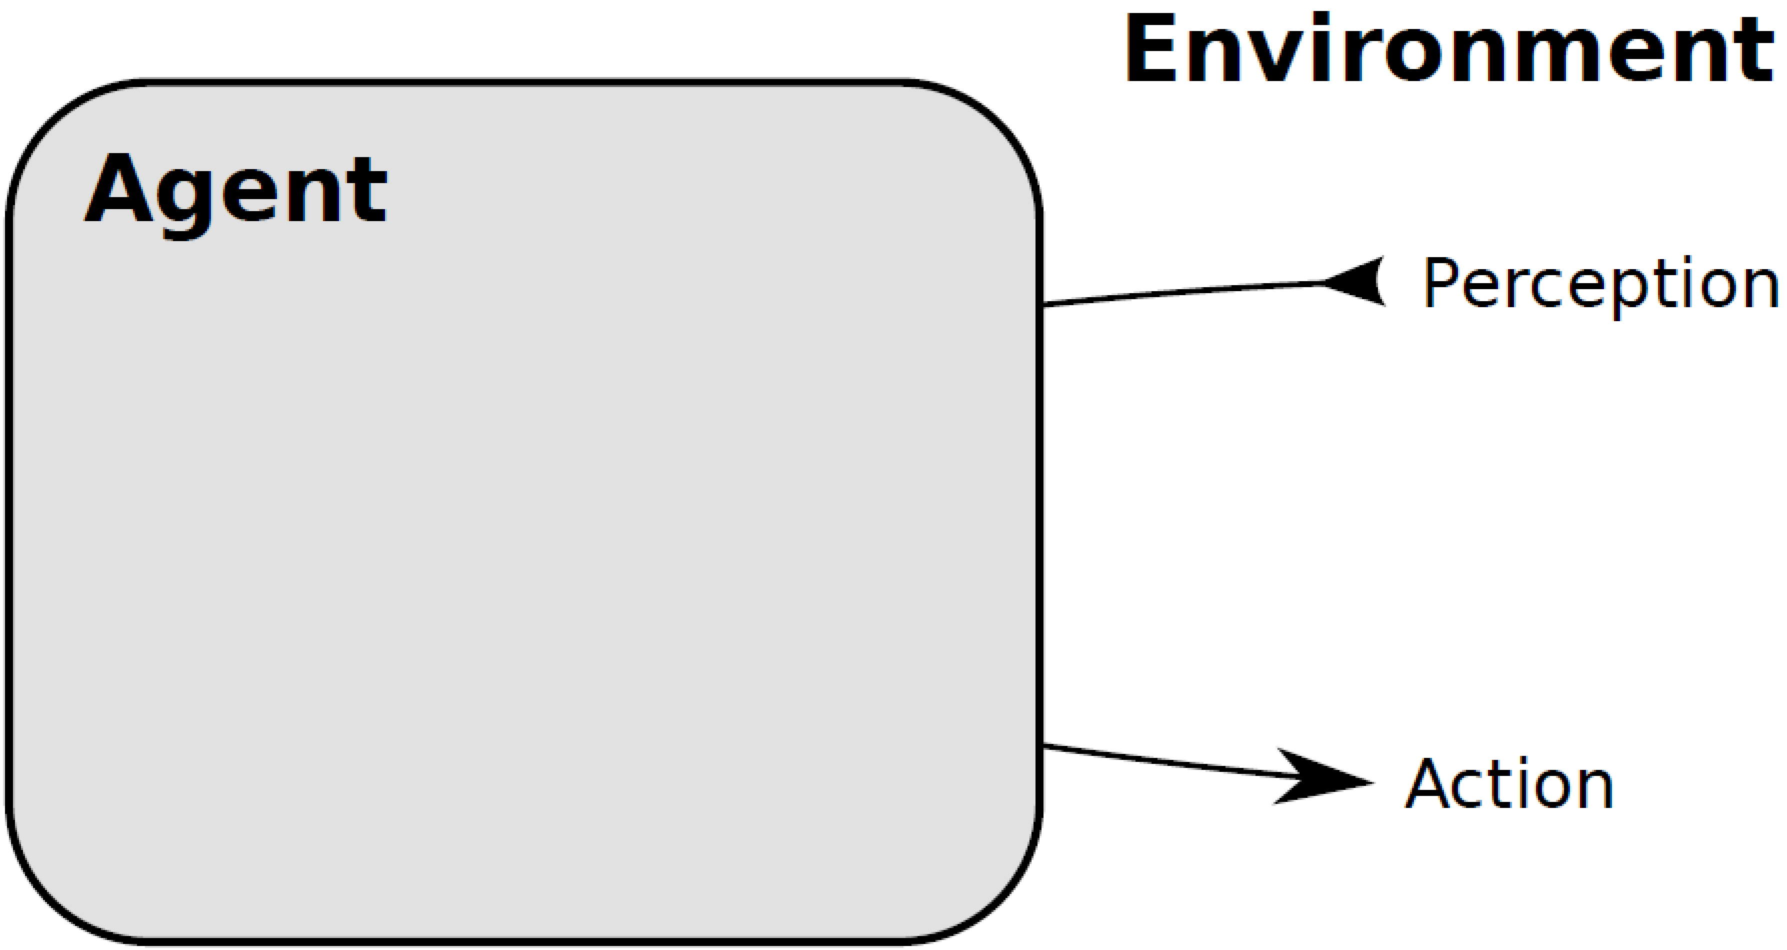

From the Agent Point of View

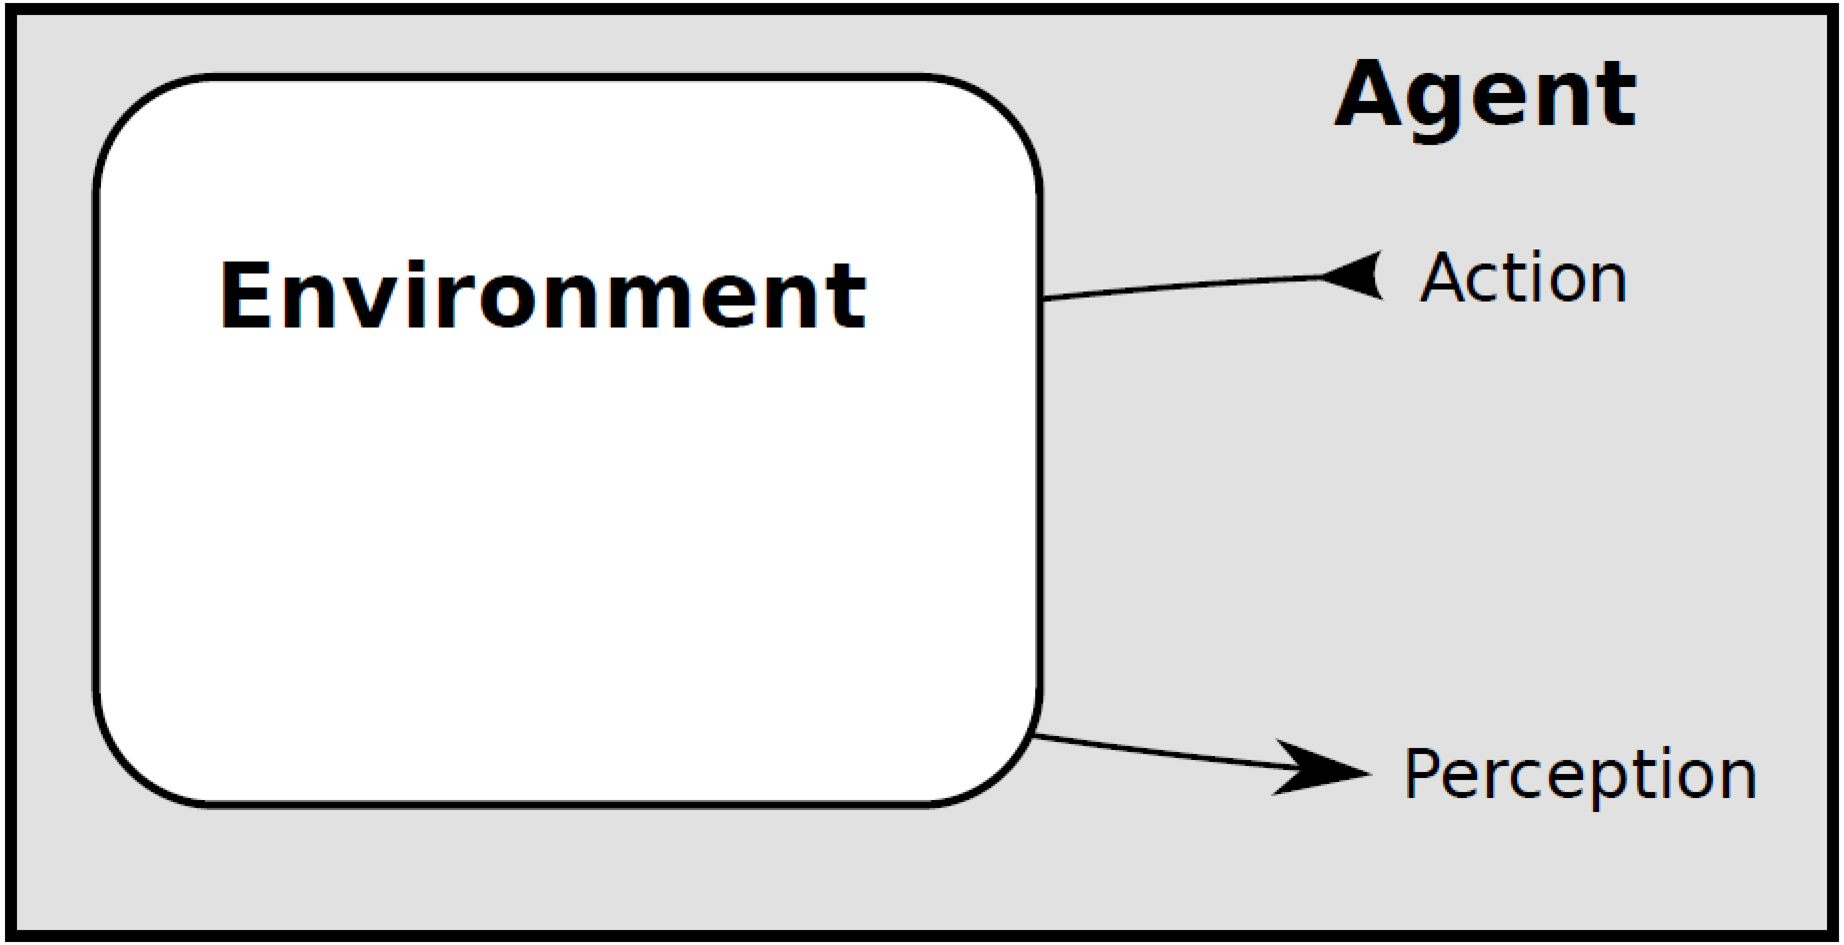

Simple Reflex Agent

Behavior: how the agent will react to the environment *perception*.

- Usually rule-based (Finite State Machine)

$$Perception \rightarrow Action$$

- May be stochastic
- Perception/Knowledge of environment is limited
- The *action* may affect the environment

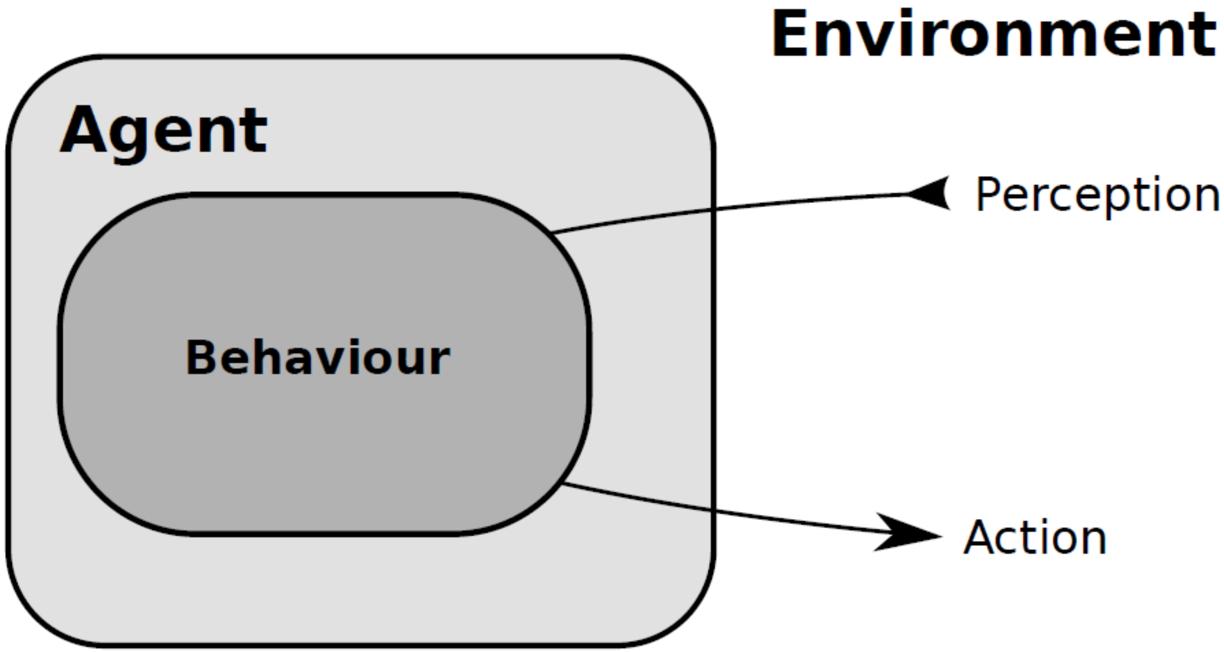

Simple Reflex Agents

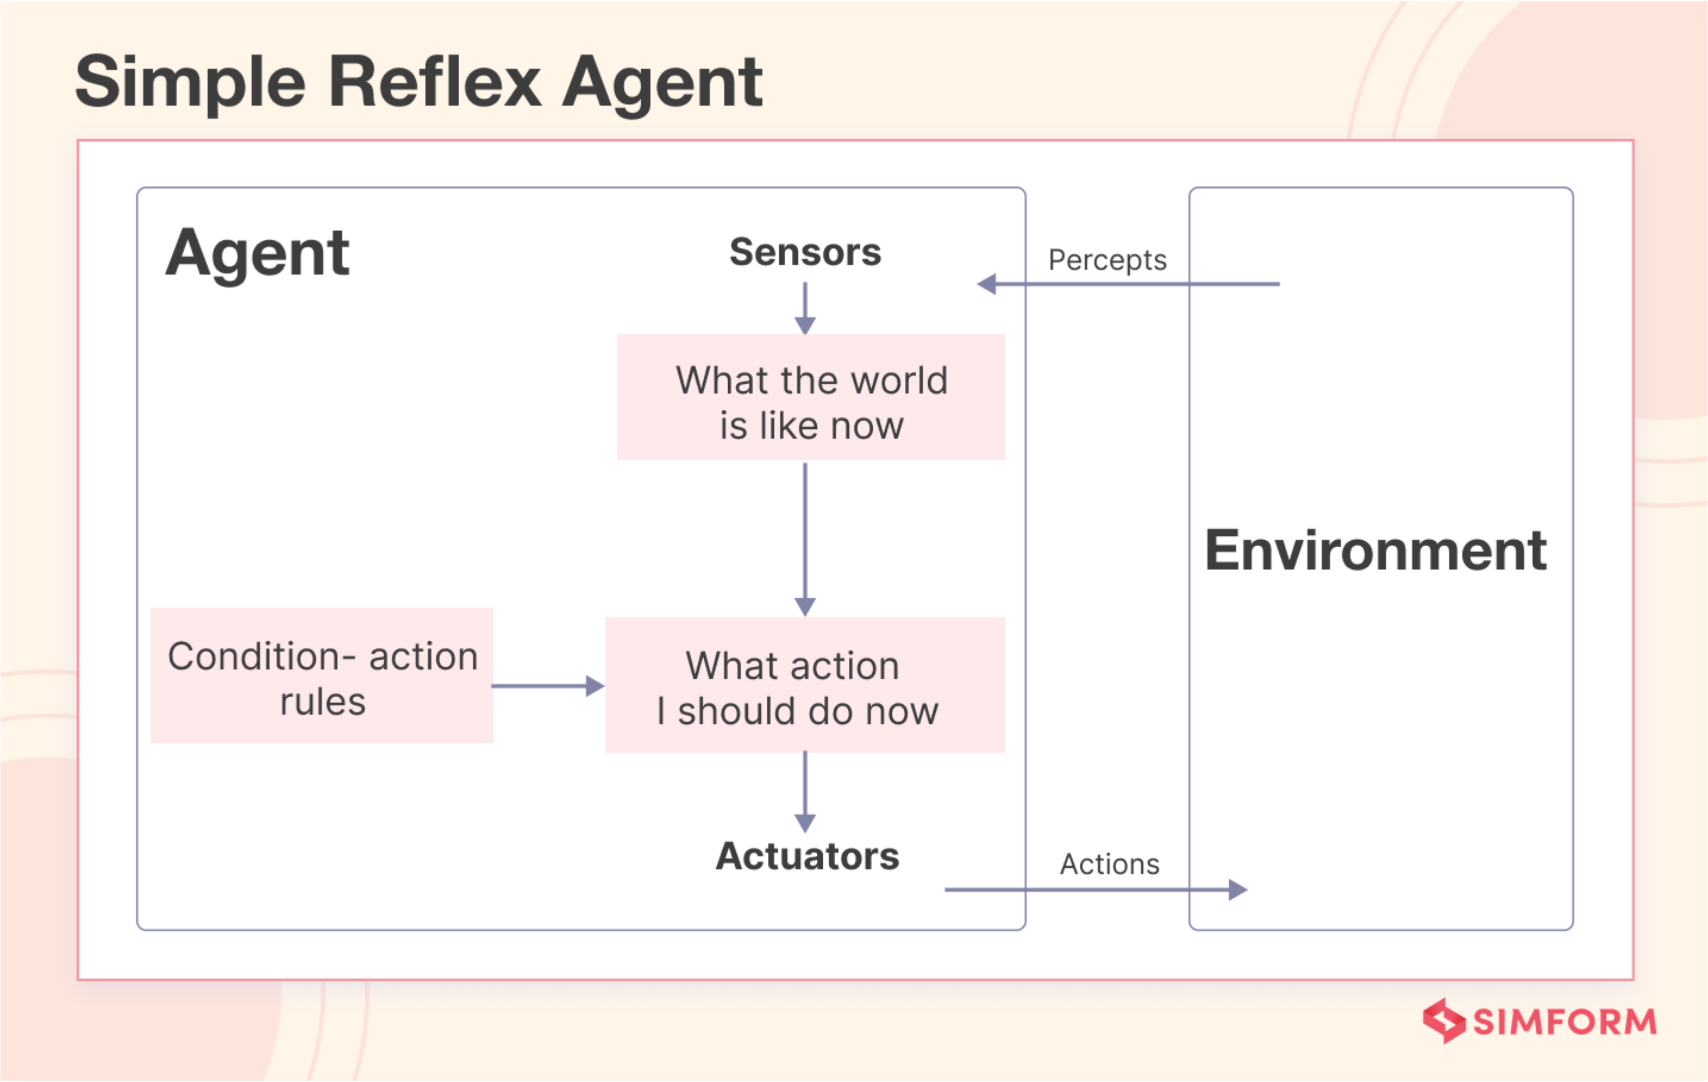

An FSM for an NPC

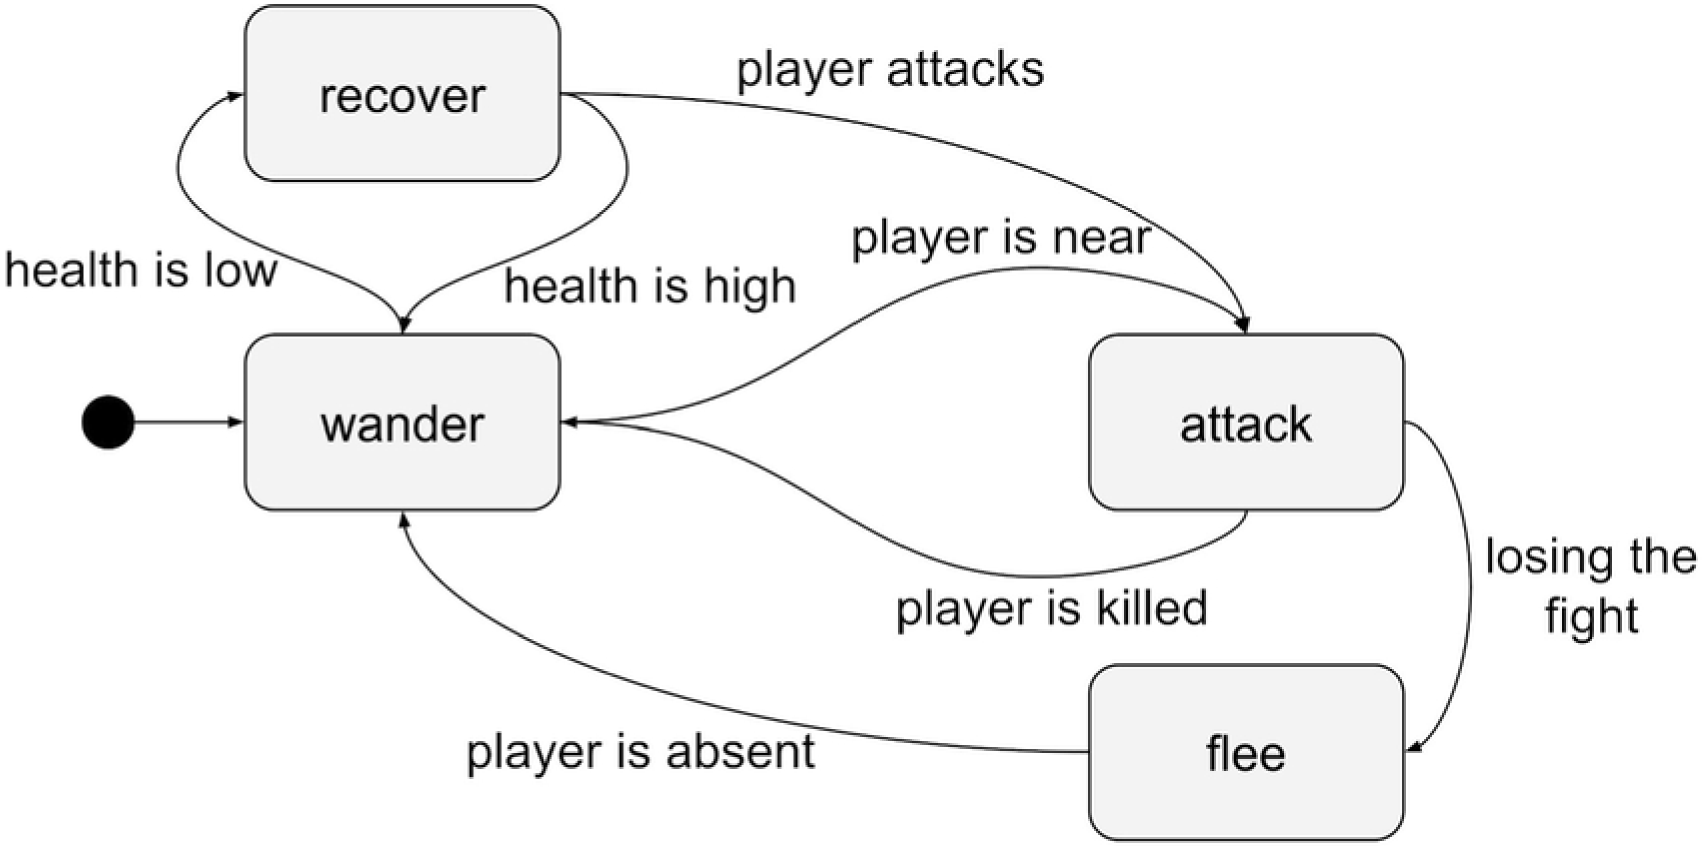

Intelligent Agent

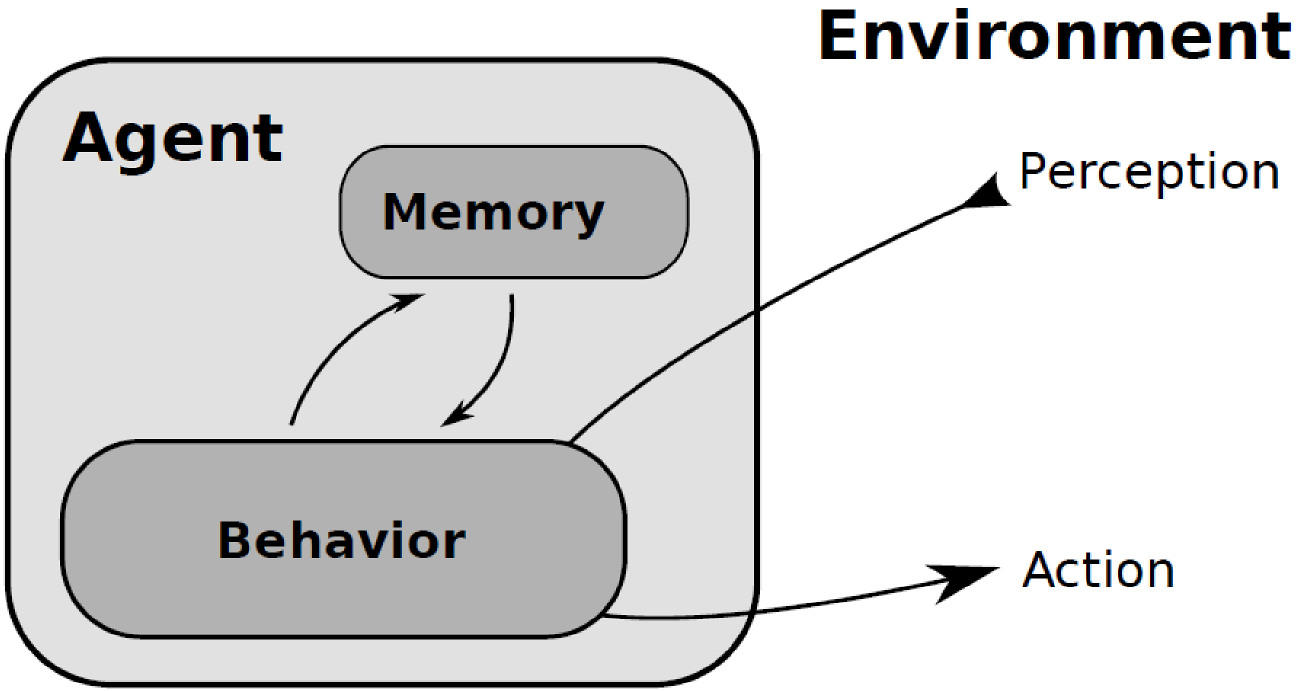

The agent has a state, which can be as simple as a boolean or as complex as it needs to be

Behavior function:

$$Perception \times State \rightarrow Action \times State$$

- The state is a kind of memory of past perceptions/actions
- The behavior depends on memory
- Hence, the agent is capable of learning

Utility-Based Agents

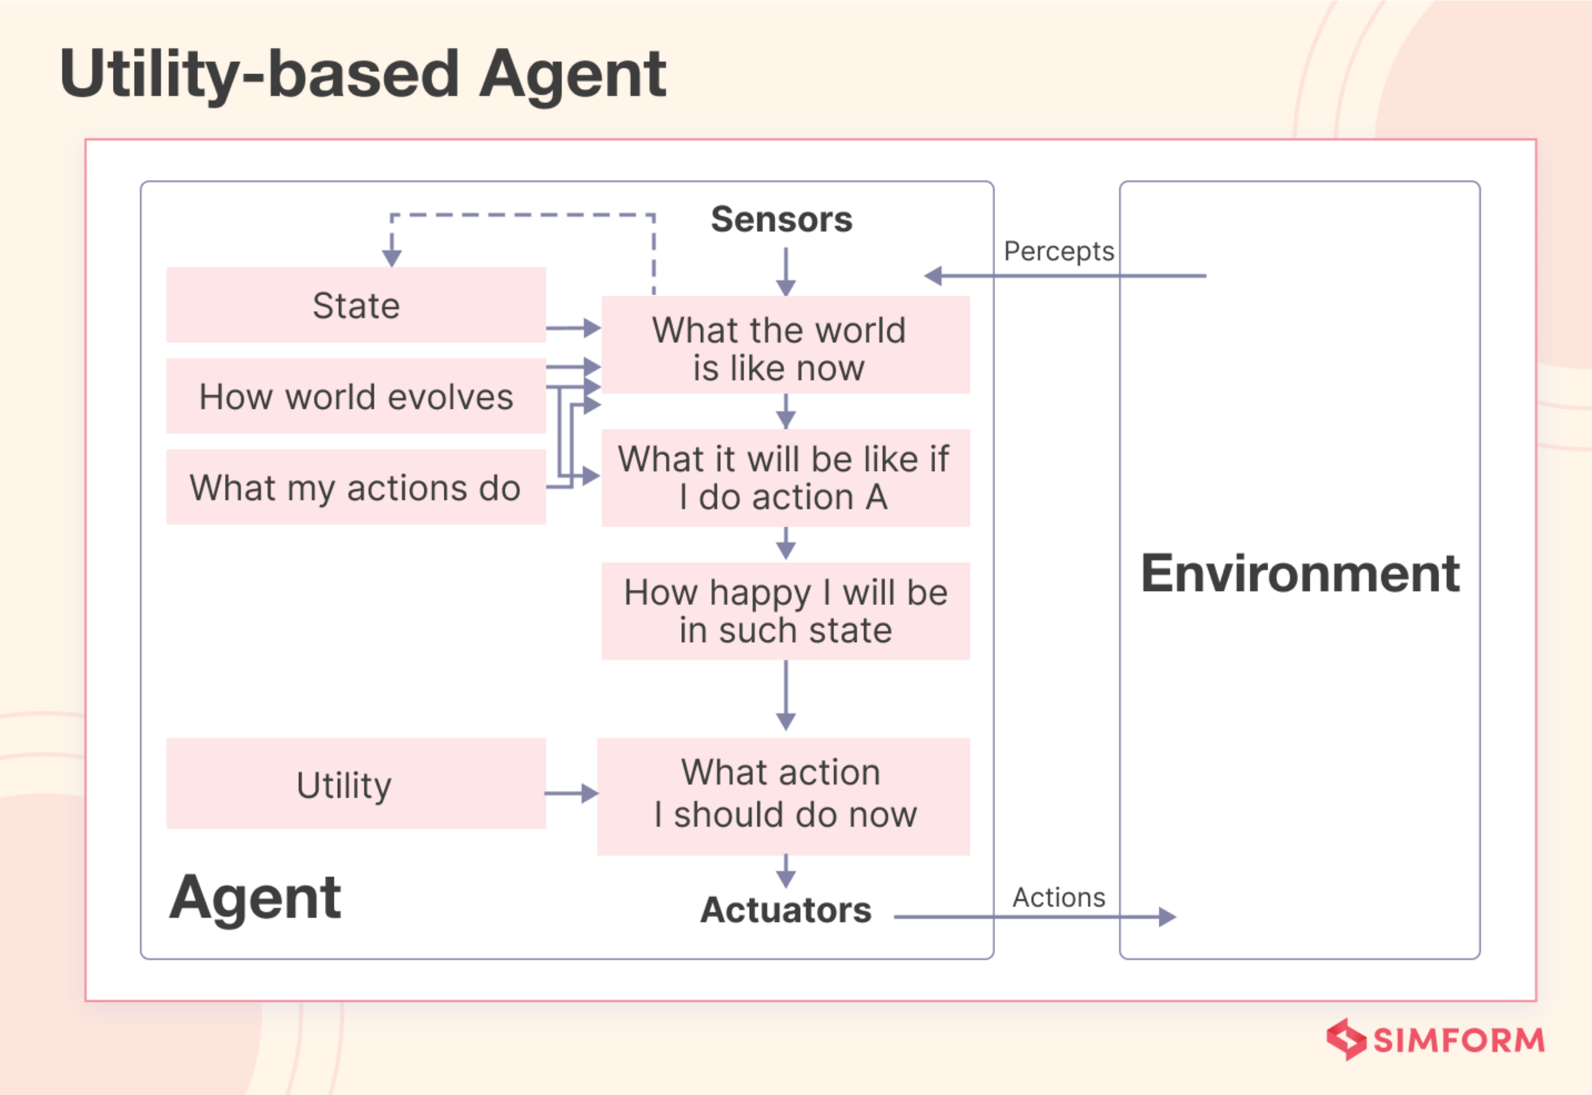

Learning Agents

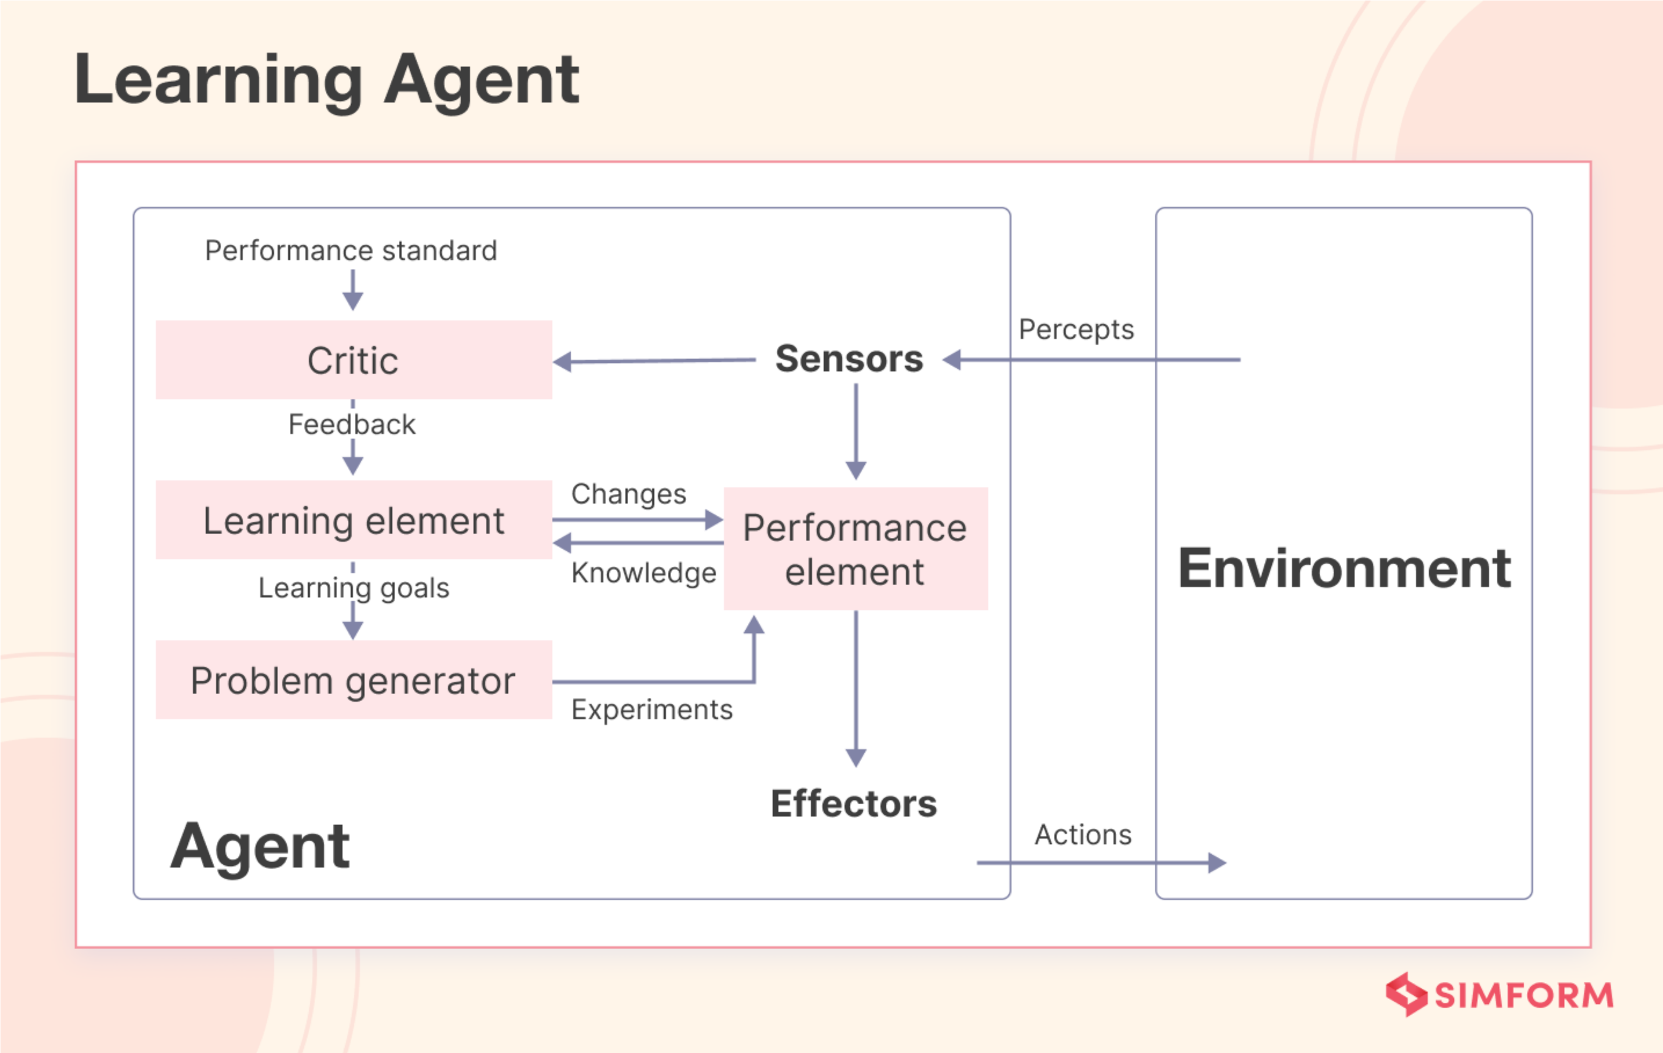

Agent Types

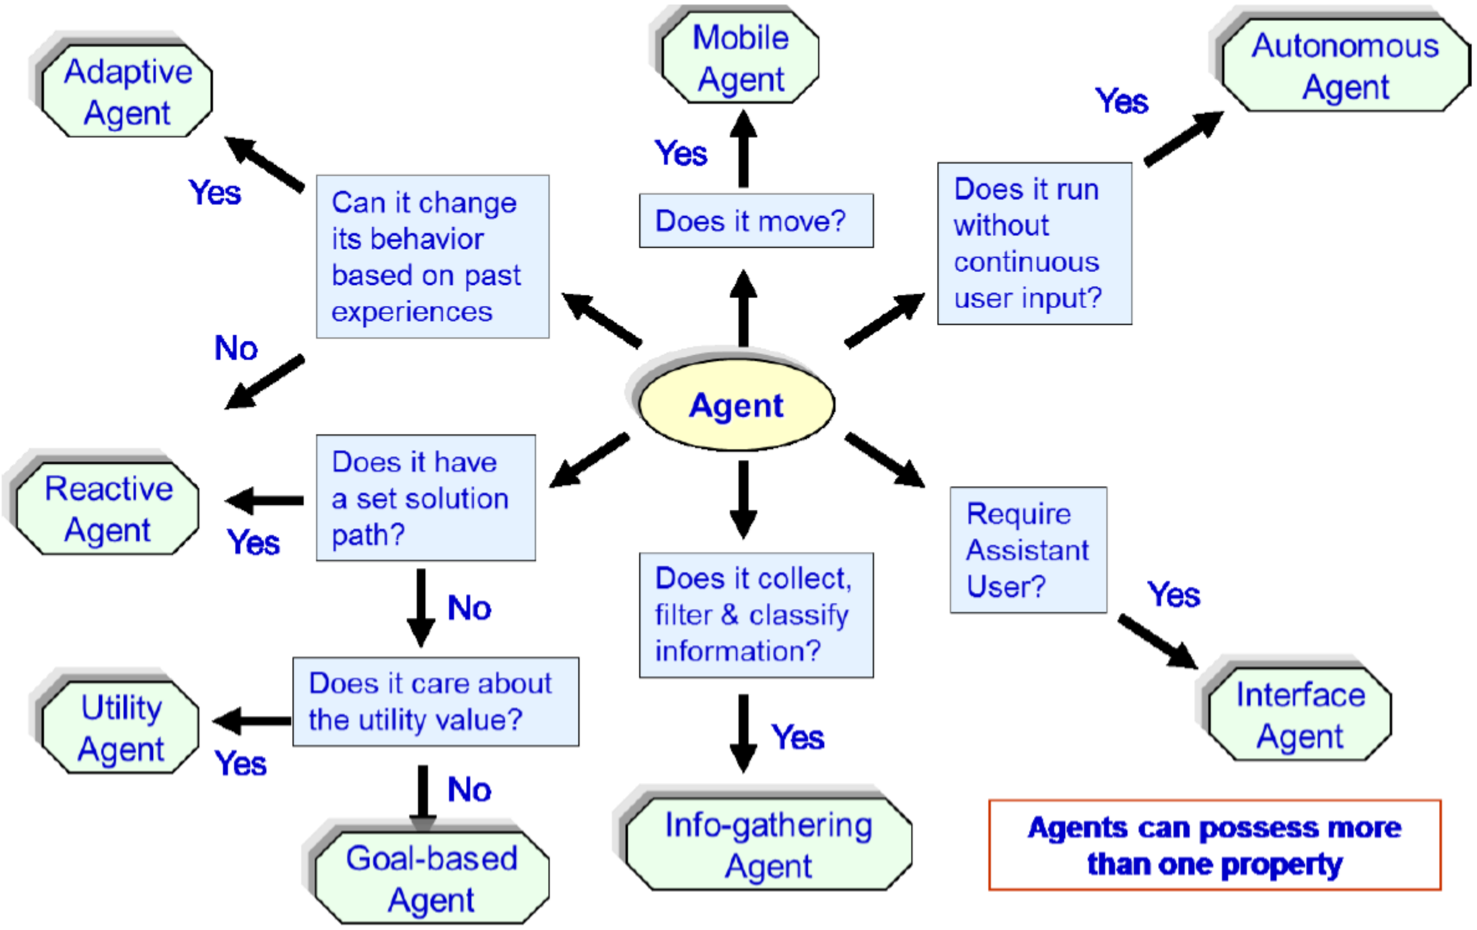
Adaptive Agent
Mobile Agent
Autonomous Agent
Yes
Yes
Yes
Can it change its behavior based on past experiences
Does it move?
Does it run without continuous user input?
No
Agent
Reactive Agent
Does it have a set solution path?
Yes
Require Assistant User?
No
Does it collect, filter & classify information?
Yes
Utility Agent
Does it care about the utility value?
Interface Agent
Yes
Yes
No
Goal-based Agent
Info-gathering Agent
Agents can possess more than one property

Environments

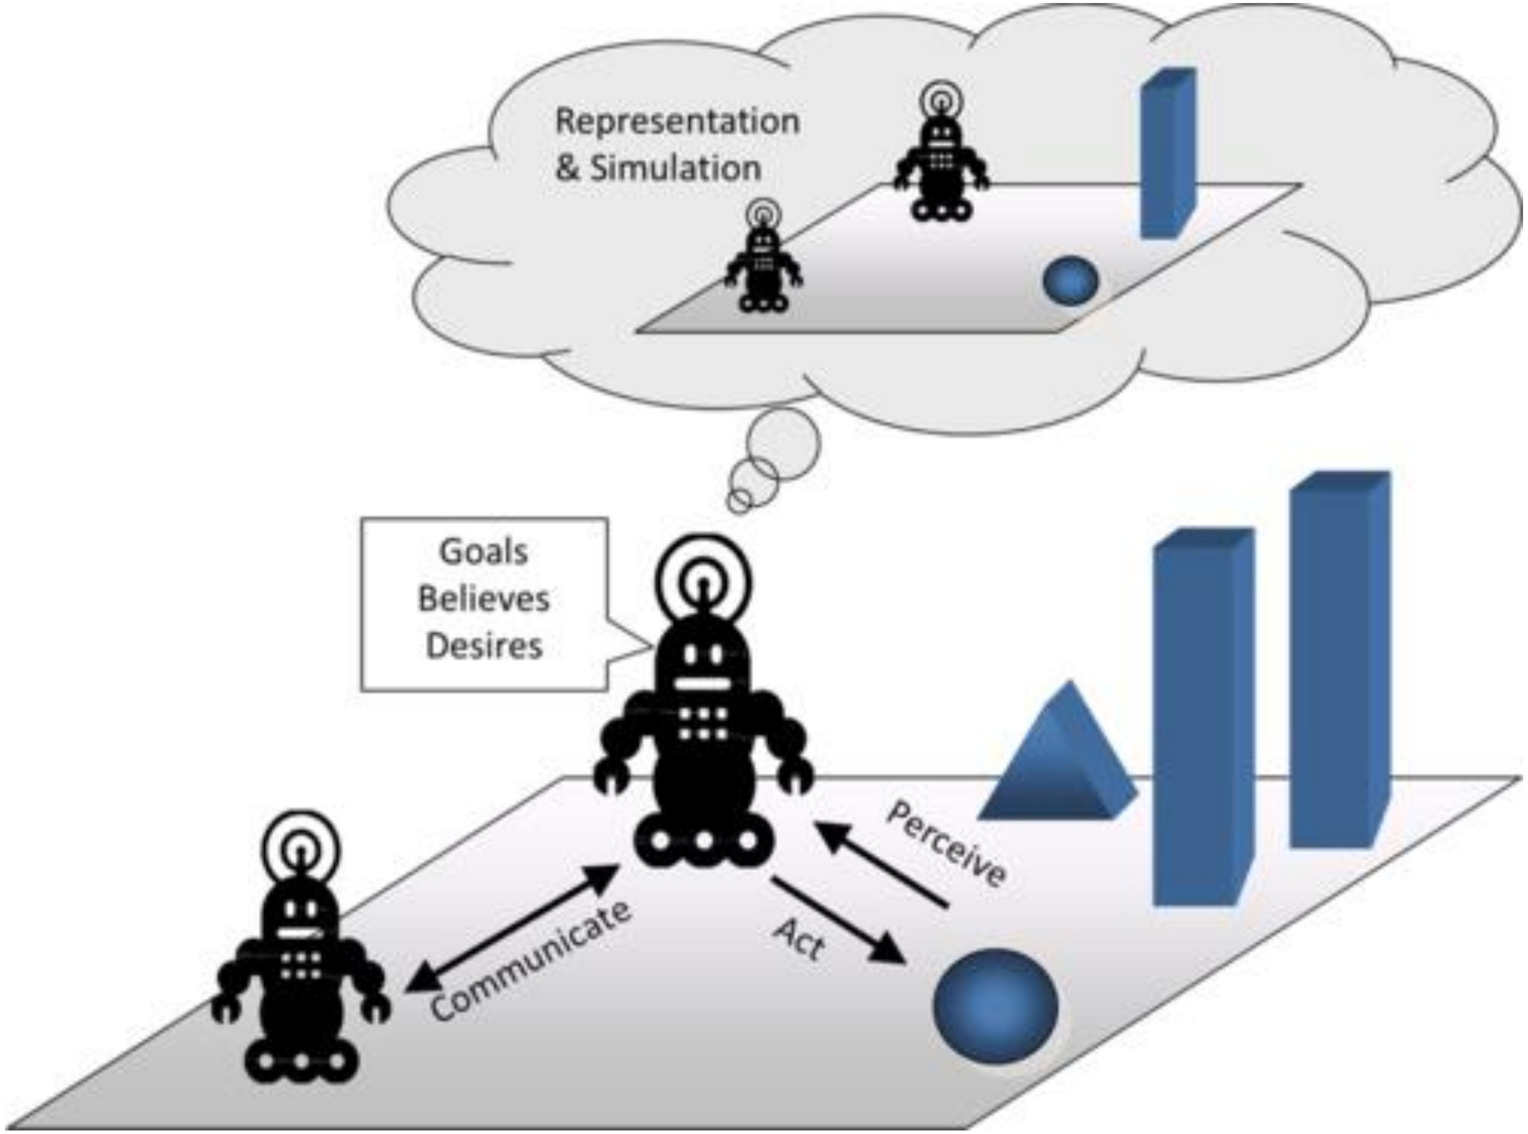

Some Types of Environments (Combinations and Variations are Possible)

Simple Genetic Alg

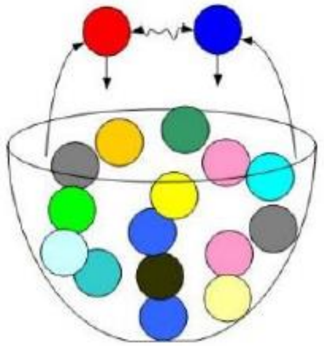

(a) "Soup" Model (Aspatial)

Some Types of Environments (Combinations and Variations are Possible)

Simple Genetic Alg

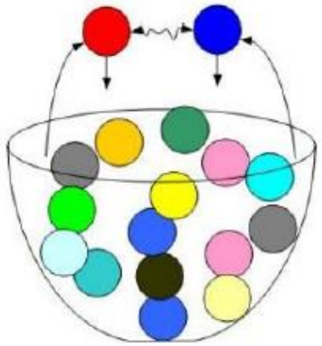

(a) “Soup” Model (Aspatial)

Lattice Gas Automaton

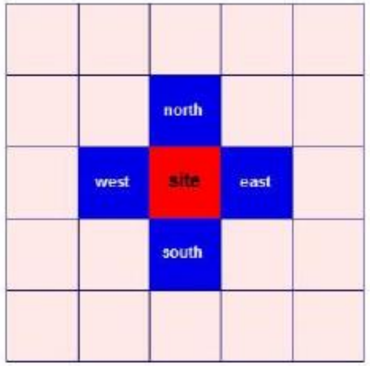

(b) Cellular Automata (von Neumann)

Some Types of Environments (Combinations and Variations are Possible)

Simple Genetic Alg

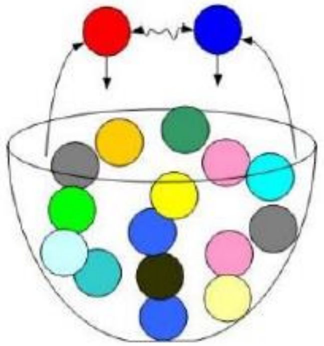

(a) "Soup" Model (Aspatial)

Lattice Gas Automaton

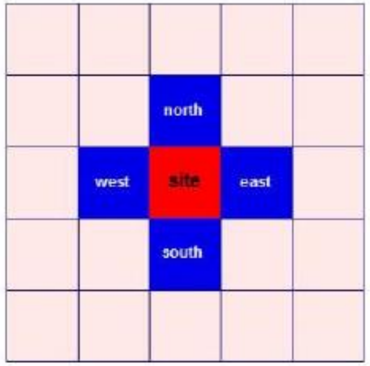

(b) Cellular Automata (von Neumann)

Flocking

(c) Euclidean Space (2-D)

Some Types of Environments (Combinations and Variations are Possible)

Simple Genetic Alg

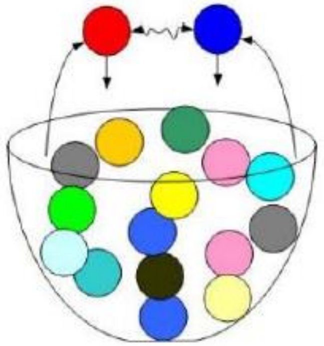

(a) "Soup" Model (Aspatial)

Lattice Gas Automaton

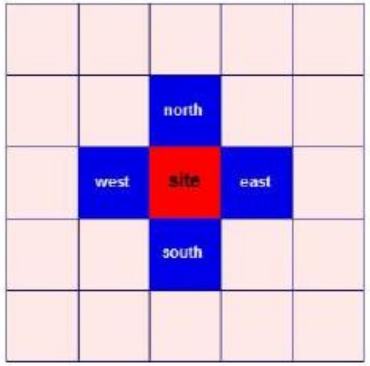

(b) Cellular Automata (von Neumann)

Flocking

(c) Euclidean Space (2-D)

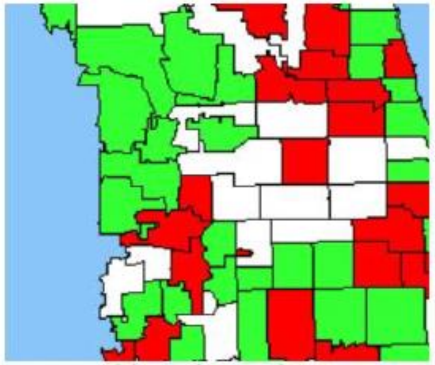

(d) Geographic Information System (GIS)

Grand Canyon Model

Some Types of Environments (Combinations and Variations are Possible)

Simple Genetic Alg

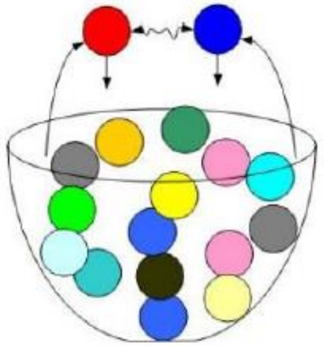

(a) “Soup” Model (Aspatial)

Lattice Gas Automaton

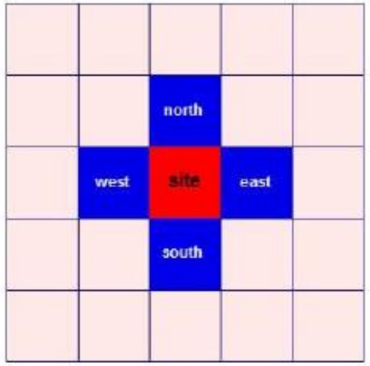

(b) Cellular Automata (von Neumann)

Flocking

(c) Euclidean Space (2-D)

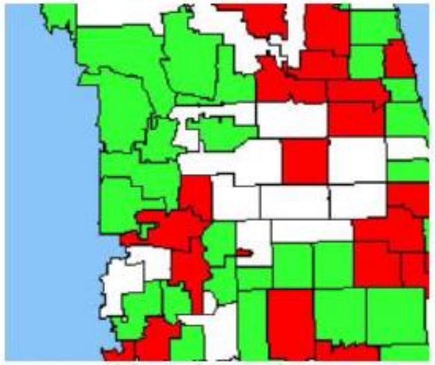

(d) Geographic Information System (GIS)

Grand Canyon Model

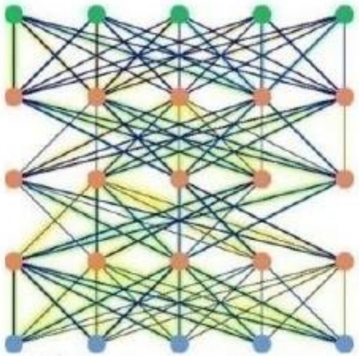

(e) Network topology

Virus on a Network

Four Types of Problems based on Agent/Environment

e-Science

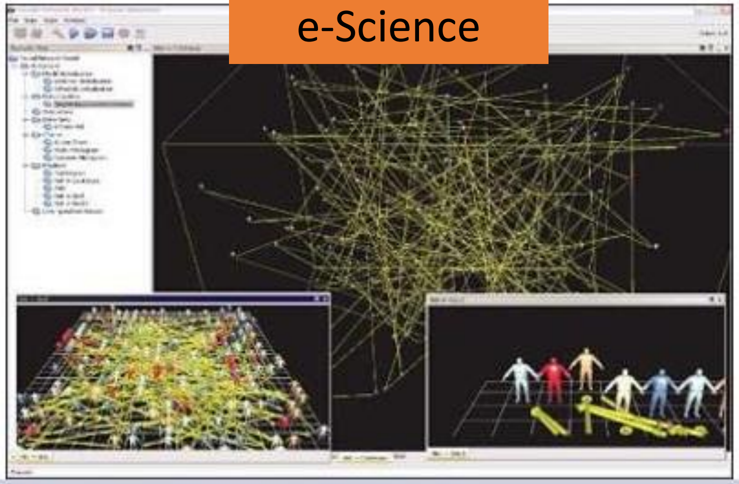

Artificial agents, artificial environment

Engineering Applications

Artificial agents, natural environment

Behavioral Experiments

Natural agents, artificial environment

Descriptive Model

Natural Agents, natural environment

Interactions

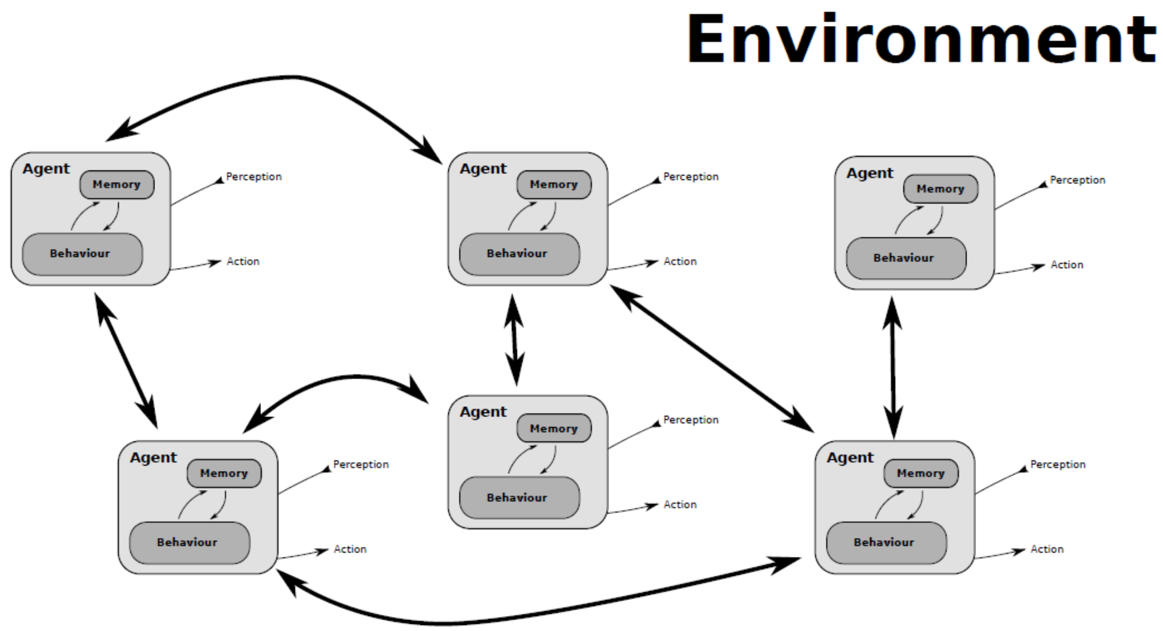

Environment Awareness

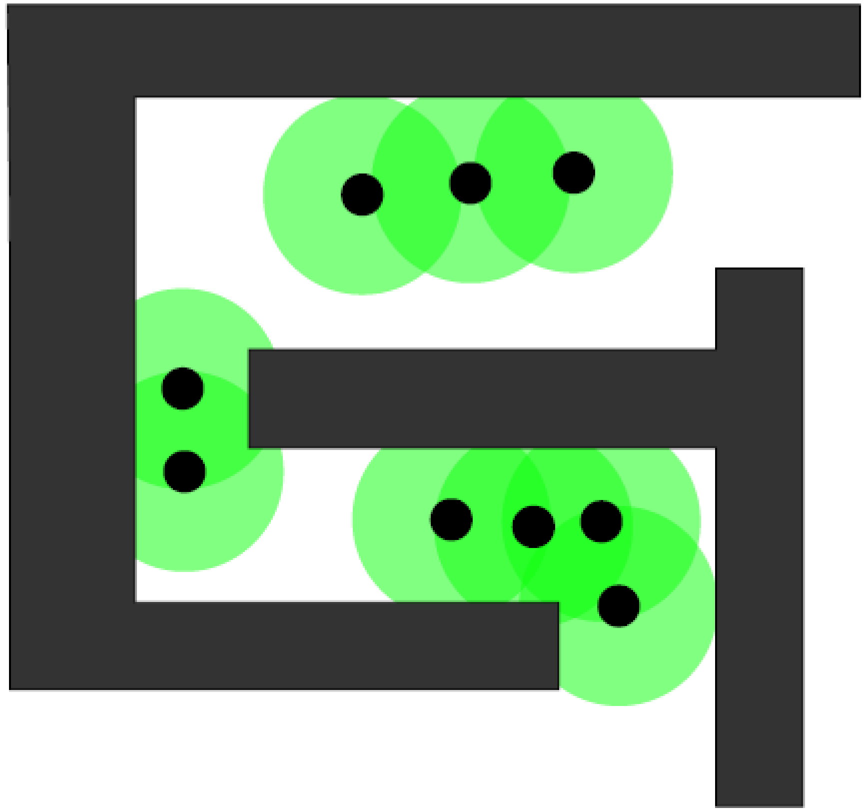

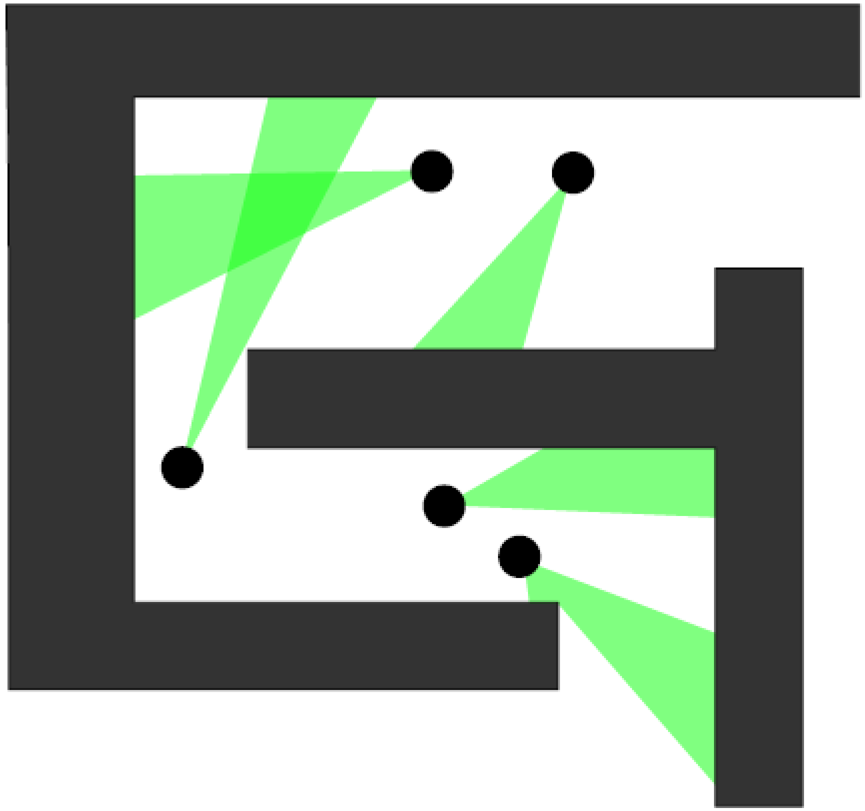

Interaction Topology

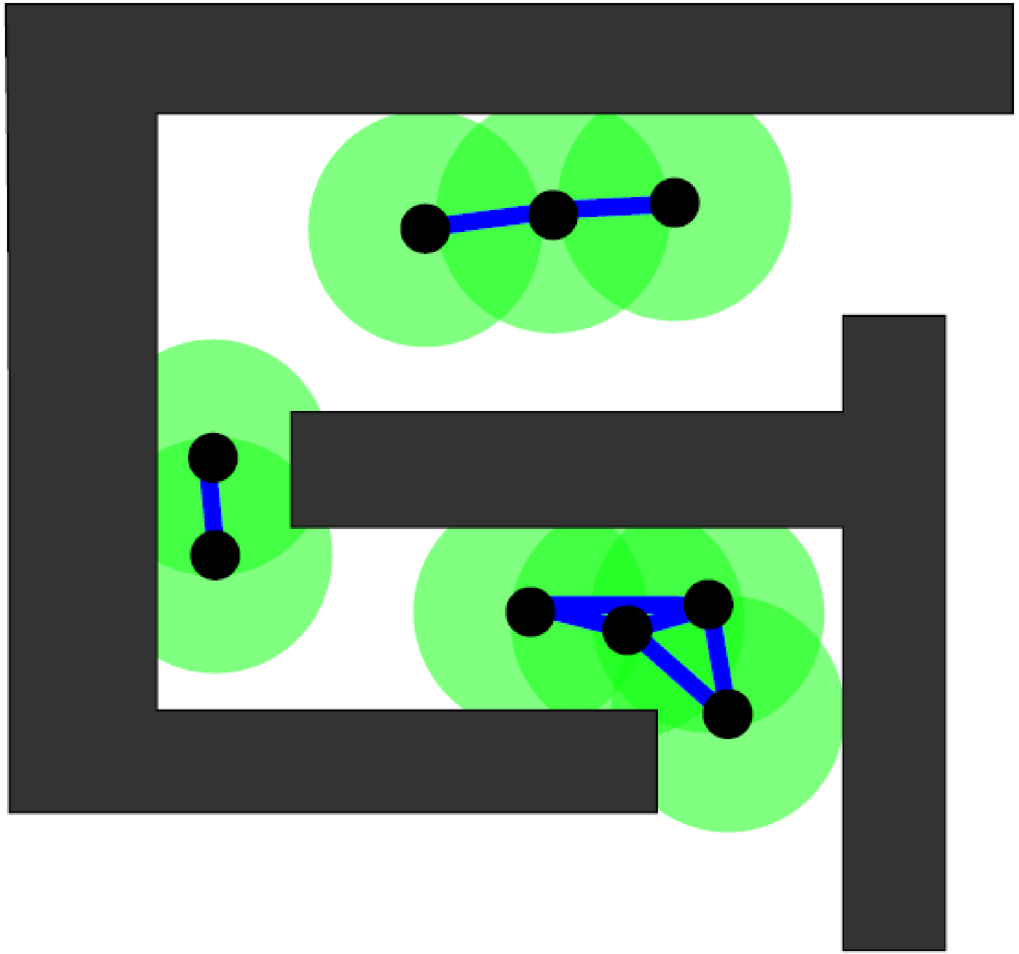

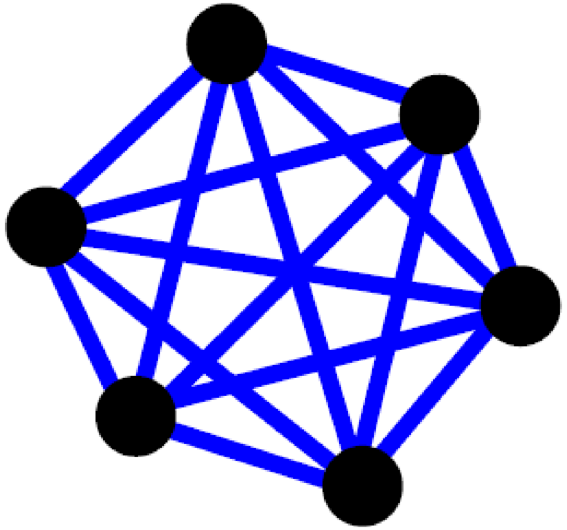

Interaction Types

- Agent – Self (e.g., reproduction)
- Environment – Self (e.g., grass)
- Agent – Agent (e.g., wolf-sheep)
- Agent – Environment (e.g., sheep-grass)
- Environment – Environment (e.g., diffusion in ants)

When to Use ABM?

Medium Numbers (Mesoscopic)

Too few agents and the “simple” may be too simple

- Game theory and ethnography work well

Too many agents and “means” may describe the system well

- Mean-field approaches and statistical descriptions

The key is that the number of agents that can affect the outcome of the system be a medium number.

Heterogeneity

- Agents can be as heterogeneous as they need to be

- Many other approaches assume homogeneity over individuals

Complex but Local Interactions

ABM can model complex interactions

- History dependent
- Property dependent

The assumption is that these are local

- No global knowledge

Rich Environments

The environment the agents interact in can be extremely rich

- Social Networks
- Geographical Systems
- The environment can even have its own agent-like rules
- …

Time

Almost all agent-based models feature time

- ABM is a model of process
- Nearly necessary
- There are exceptions
  - Solving complex equilibrium problems

Adaptation

- Adaptation is when an agent’s actions are contingent on their past history
- An agent may take different actions depending on its own past experience
- Very few modeling approaches besides ABM feature adaptive individuals

ABM vs EBM (Equation-Based Modeling)

- Many EBMs make the assumption of homogeneity
- EBMs are often continuous and not discrete
  - The nano-wolf problem (Wilson, 1998)
- EBMs require aggregate knowledge in many cases
- Ontology of EBMs is at a global level
- EBMs do not provide local detail
- EBMs are Top-Down, ABMs are Bottom-Up
- EBMs are generalizable, but restricted
- ABM can be built from analytical models, and can complement EBMs

ABM and Statistical Modeling

- Hard to link to first principles and behavioral theory
- Need to have the right kind of data
- ABM can complement by building from first principles to statistical results

ABM vs Lab Experiments

- Lab experiments can generate theory

- Lab experiments are rarely scaled up

- ABM can be created from lab experiments
  - ABM can explore macro-implications of lab experiments
  - ABM can generate new hypotheses
  - ABM can determine sensitivity of results
  - ABM can compare generative principles

Limitations

High Computational Cost

- Benefit of more insight and data to intermediate stages

Many Free Parameters

- Simply exposing parameters that other models assume

May Require Individual-Level Behavioral Knowledge

- Provides better insight

Why is ABM Resisted?

- Lack of Education about Complex Systems
- The Drunk, The Keys and The Streetlight
  - People want to search for solutions where it is easy
- Centralized and Deterministic Mindset
  - People expect their to be a central leader
  - People expect that everything happens for a “cause” and negate the possibility of chance

Uses of ABM

- Description (a simplification of a system)
- Explanation (underlying phenomena that control a system)
- Experimentation (make interventions)
- Analogy (boids to swarm of drones)

- Education (easy to understand, explore, and experiment)
- Touchstone (turns complex systems to simple rules)
- Thought Experiments (think of anything and experiment)
- Prediction (future scenarios – Very careful here…)

Back to the Ants

Termites Mo

Deneubourg’s Model (1991)

- Ants on a regular grid, with 4 directions (von Neuman neighborhood)
- Random walk, can walk over a corpse
- Sequential (asynchronous) updating scheme

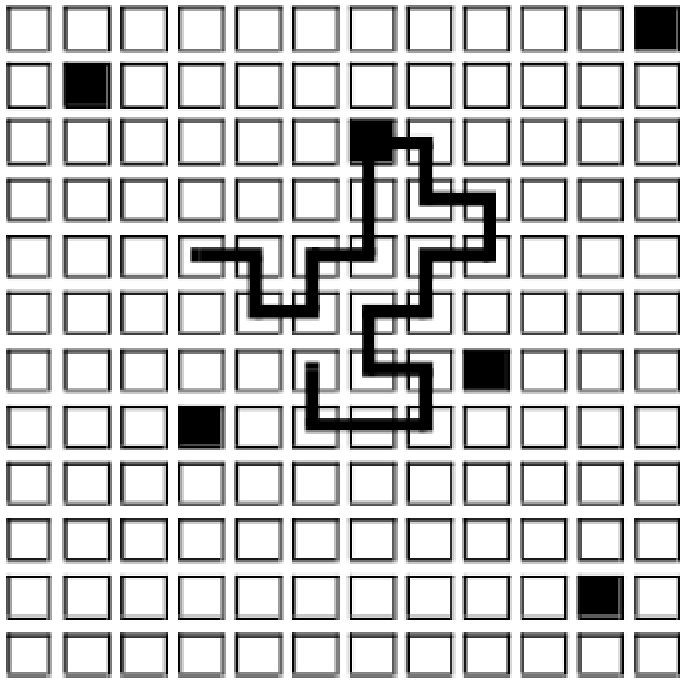

Ant Behavior

- ➢ With probability P_p, the workers pick up a corpse if it is isolated or in a small cluster
- ➢ With probability P_d, the workers deposit a corpse in large cluster of dead bodies

How does the ant evaluate the cluster size?

- Each ant has a memory M of size n:
- The memory locations indicate the state of the cells visited by the ant during the last n steps: $M(i) = 1$ if there was a corpse at time $t - i$, 0 otherwise

$$f = \sum_{i=1}^{n} M(i)$$

$$P_p = \left(\frac{k_1}{k_1 + f}\right)^2$$

$$P_d = \left(\frac{f}{f + k_2}\right)^2$$

k_1 and k_2 are model parameters.

Simulation of Smart Ants

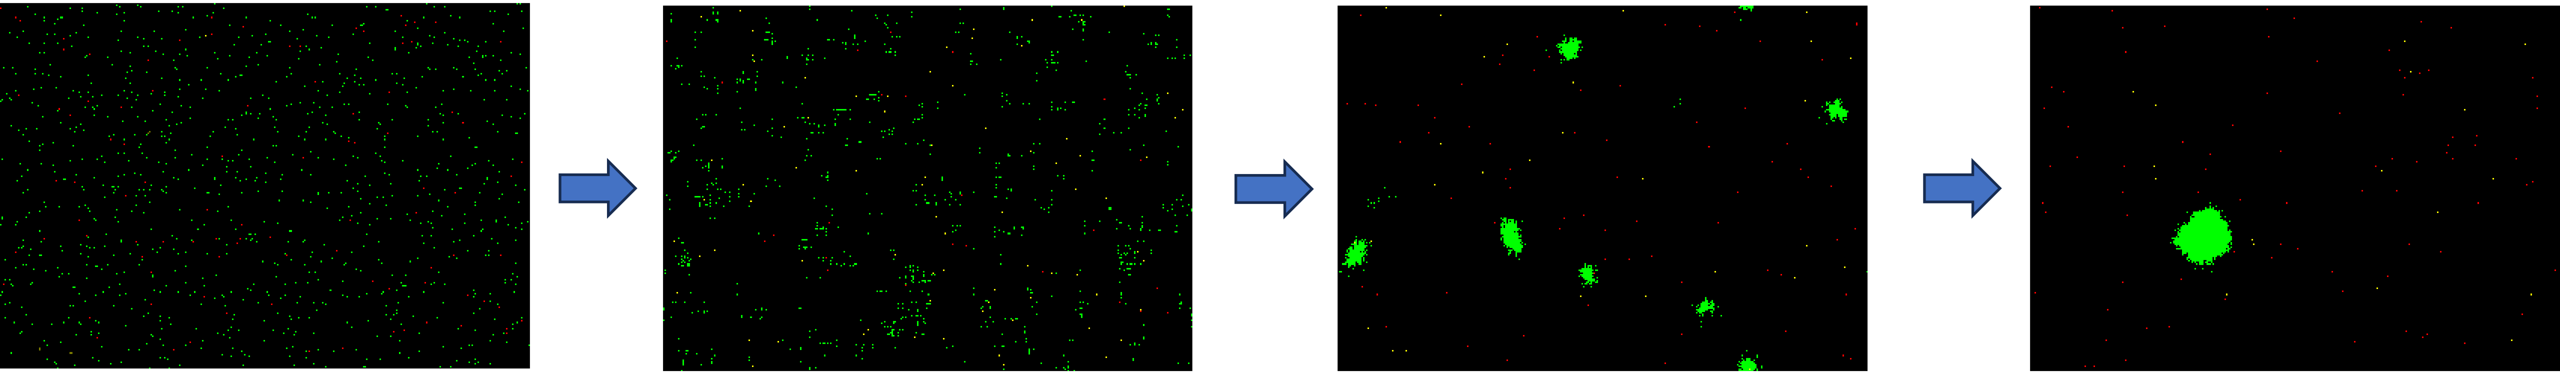

Deneubourg’s model works well

- Basic mechanism is intuitive
- But it requires a lot of "intelligence" from ants
- What about dumber ants?

Unige Model (2000)

- Regular Grid, with 8 directions (Moore neighborhood)
- Random Walk with large diffusion constant
- Asynchronous updating

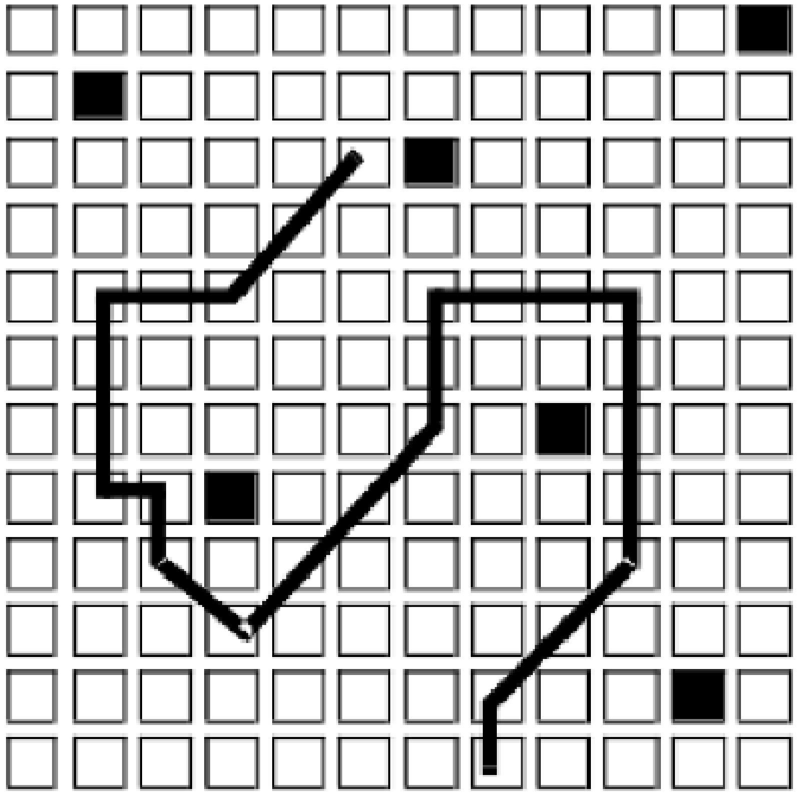

Behavior

- The ants avoids all obstacles:
  - ant corpses
  - other working ants
  - boundaries and walls

- An unloaded ant always picks a found corpse
- A loaded ant who finds another corpse always drops the carried corpse.

Simulation of Dumb Ants

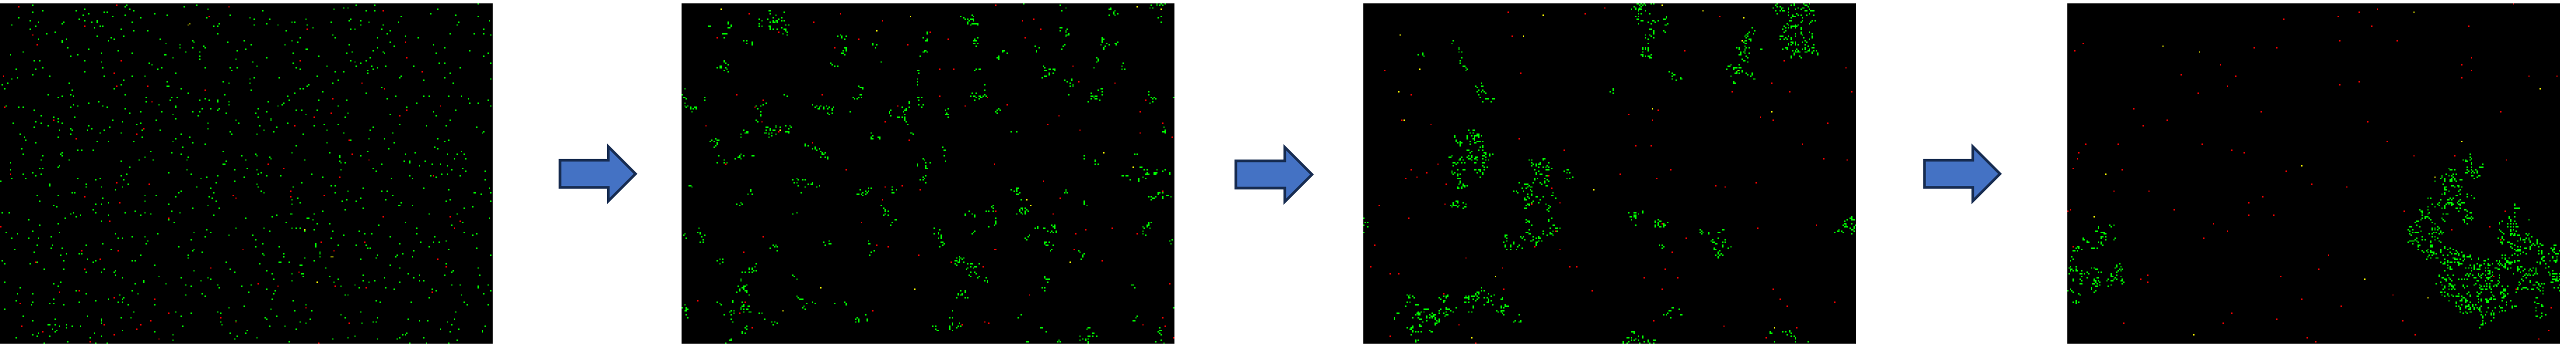

It works… but why?

- The probabilities to remove a corpse from a cluster, or to add a new corpse are the same.
- Ants make no difference between a large or a small cluster
- When a cluster is emptied it will never reappear.
- Due to fluctuations, all clusters but one will eventually reach a zero size

Quantitative Results

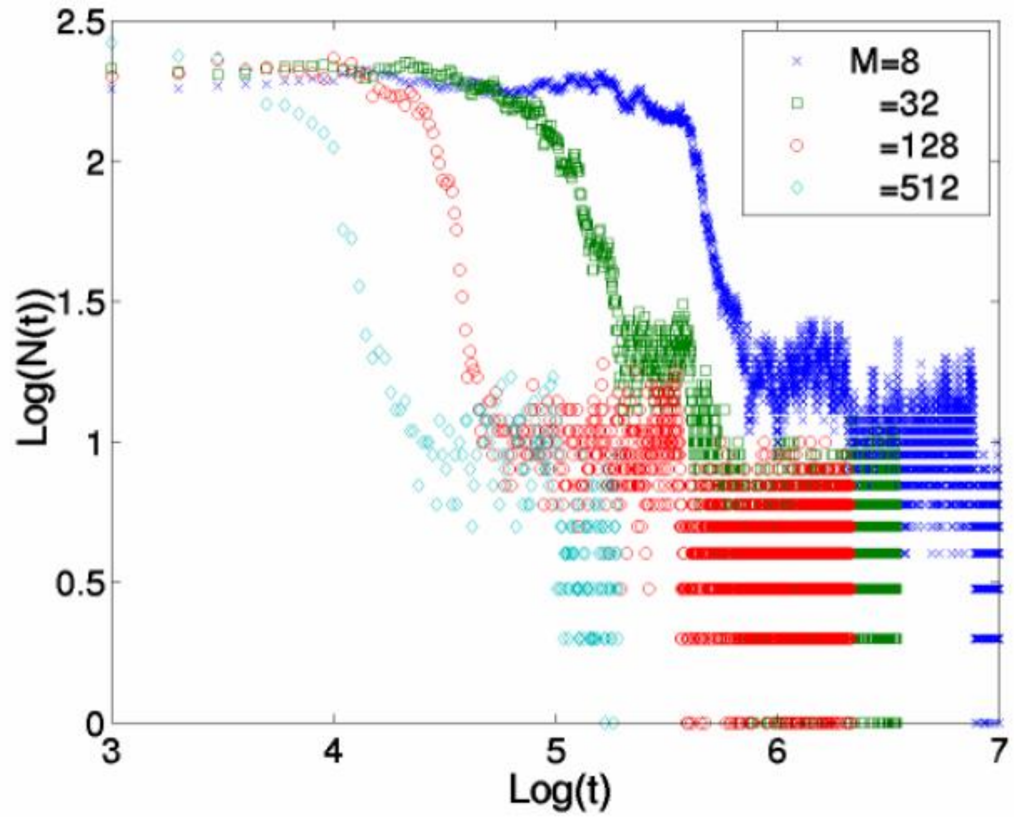

Deneubourg, with 8 directions

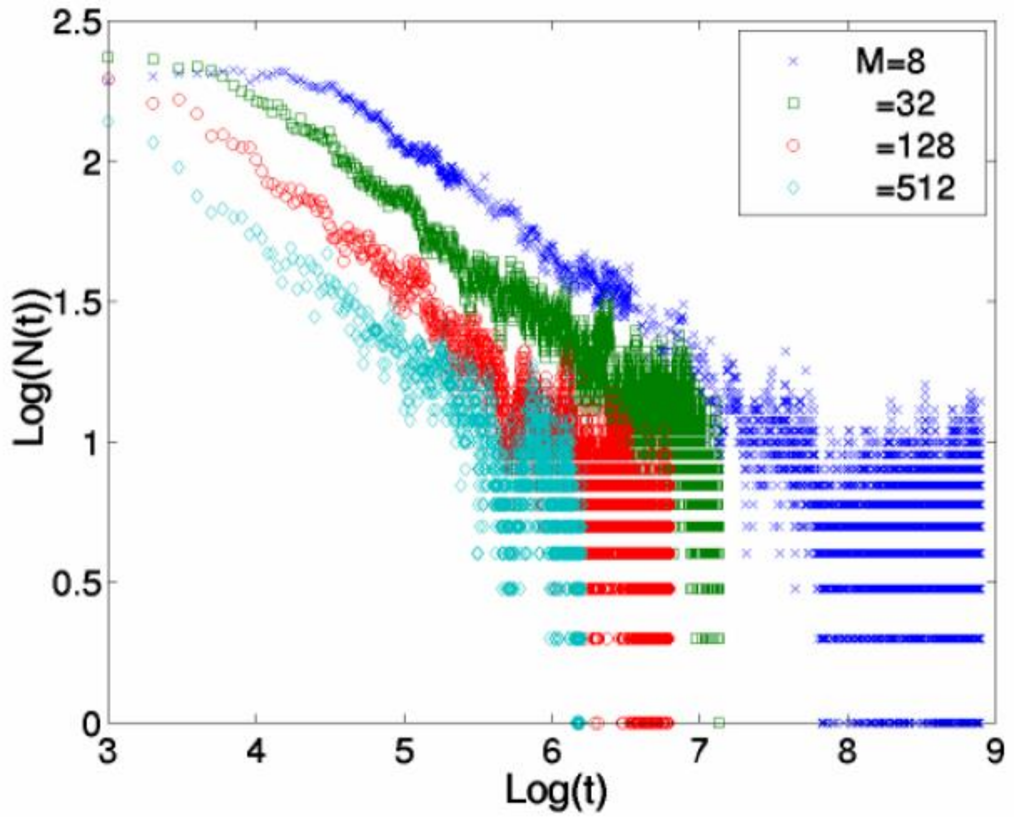

Unige

With Obstacle

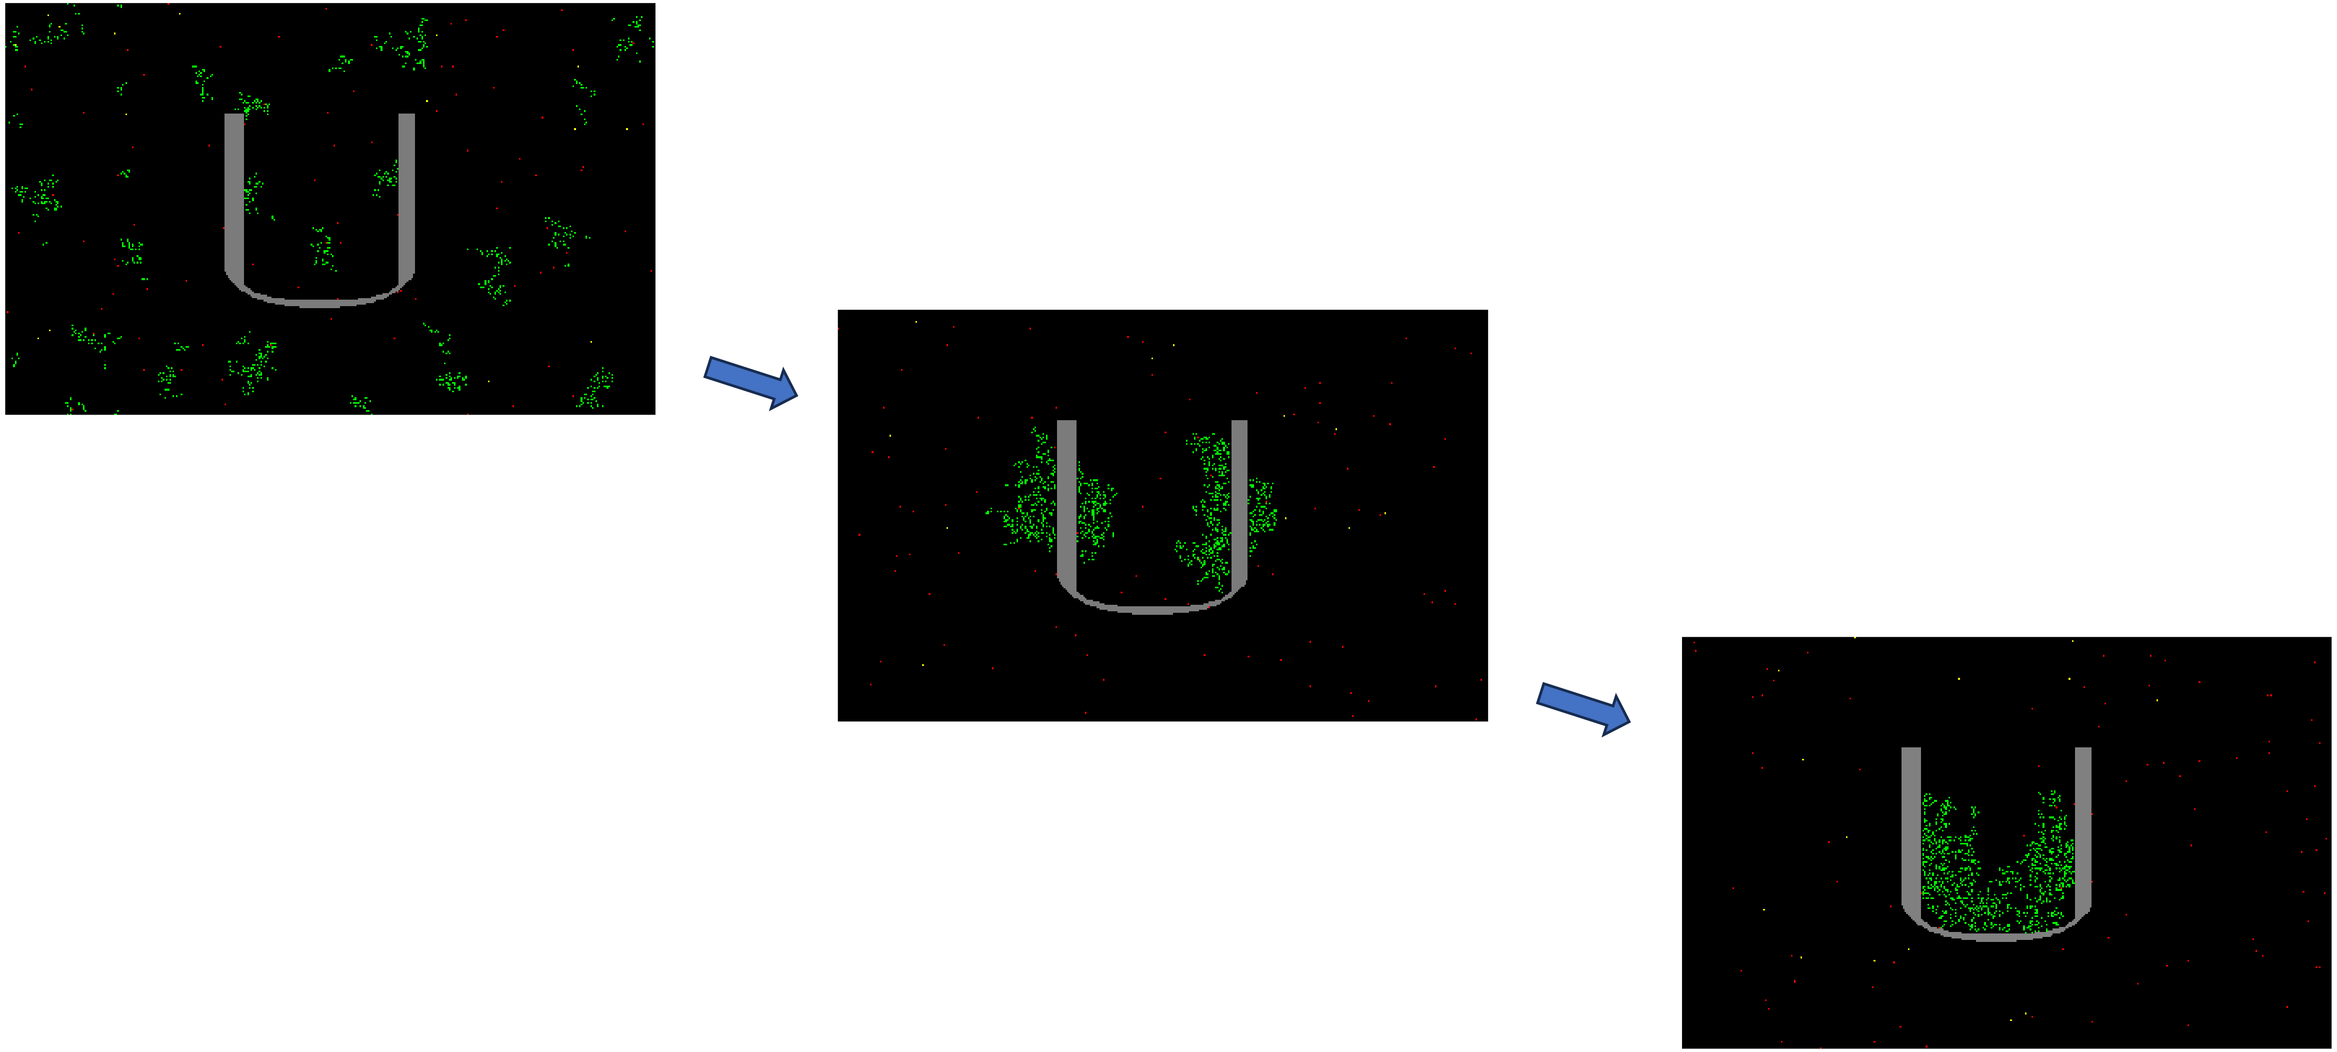

Conclusions

- Ant corps pile construction can be explained by statistical fluctuations

- Yet, intelligence speeds up the process
  - In Deneubourg's model converges $\sim 10\times$ faster (using better random walk).

- Not a collective effect, just a collaboration with a linear speedup
  - One single ant would make it, but slower
  - In both models: not a collective behavior, $N(t) = f(Mt)$

A Sketchbook for Ethics in Agent-Based Modelling

Slides taken from Andy Evans

Crime Modeling

Agent-based modelling of crime (burglary) at the city scale (Leeds in UK).

Ongoing collaboration with local police/government crime prevention partnership.

Model process

Inputs

- Real offender homes (postcodes)
- Victims from census (microsimulated ~100 household scale)

Behaviour

- Offender behaviour model (drug use, sleep, socialising, work, knowledge of city)
- Victim daily routines (sleep, socialising, work)

Prediction

- Crime locations
Journey to crime

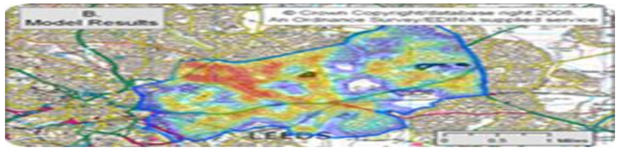

Near future

This is a likely model for all agent-based (and other) socio-economic studies, not just crime.

Inputs	Real offender data at individual level	Crowd-sourced dynamically streamed data
Behaviour	Offender behaviour model based on individuals	Individual victim daily routines estimated
Prediction	Crime locations Journey to crime	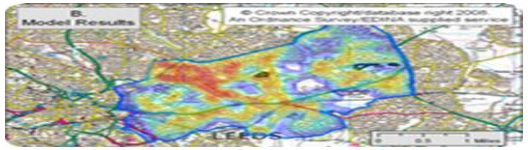

Error constraint

Dynamic data

Ethics

Inputs	Personal data	Ethics well understood
Behaviour	Validation of personal behaviour	Ethics poorly understood
Prediction	Predicted personal behaviour	Ethics rarely considered

Individuals

Crime prediction: Burglar committed to crime. (Think of Minority Report)

Sales opportunity prediction: Pregnant woman who is trying to hide it from an abusive family.

Input data

Data volunteered for another purpose / to secure other services

Data volunteered for the purpose

Data volunteered for purpose which is then perverted

Data not volunteered

Data extracted by force

Data

Are there differences between these?

Data volunteered for another purpose / to secure other services

Data volunteered for the purpose

Data volunteered for purpose which is then perverted

Unanticipated derived data

Difference depends in part on legitimacy of purpose.
Social good
Is commerce a sufficient social good?

Levels

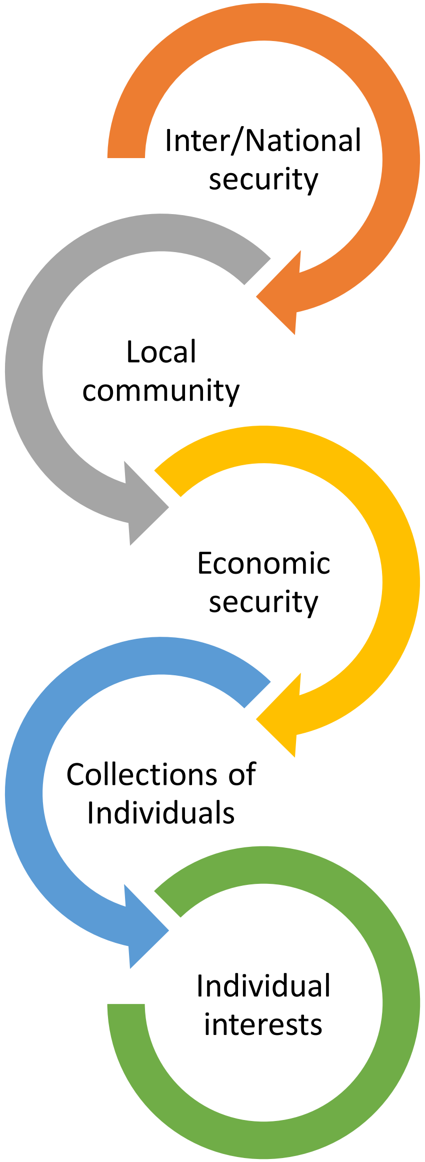

In some cases we'd weigh up by level affected.

However, in some cases there may be inviolable rights.

Behavior

Model runs to replicate current situation

Prove theory

Prove involvement

Are there cases where this wouldn't be appropriate?

When would this be appropriate with data not specifically based on this purpose?

Prediction

Predicted data

Agent-based studies offer a toolkit that could avoid the ecological fallacy

At what scale do individual concerns kick in? When identifiable? When they could be affected?

Do we have a right to contest initial data and/or predictions?

Do we have a right to know about predictions about us?

Do we have a right to be consulted on use?

Are predictions private data about us?

Predicted data

Are predictions private data about us?

Does it depend on how accurate the predictions are?

More accurate predictions are likely to compromise the previously identified rights

What if we can accurately predict otherwise private data?

Less accurate predictions = more error about us being acted on

Why does it matter that it is future data?

Isn't future data more important than past data?

Doesn't it make accuracy more important?

Actions on predictions

Do we have a right not to be predicted?

Do we replace the ecological fallacy with the fallacy of the individual?

Is that problematic?

How accurate do we have to be to act on a prediction?

What criteria do we use to make this decision? Do we draw different limits for acting on society and acting on individuals?

Issues

Do people have a right to control what their data is used for when non-voluntary?

When/should we model real individuals?

At what scale do we start worrying about individuals?

Are accurate predictions of private data the same as private data generated at a future point in time?