

## Berry Eaters: Learning Colour Concepts with Template Based Evolution

Chris Headleand<sup>1</sup>, Llyr Ap Cynedd<sup>1</sup> and William J. Teahan<sup>1</sup>

<sup>1</sup>Bangor University, Wales, UK  
c.headleand@bangor.ac.uk

### Abstract

In this paper, we propose an approach inspired by Gärdenfors conceptual spaces, to represent concepts (specifically colour) in a situated embodied agent. Through this we also extend the work of Braitenberg by proposing an alternate vehicle capable of recognising colour concepts. Furthermore, we demonstrate how it can learn to distinguish between different colours by evolving our vehicle's conceptual space using Template Based Evolution.

### Background and motivation

Categorisation and the ability to recognise and represent concepts has been described as a fundamental cognitive ability and a primary research goal of AI (Gärdenfors and Williams, 2001). The traditional approach from classical AI is to represent concepts symbolically using a top-down design process. The bottom-up subsymbolic (connectionist) approach, in contrast, emphasizes the importance of symbol grounding and the crucial role that the body has to play, by having the agent situated in the world and embodied—able to perceive and interact with the world through sensory-motor co-ordination (Pfeifer and Scheier, 2001).

Being able to recognise similarity is key to many biological phenomena, such as how simplistic agents are able to learn what food is safe to eat. Whilst prototypical items are easily represented through the symbolic approach, this does not take into account possible variations that naturally occur in real-life. Take the problem of colour recognition, for example. In natural systems, it is very unlikely to find exactly the same colour—no leaf on a tree is ever the same shade of green, or a berry is ever the same shade of red.

Gärdenfors (2004) presents an alternative methodology, called Conceptual Spaces which models knowledge in the form of regions in a geometric space made up of quality dimensions. By modeling in this way, Gärdenfors argues that an agent is able to recognise variations on a defined concept prototype. A primary motivation for this paper is to see if a Conceptual Spaces representation can be implemented in a situated, embodied agent in a manner analogous to the approach adopted by Pfeifer and Scheier (2001) in order to

demonstrate how concepts such as colour can be learnt in an unsupervised way.

### Conceptual Spaces

In his book, Gärdenfors (2004) proposes Conceptual Spaces as an alternative and possible bridge between alternative knowledge representation methods, specifically the symbolic and connectionist approaches. One of the principle problems this method seeks to address is the modelling of abstract concepts. Concepts, as a theory within cognitive science, focus on the understanding of similarity between entities, allowing for categorisation. In Gärdenfors' theory, a Conceptual Space is a representation model where each concept is represented by a region in a geometric space defined by a number of quality dimensions. For example, if we were to represent fruit, we may consider the size, shape, taste and texture as different quality dimensions. Colour is a good example as we naturally represent it dimensionally; Gärdenfors discusses the HSV (Hue, Saturation, Brightness) colour space, as it represents the phenomenal structure of the colours rather than the scientific properties. There is no claim made that this closely represents human perception; in fact there are many other dimensionally represented colour spaces which could also be applied to the same representation including RG, RGB, CMYK and NCS Wyszecki and Stiles (1982). Interestingly the eye has detection mechanisms along three distinct spectrums (Blue-violet, Green and Yellow-Red) so even at the detection level, we perceive colour dimensionally.

### Braitenberg Vehicles

Braitenberg (1986) proposed a series of hypothetical thought experiments concerning what he called "vehicles". These self operating machines were designed to show how seemingly complex behaviour that an observer agent might associate with traits such as fear and aggression could be exhibited from relatively simple sensory motor connections. We will now discuss vehicle types 2 to 4 and 7 in particular as they directly relate and have inspired our approach.




Figure 1: The direct connection type 2a vehicle turns away from the light, where as the type 2b turns towards the light.

**Types 2 and 3** A type 2 vehicle has two light sensors connected to two motors. Type 2a has a direct connection (the left sensor to the left motor and right to right) while in the 2b model the connections are crossed over. This type of design allows the vehicles to demonstrate rudimentary fear (2a) and aggression (2b) towards a light source as they will either turn towards or away from it (see figure 1). With the type 3 model, the connections are negative (inhibitory) as opposed to positive (excitatory); for example, a light source will slow a motor rather than speed it up. This produces a more gradual effect, vehicles that circle a light source, or vehicles that will stay close to a light source while exploring their environment.

The type 3c vehicle has a mix of sensors with a variety of connections allowing the vehicle to display all the characteristics of the other type 2 and type 3 vehicles. These vehicles are described as having light, heat, oxygen and bio matter sensors connected respectively in the aggressive, fear, loving and exploring formats. This type of vehicle is able to produce complex behaviour, such as appearing to show preferences, disliking certain areas and favouring others.

**Types 4 and 5** In the type 4 and 5 vehicles, threshold devices are introduced to the design. These pass-through devices limit power by allowing no output until a threshold level has been reached, or inversely by allowing everything until a threshold has been reached. In type 4 vehicles, this principle is used to simulate instinctive decisions, advancing the base type 3 vehicles. In type 5 vehicles, Braitenberg discusses how these thresholds could be used to develop a simple logic system and memory.

**Concepts** In the type 7 vehicle, Braitenberg discusses how concepts can be represented in a vehicle using a hypothetical component called Mnemotrix. This component, described as a wire, has an interesting property—its resistance is high

until an electrical current simultaneously traverses the two components it is connected between. This allows the vehicles to associate input with a specific behaviour. For example, if the oxygen sensor was high while the bio-sensor was also activated, the vehicle can “learn” to associate biology with oxygen by, quite literally, crossing its wires. While this form of association is one definition of a concept (a general idea derived from instances or occurrences), it does not describe the abstract concepts this paper seeks to explore.

## Applications

Braitenberg vehicles have been applied to a large number of experimental and practical applications, particularly as models for animal behaviour, but they have also been directly applied to robotics. For example, Yang et al. (2006) applies Braitenberg vehicles to the navigation of mobile robots and Lilienthal and Duckett (2003) experiments with using types 2, 3 and 4 Braitenberg vehicles to allow a mobile robot to locate a static odour. Salumae et al. (2012) applied the type 2b controller to the design of an underwater robot, with sensors allowing it to localise and navigate a current. Braitenberg vehicles, due to the simplicity of their design, provide an exceptional platform for experimental investigation into a variety of reactive embodied behaviours.

## Fungus Eaters

In 1962, Toda described a thought experiment involving autonomous agents known as fungus eaters (Pfeifer, 1996; Toda, 1962). This study was the origin of the concept of complete autonomous agents and one of the earliest works in embodied artificial intelligence, pre-dating Brooks by almost 30 years.

Toda describes a robot “sent to a hypothetical planet as a robot uranium miner, which sustains itself by eating fungi as its energy source”. The solitary fungus eater has a means of collection and locomotion and a means of making decisions while interacting with its environment. It must be autonomous as it cannot receive any external aid.

Toda’s original argument discusses how notions of emotion, irrationality and social behaviour can be applied to a Fungus Eater simply as a function of operating within its task environment and attempting to maintain itself by eating fungi.

## Template Based Evolution

Template Based Evolution (TBE) is a methodology for evolving reactive, subsumption based agents. This is achieved implicitly, by testing agents through trials in an environment, rather than via a fitness function. The motivation behind the development of the algorithm and methodology was the work of John Holland (1990).

## Genome

The Genome in a TBE simulation is a collection of unique attributes, each attribute taking a specific role in the behaviour of the agent. In the original paper (Headleand and Teahan, 2013), two attribute types were defined, action and trait which are usually represented numerically. This was subsequently extended in (Headleand, 2013) to include threshold attributes.

**Trait** Trait attributes represent qualities of the agent, for example, colour, speed, energy, weight etc. They can be independent of other attributes but can also affect selected behaviours; for example, a trait of speed may affect a run behaviour.

**Action** An action attribute is an index value which represents the evolved selection of an action within a list of possible actions within the subsumption. During each time step, as the agent moves through its subsumption, the output is determined by the action attribute.

**Threshold** The third class of attribute described in Headleand (2013) was for alternate styles of simulation. This is used to define the threshold where different higher layers of the template subsume lower layers. Alternatively, it may be used to define activation functions or decision processed within specific actions.

## Design of a TBE simulation

A TBE simulation is built from the bottom up, starting with the environment and ending with the individual agents. This formal methodology forces the designer to consider the ecological niche of the agents. As a result, the task environment is integral to the behaviour rather than an abstract container, which the agents are simply placed within.

**Environment** The environment is the domain where each agent's suitability is tested through successive trials. A trial is a challenge or obstacle specific to that environment, which may directly affect the agent's ability to survive (and subsequently reproduce). Trials can take multiple forms, for example if we were to simulate the behavioural evolution of herd animals, a trial may be a predator, exposure to adverse weather or the availability of food.

**Species** All agents of a species of a specific species must conform to a basic prototype. This includes the species template and the genome prototype (the amount and type of attributes). The species template is a subsumption architecture, with either outputs or conditions being determined by attributes from the agents genome. This ensures that whilst the subsumption architecture template is common across all agents in the simulation, the agents' unique genome will define their behavioural responses. (For an example, refer to figure 3).

**Agents** As the subsumption template is established within the species definition, each unique agent is only required to maintain their own genome. This allows for a moderate efficiency increase over alternative methods such as evolving an Artificial Neural Network or Markov Network where each unique agent in the simulation requires the storing of a full network.

## Applications

In Headleand and Teahan (2013), the TBE methodology was applied to simulating the behavioural evolution of migratory birds. The purpose of this study was to test the methodology but also to assess whether the predator escape tactics proposed in the literature could be evolved implicitly. The study demonstrated that the TBE methodology was capable of evolving complex behavioural and morphological changes such as wide vision and group escape tactics.

## Berry Eaters: An adapted thought experiment

In the background to this research, we described two thought experiments, Toda's fungus eaters and Braitenberg vehicles. In this section, we will discuss an adaptation that adopts aspects of each—the Berry Eaters (BE) experiment has been devised to allow exploration of abstract concept learning within reactive embodied agents.

## The Scenario

As discussed in Toda's "the Design of a Fungus Eater" (Toda, 1962), the author describes a hypothetical scenario where robots are sent to a planet where a native fungus is used as fuel. We propose an alternative to this experiment where there are two (rather than one) similar fuel source options. One fuel source will power the robot (known as edible) the other is incompatible with the robot (known as poisonous). These two alternatives are distinguished by one key feature—colour—for the purpose of this study. To add additional complexity, the designers of the eaters will have no knowledge of these features during the design process.

In order to differentiate our thought experiment between this and the fungus eaters experiment, the fuel source will be berries.

## Design of a Berry Eater

The Braitenberg vehicles demonstrate how seemingly complex behaviours can emerge from simple sensory motor couplings. The insight gained from these experiments provides us with possible theories as to how simple creatures, such as ants, can demonstrate complex behaviour despite limited cognitive capacity.

As proposed for our experiment, there are two types of berries in this experiment, edible and poisonous. For the purpose of discussion, let's say that the edible berries are red and the poisonous berries are blue. On face value, recognition in this case seems relatively straight forward. However,


consider the colour red. Although we may have a clear idea of what constitutes a prototypical red, in fact there are many of different shades of red due to slight variations in the hue, saturation and brightness. When we consider these variations, we are faced with boundary problems such as defining the point that red stops and pink begins, for example. As humans who can easily recognize colour differences, clearly we have an abstract conceptual understanding that allows us to overcome these problems. It is also reasonable to assume that this ability is adaptive, developing as new information becomes available.

**Basic Design – BE Type 1** Using a Brightenberg vehicle as a basis for the design, we can combine the qualities of a type 2a and 2b vehicle to create a hybrid that would turn away from one artefact and towards another. However, this only considers the sensory motor coordination; it does not address in itself how such a vehicle could learn which colour of berry is edible and which is poisonous, or how to go about detecting colour in the first place.

The problem of how to detect colour can be overcome in the following way (we call this a BE Type 1). Each light sensor in our hybrid vehicle is replaced with three wavelength sensors either side calibrated to red, green and blue respectively. As the sensor detects a wavelength in its calibrated spectrum, it produces an electrical signal. A circuit with a preset threshold and cut-off (such as within the type 5 vehicles) can be used to only produce an electrical signal at a particular value along that spectrum. We can take the three outputs from these threshold components into a set of switches, allowing current to pass through only if all three threshold devices are triggered simultaneously (for discussion's sake, let's call this component an activator). The combination of these components allows us, as the designer, to set the BE to respond to a specific colour (see figure 2.(a) where we have illustrated the circuitry required to get the BE to move towards a colour). Furthermore, we can include two sets of activators, one for a specific colour to avoid and another for a colour to be attracted to.

This would result in a Berry Eater that can move towards and away from two different coloured objects. However, this requires that the objects match an exact colour. This current configuration does not allow for the variation in the colour of the berries we discussed earlier; even an altering of the light conditions would render this particular design of berry eater incapable of operation.

**Advanced Design – BE Type 2** Our more advanced design of a berry eater (BE Type 2) is capable of recognising variation in colour by incorporating a conceptual spaces inspired model of representation. In the previous BE Type 1 design, we had a single threshold/cut-off circuit. However, for our more advanced design, we propose using two circuits, a threshold that only allows signals to pass through



(a) The basic Berry Eater design, equipped to respond to a single colour. A duplication of these components with the wires wired directly instead of crossed over would allow the BE to move away from another colour.

(b) The advanced Berry Eater capable of representing colour concepts. Notice the threshold component triggered only if a signal is above a specific value and the cut-off component that allows all signals below a specific threshold.


Figure 2: Design of a Berry Eater.

if they are over a certain value followed by a separate cut-off that only allows signals below a certain value to pass through. Using this system, it is possible to define a RGB conceptual space type controller that is able to encode variations around a colour prototype in order to represent colour concepts; see figure 2.(b). In addition, through the unsupervised evolutionary process described below, the agent is able to learn the colour of the two distinct classes of berries, those that provide fuel and those that will destroy it.

For the simulation of our final Berry Eater design, we first generate a population of Berry Eaters with randomly assigned threshold values and deploy the agents within the environment. Before new BE's are manufactured, the threshold values of two surviving agents (the parents) are recombined and mutated to produce the thresholds of a new agent (the offspring). To help ensure that the parents are successful agents, we only breed from two parents who have survived for a given amount of time (a maturity measure). The agents unable to find edible berries will have run out of power (due to starvation), and agents unable to differentiate between edible and poisonous berries will also have been destroyed, ruling them out for reproduction.

The survival and success of many species relies on the interaction of multiple agents to produce new offspring. However, the BE design does not contain any mechanisms for communication other than breeding and each BE does not even have the ability to sense each other's presence. While we can predict the possible behaviour of an individual BE, it is significantly harder to imagine what emergent qualities a colony of berry eaters may exhibit. To explore this question, we performed the simulation as described below.

### Genome to Berry Eater Phenome Mapping



#### TBE Template (left sensor)

```

sR      = sensor-Red;
sG      = sensor-Green;
sB      = sensor-Blue;
R,G,B   = false;

//check sensor reading is within
//conceptual space defined by genome
if(sR > g1 && sR < g4) {R = true}
if(sG > g2 && sG < g5) {G = true}
if(sB > g3 && sB < g6) {B = true}

//output to motor
if( R&& G && B){set right-motor + 1}
else {set right-motor - 1}

```

Figure 3: An illustration of the genome attribute to BE phenome mapping for the “Box Representation” approach, allowing the Berry Eater to be attracted towards a colour, and pseudo-code implementation of the left sensor template. Note that these components would need to be duplicated with the motor wires uncrossed for the Berry Eater to also avoid a colour.

## Experimental Results

In this section, we will discuss the design of the Berry Eaters simulation and the experimental results. The design focuses on the TBE methodology, starting with the environment and ending with the individual agents. The purpose of this study was to explore two principle research questions:

1. Could a Conceptual Spaces inspired approach be implemented within a reactive embodied agent to provide a mechanism for representing colour concepts?
2. What behaviours would a group of Berry Eaters produce when placed into an environment?

### Environment

The environment that we simulated in the NetLogo programming language is a wrapped world 50 units high by 50 units wide. Each unit, known as a patch in NetLogo, is large

enough to fit a berry in its entirety, though multiple berries may be stacked on top of each other.

Within the environment, there are the poisonous berries described previously, and any agent who eats a poisonous berry is destroyed. All agents will starve if they are unable to eat within a specific time period (set at 360 ticks within the simulation). To be able to breed, all agents will also have to reach a maturity which is set to 400 ticks within the simulation requiring breeding agents to have eaten a berry. These values were selected due to the size of the environment, 360 ticks being considered enough opportunity to find a food source. 400 ticks was selected as it forced a BE to eat at least once before it was able to breed (to exclude unsuccessful agents).

### Species

In the previous section, we presented a Braitenberg-style design of the Berry Eater. This design is implemented as a subsumption architecture within the species template (see figure 3). In the pseudo-code shown in the figure for the left sensor (attraction), the RGB readings are checked to see if the detected colour falls within the conceptual space defined by the genome attributes, if they do the right motor speed is increased. For this study the population size was fixed to 500 vehicles, each new vehicle had a mutation chance set to 5%.

The attributes within the genome definition represent the threshold values for the colour detection. This requires six attributes for positive (edible) detection and a further six attributes for negative detection (poison).

### Agent

The agents are Berry Eaters that maintain their own unique genome of attributes. The agents respond to their environment based on the signals they receive from their virtual sensors. This is illustrated in figure 3. If the agents receive no signals, they will simply move forward, wandering slightly to simulate friction/slip of their wheels on the ground as with traditional Braitenberg Vehicles.

**Attributes** For the purpose of this study, threshold attributes have been used to define the conceptual spaces as described in an earlier section on Template Based Evolution and in (Headleand, 2013).

**Representation** The BE Type 2 design described one method of defining a conceptual space, using upper and lower thresholds along each dimension to construct a region in 3D space. In this section, we discuss how this box is implemented and present an alternative, sphere-based representation which uses a single point in RGB space, and a radius which defines the region.

*Box representation:* The box representation uses a lower and upper threshold to define three acceptable ranges along three




Figure 4: An illustration of two conceptual regions defined around their central colour prototypes, for both the box and spherical representations.

dimensions, creating a box in 3D space. This method is similar to the original conceptual spaces theory proposed by Peter Gärdenfors. In this method, any color which falls within this region is considered conceptually similar to the learned prototype. To implement the box representation in TBE, six threshold attributes were used to define a shape in three dimensional RGB space.

*Spherical representation:* An alternative to the box representation was also implemented. In this version, a single point in 3D space, the colour prototype, was learnt by the Berry Eater. In addition, the Berry Eater also learnt a single value representing an acceptable distance from the prototype, creating a sphere in 3D space. In this representation, any colour which fell within this sphere was considered conceptually similar to the learnt prototype. To implement this method in TBE, three threshold attributes were used to represent the position within the space and a fourth attribute represented the distance of the radius of the sphere.

**Colour Space** In this study, we have chosen to represent the RGB colour space. Although HSV more closely matches Gärdenfors original description, adopting RGB has the advantage that RGB sensors are readily available, and therefore real-world Berry Eaters can be developed in future to see how the simulated and real-world equivalents compare.

## Experimental Setup

We conducted several experiments in order to help answer the two research questions posed above. Six different experimental setups were designed (detailed in Table 1), and for each of these experiments, we implemented both the box and the spherical conceptual space representation (experiment variants (a) and (b)). For all experimental setups (12 in total), we ran 100 simulations (1200 simulations in total). Each simulation was allowed to run for 5000 ticks, at which point the results were recorded.




Figure 5: Berry eaters demonstrating feeding frenzy behaviour: in the middle of the left side of the image, we can see a feeding frenzy about to reach its conclusion. At the edible (green) clump towards the bottom of the centre of the image, we can see another beginning to form.

## Group Behaviour Patterns

Various behaviours were visibly apparent during the simulations as discussed below. Although the agents have no knowledge of each other, the simple environment to agent feedback has resulted in visually complex group behaviours.

In this section, we have made no differentiation between which conceptual representation was used (sphere or geometric) as this did not make any difference in the behavioural outcome.

**Feeding Frenzy** When the berry-growing algorithm was set to generate a clump of berries, the BEs in the simulation regularly demonstrated a feeding-frenzy behaviour. This is where the vehicles would approach from multiple directions, swarming around a particular food source. The swarm would get larger and denser as the berries were eaten and re-grown in different locations within the clump, until eventually all the berries were eaten and a clump formed elsewhere.

In the middle of the left side of figure 5, we can see the final results of a feeding frenzy, with agents in a dense swarm covering a diminishing edible clump. Additionally in the center of the image towards the bottom, we can see a second beginning to form, with agents beginning to swarm from above and left.

**Path Following** If the berry growth algorithm was set to random, the BEs would end up moving along similar paths. These paths resemble ant trails, and therefore provide the impression that the agents are performing a “follow the leader” type behaviour. However, the agents have simply

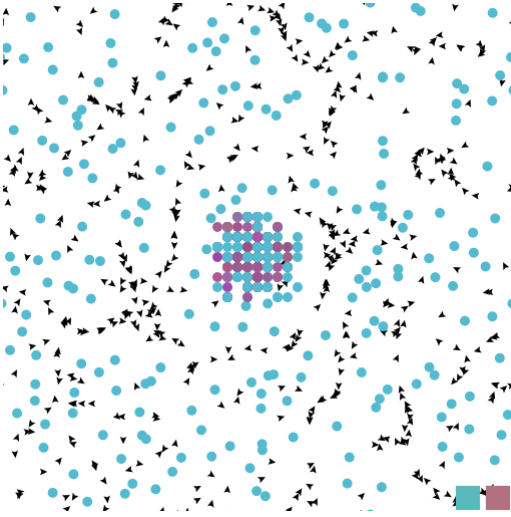


Figure 6: Berry eaters demonstrating path following behaviour: notice how the agents in the simulation appear to be following each other on a path between the poisonous berries. Around the clump in the centre, we can observe the apprehensive behaviour mentioned later in section .

learnt the correct colours and are navigating the environment by being repelled away from the poisonous berries and being attracted to the edible ones. This results in reactive path-following behaviour based on the placement of the berries in the simulation (see figure 6).

**Pseudo-Flocking** Often the agents would demonstrate a behaviour that resembled flocking. Indeed, this flock would move in unison for a significant amount of time before being broken up by the presence of food or poisonous berries to avoid. A group of agents would be attracted towards a series of berries (often forming a path). When the available berries were eaten, the group of agents would begin to align. As they were subsequently repelled or attracted to other berries along their route, the flock would maintain its structure until a large food source or clump of poisonous berries forced the flock to disperse.

**Apprehension** If the berry growth algorithm was set to clump, occasionally the poisonous and edible berries would grow together. This caused interesting apprehensive-like group behaviours to appear as a result. The BEs would approach from the outside, but once there they would twitch and turn in multiple directions. Occasionally one would rush forward before quickly returning back to the outside. In figure 7, you can see that the berry eating trend has begun to taper at the top of the graph. There is also a significant step a little lower. Both these drops in the berries being eaten have occurred because of the described behaviour.

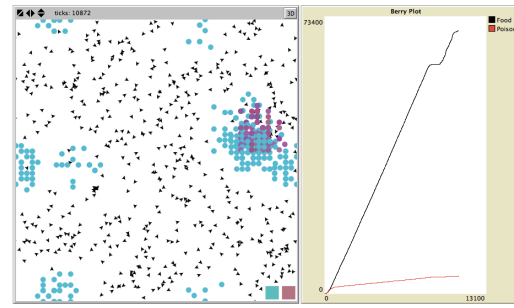


Figure 7: Berry Eaters demonstrating apprehensive behaviour: on the graph, we can see one step where this situation has already taken place and at the top of the graph we can see another one beginning to form. A few agents can be seen exploiting a few solitary edible berries that have grown between the poisonous clumps. However, the rest of the colony has avoided the clump creating a small exclusion zone around the area.

## Colour Concept Acquisition

The problem of colour concept acquisition is the core focus of this study. As described above, two possible methods of colour conceptual space representation have been explored: the first is where the concept is represented geometrically in 3D space as a sphere (or as a hypersphere in higher dimensions), and the second is where the concept is represented in 3D space as a box. We also discussed two berry growing algorithms, one which caused the berries to grow in clumps and another which caused the berries to grow randomly within the environment.

A colour has been considered ‘successfully learnt’ if the average colour that the agents have learnt has a Euclidean distance of less than some threshold distance from the colour prototype which is randomly generated at the beginning of the experiment. (In our simulation, we set the threshold to 20 because two colours with a Euclidean distance of less than 20 were visually of the same prototypical colour for a human observer).

The 12 different experimental setups are presented in table 1. In the second column, we have stated which berry growing algorithm was used: either Clump or Random. The column labelled ‘Rep.’ details the concept representation method that was used—either box, or spherical. In the columns labelled ‘E’ and ‘P’, we detail the number of edible (‘E’) and poisonous (‘P’) berries in the simulation. In the column labelled ‘LE’, we list the number of unique simulations where the agents successfully learnt the colour of the edible berries. In the column labelled “LP”, we list the percentage of experiments where the agents successfully learnt the colour of the poisonous berries.



Setup	Growth	Rep.	E	P	LE	LP
1.(a)	Clump	Box	300	300	71	62
1.(b)	Clump	Sphere	300	300	75	64
2.(a)	Random	Box	300	300	61	56
2.(b)	Random	Sphere	300	300	73	45
3.(a)	Clump	Box	150	450	94	87
3.(b)	Clump	Sphere	150	450	96	99
4.(a)	Random	Box	150	450	71	63
4.(b)	Random	Sphere	150	450	85	92
5.(a)	Clump	Box	450	150	63	48
5.(b)	Clump	Sphere	450	150	67	40
6.(a)	Random	Box	450	150	97	22
6.(b)	Random	Sphere	450	150	78	11

Table 1: Results from various setups of the Berry Eaters simulation.

### Analysis

From table 1, we can see that the agents were typically successful in learning the colours in the simulation. Also, typically the spherical representation produced superior results. However, setups 6.(a) and 6.(b) produced confusing results: despite the fact that the environment was still populated with poisonous berries, the agents seemingly did not learn the poisonous colour although observations of these experimental setups showed that the agents still avoided poisonous berries.

For setup 6.(b) (spherical representation), we found that the agents developed an interesting coping mechanism. Instead of learning the specific poisonous colour prototype, the BEs learnt to avoid large but inaccurate conceptual space, however due to the size of this region it typically included the poisonous colour.

### Conclusion

In this paper, we have discussed using a geometric representation as a viable method of representing concepts. This was demonstrated through the Berry Eaters experiment where virtual agents were able to evolve a simple conceptual understanding of colour which directly influenced their behaviour. We explored how these artificial life forms could be built following the same format proposed by Braitenberg and tested them through simulation.

Despite having no awareness of each other, the Berry Eaters acted in such a way that gave the appearance of group behaviour simply by following common rules and reacting to the environment. The simplest form of this pseudo-group behaviour was agents moving in the same direction, either along paths or in clusters. However, we also observed flocking-like behaviour, which exhibited many of the emergent qualities visible in more complex models such as Boids (Reynolds, 1987). This type of behaviour is usually only observed in models where the agents are aware of each other;

however, our simulation demonstrates that an environment based approach could generate comparable behaviour.

Typically, the spherical representation was more successful in achieving the goal of colour concept learning, although there was evidence of over-fitting and this approach is also further from Gärdenfors original theoretical framework. Future studies will explore other methods of geometric representation to determine whether further improvements can be gained.

### Acknowledgements

Chris Headleand would like to thank Fujitsu for their ongoing financial support.

### References

- Braitenberg, V. (1986). *Vehicles: Experiments in synthetic psychology*. MIT press.
- Gärdenfors, P. (2004). *Conceptual spaces: The geometry of thought*.
- Gärdenfors, P. and Williams, M.-A. (2001). Reasoning about categories in conceptual spaces. In *IJCAI*, volume 1, pages 385–392.
- Headleand, C. J. (2013). Bio-inspired methods of behaviour generation. Master’s thesis.
- Headleand, C. J. and Teahan, W. J. (2013). Template based evolution. In *Proceeding of the fifteenth annual conference companion on Genetic and evolutionary computation conference companion*, pages 1383–1390. ACM.
- Holland, J. H. et al. (1990). Echo: Explorations of evolution in a miniature world.
- Lilienthal, A. and Duckett, T. (2003). Experimental analysis of smelling braitenberg vehicles. *environment*, 5:10.
- Pfeifer, R. (1996). Building fungus eaters: Design principles of autonomous agents. *From animals to animats*, 4:3–12.
- Pfeifer, R. and Scheier, C. (2001). *Understanding Intelligence*. MIT press.
- Reynolds, C. W. (1987). Flocks, herds and schools: A distributed behavioral model. In *ACM SIGGRAPH Computer Graphics*, volume 21, pages 25–34. ACM.
- Salumae, T., Ranó, I., Akanyeti, O., and Krusmaa, M. (2012). Against the flow: A braitenberg controller for a fish robot. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 4210–4215. IEEE.
- Toda, M. (1962). The design of a fungus-eater: A model of human behavior in an unsophisticated environment. *Behavioral Science*, 7(2):164–183.
- Wyszecki, G. and Stiles, W. S. (1982). *Color science*, volume 8. Wiley New York.
- Yang, X., Patel, R. V., and Moallem, M. (2006). A fuzzy-braitenberg navigation strategy for differential drive mobile robots. *Journal of Intelligent and Robotic Systems*, 47(2):101–124.